AstroLLaMA adapts a general language model to astronomy abstracts.

arXiv:2309.06126v1 [astro-ph.IM] 12 Sep 2023

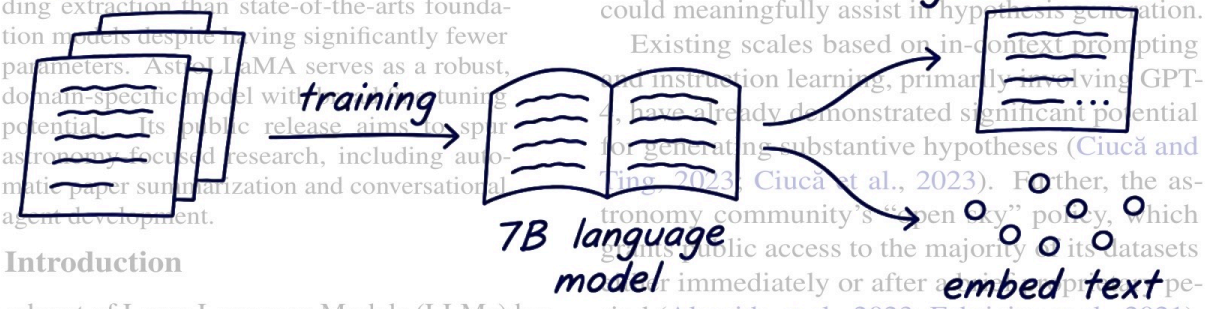
AstroLLaMA : Towards Specialized Foundation Models in Astronomy

Tuan Dung Nguyen[1, 2*], Yuan-Sen Ting[2, 3*], Ioana Ciucă[2*]
Charles O'Neill[2†], Ze-Chang Sun[4†], Maja Jabłońska[2†], Sandor Kruk[5†]
Ernest Perkowski[5], Jack Miller[2], Jason Li[6], Josh Peek[7] Kartheik Iyer[8],
Tomasz Różański[2,9], Pranav Khetarpal[10], Sharaf Zaman[2], David Brodrick[2]
Sergio J. Rodríguez Méndez[2], Thang Bui[2], Alyssa Goodman[11], Alberto Accomazzi[12],
Jill Naiman[13], Jesse Cranney[2], Kevin Schawinski[14], UniverseTBD
[1]University of Pennsylvania, United States [2]Australian National University, Australia
[3]Ohio State University, United States [4]Tsinghua University, China
[5]European Space Astronomy Centre, Spain
[6]Learning Machines, Australia [7]Space Telescope Science Institute, United States
[8]Columbia University, United States [9]Wrocław University, Poland
[10]Indian Institute of Technology Delhi, India [11]Harvard University, United States
[12]NASA Astrophysics Data System, Harvard & Smithsonian, United States
[13]University of Illinois at Urbana-Champaign [14]Modulos AG

Abstract

Large language models excel in many human-language tasks but often falter in highly specialized domains like scholarly astronomy. To bridge this gap, we introduce AstroLLaMA, a 7-billion-parameter model fine-tuned from LLaMA-2 on over 300,000 astronomy abstracts from arXiv. Optimized for traditional causal language modeling, AstroLLaMA achieves a 30% lower perplexity than Llama-2, showing marked domain adaptation. Our model generates more insightful and scientifically relevant text completions and embedding extraction than state-of-the-arts foundation models despite having significantly fewer parameters. AstroLLaMA serves as a robust, domain-specific model with broad fine-tuning potential. Its public release aims to spur astronomy-focused research, including automatic paper summarization and conversational agent development.

The training examples are abstracts, not images or full papers.



Specialization changes the language patterns the model learns; it does not certify the facts it generates.

1 Introduction

The advent of Large Language Models (LLMs) has sparked interdisciplinary interest thanks to a confluence of factors: enhanced computational power, bigger datasets, leaps in computation power, and breakthroughs in neural architectures. Flagship models like GPT-4 (OpenAI, 2023), PaLM (Chowdhery et al., 2022; Anil et al., 2023), LLaMA (Touvron et al., 2023; Meta, 2023) have exhibited exceptional versatility in a variety of tasks from logical reasoning and comprehension to creative writing, often accomplished via methods like prompting, fine-tuning, and human-in-the-loop reinforcement learning.

The astronomy discipline presents both a unique challenge and a fertile ground for the application of LLMs. First, the corpus of scholarly texts in astronomy likely constitutes but a minuscule portion of the data on which generic LLMs are trained, resulting in limitations like hallucinations in favor of more "generic" responses. Second, the nature of astronomical research often involves cross-disciplinary insights due to universally applicable physical processes. Whereas [illegible] could meaningfully assist in hypothesis generation.

Existing scales based on in-context prompting and instruction learning, primarily involving GPT-4, have already demonstrated significant potential for generating substantive hypotheses (Ciucă and Ting, 2023; Ciucă et al., 2023). Further, the astronomy community's "open sky" policy, which grants public access to the majority of its datasets either immediately or after a [illegible] period (Almeida et al., 2023; Fabricius et al., 2021), pairs well with the wealth of resources available in archives like NASA's Astrophysics Data System (Accomazzi et al., 2015; Borgman and Wofford, 2021). Such an open-access policy can facilitate deep engagement with the astronomical literature.

Despite their general capabilities, LLMs frequently lag behind specialized, smaller models in domain-specific applications. This disparity stems from two primary factors: (i) the eclectic nature of the training datasets, which dilutes the focus on specialized subjects, and (ii) the design ethos

*Lead contribution. Email: joshtn@seas.upenn.edu
†Major contribution.

LLaMA-2 → astronomy abstracts → AstroLLaMA

Repeated passes through the abstracts improve next-token prediction.

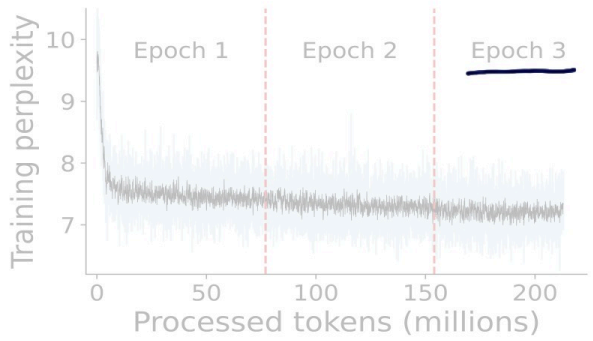
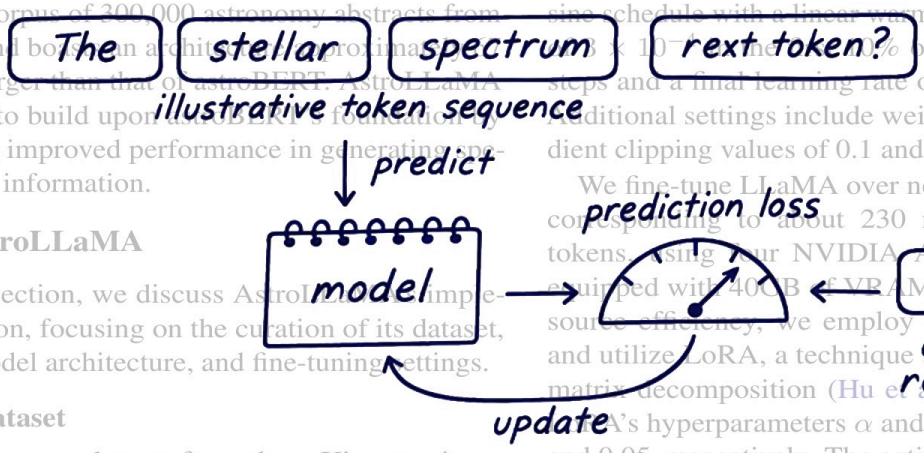


Figure 1: Learning curve of AstroLLaMA during its fine-tuning on the arXiv astrophysics dataset. The Fig.tracks the evolution of perplexity, a measure of the model's next-token prediction performance. The light blue curve shows the training perplexity at each AdamW update step, while the dark black curve provides a smoothed average taken over 10-step intervals.

of LLMs as "foundation models" meant for subsequent fine-tuning tailored to specific tasks. The existing landscape for fine-tuned LLMs in astronomy remains limited, however. To our knowledge, the only existing specialized model is astroBERT (Grezes et al., 2021), which has 110 million parameters, trained on nearly 400,000 ADS papers. But as an non-generative model, the utility of astroBERT remains limited to discriminative tasks. Motivated by these gaps, we present AstroLLaMA, a state-of-the-art generative language model fine-tuned from LLaMA-2. Our model leverages a corpus of 300,000 astronomy abstracts from arXiv and boasts an architecture approximately 30 times larger than that of astroBERT. AstroLLaMA aspires to build upon astroBERT's foundation by offering improved performance in generating specialized information.

We train on the preceding words to predict the next token.



2 AstroLLaMA

In this section, we discuss AstroLLaMA's implementation, focusing on the curation of its dataset, base model architecture, and fine-tuning settings.

2.1 Dataset

We derive our dataset from the arXiv repository, available on Kaggle.[†] Our curated subset focuses on papers classified under the astrophysics category (`astro-ph`), resulting in a collection of 326,238 articles spanning from April 1992 to July 2023. We extract the these papers' abstracts to form a corpus consisting of approximately 95 million tokens. The median length of these abstracts is 291 tokens. To enable effective model evaluation, we randomly designate 20% of this curated dataset for testing.

[†] https://www.kaggle.com/Cornell-University/arxiv

2.2 Base Model

Our base model is LLaMA-2, a 6.7 billion-parameter model developed by Meta (Meta, 2023). Originally trained on a corpus containing 2 trillion tokens, LLaMA-2 features a context window of 4,096 tokens. For tokenization, the model employs a bytepair encoding strategy (Sennrich et al., 2016; Kudo and Richardson, 2018), incorporating a vocabulary set of 32,000 unique tokens.

A fifth of the abstracts is held out for testing.

2.3 Fine-Tuning Settings

For the fine-tuning phase, we rely on our curated training set described in Section 2.1, which includes 77 million tokens. Special `[BOS]` (Beginning Of Sequence) and `[EOS]` (End Of Sequence) tokens are prepended and appended to each training sequence. These sequences are then concatenated and divided into fixed-length chunks, each comprising 512 tokens.

The fine-tuning process follows the causal language modeling objective employed during the model's pre-training phase. We use the AdamW optimizer (Loshchilov and Hutter, 2018) with hyperparameters $\beta_1 = 0.9, \beta_2 = 0.95, \epsilon = 10^{-5}$ and a batch size of 32. The learning rate follows a cosine schedule with a linear warmup to a peak value of 3×10^{-4} after the first 1% of the optimization steps and a final learning rate of 10% of its peak. Additional settings include weight decay and gradient clipping values of 0.1 and 1.0, respectively.

We fine-tune LLaMA over nearly three epochs, corresponding to about 230 million processed tokens, using four NVIDIA A100 GPUs, each equipped with 40GB of VRAM. To maximize resource efficiency, we employ 4-bit quantization and utilize LoRA, a technique based on low-rank matrix decomposition (Hu et al., 2021). We set LoRA's hyperparameters α and dropout rate to 32 and 0.05, respectively. The entire process is facilitated through the Hugging Face Python library.

2.4 Fine-Tuning Evaluation

Fig. 1 depicts the performance of AstroLLaMA during its fine-tuning phase. Here, we present perplexity, a commonly used metric for evaluating causal language models. Perplexity is defined as

Lower perplexity means less surprise at the text, not a guarantee that a statement is true.

A plausible abstract completion can still invent scientific details.

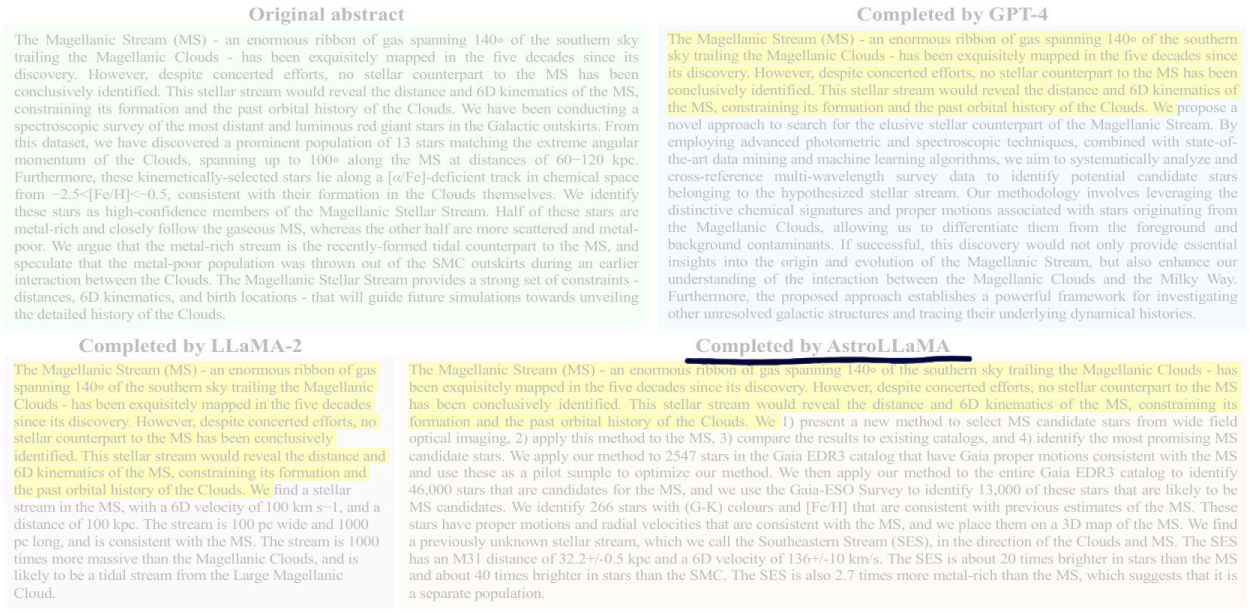


Figure 2: Completion of an abstract from the arXiv database (ID: 2306.15719) using three different models: GPT-4, LLaMA-2, and AstroLLaMA. Each model is prompted with the same short text snippet, highlighted in their respective boxes. GPT-4 tends to produce more generic statements, lacking domain-specific nuance. AstroLLaMA demonstrates the most robust completion, offering more relevant concepts and deeper insights specific to the field of astronomy, thus significantly outperforming LLaMA-2 and GPT-4.

These are illustrative completions, not verified measurements.

the exponentiation of the training loss, with lower values indicating a better fit.

Our initial observations reveal that LLaMA-2 performs suboptimally on our dataset, with an average perplexity close to 10. By the conclusion of three epoch, AstroLLaMA achieves an average perplexity of 6.55. This represents a 32.5% reduction in perplexity compared to the base LLaMA-2 model, signifying a substantial improvement in the model's predictive accuracy.

3 Results

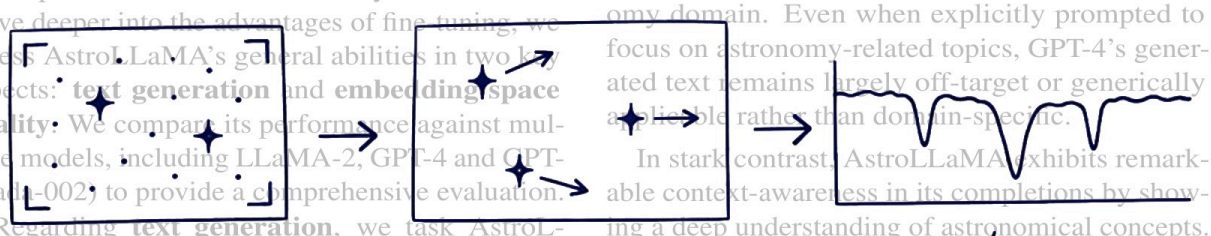
As illustrated in the previous section, AstroLLaMA outperforms its non-fine-tuned counterpart, LLaMA-2, in terms of context-aware token prediction within astronomy abstracts. To delve deeper into the advantages of fine-tuning, we assess AstroLLaMA's general abilities in two key aspects: **text generation** and **embedding space quality**. We compare its performance against multiple models, including LLaMA-2, GPT-4 and GPT-3 (ada-002) to provide a comprehensive evaluation.

Regarding **text generation**, we task AstroLLaMA and LLaMA-2 with completing various astronomy-related abstracts, an example of which is presented in Fig. 2. Each model is given the first few sentences of an abstract as a prompt, allowing us to gauge its ability to comprehend the context and generate a meaningful continuation. For GPT-4, we utilize ChatGPT and specifically prompt it to limit the completion to a single paragraph. AstroLLaMA and LLaMA-2 are deployed using standard sampling methods, with the temperature set to 0.3 and a maximum new tokens limit of 1,024. We find that altering the temperature setting does not substantively improve LLaMA-2's results.

Our observations largely echo the patterns depicted in Fig. 2. LLaMA-2 often deviates from the intended context after generating only a short and often off-topic continuation, resulting in inferior completions. While GPT-4 produces more coherent text, its responses are too generic to capture the nuanced understanding required in the astronomy domain. Even when explicitly prompted to focus on astronomy-related topics, GPT-4's generated text remains largely off-target or generically applicable rather than domain-specific.

In stark contrast, AstroLLaMA exhibits remarkable context-awareness in its completions by showing a deep understanding of astronomical concepts. For example, in Fig. 2, AstroLLaMA comprehends that an effective search for stars in the Magellanic Stream involves a three-step process: initial wide-field imaging, followed by refinement using

A plausible search strategy



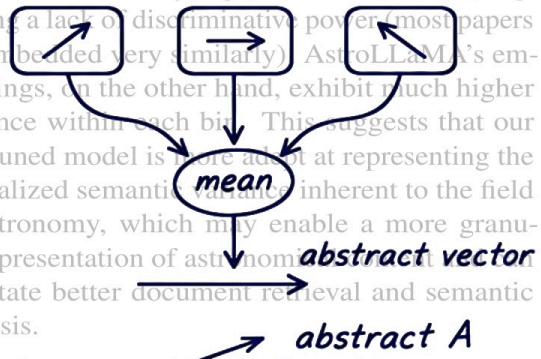
The completion recovers a sensible search sequence, yet its numerical claims can still be invented.

More varied embeddings do not remove the need to check factual accuracy.

metric data from Gaia, and then further curation with spectroscopic data. The model also understands Gaia-ESO is surveying the southern sky and hence can observe (part of) the Magellanic Stream. It also demonstrates nuanced knowledge of the Magellanic Stream, understanding the importance of bifurcation within the stream. As a result, it appropriately completes the text by discussing the southeast stream and exploring metallicity differences to ascertain their origins.

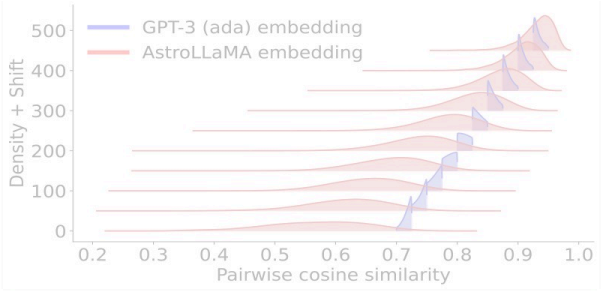
Regarding **embedding space quality**, we assess models' ability to reflect semantic similarities among astronomy texts. We randomly choose 10,000 abstracts from our dataset and embed them using AstroLLaMA and GPT-3. Specifically, we use OpenAI's API to invoke the text embedding function for GPT-3 (ada-002). To get text embeddings from AstroLLaMA, we pass an input through the model and extract its final hidden states, which contain embeddings for all tokens in the input. Then, we omit the `[BOS]` token and take the average of all other tokens' embeddings to get the final result. Finally, for each pair of abstracts we calculate their cosine similarity (the normalised dot product) between on their vector embeddings.

Average the token states to represent one abstract; compare directions to compare texts.



The top panel of Fig. 3 presents the distribution of these pairwise similarities for the two embedding methods. We find that the embeddings by GPT-3 are overly generic with similarities clustering around relatively high values of 0.7–0.9, suggesting a lack of discriminative power (most papers are embedded very similarly). AstroLLaMA's embeddings, on the other hand, exhibit much higher variance within each bin. This suggests that our fine-tuned model is more adept at representing the specialized semantic variance inherent to the field of astronomy, which may enable a more granular representation of astronomical texts and thus facilitate better document retrieval and semantic analysis.

The bottom panel of Fig. 3 provides two representative examples where AstroLLaMA and GPT-3 classifications diverge. In the first example, GPT-3 fixates on the keyword "magnetized," resulting in an inflated similarity score, despite the contexts being markedly different. AstroLLaMA, on the other hand, successfully distinguishes between these disparate contexts. In the second example, AstroLLaMA accurately identifies that the study of Spitzer is closely related to star formation. GPT-3, however, fails to make this connection due to the absence of matching keywords.



Paper 1: Astrophysical gyrokinetics: kinetic and fluid turbulent cascades in magnetized weakly collisional plasma
Paper 2: Comment on modified Coulomb law in a strongly magnetised vaccum
GPT-3 cosine similarity score: **78.5%**
AstroLLaMa cosine similarity score: 36.3%

Paper 1: A Spitzer census of the IC 348 nebula
Paper 2: Sequential and spontaneous star formation around the mid-infrared halo HII region KR 14
GPT-3 cosine similarity score: 82.4%
AstroLLaMa cosine similarity score: **92.8%**

Figure 3: *Top:* Distribution of pairwise cosine similarities among 10,000 randomly selected abstracts from our corpus, divided into 10 equal bins based on similarity levels from GPT-3. *Bottom:* Two representative examples illustrating divergent cosine similarity values when comparing AstroLLaMA and GPT-3 embeddings.

4 Limitations and Future Directions

In this work, we introduce AstroLLaMA, a 7-billion-parameter language model fine-tuned on a dataset encompassing over 300,000 abstracts from astronomical research papers. Compared to its base model, LLaMA-2, and even GPT-4, a current state-of-the-art general LLM, AstroLLaMA exhibits marked improvements in generating high-quality abstracts with a competent grasp of relevant information in this literature.

AstroLLaMA is not without limitations, nevertheless. The most salient is the model's knowledge gaps in certain areas of astronomy: in Fig. 2, AstroLLaMA's estimation of potential star candidates from Gaia-ESO data is notably inaccurate. To address such issues, we are in the process of enriching AstroLLaMA's training set with not just abstracts but the full LaTeX sources of existing astronomy articles, thereby expanding the token count by approximately two orders of magnitude. Another concern lies in the model's tendency to generate hallucinated or fictitious numerical data, an issue likely attributed to our focus on reducing perplexity rather than explicitly steering the model towards factual accuracy. The release of AstroLLaMA aims to facilitate community engagement, both for addressing these inaccuracies and for refining its bal-

A wider spread of similarities is suggestive, but these examples are not a factuality benchmark.

The released model is a starting point for further adaptation and evaluation.

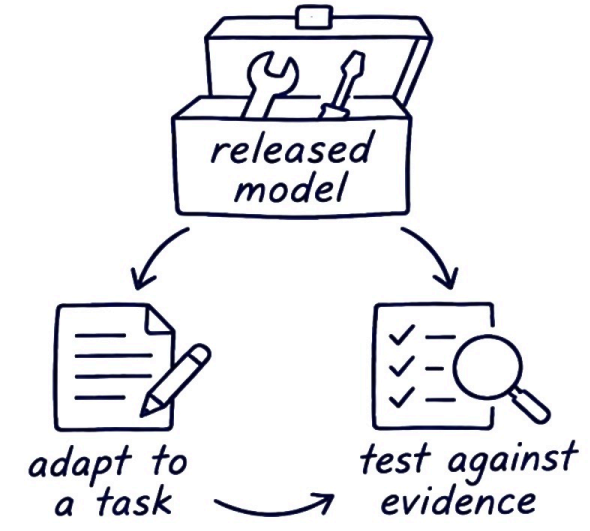
ance between "faithfulness" (respecting scientific evidence and accuracy) and "creativity" (being able to come up with interesting hypotheses).

AstroLLaMA stands as a compelling prototype for specialized LLMs in astronomy, showing superior context-aware capabilities compared to GPT-4 despite having much fewer parameters. It not only paves the way for improved performance in tasks like question-answering, scientific summarization and hypothesis generation but applies also to multi-modal models (Liu et al., 2023). We have made the AstroLLaMA's weights and its training data publicly available[†] for researchers interested in leveraging LLMs for astronomy-centric applications. Along with this, we are establishing various "playgrounds" on Hugging Face to invite interested readers to further adapt and refine this robust starting point for a variety of relevant downstream tasks.

Acknowledgments

We are deeply grateful to the Microsoft Accelerate Foundation Models Research Initiative for enabling us to fast-track our project. Thanks to advanced AI platform from Microsoft Research, we have been able to significantly expedite our efforts in using language models to analyze astronomical literature.

Open weights let us adapt the model; trustworthy answers still require independent checks.



[†] https://huggingface.co/universeTBD/astrollama

Ethics Statement

We obtain the pre-trained weights for LLaMA-2 from Meta, which offers these models for download on Hugging Face. The arXiv dataset used in this paper is publicly available on Kaggle. While we have demonstrated that AstroLLaMA is capable of generating high-quality, relevant abstracts for astronomical research papers, we have noted that it has the potential to generate inaccurate data and measurements. This should serve as a caution for researchers aiming to use this model for downstream tasks, and we invite the adoption of alignment strategies in future work to ameliorate this issue.

References

Google AI PaLM 2. https://ai.google/discover/palm2/.

A. Accomazzi, M. J. Kurtz, E. A. Henneken, R. Chyla, J. Luker, C. S. Grant, D. M. Thompson, A. Holachek, R. Dave, and S. S. Murray. 2015. ADS: The Next Generation Search Platform. In *Open Science at the Frontiers of Librarianship*, volume 492 of *Astronomical Society of the Pacific Conference Series*, page 189.

Andrés Almeida, Scott F. Anderson, Maria Argudo-Fernández, Carles Badenes, Kat Barger, Jorge K. Barrera-Ballesteros, Chad F. Bender, Erika Benitez, Felipe Besser, Dmitry Bizyaev, Michael R. Blanton, John Bochanski, Jo Bovy, William Nielsen Brandt, Joel R. Brownstein, Johannes Buchner, Esra Bulbul, Joseph N. Burchett, Mariana Cano Díaz, Joleen K. Carlberg, Andrew R. Casey, Vedant Chandra, Brian Cherinka, Cristina Chiappini, Abigail A. Coker, Johan Comparat, Charlie Conroy, Gabriella Contardo, Arlin Cortes, Kevin Covey, Jeffrey D. Crane, Katia Cunha, Collin Dabbieri, James W. Davidson Jr. au2, Megan C. Davis, Nathan De Lee, José Eduardo Méndez Delgado, Sebastian Demasi, Francesco Di Mille, John Donor, Peter Dow, Tom Dwelly, Mike Eracleous, Jamey Eriksen, Xiaohui Fan, Emily Farr, Sara Frederick, Logan Fries, Peter Frinchaboy, Boris T. Gaensicke, Junqiang Ge, Consuelo González Ávila, Katie Grabowski, Catherine Grier, Guillaume Guiglion, Pramod Gupta, Patrick Hall, Keith Hawkins, Christian R. Hayes, J. J. Hermes, Lorena Hernández-García, David W. Hogg, Jon A. Holtzman, Hector Javier Ibarra-Medel, Alexander Ji, Paula Jofre, Jennifer A. Johnson, Amy M. Jones, Karen Kinemuchi, Matthias Kluge, Anton Koekemoer, Juna A. Kollmeier, Marina Kounkel, Dhanesh Krishnarao, Mirko Krumpe, Ivan Lacerna, Paulo Jakson Assuncao Lago, Chervin Laporte, Ang Liu, Chao Liu, Xin Liu, Alexandre Roman Lopes, Matin Macktoobian, Viktor Malanushenko, Dan Maoz, Thomas Masseron, Karen L. Masters,

Gal Matijevic, Aidan McBride, Ilija Medan, Andrea Merloni, Sean Morrison, Natalie Myers, Szabolcs Mészáros, C. Alenka Negrete, David L. Nidever, Christian Nitschelm, Audrey Oravetz, Daniel Oravetz, Kaike Pan, Yingjie Peng, Marc H. Pinsonneault, Rick Pogge, Dan Qiu, Anna Barbara de Andrade Queiroz, Solange V. Ramirez, Hans-Walter Rix, Daniela Fernández Rosso, Jessie Runnoe, Mara Salvato, Sebastian F. Sanchez, Felipe A. Santana, Andrew Saydjari, Conor Sayres, Kevin C. Schlaufman, Donald P. Schneider, Axel Schwope, Javier Serna, Yue Shen, Jennifer Sobeck, Ying-Yi Song, Diogo Souto, Taylor Spoo, Keivan G. Stassun, Matthias Steinmetz, Ilya Straumit, Guy Stringfellow, José Sánchez-Gallego, Manuchehr Taghizadeh-Popp, Jamie Tayar, Ani Thakar, Patricia B. Tissera, Andrew Tkachenko, Hector Hernandez Toledo, Benny Trakhtenbrot, Jose G. Fernandez Trincado, Nicholas Troup, Jonathan R. Trump, Sarah Tuttle, Natalie Ulloa, Jose Antonio Vazquez-Mata, Pablo Vera Alfaro, Sandro Villanova, Stefanie Wachter, Anne-Marie Weijmans, Adam Wheeler, John Wilson, Leigh Wojno, Julien Wolf, Xiang-Xiang Xue, Jason E. Ybarra, Eleonora Zari, and Gail Zasowski. 2023. The eighteenth data release of the sloan digital sky surveys: Targeting and first spectra from sdss-v.

Christine L. Borgman and Morgan F. Wofford. 2021. From Data Processes to Data Products: Knowledge Infrastructures in Astronomy. *arXiv e-prints*, page arXiv:2109.01707.

Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, Parker Schuh, Kensen Shi, Sasha Tsvyashchenko, Joshua Maynez, Abhishek Rao, Parker Barnes, Yi Tay, Noam Shazeer, Vinodkumar Prabhakaran, Emily Reif, Nan Du, Ben Hutchinson, Reiner Pope, James Bradbury, Jacob Austin, Michael Isard, Guy Gur-Ari, Pengcheng Yin, Toju Duke, Anselm Levskaya, Sanjay Ghemawat, Sunipa Dev, Henryk Michalewski, Xavier Garcia, Vedant Misra, Kevin Robinson, Liam Fedus, Denny Zhou, Daphne Ippolito, David Luan, Hyeontaek Lim, Barret Zoph, Alexander Spiridonov, Ryan Sepassi, David Dohan, Shivani Agrawal, Mark Omernick, Andrew M. Dai, Thanumalayan Sankaranarayana Pillai, Marie Pellat, Aitor Lewkowycz, Erica Moreira, Rewon Child, Oleksandr Polozov, Katherine Lee, Zongwei Zhou, Xuezhi Wang, Brennan Saeta, Mark Diaz, Orhan Firat, Michele Catasta, Jason Wei, Kathy Meier-Hellstern, Douglas Eck, Jeff Dean, Slav Petrov, and Noah Fiedel. 2022. PaLM: Scaling Language Modeling with Pathways.

Ioana Ciucă and Yuan-Sen Ting. 2023. Galactic ChitChat: Using Large Language Models to Converse with Astronomy Literature. *arXiv e-prints*, page arXiv:2304.05406.

Ioana Ciucă, Yuan-Sen Ting, Sandor Kruk, and Kartheik Iyer. 2023. Harnessing the Power of Adversarial Prompting and Large Language Models for Robust Hypothesis Generation in Astronomy. *arXiv e-prints*, page arXiv:2306.11648.

C. Fabricius, X. Luri, F. Arenou, C. Babusiaux, A. Helmi, T. Muraveva, C. Reylé , F. Spoto, A. Vallenari, T. Antoja, E. Balbinot, C. Barache, N. Bauchet, A. Bragaglia, D. Busonero, T. Cantat-Gaudin, J. M. Carrasco, S. Diakité, M. Fabrizio, F. Figueras, A. Garcia-Gutierrez, A. Garofalo, C. Jordi, P. Kervella, S. Khanna, N. Leclerc, E. Licata, S. Lambert, P. M. Marrese, A. Masip, P. Ramos, N. Robichon, A. C. Robin, M. Romero-Gómez, S. Rubele, and M. Weiler. 2021. Gaia early data release 3. *Astronomy & Astrophysics*, 649:A5.

Felix Grezes, Sergi Blanco-Cuaresma, Alberto Accomazzi, Michael J. Kurtz, Golnaz Shapurian, Edwin Henneken, Carolyn S. Grant, Donna M. Thompson, Roman Chyla, Stephen McDonald, Timothy W. Hostetler, Matthew R. Templeton, Kelly E. Lockhart, Nemanja Martinovic, Shinyi Chen, Chris Tanner, and Pavlos Protopapas. 2021. Building astroBERT, a language model for Astronomy & Astrophysics.

Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2021. LoRA: Low-Rank Adaptation of Large Language Models.

Taku Kudo and John Richardson. 2018. SentencePiece: A simple and language independent subword tokenizer and detokenizer for Neural Text Processing. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 66–71.

Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2023. Visual instruction tuning.

Ilya Loshchilov and Frank Hutter. 2018. Decoupled Weight Decay Regularization. In *International Conference on Learning Representations*.

Meta. 2023. Llama 2: Open Foundation and Fine-Tuned Chat Models | Meta AI Research. https://ai.meta.com/research/publications/llama-2-open-foundation-and-fine-tuned-chat-models/.

OpenAI. 2023. GPT-4 Technical Report.

Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. Neural machine translation of rare words with subword units. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1715–1725.

Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurelien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. 2023. LLaMA: Open and Efficient Foundation Language Models.