



OPEN

# Achieving GPT-4o level performance in astronomy with a specialized 8B-parameter large language model

Tijmen de Haan<sup>1✉</sup>, Yuan-Sen Ting<sup>2,3</sup>, Tirthankar Ghosal<sup>4</sup>, Tuan Dung Nguyen<sup>5</sup>, Alberto Accomazzi<sup>6</sup>, Azton Wells<sup>7</sup>, Nesar Ramachandra<sup>7</sup>, Rui Pan<sup>8</sup> & Zechang Sun<sup>9</sup>

AstroSage-Llama-3.1-8B is a domain-specialized natural-language AI assistant tailored for research in astronomy, astrophysics, cosmology, and astronomical instrumentation. Trained on the complete collection of astronomy-related arXiv papers from 2007 to 2024 along with millions of synthetically-generated question-answer pairs and other astronomical literature, AstroSage-Llama-3.1-8B demonstrates remarkable proficiency on a wide range of questions. AstroSage-Llama-3.1-8B scores 80.9% on the AstroMLab-1 benchmark, greatly outperforming all models—proprietary and open-weight—in the 8-billion parameter class, and performing on par with GPT-4o. This achievement demonstrates the potential of domain specialization in AI, suggesting that focused training can yield capabilities exceeding those of much larger, general-purpose models. AstroSage-Llama-3.1-8B is freely available, enabling widespread access to advanced AI capabilities for astronomical education and research.

**Keywords** AI assistant, Large-language model, Continued pretraining, Supervised fine-tuning

Large-language model (LLM) assistants are rapidly gaining traction across all sectors of knowledge work worldwide. In astronomy, these models are used for providing factual information, as programming assistants, for brainstorming ideas, and for providing explanations tailored to the level of understanding or preferred style of the user. LLMs exhibit a remarkable robustness, often delivering useful outputs even when the input is malformed, lacks context, or contains inaccuracies.

Despite their potential, the development of specialized LLMs has been limited due to their recent emergence and the substantial resources required for training. Previous studies<sup>1–4</sup> have shown that models narrowly tailored to a specific domain can perform on par with, or even exceed, much larger general-purpose models. This suggests that a large, highly domain-specific model could achieve state-of-the-art performance.

In astronomy, however, high-performing specialized language models have not yet been achieved. While models like AstroLLaMA<sup>5,6</sup> have gained attention, they lack comprehensive benchmarking of their astronomical knowledge recall capabilities. Recent studies<sup>7</sup> have shown that many of these models, due to limited specialized training data and fine-tuning for instruction-following, suffer from either catastrophic forgetting or an inability to follow precise question-answering instructions, often performing worse than their baseline models (in this case, the Llama models).

Building on the previous efforts of COSMOSAGE<sup>8</sup> and AstroLLaMA, we have developed AstroSage-Llama-3.1-8B, a natural language assistant specialized in astronomy, astrophysics, cosmology, and astronomical instrumentation. For the remainder of this paper, we will refer to these subdomains collectively as “astronomy”.

<sup>1</sup>International Center for Quantum-field Measurement Systems for Studies of the Universe and Particles (QUP-WPI), High Energy Accelerator Research Organization (KEK), Tsukuba, Ibaraki, Japan. <sup>2</sup>Department of Astronomy, The Ohio State University, Columbus, OH, USA. <sup>3</sup>Center for Cosmology and AstroParticle Physics (CCAPP), The Ohio State University, Columbus, OH, USA. <sup>4</sup>National Center for Computational Sciences, Oak Ridge National Laboratory, Oak Ridge, TN, USA. <sup>5</sup>Department of Computer and Information Science, University of Pennsylvania, Philadelphia, PA, USA. <sup>6</sup>Center for Astrophysics, Harvard & Smithsonian, Cambridge, MA, USA. <sup>7</sup>Computational Science Division, Argonne National Laboratory, Lemont, IL, USA. <sup>8</sup>Department of Computer Science and Engineering, Hong Kong University of Science and Technology, Kowloon, Hong Kong. <sup>9</sup>Department of Astronomy, Tsinghua University, Beijing, People's Republic of China. ✉email: tijmen.dehaan@gmail.com

Through the use of a substantially more extensive and well-curated training dataset, we demonstrate for the first time that our specialized language model significantly outperforms baseline models in downstream tasks.

In the long term, we envision an agentic research assistant capable of autonomously conducting literature reviews, identifying relevant hypotheses, carrying out data analysis, and even formulating new research questions. The development of such scientific agents (LLMs capable of solving scientific problems end-to-end) is already a rapidly growing field in astronomy. Recent studies have shown promising results in automating research tasks, such as analyzing James Webb Space Telescope data through multi-agent collaboration and self-play reinforcement learning<sup>9</sup>. However, these studies have been largely constrained by the substantial costs associated with utilizing proprietary models at high volume.

Realizing this level of agency will require extensive experimentation and careful optimization. Given the substantial compute costs and data requirements inherent in large-scale model training, keeping the model size manageable while maintaining high performance is crucial. Our approach, demonstrated through astronomy knowledge recall, shows that specialized models can achieve state-of-the-art performance in specific domains. This not only makes the development of advanced research assistants more feasible but also ensures their accessibility to a wider range of institutions and researchers, potentially transforming the landscape of astronomical research and education.

## Continued pretraining

For AstroSage-Llama-3.1-8B, we selected Meta's Llama-3.1 8-billion parameter model<sup>10</sup> as our foundation model. This base model was chosen for both its strong general-purpose capabilities and its availability under the permissive Llama 3.1 Community License. Furthermore, among models in the 8-billion parameter class, it demonstrated superior performance in astronomical knowledge recall compared to both general-purpose models<sup>11</sup> and specialized astronomical LLMs<sup>7</sup>, making it an ideal baseline.

To begin the development process, we first focused on curating, obtaining, and cleaning a continued pretraining (CPT) dataset.

## Dataset preparation

The scaling laws of<sup>2</sup> show that model capability increases predictably with increased training data volume and computational resources. More recently, it has been found<sup>12–15</sup> that the power-law index of these scaling laws depends on data quality. Therefore, our general approach to assembling a corpus was to focus on maintaining a high quality threshold, maximizing data volume at that quality level.

We employed a multi-faceted approach to create a comprehensive, high-quality, high-variety CPT corpus. The primary components of our dataset include:

- Approximately 250,000 arXiv preprints from 2007–2024 with primary or cross-listed categories in astro-ph (astrophysics) or gr-qc (general relativity and quantum cosmology). For preprints with multiple versions, we selected the most recent. We deliberately excluded the Annual Review of Astronomy and Astrophysics papers used in<sup>11</sup>, AstroMLab-1 to generate the benchmark questions, ensuring our evaluation would test the model's ability to generalize knowledge rather than recall specific source materials.
- Relevant articles from a depth-2 search through Wikipedia's astronomy and astrophysics categories. We include articles that directly belong to these categories or their immediate subcategories, then further add linked articles. We perform keyword-based exclusion on a few tens of custom keywords such as “deity”, “cyberspace”, “mythological”, “sports” and so on, yielding a total of nearly 30,000 articles.
- A selection of ~800 textbooks that are available as PDFs or ebooks online.

For the vast majority of the dataset, the rendered PDF files were converted to Mathpix Markdown using Nougat OCR<sup>16</sup>. For the remaining sources, the data was either already in Markdown format, or was left as plain text.

## Pretraining run

The pretraining was conducted on the ORNL OLCF Frontier 1.6 exaflop supercomputer, leveraging its substantial computational resources. We used 184 nodes simultaneously, each of which is equipped with four AMD MI250X, which in turn have 2 Graphics Compute Dies (GCDs) each for a total of 8 GCDs per node, or a total of 1472 GCDs.

Further statistics and our choices of pretraining hyperparameters are summarized in Table 1.

The Llama-3.1 tokenizer, which uses a variant of tiktoken<sup>17</sup> with UTF-flexible encoding, was sufficient for our purposes, so we did not introduce any astronomy-specific tokens to the vocabulary. We considered incorporating arXiv identifiers (in the format arXiv:YYMM.numberV) for papers in the training set but ultimately decided against it due to the substantial increase in vocabulary size this would entail. Future work may explore expanding the vocabulary to include numerical representations relevant to astronomy, such as common quantities and units.

During training, we tracked the loss function and step sizes. Minimal hyperparameter tuning was performed, relying on the hyperparameters from<sup>8</sup> for all parameters other than the learning rate schedule. The warmup period of 40 samples was determined empirically from the loss curve of an earlier, canceled CPT run. The peak learning rate was extrapolated from smaller runs, but remained problematic, as Frontier requires high levels of parallelization with short wall times. This caused our initial runs to suffer from either insufficient learning due to a low peak learning rate, or catastrophic exploding gradients due to an excessively high peak learning rate. The final run described in this work used a tuned learning schedule with a peak learning rate as high as possible but still allowing convergence. In future efforts, we plan to further optimize the efficiency of the training procedure along the lines of the work presented in<sup>18</sup>, incorporating tensor, pipeline, and data parallelism through libraries

Hyperparameter	Value
Sequence length	8192 tokens
Micro batch size	3
Epochs	2
Learning rate	1.5e−4
Base model	Llama 3.1 8B
Optimizer	AdamW
Adam beta2	0.95
Adam epsilon	1e−5
Learning rate schedule	Constant with quadratic warmup
Max gradient norm	3.0
Weight decay	0.001
Warmup steps	40
Precision	BF16
FSDP	Full shard, auto wrap
Resource	Value
Plaintext filesize	19.9 GB
Token count	3.3 billion
Nodes	184
GCDs (effective # of GPUs)	1472
Training wall time	10 h
Total time spent	11.5 h
Effective GPU-hours spent	14,872 h
VRAM usage	96% of 64 GB/GCD

**Table 1.** Summary of pretraining hyperparameters and resource usage.

such as Megatron-Deepspeed<sup>19,20</sup>. We also aim to request a longer walltime with lower parallelization factor in order to be able to reduce loss and improve downstream performance.

### Dataset cleaning

We followed the cleaning procedures from<sup>8</sup>, including a perplexity-based cleaning approach. This method first splits the corpus into individual paragraphs and calculates their respective perplexity scores. Perplexity measures how well a language model can predict a given text sequence. Lower perplexity indicates text that follows expected patterns of natural language, while very high perplexity values often signal anomalous or corrupted content. Outliers with such high perplexity frequently stem from OCR errors, malformed text, or non-prose content such as tables.

Based on the distribution of perplexity scores (Fig. 1), as well as manual inspection of some examples, we established a threshold that excluded the top 2% of paragraphs with the highest perplexity scores. We then reconstructed each document using only the paragraphs below this threshold. This cleaning procedure removed approximately 2% of the total data volume. Figure 1 illustrates this process by showing the distribution of perplexity scores and our chosen threshold.


### Supervised fine-tuning

To improve the model's ability to follow instructions and answer questions effectively, we performed supervised fine-tuning (SFT). In this process, the model was trained to predict appropriate responses to given prompts, learning from a collection of high-quality example conversations. Below, we describe our approach to generating and curating the SFT dataset.

### SFT dataset

<sup>7</sup> identified a critical limitation in the AstroLLaMA series of specialized astronomical LLMs: their inability to outperform even their own starting base models, partly due to inadequate SFT. While these models showed marginal improvements in basic next-token prediction tasks, they performed worse than their baseline models on instructional question-and-answer (Q & A) tasks, even for straightforward astronomical knowledge recall. This shortcoming fundamentally undermines the purpose of specialized training. Therefore, in our study, we paid particular attention to the SFT process, generating training datasets orders of magnitude larger than previously available in astronomy.

The largest component of our SFT dataset consists of Q & A pairs. Using the method from<sup>8</sup>, we generated over 11 million synthetic Q & A pairs from papers in our CPT dataset. These Q & A pairs were then evaluated using an LLM based on four criteria:



**Fig. 1.** Histogram of paragraph-level perplexity used for dataset cleaning. The perplexity is calculated as  $\exp(-\langle \ln P \rangle)$ , where  $\langle \ln P \rangle$  is the average log probability per token in each paragraph. The red dashed line indicates the manually chosen threshold used to filter out high-perplexity paragraphs, which comprised approximately 2% of the total data volume. Paragraphs with perplexity above this threshold were removed from the dataset. Note that the highest observed perplexity values (around  $10^{12}$ ) extend beyond the right edge of the plot.

1. **Correctness:** The factual accuracy of the answer in relation to the question, ensuring that each Q & A pair adheres to current scientific understanding and accurately reflects the information presented in the source material.
2. **Stand-alone:** The ability of each Q & A pair to be understood in isolation, without needing additional context beyond the content provided. This ensures that an astronomy expert could answer the question based solely on the information within the pair.
3. **Pertinence:** The relevance and importance of the question to researchers or students in astronomy. We want the questions to be ones that a professional in the field might ask or find valuable for deeper understanding.
4. **Overall Quality:** a holistic assessment of the Q & A pair. This criterion is intended to capture aspects that may not be explicitly covered by correctness, stand-alone completeness, or pertinence but contribute to the overall usefulness of the Q & A pair.

The LLM was presented with several hand-written examples of ratings in the form Question–Answer–Score–Explanation, followed by the Question–Answer pair to be judged. The model was then asked to complete the score values. A small number of resulting scores were verified and confirmed to be sufficiently accurate. Only Q & A pairs with perfect scores in all four categories were kept, resulting in 8.8 million high-quality Q & A pairs. Below are a few examples.

- **Q:** What are the three distinct phases of galaxy formation identified in the EAGLE simulation, and how are they primarily determined? **A:** The three distinct phases of galaxy formation identified are:
  1. Disorganized growth in low-mass galaxies (stellar mass less than approximately  $10^{9.5} M_{\odot}$ ), characterized by random stellar motions and dominated by in-situ star formation and mergers;
  2. Evolution towards a disc-dominated morphology in intermediate-mass galaxies (stellar mass between  $10^{9.5} M_{\odot}$  and  $10^{10.5} M_{\odot}$ ), with bulge formation associated with mergers;
  3. High-mass galaxies (stellar mass greater than approximately  $10^{10.5} M_{\odot}$ ) experience a slowdown in in-situ star formation, transitioning to spheroidal morphologies driven by the accretion of ex-situ stars. These phases are primarily determined by stellar mass rather than redshift.
- **Q:** What are the expected contributions of AGN feedback to the heating and dynamics of the hot ISM? **A:** AGN feedback is anticipated to play a significant role in regulating the cooling of the hot ISM, with energy released during accretion often exceeding starburst-driven heating and driving turbulence and outflows that influence the dynamics of the ISM.
- **Q:** What are the applications of the Augmented Lagrangian Perturbation Theory (ALPT) in cosmology? **A:** ALPT can be used for setting initial conditions for N-body simulations, generating mock galaxy catalogs, analyzing cosmic structures, and reconstructing primordial density fluctuations for baryon acoustic oscillation studies.
- **Q:** What are the three effects that anisotropies in the optical depth produce in the CMB? **A:** Anisotropies in the optical depth produce three effects in the CMB: (i) screening of the temperature and polarization fluctuations observed today by an overall factor of  $\exp(-\tau(\hat{n}))$ ; (ii) Thomson scattering, producing new E-mode polarization from the local temperature quadrupole, with subdominant contributions to B-modes

from reionization effects; and (iii) new temperature anisotropy generated from the radial motion of ionized bubbles relative to the observer.

We also included a filtered version of Infinity-Instruct-7M<sup>21</sup>, keeping only entries with at least 70% alphanumeric characters, as well as filtering out entries with certain keywords. The inclusion of this dataset was to ensure that AstroSage-Llama-3.1-8B would gain instruction-following abilities such as multi-turn conversation.

Additionally, we generated synthetic summaries for all of the papers in the CPT dataset. The SFT dataset for these summaries consists of a user prompt asking to summarize a certain paper, with the assistant completion being a small preamble followed by the summary. The user prompt was created through a series of random choices about the way the question is asked, a small preamble, and the way in which the paper is referenced, yielding a high variety of question styles.

Furthermore, we generated a metadata-based dataset, where again a series of custom rules and random selections result in diverse questions about titles, dates of publication, arXiv IDs, and first author names from the papers in the CPT dataset. This was included in an effort to memorize the paper metadata so that users can reference papers in their conversations with AstroSage-Llama-3.1-8B. However, we find that after training, the model is unable to reliably answer metadata-based questions such as producing a first author name from an arXiv ID. This failure to memorize could be ameliorated in future training runs by increasing the learning rate or number of epochs.

These datasets were combined with five further datasets which were assembled by hand from various sources on the web. The combined dataset comprised approximately 2 billion tokens.

### SFT procedure

The SFT process was also conducted on the Frontier supercomputer with a configuration summarized in Table 2.

Throughout the fine-tuning process, we again monitored loss and step sizes, which are shown in Fig. 2 (CPT curves look similar). The learning was—like in “Pretraining run” section—limited by the maximum wall-time allowed by the Frontier HPC system, which for 184 nodes or more is a maximum of 12 hours. With the learning rate as high as it could comfortably be set to avoid exploding gradients, this limitation on walltime strongly limited the number of steps that could be taken and therefore the final loss that could be obtained.

### Model merging

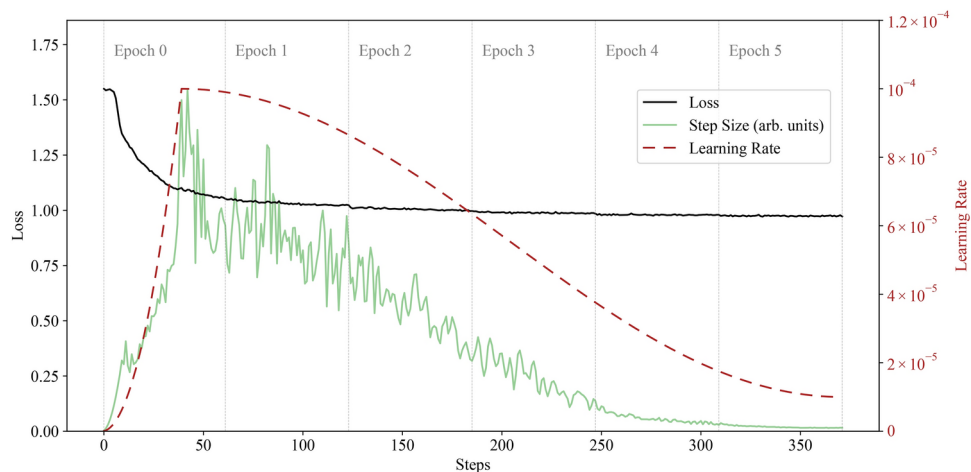
Model merging, also known as parameter averaging, has emerged as a powerful technique for combining capabilities of multiple expert models into a single language model<sup>22,23</sup>. While our CPT+SFT procedure significantly improved the model’s astronomical knowledge recall in few-shot prompts, we observed that performance in conversational Q & A scenarios such as multi-turn conversations and instructions regarding the output style still fell slightly short of optimal. This challenge likely stems from the fact that the “instruct” version of Llama-3.1-8B provided by Meta underwent substantially more extensive supervised fine-tuning than what we could achieve as an academic group. We found that merging our specialized model with Meta’s instruct model significantly improved these conversational capabilities.

To create the final version of AstroSage-Llama-3.1-8B, we employed MERGEKIT<sup>24</sup>, using the DARE-TIES method to combine our SFT-trained model described in “Supervised fine-tuning” section with Meta-Llama-3.1-8B-Instruct<sup>10</sup>. The merge was performed at full density, BF16 precision, and with the weight parameters set to 0.75 and 0.25 for AstroSage-Llama-3.1-8B-SFT and Meta-Llama-3.1-8B-Instruct, respectively.

The resulting merged model exhibits enhanced instruction-following capabilities and improved performance on the AstroMLab-1 multiple-choice question benchmark in both few-shot and structured output scenarios. To determine whether these improvements stemmed from enhanced instruction-following rather than additional astronomical knowledge, we conducted a control experiment. We fine-tuned a separate version of the CPT+SFT model on unrelated multiple-choice questions using identical output formatting. This control model achieved near-identical scores on the AstroMLab-1 benchmark in the structured output scenario without any merging,

Hyperparameter	Value
Epochs	6
Learning rate	1e−4
Base model	CPT model
Learning rate schedule	Cosine with quadratic warmup
Weight decay	0.0
Resource	Value
Plaintext filesize	9.8 GB
Token count	2.0 billion
Training wall time	9.5 h
Total time spent	11.5 h
Effective GPU-hours spent	13,738 h

**Table 2.** Summary of supervised fine-tuning hyperparameters and resource usage. Parameters that are not stated here were kept the same as in Table 1.



**Fig. 2.** Supervised fine-tuning loss curve. The learning rate schedule, shown as the dashed brick red curve, follows a quadratic warmup followed by a cosine schedule that ends at 10% of the peak learning rate. The peak learning rate was chosen to prevent exploding gradients. The loss curve in black shows no significant signs of overfitting, as evidenced by minimal discontinuities at epoch boundaries. The green curve represents the L2 norm of the parameter update, shown in arbitrary units.

suggesting that the process of merging in a small fraction of the Meta-Llama-3.1-8B-Instruct model weights transferred general question-answering capabilities rather than domain-specific knowledge.

To recap, we began with Meta-Llama-3.1-8B as our base model, then performed CPT on a large corpus of astronomy literature to instill domain knowledge. This was followed by SFT using carefully curated instruction-response pairs to improve task performance and instruction following. Finally, we merged the resulting model with Meta-Llama-3.1-8B-Instruct to enhance general instruction-following capabilities while preserving the astronomical expertise, resulting in our final model which we are releasing as AstroSage-Llama-3.1-8B.

## Evaluation

We have performed three types of quantitative evaluation. The first is an evaluation of AstroSage-Llama-3.1-8B on a multiple-choice question benchmark, the second on several general-purpose benchmarks, and the third direct preference testing against a similar model that was not specialized in astronomy.

### Multiple-choice question benchmark

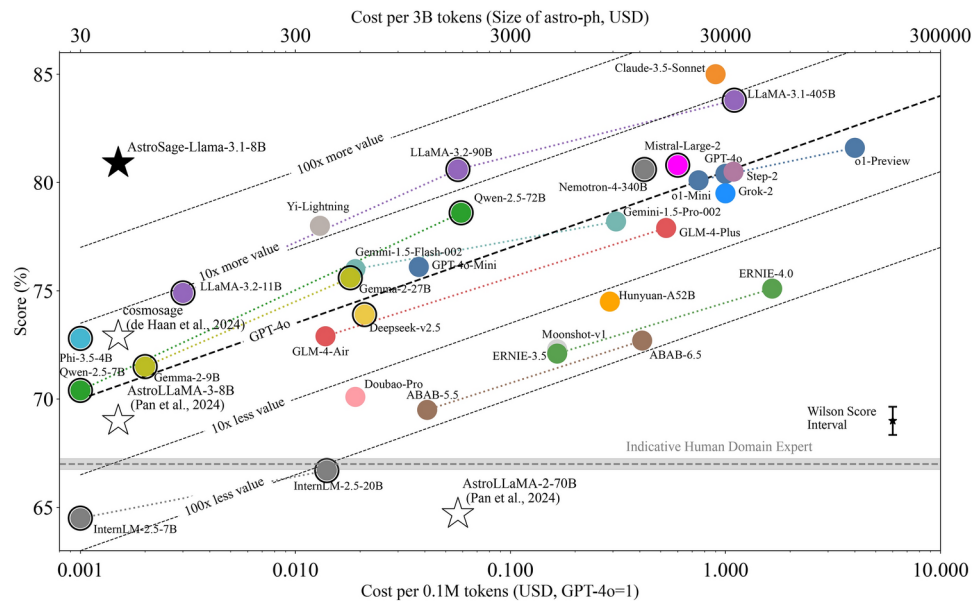
To evaluate AstroSage-Llama-3.1-8B's performance, we employed the multiple-choice question benchmark from the first paper in this series<sup>11</sup>, AstroMLab 1. This benchmark consists of diverse astronomy-related questions generated from selected Annual Review of Astronomy and Astrophysics (ARAA) papers and remains, to our knowledge, the only comprehensive astronomy-specific benchmarking effort available. We refer interested readers to the original paper for detailed benchmark specifications.

Importantly, we deliberately excluded the ARAA papers from AstroSage-Llama-3.1-8B's training dataset. This strategic exclusion enables us to evaluate the model's broader understanding of astronomical concepts rather than its ability to recall specific information from the source materials. This approach helps ensure that the benchmark scores reflect AstroSage-Llama-3.1-8B's genuine comprehension of astronomy rather than mere memorization of the content used to create the questions.

Our choice to primarily evaluate AstroSage-Llama-3.1-8B with a knowledge-based benchmark was motivated by two key factors. First, this benchmark represents the only extensively tested and human-vetted dataset available for astronomical knowledge assessment. Second, while astronomical knowledge recall represents just one aspect of LLM capabilities, it serves as a critical foundation for more advanced applications such as scientific agents. The primary goal is to demonstrate that proper fine-tuning of a relatively small model can significantly improve performance on a specific task—an achievement not previously demonstrated in astronomy.

The performance score is calculated as the fraction of correctly answered multiple-choice questions in the benchmarking dataset. The resulting scores are shown in Fig. 3, where round symbols represent scores for cutting-edge proprietary and open-weight models. The open-weight models are also marked with an outer circle. The x-axis displays the cost per  $10^5$  tokens, a metric chosen based on practical applications: in the first (and to our knowledge, only) implementation of astronomical agents<sup>9</sup>, analyzing a celestial source's spectral energy distribution from James Webb Space Telescope data requires approximately  $10^5$  tokens. The top x-axis shows costs scaled to 3B ( $3 \times 10^9$ ) tokens, roughly equivalent to the entire astro-ph section of the arXiv. For proprietary models, we use current token costs (averaging input and output costs where they differ), while open-weight model costs are estimated based on typical pricing of commercial GPU platforms.

Specialized astronomical LLMs are denoted by star symbols, except for the first AstroLLaMA model<sup>5</sup>, whose score falls below the plot's lower limit. The bottom right panel shows the typical uncertainty (calculated using the Wilson score interval), demonstrating that our dataset of 4425 multiple-choice questions provides sufficiently



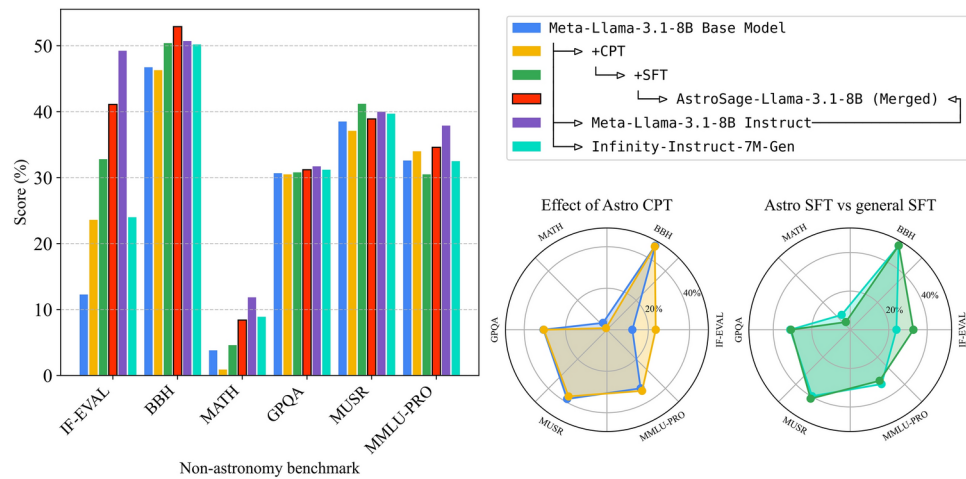
**Fig. 3.** Performance comparison on the AstroMLab 1 benchmark, which contains 4425 high-quality, human-verified multiple-choice questions across astronomy, astrophysics, cosmology, and instrumentation. The gray shaded bar spans the performance of the two human domain experts. We present updated results as of November 2024, incorporating both cutting-edge proprietary and open-weight models. Open-weight models are circled. AstroSage-Llama-3.1-8B outperforms all other models in the 8-billion parameter class and achieves performance comparable to OpenAI’s latest models, including GPT-4o, while Claude-3.5-Sonnet maintains the highest performance overall. The diagonal dashed lines represent cost-efficiency trade-offs as determined in AstroMLab 1 (see text for details). The Wilson Score interval shows the typical uncertainty in the score due to the finite number of questions. Star symbols indicate all published specialized LLMs for astronomy to our knowledge. Previously, these specialized models often failed to outperform their baseline models in astronomical recall due to various training limitations. AstroSage-Llama-3.1-8B represents a significant advancement in specialized astronomical LLMs, demonstrating that extensive data curation, massive continued pre-training and supervised fine-tuning, and model merging techniques can substantially improve performance on specific astronomical tasks. This result highlights the effectiveness of domain specialization even in relatively smaller models.

small sampling noise to establish robust performance differences. We have updated all scores using the latest model versions following the methodology from<sup>11</sup>, AstroMLab 1.

The diagonal dashed lines represent a universal cost-efficiency trade-off observed across major model series (e.g. Llama, GPT, GLM) that simultaneously released models at multiple sizes. We consistently observe a 3.5-point improvement in performance for every 10-fold increase in cost across model families. Each dashed line represents this equivalent trade-off, offset by 3.5 percentage points (equivalent to a 10-fold gain in cost-effectiveness). Despite similar performance on general benchmarks, cutting-edge models can differ by up to 1000-fold in cost-effectiveness on astronomical tasks, highlighting the importance of specialized astronomical benchmarks for evaluating performance on niche technical domains.

To establish a human performance baseline, two domain experts from our team independently completed a random subset of benchmark questions under controlled conditions. The experts answered on the order of one hundred questions each, taking around 30 seconds per question. No external references, web searches, or language model assistance were used. Both experts achieved remarkably consistent scores of approximately 68%, which we designate as the “Indicative Human Domain Expert” score. The fact that most evaluated LLMs significantly surpassed this baseline demonstrates both the benchmark’s comprehensive scope and difficulty, while highlighting the remarkable capabilities of current LLMs in capturing and applying complex astronomical knowledge.

As previously noted in<sup>7</sup>, AstroMLab 2, existing specialized astronomical LLMs (shown as open stars in Fig. 3) fail to outperform baseline models of comparable parameter size. In many cases, suboptimal specialization techniques actually led to performance degradation. In contrast, AstroSage-Llama-3.1-8B, despite its modest size of 8 billion parameters, achieved an accuracy of 80.9% on this benchmark-comparable to OpenAI’s latest flagship models (GPT-4o: 80.4%) and the best 90B-parameter open-weight Meta-Llama models (80.6%). This performance is particularly notable because AstroSage-Llama-3.1-8B achieves these results at approximately one-thousandth the inference cost of proprietary models and one-hundredth the cost of open-weight models. Furthermore, it demonstrates an 8-point improvement over its baseline model, Meta-Llama-3.1-8B (72.9%). To our knowledge, this represents the first demonstration of a specialized astronomical LLM achieving objectively verified improvements through model fine-tuning.



**Fig. 4.** Performance comparison across general language model benchmarks. The bar chart shows a comparison of model performances on standard benchmarks. The left radar chart visualizes the effect of continued pretraining. The right radar chart compares our model post-SFT against BAAI/Infinity-Instruct-7M-Gen-LLaMA3\_1-8B, which was trained on the same base model and SFT data but without our astronomy-specific data. Despite optimization for astronomy tasks, the merged model we are releasing as AstroSage-Llama-3.1-8B maintains strong general capabilities in reasoning, mathematics, and coding, demonstrating that domain specialization did not come at the cost of other abilities.

### General-purpose benchmarks

To ensure our domain specialization did not compromise general capabilities, we evaluated AstroSage-Llama-3.1-8B across a comprehensive suite of standard language model benchmarks. These include IF-EVAL<sup>25</sup> (instruction following), BBH<sup>26</sup> (binary hypothesis testing), MATH<sup>27</sup> (mathematical reasoning), GPQA<sup>28</sup> (graduate-level science questions), MUSR<sup>29</sup> (real-world decision-making scenarios), and MMLU-PRO<sup>30</sup> (an expanded version of MMLU with more challenging reasoning questions). As shown in Fig. 4, our CPT+SFT model (green, initialized from the Llama-3.1 base model) initially performed below the Llama-3.1 instruct model (purple) on five out of the six non-astronomy benchmarks. This was expected, given that Meta's proprietary SFT dataset for their instruct model likely far exceeds what is feasible for an academic research group to reproduce. The merging procedure, pulling in only 25% of its weight from Meta-Llama-3.1-8B-Instruct, allowed us to recover much of this performance deficit.

Crucially, this performance recovery through model merging did not compromise AstroSage-Llama-3.1-8B's astronomical expertise—it maintained its 8-point improvement (representing more than 100-fold increase in cost-effectiveness) on astronomical Q & A tasks while largely preserving capabilities across most general benchmarks. The only notable performance decrease occurred in IF-EVAL, which tests instruction following. This limited decline is unsurprising, as instruction following remains one of the more brittle capabilities in language models and likely heavily depends on the proprietary training data used in Meta's instruct model. In fact, when comparing AstroSage-Llama-3.1-8B to BAAI/Infinity-Instruct-7M-Gen-LLaMA3\_1-8B, the latter shows an even more severe performance deficit, highlighting how our refined training strategy and expanded SFT dataset represent crucial improvements. Ultimately, our model merging approach successfully preserved most general capabilities without sacrificing the gained astronomical expertise. This balance is essential, as it enables AstroSage-Llama-3.1-8B to engage in natural conversations and assist with broader tasks while excelling in astronomy-specific applications.

### Human blind rankings

Simultaneously with<sup>8</sup>, we performed a blinded preference ranking of AstroSage-Llama-3.1-8B against Meta-Llama-3.1-8B-Instruct. Fifteen questions were written about diverse areas of cosmology, spanning layperson to professional-level difficulty. Independent evaluators compared responses from AstroSage-Llama-3.1-8B and Meta-Llama-3.1-8B-Instruct, with both models receiving identical system prompts tailored to the cosmology context.

Responses were presented in randomized order to the evaluators, who rated their quality without knowledge of the source. The evaluation involved three evaluators, who preferred the AstroSage-Llama-3.1-8B answers in 73% of cases, reflecting a statistically significant preference for AstroSage-Llama-3.1-8B.

### Availability

To promote reproducibility and advance the field of domain-specific AI assistants, we are making AstroSage-Llama-3.1-8B freely available under the highly permissive Llama 3.1 Community License. The full model weights can be accessed and downloaded from our project repository on Hugging Face: <https://huggingface.co/AstroMLab/AstroSage-8B> in either PyTorch or safetensors format.

The code used to prepare the datasets and perform the training will be made available upon reasonable request. The synthetically generated Q & A pairs are currently supporting research and development of our next-generation model. After the completion of our research trajectory, we anticipate a full release of our synthetic data set. For inquiries regarding data usage, collaboration, or access to specific subsets of our work, interested parties are encouraged to contact the corresponding author

By making AstroSage-Llama-3.1-8B widely available, we aim to foster collaboration and innovation in the astronomy community. We encourage researchers to build upon our work and contribute to the ongoing development of specialized AI assistants for scientific domains.

## Discussion and future work

This work demonstrates the potential of specialized language models in astronomy through a systematic approach to model development and evaluation. While previous efforts like<sup>7</sup> laid important groundwork in domain-specific modeling, the field has faced persistent challenges in achieving performance gains over baseline models, especially in instruction-following tasks. Our multi-stage training process—combining continued pretraining, extensive supervised fine-tuning, and strategic model merging—addresses these challenges, achieving a notable improvement over the baseline model.

These results demonstrate that powerful AI assistants can be developed with relatively small language models when sufficiently specialized. Despite its modest size of 8 billion parameters, AstroSage-Llama-3.1-8B achieves performance comparable to latest flagship models at a fraction of the cost—approximately one-thousandth of proprietary models and one-hundredth of open-weight models. This remarkably favorable performance-to-parameter ratio suggests even greater potential for improvement through scaling. Given access to the necessary computational resources, we plan to apply our successful CPT/SFT procedure to a 70B-class model to pursue state-of-the-art astronomy-specific performance.

Beyond the performance achievements, our work establishes a more systematic approach to model evaluation in astronomy. Through tailored astronomy-specific benchmarking in, we provide a more rigorous and transparent assessment than previously available. However, significant challenges remain in comprehensive model evaluation. The field currently lacks standardized, astronomy-specific benchmarks capable of assessing understanding across the full spectrum of astronomical tasks, particularly in exact problem-solving capabilities like those tested in ScienceAgentBench<sup>31</sup>. This limitation restricts our ability to validate comparisons in more direct scientific agent contexts.

The constraints of an 8B-parameter model also become apparent in certain scenarios. While AstroSage-Llama-3.1-8B demonstrates impressive performance in subjective testing, the AstroMLab-1 benchmark, and general benchmarks, it encounters natural limitations in memory capacity and reasoning depth. Particularly challenging are questions requiring complex multi-step reasoning or sophisticated calculations, where larger general-purpose models still maintain an advantage.

To address these limitations, our future work will pursue several complementary directions. While scaling up model size remains a primary goal, we will also focus on developing more specialized benchmarking tools and exploring retrieval-augmented generation for improved knowledge access. Additional initiatives include creating multilingual astronomy assistants, implementing mechanisms for real-time knowledge updates, and providing public inference capabilities.

The broader implications of this work extend well beyond its immediate achievements. AstroSage-Llama-3.1-8B serves as a compelling proof of concept for highly specialized, smaller-scale language models in astronomy. Our approach of extensive data curation, continued pretraining, and careful supervised fine-tuning demonstrates how domain-specific expertise can be enhanced while preserving general capabilities. As the field progresses toward agentic research assistants capable of autonomous literature review, data analysis, and hypothesis generation, the need for affordable, highly competent domain-specific models will only grow. While challenges remain, AstroSage-Llama-3.1-8B charts a promising course for developing the next generation of specialized scientific AI assistants, potentially transforming how we approach astronomical research and education.

## Data availability

The datasets generated and/or analyzed during the current study are not publicly available due to their ongoing use in research and development of our next-generation model, but are available from the corresponding author on reasonable request.

Received: 3 December 2024; Accepted: 2 April 2025

Published online: 21 April 2025

## References

1. Fu, X.-Y. Laskar, M. T. R. Khasanova, E., Chen, C. & TN, S. B. Tiny Titans: Can Smaller Large Language Models Punch Above Their Weight in the Real World for Meeting Summarization? [arXiv:2402.00841](https://arxiv.org/abs/2402.00841) (2024).
2. Hoffmann, J. et al. Training Compute-Optimal Large Language Models. [arXiv:2203.15556](https://arxiv.org/abs/2203.15556) (2022).
3. Schick, T. & Schütze, H. It's Not Just Size That Matters: Small Language Models Are Also Few-Shot Learners. [arXiv:2009.07118](https://arxiv.org/abs/2009.07118) (2021).
4. Turc, I. Chang, M.-W. Lee, K. & Toutanova, K. Well-Read Students Learn Better: On the Importance of Pre-training Compact Models. [arXiv:1908.08962](https://arxiv.org/abs/1908.08962) (2019).
5. Nguyen, T. D. et al. AstroLLaMA: Towards Specialized Foundation Models in Astronomy. [arXiv:2309.06126v1](https://arxiv.org/abs/2309.06126v1) (2023).
6. Perkowski, E. et al. Research Notes of the AAS 8, 7 (The American Astronomical Society, 2024).
7. Pan, R. et al. In *SC24-W: Workshops of the International Conference for High Performance Computing, Networking, Storage and Analysis* 87–96 (2024).
8. de Haan, T. *Astron. Comput.* **51**, 100934 (2025).

9. Sun, Z., Ting, Y.-S., Liang, Y., Duan, N., Huang, S. & Cai, Z. Interpreting Multi-band Galaxy Observations with Large Language Model-Based Agent. [arXiv:2409.14807](https://arxiv.org/abs/2409.14807) (2024).
10. Dubey, A. et al. The Llama 3 Herd of Models. [arXiv:2407.21783](https://arxiv.org/abs/2407.21783) (2024).
11. Ting, Y.-S. et al. AstroM-Lab 1: Who Wins Astronomy Jeopardy!? [arXiv:2407.11194v1](https://arxiv.org/abs/2407.11194v1) (2024).
12. Rae, J. W. et al. Scaling Language Models: Methods, Analysis & Insights from Training Gopher. [arXiv:2112.11446v2](https://arxiv.org/abs/2112.11446v2) (2021).
13. Brown, T. et al. In *Advances in Neural Information Processing Systems*, Vol. 33, 1877–1901 (Curran Associates, Inc., 2020).
14. Penedo, G. et al. The RefinedWeb Dataset for Falcon LLM: Outperforming Curated Corpora with Web Data, and Web Data Only. [arXiv:2306.01116v1](https://arxiv.org/abs/2306.01116v1) (2023).
15. Li, R., Wei, Y., Zhang, M., Yu, N., Hu, H. & Peng, H. ScalingFilter: Assessing Data Quality through Inverse Utilization of Scaling Laws. [arXiv:2408.08310v1](https://arxiv.org/abs/2408.08310v1) (2024).
16. Blecher, L., Cucurull, G., Scialom, T. & Stojnic, R. Nougat: Neural Optical Understanding for Academic Documents. [arXiv:2308.13418](https://arxiv.org/abs/2308.13418) (2023).
17. tiktoken. <https://github.com/openai/tiktoken> (2022).
18. Dash, S. et al. Optimizing Distributed Training on Frontier for Large Language Models. [arXiv:2312.12705v2](https://arxiv.org/abs/2312.12705v2) (2023).
19. Megatron-DeepSpeed. <https://github.com/deepspeedai/Megatron-DeepSpeed> (2021).
20. Smith, S. et al. Using DeepSpeed and Megatron to Train Megatron-Turing NLG 530B, A Large-Scale Generative Language Model. [arXiv:2201.11990](https://arxiv.org/abs/2201.11990) (2022).
21. Beijing Academy of Artificial Intelligence (BAAI), “BAAI/Infinity-Instruct. Datasets at Hugging Face,” <https://huggingface.co/datasets/BAAI/Infinity-Instruct> (2024).
22. Yadav, P. et al. What Matters for Model Merging at Scale? [arXiv:2410.03617](https://arxiv.org/abs/2410.03617) (2024).
23. Dassanaik-Perera, A., Waiwitlikhit, S. & Gilbai, K. Cuts and Stitches: Does Model Merging Produce Better Multitask Learners?, Stanford CS224N Default Project (Stanford, 2023).
24. Goddard, C. et al. Arcee’s MergeKit: A Toolkit for Merging Large Language Models. [arXiv:2403.13257](https://arxiv.org/abs/2403.13257) (2024).
25. Zhou, J. et al. Instruction-Following Evaluation for Large Language Models. [arXiv:2311.07911](https://arxiv.org/abs/2311.07911) (2023).
26. Suzgun, M. et al. Challenging BIG-Bench Tasks and Whether Chain-of-Thought Can Solve Them. [arXiv:2210.09261](https://arxiv.org/abs/2210.09261) (2022).
27. Hendrycks, D. et al. Measuring Mathematical Problem Solving With the MATH Dataset. [arXiv:2103.03874](https://arxiv.org/abs/2103.03874) (2021).
28. Rein, D. et al. GPQA: A Graduate-Level Google-Proof Q & A Benchmark. [arXiv:2311.12022](https://arxiv.org/abs/2311.12022) (2023).
29. Sprague, Z., Ye, X., Bostrom, K., Chaudhuri, S. & Durrett, G. MuSR: Testing the Limits of Chain-of-thought with Multistep Soft Reasoning. [arXiv:2310.16049v2](https://arxiv.org/abs/2310.16049v2) (2023).
30. Wang, Y. et al. MMLU-Pro: A More Robust and Challenging Multi-Task Language Understanding Benchmark. [arXiv:2406.01574](https://arxiv.org/abs/2406.01574) (2024).
31. Chen, Z. et al. [arXiv:2410.05080](https://arxiv.org/abs/2410.05080) (2024).

## Acknowledgements

This research used resources of the Oak Ridge Leadership Computing Facility (OLCF), which is a DOE Office of Science User Facility at the Oak Ridge National Laboratory supported by the U.S. Department of Energy under Contract No. DE-AC05-00OR22725 and support from Microsoft’s Accelerating Foundation Models Research (AFMR) program. TdH was supported by World Premier International Research Center Initiative (WPI), MEXT, Japan. YST is supported by the National Science Foundation under Grant No. 2406729. Work at Argonne National Lab is supported by UChicago Argonne LLC, Operator of Argonne National Laboratory. Argonne, a U.S. Department of Energy Office of Science Laboratory, is operated under contract no. DE-AC02-06CH11357. A special thanks goes out to Cassie Reuter and Joshua Montgomery for acting as independent evaluators.

## Author contributions

TdH performed data preparation, cleaning, training, astronomy-specific evaluation, wrote the majority of the manuscript text, and created the figures. YST performed astronomy-specific evaluation, provided computational resources, provided the first draft of the astronomy benchmark figure, and made significant contributions to the manuscript text. AW performed general-purpose evaluation. All authors read and agree with the contents of the manuscript.

## Declarations

### Competing interests

The authors declare no competing interests.

### Additional information

**Correspondence** and requests for materials should be addressed to T.H.

**Reprints and permissions information** is available at [www.nature.com/reprints](http://www.nature.com/reprints).

**Publisher’s note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2025