

Review

Not peer-reviewed version

Domain Specific Benchmarks for Evaluating Multimodal Large Language Models

[Khizar Anjum](#) , [Muhammad Arbab Arshad](#) , [Kadhim Hayawi](#) , [Efsthios Polyzos](#) , Asadullah Tariq ,
[Mohamed Adel Serhani](#) , [Laiba Batool](#) , [Brady D. Lund](#) , [Nishith Reddy Mannuru](#) , [Ravi Varma Kumar Bevara](#) ,
Taslim Mahbub , Muhammad Zeeshan Akram , [Sakib Shahriar](#) *

Posted Date: 26 May 2025

doi: 10.20944/preprints202505.1993.v1

Keywords: Large Language Models; Artificial intelligence; domain-specific analysis; Artificial General Intelligence (AGI); Benchmarking LLMs



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a Creative Commons CC BY 4.0 license, which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Review

Domain Specific Benchmarks for Evaluating Multimodal Large Language Models

Khizar Anjum ^{1,†}, Muhammad Arbab Arshad ^{2,†}, Kadhim Hayawi ³, Efstathios Polyzos ³, Asadullah Tariq ⁴, Mohamed Adel Serhani ⁵, Laiba Batool ⁶, Brady Lund ⁷, Nishith Reddy Mannuru ⁷, Ravi Varma Kumar Bevara ⁷, Taslim Mahbub ⁸, Muhammad Zeeshan Akram ⁹ and Sakib Shahriar ^{10,*}

- ¹ Rutgers University, USA
- ² Iowa State University, USA
- ³ Zayed University, UAE
- ⁴ United Arab Emirates University, UAE
- ⁵ University of Sharjah, UAE
- ⁶ NUCES, PK
- ⁷ University of North Texas, USA
- ⁸ George Washington University, USA
- ⁹ University of Louisville, USA
- ¹⁰ University of Guelph, Canada
- * Correspondence: shahrias@uoguelph.ca
- † These authors contributed equally to this work.

Abstract: Large language models (LLMs) are increasingly being deployed across disciplines due to their advanced reasoning and problem-solving capabilities. To measure their effectiveness, various benchmarks have been developed that measure aspects of LLM reasoning, comprehension, and problem-solving. While several surveys address LLM evaluation and benchmarks, a domain-specific analysis remains underexplored in the literature. This paper introduces a taxonomy of seven key disciplines, encompassing various domains and application areas where LLMs are extensively utilized. Additionally, we provide a comprehensive review of LLM benchmarks and survey papers within each domain, highlighting the unique capabilities of LLMs and the challenges faced in their application. Finally, we compile and categorize these benchmarks by domain to create an accessible resource for researchers, aiming to pave the way for advancements toward artificial general intelligence (AGI).

Keywords: large language models; artificial intelligence; domain-specific analysis; artificial general intelligence (AGI); benchmarking LLMs

1. Introduction

The rapid evolution of Multimodal Large Language Models (MLLMs) has led to remarkable general capabilities. However, achieving nuanced performance in specialized areas—the ‘last mile’—requires a dedicated focus. This paper embarks on a comprehensive survey of domain-specific benchmarks for MLLMs, contextualizing recent advancements, articulating the critical need for domain-specific evaluation, and outlining this survey’s main contributions towards fostering targeted MLLM development.

1.1. Recent Advances and the Rise of Foundational Models

Recent breakthroughs, exemplified by the launch of highly capable models like GPT-4 [1] and Gemini [2], have dramatically shifted the MLLM landscape. Beyond initial scaling in parameters and data, a key advancement lies in “instruction tuning” through techniques such as Reinforcement Learning from Human Feedback (RLHF). This has made models more adept at understanding and responding to human requests. Further capabilities include Retrieval Augmented Generation (RAG), enabling MLLMs to incorporate external, up-to-date information, and the development of agentic

AI, where models can utilize tools and APIs to perform actions in the digital or physical world. The evolution also extends to more sophisticated reasoning processes, sometimes referred to as “test-time scaling” or “thinking,” allowing models to tackle complex problems step-by-step. Multimodal capabilities, particularly in areas like image generation and understanding, continue to mature, integrating diverse data types into cohesive reasoning frameworks.

The MLLM revolution is built upon “foundational models” [3], often characterized by their critically central yet incomplete nature and their broad pre-training that allows adaptation to a wide range of downstream tasks. These models, typically based on the transformer architecture (e.g., LLaMA 3 [4], Phi-4 [5]), are initially trained via self-supervised learning on vast and diverse datasets, often encompassing trillions of tokens from text and other modalities. A core principle has been that scaling—increasing model parameters, training data, and computational resources—leads to qualitative shifts and emergent capabilities, such as improved understanding and reasoning, that were not explicitly programmed. These pre-trained foundational models then serve as a base that can be adapted through further training phases, like instruction tuning, to perform a wide array of specific tasks, making them profoundly versatile.

1.2. The Case for Domain-Specific Benchmarks

While foundational MLLMs exhibit impressive general intelligence, their efficacy often diminishes when confronted with the nuanced demands of specialized domains. This ‘last mile problem’ underscores a critical gap: general-purpose models, despite their vast training, frequently falter on tasks requiring deep, domain-specific knowledge, intricate reasoning, or precise interpretation of specialized data. For instance, even leading models like GPT-4 have demonstrated significant limitations in specialized financial question answering, achieving low accuracy on benchmarks like FinanceBench [6]. Similarly, in software engineering, strong performance on general coding tasks does not readily translate to proficiency in domain-specific contexts, as highlighted by evaluations on DomainCodeBench [7]. Furthermore, in robotics, benchmarks such as MMRo [8] reveal that while models may excel at high-level planning, they can struggle with fundamental perceptual tasks crucial for real-world interaction. These examples, which are explored in detail in subsequent sections of this paper, motivate the central thesis of this survey: the critical need for developing and analyzing domain-specific benchmarks to truly harness and evaluate the potential of MLLMs in diverse, real-world applications.

This focus on domain specificity is not limited to academic research; it is also evident in the trajectory of leading commercial MLLMs. For instance, models like Anthropic’s Claude have demonstrated strong performance in software engineering tasks, as seen on benchmarks such as SWE-bench Verified where Claude 3 Opus achieved 70.3%, slightly ahead of OpenAI’s GPT-4o (referred to as ‘o3’) at 69.1%. Conversely, on mathematical reasoning benchmarks like AIME 2025 (a proxy for high-level math contests), GPT-4o showed a significant lead with 88.9% compared to Claude 3 Opus’s 49.5%. This divergence highlights that even large foundational models exhibit varying strengths and weaknesses across different domains, reinforcing the idea that a one-size-fits-all approach has limitations and that specialized capabilities are becoming key differentiators.

Given these limitations of generalist models, this survey turns its attention to the growing ecosystem of benchmarks specifically engineered for domain-focused MLLM development. These specialized benchmarks are not merely for assessing performance within a niche; they are crucial for driving innovation and refining models to handle the unique challenges of individual fields. For example, benchmarks in robotics like MMRo [8] serve as diagnostic tools, pinpointing failures in perception that general MLLMs must overcome for effective real-world deployment. Similarly, evaluations in areas like software engineering with DomainCodeBench [7] highlight the necessity for domain-aware prompting and contextual understanding, techniques which can enhance MLLM applicability more broadly. By pushing MLLMs to master complex, domain-specific data and reasoning patterns—from understanding financial nuances to interpreting specialized engineering diagrams or medical imagery—these benchmarks also provide invaluable feedback. The insights and architectural advancements spurred by the need to succeed on these focused evaluations can, in turn, enrich

the development of more robust and versatile foundational models, contributing to overall MLLM progress.

1.2.1. Main Contributions

Our main contributions are as follows:

- Presents a pioneering and comprehensive survey of domain-specific benchmarks for Multimodal Large Language Models (MLLMs) across a wide array of disciplines.
- Highlights the critical role of developing MLLMs tailored for diverse domains as essential for addressing the ‘last mile problem’ and achieving practical efficacy in real-world applications.
- Argues that such specialized advancements not only ensure domain-specific efficacy but also supplementally contribute to the broader development and refinement of large foundational models.

2. Review Methodology & Framework

This paper presents a systematic review of domain-specific benchmarks for Multimodal Large Language Models (MLLMs). Our review examines **eight key disciplines**, which are explored in detail in Sections 3–10:

1. **Engineering** (Section 3) - Including industrial engineering, software engineering, and systems engineering
2. **Science** (Section 4) - Covering geography & remote sensing, physics & chemistry, and environmental science
3. **Technology** (Section 5) - Encompassing computer vision & autonomous systems, robotics & automation, and blockchain & cryptocurrency
4. **Mathematics** (Section 6) - Addressing mathematical reasoning, formal methods, and statistical analysis
5. **Humanities** (Section 7) - Examining social studies, arts & creativity, music, urban planning, morality & ethics, and philosophy & religion
6. **Finance** (Section 8) - Covering financial forecasting, sentiment analysis, document QA, and regulatory compliance
7. **Healthcare** (Section 9) - Focusing on medicine & healthcare and medical imaging
8. **Language Understanding** (Section 10) - Including multimodal communication, cross-lingual tasks, and reasoning

Each discipline is further broken down into specific domains, their respective sub-domains and application areas, as visually outlined in Figure 1. A core component of our analysis for each domain is a comprehensive table designed to consolidate critical information about relevant benchmarks and survey papers. These tables meticulously detail characteristics such as:

- **Scale:** Dataset size and diversity, color-coded to indicate **precise quantities** (blue), **collection descriptors** (green), or **unspecified amounts** (gray)
- **Task Type:** Nature and complexity of the evaluation task
- **Input Modality:** Input types utilized (text, image, audio, video)
- **Model:** Models evaluated in the benchmark
- **Performance:** Quantitative or qualitative results reported
- **Key Focus:** Core objectives and applications of the benchmark

Complementing the structured information in the tables, the textual discussion within each domain section offers deeper qualitative insights. This narrative focuses on elucidating overarching trends, significant advancements, and persistent challenges observed within the various application areas. By synthesizing these observations, we aim to provide a nuanced understanding that extends beyond the summarized data, highlighting the current capabilities and limitations of MLLMs in specialized contexts and explaining the significance of each referenced paper in advancing the field.



Figure 1. The domain hierarchy of the eight disciplines covered in this article. The hierarchy is based on the domains and sub-domains of each discipline, and the sub-domains are further broken down into their application areas.

Search Strategy and Scope. Our search strategy involved systematically querying major academic databases (e.g., IEEE Xplore, ACM Digital Library, Scopus) and preprint archives (e.g., arXiv) using a combination of general keywords such as “Large Language Model,” “Multimodal LLM,” “benchmark,” “evaluation,” and domain-specific terms relevant to the disciplines covered. The scope of this review encompasses studies focused on Large Language Models (LLMs), Vision Language Models (VLMs), and more broadly, Multimodal Large Language Models (MLLMs). We specifically prioritized papers that introduce new benchmarks, conduct empirical evaluations of MLLMs on existing benchmarks, or provide comprehensive analyses of model performance within particular domains. Consequently, purely theoretical papers without a significant evaluation component were generally excluded. Survey papers are selectively included within specific sub-domains if they offer a valuable consolidation of existing benchmarks or evaluation practices pertinent to that application area, thereby enriching the contextual understanding of the domain’s evaluation landscape. Our primary focus is on papers that contribute to understanding how MLLMs are evaluated and perform in various specialized fields, rather than theoretical explorations of model architectures or training methodologies alone.

3. Engineering

Large Language Models (LLMs) have emerged as transformative tools across engineering disciplines, revolutionizing how engineers approach complex problems, design systems, and manage technical knowledge. By leveraging their ability to understand and generate domain-specific content, LLMs are increasingly being integrated into engineering workflows to enhance productivity, facilitate knowledge transfer, and support decision-making processes. The following sections explore how various engineering disciplines are developing specialized benchmarks to evaluate and improve LLM performance in technical contexts, highlighting both the promising capabilities and current limitations of these models in addressing real-world engineering challenges. An overview of all the relevant benchmarks is presented in Table 1.

Table 1. Engineering Discipline: Domain-Specific Benchmarks.

Domain	Sub-domain	Benchmark	Scale	Task Type	Input Modality	Model	Performance	Key Focus
Industrial Engineering	Manufacturing & Design	DesignQA [9]	N/A	Rule Comprehension, Extraction	Text + CAD Images	Multiple	1,449 Qs	Engineering Documentation
		Freire et al. [10]	Multiple tasks	Knowledge Retrieval	Text	Multiple	97.5% factuality	Factory Documentation
		Manu-Eval [11]	22 subcategories	Manufacturing Tasks	Multi-modal	20 LLMs	N/A	Manufacturing Industry
	Quality Control	FDM-Bench [12]	Multiple tasks	FDM-specific Tasks	Text + G-code	Multiple	62% accuracy	Additive Manufacturing
		LLM4PLC [13]	Multiple tasks	Code Generation	Text	Multiple	72% pass rate	PLC Programming
Industrial Engineering	Industrial Automation	Tizaoui et al. [14]	Multiple tasks	Extractive QA	Text	6 LLMs	N/A	Process Automation
		Xia et al. [15]	Multiple tasks	Production Planning	Multi-modal	LLM Agents	N/A	Production Systems
	Supply Chain Management	Ogundare et al. [16]	Multiple tasks	Problem Solving	Text	ChatGPT	N/A	Oil & Gas Engineering
		Rahman et al. [17]	Multiple tasks	Optimization	Text	GPT-4o	95% correct	Supply Chain Optimization
Software Engineering	Knowledge Graphs & Semantic Systems	OptiGuide [18]	Multiple tasks	Query Processing	Text	LLMs	92% accuracy	Supply Chain Operations
		Raman et al. [19]	150 questions	Question Answering	Text	ChatGPT, Bard	4.95/5 accuracy	Supply Chain Management
		LLM-KG-Bench [20]	Multiple tasks	Knowledge Graph Engineering	Text	Multiple	N/A	Knowledge Graph Generation
	Code Review & Analysis	BIG-bench [21]	204 tasks	Multiple Tasks	Text	Multiple	N/A	General Capabilities
		Azanza et al. [22]	Multiple tasks	Test Generation	Text	Multiple	90% (2024)	Software Testing
Systems Engineering	System Architecture	DomainCodeBench [7]	2,400 tasks	Code Generation	Text	10 LLMs	+38% w/ context	Domain-specific Coding
		StackEval [23]	Multiple tasks	Coding Assistance	Text	Multiple	95.5% (hist.), 83% (unseen)	Code Writing & Review
	Reliability Engineering	Zhang et al. [24]	947 studies	Multiple Tasks	Text	62 LLMs	N/A	Software Engineering
		SysEngBench [25]	Multiple tasks	Systems Engineering Tasks	Text	Multiple	N/A	Systems Engineering
Systems Engineering	Reliability Engineering	Hu et al. [26]	Multiple tasks	Question Answering	Text	GPT-4, GPT-3.5	92% (CRE)	Reliability Engineering
		Platinum Benchmarks [27]	15 benchmarks	Multiple Tasks	Text	Multiple	N/A	Model Reliability
Systems Engineering	Trustworthiness Assessment	Liu et al. [28]	29 categories	Trustworthiness Assessment	Text	Multiple	N/A	LLM Trustworthiness

3.1. Industrial Engineering

In the realm of industrial engineering, LLMs are being deployed to optimize manufacturing processes, improve quality control, streamline supply chain operations, and enhance industrial automation. These applications require models to interpret technical specifications, understand engineering drawings, analyze production data, and generate actionable insights. Industrial engineering benchmarks are particularly focused on evaluating how well LLMs can bridge the gap between general language understanding and specialized domain knowledge, with emphasis on their ability to reason about physical systems, interpret multimodal inputs (text, images, CAD files), and provide reliable recommendations in safety-critical environments. The growing adoption of LLMs in this field reflects their potential to address persistent challenges in knowledge management, process optimization, and decision support across various industrial sectors.

3.1.1. Manufacturing & Design

Recent years have seen the emergence of several specialized benchmarks and studies aimed at rigorously evaluating the capabilities of large language models (LLMs) and multimodal LLMs (MLLMs) in the context of manufacturing and design [9,10,12]. The following paragraphs provide detailed explanations of key works in this area, each highlighting unique contributions, evaluation strategies, and the current state of LLM performance in addressing real-world engineering challenges.

DesignQA [9] introduces a novel multimodal benchmark for evaluating the ability of large language models (LLMs) and multimodal LLMs (MLLMs) to comprehend and apply engineering requirements from technical documentation. Developed using real-world data from the Formula SAE competition and the MIT Motorsports team, DesignQA uniquely combines textual design requirements, CAD images, and engineering drawings. The benchmark is structured into three segments—Rule Extraction, Rule Comprehension, and Rule Compliance—mirroring the core tasks engineers face in design processes. It features 1,449 questions that test models on extracting rules from lengthy documents, identifying components in CAD models, and checking design compliance with requirements. Evaluation of state-of-the-art models (GPT-4o, GPT-4, Claude-Opus, Gemini-1.0, LLaVA-1.5) revealed that, while GPT-4o generally performed best, all models struggled with reliably extracting rules, recognizing technical components, and analyzing engineering drawings. The study highlights the current limitations of MLLMs in handling complex engineering documentation and sets a foundation for future research in AI-assisted engineering design.

FDM-Bench [12] is the first comprehensive benchmark specifically designed to evaluate LLMs on tasks related to Fused Deposition Modeling (FDM) in additive manufacturing. Developed by a collaboration of researchers from the University of Illinois, Rutgers, and the University of Michigan, FDM-Bench assesses both the ability of LLMs to detect anomalies in G-code (the programming language for 3D printers) and to answer user queries at varying expertise levels. The benchmark includes both multiple-choice and open-ended questions, as well as G-code samples representing different types of print defects. Four leading LLMs—GPT-4o, Claude 3.5 Sonnet, Llama-3.1-70B, and Llama-3.1-405B—were evaluated. Results showed that closed-source models (GPT-4o, Claude)

generally outperformed open-source models in G-code anomaly detection, while Llama-3.1-405B performed comparably or better in user query response. The study demonstrates the promise of LLMs in supporting FDM 3D printing but also highlights the need for further improvements, especially in visual defect detection and domain-specific fine-tuning.

Freire et al. [10] present a practical LLM-based knowledge management system for manufacturing environments, focusing on improving information retrieval and knowledge sharing among factory operators. The system uses Retrieval Augmented Generation (RAG) to answer operator queries by pulling relevant information from factory manuals and issue reports. A user study conducted at a detergent factory revealed that the tool was easy to use and improved access to information, but users still preferred learning from human experts when possible. The study also benchmarked seven LLMs (including GPT-4, GPT-3.5, Mixtral 8x7B, Llama 2, StableBeluga2, and Guanaco variants) on their ability to answer factory-specific questions. GPT-4 achieved the highest factuality and completeness, but newer open-source models like StableBeluga2 and Mixtral 8x7B showed strong performance and lower hallucination rates, making them attractive for privacy-sensitive industrial applications. The research highlights both the potential and the challenges of deploying LLM-powered tools in manufacturing, especially regarding safety, user acceptance, and the need for human oversight.

3.1.2. Quality Control

Quality control in industrial settings presents unique challenges for LLMs, particularly in the context of code generation and validation for programmable logic controllers (PLCs). LLM4PLC [13] presents an innovative pipeline for leveraging Large Language Models (LLMs) to generate verifiable programming code for Programmable Logic Controllers (PLCs) in industrial control systems. The approach integrates external verification tools—including grammar checkers, compilers, and symbolic model verifiers—into a user-guided, iterative workflow that transforms natural language specifications into IEC 61131-3 Structured Text code. The pipeline features automated feedback loops, prompt engineering, and LoRA-based fine-tuning to iteratively improve code quality and correctness. Experimental results on a real manufacturing testbed show that the pipeline increases code pass rates from 47% to 72% and boosts expert-rated code quality, while also reducing engineering effort from hours to minutes. The study demonstrates that, with proper verification and adaptation, LLMs can bridge the gap between natural language requirements and reliable, deployable PLC code, though challenges remain in ensuring explainability and safe deployment in critical systems.

3.1.3. Industrial Automation

Recent research in industrial automation has produced a range of benchmarks and case studies that explore the integration of large language models (LLMs) into process automation, production planning, and complex engineering problem-solving. The following paragraphs summarize key contributions in this area, highlighting the development of domain-specific datasets, the use of LLM agents in digital twin systems, and the evaluation of LLMs in specialized industrial contexts.

Tizaoui et al. [14] address the lack of domain-specific datasets for process automation by introducing a novel extractive question answering (QA) benchmark. Their methodology combines automated and manual techniques to generate and annotate a diverse set of QA pairs from academic papers in process automation. The resulting dataset, which includes both original and paraphrased questions, is used to fine-tune and evaluate six encoder-only LLMs. The study demonstrates that fine-tuning and data augmentation significantly improve model performance, and that models generalize well to unseen documents within the domain. This work establishes a foundation for future industrial NLP efforts by providing a robust benchmark and evaluation framework for process automation tasks.

Xia et al. [15] propose a novel approach to integrating LLMs into automated production systems using a hierarchical framework based on the automation pyramid. Their system semantically enriches low-level production data, making it interpretable for LLM agents that generate process plans and decompose them into atomic operations executed as microservices. The authors implement a digital twin system and a multi-agent LLM architecture in a modular production facility, demonstrating the

feasibility of LLM-driven planning and control. While the prototype shows promise in generating executable commands and handling both predefined and unforeseen scenarios, the study notes challenges related to real-time performance, the need for comprehensive real-world testing, and the importance of cost-benefit analysis for industrial adoption.

Ogundare et al. [16] evaluate the capabilities and limitations of LLMs, particularly ChatGPT, in solving complex oil and gas engineering problems. The paper assesses LLM performance on mathematical modeling tasks such as fluid flow, reservoir simulation, and equation discretization, using both stepwise and aggregate quality metrics. While LLMs demonstrate strong theoretical reasoning and the ability to generate sound equations for standard problems, they struggle with non-standard scenarios, complex geometries, and computational accuracy. The authors recommend enhancing LLMs with domain-specific knowledge and physical constraints, and suggest that LLMs are most valuable for theoretical model development rather than direct numerical computation in industrial engineering.

3.1.4. Supply Chain Management

Recent advances in supply chain management (SCM) and operations research have leveraged large language models (LLMs) to automate optimization, enhance interpretability, and support both educational and professional decision-making. The following paragraphs summarize key contributions in this area, highlighting frameworks for code generation and solver integration, natural language interfaces for optimization, and comparative evaluations of generative AI tools in SCM contexts.

Rahman et al. [17] demonstrate how LLMs, specifically GPT-4o, can bridge the gap between natural language problem descriptions and mathematical optimization in supply chain management. Their two-stage framework translates SCM problems into executable GUROBI code and interprets solver outputs for user-friendly analysis, using prompt engineering and advanced reasoning techniques like Tree of Thoughts. The study shows that LLMs can generate error-free code and provide accurate interpretations for a range of transportation problems, achieving approximately 95% correctness in output interpretation. While the approach accelerates model development and enhances accessibility, the authors note limitations in handling highly complex problems and emphasize the need for expert verification and further industrial validation.

Li et al. [18] introduce OptiGuide, a privacy-preserving framework that integrates LLMs with combinatorial optimization solvers to provide natural language explanations of supply chain decisions. OptiGuide enables users to pose plain text queries, which are translated into optimization code and executed by solvers, with results interpreted back into human-readable explanations. The system is designed to work alongside existing optimization technology, preserving proprietary data and leveraging in-context learning for flexible adaptation. Deployed in Microsoft's cloud supply chain, OptiGuide achieved over 90% accuracy in real-world server placement scenarios and received positive feedback from planners and engineers. The authors highlight the importance of specificity in user queries and the need for robust application-specific safeguards.

Raman et al. [19] provide a comparative analysis of ChatGPT and Bard in the context of supply chain management education, using a dataset of 150 certified supply chain professional exam questions. The study evaluates both models on accuracy, relevance, clarity, and readability, finding that ChatGPT outperforms Bard in accuracy and relevance, while Bard demonstrates slightly better readability. Both tools are shown to be effective study aids and professional resources, capable of supporting lesson planning, practice question development, and corporate training. The authors suggest that continual model adaptation and the use of diverse datasets are essential for ensuring ethical, effective, and up-to-date AI applications in supply chain management.

3.2. Software Engineering

Software Engineering has emerged as one of the most promising application domains for LLMs, with models demonstrating remarkable capabilities in code generation, debugging, documentation, and knowledge management. Benchmarks in this field evaluate LLMs on their ability to understand

programming languages, follow software specifications, adhere to best practices, and integrate with existing codebases. These evaluations span diverse programming paradigms, languages, and development environments, reflecting the multifaceted nature of modern software engineering. While LLMs show impressive performance in generating syntactically correct code and explaining programming concepts, challenges remain in ensuring semantic correctness, security, and maintainability of the generated solutions. The following sections explore specialized benchmarks that assess LLM performance in knowledge representation and code analysis tasks that are fundamental to software engineering practice.

3.2.1. Knowledge Graphs & Semantic Systems

Recent work in knowledge graph engineering and semantic systems has led to the development of specialized benchmarks for evaluating large language models (LLMs) on tasks such as knowledge extraction, syntax correction, and semantic reasoning. The following paragraphs summarize key contributions in this area, highlighting frameworks for automated evaluation and the breadth of tasks used to probe LLM capabilities.

LLM-KG-Bench [20] introduces a modular benchmarking framework specifically designed to assess LLMs in knowledge graph engineering (KGE) tasks. The framework features automated evaluation modules for tasks such as fixing errors in Turtle files, extracting knowledge from plaintext, and generating synthetic datasets. Empirical results show that while LLMs like GPT-4 and Claude-1.3 outperform earlier models in syntax and extraction tasks, all models struggle with zero-shot knowledge graph generation and show diminishing performance as task complexity increases. LLM-KG-Bench provides a scalable, automated approach for tracking LLM performance in KGE and highlights both the potential and current limitations of LLMs in semantic web applications.

BIG-bench [21] is a large-scale, collaborative benchmark designed to evaluate and extrapolate the capabilities of LLMs across 204 diverse tasks, including linguistics, reasoning, software development, and more. The benchmark was created to address the rapid obsolescence of existing benchmarks as models improve, and it includes contributions from hundreds of researchers worldwide. Key findings indicate that while model performance and calibration improve with scale, even the largest models lag behind human experts on many tasks, and performance can be brittle and highly sensitive to prompt formulation. BIG-bench remains a living benchmark, continually updated to track the evolving capabilities and limitations of LLMs in a wide range of semantic and reasoning tasks.

3.2.2. Code Review & Analysis

Recent research in code review and analysis has produced a variety of benchmarks and frameworks to evaluate large language models (LLMs) on tasks such as code generation, debugging, test generation, and software engineering lifecycle support. The following paragraphs summarize key contributions in this area, highlighting the importance of domain-specific evaluation, continuous benchmarking, and the integration of LLMs into real-world software engineering workflows.

DomainCodeBench [7] introduces a multi-domain benchmark designed to assess LLM performance on code generation tasks across 12 application domains and 15 programming languages. The study reveals that strong performance on general coding benchmarks does not guarantee success in specialized domains, as LLMs often struggle with domain-specific libraries, algorithms, and workflows. Contextual enhancement—such as providing relevant code snippets—significantly improves results, underscoring the need for domain-aware evaluation and prompt design in practical development scenarios.

StackEval [23] presents a comprehensive benchmark for coding assistance, featuring two datasets: StackEval, which covers 925 curated Stack Overflow questions across 25 languages and four task types, and StackUnseen, which evaluates LLMs on emerging technologies and recent content. The benchmark also introduces an LLM-as-a-Judge framework to assess the quality of generated code. Results show that while top models excel on historical data, performance drops on new or unseen

problems, highlighting the need for continuous adaptation and responsible AI development in coding assistance.

Azanza et al. [22] propose a measurement framework for the continuous evaluation of LLM-based test generators in industrial environments. Validated through a longitudinal study at a software consultancy, the framework integrates with industry tools like SonarQube and tracks metrics such as code quality, coverage, and test parameterization. The study demonstrates rapid improvements in LLM-generated test quality over time, but emphasizes the ongoing need for expert oversight, prompt engineering, and systematic re-evaluation as models evolve.

Zhang et al. [24] provide a comprehensive survey of LLM applications in software engineering, categorizing 947 studies across 112 SE tasks and five phases of the software lifecycle. The survey highlights the transformative impact of LLMs on requirements, development, testing, maintenance, and management, while also discussing challenges related to evaluation, security, domain adaptation, and model optimization. The authors advocate for the creation of clean, diverse benchmarks and the integration of explainable AI techniques to advance LLM-based software engineering research and practice.

3.3. Systems Engineering

Systems engineering integrates multiple disciplines to design and manage complex systems throughout their lifecycles. Large language models (LLMs) have emerged as valuable tools in this domain, supporting critical tasks from requirements elicitation to stakeholder communication. Their natural language capabilities help bridge communication gaps between technical and non-technical stakeholders—a persistent challenge in systems engineering. Recent benchmarks like SysEngBench [25] and platinum benchmarks. [27] evaluate LLMs' effectiveness in model-based systems engineering and safety-critical verification. As shown in Table 1, LLMs accelerate documentation, enhance knowledge integration, and improve requirements management, though challenges remain in handling interdisciplinary complexity and ensuring reliability in mission-critical applications.

3.3.1. Systems Architecture

In the context of systems architecture, LLMs are evaluated on their capacity to support the design and orchestration of complex, hierarchical production systems. For example, SysEngBench [25] introduces a novel benchmark for assessing LLMs in systems engineering, with a focus on model-based systems engineering (MBSE) and the transition from document-centric to model-centric approaches. The benchmark covers ten core topic areas using multiple-choice questions derived from lecture slides and human review, and evaluates models such as Llama 2, Mistral, and Orca 2. Results show that while some models perform well in certain areas, challenges remain in requirements analysis and in handling complex, real-world systems engineering tasks. SysEngBench provides a foundation for future research, with plans to expand the benchmark to include more complex questions, multimodal inputs, and human performance comparisons.

3.3.2. Reliability Engineering

Recent work in reliability engineering has produced a range of studies and benchmarks to evaluate the trustworthiness, reliability, and alignment of large language models (LLMs) in high-stakes engineering contexts. The following paragraphs summarize key contributions in this area, highlighting empirical evaluations, the development of platinum benchmarks, and comprehensive trustworthiness frameworks.

Hu et al. [26] examine the application of LLMs in reliability engineering by testing GPT-3.5 and GPT-4 on Certified Reliability Engineer (CRE) exam questions. The study finds that prompt engineering and context-specific instructions significantly improve model performance, with GPT-4 achieving up to 92% accuracy. While LLMs show promise for information extraction, FMEA processing, and code generation, the authors caution against sole reliance on LLMs for safety-critical or regulatory decisions, emphasizing the need for human oversight and collaboration.

Platinum Benchmarks [27] addresses the gap between capability and reliability in current LLM benchmarks by proposing “platinum benchmarks”—carefully curated to minimize label errors and ambiguity. The study demonstrates that even frontier models make simple mistakes on basic tasks, and that existing benchmarks often fail to measure true reliability due to pervasive label errors. The authors advocate for rigorous reliability testing as a standard practice for deployed LLMs, similar to traditional reliability engineering.

Liu et al. [28] present a comprehensive framework for evaluating LLM trustworthiness and alignment, introducing a taxonomy of seven major categories and 29 sub-categories. Their empirical studies reveal that while more aligned models generally perform better, alignment effectiveness varies across trustworthiness dimensions, and even well-aligned models have limitations in fairness, causal reasoning, and robustness. The authors highlight the need for fine-grained, targeted alignment strategies and automated evaluation methods to ensure LLMs meet societal values and ethical standards.

4. Science

The scientific domain presents unique challenges for evaluating large language models, requiring deep domain expertise, multi-step reasoning, and the ability to work with diverse data modalities. This section explores how LLMs are being evaluated and applied across geography, physical sciences, and other scientific disciplines. An overview of all the relevant benchmarks is presented in Table 2.

Table 2. Science Discipline: Domain-Specific Benchmarks.

Domain	Sub-domain	Benchmark	Scale	Task Type	Input Modality	Model	Performance	Key Focus
Geography & Remote Sensing	Earth Observation	TEOChat [29]	N/A	Temporal EO	Image sequence	TEOChat	SOTA	Earth observation
		EarthNets [30]	500+ datasets	Earth obs	Multi-modal	Multiple	N/A	Dataset benchmark
		VLEO-Bench [31]	N/A	EO analysis	Satellite	GPT-4V	Mixed	Scene understanding
	Climate & Environment	STBench [32]	60K QA pairs	Spatio-temporal	Multi-modal	Multiple	N/A	Data mining
	Agricultural Science	AgriLLM [33]	Survey paper	Agriculture	Multi-modal	Multiple	N/A	Farming applications
	Geospatial Analysis	GEOBench-VLM [34]	10K instructions	Geospatial	Multi-modal	GPT-4V	40%	Remote sensing
		Roberts et al. [35]	N/A	Geographic	Multi-modal	GPT-4V	Human-level	Geospatial tasks
		RSUniVLM [36]	1B params	Multi-granular	Multi-modal	RSUniVLM	SOTA	Remote sensing
Physical & Chemical Sciences	Fundamental Science	INS-MM-Bench [37]	2.2K questions	Insurance	Multi-modal	Multiple	N/A	Domain tasks
		MMLU [38]	Varies by Domain	QA	Text	Multiple	Below Expert-Level	Academic knowledge
		ScienceQA [39]	Varies by Domain	QA	Text and Images	GPT3 with CoT	75.17%	High-School QA
		IsoBench [40]	75 per domain	Multimodal QA	Text and Images	Multiple	Text better than Image	General Science
		VisScience [41]	1000 per domain	Multi-modal reasoning	Text and Images	Multiple	38.2% Physics, 47.0% Chemistry	K12 assessment
		GPQA [42]	448 MCQs	Difficult MCQ	Text	Claude 3 Opus	PhD-level (60%)	Graduate-level reasoning
	Molecular & Materials Science	MM-PhyQA [43]	4500 q/a	Q/A	Text and Images	LLaVA-1.5 13b	71.65%	High-school physics
		ChemBench [44]	2700+ q/a pairs	Q/A, Reasoning	Text	Multiple	64% accuracy	General Chemistry
		ChemQA [45]	85k examples	Chemistry reasoning	Text	N/A	N/A	Multiple chemistry tasks
		ChemLLMBench [46]	Varies by task	Multiple tasks	Text	Multiple	Poor vs Chemformer	Molecular reasoning
	Applied Chemistry	SMol-Instruct [47]	3M+ pairs	Instruction QA	Text	Multiple	SOTA	Chemical Reaction
		MaScQA [48]	650 QA	QA	Text	GPT4	62% Accuracy	Materials Knowledge
Environmental Science	Climate & Weather	LLM4Mat-Bench [49]	1.9M pairs	Classification	Multiple	Multiple	N/A	Property Prediction
		PRESTO [50]	3M samples	Multiple tasks	Graphs & Text	Multiple	SOTA	Molecule-text Modeling
		DrugLLM [51]	N/A	Generation	GMR	N/A	N/A	Drug design
	Disaster Response	WeatherBench 2 [52]	1-14 day	Forecasting	Multivariate data	N/A	N/A	Medium-range forecasting
		ClimateQA [53]	254,040 VQA	VQA	Image + Text	Multiple	SOTA	Weather Event Detection
	Pollution Monitoring	WeatherQA [54]	8000 VQA	VQA	Image + Text	Multiple	Needs Improvement	Weather Event Detection
		CLLMate [55]	26,156 QA	Forecast	Multiple	Multiple	Needs Improvement	Event Forecasting
	Biodiversity & Ecology	FFD-IQA [56]	22,422 questions	VQA	Image + Text	Multiple	Needs Improvement	Safety assessment
		DisasterQA [57]	707 questions	MCQ	Text	GPT-4o	85.78%	Damage response
		VayuBuddy [58]	7 years data	Code Generation	Text	Multiple	N/A	Sensor Data Analysis
		Species-800 [59]	800 abstracts	Classification	Text	N/A	N/A	NER
		BiodivNERE [60]	2057 sentences	Classification	Text	N/A	N/A	NER

4.1. Geography & Remote Sensing

The field of geography and remote sensing has seen remarkable advancements in the application of LLMs, particularly in processing and analyzing Earth observation data, climate patterns, agricultural systems, and complex geospatial information. Recent benchmarks demonstrate significant progress in multimodal capabilities, temporal reasoning, and domain-specific applications.

4.1.1. Earth Observation

Earth observation has emerged as a critical application domain for multimodal language models, with several significant benchmarks demonstrating the potential of these models in processing satellite and aerial imagery. The field has seen rapid advancement in both dataset development and model capabilities, particularly in temporal analysis and spatial reasoning tasks.

TEOChat [29] represents a significant breakthrough as the first vision-language assistant designed specifically for temporal satellite imagery analysis. The benchmark introduces TEOChatlas, a comprehensive dataset of 245,210 temporal examples spanning diverse tasks requiring both spatial and

temporal reasoning. TEOChat demonstrates superior performance across seven task categories including temporal scene classification (75.1% accuracy on fMoW RGB), change detection, and question answering, substantially outperforming previous vision-language models and even rivaling specialist models trained for specific tasks. The model's ability to process image sequences enables sophisticated analysis of dynamic phenomena such as disaster impacts and urban development, while its strong zero-shot generalization capabilities highlight the potential of multimodal approaches to enhance earth observation analysis for applications requiring temporal understanding.

EarthNets [30] addresses the fragmentation in earth observation datasets by providing a comprehensive platform and benchmark for evaluating AI models on remote sensing data. The project systematically reviews and analyzes over 500 publicly available earth observation datasets spanning diverse domains including agriculture, land use, disaster monitoring, scene understanding, and climate change. What distinguishes EarthNets is its approach to dataset analysis across four dimensions: volume, resolution distributions, research domains, and inter-dataset correlations. The platform facilitates standardized evaluation of deep learning methods through unified dataset libraries and processing pipelines, effectively bridging the gap between remote sensing and machine learning communities. By enabling fair and consistent model comparisons, EarthNets represents a crucial step toward developing more robust foundation models for earth observation applications.

VLEO-Bench [31] presents a comprehensive benchmark for evaluating Vision-Language Models (VLMs) on Earth observation data, with a particular focus on GPT-4V's capabilities. The benchmark assesses three key areas: scene understanding, localization and counting, and change detection, covering diverse applications such as urban monitoring, disaster relief, land use, and conservation. While GPT-4V demonstrates strong performance in open-ended tasks like location understanding and image captioning, it shows limitations in spatial reasoning tasks such as object localization and counting. The benchmark includes various datasets like RSICD for captioning, BigEarthNet for land cover classification, and xBD for change detection, providing a systematic evaluation framework for VLMs in Earth observation applications. This work highlights both the potential and current limitations of VLMs in processing satellite and aerial imagery, offering valuable insights for future model development in this domain.

The Earth Observation domain demonstrates significant progress in leveraging LLMs for geospatial and remote sensing tasks. Models in this field have showcased their ability to process diverse data types, such as satellite imagery and temporal sequences, providing actionable insights for tasks like environmental monitoring and urban planning. A notable trend is the increasing focus on temporal analysis, where LLMs are used to track dynamic changes over time. Despite these successes, the inherent complexity of geospatial data, including its high dimensionality and the need for precision, often limits the accuracy of existing models. Annotation quality remains another critical bottleneck; while automated annotations reduce costs and time, they may not capture subtle domain-specific nuances, leading to potential inaccuracies. The lack of unified evaluation metrics across geospatial tasks hinders consistent benchmarking and comparison of model performance.

4.1.2. Climate & Environment

STBench [32] introduces a comprehensive framework for evaluating large language models' spatio-temporal analysis capabilities, particularly relevant for climate and environmental applications. The benchmark assesses four key dimensions: knowledge comprehension, spatio-temporal reasoning, accurate computation, and downstream applications, using over 60,000 question-answer pairs across 13 tasks.

The evaluation of 13 state-of-the-art LLMs, including GPT-4o and Gemma, reveals strong performance in knowledge comprehension and spatio-temporal reasoning tasks, with GPT-4o achieving 79.26% accuracy on POI Category Recognition. However, models face challenges in accurate computation tasks. Advanced prompting techniques like in-context learning and chain-of-thought prompting show significant improvements, with ChatGPT's accuracy on POI Identification increasing from 58.64% to 76.30%.

While STBench's comprehensive evaluation framework is valuable for climate research, the lower performance on computation tasks suggests current LLMs need further development for precise environmental data analysis. The benchmark's open-source availability facilitates its adoption for climate and environmental research applications.

4.1.3. Agricultural Science

AgriLLM [33] presents a comprehensive review of multi-modal large language models in agriculture, addressing 11 key research questions (4 general and 7 agriculture-focused) to understand their capabilities, challenges, and future directions. The survey systematically analyzes over 200 papers from multiple scientific databases, focusing on the integration of LLMs in agricultural applications.

The review highlights that multi-modal LLMs significantly enhance agricultural decision-making and image processing efficiency, particularly in areas such as crop monitoring, pest management, and yield prediction. These models demonstrate superior capabilities compared to traditional ML and DL methods, offering improved contextual understanding and predictive analytics. However, challenges remain in data privacy, computational costs, and the need for extensive domain-specific training data.

4.1.4. Geospatial Analysis

GEOBench-VLM [34] introduces a comprehensive benchmark suite specifically designed for evaluating vision-language models (VLMs) on geospatial tasks, featuring over 10,000 manually verified instructions across 31 fine-grained tasks categorized into 8 broad categories including scene understanding, object counting, localization, segmentation, and temporal analysis. The benchmark addresses key challenges in geospatial data analysis such as temporal change detection, large-scale object counting, and tiny object detection, while employing multiple-choice questions to ensure objective and automated evaluation. Evaluation of 13 state-of-the-art VLMs reveals that while existing models show potential, they face significant challenges in geospatial tasks, with the best-performing LLaVA-OneVision achieving only 41.7% accuracy on MCQs, slightly outperforming GPT-4o but still highlighting substantial room for improvement in geospatial-specific capabilities.

Roberts et al. [35] explores the geographic and geospatial capabilities of multimodal large language models (MLLMs), particularly focusing on GPT-4V's performance across various visual tasks. The study reveals that while GPT-4V demonstrates strong performance in recognizing fine details and reasoning about geographic features, it faces challenges in precise localization tasks. The evaluation of multiple MLLMs shows that model performance varies significantly by task type, with Qwen-VL and LLaVA-1.5 often outperforming GPT-4V in localization tasks, while GPT-4V excels in broader geographic understanding and reasoning. The study also highlights the models' tendency to struggle with multi-object images and their varying susceptibility to output format constraints.

RSUniVLM [36] presents a unified vision-language model specifically designed for remote sensing applications, featuring a novel Granularity-oriented Mixture of Experts architecture that enables comprehensive vision understanding across multiple granularity levels while maintaining a compact 1 billion parameter size. The model demonstrates strong performance across image-level tasks (captioning and visual question answering), region-level tasks (visual grounding and referring expression generation), and pixel-level tasks (semantic segmentation), while also supporting multi-image analysis for change detection and captioning. RSUniVLM's unified approach addresses the limitations of existing models in handling fine-grained pixel-level interpretation, which is crucial for applications like land-cover mapping and environmental protection.

INS-MMBench [37] introduces the first comprehensive benchmark for evaluating large vision-language models (LVLMs) in the insurance domain, with particular relevance to geospatial analysis through its coverage of property and agricultural insurance tasks. The benchmark comprises 2.2K multiple-choice questions across 12 meta-tasks and 22 fundamental tasks, systematically organized through a bottom-up hierarchical task definition methodology. Evaluation of 10 LVLMs, including both closed-source (GPT-4o, GPT-4V) and open-source models (LLaVA, BLIP-2), reveals significant performance variations across insurance types, with GPT-4o achieving the highest score of 72.91/100.

The benchmark's structured approach to task definition and evaluation provides valuable insights into LVLMS' capabilities in processing geospatial data for insurance applications, while highlighting challenges in domain-specific knowledge and perception accuracy.

4.2. Physical & Chemical Sciences

The physical and chemical sciences present unique challenges for evaluating large language models, requiring deep domain knowledge, multi-step reasoning, and the ability to work with diverse representations including mathematical equations, chemical structures, and scientific diagrams. Recent benchmarks assess both foundational understanding and specialized capabilities across physics, chemistry, and interdisciplinary domains.

4.2.1. Fundamental Science

One of the earlier benchmark datasets designed to evaluate the world knowledge and problem-solving abilities of large language models (LLMs) was MMLU (Massive Multitask Language Understanding) [38]. This dataset spans a wide range of topics, including foundational subjects in chemistry and physics such as mechanics and electromagnetism. It consists of multiple-choice questions sourced from high school and undergraduate exams. While MMLU has certain limitations typical of early benchmarks, it remains widely used within the research community and continues to evolve through iterative refinements [61].

Building on this, ScienceQA [39] was introduced as a large-scale multimodal benchmark, comprising approximately 21,000 multiple-choice questions across various science disciplines. Each question is annotated with supporting lectures, contextual information, and detailed explanations. A notable contribution of ScienceQA is its integration of chain-of-thought (CoT) reasoning, where the provided explanations simulate multi-step reasoning paths. This approach has been shown to improve LLM performance on science question answering tasks. The dataset is curated from elementary and high school curricula, and a significant portion of the questions are accompanied by images, enabling experiments that combine textual and visual modalities.

More recent benchmarks have focused on increasing the difficulty of evaluation tasks to better assess the capabilities of advanced models. One notable example is GPQA [42], a highly challenging benchmark comprising 448 multiple-choice questions authored by domain experts in biology, physics, and chemistry. The difficulty of the dataset is underscored by the performance gap between experts and non-experts: even individuals with or pursuing PhDs in relevant fields achieve only around 65% accuracy, while highly skilled non-expert annotators reach just 34%—despite having unlimited web access and spending over 30 minutes per question. These “Google-proof” questions are designed to resist superficial search-based strategies. At the time of release, state-of-the-art models also struggled, with a GPT-4-based baseline achieving only 39% accuracy. The authors emphasize the need for improved and scalable oversight mechanisms to match the increasing capabilities of large models.

Another recent addition is VisScience [41], a multimodal benchmark specifically developed to evaluate the visual reasoning capabilities of MLLMs across scientific domains. The benchmark includes 3,000 K-12 level questions that integrate textual and visual modalities, spanning 21 science-related subjects and categorized into five levels of difficulty. Notably, VisScience highlights performance disparities across subjects and models. For instance, Claude 3.5-Sonnet achieved 53.4% accuracy in mathematics, GPT-4o reached 38.2% in physics, and Gemini 1.5 Pro scored 47.0% in chemistry. These results suggest that even leading models face substantial challenges in visual scientific understanding, underscoring the need for further advancements in multimodal reasoning capabilities tailored to the diverse demands of STEM disciplines.

Meanwhile, IsoBench [40] introduces a novel benchmarking framework to evaluate Multimodal Foundation Models on isomorphic representations—distinct formats conveying the same underlying information. A key focus of IsoBench is on scientific domains, particularly physics and chemistry, with 75 questions from each discipline. Each question is presented in multiple modalities, including visual (e.g., diagrams), textual (descriptive captions), and mathematical representations. This setup

allows for a fine-grained analysis of how the form of input affects model performance. IsoBench reveals that current models exhibit a strong bias toward textual inputs. Notably, when evaluated across all IsoBench problems, Claude 3 Opus performs 28.7 percentage points worse when given images instead of equivalent textual descriptions. To mitigate such disparities, the authors propose prompting strategies—IsoCombination and IsoScratchPad—which encourage models to reason across, and translate between, different representations. These methods show measurable improvements, suggesting that aligning multimodal reasoning pathways can enhance model robustness. The dataset construction process is meticulous. For the visual modality, each question includes a figure paired with a textual prompt and multiple-choice options. The corresponding isomorphic textual input was manually annotated: an author provided a precise description of each figure, intentionally avoiding any additional inference beyond the content depicted. A portion of the IsoBench questions are adapted from existing benchmarks, including ScienceQA [39], ensuring continuity and relevance with prior multimodal evaluation efforts.

Focusing more specifically on the domain of physics, MM-PhyQA [43] was developed to benchmark multi-step reasoning in high school-level physics tasks under a multimodal setting. The dataset includes physics questions accompanied by relevant visual inputs—such as diagrams, graphs, or circuit illustrations—and is designed to evaluate models' ability to perform complex reasoning that goes beyond surface-level pattern recognition. To support more advanced inference, the authors also release a CoT-Prompting variant, which incorporates exemplar questions during training to encourage chain-of-thought reasoning. The study provides an in-depth analysis of several factors that impact model performance: the inclusion of additional modalities beyond text, the use of CoT prompting techniques, and the effects of fine-tuning LLMs and LMMs for domain-specific tasks like physics question answering. One of the key contributions is the introduction of Multi-Image Chain-of-Thought (MI-CoT) prompting—a method that enables models to reason over multiple images alongside textual input, thereby more closely mirroring the way humans interpret and solve visual science problems.

For inference, the authors employed zero-shot GPT-4 for baseline predictions and evaluated LLaVA and LLaVA-1.5 models, both of which were fine-tuned on MM-PhyQA to assess performance gains from domain adaptation. For text-only evaluation, they tested both base and fine-tuned versions of Mistral-7B and LLaMA2-7B, comparing results across different prompting and fine-tuning strategies. Among all evaluated approaches, MI-CoT prompting combined with a fine-tuned LLaVA-1.5 13B model achieved the strongest performance, attaining the highest accuracy of 71.65% on the test set and outperforming other models across most evaluation metrics.

The field of chemistry and its related sub-disciplines have also witnessed the emergence of dedicated benchmarks aimed at evaluating LLMs' scientific reasoning and domain expertise. One such benchmark, ChemBench [44], introduces a comprehensive and automated framework designed to assess the chemical knowledge and reasoning capabilities of state-of-the-art LLMs, comparing their outputs directly against those of trained chemists. The benchmark consists of 2,700 question-answer pairs, with each question annotated according to its topic, required skill (e.g., reasoning, calculation, factual knowledge, or intuition), and difficulty level. To establish human baselines, the authors developed a web-based interface and conducted a survey involving domain experts in chemistry. They then evaluated a range of both open- and closed-source LLMs, discovering that top-performing models, in aggregate, outperformed the best-performing human chemists in their study. However, this performance came with caveats: the models were prone to overconfident answers, particularly on simpler or foundational tasks, where basic errors were still prevalent.

Building on the momentum of multimodal reasoning in the chemical sciences, ChemQA [45] introduces a comprehensive multimodal question-answering dataset designed to evaluate the chemistry understanding of LLMs and MLLMs across a range of structured tasks. Inspired by benchmarks such as IsoBench [40] and ChemLLMBench [46], ChemQA targets practical chemistry scenarios that require both symbolic reasoning and visual interpretation. The dataset includes five types of tasks: identifying atomic composition in organic molecules, calculating molecular weights, converting between SMILES

strings and IUPAC names, generating and editing textual descriptions of molecules, and planning retrosynthetic pathways. These tasks collectively represent key challenges in computational chemistry and molecular design. ChemQA consists of approximately 85,000 examples, split into training, validation, and testing sets, and is available through Hugging Face.

ChemLLMBench [46] offers a broad and systematic evaluation of large language models across a diverse set of real-world chemistry tasks. Designed to assess LLM capabilities in practical chemical reasoning, the benchmark spans eight task types, including bidirectional name conversion between IUPAC and SMILES, molecular property prediction (e.g., solubility and toxicity), reaction yield classification, forward reaction prediction, retrosynthesis planning, molecule captioning, reagent selection, and text-driven molecule design. It draws upon widely used datasets such as BBBP, Tox21 [62], PubChem [63], USPTO [64], and ChEBI [65] to construct its task-specific benchmarks. The scale of data varies depending on the task, ranging from thousands of molecules to reactions, and includes a mix of input modalities such as textual descriptions, SMILES strings, chemical formulas, and in some cases molecular graphs or images. Tasks cover both classification and generative formats. In their evaluation of five prominent LLMs—GPT-4, GPT-3.5, Davinci-003, LLaMA, and Galactica—GPT-4 consistently outperforms the others. However, the study finds that while LLMs perform strongly in classification-based or descriptive tasks like molecule captioning, they struggle with generative tasks that require deeper chemical understanding, such as reaction prediction and molecule synthesis.

SMol-Instruct [47] is a large-scale, high-quality dataset specifically designed for instruction tuning in the chemistry domain. It encompasses over three million samples across 14 carefully curated chemistry tasks, providing a robust foundation for training and evaluating large language models in a wide range of chemical reasoning challenges. The tasks span several categories, including name conversion between IUPAC, molecular formulas, and SMILES strings; molecule description through captioning and generation; property prediction across datasets like ESOL, LIPO, BBBP, ClinTox, HIV, and SIDER; as well as reaction tasks such as forward synthesis and retrosynthesis. To assess the effectiveness of instruction tuning, the authors fine-tune several open-source LLMs under the name LlaSMol, identifying Mistral as the most effective base model for chemistry tasks. SMol-Instruct's scale is significant, containing 3.3 million samples involving 1.6 million unique molecules, selected to reflect a broad diversity in molecular size, structure, and property space. The dataset undergoes rigorous quality control, including the removal of low-quality or erroneous samples, careful data partitioning, and the canonicalization of SMILES representations to ensure consistency and reliability across tasks.

4.2.2. Molecular & Materials Science

MaScQA [48] (Materials Science Question Answering) is a benchmark specifically designed to evaluate large language models on undergraduate-level materials science, a field that bridges physics and chemistry. The dataset contains 650 questions categorized by subtopics such as material structure, properties, processing, and phase behavior, with a focus on assessing knowledge of concepts like crystallography and phase diagrams. The authors evaluate GPT-3.5 and GPT-4 using both zero-shot and chain-of-thought prompting, finding that GPT-4 performs best with an accuracy of around 62%. Notably, unlike in many other domains, chain-of-thought prompting did not yield a significant improvement in performance, suggesting that conventional prompting techniques may have limited effectiveness in highly specialized scientific fields like materials science.

Similarly, LLM4Mat-Bench [49] represents an extensive benchmark to date for assessing the capability of large language models in predicting properties of crystalline materials. The dataset comprises approximately 1.9 million crystal structures curated from ten publicly available materials science databases and spans 45 distinct material properties. It supports multiple input modalities, including crystal composition, crystallographic information files (CIFs), and textual descriptions of crystal structures, encompassing 4.7 million, 615.5 million, and 3.1 billion tokens, respectively. The benchmark is used to fine-tune and evaluate various models, such as LLM-Prop and MatBERT, across these input types. Results reveal that general-purpose LLMs often fall short in accurately predicting

materials properties, emphasizing the need for domain-adapted, instruction-tuned models specifically designed for materials science applications.

4.2.3. Applied Chemistry

A recent effort aimed at advancing multimodal modeling in synthetic chemistry is PRESTO [50] (Progressive Pretraining Enhances Synthetic Chemistry Outcomes), a framework designed to bridge the molecule-text modality gap in reaction-centric tasks. Synthetic chemistry, which focuses on designing and executing chemical reactions to produce compounds with specific properties, presents unique challenges for multimodal learning due to the complexity of reaction procedures and molecular representations. PRESTO addresses this by integrating a large-scale dataset of approximately 3 million samples, consisting of synthetic procedure descriptions and molecule name conversions. It provides a structured benchmark for evaluating various pretraining strategies and dataset configurations. The framework supports a range of tasks critical to synthesis, including reaction prediction, reaction condition prediction, reagent selection, reaction type classification, and yield regression. Through progressive pretraining techniques that promote cross-modal alignment and multi-graph understanding, PRESTO demonstrates notable improvements in the performance of multimodal LLMs across these tasks, offering a robust foundation for future work in automated reaction planning and chemical synthesis modeling.

In the domain of drug discovery, the DrugLLM [51] dataset introduces a large-scale, specialized resource designed to train LLMs for molecule generation and optimization. At its core is the Group-based Molecular Representation (GMR), which sequences molecular modifications aimed at improving specific chemical properties. By learning to predict successive molecular edits, DrugLLM models the process of rational drug design. The dataset draws from two well-established sources, ZINC [66] and ChEMBL [67], transforming tabular molecular data into structured “modification paragraphs”—narratives of molecular changes associated with over 10,000 different properties or activities. In total, the dataset includes over 25 million paragraphs and 200 million training samples, making it one of the largest resources of its kind. Empirical results show that DrugLLM exhibits strong few-shot generation capabilities, effectively proposing novel molecules with desired properties using limited input examples, positioning it as a promising tool for accelerating drug design through language-based modeling.

4.3. Environmental Science

Environmental science presents unique challenges for large language models, requiring them to reason about complex natural systems, process multimodal data from diverse sensors and satellites, and support critical decision-making in areas like climate change, disaster response, and ecological conservation. This section examines how LLMs and MLLMs are being evaluated and applied across key environmental domains including weather forecasting, disaster management, pollution monitoring, and biodiversity assessment.

4.3.1. Climate & Weather

WeatherBench 2 [52] is an update to the global, medium-range (1–14 day) weather forecasting benchmark called WeatherBench [68], designed with the aim to accelerate progress in data-driven weather modeling. WeatherBench 2 consists of an open-source evaluation framework, publicly available training, ground truth and baseline data as well as a continuously updated website with the latest metrics and state-of-the-art models. Studies have leveraged this data and used LLMs in weather forecasting. For example, CLIMATELLM [69] captures spatiotemporal dependencies via a cross-temporal and cross-spatial collaborative modeling framework that integrates Fourier-based frequency decomposition with Large Language Models (LLMs) to strengthen spatial and temporal modeling.

ClimateIQA [53] a first meteorological VQA dataset, which includes 8,760 wind gust heatmaps and 254,040 question-answer pairs covering four question types, both generated from the latest climate reanalysis data. The researchers redefines Extreme Weather Events Detection (EWED) by

framing it as a Visual Question Answering (VQA) problem, thereby introducing a more precise and automated solution. Leveraging Vision-Language Models (VLM) to simultaneously process visual and textual data, we offer an effective aid to enhance the analysis process of weather heatmaps. In the creation process, images were processed using SPOT to extract color contours and representative point coordinates. The extracted data were integrated into geographic knowledge bases to retrieve location-specific information. These data, including location, coordinates, and wind speed, were then input into predefined question-and-answer templates, resulting in the generation of 254,040 question-answer pairs.

WeatherQA [54] is a multimodal dataset designed for machines to reason about complex combinations of weather parameters and predict severe weather in real-world scenarios. Collected and processed from the NOAA Storm Prediction Center, the dataset includes 8,511 (multi-image, text) pairs of weather data from severe convective and winter storms together with expert analyses. Each pair contains rich information crucial for forecasting—the images describe the ingredients capturing environmental instability, surface observations, and radar reflectivity, and the text contains in-depth forecast analyses written by human experts. They also evaluate many vision language models (VLMs) in two tasks: (1) multi-choice QA for predicting the affected area and (2) classification of the development potential of severe convection. Results show a substantial gap between the strongest VLM, GPT4o, and human reasoning. They suggest that better training and data integration are necessary to bridge this gap.

4.3.2. Disaster Response

Natural disasters pose significant threats to human life and infrastructure. Multimodal Large Language Models (MLLMs) show promise in enhancing disaster response by processing both textual and visual data to support emergency management. However, given the high-stakes nature of disaster scenarios, rigorous evaluation of MLLMs' capabilities is essential before deployment. Recent benchmarks assess both textual reasoning and visual understanding in disaster contexts, focusing on real-world applicability.

The FFD-IQA (Freestyle Flood Disaster Image Question Answering) benchmark [56] addresses the challenge of visual question answering in flood disaster scenarios. The dataset comprises 2,058 images paired with 22,422 question-answer pairs spanning free-form, multiple-choice, and yes-no question formats. Unlike previous disaster-focused VQA benchmarks, FFD-IQA expands beyond closed-answer spaces to allow for more nuanced responses about disaster conditions. The authors also introduce a Zero-shot VQA for Flood Disaster Damage Assessment (ZFDDA) model, which leverages chain-of-thought (CoT) demonstrations to enhance reasoning capabilities without requiring pre-training on disaster-specific data. Their experimental results demonstrate that implementing CoT prompts significantly improves model accuracy, particularly for complex questions requiring multi-step reasoning. The benchmark specifically focuses on evaluating safety assessment capabilities—determining whether individuals are in danger and whether emergency services are responding adequately—making it particularly relevant for real-time disaster management applications.

Complementing the visual understanding focus of FFD-IQA, DisasterQA [57] provides a comprehensive text-based benchmark for evaluating LLMs' knowledge of disaster response protocols and best practices. The dataset consists of 707 multiple-choice questions curated from six reliable online sources, covering a wide range of disaster response topics. The authors evaluated five leading LLMs (GPT-3.5 Turbo, GPT-4 Turbo, GPT-4o, Llama 3.1, and Gemini) using four different prompting methods: standard prompting, directional stimulus prompting, chain-of-thought prompting, and few-shot prompting. Their results indicate that GPT-4o achieved the highest accuracy of 85.78% when using chain-of-thought prompting, demonstrating promising capabilities but also highlighting room for improvement before these models can be reliably deployed in critical disaster response scenarios. Notably, the study revealed that certain types of disaster response questions were consistently answered incorrectly across all models, suggesting fundamental knowledge gaps that need to be addressed. The

benchmark emphasizes the importance of not just accuracy but also appropriate confidence calibration, as overconfident yet incorrect responses could be particularly harmful in emergency situations.

4.3.3. Pollution Monitoring

The effective monitoring of environmental pollutants represents a critical application domain for multimodal language models, with direct implications for public health and environmental policy. Air pollution alone accounts for approximately 6.7 million deaths annually, making timely access to accurate pollutant data essential for mitigating exposure risks. While government agencies routinely collect extensive environmental sensor data, there exists a significant barrier between this raw information and various stakeholders who could benefit from it. VayuBuddy [58] addresses this challenge through an LLM-powered chatbot system that bridges the gap between technical sensor data and diverse stakeholders. Using seven years of air quality data from Indian government sensors, the benchmark evaluates models' ability to generate Python code that analyzes structured sensor data in response to natural language queries about air pollution. The system assesses various LLMs including Llama, Mixtral, Codestral, and Gemma across 45 diverse question-answer pairs. While current models show promising performance in generating accurate code for data analysis, significant challenges remain in handling complex queries and ensuring consistent performance across different types of environmental analysis tasks. The benchmark demonstrates the potential of MLLMs to democratize access to environmental data by enabling stakeholders to derive meaningful insights without specialized technical expertise.

4.3.4. Biodiversity & Ecology

The monitoring and analysis of biodiversity data is crucial for ecological research and conservation. With rapidly growing scientific literature in this field, automated tools are needed to extract information about species and their relationships from text. Named Entity Recognition (NER) and Relation Extraction (RE) face challenges due to complex taxonomic nomenclature and specialized vocabulary. MLLMs show promise in transforming biodiversity informatics by automating information extraction from scientific literature and enabling comprehensive ecological analysis.

The Species-800 (S800) corpus [59] provides a gold-standard dataset for evaluating named entity recognition of species and other taxonomic mentions in biomedical text. The corpus comprises 800 manually annotated PubMed abstracts spanning eight taxonomic categories, including viruses, bacteria, plants, fungi, and various animal groups. Unlike previous taxonomic datasets, S800 represents a diverse range of organisms and publication sources, offering insights into which types of taxonomic names are particularly challenging for models to identify correctly. The benchmark was developed to support the SPECIES tagger, a dictionary-based approach to taxonomic NER that combines the NCBI Taxonomy with expanded synonyms and variant forms. The authors' evaluation demonstrates that their approach achieves comparable accuracy to existing tools while being more than an order of magnitude faster, making it suitable for processing large document collections. The S800 corpus reveals category-specific performance differences, with virus names posing particular challenges compared to other taxonomic groups. This benchmark not only supports the development of more efficient taxonomic NER tools but also facilitates the creation of resources like ORGANISMS, a web-based platform that provides access to pre-computed species tagging results for the entire Medline database.

Expanding beyond species identification, the BiodivNERE dataset [60] provides a more comprehensive benchmark for both named entity recognition and relation extraction in biodiversity literature. This gold-standard corpus includes 2,057 manually annotated sentences covering six entity classes and 17 relation types derived from biodiversity research questions and ontologies. The dataset was constructed through a rigorous process involving biodiversity experts, with annotation guidelines designed to ensure consistency and accuracy. The annotation schema incorporates a wide range of entities relevant to biodiversity research, including not just organisms but also habitats, environmental features, materials, processes, and qualities. The relation types capture important ecological and evolutionary connections between these entities, such as "lives_in" relationships between organisms

and habitats or “affects” relationships between processes and organisms. By providing benchmarks for both NER and RE tasks, BiodivNERE enables the development of more sophisticated tools for analyzing the rapidly growing biodiversity literature. The authors’ evaluation reveals high inter-annotator agreement and provides detailed statistics on entity and relation distributions, offering valuable insights into the challenges of information extraction in this domain. Together with S800, these benchmarks highlight the complexity of biodiversity information extraction and provide essential resources for training and evaluating MLLMs in ecological applications.

5. Technology

Technology domains represent critical application areas where MLLMs are transforming capabilities across industries. These domains leverage multimodal reasoning to enhance perception, decision-making, and human-machine interaction in complex technical environments. An overview of all the relevant benchmarks is presented in Table 3.

Table 3. Technology Discipline: Domain-Specific Benchmarks.

Domain	Sub-domain	Benchmark	Scale	Task Type	Input Modality	Model	Performance	Key Focus
Computer Vision & Autonomous Systems	Visual Understanding	Rank2Tell [70]	N/A	Importance ranking	Multi-modal	N/A	N/A	Object importance
		NuInstruct [71]	91K pairs	Multi-view QA	Video + BEV	BEV-InMLLM	+9%	Holistic understanding
		NuScenes-QA [72]	460K QA pairs	Multi-modal VQA	Image + LiDAR	Multiple	N/A	Multi-frame VQA
		Cambrian-1 [73]	20+ encoders	Vision-centric	Multi-modal	Cambrian-1	SOTA	Visual grounding
	Autonomous Driving	HAZARD [74]	3 scenarios	Decision making	Multi-modal	LLM agent	N/A	Disaster response
		LLM4Drive [75]	Survey paper	System review	Multi-modal	Multiple	N/A	System architecture
		WTS [76]	1.2k events	Video analysis	Video + Text	VideoLLM	N/A	Pedestrian safety
		DRAMA [77]	17,785 scenarios	Risk assessment	Video + Objects	N/A	N/A	Risk localization
		GenAD [78]	2000+ hours	Video prediction	Video + Text	GenAD	N/A	Generalized prediction
		Reason2Drive [79]	600K+ pairs	Chain reasoning	Video + Text	VLMs	N/A	Interpretable reasoning
		DriveLM [80]	N/A	Graph VQA	Multi-modal	DriveLM-Agent	N/A	End-to-end driving
		NuPrompt [81]	35,367 prompts	Object tracking	Multi-view	PromptTrack	N/A	Language prompts
Robotics & Automation	Robot Control	MMRo [8]	N/A	Manufacturing	Multi-modal	Multiple	N/A	Manufacturing
	Process Automation	DesignQA [9]	N/A	Design	Multi-modal	Multiple	N/A	Design
Blockchain & Cryptocurrency	Smart Contract Analysis	Web3Bugs [82]	Multiple versions	Security classification	Code (Solidity)	GPT-4 (prompted)	60% false-positive reduction	Bug detection
		LLM-SmartAudit [83]	Two datasets	Multi-agent audit	Code + Documentation	LLM-SmartAudit	Outperforms conventional tools	Comprehensive auditing
	Market Analysis	ACFEX [84]	118 contracts	Auto-repair	Code (Solidity)	GPT-4 + RBAC	94.9% fix rate	Patch generation
		CryptoNews [85]	3,200 articles	Sentiment analysis	News text	FT-GPT-4	86.7% accuracy	Market sentiment
		Ethereum Prices [86]	5 years daily	Price forecasting	Time-series	GPT-2 (FT)	SOTA MSE	Short-term pred.
Blockchain & Cryptocurrency	LLM-Trading [87]	Multi-modal	Strategy reasoning	Text + Indicators	GPT-4 agent	jPnL in sim.		Fact-subjectivity reasoning
	Fraud Detection	BLOCKGPT [88]	M txns	Anomaly detection	Txn graph	BLOCKGPT	40% attack detection	Real-time monitoring
	Governance	Crypto Legal Cases [89]	SEC files	Legal reasoning	Case text	GPT-3.5	Mixed	Compliance assistance

5.1. Computer Vision & Autonomous Systems

Autonomous systems rely on computer vision to understand their environment and make quick decisions. They use tasks like object detection, scene segmentation, and predicting movement to stay aware and safe. With the rise of multimodal large language models (MLLMs), these systems can now combine images and text to improve reasoning, planning, and interaction with humans.

Recent research shows that foundational computer vision tasks such as semantic segmentation, object tracking, and depth estimation are being reimaged for autonomous driving using grid-centric and multimodal approaches. Furthermore, by fusing LiDAR, video, and natural language, Multimodal Large Language Models enable interpretable, instruction-based decision-making for autonomous systems [90,91]. This section is divided into two key subdomains: Visual Understanding and Autonomous Driving, each highlighting a different facet of how computer vision integrates into autonomous systems through MLLMs.

5.1.1. Visual Understanding

Visual understanding is foundational for deploying Multimodal Large Language Models (MLLMs) in autonomous driving, where systems must interpret complex, dynamic scenes using visual, spatial, and temporal cues. While early benchmarks focused on general-purpose Visual Question Answering (VQA), recent efforts prioritize domain-specific reasoning, particularly concerning safety, planning, and contextual awareness. Emerging datasets now evaluate how effectively MLLMs can ground language in real-world driving scenes.

A key shift involves an emphasis on spatial-temporal reasoning. NuScenes-QA introduces a large-scale benchmark with over 460,000 question-answer pairs built from RGB and LiDAR inputs across multiple frames. It probes dynamic object tracking, spatial relationships, and causal inference in traffic scenarios [72]. Building on this, NuInstruct adopts an instruction-driven format with 91,000 prompts

covering 17 driving subtasks, incorporating multi-view video, temporal dynamics, and bird's-eye-view (BEV) representations. The benchmark includes the BEV-InMLLM model, which aligns BEV features with LLM reasoning to simulate human-like decision-making logic [71]. Together, these benchmarks highlight a shift from static image-based VQA to structured, temporally-aware evaluations.

Another major advancement focuses on agent prioritization and explainability. Rank2Tell presents a multimodal dataset where models must not only identify traffic participants but also rank them by importance and provide natural language justifications. It evaluates whether models can reason about risk, attention, and intent, mirroring human driver cognition and interpretability needs [70].

In addition, Cambrian-1 focuses on the internal visual representation quality of MLLMs rather than downstream performance. It converts classical computer vision tasks into VQA format and introduces the Spatial Vision Aggregator (SVA) to fuse high-resolution visual features into LLMs. Results show that visual grounding—rather than language reasoning—is often the limiting factor in MLLM performance, especially in tasks requiring spatial disambiguation or high-detail understanding [92].

Collectively, these benchmarks reflect a trend toward more holistic and interpretable visual understanding spanning perception, spatial reasoning, and explanation. However, current models still struggle with BEV-based and multi-view reasoning, high-resolution perception, and agent prioritization. Evaluation protocols remain inconsistent, hindering cross-benchmark comparisons. Table 3 summarizes the design goals and task types of each benchmark.

Going forward, unified evaluation strategies combining spatial fidelity, multi-agent reasoning, and instruction-based interaction are needed. Innovations like modular vision encoders, BEV fusion, and justification modules may be essential for real-world deployment of MLLMs in safety-critical driving environments.

5.1.2. Autonomous Driving

The integration of Multimodal Large Language Models (MLLMs) into autonomous driving systems has led to growing interest in evaluating how these models perceive, reason, and act within highly dynamic and safety-critical environments. Existing benchmarks explore a range of capabilities, from low-level perception to high-level reasoning and decision-making, offering diverse perspectives on how MLLMs interact with visual and textual driving data. Rather than focusing on single datasets or metrics, recent works emphasize the need for holistic evaluation frameworks that test the robustness, explainability, and generalization ability of MLLMs under real-world constraints.

Several benchmarks focus on fine-grained perception and the ability of MLLMs to generate semantically meaningful representations of complex traffic scenes. The WTS dataset offers a pedestrian-centric benchmark that combines multi-view videos with synchronized 2D/3D spatial data, textual event descriptions, and 3D gaze annotations to capture nuanced pedestrian and vehicle behaviors [76]. It introduces LLMScorer, a semantic similarity-based evaluation metric tailored to long, structured video captions, overcoming the limitations of traditional n-gram metrics. This work highlights the importance of understanding human behaviors in real-time, particularly when interpreting intent in dense, urban scenarios. Complementing this, the DRAMA dataset introduces a joint framework for risk localization and captioning, where annotated driving scenarios focus on identifying critical objects and explaining the reasoning behind risk predictions [77]. DRAMA pushes beyond classification by integrating visual question answering with natural language reasoning, providing both open-ended and structured evaluation formats. These datasets mark a shift from pure perception to context-aware safety reasoning, bridging the gap between detection and interpretable understanding.

A second group of benchmarks focuses more explicitly on causal reasoning and structured decision processes. The Reason2Drive benchmark offers over 600,000 video-text pairs that model the full driving pipeline—perception, prediction, and reasoning—through chain-of-thought annotations derived from datasets such as nuScenes, Waymo, and ONCE [79]. It proposes a novel evaluation metric, ADRScore, which captures the logical quality of reasoning chains beyond surface-level textual similarity. Similarly, DriveLM introduces Graph Visual Question Answering (GVQA) as a proxy

for evaluating object-centric reasoning in autonomous driving, incorporating perception, behavior, and planning questions linked through graph-structured dependencies [80]. Unlike single-turn QA datasets, DriveLM supports multi-step logic grounded in object interactions and scenario evolution. Together, these benchmarks reveal a shift toward evaluating models not only on recognition accuracy but also on their ability to construct interpretable, stepwise justifications for driving decisions—an essential requirement for deployment in high-stakes applications.

Another set of works explores the use of natural language prompts to drive perception and planning. NuPrompt extends the nuScenes dataset with over 40,000 prompt-object trajectory pairs, allowing evaluation of MLLMs' ability to ground language instructions to multi-object, multi-view 3D scenes [81]. Unlike earlier prompt-based vision tasks that focus on static images or individual objects, NuPrompt supports object tracking across frames and views, enabling evaluation of temporal and spatial grounding at scale. This benchmark introduces a new class of evaluation tasks that blend natural language interaction with multimodal tracking and planning. GenAD further expands this vision by introducing a generative video prediction model trained on more than 2,000 hours of driving footage, called OpenDV-2K [93]. The model predicts future driving scenes given textual instructions and past observations, enabling zero-shot generalization to new environments, weather conditions, and sensor configurations. By treating driving as a predictive task, GenAD demonstrates that generative modeling can support both action planning and environment simulation across unseen conditions.

Beyond perception and planning, the HAZARD benchmark introduces dynamic environments such as fires, floods, and wind, challenging MLLMs to adapt to non-static, unfolding disasters [74]. This simulation-based benchmark pushes the boundaries of embodied AI by requiring models to perform real-time decision-making in unstable settings. An LLM API is provided to convert raw visual and memory data into semantic text representations, allowing models to plan and act through natural language-based policies. This opens a new line of research into embodied reasoning under temporal stress and environmental uncertainty, areas often overlooked in conventional benchmarks. Finally, the LLM4Drive survey offers a meta-perspective on this evolving field, identifying key trends, limitations, and future opportunities in applying LLMs to autonomous systems [75]. It discusses how modular perception-prediction-planning pipelines are increasingly being replaced by unified, prompt-driven architectures capable of open-world reasoning, few-shot learning, and contextual adaptation. The survey also raises concerns about interpretability, sim-to-real generalization, and the lack of standardized evaluation metrics across tasks.

Together, these benchmarks form a comprehensive and multi-layered foundation for evaluating MLLMs in autonomous driving. While they vary in focus from pedestrian intent detection to graph-based reasoning and disaster-aware planning, they all highlight the need for interpretable, robust, and generalizable models. A clear research gap persists in unified evaluation protocols and cross-dataset comparability, particularly in reasoning accuracy and safety-critical contexts. As LLMs and MLLMs become integral to autonomous systems, future work must address these limitations by integrating standardized, high-stakes, and explainable evaluation practices that reflect real-world complexity and operational constraints.

5.2. Robotics & Automation

Multimodal Large Language Models (MLLMs) are playing a bigger role in robotics and automation because they can combine language, vision, and reasoning. In robotics, they help plan and understand tasks, while in automation, they support decision-making using visual and language inputs. However, to properly evaluate their performance, we need strong, domain-specific tests that go beyond basic results. This section focuses on two key areas—robot control and process automation, each with different needs for perception, planning, and safety.

5.2.1. Robot Control

Recent efforts to assess the viability of MLLMs as “brains” for in-home robots have led to the development of the MMRo benchmark, which provides a structured evaluation of MLLMs across four

robotics-specific dimensions: perception, task planning, visual reasoning, and safety [8]. Unlike traditional VQA benchmarks that focus on general knowledge, MMRo targets domain-relevant scenarios such as object manipulation, spatial awareness, and hazard recognition. Each task is framed through real-world or synthetic images and paired with both multiple-choice and open-ended questions, allowing fine-grained analysis of model behavior.

Experimental results reveal that while top-performing models like GPT-4V and Gemini-Pro excel in high-level reasoning tasks (e.g., object function recognition or plan sequencing), they struggle with basic perceptual capabilities such as recognizing shapes or identifying object materials. These deficits raise concerns about their reliability in real-world deployment. The benchmark also identifies significant inconsistencies across models in safety-related tasks. For example, correctly identifying when an object is hot or sharp is critical for safe manipulation and is a challenge for both commercial and open-source MLLMs. This indicates that while MLLMs may enable abstract reasoning, they currently lack the precision and robustness required for embodied control. The MMRo benchmark thus serves not only as an evaluation tool but also as a diagnostic framework, highlighting specific failure modes that must be addressed for MLLMs to serve as viable cognitive cores in robotics [8].

5.2.2. Process Automation

In the broader field of automation, MLLMs are also being evaluated for tasks like Earth observation, disaster monitoring, and environmental assessment. A recent benchmark evaluates GPT-4V on high-resolution satellite imagery to test its performance on scene understanding, object counting, and change detection [31]. The study focuses on zero-shot generalization—an essential capability for real-world automation—since models are not fine-tuned on remote sensing data. GPT-4V performs well on descriptive tasks such as image captioning and landmark recognition but shows limited effectiveness in spatially demanding tasks like counting buildings or identifying changes in flood-prone regions.

The benchmark also highlights challenges in grounding language to dense, structured data. For instance, while GPT-4V can describe scenes fluently, it often misinterprets spatial relationships and object quantities, shortcomings that would severely impact downstream decision-making in infrastructure or crisis management. Although promising as natural language front-ends for automated systems, MLLMs currently fall short of the precision required for scalable, real-world automation. The study emphasizes the need for better alignment between language models and geospatial reasoning modules to ensure safe and actionable outcomes.

5.3. Blockchain & Cryptocurrency

Blockchain and cryptocurrency technologies present unique challenges for multimodal AI systems, requiring specialized understanding of code, finance, and decentralized systems. This section examines how MLLMs are being evaluated and applied across smart contract security, market analysis, fraud detection, and governance tasks.

5.3.1. Brief Introduction

The blockchain and cryptocurrency domain poses unique challenges and opportunities for applying large language models. Smart contracts underpin decentralized finance and applications, carrying billions in value and demanding high security. Likewise, cryptocurrency markets are heavily driven by public sentiment and complex token economies. LLMs have begun to serve as powerful tools in this domain, from auditing smart contract code to interpreting market sentiment, due to their human-like language and code understanding [82,88]. This section surveys how both text-only and multimodal LLMs are being leveraged across blockchain security, crypto market analysis, and decentralized decision-making. We highlight key application areas – smart contract analysis, sentiment forecasting, fraud detection, and governance – and discuss emerging trends and challenges.

5.3.2. Smart Contract Analysis and Auditing

One critical application of LLMs in blockchain is **smart contract auditing**. Smart contracts are self-executing programs (often in Solidity) that manage digital assets, and vulnerabilities in them have led to major financial losses [82]. Traditional static analyzers and fuzzers detect some bugs but often miss complex logic flaws. LLMs offer a new angle by understanding code semantics and even natural-language documentation of contracts. Early studies using GPT-3.5/4 to detect Solidity bugs found encouraging recall but alarmingly high false-positive rates. For instance, one analysis noted that GPT-4-based detectors achieved relatively high recall on known vulnerability types but with precision trade-offs, especially on newer Solidity versions [82]. Recent research tackles these limits: carefully engineered prompts can cut false alarms by over 60%, and fine-tuning or combining LLMs with symbolic analysis significantly boosts accuracy. Novel frameworks like **LLM-SmartAudit** employ multiple specialized LLM agents in a cooperative manner to audit code, outperforming all conventional tools on broad vulnerability benchmarks [83]. These LLM-driven auditors not only catch standard bugs (reentrancy, arithmetic errors) but even complex logic or business rule violations that rule-based scanners overlooked. Beyond detection, LLMs have also been used to **repair vulnerabilities** by suggesting fixes in code. Early results show GPT-4 can propose patches for certain security bugs (e.g. access control flaws) when guided with the right context [88]. Likewise, developers are using code-aware LLMs (e.g. GitHub Copilot or ChatGPT) to *generate* smart contracts and unit tests. This accelerates development but comes with the risk of the model introducing subtle bugs [82]. Overall, LLMs in smart contract analysis demonstrate **promising capabilities** in reasoning about code and security; the trend is toward hybrid approaches (LLM + static analysis + fuzzing) to mitigate errors. Key challenges remain in keeping LLMs up-to-date with evolving languages (e.g. Solidity changes that confused earlier models) and guaranteeing reliability to a level suitable for high-stakes financial code.

5.3.3. Market Analysis and Sentiment Forecasting

Cryptocurrency markets are notoriously volatile and driven by news, social media, and investor sentiment. LLMs have begun playing a role in **crypto market analysis** by extracting and forecasting sentiment signals from text. Financial NLP was an early success area for domain-specific language models (e.g. FinBERT for stock sentiment), and similar approaches are now applied to crypto [94]. A recent study fine-tuned GPT-4 on **cryptocurrency news articles** to classify sentiment, achieving ~86.7% accuracy – slightly surpassing finance-tuned BERT models (FinBERT ~84.3%) on a test set of news headlines [85]. Such models can discern subtle tone in news (e.g. regulatory announcements or macroeconomic signals) and predict short-term market reactions. In practical terms, positive or negative sentiment scores produced by LLMs serve as features for trading algorithms or risk management, often improving predictive power over using price data alone. Beyond news, researchers are leveraging LLMs to **analyze social media** (tweets, Reddit discussions) for crypto sentiment, which is a multimodal challenge: understanding slang, memes, and the context of discussions requires advanced language comprehension.

LLMs are also being explored as direct **time-series forecasters** for crypto prices. In a novel approach, Makri *et al.* (2025) repurposed a pre-trained language model for Ethereum price prediction by treating historical price sequences as the input “language” [86]. By freezing most of the language model’s layers and fine-tuning it on a limited sequence of price data, they achieved state-of-the-art accuracy in short-term forecasting (outperforming classic statistical models on MSE and RMSE). This suggests LLM architectures can capture complex temporal patterns, possibly learning latent market dynamics. Recent research on fact-subjectivity-aware reasoning agents shows that GPT-4–powered analytics can track sentiment shifts across multiple social platforms, providing real-time alerts to fear-or-hype cycles [87]. Moving forward, researchers suggest **multimodal fusion** of textual sentiment with numerical price data to further improve crypto forecasts. An LLM could, for example, read news headlines and concurrently analyze price charts, combining qualitative and quantitative signals—an approach aligned with how human analysts operate. The trend in this subdomain is

toward **integrating diverse data modalities** (news, social chatter, on-chain metrics, price history) using LLMs as the unifying analytical engine. Key challenges include the need for up-to-date information (crypto news cycles are rapid), handling the noise and intentional manipulation in social media data, and ensuring that models do not learn spurious correlations. Nevertheless, early benchmarks show that domain-tuned LLMs can significantly enhance crypto sentiment analysis and market prediction accuracy [85].

5.3.4. Fraud Detection and Anomaly Detection

Blockchain's open ledger provides a rich data source for analyzing transactions, but spotting illicit or fraudulent activities in massive transaction streams is a needle-in-haystack problem. LLMs have emerged as powerful **anomaly detectors** in this context by learning the normal patterns of blockchain transactions and flagging outliers. Unlike rule-based fraud systems (which rely on known signatures of scams), an LLM can generalize from contextual cues and detect **previously unseen attack patterns** [88]. For example, *BLOCKGPT* is a Transformer-based model trained on millions of Ethereum transactions (with a special focus on known attack cases). It operates as an IDS (Intrusion Detection System) for blockchain, monitoring incoming transfers and smart contract calls in real-time. Impressively, *BLOCKGPT* was able to correctly identify 49 out of 124 real exploit transactions among the top anomalies in testing, catching complex attacks that would evade simpler detectors. Moreover, it maintained high throughput – analyzing ~2,284 transactions per second – indicating viability for deployment in live networks. The success of such models comes from a **proactive, learning-based approach**: the LLM encodes a broad notion of “normal” vs “suspicious” behavior rather than any single indicator. Other efforts use LLMs to examine blockchain data for money laundering patterns, Ponzi schemes, or phishing attempts by interpreting transaction histories as narratives. Multimodal extensions are also possible here (e.g. combining on-chain data with darknet forum text where criminals coordinate).

A related application is **log analysis in distributed networks**, where LLMs parse system logs or protocol messages to detect anomalies or predict failures. For instance, *LogGPT* and similar models have been applied to general IT system logs, demonstrating that an LLM can outperform traditional anomaly detectors by understanding semantic context in events [95,96]. In blockchain, this means beyond financial fraud, LLMs can help spot network malfunctions or security breaches by reading the sequence of node logs or consensus messages. The trend in anomaly detection is leaning toward *real-time LLM-driven monitors* that continuously learn from new data (staying adaptive to new hacking strategies). Challenges include scalability (running large models continuously can be costly, though lighter fine-tuned models or distillations are being studied) and the availability of labeled anomalies for training. There is also a risk of adversarial attempts to fool LLM-based detectors, prompting research into robustifying these models for security. Nonetheless, early deployments like *BLOCKGPT* highlight that LLMs bring a **significant advantage in flexibility and depth of understanding** for blockchain fraud detection, potentially making decentralized ecosystems safer.

5.3.5. Decentralized Governance and Decision-Making

Decentralized communities (e.g. DAO governance forums, crypto social platforms) generate large volumes of unstructured text – proposals, discussions, votes – which LLMs are well-suited to analyze and support. One emerging area is using LLMs as **participants or assistants in blockchain governance**. Research by Trozze *et al.* explored GPT-3.5's ability to assist in legal reasoning for crypto-related cases, like identifying regulatory violations in a token sale scenario [89]. The LLM could correctly suggest some applicable laws, though it missed others, highlighting that current models still struggle with complex legal logic. Interestingly, the same work showed ChatGPT can draft legal documents (e.g. a lawsuit complaint) of comparable quality to human lawyers – mock jurors could not distinguish AI-written complaints from human-written ones. This suggests LLMs could aid in writing proposals, bylaws, or explaining governance policies in crypto communities, improving efficiency. Similarly, **community moderation** is being augmented by LLMs: Axelsen *et al.*

(2023) demonstrate that ChatGPT-based systems can enforce forum rules and filter toxic content in online communities [88]. In a DAO context, this means automating the removal of spam or off-topic posts and summarizing member comments, ensuring healthier deliberation. Another study proposes an **LLM-driven blockchain governance framework**, where an LLM helps classify and route proposals, and even provides recommendations by drawing on past proposals and outcomes. These AI “advisors” in governance could democratize knowledge – e.g. a DAO member could query an LLM about the implications of a complex tokenomics change and get an accessible summary rather than reading a 50-page whitepaper.

While promising, **decentralized decision-making support** via LLMs faces challenges. One issue is **domain understanding** – models must grasp technical and economic context (for example, the nuances of a liquidity pool’s tokenomics) to give sound advice. Misunderstandings could lead to flawed recommendations. There are also concerns about trust and autonomy: current research cautions that LLM assistants should remain advisory, not autonomous decision-makers, until they prove reliability [88]. Bias is another worry; an LLM trained predominantly on English crypto discussions might undervalue perspectives from other language communities, affecting truly global projects. Privacy and confidentiality matter as well, especially for legal and financial advice – models need to handle sensitive governance data without leaking. Despite these hurdles, the trend is that **LLMs are increasingly integrated into the workflows of crypto communities** – from drafting governance proposals and smart contracts, to providing real-time legal/compliance feedback, to guiding newcomers through educational tools (e.g. *GPTutor* for blockchain finance FAQs). In sum, LLMs act as amplifiers of human expertise in decentralized settings, but careful oversight is required. As models improve in factual accuracy and context-awareness, we can expect them to take on larger supportive roles in decentralized decision-making, potentially enabling more informed and inclusive governance.

5.3.6. Section Conclusion

General trends: Across blockchain and cryptocurrency applications, LLMs are proving to be versatile allies – excelling at synthesizing code and language, which is crucial in a domain that blends software with finance and law. We observe a surge of domain-specific LLM efforts in 2023–2025 tackling everything from smart contract security to market prediction. A common theme is integration: rather than using LLMs in isolation, the best results come from combining them with other tools (e.g. static analyzers in security, or financial models in trading) and modalities (text, code, numeric data). Multimodal LLMs, which can handle not just text but also sources like transaction graphs or time-series, are on the horizon to address the rich data in crypto ecosystems.

Challenges: Despite progress, significant challenges remain. Ensuring **accuracy and reliability** is paramount – a single hallucinated output in a contract audit or an erroneous trading signal can cause large losses. This drives the need for benchmark datasets and rigorous evaluation: as noted, many studies currently use bespoke datasets (Table 6), making it hard to compare models. Creating standardized, community-accepted **benchmarks for LLM4Crypto** tasks (for vulnerability detection, fraud detection, etc.) will be important for measuring progress. Scalability is another concern: blockchain data is continuously growing, and LLM inference can be costly; efficient fine-tuning and model compression will be needed for practical deployment. **Privacy and ethical considerations** also come to the forefront. In decentralized finance, an LLM might inadvertently reveal sensitive information from training data, or contribute to market manipulation if misused for fake news generation – calling for careful policy and perhaps on-chain verification of AI outputs.

Nonetheless, the **future outlook** is optimistic. LLMs are rapidly improving their capacity to reason about structured data (like code and transactions) and can be expected to become even more accurate with domain adaptation. We anticipate LLMs playing a key role in **automated blockchain development lifecycles** – for example, an AI assistant that can write a smart contract, formally verify it, deploy it, and monitor its transactions, all through natural language instructions. In crypto finance, LLMs might evolve into personalized financial advisors that understand a user’s portfolio and the market sentiment simultaneously. The interdisciplinary nature of this domain means collaboration

between AI researchers, security experts, economists, and the crypto community is vital. In conclusion, LLMs are catalyzing a new wave of intelligence in blockchain technology, **enhancing security, transparency, and decision-making**. As research converges on better benchmarks and methods, we move closer to realizing trustworthy AI-driven systems in decentralized finance – a step that could bolster confidence and accelerate innovation in the blockchain space.

6. Mathematics

Mathematics benchmarks evaluate LLMs’ capabilities in formal reasoning, theorem proving, and quantitative analysis across various mathematical domains. These benchmarks test both computational accuracy and the depth of conceptual understanding. An overview of all the relevant benchmarks is presented in Table 4.

Table 4. Mathematics Discipline: Domain-Specific Benchmarks.

Domain	Sub-domain	Benchmark	Scale	Task Type	Input Modality	Model	Performance	Key Focus
Mathematical Reasoning	Problem Solving	MathVerse [97]	15K samples	Visual math	Image + Text	GPT-4V	N/A	Diagram interpretation
		HARDMath [98]	366 problems	Graduate-level	Text	GPT-4	43.8%	Advanced reasoning
		MMIQC [99]	N/A	Competition math	Text	Qwen-72B-MMIQC	45.0%	Question composition
		MathChat [100]	N/A	Interactive	Text + Dialogue	N/A	N/A	Multi-turn reasoning
		GSM-MC [101]	N/A	Multiple-choice	Text	Multiple	30x faster	Efficient evaluation
		PROBLEMATHIC [102]	N/A	Robustness	Text	Llama-2	+8%	Noise handling
Formal Methods	Theorem Proving	LeanDojo [103]	98,734 theorems	Theorem proving	Text (code)	GPT-4	N/A	Formal verification
Statistical Analysis	Data Analysis	KnowledgeMath [104]	1,259 problems	Finance MWP's	Text + Tables	GPT-4	45.4%	Domain knowledge

6.1. Mathematical Reasoning – Problem Solving

MATHVERSE is a benchmark specifically designed to evaluate the mathematical reasoning capabilities of language models in situations that involve interpreting complex visual content, such as diagrams and geometric figures [97]. Unlike the traditional benchmarks that rely only on text-based prompts, MATHVERSE focuses on multimodal reasoning, where both textual descriptions and visual elements must be synthesized to solve complex problems. This kind of reasoning mirrors the real-world mathematical problem-solving skills exhibited by learners, where visuals often play a crucial role. The benchmark’s structure highlights a significant gap in current large language model design, where many of these models underperform on visual-based tasks, indicating a lack of robust visual understanding in even the most advanced systems.

To support a deeper analysis of how models handle visual versus textual reasoning, MATHVERSE introduces six controlled variations of each problem by iteratively modifying or removing textual and diagrammatic elements. This design enables researchers to isolate the contributions of each modality to the model’s final answer. Interestingly, the evaluation reveals that some LLMs achieve higher accuracy when diagrams are excluded, suggesting that these models may be relying on textual cues and redundancies rather than truly interpreting the visual information. This observation raises concerns about the depth of multimodal reasoning in current models and emphasizes the need for improved integration of visual understanding within language models to handle authentic, diagram-rich mathematical tasks effectively [97].

HARDMath [98] represents a significant step forward in evaluating the mathematical reasoning capabilities of large language models by shifting the focus from elementary problem-solving to rigorous, research-grade challenges. Unlike earlier datasets that predominantly featured deterministic computations or straightforward symbolic manipulation, HARDMath introduces tasks that demand a deeper conceptual understanding and the ability to navigate open-ended problem spaces. This shift is critical in assessing a model’s aptitude for graduate-level reasoning, where precision, abstraction, and strategic thinking are essential. By incorporating complex domains such as asymptotic behavior and approximation theory—areas where closed-form solutions are often unavailable—HARDMath pushes models to emulate the thought processes of human mathematicians rather than simply following formulaic procedures. HARDMath’s [98] design not only increases the difficulty of tasks but also enhances the transparency and reproducibility of model evaluation. The inclusion of automatically generated, step-by-step solutions allows researchers to track the reasoning paths models take, making

it easier to diagnose errors and understand performance bottlenecks. This structured approach to evaluation provides a more granular insight into a model's strengths and weaknesses, facilitating targeted improvements. As a result, HARDMath not only raises the bar for what it means to "understand" mathematics in the context of AI but also offers a robust framework for future developments in mathematically capable language models.

MMIQC [99] introduces a high-quality dataset for competition-style mathematics, generated through the novel Iterative Question Composing (IQC) method, which refines math problems in stages and employs LLM-based rejection sampling to ensure clarity, difficulty, and pedagogical value. This iterative augmentation process produces problems that closely resemble real math competition questions, offering both diversity and depth. Fine-tuning large open-source models such as Qwen-72B on MMIQC results in substantial performance improvements, with reported gains of over 8% on the MATH benchmark—highlighting the effectiveness of domain-specific data in enhancing advanced reasoning abilities and setting a new standard for dataset generation in mathematical problem solving.

The MathChat benchmark [100] offers a new measure by which to assess the depth and adaptability of LLMs when confronted with mathematical reasoning tasks that unfold over multiple conversational turns. These types of problems both test a model's computational accuracy and its ability to manage context, interpret user intent over time, and generate intermediate steps that guide the reasoning process. Unlike the traditional benchmarks, which focus on one-shot question-answering, MathChat emphasizes the need for LLMs to engage in dynamic, multi-step dialogue—mirroring how real users might approach complex math queries in practice. This includes the model's capacity to ask clarifying questions, revise earlier reasoning when new information emerges, and maintain consistency across a series of interdependent prompts.

In their report on the MathChat benchmark, Liang et al. [100] present a comprehensive analysis of how various LLMs handle these challenges, revealing a substantial performance gap among them. While GPT-4o demonstrated strong contextual retention and error-correction capabilities, particularly in maintaining logical consistency over multiple exchanges, other models struggled with either misinterpreting follow-up queries or failing to revise incorrect prior responses. These findings shine a light on the importance of interactive reasoning benchmarks like MathChat in advancing the development of more robust and conversationally intelligent LLMs, especially in domains such as mathematics where accuracy, adaptability, and sustained reasoning are critical.

The Grade School Math - Multiple Choice (GSM-MC) benchmark measures both the accuracy and speed of LLMs when solving grade school-level math problems presented in a multiple-choice format. Proposed by Zhang et al. [101] as an enhancement of the GSM8K dataset, GSM-MC introduces a structured question format that not only tests computational correctness but also the models' ability to discern between correct answers and plausible distractors—a key challenge in educational and standardized testing contexts. Unlike traditional benchmarks that focus solely on accuracy, GSM-MC incorporates timing metrics to assess computational efficiency, making it particularly relevant for applications where rapid inference is essential. Additionally, this benchmark serves as a diagnostic tool for identifying and addressing known issues in LLM behavior, such as biases toward certain response options, a problem previously noted by Pezeshkpour and Hruschka [105]. By utilizing randomized and well-balanced multiple-choice sets, GSM-MC provides a more nuanced understanding of a model's reasoning capabilities, helping distinguish genuine problem-solving ability from reliance on statistical patterns, and offering valuable guidance for future improvements in model training and evaluation methodologies.

PROBLEMATHIC [102] is a benchmark evaluation dataset developed by researchers at Arizona State University and Georgia Institute of Technology to examine the robustness of LLMs when solving real-world math word problems, including both simple problems that require a single operation as well as more complex, multi-step problems. Unlike the traditional mathematical reasoning benchmarks, such as MathChat, which primarily test a model's ability to compute correct answers, PROBLEMATHIC is unique in emphasizing the importance of parsing through irrelevant or misleading information -

often presented in adversarial prompts - to assess whether models can identify and focus on the facts necessary to solve a problem. This design aligns with the cognitive demands placed on human problem solvers and provides a more realistic evaluation of an LLM's reasoning abilities. The benchmark reveals a significant vulnerability in current models, with performance dropping by up to 26% when adversarial noise is introduced, underscoring the difficulty LLMs face in distinguishing between relevant and irrelevant information. However, the dataset also highlights the potential of targeted training approaches: fine-tuning on adversarially noisy examples improves model robustness by approximately 8%, illustrating the value of noise-aware training strategies and the need for more sophisticated reasoning mechanisms in LLM development.

6.2. Formal Methods – Theorem Proving

LeanDojo [103] serves as a comprehensive and challenging benchmark for testing the capabilities of large language models in the domain of mathematical theorem proving. By providing 98,734 theorems and proofs extracted from a vast collection of 3384 Lean proof assistant files, it allows developers to assess how well their models perform on tasks related to formal mathematics. The Lean proof assistant itself is an open-source tool widely used by researchers to formalize and verify mathematical proofs. The open availability of LeanDojo as a Lean playground—where all the necessary data, resources, and evaluation criteria are included—ensures that researchers and developers have a consistent and transparent framework to compare the performance of different LLMs in proving theorems. One of the key insights from the benchmark's initial results is that many current LLMs struggle significantly with mathematical reasoning tasks, as seen in the relatively low success rate of models like ReProver, which only correctly proved 51.2% of the theorems.

The LeanDojo benchmark's structure has made it a valuable tool for advancing the field of automated theorem proving. Its ability to enable models to interact directly with the Lean proof assistant allows for a more nuanced evaluation of model performance, particularly in stepwise proof construction. Additionally, LeanDojo's retrieval-augmented model, ReProver, highlights the importance of premise selection—demonstrating that effective identification of relevant premises can substantially boost a model's theorem proving accuracy. Furthermore, the benchmark's design includes a split that tests models on previously unseen theorems requiring novel premises, which ensures that the evaluation process measures the model's ability to generalize beyond what it has encountered during training. This approach not only promotes a deeper understanding of how well models can handle mathematical reasoning but also supports the broader goal of improving LLM performance in structured, formal problem-solving environments [103,106].

6.3. Statistical Analysis – Data Analysis

KnowledgeMath, also referred to in some publications as FinanceMath, was introduced by Zhao et al. [104] as a comprehensive benchmark designed to assess the capabilities of large language models in solving complex mathematical challenges related to the financial sector. The benchmark includes over 1200 problems relating to a range of topics, including financial mathematics and statistical analysis, presented in diverse formats such as textual descriptions and graphical tables. Approximately 40% of these tasks require tabular reasoning, demanding a high level of analytical thinking and problem-solving. To support more detailed insights into the models' reasoning abilities, each problem is paired with executable Python-based solutions, which enable code-level analysis and verification of the model's approach. This structure makes it possible to directly compare LLM performance against human experts, who achieve an average accuracy of 94%. The benchmark thus provides a rigorous standard for evaluating model performance.

At the time that KnowledgeMath was released, even the most advanced models, such as GPT-4, were found to underperform significantly, achieving less than 50% accuracy on the benchmark. This gap in performance quality demonstrated the limitations of even state-of-the-art models in handling specialized tasks requiring advanced mathematical reasoning and financial domain knowledge. However, recent developments have led to the creation of fine-tuned models that show promise

in improving performance. For instance, models like GPT-4o and Gemini, when augmented with Chain-of-Thought and Program-of-Thought prompting techniques, have demonstrated considerable improvements, with GPT-4o reaching an accuracy of 60.9%. KnowledgeMath’s integration of a domain knowledge base encompassing over 860 finance-related concepts further enhances the LLMs’ ability to retrieve relevant information, contributing to their improved performance. Yuan et al. [107] highlight these advances, showcasing how specialized prompting strategies and domain knowledge integration are key to refining LLM capabilities in the financial sector.

7. Humanities

The Humanities discipline encompasses diverse fields that examine human culture, society, and creative expression. Recent benchmarks in this area evaluate LLMs’ capabilities in understanding social dynamics, artistic creation, and ethical reasoning. An overview of all the relevant benchmarks is presented in Table 5.

Table 5. Humanities Discipline: Domain-Specific Benchmarks.

Domain	Sub-domain	Benchmark	Scale	Task Type	Input Modality	Model	Performance	Key Focus
Social Studies	Social Media Analysis	MM-SOC [108]	Multiple	Social media	Multi-modal	Multiple	N/A	Platform analysis
		HOTVCOM [109]	10K videos	Comment generation	Video + Text	Multiple	0.42 ROUGE-L	User engagement
		XMeCap [110]	10K memes	Caption generation	Multi-image	Multiple	0.31 BLEU	Meme creation
	Bias and Privacy	Priv-IQ [111]	Multiple	Privacy intelligence	Multi-modal	Multiple	N/A	Privacy evaluation
		LLM-Bias [112]	Survey paper	Bias analysis	Text	Multiple	N/A	Fairness evaluation
Arts & Creativity	Sociocultural	CultureVLM [113]	100K images	Cultural understanding	Image + Text	Multiple	N/A	Cultural diversity
		TimeTravel [114]	Multiple	Historical reasoning	Multi-modal	Multiple	N/A	Cultural heritage
	Emotional	EmoBench-M [115]	Multiple	Emotion cognition	Multi-modal	Multiple	N/A	Emotional intelligence
		EmotionQueen [116]	1,000 dialogues	Empathy evaluation	Text	Multiple	N/A	Empathetic response
Music	Creative Tasks	EditWorld [117]	Multiple	Image editing	Image + Text	Multiple	N/A	World dynamics
		LLM-Narrative [118]	Survey paper	Storytelling	Text	Multiple	N/A	Narrative generation
		WenMind [119]	Multiple	Creative expression	Multi-modal	Multiple	N/A	Emotional creativity
Urban Planning	Music Intelligence	ZIQI-Eval [120]	14K items	Music knowledge	Text	GPT-4	58.68% F1	Comprehensive music
		MuChoMusic [121]	1,187 questions	Audio-language	Audio + Text	Multiple	N/A	Music comprehension
		Music-LLM [122]	Multiple	Symbolic music	Text	Multiple	N/A	Music reasoning
		MER-Benchmark [123]	Multiple	Entity recognition	Text	Multiple	N/A	Music entities
		MuChin [124]	Multiple	Music description	Text	Multiple	N/A	Chinese music
Morality & Ethics	City Environment	CityEQA [125]	Multiple	Embodied QA	3D + Text	Multiple	60.7% human	Urban navigation
		TransGames [126]	Multiple	Transportation	Multi-modal	Multiple	N/A	Traffic analysis
		LLM-Transport [127]	Survey paper	Transportation	Multi-modal	Multiple	N/A	System integration
	Urban Intelligence	Urban-FM [128]	Survey paper	Urban modeling	Multi-modal	Multiple	N/A	Foundation models
		UrbanPlanBench [129]	Multiple	Urban planning	Text	Multiple	N/A	Planning knowledge
Philosophy and Religion	Moral Reasoning	CityBench [130]	13 cities	Urban simulation	Multi-modal	Multiple	N/A	City management
		MoralBench [131]	Multiple	Moral evaluation	Text	Multiple	N/A	Ethical reasoning
		M ² oralBench [132]	Multiple	Multimodal ethics	Multi-modal	Multiple	N/A	Visual morality
		Greatest-Good [133]	Multiple	Utilitarian reasoning	Text	Multiple	N/A	Moral dilemmas
		Value-Alignment [134]	Multiple	Value systems	Text	Multiple	N/A	Alignment evaluation
Computational philosophy	Religion & Theology	Self-Awareness [135]	Multiple	Self-knowledge	Text	Multiple	N/A	Model capabilities
		InterIDEAS [136]	45,000 pages	Intertextual link discovery	Text	Custom pipeline	85-91% accuracy	NLP-philosophy bridges
Religion and Theology	Religion & Chatbots [137]	Religion & Chatbots [137]	N/A	Philosophical reasoning	Text	Multiple LLMs	N/A	Theological reasoning

7.1. Social Studies

Social Studies benchmarks assess LLMs’ ability to analyze social phenomena, cultural contexts, and human interactions. These evaluations reveal both promising capabilities and significant limitations in understanding complex social dynamics.

7.1.1. Social Media Analysis

Given the multimodal nature of social media user activity, the MM-Soc benchmark was introduced to evaluate the performance of multimodal LLMs (MLLMs) in diverse social media tasks across domains [108]. Evaluation of various tasks like misinformation detection and image description show notable limitations of MLLMs, particularly in zero-shot conditions, which are on par with random guesses. New datasets explore how to generate “hot” or popular comments for user engagement, reflecting a shift toward more multimodal cues [109]. XMeCap addresses meme creation where multiple images must collectively convey humor or narrative [110]. Rather than handling single-image prompts, the authors dissect sub-image relationships and fuse them into a coherent caption. They also note the challenges in maintaining comedic timing and semantic cohesion when diverse visuals are at play.

7.1.2. Bias and Privacy

On the ethical side, comprehensive surveys map persistent concerns such as bias, privacy, and intellectual property while shedding light on emerging dilemmas like truthfulness and hallucinations

[138]. From a privacy perspective, the ability of LLMs to comprehend or demonstrate privacy intelligence is vital both in terms of ethics and its ability to serve as a privacy-preserving tool. To this end, Priv-IQ was proposed as the first multimodal privacy evaluating benchmark evaluating LLMs across eight privacy competencies like named entity recognition, multilingual privacy understanding, and knowledge of privacy regulations [111]. Evaluations across state-of-the-art LLMs show that despite promising results, LLMs must improve in various competencies, including awareness of context and multilingual proficiency. Gallegos et al. [112] conducted a comprehensive review of bias and fairness in LLMs by introducing three taxonomies of metrics, datasets, and techniques for evaluating and mitigating bias.

7.1.3. Sociocultural Understanding

Multi-hop multimodal claim verification (MMCV) has emerged to address the rising complexity of cross-referencing textual, tabular, and visual information. Due to the rise in AI-generated content, automated fact-checking and verification are necessary. To this end, a large-scale claim verification dataset was introduced that reveals the challenges faced by MLLMs [139]. Claim verification requires robust reasoning capabilities and reducing the risk of superficial or one-shot model conclusions. Given the importance of peer review for scientific publication, frameworks to simulate peer reviews and understand decisions are needed. To this end, Jin et al. [140] introduced an LLM-based agent peer review simulation framework. The authors found that reviewers' bias plays a notable role in peer review decisions, further explained by social influence theory. Recent multi-object sentiment analysis research demonstrates the growing complexity of multimodal tasks. Benchmarks like MOSABench [141] highlight the challenges LLMs face when sentiment signals spread across multiple objects in an image. This points to the need for more advanced multi-object detection and context modeling. In parallel, multicultural understanding via large-scale datasets (like CultureVLM [113]) shows increasing efforts to extend LLMs and VLMs beyond Western-centric data. These initiatives are designed to handle region-specific nuances (both textual and visual) so that models can adapt to diverse cultural norms and artifacts. TimeTravel brings attention to the extensive breadth and depth of historical artifacts; the framework categorizes and comprehends cultural pieces from multiple civilizations [114]. Its core ambition is to see if AI can go beyond superficial pattern recognition to engage in historical reasoning to contribute to the digital preservation of cultural heritage. A survey of cultural representation and inclusion in LLMs reveals gaps in notable areas like semantic domains and aboutness and the need for analyzing cultural misrepresentation and underrepresentation [142].

7.1.4. Emotional Understanding

Studies on emotion cognition indicate that despite notable strides, LLMs struggle with emotional subtlety and interpretability; ongoing challenges persist in context-sensitive emotion generation and classification [143]. EmoBench-M introduces a novel viewpoint on how multimodal models manage intricate emotional cues across video, audio, and text, moving beyond simple sentiment tasks by incorporating more realistic and contextual scenarios [115]. EmotionQueen focuses on aspects of empathy in large language models, particularly targeting how well they identify user motivations and handle implicit emotional contexts in conversational settings [116].

7.2. Arts & Creativity

Recent advancements in MLLMs have reshaped creative arts by introducing tools for artists and creators. EditWorld enables diffusion models to simulate realistic world dynamics based on human-like instructions [117]. The authors curated a dataset with multimodal scenarios and used pretrained models like GPT-3.5. This approach enhanced models' capability to perform complex and scenario-aware editing tasks, surpassing traditional methods constrained by simpler operations. Recent research has extended the creative potential of MLLMs into the literary domain. For instance, the development of language-model-guided narrative generation frameworks that maintain long-term coherence and thematic depth, enabling novel storytelling approaches and complex narrative

structures [118]. These approaches leverage deep semantic understanding and contextual continuity, pushing creative boundaries beyond conventional language generation. Integrating generative AI techniques in interactive storytelling further demonstrates progress in creative domains. Specifically, advanced generative models have been employed to dynamically generate content within interactive narratives to enhance immersion and personalization of storytelling experiences [144]. These AI-driven methods enable real-time adaptation of narrative elements based on user interactions. The potential of generative models is demonstrated to facilitate creative expression and actively engage participants in co-creative processes. Moreover, researchers explored the intersection of creative AI and emotional intelligence, developing systems that effectively manage emotional expressions and reactions in generated artworks and narratives [119].

7.3. Music

A growing body of work is dedicated to benchmarking and evaluating the musical intelligence of LLMs, revealing insights and limitations. Focusing on symbolic music tasks, Zhou et al. [122] analyzed multi-step music reasoning and highlight that while LLMs like GPT-4 and LLaMA2 can generate structured outputs, they fail to synthesize motifs and forms. This indicates a deficiency in compositional logic of LLMs. Hachmeier and Jäschke [123] explored music entity recognition in user-generated content. The authors noted that LLMs with in-context learning outperform smaller models but remain sensitive to prior exposure; therefore, there is room for improvements in generalization and mitigating hallucination. ZIQI-Eval was introduced as a large-scale music benchmark spanning over 14,000 items across ten thematic categories, from theory and performance to world music traditions [120]. Despite its breadth, results show all evaluated LLMs perform poorly, with GPT-4 scoring a modest average F1 of 58.68. This work demonstrates the gap between language-based reasoning and musical comprehension. A Chinese benchmark for colloquial music description reveals a dual challenge: LLMs must not only understand professional structures but also resonate with public sentiment [124]. MuChoMusic [121] can be used to evaluate audio-language models across 1,187 questions spanning knowledge and reasoning. The findings expose models' over-reliance on text cues and highlight weaknesses in cultural, lyrical, and temporal musical interpretation.

7.4. Urban Planning

Urban planning benchmarks evaluate LLMs' ability to understand and reason about city environments and transportation systems. These assessments reveal strengths in static geospatial tasks but expose challenges in real-time decision-making and interactive planning.

7.4.1. City Environment and Transportation

Urban intelligence and transportation modeling have emerged for evaluating LLM across multimodal, dynamic, and planning-oriented tasks. To this end, Zhao et al. [125] proposed a novel benchmark for embodied question answering in urban environments. The CityEQA benchmark requires agents to interpret spatial landmarks, navigate complex 3D cityscapes, and answer open-ended questions using a hierarchical framework. Despite achieving 60.7% of human-level accuracy, LLMs face challenges with long-horizon spatial reasoning and fine-grained urban perception. Meanwhile, TransportationGames is a comprehensive benchmark for assessing LLMs and MLLMs on transportation tasks such as traffic regulation comprehension, accident analysis, and emergency plan generation [126]. The findings reveal that while text-based LLMs perform moderately well on memorization and understanding, multimodal tasks like sign recognition or road occupation detection expose gaps in practical reasoning and visual integration. Nie et al. [127] examined the role of LLMs in transportation systems through a conceptual framework and survey. The authors categorize LLM functionalities into four roles (information processors, knowledge encoders, component generators, and decision facilitators) across sensing, learning, modeling, and managing tasks. Applications of LLMs include traffic prediction, autonomous driving, and safety analytics, and LLMs can act as adaptive agents in bridging fragmented pipelines and enabling context-aware urban mobility.

7.4.2. Urban Intelligence

Recent works have explored the promise of LLMs in urban intelligence. Zhang et al. [128] proposed Urban Foundation Models, which are pre-trained on multimodal urban data. The authors introduced a taxonomy for classifying tasks and highlighting potential in domains like transportation, public safety, and environmental monitoring. The survey emphasizes the unique integration, spatio-temporal reasoning, and privacy challenges faced in building generalized urban intelligence models. UrbanPlanBench is a domain-specific benchmark evaluating LLMs across three key perspectives of urban planning: principles, professional knowledge, and regulations [129]. Despite leveraging a large instruction-tuned dataset, results reveal a gap between model performance and human-certified urban planners, especially in regulation-related tasks. CityBench provides a simulation-based evaluation framework to test LLMs on multimodal urban tasks such as traffic control, mobility prediction, and street navigation [130]. Across 13 cities and multiple LLMs, results show strengths in static geospatial tasks but expose challenges in real-time decision-making and interactive planning.

7.5. Morality and Ethics

Across moral reasoning benchmarks, LLMs can identify ethical nuances but differ significantly in depth and consistency ([131–133,145,146]). While some models, particularly the larger ones, tend to align with human intuitions (rejecting harm or maximizing collective good), others exhibit inconsistent judgment or require heavy prompting to produce coherent moral stances. The multimodal expansion of moral tasks further exposes how vision-based reasoning lags behind textual understanding, as classifying moral foundations from images adds complexity that current systems struggle to handle reliably. Advanced LLMs often outperform translation-based methods on cross-lingual tasks, yet cultural subtleties and domain shifts (like political discourse) remain stumbling blocks. Many studies converge on a similar pattern: top-tier LLMs handle general moral classification but break down on domain-specific or more nuanced ethical questions. Meanwhile, value alignment emerges as a multi-dimensional problem; no single model excels equally across all moral dimensions or value systems [134]. Self-awareness-related tasks indicate that even the best LLMs can misstate their abilities, underscoring the tension between impressive moral or social reasoning and weak self-knowledge [135].

7.6. Philosophy and Religion

Within broader philosophical and religious contexts, large corpora of historical and theological texts present a formidable test for advanced reasoning. In tasks like identifying intertextual links across extensive philosophical writings, a well-tuned LLM can approach expert-level accuracy [136]. However, the success depends on domain-specific customization, showing that general-purpose chatbots need targeted fine-tuning for sophisticated tasks. Religious and theological queries demand interpretive depth ranging from layered scriptural exegesis to abstract metaphysical deduction, and current chatbot performance is inconsistent. They exhibit enough creativity to outline arguments or interpret symbolic language but often require significant prodding to attempt complex logic (and may still fail to see ramifications) [137]. This domain reveals how advanced tasks like bridging theology and ontology test LLMs' ability to go beyond surface-level summarization into contextually rich reasoning.

7.7. Section Conclusion

Across the creative and civic domains, LLMs show remarkable progress but still fall short of expert-level reasoning and multimodal fluency ([118,144]). From short-story composition to Chinese classical poetry, specialized benchmarks ([119]) highlight the challenge of capturing nuanced cultural or artistic elements. In music, despite LLMs struggle when deep, domain-specific reasoning is required, particularly for advanced symbolic analysis or tasks that rely on actual audio content ([122,123]). These works show that many models are prone to rely heavily on textual features rather than truly parsing or understanding non-text modalities like audio or symbolic notation. Across urban planning benchmarks, a common thread is the difficulty of bridging advanced domain knowledge (especially

regulatory constraints, real-world city contexts, and multi-step reasoning) with the more general language abilities of LLMs ([125,129,130]). While many models perform decently on basic city-related questions or straightforward tasks, their performance deteriorates when confronted with intricate traffic control, urban policy, and fine-grained geographical details. Across the humanities, although LLMs have begun to engage with human-centered domains of culture, arts, and urban systems, their ability to navigate nuanced and multimodal problems remains constrained. Continued advancements in fine-tuning, multimodal integration, and domain adaptation will be essential to close the performance gap.

8. Finance

The finance domain presents unique challenges and opportunities for LLMs, requiring models to process diverse data types while maintaining high accuracy in high-stakes environments. This section explores how LLMs are being applied and evaluated across financial forecasting, sentiment analysis, document processing, and regulatory compliance. An overview of all the relevant benchmarks is presented in Table 6.

Table 6. Finance Discipline: Domain-Specific Benchmarks.

Domain	Sub-domain	Benchmark	Scale	Task Type	Input Modality	Model	Performance	Key Focus
Finance	Forecasting	FLUE [94]	5 tasks	Sentiment, NER, classification	Text	FLANG-BERT, FinBERT	+3-5% F1	Financial text processing
		FinMA / PIXIU [147]	9 datasets	Classification, QA	Text + Structured	FinMA	+10-37% F1	Finance-specific fine-tuning
	QA (Numerical)	FinanceMath [104]	1,259 items	Math reasoning	Text + Tables	GPT-4	60.9% accuracy	Financial numerical reasoning
		FinQA [148]	8,281 pairs	Multi-step QA	Text + Tables	RoBERTa + executor	65% exec acc.	Semi-structured math QA
	QA (Long-context)	DocFinQA [149]	123k tokens	Document-level QA	Text + Tables	GPT-4	67% accuracy	Long-context computation
	QA (Conversational)	ConvFinQA [150]	3,892 dialogues	Conversational QA	Text + Tables	FinQANet	68.9% exec acc.	Dialogue reasoning
	QA (Factual)	FinanceBench [6]	10,231 items	Retrieval QA	Text + Retrieval	GPT-4 + retriever	19% correct	Hallucination resistance
	ESG/Climate	ClimateBERT [151]	2M paras	Text classification	Text	DistilROBERTa	46-48% entropy loss	ESG and greenwashing
	Info Extraction	FinRED [152]	2,400 articles	Entity/relation extraction	Text	FinBERT, RoBERTa	87% F1	Structured knowledge extraction
	Multilingual QA	Golden Touchstone [153]	12 datasets	Multi-task NLP	Text (bilingual)	GPT-4, FinGPT	Strong overall	Robust multilingual benchmarking

8.1. Introduction

The domain of finance is particularly critical for deploying LLMs due to the high stakes of financial decision-making and the complexity of domain data, particularly considering the interaction of the various variables, but also the fact that, from a bottom-up perspective, these variables aggregate individual human behaviour [154]. Financial applications demand accurate understanding of diverse modalities, ranging from textual news and reports to structured balance sheets, time-series market data, charts, and potentially extending even to audio (e.g. earnings call recordings). Unlike general text domains, financial analysis often requires reasoning over both narrative disclosures and quantitative data [149]. Success in this domain can inform trading strategies, risk modelling, economic policy, macroprudential regulation and behavioural analysis for investor profiling [155-159], but errors and false prediction can carry costly consequences. Early studies show that even advanced models like GPT-4 falter on straightforward financial questions [6,160], despite the model’s apparent success in the academic realm [161]. This dual nature that comprises rich multi-modal information and high reliability requirements makes economics and finance an interesting application domain for multimodal LLMs.

8.2. LLM Applications and Benchmarks in Finance

Financial applications of LLMs span multiple areas including market forecasting, sentiment analysis, document processing, and regulatory compliance. The following sections examine how specialized benchmarks evaluate model performance across these critical financial tasks.

8.2.1. Financial Forecasting and Trading

Forecasting asset prices and market movements is a core application area for LLMs in the financial domain. This space has traditionally been dominated by time-series econometrics and, somewhat more recently, machine learning models, but is currently seeing growing interest in the integration of unstructured textual data with structured numerical inputs to capture market-relevant signals [162]. Multimodal LLMs offer new capabilities in this regard, enabling models to combine textual sources, such as financial news, earnings calls and analyst reports, with tabular or time-series data for predictive

purposes [163] This fusion of multimodal input aligns with the behavioural finance literature, which recognises that investor decision-making is shaped not only by quantitative indicators but also by narratives and sentiment embedded in public communications [156,162,164,165].

Recent approaches solve the problem of multiple data sources by using multi-modal pipelines. The Ploutos approach of Tong et al. [92] introduces a financial LLM framework with separate “experts” for text and numerical time-series, which are then combined to generate an interpretable trading rationale, which can support analysts’ decisions. Their results show that such hybrid LLM systems can outperform prior state-of-the-art trading algorithms on both prediction accuracy and explanation quality. Similarly, open-source initiatives like FinGPT [166] and proprietary models like BloombergGPT [167] have demonstrated data-centric techniques for algorithmic trading and robo-advising by automatically curating large volumes of financial data and fine-tuning models. Finally, smaller versions, like FinMA and its accompanying benchmark suite PIXIU [147], further contribute to this space. This model has been trained on a curated corpus of Chinese and English financial data and has shown strong performance in financial sentiment analysis, document classification and QA tasks, validating the importance of domain alignment in financial forecasting applications. On the computational side, an emerging benchmark is FinanceMath [104], which evaluates LLMs on domain-specific math word problems rooted in financial contexts.

All these models leverage vast financial news archives, price data and even central bank statements. Multimodality has recently extended to audio, with the newly introduced audio-text model, AT-FinGPT, that can fuse audio from earnings calls with transcript summaries for risk forecasting [168]. By capturing not only what executives say but how they say it (tone, sentiment), such models aim to predict volatility or default risk more accurately. The model’s early results in financial risk prediction confirm that integrating textual analysis with audio features calls can improve predictive power over text alone.

Nonetheless, there are persistent limitations, with interpretability being the most prominent among them. Golden Touchstone [153], a new bilingual financial LLM benchmark, explicitly evaluates not only accuracy but also a model’s ability to process complex information and provide useful explanations across tasks. Temporal robustness also remains a challenge, as models trained on historical language patterns may not generalise well to sudden macroeconomic shifts or policy regime changes. This is particularly relevant for markets affected by crisis events, even though unstructured data can often support capturing these crises [169]. Moreover, hallucination, outdated training data and, more importantly, misalignment with economic theory can lead to faulty or overconfident predictions, making it essential to pair LLMs with expert-guided guardrails.

8.2.2. Sentiment Analysis and Market Signals

Another active area for LLMs is sentiment analysis, where LLMs gauge the mood of investors or the tone of communications to inform investment decisions. Financial markets are highly sensitive to sentiment in news, social media (e.g. Reddit, Twitter) and analyst reports [170–173]. Domain-specific models like FinBERT (a BERT variant trained on finance text) were early successes in parsing financial sentiment [174], building on the work of Loughran and McDonald [175]. More recently, comprehensive benchmarks such as FLUE (Financial Language Understanding Evaluation) have catalysed progress [94]. FLUE is able to handle multiple tasks, including financial phrase sentiment classification, headline categorisation and Named Entity Recognition (NER), thus providing a suite of evaluation data to fine-tune and test LLMs on finance-specific language. Models fine-tuned on these datasets can discern subtle sentiment cues, like differentiating optimistic vs. cautious tone in an earnings press release. For example, on the FiQA sentiment analysis task, a finance-tuned model (FLANG-BERT) significantly outperforms generic BERT, highlighting the value of domain adaptation [94].

Such sentiment-aware LLMs have been applied to investor social media as well to quantify bullish or bearish signals. In algorithmic trading, these sentiment signals are often combined with price data in a form of multimodal fusion where qualitative sentiment augments quantitative models. There is also growing interest in behavioural finance modelling using LLMs, an example of which is analysing

the narrative sentiment of central bank communications to predict market reactions [176]. This bridges into behavioural finance, where LLMs help model how biases and emotions in text might foreshadow irrational market moves. Sentiment analysis in finance demonstrates a success story for LLMs, since by training on domain-specific corpora and benchmarks, LLMs can achieve human-like interpretation of financial tone and emotion, providing valuable features for trading and risk assessment.

8.2.3. Document Question Answering and Analysis

Financial professionals frequently need to query large corpora of documents to extract specific information or perform due diligence. LLMs have been applied to document question answering in this context, permitting integration of text with tabular data from financial statements. A series of benchmark datasets have driven progress here. FinQA [148] introduced a series of expert-written QA pairs on financial reports, focusing on numerical reasoning with both textual and tabular evidence. Each question in FinQA is accompanied by a small Python computation, ensuring that models not only find facts but also perform calculations for full explainability. Models are evaluated on answer accuracy and their ability to generate the correct reasoning program.

Building on this, DocFinQA [149] extended the task to full annual reports, massively increasing context length. This long-context setting reflects realistic analyst scenarios, such as browsing an entire filing for answers, and represents an apposite test of an LLM's ability to perform retrieval and reasoning over hundreds of pages. Not surprisingly, it challenges even state-of-the-art models, as experiments found that GPT-4's error rate roughly doubled when moving from the short snippets of FinQA to the full-document DocFinQA setting. The need to handle both natural language and semi-structured data (tables, financial figures) in a long context makes this a tough multimodal problem.

Another variant is ConvFinQA, which transforms FinQA into a conversational QA format [150] and can evaluate an LLM's ability to perform step-by-step reasoning in an interactive setting, thus mimicking an analyst probing a report iteratively. Beyond academic datasets, industry has also recognised the need for robust QA evaluation with FinanceBench [6]. FinanceBench questions are designed to be answerable from filings with minimal ambiguity, thus providing essentially a minimum competency test for LLMs in finance. Yet, when state-of-the-art models (GPT-4, LLaMA-2, etc.) were evaluated with retrieval augmented techniques, most struggled, with GPT-4 providing correct responses to only 19% of questions and often refusing or hallucinating answers [6]. These QA benchmarks (the FinQA family and FinanceBench) have driven the progress but also emphasises the challenges. LLMs fine-tuned on financial documents can answer many questions and even perform calculations, but ensuring accuracy and reliability at scale (especially for compliance-critical scenarios) demands increased model power but also consistency and precision.

8.2.4. Regulatory Compliance, ESG and Behavioural Finance

Beyond core trading and analysis tasks, LLMs are being explored in less mainstream but equally important subdomains of finance, with one key being regulatory compliance. To ensure that they meet the required standards, institutions must parse complex regulations, policies and legal filings. LLMs can assist by answering questions about compliance requirements or by flagging sections of a legal document relevant to a query. While there is still a lack of large public benchmarks in this exact area, the capabilities overlap with the financial QA already discussed as well as with legal NLP tasks. As with most topics in the Finance domain, ensuring explainability is crucial here, since compliance officers need to trust and verify the sources of the answer provided by the LLM. As a result, there is a strong need for LLMs that can cite evidence and reason transparently.

Another growing field is ESG (Environmental, Social, Governance) and climate finance. Specialised models like ClimateBERT [151] have been developed to handle climate-related financial text. ClimateBERT is a transformer model refined on sustainability reports and climate disclosures, enabling it to detect language about climate risks and commitments [177]. Tasks in this area include classifying statements in annual reports as substantive or "greenwashing", extracting metrics related to carbon emissions or summarising climate risk disclosures for investors. Multimodal aspects appear when

combining text with data like emissions figures or when analysing satellite imagery alongside reports for due diligence, though such extreme multimodality is still at early stages.

Lastly, behavioural finance, which studies the psychological factors in investor decisions, offers intriguing opportunities for LLMs. Researchers have begun using LLMs to model investor sentiment and biases or even to simulate market participants. For example, an LLM agent could be assigned the persona of an “overconfident investor” versus a “loss-averse investor” to see how their responses differ when presented with the same news, potentially uncovering bias-driven divergences. While this is still exploratory, it builds on the idea of agent-based financial simulations that are able to uncover important insights regarding individual behaviour and how it can impact government policies [178,179]. A collaborative project by the Stevens Institute and others has started testing LLM agents in human-centred financial decision-making, essentially using LLMs to augment or simulate the role of financial advisors and investors [180]. In this manner, domain-specific LLM research in this area aims to expand the scope of what such models can do, by moving from pure “number-crunching” and fact-finding into interpreting the human side of finance.

8.3. Section Conclusion

Across these subdomains, several clear trends have emerged. First, we see a proliferation of finance-specific LLMs and fine-tuned models. Large proprietary models, like BloombergGPT and open-source counterparts, like FinGPT and FinMA/PIXIU, demonstrate the value of ingesting domain data at scale, continually integrating fresh financial data to keep models current. Moreover, smaller models fine-tuned via instruction on financial tasks (e.g. FinMA on the FLUE benchmark tasks) have made domain adaptation accessible without training from scratch. Second, there is growing emphasis on integrating multimodal capabilities with structured data handling, particularly in tasks that involve numerical reasoning. In financial tasks that require multi-step calculations or querying large documents, such as those found in benchmarks like DocFinQA, models that combine LLMs with external tools tend to perform more accurately and reliably than those relying on language understanding alone. Third, interpretability and explainability are recognised as essential features. In high-stakes domains, such as Finance, a black-box answer is often not enough as regulators and practitioners demand reasoning traces. Approaches like program annotations in FinQA or rationale generation in Ploutos address this by design, and future benchmarks are likely to include explainability as a first-class metric.

Alongside these successes, however, significant challenges persist. Access to high-quality financial data remains a significant challenge. Although the volume of financial information has increased substantially, much of it is proprietary or non-public, which restricts its use in open academic research [166]. At the same time, publicly available data sources may contain misinformation, inconsistencies or unverified content, particularly when sourced from social media or news aggregators. In addition, alignment with domain experts and industry regulations poses another hurdle, since an LLM’s answers must not only be correct, but also compliant with legal and ethical standards. Additionally, hallucination and factual accuracy are also important concerns. This findings of FinanceBench regarding widespread mistakes even by top models underscore that current LLMs can be unreliable [6]. In applications like finance where factual errors can lead to financial loss or legal liability, this creates serious issues and hinders adoption. Temporal robustness is also a unique challenge, as financial models need to stay current with evolving market conditions, new regulations and recent company disclosures, since an LLM trained on last year’s data might miss critical context. Finally, efficiency and scalability are practical issues, since handling lengthy documents or streaming data in real-time can be computationally intensive and latency is a concern in trading environments.

Looking forward, several avenues could further advance domain-specific LLM performance in finance. One is the creation of synthetic and privacy-preserving datasets. For example, using simulators or generative techniques can help produce realistic financial data and scenarios that can augment training without risking real-world consequences. Another is developing standardised evaluation metrics for explainability and reasoning transparency, so that models are not only judged on accuracy

but also on the faithfulness of their justifications. Finally, integrating insights from economics and behavioural finance into model design could prove fruitful. By embedding theoretical constraints or known bias patterns, we might guide LLMs to make decisions that are both rational and cognisant of human irrationality. The trajectory in this domain suggests that finance will remain at the forefront of domain-specific LLM evaluation both because of its challenges and because of the tangible rewards of getting it right.

9. Healthcare

Large Language Models (LLMs) have emerged as powerful tools in healthcare, with applications spanning clinical decision support, medical research, and patient care. Recent benchmarks demonstrate both the promise and limitations of these models across diverse medical domains. In clinical settings, LLMs show potential for assisting with diagnosis, treatment planning, and medical documentation [181,182]. For bioinformatics applications, models are being evaluated on complex tasks like genomic analysis and drug discovery [183,184]. In medical imaging, specialized architectures are tackling challenges in radiology, pathology, and 3D analysis [185,186]. An overview of all the relevant benchmarks is presented in Table 7.

Table 7. Healthcare Discipline: Domain-Specific Benchmarks.

Domain	Sub-domain	Benchmark	Scale	Task Type	Input Modality	Model	Performance	Key Focus
Medicine & Healthcare	Clinical Applications	GMAI-MMBench [181]	284 datasets	VQA	Multi-modal	GPT-4o	53.96%	Comprehensive medical AI
		Asclepius [182]	15 specialties	Multi-task	Multi-modal	Multiple	N/A	Specialty evaluation
		MultiMed [187]	2.56M samples	Multi-task	Multi-modal	Multiple	N/A	Cross-modality
		MediConfusion [188]	Paired images	VQA	Multi-modal	Multiple	Below random	Model reliability
		CTBench [189]	2.5K samples	Multi-task	Text	GPT-4	72.3%	Clinical trial design
	Bioinformatics	PGxQA [183]	110K+ questions	QA	Text	GPT-4o	68.4%	Pharmacogenomics
		GenoTEx [184]	Multiple datasets	Gene analysis	Text + Data	GenoAgent	65-85%	Gene expression
		MedCalc-Bench [190]	2K problems	Calculation	Text + Numbers	GPT-4	57.8%	Medical calculations
Medical Imaging	Bio-Benchmark [191]	30 tasks	Multi-task	Text	GPT-4o	Varies by task	Bioinformatics NLP	
	Radiology & 3D Analysis	M3D [185]	120K img pairs	Multi-task	3D + Text	M3D-LaMed	SOTA	3D medical imaging
		TriMedLM [192]	Multi-modal	Image + Text	TriMedLM	N/A	Trimodal integration	
		RadGPT [186]	9.2K CT scans	Report gen	3D CT + Text	RadGPT	80-97% sensitivity	Tumor detection
	Pathology/Microscopy	Micro-Bench [193]	3.2K images	VQA	Image + Text	GPT-4V	54.7%	Microscopy understanding
		PathMMU [194]	33K questions	Multiple choice	Image + Text	GPT-4V	49.8% vs 71.8% human	Pathology expertise
		MicroVQA [195]	2.8K QA pairs	VQA	Image + Text	GPT-4V	48.2%	Scientific microscopy
		SlideChat [196]	Large-scale	Dialogue	Image + Text	SlideChat	N/A	Interactive analysis
μ -Bench [197]		Multiple tasks	Multi-task	Image + Text	Multiple	Varies by task	Microscopy analysis	

9.1. Medicine & Healthcare

The healthcare domain has seen extensive development of specialized benchmarks across clinical medicine, bioinformatics, and medical imaging. In clinical applications, benchmarks like GMAI-MM Bench [181] and Asclepius [182] evaluate models across diverse medical specialties and tasks. For bioinformatics, datasets like PGxQA [183] and GenoTEX [184] assess capabilities in genomics and pharmacology. Medical imaging introduces unique challenges, with benchmarks like M3D [185] and RadGPT [186] testing 3D analysis capabilities. Let us examine each subdomain in detail.

9.1.1. Clinical Applications

Large Language Models (LLMs) are increasingly being applied in clinical healthcare, offering potential to assist medical professionals across diagnostic support, decision-making, and information access. Recent benchmarks have focused on evaluating how effectively multimodal LLMs can handle clinical tasks, revealing both promising capabilities and concerning limitations. These evaluations are critical for determining whether such models can safely support healthcare workflows, where errors can have serious consequences. The benchmarks in this domain test models across diverse medical specialties, task types, and input modalities, providing a comprehensive assessment of their clinical applicability and reliability.

GMAI-MMBench [181] represents the most comprehensive medical AI benchmark to date, comprising 284 datasets across 38 medical image modalities, 18 clinical tasks, and 18 medical departments. The benchmark features a well-categorized structure with four different perceptual granularities (image, box, mask, and contour levels), allowing for detailed assessment of model capabilities. Evaluations of 50 large vision-language models, including both open-source and proprietary systems, revealed that even the best performing model (GPT-4o) achieved only 53.96% accuracy, highlighting

significant room for improvement. The benchmark identified five key insufficiencies in current models: uneven task performance across different clinical scenarios, unbalanced department performance, inconsistent perceptual robustness, poor instruction following in medical contexts, and difficulties with multiple-choice questions.

Asclepius [182] offers a specialized benchmark covering 15 medical specialties, designed to evaluate models' ability to perform across diverse clinical domains. This benchmark addresses the critical need for comprehensive evaluation across specialized medical fields, as performance in one specialty may not generalize to others. The evaluation identifies significant variations in model capabilities across different medical areas, highlighting the challenge of developing systems that maintain consistent performance across the spectrum of clinical practice.

MultiMed [187] provides a massive dataset of 2.56 million samples to assess models' capacity for cross-modality reasoning in medical contexts. This benchmark is particularly valuable for evaluating how effectively models can integrate information across different data types (such as images, text, and structured data), which is essential for many clinical workflows that require synthesis of diverse information sources. The large scale of this dataset enables robust assessment of model generalization capabilities across varied medical scenarios.

MediConfusion [188] takes a unique approach by specifically probing the failure modes of medical multimodal LLMs. This benchmark exposes a concerning limitation: state-of-the-art models are easily confused by image pairs that are visually dissimilar and clearly distinct to medical experts. All evaluated models—both open-source and proprietary—performed below random guessing on this benchmark, raising serious questions about their reliability for clinical deployment. The study identified common patterns of model failure, including difficulties distinguishing between normal anatomy and pathology, interpreting lesion signal characteristics, analyzing vascular conditions, and identifying medical devices. These findings highlight fundamental limitations in current models' visual understanding capabilities.

CTBench [189] focuses specifically on clinical trial design, containing 2,500 samples to evaluate models' ability to assist in this critical aspect of medical research. GPT-4 achieved a relatively strong performance of 72.3% on this benchmark, suggesting potential utility in supporting the complex process of clinical trial planning. This specialized benchmark addresses a specific clinical workflow where LLMs might provide significant value by accelerating research design and improving protocol quality.

9.1.2. Bioinformatics

Large Language Models offer transformative potential in bioinformatics by automating complex biological data analysis, interpreting genetic information, and accelerating biomedical research. Recent benchmarks in this domain evaluate LLMs' capabilities in processing specialized biological knowledge and performing data-intensive computational tasks. These evaluations reveal that while state-of-the-art models can achieve reasonable performance on specific bioinformatics tasks, there remains a substantial gap between current abilities and the expertise required for clinical implementation. The benchmarks in this area focus on assessing various capabilities essential to bioinformatics, from knowledge retrieval to data manipulation, providing insights into the strengths and limitations of applying LLMs to this highly specialized domain.

PGxQA [183] focuses on pharmacogenomics, evaluating LLMs on their ability to answer questions about genetics-guided treatment from the perspectives of clinicians, patients, and researchers. The benchmark contains over 110,000 questions covering areas such as allele definition, frequency, diplotype-to-phenotype mapping, and drug-gene interactions. Evaluation of models including GPT-4o, Llama 3, and Gemini 1.5 Pro revealed that while newer models like GPT-4o (68.4% accuracy) significantly outperform older generations, they still fall short of the standards required for clinical use. The study highlights particular difficulties with questions related to allele definitions and functions, while showing stronger performance on queries involving drug and gene entities. This benchmark addresses

a critical gap in pharmacogenetics implementation, where the lack of education and awareness among clinicians represents a major barrier to the adoption of precision medicine approaches.

GenoTEX [184] introduces a benchmark for evaluating LLM-based methods in automating gene expression data analysis. The benchmark provides analysis code and results for solving gene-trait association problems, covering key steps including dataset selection, preprocessing, and statistical analysis in a pipeline that adheres to computational genomics standards. The benchmark introduces three core tasks for evaluation: dataset selection, data preprocessing, and statistical analysis, each with corresponding metrics. To establish baselines, the researchers present GenoAgent, a team of collaborative LLM-based agents that adopt a multi-step programming workflow with flexible self-correction. Experimental results showed that GenoAgent performs well on certain tasks, achieving 65-85% success across various metrics, demonstrating the potential of LLM-based methods in genomic data analysis. However, the evaluation also highlights limitations in handling complex tasks like clinical feature extraction and statistical analysis, identifying areas for further improvement.

MedCalc-Bench [190] presents a benchmark consisting of 2,000 medical calculation problems designed to test LLMs' ability to perform precise quantitative analyses in healthcare contexts. The benchmark evaluates models on tasks requiring both medical knowledge and mathematical reasoning, with GPT-4 achieving a moderate performance of 57.8%. These results highlight the challenge of combining domain expertise with computational accuracy, a critical requirement for clinical decision support. The benchmark addresses an important aspect of healthcare applications where errors in calculation could lead to serious consequences for patient care.

Bio-Benchmark [191] offers a comprehensive framework for evaluating LLMs across 30 key bioinformatics tasks spanning seven domains: protein, RNA, RNA-binding protein (RBP), drug, electronic health records (EHR), medical Q&A, and traditional Chinese medicine (TCM). The researchers evaluated six mainstream LLMs, including GPT-4o and Llama-3.1-70b, using both zero-shot and few-shot Chain-of-Thought settings without fine-tuning to reveal their intrinsic capabilities. The evaluation utilized BioFinder, a new tool developed for extracting answers from LLM responses with approximately 30% greater accuracy than existing methods. Results revealed that few-shot learning significantly improved performance across multiple tasks, with some models seeing up to 20x improvement in protein species prediction and substantial gains in drug design and TCM tasks. Performance varied considerably across different bioinformatics subdomains, with GPT-4o generally performing best but different models excelling in specific areas—Mistral-large-2 led in protein species prediction (82%) and drug design (91%), while Llama-3.1-70b performed strongly in RNA function prediction (89%).

9.2. Medical Imaging

Medical imaging benchmarks have emerged as a critical frontier in evaluating multimodal LLMs, particularly for their ability to handle complex 3D volumetric data and specialized imaging modalities. Recent evaluations reveal both promising advances and significant challenges - while models like RadGPT [186] achieve strong performance in tumor detection (80-97% sensitivity), even state-of-the-art systems struggle with basic visual reasoning tasks in pathology. The benchmarks in this domain are increasingly focusing on clinically-relevant capabilities, from 3D spatial understanding to interactive analysis workflows that mirror real clinical practice.

9.2.1. Radiology & 3D Analysis

Medical imaging analysis presents unique challenges for multimodal large language models, particularly when dealing with 3D volumetric data common in modalities like CT and MRI scans. Recent benchmarks in this domain evaluate how well models can interpret complex spatial relationships, generate accurate clinical descriptions, and support diagnostic decision-making from 3D medical imagery. These evaluations reveal both the promise of extending MLLMs beyond 2D images and the significant challenges in handling the increased complexity of volumetric data. Benchmarks in this area assess diverse capabilities including anatomical understanding, report generation, visual question answering, and diagnostic accuracy.

M3D [185] represents a significant advancement in evaluating 3D medical image analysis capabilities through MLLMs. This benchmark comprises a large-scale dataset (M3D-Data) containing 120,000 image-text pairs and 662,000 instruction-response pairs specifically tailored for various 3D medical tasks. The benchmark evaluates models across eight diverse tasks in five major categories: image-text retrieval, report generation, visual question answering (both open-ended and closed-ended formats), positioning (including referring expression comprehension and generation), and segmentation (both semantic and referring expression segmentation). The researchers also introduce M3D-LaMed, a versatile multimodal model for 3D medical image analysis that integrates a 3D Vision Transformer as the vision encoder, a 3D spatial pooling perceiver for reducing embedding dimensions, and the LLaMA-2-7B model as the base LLM. Comprehensive evaluation demonstrates that M3D-LaMed outperforms existing solutions across the benchmark tasks, establishing a new state-of-the-art for 3D medical image understanding, while also showing promising capabilities in answering out-of-distribution questions.

TriMedLM [192] introduces a trimodal approach to medical imaging analysis, combining conventional image data with additional modalities to enhance diagnostic capabilities. This benchmark evaluates models on their ability to integrate information across these modalities, recognizing that medical diagnosis often requires synthesis of diverse data types. The model architecture employs a specialized design to handle the trimodal inputs, enabling more comprehensive analysis than is possible with unimodal approaches. Evaluation results demonstrate that this integrated approach yields improvements in diagnostic accuracy and completeness of analysis, highlighting the value of multimodal integration in medical imaging applications.

RadGPT [186] presents an anatomy-aware vision-language AI agent for generating detailed reports from CT scans, with a particular focus on tumor analysis. The system first segments tumors (including benign cysts and malignant tumors) and their surrounding anatomical structures, then transforms this information into both structured and narrative reports containing detailed information on tumor size, shape, location, attenuation, volume, and interactions with surrounding blood vessels and organs. Evaluation on unseen hospitals demonstrated strong performance in tumor detection, with high sensitivity/specificity for small tumors (<2 cm): 80/73% for liver tumors, 92/78% for kidney tumors, and 77/77% for pancreatic tumors. For larger tumors, sensitivity ranged from 89% to 97%, significantly surpassing the state-of-the-art in abdominal CT report generation. The researchers also created AbdomenAtlas 3.0, the first publicly available image-text 3D medical dataset, comprising over 1.8 million text tokens and 2.7 million images from 9,262 CT scans, including 2,947 tumor scans/reports of 8,562 tumor instances. This resource provides valuable capabilities including localization of tumors in organ sub-segments, determination of pancreatic tumor stage, and individual analyses of multiple tumors—a feature rarely found in human-made reports.

9.2.2. Pathology/Microscopy

Pathology and microscopy represent specialized areas of medical imaging that involve the analysis of cellular and tissue-level structures at various magnifications. Multimodal large language models applied to these domains must interpret fine-grained visual features that are substantially different from natural images or radiological scans, requiring both deep domain knowledge and sophisticated visual reasoning. Recent benchmarks in this area evaluate how effectively MLLMs can perform tasks ranging from basic microscopic image interpretation to complex diagnostic reasoning. These evaluations consistently reveal a significant performance gap between current state-of-the-art models and human expert pathologists, highlighting the challenges of adapting general-purpose vision-language models to this specialized domain.

Micro-Bench [193] introduces a benchmark containing 3,200 microscopy images paired with visual question answering tasks to evaluate models' understanding of microscopic structures and features. In evaluations, GPT-4V achieved 54.7% accuracy, demonstrating moderate capabilities but still falling significantly short of expert-level performance. The benchmark specifically tests models' ability to interpret fine-grained visual features unique to microscopy images, including cellular morphology, tissue architecture, and subcellular structures that require specialized domain knowledge to properly

analyze. This benchmark addresses the critical need for evaluating models on their ability to interpret scientific imagery that differs substantially from natural images seen in general visual reasoning tasks.

PathMMU [194] represents the largest expert-validated pathology benchmark for evaluating multi-modal models, comprising 33,428 multiple-choice questions and 24,067 pathology images from diverse sources. Each question is accompanied by an explanation for the correct answer. The benchmark was constructed using a three-step process: data collection from diverse sources (PubMed, educational videos, pathology textbooks, expert Twitter posts), detailed description generation using GPT-4V to enhance image captions, and question generation with rigorous expert validation by seven pathologists who manually reviewed approximately 12,000 questions in the validation and test sets. Evaluation of 18 state-of-the-art models revealed that even the most advanced model, GPT-4V, achieved only 49.8% accuracy, compared to 71.8% achieved by human pathologists—a gap of over 20 percentage points. The benchmark also revealed concerning trends: 15 out of 18 tested models achieved no more than 40% accuracy, models demonstrated unexpected robustness to image corruption (raising questions about whether they actually utilized visual information), and text-only models performed surprisingly well, suggesting models may take shortcuts rather than performing proper image analysis.

MicroVQA [195] provides 2,800 question-answer pairs focused specifically on scientific microscopy, evaluating models' capability to reason about microscopic structures and processes across biological, medical, and materials science applications. GPT-4V achieved a modest 48.2% accuracy on this benchmark, highlighting the difficulty of interpreting specialized scientific imagery. The benchmark is designed to test not only recognition of microscopic structures but also reasoning about their function, relationships, and significance in scientific contexts. This evaluation provides insights into the current limitations of multimodal models when applied to specialized scientific domains that require both visual understanding and domain-specific knowledge.

SlideChat [196] introduces a dialogue-based approach to pathology slide analysis, enabling interactive exploration of pathological specimens. The system supports conversational interaction where users can ask questions about specific regions or features of a slide, approximating the interactive workflow of pathologists who iteratively examine different areas of a specimen at varying magnifications. This benchmark evaluates models on their ability to maintain context through a dialogue, accurately describe pathological features, and respond appropriately to follow-up questions—capabilities that more closely match how pathologists actually work with digital slides than single-turn question answering. The specialized SlideChat model demonstrates the value of architecture and training approaches tailored specifically for pathology applications.

μ -Bench [197] offers a comprehensive evaluation framework for microscopy image analysis across multiple tasks, including classification, segmentation, and visual question answering. The benchmark includes diverse microscopy modalities and sample types, providing a holistic assessment of models' capabilities in this domain. Evaluation results vary significantly across tasks and model architectures, with no single approach excelling across all microscopy analysis scenarios. This multi-task benchmark highlights the complexity of microscopy analysis and the need for models that can flexibly adapt to different analytical requirements, from basic feature identification to complex interpretative reasoning.

Medical imaging benchmarks consistently reveal the substantial gap between current model performance and expert-level capabilities, with even top-performing models like GPT-4V achieving accuracy well below human pathologists. A notable trend across these evaluations is models' tendency to rely more on textual information than properly analyzing visual content, taking shortcuts rather than performing true visual reasoning. Despite these limitations, specialized architectures like M3D-LaMed and RadGPT demonstrate the value of domain-specific training and architectural approaches, particularly for 3D data analysis. Future development should focus on improving fine-grained visual feature extraction, enhancing multi-scale analysis capabilities, and developing more rigorous evaluation protocols that prevent models from exploiting textual shortcuts.

10. Language Understanding

Language understanding benchmarks evaluate models’ capabilities across diverse linguistic tasks, from multimodal integration to complex reasoning. These benchmarks reveal both significant progress and persistent challenges in developing truly comprehensive language understanding systems. An overview of all the relevant benchmarks is presented in Table 8.

Table 8. Language Understanding Discipline: Domain-Specific Benchmarks.

Domain	Sub-domain	Benchmark	Scale	Task Type	Input Modality	Model	Performance	Key Focus
Multimodal Communication	Visual-Language Integration	LongLLaVA [198]	1000 images	Long-context	Multi-modal	LongLLaVA	N/A	Hybrid architecture
		LLaVA-OneVision [199]	N/A	Multi-scenario	Multi-modal	LLaVA	N/A	Cross-modal transfer
		KOSMOS-1 [200]	Web-scale	Multi-task	Multi-modal	KOSMOS-1	N/A	Cross-modal transfer
		KOSMOS-2 [201]	Large-scale	Multi-task	Multi-modal	KOSMOS-2	N/A	Visual grounding
Cross-lingual Tasks	Multilingual Models	ChatGLM [202]	10T tokens	Multi-task	Multi-modal	GLM-4	Matches GPT-4	Chinese-English LLM
		CVLUE [203]	30K+ samples	Multi-task	Multi-modal	X ² VLM	36-55%	Chinese VL understanding
Reasoning	Problem Solving	ToT [204]	N/A	Problem solving	Text	GPT-4	74%	Deliberate reasoning
	Domain-Specific Applications	AgEval [205]	12 tasks	Plant phenotyping	Multi-modal	Multiple	46-73% F1	Agricultural tasks

10.1. Multimodal Communication

Recent developments in visual-language integration have focused on addressing key challenges in multimodal large language models (MLLMs), particularly in handling long-context scenarios and improving cross-modal transfer capabilities. The benchmarks in this domain demonstrate significant progress in expanding the contextual understanding and efficiency of MLLMs.

A notable advancement in long-context processing is demonstrated by LongLLaVA [198], which introduces a hybrid architecture combining Mamba and Transformer blocks. This approach enables processing of up to 1000 images on a single GPU while maintaining competitive performance. The architecture addresses the dual challenges of performance degradation with increased image count and computational efficiency, representing a step forward in scaling MLLMs for practical applications.

Cross-modal transfer has emerged as another critical focus area, as evidenced by LLaVA-OneVision [199] and KOSMOS-1 [200]. LLaVA-OneVision demonstrates the feasibility of a unified model capable of handling multiple visual scenarios, including single-image, multi-image, and video understanding. This approach enables strong transfer learning across different modalities, particularly from images to videos. Similarly, KOSMOS-1 showcases the potential of web-scale training for developing MLLMs that can perceive general modalities and perform few-shot learning and instruction following.

As shown in Table below, these benchmarks vary in their scale and focus. While LongLLaVA provides a quantified dataset of 1000 images, KOSMOS-1 operates at a web-scale level. The progression from single-task to multi-task and multi-scenario capabilities indicates the field’s movement toward more versatile and robust models. However, challenges remain in balancing computational efficiency with model performance and ensuring effective knowledge transfer across different modalities.

10.2. Cross-lingual Tasks

In the domain of multilingual models, ChatGLM [202] represents a significant advancement in cross-lingual capabilities. The model family, particularly GLM-4, demonstrates the potential of large-scale multilingual training with a corpus of 10 trillion tokens primarily in Chinese and English, along with coverage of 24 additional languages. As shown in Table 8, ChatGLM achieves performance comparable to GPT-4 across various benchmarks while excelling specifically in Chinese language tasks. The model’s success in matching GPT-4’s performance while maintaining strong multilingual capabilities suggests the effectiveness of their multi-stage post-training process, which combines supervised fine-tuning with learning from human feedback.

CVLUE [203] addresses a critical gap in vision-language understanding evaluation for Chinese culture. Unlike existing Chinese VL datasets that often use Western-centric images from English datasets, CVLUE features images specifically selected by Chinese native speakers to represent Chinese cultural contexts. The benchmark comprises four distinct tasks: image-text retrieval, visual question answering, visual grounding, and visual dialogue. Evaluations of multilingual VLMs on CVLUE

reveal significant performance gaps between English and Chinese vision-language understanding capabilities, with models scoring 20-35% lower on Chinese tasks compared to English counterparts. The benchmark's category-level analysis demonstrates that fine-tuning on Chinese culture-related visual-language data substantially improves model performance, particularly for culturally specific concepts.

10.3. Reasoning

The field of reasoning has seen methodological innovations in how language models approach complex problem-solving tasks. A notable advancement in this domain is the Tree of Thoughts (ToT) framework [204], which extends beyond traditional sequential reasoning approaches. As indicated in Table 8, ToT demonstrates substantial improvements in deliberate problem solving, achieving a 74% success rate on specific tasks where conventional methods struggle. The framework enables language models to explore multiple reasoning paths simultaneously, implement strategic lookahead, and perform self-evaluation of intermediate decisions. This approach particularly excels in tasks requiring non-trivial planning or search capabilities, such as mathematical puzzles and creative writing tasks, where initial decisions significantly impact the final outcome. The success of ToT suggests the importance of structured exploration and deliberate decision-making in enhancing language models' problem-solving capabilities.

10.4. Domain-Specific Applications

AgEval [205] introduces a comprehensive benchmark for evaluating Vision Language Models (VLMs) on specialized agricultural tasks, particularly plant stress phenotyping. The benchmark encompasses 12 diverse plant stress tasks across identification, classification, and quantification categories, assessing zero-shot and few-shot in-context learning performance of models like Claude, GPT, Gemini, and LLaVA. Results demonstrate VLMs' rapid adaptability to specialized tasks, with F1 scores increasing from 46.24% to 73.37% in 8-shot settings. The study also quantifies performance disparities across plant stress classes using coefficient of variation metrics (ranging from 26.02% to 58.03%) and shows that strategic example selection improves model reliability by up to 15.38%. AgEval establishes that VLMs, with minimal few-shot examples, can serve as viable alternatives to traditional specialized models in agricultural applications.

11. Conclusions and Future Work

This survey has underscored the indispensable role of domain-specific benchmarks in guiding Multimodal Large Language Models (MLLMs) beyond general proficiency to achieve the critical 'last mile problem' of efficacy in specialized applications. While foundational models show broad capabilities, their performance often diminishes when faced with the nuanced demands of fields like finance, where even top models like GPT-4 struggle on benchmarks such as FinanceBench [6], or in medicine, where models can be confused by visually distinct images as seen in MediConfusion [188]. The path forward necessitates a concerted effort towards greater standardization and comparability in benchmark design and evaluation protocols. Future benchmarks must evolve to incorporate comprehensive metrics beyond mere accuracy, assessing robustness, efficiency, and safety. They also need to become more dynamic, or "living," to keep pace with rapid MLLM advancements, and must challenge models with increasing multimodal complexity, noisy real-world data, and tasks that demand deeper, more causal reasoning than currently evaluated—moving beyond the current limitations where, for instance, even advanced models achieve only moderate success on comprehensive medical benchmarks like GMAI-MMBench [181].

Beyond refining evaluation quality, future work must focus on enhancing the broader applicability and trustworthiness of MLLMs. This includes a critical emphasis on ethical considerations and societal impact, developing benchmarks that rigorously assess for biases, privacy vulnerabilities, and the potential for misuse, particularly in sensitive areas like social media analysis as highlighted by concerns in MM-Soc [108] and Priv-IQ [111], or within autonomous systems. The 'black box' nature

of MLLMs remains a significant hurdle; thus, fostering interpretability and explainability through benchmark design is paramount, especially in high-stakes domains like healthcare where models sometimes resort to textual shortcuts rather than true visual reasoning, as observed in pathology benchmarks like PathMMU [194]. Furthermore, advancing MLLM utility involves exploring cross-domain knowledge transfer and integrating human-in-the-loop evaluations for tasks requiring subjective or nuanced judgment, prevalent in the humanities and arts where automated metrics are often insufficient. Addressing these multifaceted challenges will require a holistic approach to benchmark development, guiding MLLMs to be not only powerful but also reliable, interpretable, and ethically aligned with societal values as they transform a myriad of applications.

References

1. Achiam, J.; Adler, S.; Agarwal, S.; Ahmad, L.; Akkaya, I.; Aleman, F.L.; Almeida, D.; Altenschmidt, J.; Altman, S.; Anadkat, S.; et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774* **2023**.
2. Team, G.; Anil, R.; Borgeaud, S.; Alayrac, J.B.; Yu, J.; Soricut, R.; Schalkwyk, J.; Dai, A.M.; Hauth, A.; Millican, K.; et al. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805* **2023**.
3. Bommasani, R.; Hudson, D.A.; Adeli, E.; Altman, R.; Arora, S.; von Arx, S.; Bernstein, M.S.; Bohg, J.; Bosselut, A.; Brunskill, E.; et al. On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258* **2021**.
4. Grattafiori, A.; Dubey, A.; Jauhri, A.; Pandey, A.; Kadian, A.; Al-Dahle, A.; Letman, A.; Mathur, A.; Schelten, A.; Vaughan, A.; et al. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783* **2024**.
5. Abdin, M.; Aneja, J.; Behl, H.; Bubeck, S.; Eldan, R.; Gunasekar, S.; Harrison, M.; Hewett, R.J.; Javaheripi, M.; Kauffmann, P.; et al. Phi-4 technical report. *arXiv preprint arXiv:2412.08905* **2024**.
6. Islam, P.; Kannappan, A.; Kiela, D.; Qian, R.; Scherrer, N.; Vidgen, B. FinanceBench: A new benchmark for financial question answering, [2311.11944]. *arXiv preprint*.
7. Zheng, D.; Wang, Y.; Shi, E.; Liu, X.; Ma, Y.; Zhang, H.; Zheng, Z. Top General Performance = Top Domain Performance? DomainCodeBench: A Multi-domain Code Generation Benchmark, 2025, [arXiv:cs.SE/2412.18573].
8. Li, J.; Zhu, Y.; Xu, Z.; Gu, J.; Zhu, M.; Liu, X.; Liu, N.; Peng, Y.; Feng, F.; Tang, J. Mmro: Are multimodal llms eligible as the brain for in-home robotics? *arXiv preprint arXiv:2406.19693* **2024**.
9. Doris, A.C.; Grandi, D.; Tomich, R.; Alam, M.F.; Ataei, M.; Cheong, H.; Ahmed, F. DesignQA: A Multimodal Benchmark for Evaluating Large Language Models' Understanding of Engineering Documentation. *Journal of Computing and Information Science in Engineering* **2024**, *25*, 021009, [https://asmedigitalcollection.asme.org/computingengineering/article-pdf/25/2/021009/7419451/jcise_25_2_021009.pdf]. https://doi.org/10.1115/1.4067333.
10. Kernan Freire, S.; Wang, C.; Foosherian, M.; Wellsandt, S.; Ruiz-Arenas, S.; Niforatos, E. Knowledge sharing in manufacturing using LLM-powered tools: user study and model benchmarking. *Frontiers in Artificial Intelligence* **2024**, *7*, 1293084.
11. Liu, X.; Yang, S.; Dong, X.; Rong, H.; Fu, B. Manu-Eval: A Chinese Language Understanding Benchmark for Manufacturing Industry. In Proceedings of the China Conference on Knowledge Graph and Semantic Computing. Springer, 2024, pp. 309–317.
12. Eslaminia, A.; Jackson, A.; Tian, B.; Stern, A.; Gordon, H.; Malhotra, R.; Nahrstedt, K.; Shao, C. FDM-Bench: A Comprehensive Benchmark for Evaluating Large Language Models in Additive Manufacturing Tasks. *arXiv preprint arXiv:2412.09819* **2024**.
13. Fakihi, M.; Dharmaji, R.; Moghaddas, Y.; Quiros, G.; Ogundare, O.; Al Faruque, M.A. LLM4PLC: Harnessing Large Language Models for Verifiable Programming of PLCs in Industrial Control Systems. In Proceedings of the Proceedings of the 46th International Conference on Software Engineering: Software Engineering in Practice, New York, NY, USA, 2024; ICSE-SEIP '24, p. 192–203. https://doi.org/10.1145/3639477.3639743.
14. Tizaoui, T.; Tan, R. Towards a benchmark dataset for large language models in the context of process automation. *Digital Chemical Engineering* **2024**, p. 100186.
15. Xia, Y.; Zhang, J.; Jazdi, N.; Weyrich, M. Incorporating Large Language Models into Production Systems for Enhanced Task Automation and Flexibility. *arXiv preprint arXiv:2407.08550* **2024**.
16. Ogundare, O.; Madasu, S.; Wiggins, N. Industrial Engineering with Large Language Models: A case study of ChatGPT's performance on Oil & Gas problems. In Proceedings of the 2023 11th International Conference on Control, Mechatronics and Automation (ICCMA). IEEE, 2023, pp. 458–461.

17. Rahman, S.A.; Chawla, S.; Yaqot, M.; Menezes, B. Leveraging Large Language Models for Supply Chain Management Optimization: A Case Study. In Proceedings of the International Conference on Innovative Intelligent Industrial Production and Logistics. Springer, 2024, pp. 175–197.
18. Li, B.; Mellou, K.; Zhang, B.; Pathuri, J.; Menache, I. Large language models for supply chain optimization. *arXiv preprint arXiv:2307.03875* **2023**.
19. Raman, R.; Sreenivasan, A.; Suresh, M.; Gunasekaran, A.; Nedungadi, P. AI-driven education: a comparative study on ChatGPT and Bard in supply chain management contexts. *Cogent Business & Management* **2024**, *11*, 2412742.
20. Meyer, L.P.; Frey, J.; Junghanns, K.; Brei, F.; Bulert, K.; Gründer-Fahrer, S.; Martin, M. Developing a Scalable Benchmark for Assessing Large Language Models in Knowledge Graph Engineering, 2023, [[arXiv:cs.AI/2308.16622](https://arxiv.org/abs/cs.AI/2308.16622)].
21. bench authors, B. Beyond the Imitation Game: Quantifying and extrapolating the capabilities of language models. *Transactions on Machine Learning Research* **2023**.
22. Azanza, M.; Lamancha, B.P.; Pizarro, E. Tracking the Moving Target: A Framework for Continuous Evaluation of LLM Test Generation in Industry, 2025, [[arXiv:cs.SE/2504.18985](https://arxiv.org/abs/cs.SE/2504.18985)].
23. Shah, N.; Genc, Z.; Araci, D. StackEval: Benchmarking LLMs in Coding Assistance. *Advances in Neural Information Processing Systems* **2024**, *37*, 36976–36994.
24. Zhang, Q.; Fang, C.; Xie, Y.; Zhang, Y.; Yang, Y.; Sun, W.; Yu, S.; Chen, Z. A Survey on Large Language Models for Software Engineering, 2024, [[arXiv:cs.SE/2312.15223](https://arxiv.org/abs/cs.SE/2312.15223)].
25. Bell, R.; Longshore, R.; Madachy, R. Introducing SysEngBench: A novel benchmark for assessing large language models in systems engineering. Technical report, Acquisition Research Program, 2024.
26. Hu, Y.; Goktas, Y.; Yellamati, D.D.; De Tassigny, C. The use and misuse of pre-trained generative large language models in reliability engineering. In Proceedings of the 2024 Annual Reliability and Maintainability Symposium (RAMS). IEEE, 2024, pp. 1–7.
27. Vendrow, J.; Vendrow, E.; Beery, S.; Madry, A. Do Large Language Model Benchmarks Test Reliability? *arXiv preprint arXiv:2502.03461* **2025**.
28. Liu, Y.; Yao, Y.; Ton, J.F.; Zhang, X.; Guo, R.; Cheng, H.; Klochkov, Y.; Taufiq, M.F.; Li, H. Trustworthy LLMs: a Survey and Guideline for Evaluating Large Language Models' Alignment, 2024, [[arXiv:cs.AI/2308.05374](https://arxiv.org/abs/cs.AI/2308.05374)].
29. Irvin, J.A.; Liu, E.R.; Chen, J.C.; Dormoy, I.; Kim, J.; Khanna, S.; Zheng, Z.; Ermon, S. TEOChat: A Large Vision-Language Assistant for Temporal Earth Observation Data, 2024. _eprint: 2410.06234.
30. Xiong, Z.; Zhang, F.; Wang, Y.; Shi, Y.; Zhu, X.X. EarthNets: Empowering artificial intelligence for Earth observation. *IEEE Geoscience and Remote Sensing Magazine* **2024**, pp. 2–36. <https://doi.org/10.1109/MGRS.2024.3466998>.
31. Zhang, C.; Wang, S. Good at captioning, bad at counting: Benchmarking gpt-4v on earth observation data. *arXiv preprint arXiv:2401.17600* **2024**.
32. Li, W.; Yao, D.; Zhao, R.; Chen, W.; Xu, Z.; Luo, C.; Gong, C.; Jing, Q.; Tan, H.; Bi, J. STBench: Assessing the Ability of Large Language Models in Spatio-Temporal Analysis. *arXiv preprint arXiv:2406.19065* **2024**.
33. Sapkota, R.; Qureshi, R.; Hassan, S.Z.; Shutske, J.; Shoman, M.; Sajjad, M.; Dharejo, F.A.; Paudel, A.; Li, J.; Meng, Z.; et al. Multi-modal LLMs in agriculture: A comprehensive review. *Authorea Preprints* **2024**. Publisher: Authorea.
34. Danish, M.S.; Munir, M.A.; Shah, S.R.A.; Kuckreja, K.; Khan, F.S.; Fraccaro, P.; Lacoste, A.; Khan, S. GEOBench-VLM: Benchmarking Vision-Language Models for Geospatial Tasks, 2024. _eprint: 2411.19325.
35. Roberts, J.; Lüddecke, T.; Sheikh, R.; Han, K.; Albanie, S. Charting new territories: Exploring the geographic and geospatial capabilities of multimodal llms. In Proceedings of the Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2024, pp. 554–563.
36. Liu, X.; Lian, Z. RSUniVLM: A Unified Vision Language Model for Remote Sensing via Granularity-oriented Mixture of Experts, 2024. _eprint: 2412.05679.
37. Lin, C.; Lyu, H.; Xu, X.; Luo, J. INS-MMBench: A Comprehensive Benchmark for Evaluating LVLMS' Performance in Insurance. *arXiv preprint arXiv:2406.09105* **2024**.
38. Hendrycks, D.; Burns, C.; Basart, S.; Zou, A.; Mazeika, M.; Song, D.; Steinhardt, J. Measuring massive multitask language understanding. *arXiv preprint arXiv:2009.03300* **2020**.
39. Lu, P.; Mishra, S.; Xia, T.; Qiu, L.; Chang, K.W.; Zhu, S.C.; Tafjord, O.; Clark, P.; Kalyan, A. Learn to explain: Multimodal reasoning via thought chains for science question answering. *Advances in Neural Information Processing Systems* **2022**, *35*, 2507–2521.

40. Fu, D.; Guo, R.; Khalighinejad, G.; Liu, O.; Dhingra, B.; Yogatama, D.; Jia, R.; Neiswanger, W. Isobench: Benchmarking multimodal foundation models on isomorphic representations. *arXiv preprint arXiv:2404.01266* **2024**.
41. Jiang, Z.; Yang, Z.; Chen, J.; Du, Z.; Wang, W.; Xu, B.; Tang, J. VisScience: An Extensive Benchmark for Evaluating K12 Educational Multi-modal Scientific Reasoning. *arXiv preprint arXiv:2409.13730* **2024**.
42. Rein, D.; Hou, B.L.; Stickland, A.C.; Petty, J.; Pang, R.Y.; Dirani, J.; Michael, J.; Bowman, S.R. Gpqa: A graduate-level google-proof q&a benchmark. In Proceedings of the First Conference on Language Modeling, 2024.
43. Anand, A.; Kapuriya, J.; Singh, A.; Saraf, J.; Lal, N.; Verma, A.; Gupta, R.; Shah, R. Mm-phyqa: Multimodal physics question-answering with multi-image cot prompting. In Proceedings of the Pacific-Asia Conference on Knowledge Discovery and Data Mining. Springer, 2024, pp. 53–64.
44. Mirza, A.; Alampara, N.; Kunchapu, S.; Ríos-García, M.; Emoekabu, B.; Krishnan, A.; Gupta, T.; Schilling-Wilhelmi, M.; Okereke, M.; Aneesh, A.; et al. Are large language models superhuman chemists? *arXiv preprint arXiv:2404.01475* **2024**.
45. Zhu, S.; Liu, X.; Khalighinejad, G. ChemQA: a Multimodal Question-and-Answering Dataset on Chemistry Reasoning. <https://huggingface.co/datasets/shangzhu/ChemQA>, 2024.
46. Guo, T.; Nan, B.; Liang, Z.; Guo, Z.; Chawla, N.; Wiest, O.; Zhang, X.; et al. What can large language models do in chemistry? a comprehensive benchmark on eight tasks. *Advances in Neural Information Processing Systems* **2023**, 36, 59662–59688.
47. Yu, B.; Baker, F.N.; Chen, Z.; Ning, X.; Sun, H. Llamol: Advancing large language models for chemistry with a large-scale, comprehensive, high-quality instruction tuning dataset. *arXiv preprint arXiv:2402.09391* **2024**.
48. Zaki, M.; Krishnan, N.; et al. Mascqa: A question answering dataset for investigating materials science knowledge of large language models. *arXiv preprint arXiv:2308.09115* **2023**.
49. Rubungo, A.N.; Li, K.; Hattrick-Simpers, J.; Dieng, A.B. LLM4Mat-bench: benchmarking large language models for materials property prediction. *arXiv preprint arXiv:2411.00177* **2024**.
50. Cao, H.; Shao, Y.; Liu, Z.; Liu, Z.; Tang, X.; Yao, Y.; Li, Y. Presto: progressive pretraining enhances synthetic chemistry outcomes. *arXiv preprint arXiv:2406.13193* **2024**.
51. Liu, X.; Guo, Y.; Li, H.; Liu, J.; Huang, S.; Ke, B.; Lv, J. DrugLLM: Open Large Language Model for Few-shot Molecule Generation. *arXiv preprint arXiv:2405.06690* **2024**.
52. Rasp, S.; Hoyer, S.; Merose, A.; Langmore, I.; Battaglia, P.; Russell, T.; Sanchez-Gonzalez, A.; Yang, V.; Carver, R.; Agrawal, S.; et al. Weatherbench 2: A benchmark for the next generation of data-driven global weather models. *Journal of Advances in Modeling Earth Systems* **2024**, 16, e2023MS004019.
53. Chen, J.; Zhou, P.; Hua, Y.; Chong, D.; Cao, M.; Li, Y.; Yuan, Z.; Zhu, B.; Liang, J. Vision-language models meet meteorology: Developing models for extreme weather events detection with heatmaps. *arXiv preprint arXiv:2406.09838* **2024**.
54. Ma, C.; Hua, Z.; Anderson-Frey, A.; Iyer, V.; Liu, X.; Qin, L. WeatherQA: Can Multimodal Language Models Reason about Severe Weather? *arXiv preprint arXiv:2406.11217* **2024**.
55. Li, H.; Wang, Z.; Wang, J.; Lau, A.K.H.; Qu, H. Cllmate: A multimodal llm for weather and climate events forecasting. *arXiv preprint arXiv:2409.19058* **2024**.
56. Sun, Y.; Wang, C.; Peng, Y. Unleashing the potential of large language model: Zero-shot vqa for flood disaster scenario. In Proceedings of the Proceedings of the 4th International Conference on Artificial Intelligence and Computer Engineering, 2023, pp. 368–373.
57. Rawat, R. Disasterqa: A benchmark for assessing the performance of llms in disaster response. *arXiv preprint arXiv:2410.20707* **2024**.
58. Patel, Z.B.; Bachwana, Y.; Sharma, N.; Guttikunda, S.; Batra, N. VayuBuddy: an LLM-Powered Chatbot to Democratize Air Quality Insights. *arXiv preprint arXiv:2411.12760* **2024**.
59. Pafilis, E.; Frankild, S.P.; Fanini, L.; Faulwetter, S.; Pavloudi, C.; Vasileiadou, A.; Arvanitidis, C.; Jensen, L.J. The SPECIES and ORGANISMS resources for fast and accurate identification of taxonomic names in text. *PloS one* **2013**, 8, e65390.
60. Abdelmageed, N.; Löffler, F.; Feddoul, L.; Algergawy, A.; Samuel, S.; Gaikwad, J.; Kazem, A.; König-Ries, B. BiodivNERE: Gold standard corpora for named entity recognition and relation extraction in the biodiversity domain. *Biodiversity Data Journal* **2022**, 10, e89481.
61. Gema, A.P.; Leang, J.O.J.; Hong, G.; Devoto, A.; Mancino, A.C.M.; Saxena, R.; He, X.; Zhao, Y.; Du, X.; Madani, M.R.G.; et al. Are We Done with MMLU? *arXiv preprint arXiv:2406.04127* **2024**.

62. Richard, A.M.; Huang, R.; Waidyanatha, S.; Shinn, P.; Collins, B.J.; Thillainadarajah, I.; Grulke, C.M.; Williams, A.J.; Lougee, R.R.; Judson, R.S.; et al. The Tox21 10K compound library: collaborative chemistry advancing toxicology. *Chemical Research in Toxicology* **2020**, *34*, 189–216.
63. Kim, S.; Thiessen, P.A.; Bolton, E.E.; Chen, J.; Fu, G.; Gindulyte, A.; Han, L.; He, J.; He, S.; Shoemaker, B.A.; et al. PubChem substance and compound databases. *Nucleic acids research* **2016**, *44*, D1202–D1213.
64. Jin, W.; Coley, C.; Barzilay, R.; Jaakkola, T. Predicting organic reaction outcomes with weisfeiler-lehman network. *Advances in neural information processing systems* **2017**, *30*.
65. Edwards, C.; Zhai, C.; Ji, H. Text2mol: Cross-modal molecule retrieval with natural language queries. In Proceedings of the Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, 2021, pp. 595–607.
66. Irwin, J.J.; Tang, K.G.; Young, J.; Dandarchuluun, C.; Wong, B.R.; Khurelbaatar, M.; Moroz, Y.S.; Mayfield, J.; Sayle, R.A. ZINC20—a free ultralarge-scale chemical database for ligand discovery. *Journal of chemical information and modeling* **2020**, *60*, 6065–6073.
67. Davies, M.; Nowotka, M.; Papadatos, G.; Dedman, N.; Gaulton, A.; Atkinson, F.; Bellis, L.; Overington, J.P. ChEMBL web services: streamlining access to drug discovery data and utilities. *Nucleic acids research* **2015**, *43*, W612–W620.
68. Rasp, S.; Dueben, P.D.; Scher, S.; Weyn, J.A.; Mouatadid, S.; Thuerey, N. WeatherBench: a benchmark data set for data-driven weather forecasting. *Journal of Advances in Modeling Earth Systems* **2020**, *12*, e2020MS002203.
69. Li, S.; Yang, W.; Zhang, P.; Xiao, X.; Cao, D.; Qin, Y.; Zhang, X.; Zhao, Y.; Bogdan, P. ClimateLLM: Efficient Weather Forecasting via Frequency-Aware Large Language Models. *arXiv preprint arXiv:2502.11059* **2025**.
70. Sachdeva, E.; Agarwal, N.; Chundi, S.; Roelofs, S.; Li, J.; Kochenderfer, M.; Choi, C.; Dariush, B. Rank2tell: A multimodal driving dataset for joint importance ranking and reasoning. In Proceedings of the Proceedings of the IEEE/CVF winter conference on applications of computer vision, 2024, pp. 7513–7522.
71. Ding, X.; Han, J.; Xu, H.; Liang, X.; Zhang, W.; Li, X. Holistic autonomous driving understanding by bird's-eye-view injected multi-modal large models. In Proceedings of the Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2024, pp. 13668–13677.
72. Qian, T.; Chen, J.; Zhuo, L.; Jiao, Y.; Jiang, Y.G. NuScenes-QA: A Multi-Modal Visual Question Answering Benchmark for Autonomous Driving Scenario. *Proceedings of the AAAI Conference on Artificial Intelligence* **2024**, *38*, 4542–4550. Number: 5, <https://doi.org/10.1609/aaai.v38i5.28253>.
73. Tong, S.; Brown, E.; Wu, P.; Woo, S.; Middepogu, M.; Akula, S.C.; Yang, J.; Yang, S.; Iyer, A.; Pan, X.; et al. Cambrian-1: A fully open, vision-centric exploration of multimodal llms. *arXiv preprint arXiv:2406.16860* **2024**.
74. Zhou, Q.; Chen, S.; Wang, Y.; Xu, H.; Du, W.; Zhang, H.; Du, Y.; Tenenbaum, J.B.; Gan, C. HAZARD Challenge: Embodied Decision Making in Dynamically Changing Environments. *Arxiv* **2024**. Univ Massachusetts Amherst Peking Univ MIT MIT.
75. Yang, Z.; Jia, X.; Li, H.; Yan, J. LLM4Drive: A Survey of Large Language Models for Autonomous Driving. *Arxiv* **2024**. OpenDriveLab.
76. Kong, Q.; Kawana, Y.; Saini, R.; Kumar, A.; Pan, J.; Gu, T.; Ozao, Y.; Opra, B.; Sato, Y.; Kobori, N. WTS: A Pedestrian-Centric Traffic Video Dataset for Fine-Grained Spatial-Temporal Understanding. 2025, Vol. 15134, pp. 1–18. Woven Toyota, https://doi.org/10.1007/978-3-031-73116-7_1.
77. Malla, S.; Choi, C.; Dwivedi, I.; Choi, J.H.; Li, J. Drama: Joint risk localization and captioning in driving. In Proceedings of the Proceedings of the IEEE/CVF winter conference on applications of computer vision, 2023, pp. 1043–1052.
78. Yang, J.; Gao, S.; Qiu, Y.; Chen, L.; Li, T.; Dai, B.; Chitta, K.; Wu, P.; Zeng, J.; Luo, P.; et al. Generalized Predictive Model for Autonomous Driving. 2024, pp. 14662–14672.
79. Nie, M.; Peng, R.; Wang, C.; Cai, X.; Han, J.; Xu, H.; Zhang, L. Reason2Drive: Towards Interpretable and Chain-Based Reasoning for Autonomous Driving. In Proceedings of the Computer Vision – ECCV 2024; Leonardi, A.; Ricci, E.; Roth, S.; Russakovsky, O.; Sattler, T.; Varol, G., Eds., Cham, 2025; pp. 292–308. https://doi.org/10.1007/978-3-031-73347-5_17.
80. Sima, C.; Renz, K.; Chitta, K.; Chen, L.; Zhang, H.; Xie, C.; Beißwenger, J.; Luo, P.; Geiger, A.; Li, H. DriveLM: Driving with Graph Visual Question Answering. In Proceedings of the Computer Vision – ECCV 2024; Leonardi, A.; Ricci, E.; Roth, S.; Russakovsky, O.; Sattler, T.; Varol, G., Eds., Cham, 2025; pp. 256–274. https://doi.org/10.1007/978-3-031-72943-0_15.
81. Wu, D.; Han, W.; Wang, T.; Liu, Y.; Zhang, X.; Shen, J. Language Prompt for Autonomous Driving, 2023. arXiv:2309.04379 [cs], <https://doi.org/10.48550/arXiv.2309.04379>.

82. Xiao, Z.; Wang, Q.; Pearce, H.; Chen, S. Logic Meets Magic: LLMs Cracking Smart Contract Vulnerabilities, 2025, [arXiv:cs.CR/2501.07058].
83. Wei, Z.; Sun, J.; Zhang, Z.; Zhang, X.; Li, M.; Hou, Z. LLM-SmartAudit: Advanced Smart Contract Vulnerability Detection, 2024, [arXiv:cs.CR/2410.09381].
84. Zhang, L.; Li, K.; Sun, K.; Wu, D.; Liu, Y.; Tian, H.; Liu, Y. ACFIX: Guiding LLMs with Mined Common RBAC Practices for Context-Aware Repair of Access Control Vulnerabilities in Smart Contracts, 2024, [arXiv:cs.SE/2403.06838].
85. Roumeliotis, K.I.; Tselikas, N.D.; Nasiopoulos, D.K. LLMs and NLP Models in Cryptocurrency Sentiment Analysis: A Comparative Classification Study. *Big Data and Cognitive Computing* **2024**, *8*, 63. <https://doi.org/10.3390/bdcc8060063>.
86. Makri, E.; Palaiokrassas, G.; Bouraga, S.; Polychroniadou, A.; Tassioulas, L. Ethereum Price Prediction Employing Large Language Models for Short-term and Few-shot Forecasting, 2025, [arXiv:cs.AI/2503.23190].
87. Wang, Q.; Gao, Y.; Tang, Z.; Luo, B.; Chen, N.; He, B. Exploring LLM Cryptocurrency Trading Through Fact-Subjectivity Aware Reasoning, 2025, [arXiv:cs.MA/2410.12464].
88. He, Z.; Li, Z.; Yang, S.; Ye, H.; Qiao, A.; Zhang, X.; Luo, X.; Chen, T. Large Language Models for Blockchain Security: A Systematic Literature Review, 2025, [arXiv:cs.CR/2403.14280].
89. Trozze, A.; Davies, T.; Kleinberg, B. Large Language Models in Cryptocurrency Securities Cases: Can a GPT Model Meaningfully Assist Lawyers?, 2024, [arXiv:cs.AI/2308.06032].
90. Cui, C.; Ma, Y.; Cao, X.; Ye, W.; Zhou, Y.; Liang, K.; Chen, J.; Lu, J.; Yang, Z.; Liao, K.D.; et al. A survey on multimodal large language models for autonomous driving. In Proceedings of the Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision, 2024, pp. 958–979.
91. Shi, Y.; Jiang, K.; Li, J.; Qian, Z.; Wen, J.; Yang, M.; Wang, K.; Yang, D. Grid-centric traffic scenario perception for autonomous driving: A comprehensive review. *IEEE Transactions on Neural Networks and Learning Systems* **2024**.
92. Tong, P.; Brown, E.; Wu, P.; Woo, S.; IYER, A.J.V.; Akula, S.C.; Yang, S.; Yang, J.; Middepogu, M.; Wang, Z.; et al. Cambrian-1: A fully open, vision-centric exploration of multimodal llms. *Advances in Neural Information Processing Systems* **2024**, *37*, 87310–87356.
93. Yang, J.; Gao, S.; Qiu, Y.; Chen, L.; Li, T.; Dai, B.; Chitta, K.; Wu, P.; Zeng, J.; Luo, P.; et al. Generalized predictive model for autonomous driving. In Proceedings of the Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2024, pp. 14662–14672.
94. Shah, R.S.; Chawla, K.; Eidnani, D.; Shah, A.; Du, W.; Chava, S.; Raman, N.; Smiley, C.; Chen, J.; Yang, D. When flue meets flang: Benchmarks and large pre-trained language model for financial domain, [2211.00083]. arXiv preprint.
95. Guan, W.; Cao, J.; Qian, S.; Gao, J.; Ouyang, C. LogLLM: Log-based Anomaly Detection Using Large Language Models, 2025, [arXiv:cs.SE/2411.08561].
96. Sinha, R.; Elhafsi, A.; Agia, C.; Foutter, M.; Schmerling, E.; Pavone, M. Real-Time Anomaly Detection and Reactive Planning with Large Language Models, 2024, [arXiv:cs.RO/2407.08735].
97. Zhang, R.; Jiang, D.; Zhang, Y.; Lin, H.; Guo, Z.; Qiu, P.; Zhou, A.; Lu, P.; Chang, K.W.; Qiao, Y.; et al. Mathverse: Does your multi-modal llm truly see the diagrams in visual math problems? In Proceedings of the European Conference on Computer Vision. Springer, 2024, pp. 169–186.
98. Fan, J.; Martinson, S.; Wang, E.Y.; Hausknecht, K.; Brenner, J.; Liu, D.; Peng, N.; Wang, C.; Brenner, M. HARDMATH: A Benchmark Dataset for Challenging Problems in Applied Mathematics. In Proceedings of the The 4th Workshop on Mathematical Reasoning and AI at NeurIPS'24, 2024.
99. Liu, H.; Zhang, Y.; Luo, Y.; Yao, A.C.C. Augmenting Math Word Problems via Iterative Question Composing. *Arxiv* **2024**. Shanghai Qizhi Inst Beijing Univ Posts & Telecommun.
100. Liang, Z.; Yu, D.; Yu, W.; Yao, W.; Zhang, Z.; Zhang, X.; Yu, D. Mathchat: Benchmarking mathematical reasoning and instruction following in multi-turn interactions. *arXiv preprint arXiv:2405.19444* **2024**.
101. Zhang, Z.; Xu, L.; Jiang, Z.; Hao, H.; Wang, R. Multiple-choice questions are efficient and robust llm evaluators. 2024d. URL <https://doi.org/10.48550/arXiv.2405>.
102. Anantheswaran, U.; Gupta, H.; Scaria, K.; Verma, S.; Baral, C.; Mishra, S. Cutting Through the Noise: Boosting LLM Performance on Math Word Problems. *arXiv preprint arXiv:2406.15444* **2024**.
103. Yang, K.; Swope, A.; Gu, A.; Chalamala, R.; Song, P.; Yu, S.; Godil, S.; Prenger, R.J.; Anandkumar, A. Leandojo: Theorem proving with retrieval-augmented language models. *Advances in Neural Information Processing Systems* **2023**, *36*, 21573–21612.

104. Zhao, Y.; Liu, H.; Long, Y.; Zhang, R.; Zhao, C.; Cohan, A. FinanceMATH: Knowledge-Intensive Math Reasoning in Finance Domains. In Proceedings of the Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), 2024, pp. 12841–12858.
105. Pezeshkpour, P.; Hruschka, E. Large Language Models Sensitivity to The Order of Options in Multiple-Choice Questions. In Proceedings of the Findings of the Association for Computational Linguistics: NAACL 2024; Duh, K.; Gomez, H.; Bethard, S., Eds., Mexico City, Mexico, 2024; pp. 2006–2017. <https://doi.org/10.18653/v1/2024.findings-naacl.130>.
106. LaBelle, E. Monte Carlo Tree Search Applications to Neural Theorem Proving. PhD thesis, Massachusetts Institute of Technology, 2024.
107. Yuan, Z.; Wang, K.; Zhu, S.; Yuan, Y.; Zhou, J.; Zhu, Y.; Wei, W. Finllms: A framework for financial reasoning dataset generation with large language models. *IEEE Transactions on Big Data* **2024**.
108. Jin, Y.; Choi, M.; Verma, G.; Wang, J.; Kumar, S. Mm-soc: Benchmarking multimodal large language models in social media platforms. *arXiv preprint arXiv:2402.14154* **2024**.
109. Chen, Y.; Yan, S.; Guo, Q.; Jia, J.; Li, Z.; Xiao, Y. HOTVCOM: Generating Buzzworthy Comments for Videos. In Proceedings of the Findings of the Association for Computational Linguistics: ACL 2024; Ku, L.W.; Martins, A.; Srikumar, V., Eds., Bangkok, Thailand, 2024; pp. 2198–2224. <https://doi.org/10.18653/v1/2024.findings-acl.130>.
110. Chen, Y.; Yan, S.; Zhu, Z.; Li, Z.; Xiao, Y. XMeCap: Meme Caption Generation with Sub-Image Adaptability. In Proceedings of the Proceedings of the 32nd ACM International Conference on Multimedia, New York, NY, USA, 2024; MM '24, pp. 3352–3361. <https://doi.org/10.1145/3664647.3681332>.
111. Shahriar, S.; Dara, R. Priv-IQ: A Benchmark and Comparative Evaluation of Large Multimodal Models on Privacy Competencies. *AI* **2025**, 6, 29. <https://doi.org/10.3390/ai6020029>.
112. Gallegos, I.O.; et al. Bias and Fairness in Large Language Models: A Survey. *Comput. Linguist.* **2024**, 50, 1097–1179. https://doi.org/10.1162/coli_a_00524.
113. Liu, S.; et al. CultureVLM: Characterizing and Improving Cultural Understanding of Vision-Language Models for over 100 Countries. *arXiv* **2025**, [arXiv:2501.01282]. <https://doi.org/10.48550/arXiv.2501.01282>.
114. Ghaboura, S.; et al. Time Travel: A Comprehensive Benchmark to Evaluate LMMs on Historical and Cultural Artifacts. *arXiv* **2025**, [arXiv:2502.14865]. <https://doi.org/10.48550/arXiv.2502.14865>.
115. Hu, H.; et al. EmoBench-M: Benchmarking Emotional Intelligence for Multimodal Large Language Models. *arXiv* **2025**, [arXiv:2502.04424]. <https://doi.org/10.48550/arXiv.2502.04424>.
116. Chen, Y.; Yan, S.; Liu, S.; Li, Y.; Xiao, Y. EmotionQueen: A Benchmark for Evaluating Empathy of Large Language Models. In Proceedings of the Findings of the Association for Computational Linguistics: ACL 2024; Ku, L.W.; Martins, A.; Srikumar, V., Eds., Bangkok, Thailand, 2024; pp. 2149–2176. <https://doi.org/10.18653/v1/2024.findings-acl.128>.
117. Yang, L.; et al. EditWorld: Simulating World Dynamics for Instruction-Following Image Editing. *CoRR* **2024**. Accessed: Feb. 24, 2025.
118. He, Y.; et al. LLMs Meet Multimodal Generation and Editing: A Survey. *CoRR* **2024**. Accessed: Feb. 25, 2025.
119. Cao, J.; Liu, Y.; Shi, Y.; Ding, K.; Jin, L. WenMind: A Comprehensive Benchmark for Evaluating Large Language Models in Chinese Classical Literature and Language Arts. *Adv. Neural Inf. Process. Syst.* **2025**, 37, 51358–51410.
120. Li, J.; et al. The Music Maestro or The Musically Challenged, A Massive Music Evaluation Benchmark for Large Language Models. In Proceedings of the Findings of the Association for Computational Linguistics: ACL 2024; Ku, L.W.; Martins, A.; Srikumar, V., Eds., Bangkok, Thailand, 2024; pp. 3246–3257. <https://doi.org/10.18653/v1/2024.findings-acl.194>.
121. Weck, B.; Manco, I.; Benetos, E.; Quinton, E.; Fazekas, G.; Bogdanov, D. MuChoMusic: Evaluating Music Understanding in Multimodal Audio-Language Models. *CoRR* **2024**. Accessed: Feb. 25, 2025.
122. Zhou, Z.; et al. Can LLMs 'Reason' in Music? An Evaluation of LLMs' Capability of Music Understanding and Generation. *CoRR* **2024**. Accessed: Feb. 25, 2025.
123. Hachmeier, S.; Jäschke, R. A Benchmark and Robustness Study of In-Context-Learning with Large Language Models in Music Entity Detection. In Proceedings of the Proceedings of the 31st International Conference on Computational Linguistics; Rambow, O.; Wanner, L.; Apidianaki, M.; Al-Khalifa, H.; Eugenio, B.D.; Schockaert, S., Eds., Abu Dhabi, UAE, 2025; pp. 9845–9859. Accessed: Feb. 25, 2025.
124. Wang, Z.; et al. MuChin: a chinese colloquial description benchmark for evaluating language models in the field of music. In Proceedings of the Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence, Jeju, Korea, 2024; IJCAI '24, pp. 7771–7779. <https://doi.org/10.24963/ijcai.2024/860>.

125. Zhao, Y.; et al. CityEQA: A Hierarchical LLM Agent on Embodied Question Answering Benchmark in City Space. *arXiv* **2025**, [arXiv:2502.12532]. <https://doi.org/10.48550/arXiv.2502.12532>.
126. Zhang, X.; et al. TransportationGames: Benchmarking Transportation Knowledge of (Multimodal) Large Language Models. *CoRR* **2024**. Accessed: Feb. 24, 2025.
127. Nie, T.; Sun, J.; Ma, W. Exploring the Roles of Large Language Models in Reshaping Transportation Systems: A Survey, Framework, and Roadmap. *arXiv* **2025**, [arXiv:2503.21411]. <https://doi.org/10.48550/arXiv.2503.21411>.
128. Zhang, W.; Han, J.; Xu, Z.; Ni, H.; Liu, H.; Xiong, H. Urban Foundation Models: A Survey. In Proceedings of the Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, New York, NY, USA, 2024; KDD '24, pp. 6633–6643. <https://doi.org/10.1145/3637528.3671453>.
129. Zheng, Y.; et al. UrbanPlanBench: A Comprehensive Assessment of Urban Planning Abilities in Large Language Models **2024**. Accessed: Feb. 25, 2025.
130. Feng, J.; et al. CityBench: Evaluating the Capabilities of Large Language Model as World Model. *CoRR* **2024**. Accessed: Feb. 25, 2025.
131. Ji, J.; Chen, Y.; Jin, M.; Xu, W.; Hua, W.; Zhang, Y. MoralBench: Moral Evaluation of LLMs. *CoRR* **2024**. Accessed: Feb. 27, 2025.
132. Yan, B.; Zhang, J.; Chen, Z.; Shan, S.; Chen, X. M³oralBench: A MultiModal Moral Benchmark for LVLMS. *arXiv* **2024**, [arXiv:2412.20718]. <https://doi.org/10.48550/arXiv.2412.20718>.
133. Marraffini, G.F.G.; Cotton, A.; Hsueh, N.F.; Fridman, A.; Wisznia, J.; Corro, L.D. The Greatest Good Benchmark: Measuring LLMs' Alignment with Utilitarian Moral Dilemmas. In Proceedings of the Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing; Al-Onaizan, Y.; Bansal, M.; Chen, Y.N., Eds., Miami, Florida, USA, 2024; pp. 21950–21959. <https://doi.org/10.18653/v1/2024.emnlp-main.1224>.
134. Yao, J.; et al. Value Compass Leaderboard: A Platform for Fundamental and Validated Evaluation of LLMs Values. *arXiv* **2025**, [arXiv:2501.07071]. <https://doi.org/10.48550/arXiv.2501.07071>.
135. Li, Y.; Huang, Y.; Lin, Y.; Wu, S.; Wan, Y.; Sun, L. I Think, Therefore I am: Benchmarking Awareness of Large Language Models Using AwareBench. *arXiv* **2024**, [arXiv:2401.17882]. <https://doi.org/10.48550/arXiv.2401.17882>.
136. Yang, Y.; Xu, Y.; Huang, C.; Jurgensen, J.; Hu, H. InterIDEAS: An LLM and Expert-Enhanced Dataset for Philosophical Intertextuality **2024**. Accessed: Feb. 27, 2025.
137. Trepczyński, M. Religion, Theology, and Philosophical Skills of LLM–Powered Chatbots. *Disput. Philos. Int. J. Philos. Relig.* **2023**, *25*, 19–36.
138. Deng, C.; et al. Deconstructing The Ethics of Large Language Models from Long-standing Issues to New-emerging Dilemmas. *CoRR* **2024**. Accessed: Feb. 24, 2025.
139. Wang, H.; et al. Piecing It All Together: Verifying Multi-Hop Multimodal Claims. In Proceedings of the Proceedings of the 31st International Conference on Computational Linguistics; Rambow, O.; Wanner, L.; Apidianaki, M.; Al-Khalifa, H.; Eugenio, B.D.; Schockaert, S., Eds., Abu Dhabi, UAE, 2025; pp. 7453–7469. Accessed: Feb. 24, 2025.
140. Jin, Y.; et al. AgentReview: Exploring Peer Review Dynamics with LLM Agents. In Proceedings of the Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing; Al-Onaizan, Y.; Bansal, M.; Chen, Y.N., Eds., Miami, Florida, USA, 2024; pp. 1208–1226. <https://doi.org/10.18653/v1/2024.emnlp-main.70>.
141. Song, S.; et al. MOSABench: Multi-Object Sentiment Analysis Benchmark for Evaluating Multimodal Large Language Models Understanding of Complex Image. *arXiv* **2024**, [arXiv:2412.00060]. <https://doi.org/10.48550/arXiv.2412.00060>.
142. Adilazuarda, M.F.; et al. Towards Measuring and Modeling 'Culture' in LLMs: A Survey. In Proceedings of the Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing; Al-Onaizan, Y.; Bansal, M.; Chen, Y.N., Eds., Miami, Florida, USA, 2024; pp. 15763–15784. <https://doi.org/10.18653/v1/2024.emnlp-main.882>.
143. Chen, Y.; Xiao, Y. Recent Advancement of Emotion Cognition in Large Language Models. *CoRR* **2024**. Accessed: Feb. 24, 2025.
144. Chakrabarty, T.; Laban, P.; Agarwal, D.; Muresan, S.; Wu, C.S. Art or Artifice? Large Language Models and the False Promise of Creativity. In Proceedings of the Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems, New York, NY, USA, 2024; CHI '24, pp. 1–34. <https://doi.org/10.1145/3613904.3642731>.

145. Cheng, C.Y.I.; Hale, S.A. Beyond English: Evaluating Automated Measurement of Moral Foundations in Non-English Discourse with a Chinese Case Study. *arXiv* **2025**, [arXiv:2502.02451]. <https://doi.org/10.48550/arXiv.2502.02451>.
146. Bulla, L.; De Giorgis, S.; Mongiovì, M.; Gangemi, A. Large Language Models meet moral values: A comprehensive assessment of moral abilities. *Comput. Hum. Behav. Rep.* **2025**, *17*, 100609. <https://doi.org/10.1016/j.chbr.2025.100609>.
147. Xie, Q.; Han, W.; Zhang, X.; Lai, Y.; Peng, M.; Lopez-Lira, A.; Huang, J. Pixiu: A large language model, instruction data and evaluation benchmark for finance, [2306.05443]. arXiv preprint.
148. Chen, Z.; Chen, W.; Smiley, C.; Shah, S.; Borova, I.; Langdon, D.; Moussa, R.; Beane, M.; Huang, T.H.; Routledge, B.; et al. FinQA: A dataset of numerical reasoning over financial data, [2109.00122]. arXiv preprint.
149. Reddy, V.; Koncel-Kedziorski, R.; Lai, V.D.; Krumdick, M.; Lovering, C.; Tanner, C. DocFinQA: A long-context financial reasoning dataset, [2401.06915]. arXiv preprint.
150. Chen, Z.; Li, S.; Smiley, C.; Ma, Z.; Shah, S.; Wang, W.Y. ConvFinQA: Exploring the chain of numerical reasoning in conversational finance question answering, [2210.03849]. arXiv preprint.
151. Webersinke, N.; Kraus, M.; Bingler, J.A.; Leippold, M. Climatebert: A pretrained language model for climate-related text, [2110.12010]. arXiv preprint.
152. Sharma, S.; Nayak, T.; Bose, A.; Meena, A.K.; Dasgupta, K.; Ganguly, N.; Goyal, P. FinRED: A dataset for relation extraction in financial domain. In Proceedings of the Companion Proceedings of the Web Conference 2022, 2022, pp. 595–597.
153. Wu, X.; Liu, J.; Su, H.; Lin, Z.; Qi, Y.; Xu, C.; Su, J.; Zhong, J.; Wang, F.; Wang, S.; et al. Golden Touchstone: A Comprehensive Bilingual Benchmark for Evaluating Financial Large Language Models, [2411.06272]. arXiv preprint.
154. Subrahmanyam, A. Behavioural finance: A review and synthesis. *European Financial Management* **2008**, *14*, 12–29.
155. Rubbaniy, G.; Khalid, A.A.; Syriopoulos, K.; Polyzos, E. Dynamic returns connectedness: Portfolio hedging implications during the COVID-19 pandemic and the Russia–Ukraine war. *Journal of Futures Markets* **2024**, *44*, 1613–1639.
156. Li, S.; Hoque, H.; Liu, J. Investor sentiment and firm capital structure. *Journal of Corporate Finance* **2023**, *80*, 102426.
157. Karadima, M.; Louri, H. Economic policy uncertainty and non-performing loans: The moderating role of bank concentration. *Finance Research Letters* **2021**, *38*, 101458.
158. Hodbod, A.; Huber, S.; Vasilev, K. Sectoral risk-weights and macroprudential policy. *Journal of Banking and Finance* **2020**, *112*, 105336.
159. Danielsson, J.; James, K.R.; Valenzuela, M.; Zer, I. Model risk of risk models. *Journal of Financial Stability* **2016**, *23*, 79–91.
160. Oehler, A.; Horn, M. Does ChatGPT provide better advice than robo-advisors? *Finance Research Letters* **2024**, *60*, 104898.
161. Dowling, M.; Lucey, B. ChatGPT for (finance) research: The Bananarama conjecture. *Finance Research Letters* **2023**, *53*, 103662.
162. Polyzos, E.; Fotiadis, A.; Huan, T.C. The asymmetric impact of Twitter sentiment and emotions: Impulse response analysis on European tourism firms using micro-data. *Tourism Management* **2024**, *104*, 104909.
163. Kalamara, E.; Turrell, A.; Redl, C.; Kapetanios, G.; Kapadia, S. Making text count: economic forecasting using newspaper text. *Journal of Applied Econometrics* **2022**, *37*, 896–919.
164. Herrera, G.P.; Constantino, M.; Su, J.J.; Naranpanawa, A. Renewable energy stocks forecast using Twitter investor sentiment and deep learning. *Energy Economics* **2022**, *114*, 106285.
165. Dicks, D.; Fulghieri, P. Uncertainty, investor sentiment, and innovation. *The Review of Financial Studies* **2021**, *34*, 1236–1279.
166. Liu, X.Y.; Wang, G.; Yang, H.; Zha, D. Fingpt: Democratizing internet-scale data for financial large language models. *arXiv preprint arXiv:2307.10485* **2023**.
167. Wu, S.; Irsoy, O.; Lu, S.; Dabrovolski, V.; Dredze, M.; Gehrmann, S.; Kambadur, P.; Rosenberg, D.; Mann, G. BloombergGPT: A large language model for finance, [2303.17564]. arXiv preprint.
168. Liu, Y.; Bu, N.; Li, Z.; Zhang, Y.; Zhao, Z. AT-FinGPT: Financial risk prediction via an audio-text large language model. *Finance Research Letters* **2025**, p. 106967.

169. Polyzos, E. Inflation and the war in Ukraine: Evidence using impulse response functions on economic indicators and Twitter sentiment. *Research in International Business and Finance* **2023**, *66*, 102044.
170. Long, S.; Lucey, B.; Xie, Y.; Yarovaya, L. I just like the stock. *The role of Reddit sentiment in the GameStop share rally*. *Financial Review* **2023**, *58*, 19–37.
171. Lyócsa, Š.; Baumöhl, E.; Výrost, T. YOLO trading: Riding with the herd during the GameStop episode. *Finance Research Letters* **2022**, *46*, 102359.
172. Polyzos, E.; Samitas, A.; Kampouris, I. Quantifying market efficiency: Information dissemination through social media. *Available at SSRN* **2022**, 4082899.
173. Vasileiou, E. Does the short squeeze lead to market abnormality and antileverage effect? Evidence from the GameStop case. *Journal of Economic Studies* **2022**, *49*, 1360–1373.
174. Huang, A.H.; Wang, H.; Yang, Y. FinBERT: A large language model for extracting information from financial text. *Contemporary Accounting Research* **2023**, *40*, 806–841.
175. Loughran, T.; McDonald, B. When is a liability not a liability? Textual analysis, dictionaries, and 10-Ks. *The Journal of Finance* **2011**, *66*, 35–65.
176. Nițoi, M.; Pochea, M.M.; Radu, Ș.C. Unveiling the Sentiment Behind Central Bank Narratives: A Novel Deep Learning Index. *Journal of Behavioral and Experimental Finance* **2023**, *38*, 100809. <https://doi.org/10.1016/j.jbef.2023.100809>.
177. Schimanski, T.; Reding, A.; Reding, N.; Bingler, J.A.; Kraus, M.; Leippold, M. Bridging the gap in ESG measurement: Using NLP to quantify environmental, social, and governance communication. *Finance Research Letters* **2024**, *61*, 104979.
178. Polyzos, E.; Samitas, A.; Katsaiti, M.S. Who is unhappy for Brexit? A machine-learning, agent-based study on financial instability. *International Review of Financial Analysis* **2020**, *72*, 101590.
179. Polyzos, E.; Abdulrahman, K.; Christopoulos, A. Good management or good finances? An agent-based study on the causes of bank failure. *Banks & Bank Systems* **2018**, *13*, Iss. 3, 95–105.
180. Stevens Institute of Technology. Applying Large Language Models to Financial Decision-Making. <https://www.stevens.edu/news/applying-large-language-models-to-financial-decision-making>, 2023. [Accessed 3 March 2025].
181. Ye, J.; Wang, G.; Li, Y.; Deng, Z.; Li, W.; Li, T.; Duan, H.; Huang, Z.; Su, Y.; Wang, B.; et al. Gmai-mmbench: A comprehensive multimodal evaluation benchmark towards general medical ai. *Advances in Neural Information Processing Systems* **2024**, *37*, 94327–94427.
182. Liu, J.; Wang, W.; Su, Y.; Huan, J.; Chen, W.; Zhang, Y.; Li, C.Y.; Chang, K.J.; Xin, X.; Shen, L.; et al. A Spectrum Evaluation Benchmark for Medical Multi-Modal Large Language Models. *arXiv preprint arXiv:2402.11217* **2024**.
183. Keat, K.; Venkatesh, R.; Huang, Y.; Kumar, R.; Tuteja, S.; Sangkuhl, K.; Li, B.; Gong, L.; Whirl-Carrillo, M.; Klein, T.E.; et al. PGxQA: A Resource for Evaluating LLM Performance for Pharmacogenomic QA Tasks. In *Proceedings of the Biocomputing 2025: Proceedings of the Pacific Symposium*. World Scientific, 2024, pp. 229–246.
184. Liu, H.; Wang, H. GenoTEX: A Benchmark for Evaluating LLM-Based Exploration of Gene Expression Data in Alignment with Bioinformaticians. *arXiv preprint arXiv:2406.15341* **2024**.
185. Bai, F.; Du, Y.; Huang, T.; Meng, M.Q.H.; Zhao, B. M3d: Advancing 3d medical image analysis with multi-modal large language models. *arXiv preprint arXiv:2404.00578* **2024**.
186. Bassi, P.R.A.S.; Yavuz, M.C.; Wang, K.; Chen, X.; Li, W.; Decherchi, S.; Cavalli, A.; Yang, Y.; Yuille, A.; Zhou, Z. RadGPT: Constructing 3D Image-Text Tumor Datasets, 2025, [[arXiv:eess.IV/2501.04678](https://arxiv.org/abs/2501.04678)].
187. Mo, S.; Liang, P.P. MultiMed: Massively Multimodal and Multitask Medical Understanding. *arXiv preprint arXiv:2408.12682* **2024**.
188. Sepehri, M.S.; Fabian, Z.; Soltanolkotabi, M.; Soltanolkotabi, M. MediConfusion: Can you trust your AI radiologist? Probing the reliability of multimodal medical foundation models. *arXiv preprint arXiv:2409.15477* **2024**.
189. Neehal, N.; Wang, B.; Debopadhyaya, S.; Dan, S.; Murugesan, K.; Anand, V.; Bennett, K.P. Ctbench: A comprehensive benchmark for evaluating language model capabilities in clinical trial design. *arXiv preprint arXiv:2406.17888* **2024**.
190. Khandekar, N.; Jin, Q.; Xiong, G.; Dunn, S.; Applebaum, S.; Anwar, Z.; Sarfo-Gyamfi, M.; Safranek, C.; Anwar, A.; Zhang, A.; et al. Medcalc-bench: Evaluating large language models for medical calculations. *Advances in Neural Information Processing Systems* **2024**, *37*, 84730–84745.

191. Jiang, J.; Chen, P.; Wang, J.; He, D.; Wei, Z.; Hong, L.; Zong, L.; Wang, S.; Yu, Q.; Ma, Z.; et al. Benchmarking Large Language Models on Multiple Tasks in Bioinformatics NLP with Prompting. *arXiv preprint arXiv:2503.04013* **2025**.
192. Chen, L.; Han, X.; Lin, S.; Mai, H.; Ran, H. TriMedLM: Advancing Three-Dimensional Medical Image Analysis with Multi-Modal LLM. In Proceedings of the 2024 IEEE International Conference on Bioinformatics and Biomedicine (BIBM), 2024, pp. 4505–4512. <https://doi.org/10.1109/BIBM62325.2024.10822809>.
193. Lozano, A.; Nirschl, J.; Burgess, J.; Gupte, S.R.; Zhang, Y.; Unell, A.; Yeung, S. Micro-bench: A microscopy benchmark for vision-language understanding. *Advances in Neural Information Processing Systems* **2024**, *37*, 30670–30685.
194. Sun, Y.; Wu, H.; Zhu, C.; Zheng, S.; Chen, Q.; Zhang, K.; Zhang, Y.; Wan, D.; Lan, X.; Zheng, M.; et al. Pathmmu: A massive multimodal expert-level benchmark for understanding and reasoning in pathology. In Proceedings of the European Conference on Computer Vision. Springer, 2024, pp. 56–73.
195. Burgess, J.; Nirschl, J.J.; Bravo-Sánchez, L.; Lozano, A.; Gupte, S.R.; Galaz-Montoya, J.G.; Zhang, Y.; Su, Y.; Bhowmik, D.; Coman, Z.; et al. MicroVQA: A Multimodal Reasoning Benchmark for Microscopy-Based Scientific Research, 2025, [\[arXiv:cs.CV/2503.13399\]](https://arxiv.org/abs/2503.13399).
196. Chen, Y.; Wang, G.; Ji, Y.; Li, Y.; Ye, J.; Li, T.; Hu, M.; Yu, R.; Qiao, Y.; He, J. SlideChat: A Large Vision-Language Assistant for Whole-Slide Pathology Image Understanding, 2025, [\[arXiv:cs.CV/2410.11761\]](https://arxiv.org/abs/2410.11761).
197. Lozano, A.; Nirschl, J.; Burgess, J.; Gupte, S.R.; Zhang, Y.; Unell, A.; Yeung-Levy, S. μ -Bench: A Vision-Language Benchmark for Microscopy Understanding, 2024, [\[arXiv:cs.CV/2407.01791\]](https://arxiv.org/abs/2407.01791).
198. Wang, X.; Song, D.; Chen, S.; Zhang, C.; Wang, B. LongLLaVA: Scaling Multi-modal LLMs to 1000 Images Efficiently via a Hybrid Architecture. *arXiv preprint arXiv:2409.02889* **2024**.
199. Li, B.; Zhang, Y.; Guo, D.; Zhang, R.; Li, F.; Zhang, H.; Zhang, K.; Zhang, P.; Li, Y.; Liu, Z.; et al. Llava-onevision: Easy visual task transfer. *arXiv preprint arXiv:2408.03326* **2024**.
200. Huang, S.; Dong, L.; Wang, W.; Hao, Y.; Singhal, S.; Ma, S.; Lv, T.; Cui, L.; Mohammed, O.K.; Patra, B.; et al. Language Is Not All You Need: Aligning Perception with Language Models, 2023, [\[arXiv:cs.CL/2302.14045\]](https://arxiv.org/abs/2302.14045).
201. Peng, Z.; Wang, W.; Dong, L.; Hao, Y.; Huang, S.; Ma, S.; Wei, F. Kosmos-2: Grounding Multimodal Large Language Models to the World, 2023, [\[arXiv:cs.CL/2306.14824\]](https://arxiv.org/abs/2306.14824).
202. GLM, T.; Zeng, A.; Xu, B.; Wang, B.; Zhang, C.; Yin, D.; Zhang, D.; Rojas, D.; Feng, G.; Zhao, H.; et al. Chatglm: A family of large language models from glm-130b to glm-4 all tools. *arXiv preprint arXiv:2406.12793* **2024**.
203. Wang, Y.; Liu, Y.; Yu, F.; Huang, C.; Li, K.; Wan, Z.; Che, W.; Chen, H. Cvlue: A new benchmark dataset for chinese vision-language understanding evaluation. In Proceedings of the Proceedings of the AAAI Conference on Artificial Intelligence, 2025, Vol. 39, pp. 8196–8204.
204. Yao, S.; Yu, D.; Zhao, J.; Shafran, I.; Griffiths, T.; Cao, Y.; Narasimhan, K. Tree of thoughts: Deliberate problem solving with large language models. *Advances in Neural Information Processing Systems* **2024**, *36*.
205. Arshad, M.A.; Jubery, T.Z.; Roy, T.; Nassiri, R.; Singh, A.K.; Singh, A.; Hegde, C.; Ganapathysubramanian, B.; Balu, A.; Krishnamurthy, A.; et al. Leveraging Vision Language Models for Specialized Agricultural Tasks. In Proceedings of the 2025 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV). IEEE, 2025, pp. 6320–6329.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.