

TeXpert: A Multi-Level Benchmark for Evaluating L^AT_EX Code Generation by LLMs

Sahil Kale^{1*} Vijaykant Nadadur¹

¹Knowledgeverse AI

{sahil, vrn}@k-v.ai

Abstract

LaTeX’s precision and flexibility in typesetting have made it the gold standard for the preparation of scientific documentation. Large Language Models (LLMs) present a promising opportunity for researchers to produce publication-ready material using LaTeX with natural language instructions, yet current benchmarks completely lack evaluation of this ability. By introducing TeXpert, our benchmark dataset with natural language prompts for generating LaTeX code focused on components of scientific documents across multiple difficulty levels, we conduct an in-depth analysis of LLM performance in this regard and identify frequent error types. Our evaluation across open and closed-source LLMs highlights multiple key findings: LLMs excelling on standard benchmarks perform poorly in LaTeX generation with a significant accuracy drop-off as the complexity of tasks increases; open-source models like DeepSeek v3 and DeepSeek Coder strongly rival closed-source counterparts in LaTeX tasks; and formatting and package errors are unexpectedly prevalent, suggesting a lack of diverse LaTeX examples in the training datasets of most LLMs. Our dataset, code, and model evaluations are available on GitHub.¹

1 Introduction

LaTeX is a highly versatile and widely adopted document preparation system built over the TeX typesetting program (LaTeX). With research-specific advantages including robust handling of mathematical equations, simple formatting commands, and straightforward management of references, it is a popular choice to produce publication-ready scientific material (Bos and McCurley, 2023).

The recent emergence of LLMs across various applications (García-Ferrero et al., 2024; Sherifi

et al., 2024; Zhao et al., 2024) coupled with improved instruction-following ability (Yin et al., 2023; He et al., 2024) prompts an essential research question: "Can LLMs generate publication-ready LaTeX code for components of scientific documents from natural language instructions?". Through this research, we aim to evaluate the capability of LLMs in generating syntactically and logically accurate LaTeX code (which we refer to as accurate LaTeX code generation or simply LaTeX generation) and analyse the main types of errors they encounter.

While certain aspects of LaTeX code generation with LLMs, especially for mathematical content (Zou et al., 2024; Zhang et al., 2024), have been significantly studied, a comprehensive study of LLMs’ LaTeX generation ability for various components commonly used in scientific documents (such as tables, figures, bibliography, etc.) remains unexplored. We believe a comprehensive benchmark for evaluating LLMs on LaTeX generation offers two key benefits: analysing common errors LLMs make in generating LaTeX code can provide format and error-based hints for flagging AI-generated research material (Chamezopoulos et al., 2024), and delineating the complexity of LaTeX tasks that LLMs can reliably perform can greatly reduce researchers’ effort on formatting and typesetting.

In this work, we evaluate a diverse range of closed-source and open-source LLMs on their LaTeX generation capabilities. The main contributions of this paper can be stated as follows:

1. We introduce TeXpert, a benchmark designed to evaluate LLMs in generating accurate LaTeX code from natural language instructions, focused on commands in scientific documents
2. We evaluate popular open and closed-source LLMs on TeXpert by computing the success rate across three difficulty classes
3. We provide comprehensive insights pertaining to LLM limitations in LaTeX generation and identify frequent error types

*Corresponding author. Email: sahil@k-v.ai

¹<https://github.com/knowledge-verse-ai/TeXpert>

2 Related Work

Existing works on the evaluation of LLMs treat LaTeX-based tasks only as a peripheral component or limit their scope to specific output formats. The ability of LLMs to generate mathematical LaTeX equations from various sources has been explored in datasets like MATH (Hendrycks et al., 2021), MathBridge (Jung et al., 2024) and STEM-POM (Zou et al., 2024). Similarly, the STRUC-BENCH dataset (Tang et al., 2024) contains natural language inputs to test LLMs’ LaTeX generation ability specific only to tabular content. The im2latex-100k dataset (Deng et al., 2017) also focuses on the narrow aspect of testing the ability of LLMs to convert images of mathematical formulae into LaTeX code, while Image2struct (Roberts et al., 2024) includes testing vision-language models in extracting structured LaTeX information from images.

A straightforward idea to evaluate the natural language to LaTeX ability of LLMs would be to generate free-to-use LaTeX templates² representing various document styles and formats using textual queries. However, these templates are often too large to be directly generated by large language models (LLMs) and are constrained only to a standard set of basic commands, limiting their applicability in this research. Several instruction-following benchmarks for LLMs evaluate their ability to follow natural language commands (Qin et al., 2024; Chen et al., 2024); however, there is a notable absence of datasets specifically designed to assess models in LaTeX code generation for scientific material.

Identifying and acting upon this need, we present TeXpert, an organised dataset designed to evaluate LLMs’ capability to generate syntactically and logically correct LaTeX code from textual descriptions, focused on scientific document components.

3 Dataset Construction

To assess LLMs’ capability to convert unstructured textual descriptions to LaTeX code, we build a benchmark dataset by following the process described in Figure 1. The process involves two major steps:

Collecting atomic LaTeX commands: We begin by systematically analyzing a range of data sources and scientific document templates to collect atomic LaTeX commands (details of sources

²<https://www.overleaf.com/latex/templates>




Figure 1: Process used to construct TeXpert, along with the dataset schema

Category	Atomic commands	Example
Text Formatting	86	<code>\textbf</code>
Equations and Symbols	83	<code>\arcsin</code>
Document Structure	75	<code>\subsubsection{ }</code>
Citation and References	39	<code>\bibliographystyle {style}</code>
Tables and Figures	36	<code>\cellcolor{color}</code>
Total	319	

Table 1: Details of the atomic LaTeX commands used to build TeXpert

and methodology are provided in Appendix A.1). These atomic commands, representing the minimal functional units commonly used in scientific writing and typically consisting of a backslash followed by a keyword and optional arguments, were extracted to form the basis of our dataset. The commands were then classified into 5 categories based on their purpose, as shown in Table 1. By adding an extra base step of collecting atomic commands commonly found in scientific formats, we regulate the scope of our final dataset containing LaTeX code generation tasks.

Generating TeXpert using atomic LaTeX commands: We curate a structured benchmark dataset containing natural language instructions for generating LaTeX code for various elements of scientific content using a combination of manual effort and LLM-based command generation. We build our dataset incrementally (while restricting the domain to atomic commands collected in the previous step to ensure specificity to scientific document components) using three different classes, namely Simple, Average and Hard, by increasing the complexity of tasks, the number of distinct atomic commands

and components of scientific documents needed, adding package requirements, and so on.

In order to classify the final task complexity as Simple, Average or Hard, we use specific constraints based on the number of commands, packages and components, precise description of which, along with a few examples, is found in Table 6 in Appendix A.2. With a focus on a small but high-quality dataset, we manually verify every row across all three classes in our dataset to ensure clear LaTeX generation requirements and consistency with the difficulty constraints. Our final dataset, named TeXpert, thus contains instructions and a classification label based on difficulty. After experimentation, we also add columns with a LaTeX code satisfying all requirements (if generated by any LLM) for future fine-tuning, along with the LLM that generated this correct code, resulting in the final schema in Figure 1. Statistics of our dataset are shown in Table 2.

4 Experimental Setup

We utilise a systematic evaluation framework to assess LLMs’ ability to generate syntactically correct LaTeX code from natural language prompts using the TeXpert dataset. We experiment with a wide range of open-source LLMs including Mistral Large 24.11 (AI, 2024b), Codestral (AI, 2024a), DeepSeek V3 (DeepSeek-AI, 2024), and DeepSeek Coder 33b (Guo et al., 2024) as well as multiple high-performance closed-source models including GPT-4o (OpenAI, 2024b), GPT-4o-mini (OpenAI, 2024a), Gemini 1.5 Flash (Team, 2024), Claude 3.5 Sonnet (Anthropic, 2024) and Grok 2-1212 (xAI, 2024).

For each sample across the three difficulty levels in TeXpert, we provide the LLM with a prompt containing task instructions for LaTeX code generation (provided in Figure 5 in Appendix B). During generation, model parameters were set to predetermined values to ensure deterministic outputs, as detailed in Table 11 in Appendix B. Detailed model configurations are provided in Section B.3 in Appendix B. Rule-based extraction techniques are used to extract the LaTeX code from the response.

We then evaluate each LLM’s response with GPT-4o as a judge, using a predefined evaluation prompt (refer to Figure 4 in Appendix B) to compute success rates and classify error types (described in Table 7 in Appendix B). The evaluation

prompt was iteratively refined through manual spot checks of evaluation outputs, focusing on clarity, correctness, and alignment with evaluation criteria. This process continued until the prompt consistently yielded reliable and interpretable results, as per our judgment. For the hard set, we also provide manually generated and verified LaTeX code as a reference during evaluation, to help identify all requirements of the task. To mitigate potential evaluation bias from using the same model family as the judge, we use DeepSeek v3 as an evaluator for GPT-4o and GPT-4o-mini.

5 Result Discussion

The accuracy of LaTeX generation for scientific documents across difficulty classes is presented in Table 3 and visualised in Figure 2. The overall distribution of error types across all difficulty levels is presented in Table 4 and Figure 3, while individual error distributions for Simple, Average, and Hard difficulty classes are also provided in Tables 8, 9 and 10 in Appendix B.2, respectively. From Table 3, we can infer that GPT-4o outshines all other LLMs in LaTeX code generation, closely followed by DeepSeek v3. DeepSeek Coder 33b provides the best performance on the most complex tasks.




Figure 2: Overall accuracy for LaTeX generation tasks by various LLMs

LaTeX generation tasks expose fundamental LLM shortcomings: Even models that perform highly on other benchmarks like GPT-4o and Mistral Large fail to achieve over 80% and 60% accuracy in simple and average sets, respectively. This reveals a critical capability gap in using LLMs for formatting scientific documents in LaTeX, most likely due to the scarcity of LaTeX examples in training datasets.

Hard LaTeX tasks reveal a universal limitation

Difficulty Class	No. of samples	Average length of textual instructions	Average no. of atomic LaTeX commands	Average no. of extra LaTeX packages
Simple	250	115.8 \pm 24 characters	10.9 \pm 7.2	0.5 \pm 0.8
Average	150	299.1 \pm 85.7 characters	51.2 \pm 29.2	3.6 \pm 2.4
Hard	40	558.4 \pm 216.7 characters	85.9 \pm 31.0	6.6 \pm 2.0

Table 2: Statistics of the TeXpert dataset, organised by difficulty class

Model	Accuracy %			
	Simple	Average	Hard	Overall
Closed-Source Models				
GPT-4o-mini	62.4	45.3	5	51.4
GPT-4o	78.8	58.7	15	66.1
Claude-3.5 Sonnet	62.8	56.7	0	55.0
Gemini 1.5 Flash	53.6	33.3	0	41.8
Grok 2 1212	62.4	52.0	5	53.6
Open-Source Models				
Mistral Large 24.11	64.4	59.33	10	57.7
Codestral 22B	60.8	41.3	0	48.6
DeepSeek V3	71.2	58.7	10	61.4
DeepSeek Coder 33b	69.2	58.0	17.5	60.7

Table 3: Main accuracy results (in %). Values in bold indicate the best accuracy for each difficulty class

Model	Error Types in %				
	CE	SE	LE	PE	FE
Closed-Source Models					
GPT-4o-mini	0.0	1.3	53.7	23.2	21.7
GPT-4o	0.0	2.1	59.1	15.2	23.6
Claude-3.5 Sonnet	0.0	5.3	44.3	29.9	20.6
Gemini 1.5 Flash	3.4	2.0	52.3	21.6	20.8
Grok-2 1212	1.2	5.3	46.5	25.5	21.4
Open-Source Models					
Mistral Large 24.11	0.0	2.5	53.0	20.8	23.7
Codestral 22B	0.6	2.8	52.5	18.8	25.3
DeepSeek V3	1.2	3.8	54.3	18.7	22.0
DeepSeek Coder 33b	0.4	2.6	54.0	20.5	22.5

Table 4: Overall error distribution for LaTeX generation tasks by various LLMs. CE = Capability Error, SE = Syntax Error, LE = Logical Error, PE = Package Error, FE = Formatting Error




Figure 3: Error distribution for LaTeX generation tasks by various LLMs

across models: Accuracy across the Simple and Average sets remains consistent across models, however, models show a dramatic performance cliff on hard tasks, with Claude and Gemini completely failing. This consistent degradation pattern clearly shows a threshold on the number and complexity of instructions for LaTeX generation using LLMs, which can be presumed to lie between instruction statistics for the Average and Hard sets in Table 2.

Open-source models strongly rival closed-source ones in LaTeX generation: Open-source models like DeepSeek V3 and DeepSeek Coder 33b perform well on par with frontier closed-source models like GPT-4o and Claude-3.5-Sonnet in overall accuracy with minimal capability errors as well. Notably, DeepSeek Coder 33b greatly outperforms Claude 3.5 Sonnet and Grok 2 in the Hard set. This demonstrates the potential of open-source models to provide powerful yet cost-effective alternatives.

6 Error Analysis

In this section, we provide a brief analysis of the most common error types and probable sources during LaTeX generation by LLMs. From our perspective, most powerful LLMs still struggle to provide error-free code due to basic oversights like missing packages and unfaithful instruction following. It is encouraging to see minimal capability errors and syntax errors. We leave an in-depth analysis of the root cause of errors to the future scope.

Logical errors dominate: Logical errors consistently account for the majority of issues across LLMs, highlighting struggles to fully satisfy task requirements. In all the cases we analysed, the most pronounced errors across all model variants were focused on missed instructions and wrong structural placement, especially in GPT-4o-mini and all open-source models. Similarly, error clustering in multiple equation and table generation tasks indicates that LLMs like DeepSeek v3 and Mistral Large struggle with maintaining long-range consistency. We believe these errors likely arise from weak structural understanding inherent in LLMs, limited exposure to LaTeX context, and misalignment between pretraining tasks and formal document generation.

Frequent formatting lapses: Notably, formatting errors occur far more frequently than we anticipated in all the LLMs we experimented with. Analysis of the evaluations reveals that these errors primarily involve incorrect environment selection and malformed tables or captions accompanying large tables or figures. Such issues indicate limited structural understanding and inadequate grounding in LaTeX syntax, even in larger models like DeepSeek v3 and GPT-4o, showing that scale alone is not the solution. We speculate that these errors stem from a scarcity of training data and examples specifically addressing table formatting and related constructs.

Package errors are concerning: Package errors are prominently caused by improper or incomplete inclusion and configuration of essential LaTeX packages, especially bibliography-related ones, most prominent in Claude 3.5 Sonnet. GPT-4o has the lowest share of missing packages, showing encouraging signs that more inclusive training data might mitigate this issue, although Codestral’s minimal package error rate also suggests potential for alternative approaches to reduce them further. Additionally, the use of non-standard or incompatible packages, especially in DeepSeek and Mistral models, is concerning and may point to LLMs hallucinating or making up packages to fill reasoning gaps. Overall, package issues suggest a fundamental gap in dependency management and environment consistency within LaTeX code generated by LLMs.

7 Conclusion

We curate TeXpert, a comprehensive benchmark designed to challenge LLMs to evaluate their La-

TeX code generation capability from natural language prompts. Our dataset consists of a total of 440 high-quality samples, organised by difficulty. Our findings reveal that LaTeX generation is still an underperforming skill in LLMs and that there is a need to include LaTeX package details and complex layouts in the training data for LLMs to improve their capability in this task. By making the code and dataset for TeXpert publicly available, we hope to support and encourage further research within the community.

Limitations

Our research marks a significant step forward in providing a benchmark for evaluating the LaTeX generation capabilities of LLMs. However, we acknowledge the limitations of our work as follows:

- **Limited dataset size:** The Hard set’s restricted size of 40 samples is a possible challenge in the generalisability of our findings. To address this, we encourage future work to increase the number and complexity of hard examples to broaden the benchmark’s effectiveness.
- **Fine-tuning models and improved prompts:** Using our dataset to fine-tune models and reduce logical and package errors in LaTeX-based tasks is another straightforward extension to our work, along with checking advanced prompting structures for performance improvements.
- **Additional LaTeX sources and applications:** While our work focuses on generating LaTeX code for only scientific documents, incorporating sources and tasks for other document types, such as resumes and books, would broaden the research scope.

References

- Asma Ben Abacha, Wen wai Yim, Yajuan Fu, Zhaoyi Sun, Meliha Yetisgen, Fei Xia, and Thomas Lin. 2025. [Medec: A benchmark for medical error detection and correction in clinical notes](#). *Preprint*, arXiv:2412.19260.
- Mistral AI. 2024a. [Codestral](#). Accessed: 2025-01-05.
- Mistral AI. 2024b. [Mistral large](#). Accessed: 2025-01-05.
- Anthropic. 2024. [Model card claude 3 addendum](#). Accessed: 2025-01-05.

- Joppe W. Bos and Kevin S. McCurley. 2023. [LaTeX, metadata, and publishing workflows](#). *Preprint*, arXiv:2301.08277.
- Savvas Chamezopoulos, Drahomira Herrmannova, Anita De Waard, Drahomira Herrmannova, Domenic Rosati, and Yury Kashnitsky. 2024. [Overview of the DagPap24 shared task on detecting automatically generated scientific paper](#). In *Proceedings of the Fourth Workshop on Scholarly Document Processing (SDP 2024)*, pages 7–11, Bangkok, Thailand. Association for Computational Linguistics.
- Xinyi Chen, Baohao Liao, Jirui Qi, Panagiotis Eustratiadis, Christof Monz, Arianna Bisazza, and Maarten de Rijke. 2024. [The SIFo benchmark: Investigating the sequential instruction following ability of large language models](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 1691–1706, Miami, Florida, USA. Association for Computational Linguistics.
- DeepSeek-AI. 2024. [Deepseek-v3 technical report](#). *Preprint*, arXiv:2412.19437.
- Yuntian Deng, Anssi Kanervisto, Jeffrey Ling, and Alexander M. Rush. 2017. [Image-to-markup generation with coarse-to-fine attention](#). *Preprint*, arXiv:1609.04938.
- Iker García-Ferrero, Rodrigo Agerri, Aitziber Atutxa Salazar, Elena Cabrio, Iker de la Iglesia, Alberto Lavelli, Bernardo Magnini, Benjamin Molinet, Johana Ramirez-Romero, German Rigau, Jose Maria Villa-Gonzalez, Serena Villata, and Andrea Zaninello. 2024. [MedMT5: An open-source multilingual text-to-text LLM for the medical domain](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 11165–11177, Torino, Italia. ELRA and ICCL.
- Daya Guo, Qihao Zhu, Dejian Yang, Zhenda Xie, Kai Dong, Wentao Zhang, Guanting Chen, Xiao Bi, Y. Wu, Y. K. Li, Fuli Luo, Yingfei Xiong, and Wenfeng Liang. 2024. [Deepseek-coder: When the large language model meets programming – the rise of code intelligence](#). *Preprint*, arXiv:2401.14196.
- Qianyu He, Jie Zeng, Qianxi He, Jiaqing Liang, and Yanghua Xiao. 2024. [From complex to simple: Enhancing multi-constraint complex instruction following ability of large language models](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 10864–10882, Miami, Florida, USA. Association for Computational Linguistics.
- Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. 2021. [Measuring mathematical problem solving with the math dataset](#). *Preprint*, arXiv:2103.03874.
- Kyudan Jung, Sieun Hyeon, Jeong Youn Kwon, Nam-Joon Kim, Hyun Gon Ryu, Hyuk-Jae Lee, and Jaeyoung Do. 2024. [Mathbridge: A large corpus dataset for translating spoken mathematical expressions into latex formulas for improved readability](#). *Preprint*, arXiv:2408.07081.
- LaTeX. An introduction to LaTeX. <https://www.latex-project.org/about/>. Accessed: 2025-01-04.
- OpenAI. 2024a. [Gpt-4o mini: Advancing cost-efficient intelligence](#). Accessed: 2025-01-05.
- OpenAI. 2024b. [Gpt-4o system card](#). Accessed: 2025-01-05.
- Yiwei Qin, Kaiqiang Song, Yebowen Hu, Wenlin Yao, Sangwoo Cho, Xiaoyang Wang, Xuansheng Wu, Fei Liu, Pengfei Liu, and Dong Yu. 2024. [InFoBench: Evaluating instruction following ability in large language models](#). In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 13025–13048, Bangkok, Thailand. Association for Computational Linguistics.
- Josselin Somerville Roberts, Tony Lee, Chi Heem Wong, Michihiro Yasunaga, Yifan Mai, and Percy Liang. 2024. [Image2struct: Benchmarking structure extraction for vision-language models](#). *Preprint*, arXiv:2410.22456.
- Betim Sherifi, Khaled Slhoub, and Fitzroy Nembhard. 2024. [The potential of llms in automating software testing: From generation to reporting](#). *Preprint*, arXiv:2501.00217.
- Xiangru Tang, Yiming Zong, Jason Phang, Yilun Zhao, Wangchunshu Zhou, Arman Cohan, and Mark Gestein. 2024. [Struc-bench: Are large language models really good at generating complex structured data?](#) *Preprint*, arXiv:2309.08963.
- Gemini Team. 2024. [Gemini 1.5: Unlocking multi-modal understanding across millions of tokens of context](#). *Preprint*, arXiv:2403.05530.
- xAI. 2024. [Grok 2](#). Accessed: 2025-01-05.
- Wenpeng Yin, Qinyuan Ye, Pengfei Liu, Xiang Ren, and Hinrich Schütze. 2023. [LLM-driven instruction following: Progresses and concerns](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing: Tutorial Abstracts*, pages 19–25, Singapore. Association for Computational Linguistics.
- Renrui Zhang, Dongzhi Jiang, Yichi Zhang, Haokun Lin, Ziyu Guo, Pengshuo Qiu, Aojun Zhou, Pan Lu, Kai-Wei Chang, Peng Gao, and Hongsheng Li. 2024. [Mathverse: Does your multi-modal llm truly see the diagrams in visual math problems?](#) *Preprint*, arXiv:2403.14624.
- Yiyun Zhao, Prateek Singh, Hanoz Bhatena, Bernardo Ramos, Aviral Joshi, Swaroop Gadiyaram, and Saket Sharma. 2024. [Optimizing LLM based retrieval augmented generation pipelines in the financial domain](#). In *Proceedings of the 2024 Conference of the North*

American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 6: Industry Track), pages 279–294, Mexico City, Mexico. Association for Computational Linguistics.

Jiaru Zou, Qing Wang, Pratyush Thakur, and Nickvash Kani. 2024. *Stem-pom: Evaluating language models math-symbol reasoning in document parsing*. Preprint, arXiv:2411.00387.

A Curation of TeXpert - Additional Details

A.1 Data Collection and Sources

To build the core of our TeXpert dataset, we manually extracted atomic commands from the Overleaf documentation listed in row 1 of Table 5 and from 25 documents each in LaTeX template repositories given in rows 2 and 3 of Table 5. This approach ensured a diverse range of document formats and LaTeX commands commonly used in scientific materials. For each document, a Python script using regular expressions was used to extract atomic LaTeX commands. These commands were then manually verified and grouped into five categories based on their function, as shown in Table 1. This process was intended to focus the dataset on commonly used LaTeX elements in scientific writing.

A.2 Difficulty constraints

Table 6 shows the constraints followed while classifying samples into difficulty classes (Simple/Average/Hard) during the generation of tasks in the TeXpert dataset. A randomly chosen example from each set is also provided for reference.

B Experimentation - Additional Details

B.1 Prompts

The prompts used during experimentation to evaluate responses using GPT-4o/DeepSeek v3 as a judge and to generate LaTeX code using natural language instructions and are given in Figures 4 and 4 respectively.

B.2 Error descriptions and distribution

Details of error types along with examples are given in Table 7. Additionally, the individual error distributions for Simple, Average, and Hard difficulty classes for each LLM are given in Tables 8, 9 and 10 respectively.

B.3 Model parameters

We report the generation parameters for all models used in our experiments to ensure transparency and reproducibility. All models were accessed through provider APIs, and the common parameter settings used across all models (except Anthropic models) are listed in Table 11. The model sizes of all closed-source models are approximate and taken from Abacha et al. (2025).

OpenAI Models: We run our experiments on two flagship models, GPT-4o (~200B parameters) and GPT-4o-mini (~8B parameters). We use the OpenAI Python SDK to access the models via API, specifying `seed=1234` and `n=1` along with the parameter values listed in Table 11, to ensure maximum determinism in responses. All other parameters are kept to default values.

DeepSeek Models: We use two recently released models, DeepSeek v3 (~671B parameters) and DeepSeek Coder (~33B parameters). DeepSeek models were accessed using the OpenAI Python SDK by specifying the DeepSeek URL endpoint and authentication details. Here too, we set `seed=1234` and `n=1` along with the parameter values listed in Table 11 during experimentation, keeping the rest to default values.

Mistral Models: We experiment with two powerful models, Mistral-Large-Instruct-2411 (~123B parameters) and Codestral-22B-v0.1 (~22B parameters). Both models were accessed using the official API in Mistral Python SDK, with an extra parameter `random_seed=1234` along with values in Table 11, with the rest as default.

Google AI Models: The Gemini 1.5 flash model was accessed using the official Google Generative AI Python SDK. Within the Generation Config, we set parameters values to those mentioned in Table 11, along with `candidate_count=1` and the rest as default.

xAI Models: We use a recently released Grok-2-1212 model by xAI, accessed using the OpenAI Python SDK by specifying the xAI endpoint. Here too, we set `seed=1234` and `n=1` along with the parameter values listed in Table 11 during experimentation, keeping the rest to default values.

Anthropic Models: The Claude 3.5 Sonnet model (~175B parameters) was accessed via the official Anthropic Python SDK. Due to limited configurable parameters, only `temperature=0.0`, `top_p=1`, and `max_tokens=8096` were explicitly set, with all other settings left at their defaults.

Data Source	URL
Overleaf Documentation	https://www.overleaf.com/learn
Overleaf Academic Journal Templates	https://www.overleaf.com/latex/templates/tagged/academic-journal
LaTeX Templates (Creodocs)	https://www.latextemplates.com/cat/academic-journals

Table 5: Primary sources used for collecting atomic LaTeX commands

Difficulty Class	Length of textual instructions	No. of atomic LaTeX Commands	No. of extra LaTeX packages	No. of specific formatting instructions (for tables, figures, etc.)	Example
Simple	<200 characters	10–20	<2	<2	Create a document with centered text in one block and justified text in another block.
Average	200–500 characters	12–80	2–5	2–5	Create a document with two sections. The first section should contain an aligned set of equations. The second section should contain a centered table, and the table should reference a figure placed in the first section. Your task is to produce a scientific research paper for arXiv that has a title page with author names, abstract and keywords, table of contents, and several sections. Add a 3x3 table that has lists in the second column, and figures with bold captions in last column. On every page except the first, add a footer with a signature image. Add an appendix that includes a table with header row entirely merged. Finally, add a custom bibliography.
Hard	500+ characters	80+	5+	5+	

Table 6: Description of constraints used during classification of tasks in TeXpert with a few examples

Error Type	Description	Examples
Capability Error	The LLM fails or denies to provide a valid response or says the task is out of its capability.	<ul style="list-style-type: none"> • LLM responds with: "Sorry, I cannot..." • LLM does not include any code in response
Syntax Error	The code generated by the LLM does not follow valid LaTeX syntax.	<ul style="list-style-type: none"> • Missing closing braces • Unescaped special characters
Logical Error	Mismatches between user instructions and the code logic, i.e., requirements given in natural language are not satisfied by the LaTeX code.	<ul style="list-style-type: none"> • Table headers omitted when explicitly requested • Missing components in code
Package Error	Required LaTeX packages are missing or commands do not match the document type.	<ul style="list-style-type: none"> • Using <code>\includegraphics</code> without importing the <code>graphicx</code> package
Formatting & Referencing Error	Layout issues like inconsistent alignment, font size, or spacing; improper formatting for cross-references, citations, or labels.	<ul style="list-style-type: none"> • Misaligned tables with inconsistent widths • Using <code>\ref{sec:1}</code> without defining <code>\label{sec:1}</code>

Table 7: Description and examples of error types used during evaluation of generated LaTeX code by LLMs

Prompt for evaluating LaTeX code

You are a judge tasked with evaluating the given LATEX CODE against the TASK INSTRUCTIONS.

Analyse the LaTeX code thoroughly to determine if it satisfies the instructions completely and without errors.

Check every line and command with utmost care for the following types of errors:

1. Capability Error: The LLM fails or denies to provide a valid response or says the task is out of its capability.

Examples:

- LLM responds with: "Sorry, I cannot..."
- LLM does not include any code in response at all.

2. Syntax Error: The code does not follow valid LaTeX syntax.

Examples:

- Missing closing braces.
- Unescaped special characters.

3. Logical Error: The code does not fulfill all requirements as given in the task instructions.

Examples:

- Table headers omitted when explicitly requested.
- Missing components or incorrect logic in the code.

4. Package Error: Necessary LaTeX packages for the code are missing, or commands used are incompatible with the document type.

Examples:

- Using `\includegraphics` without importing the `graphicx` package.
- Using `\chapter` in a document class that does not support chapters.

5. Formatting Errors: Errors in layout, alignment, spacing, or referencing. But, make sure to ignore errors where the LaTeX code is missing the `'references.bib'` file, it is not needed anywhere!

Examples:

- Misaligned tables with inconsistent column widths.
- Using custom references like `\ref{sec:1}` without defining `\label{sec:1}`.

YOUR TASK:

Evaluate the LATEX CODE for errors of the above types. If no errors are present, return 'error' as 'No' in below format. Provide your evaluation only and only in the following Python dictionary format:

```

'''{ "error": "Yes" or "No",
    "error_types": list of error types if error is "Yes", else empty list,
    "description": "description of error(s) encountered if error is "Yes", else empty string" }
'''

```

Figure 4: System prompt used to evaluate LaTeX code generated by LLMs using GPT-4o/DeepSeek v3 as-a-judge

Prompt for LaTeX code generation

You are a helpful LaTeX code assistant.

Your main job is to produce syntactically correct and logically accurate LaTeX code based on instructions given as TASK INSTRUCTIONS.

Make sure to follow each and every small instruction given in the task instructions.

The LaTeX code need not always be for an entire document if not specified, you can give code only according to the parts required.

Your output should be enclosed within ````latex` and ````` only.

Do not give any output other than the LaTeX code, please!

Figure 5: System prompt used to generate LaTeX code using LLMs for given textual instructions

Model	Error Types in % - Simple Set				
	CE	SE	LE	PE	FE
Closed-Source Models					
GPT-4o mini	0	0	57	37.4	5.6
GPT-4o	0	6.3	66.67	17.5	9.5
Claude-3.5 Sonnet	0	13.6	36.8	41.6	8
Gemini 1.5 Flash	5.7	2.9	55	27.9	8.6
Grok-2 1212	3.6	9.9	43.2	41.4	1.8
Open-Source Models					
Mistral Large 24.11	0	7.6	55.5	28.6	8.4
Codestral 22B	1.7	5	60.3	22.3	10.7
DeepSeek V3	3.5	5.9	52.9	30.6	7.1
DeepSeek Coder 33b	0	2.2	54.4	36.7	6.7

Table 8: Error distribution for LaTeX generation tasks from the Simple set by various LLMs. CE = Capability Error, SE = Syntax Error, LE = Logical Error, PE = Package Error, FE = Formatting Error

Model	Error Types in % - Average Set				
	CE	SE	LE	PE	FE
Closed-Source Models					
GPT-4o mini	0	1.6	55.9	11.8	30.7
GPT-4o	0	0	58.3	15.7	26
Claude-3.5 Sonnet	0	1	48.5	24.8	25.7
Gemini 1.5 Flash	4.6	0.7	53.6	16.3	24.8
Grok-2 1212	0	3.6	54.5	17.4	24.5
Open-Source Models					
Mistral Large 24.11	0	0	54.1	17.3	28.6
Codestral 22B	0	1.3	53.9	16.2	28.6
DeepSeek V3	0	2.2	60	11.1	26.7
DeepSeek Coder 33b	1.1	4.2	56.8	7.4	30.5

Table 9: Error distribution for LaTeX generation tasks from the Average set by various LLMs. CE = Capability Error, SE = Syntax Error, LE = Logical Error, PE = Package Error, FE = Formatting Error

Generation Parameters	
temperature	0.0
top_p	1.0
max_tokens	8096
frequency_penalty	0.0
presence_penalty	0.0

Table 11: Generation parameters used across all models

Model	Error Types in % - Hard Set				
	CE	SE	LE	PE	FE
Closed-Source Models					
GPT-4o mini	0	2.4	48.2	20.5	28.9
GPT-4o	0	0	52.3	12.3	35.4
Claude-3.5 Sonnet	0	1.2	47.6	23.2	28
Gemini 1.5 Flash	0	2.4	48.2	20.5	28.9
Grok-2 1212	0	2.5	41.8	17.7	38
Open-Source Models					
Mistral Large 24.11	0	0	49.3	16.4	34.2
Codestral 22B	0	2.2	43.3	17.8	36.7
DeepSeek V3	0	3.2	50	14.5	32.3
DeepSeek Coder 33b	0	1.4	50.7	17.4	30.4

Table 10: Error distribution for LaTeX generation tasks from the Hard set by various LLMs. CE = Capability Error, SE = Syntax Error, LE = Logical Error, PE = Package Error, FE = Formatting Error