

Localizing AI: Evaluating Open-Weight Language Models for Languages of Baltic States

Jurgita Kapočiūtė-Dzikienė^{1,2}, Toms Bergmanis^{3,4}, Mārcis Pinnis^{3,4}

¹ Tilde IT, Lithuania

² Faculty of Informatics, Vytautas Magnus University, Lithuania

³ Tilde, Latvia

⁴ Faculty of Computing, University of Latvia

{name.surname}@tilde.com

Abstract

Although large language models (LLMs) have transformed our expectations of modern language technologies, concerns over data privacy often restrict the use of commercially available LLMs hosted outside of EU jurisdictions. This limits their application in governmental, defence, and other data-sensitive sectors. In this work, we evaluate the extent to which locally deployable open-weight LLMs support lesser-spoken languages such as Lithuanian, Latvian, and Estonian. We examine various size and precision variants of the top-performing multilingual open-weight models, Llama 3, Gemma 2, Phi, and NeMo, on machine translation, multiple-choice question answering, and free-form text generation. The results indicate that while certain models like Gemma 2 perform close to the top commercially available models, many LLMs struggle with these languages. Most surprisingly, however, we find that these models, while showing close to state-of-the-art translation performance, are still prone to lexical hallucinations with errors in at least 1 in 20 words for all open-weight multilingual LLMs.

1 Introduction

Since the fall of 2022, OpenAI and other big tech companies have transformed LLMs from an obscure technology little known outside the academic circles to a major household name. Key to this was the LLMs' ability to perform tasks specified in free-form instructions, making them excel as NLP tools¹. Furthermore, these models

can learn during inference from relevant examples provided as inputs, making them adaptable to new requirements or even tasks. Moslem et al. (2023) showed that in such a setting, GPT-3, provided with relevant translation examples, outperforms machine translation systems of major companies, including Google, DeepL, and ModernMT.

However, data privacy concerns often constrain the use of commercially available LLMs hosted outside EU jurisdiction, limiting their application in governmental, defence, and data-sensitive private sectors. Fine-tuning and operational deployment of adapted models can incur prohibitive costs in the case of commercially available LLMs, emphasizing the need for sovereign AI solutions—locally deployable alternatives that ensure security, control, and compliance. Recently, many powerful alternatives to the commercially available online LLMs have emerged (Jiang et al., 2023; Dubey et al., 2024; Team et al., 2024; Mesnard et al., 2024; Abdin et al., 2024). Although many of these LLMs officially support only a handful of languages with a large speaker base, their training data often incorporate texts from many other languages. Therefore, in practice, these languages receive some degree of support. However, the extent to which these languages are supported, to the best of our knowledge, still needs to be evaluated.

In this work, we aim to answer the question of to what extent, if at all, several popular open-weight models support Lithuanian, Latvian, and Estonian. All three languages have relatively small speaker bases and thus are unlikely to be focal points of major multilingual open-weight LLMs. We examine variants of Meta's Llama 3, Google's Gemma2, Mistral's NeMo, and Microsoft's Phi3 in their performance in multiple-choice question answering (MCQA) and machine translation (MT). We also manually assess the text quality generated by these models by identifying

¹<https://artificialanalysis.ai/leaderboards/models>.

the rate of incorrect words when answering open-ended questions.

We find that while some models like Gemma 2 nearly match the performance of top commercial models, many LLMs struggle with these languages. Surprisingly, even those models that approach state-of-the-art translation capabilities are still susceptible to lexical hallucinations.

2 Experimental Setting

We evaluate LLMs on multiple-choice question answering, machine translation, and text generation quality in open-ended question answering. While our experiments focus on the model performance for three languages—Lithuanian, Latvian, and Estonian—each with a speaker base under 3 million, we also include results for Czech and English for comparison purposes. We have chosen to evaluate models that consistently appear on various leaderboards. Specifically, we assess the **8.03B** and **70.6B** parameter versions of **Llama 3** and **Llama 3.1**, as well as the **3.21B** parameter version of **Llama 3.2** (Dubey et al., 2024) from Meta; the **9.24B** and **27.2B** parameter versions of **Gemma 2** (Team et al., 2024; Mesnard et al., 2024) by Google; the **3.8B** and **14B** versions of **Phi 3** by Microsoft (Abdin et al., 2024); and the **12.2B** parameter **NeMo** by Mistral AI (Jiang et al., 2023). To provide context for our experiments, we include online models by OpenAI such as **GPT-3.5 Turbo** and **GPT-4o** (OpenAI et al., 2024) and **DeepL** machine translation systems. In experiments assessing the quality of Lithuanian text generation, we incorporate the Lithuanian language-specific fine-tuned versions of Llama 2 (Touvron et al., 2023) with **7B** and **13B** parameters, developed by Neurotechnology – **Lt-Llama 2** (Nakvosas et al., 2024).

We run LLMs on our local hardware using the default inference parameters of the Ollama platform², which offers several levels of precision for model quantization: 4bit, 8bit, and full-precision – 16bit. By default, we use **4bit precision** in all our experiments, albeit at the cost of some performance degradation. We also evaluate the performance drop due to quantization by contrasting the results of quantized models with their full-precision counterparts.

²<https://ollama.com/>

Machine Translation For MT experiments, we use the FLORES-200 benchmark dataset (Goyal et al., 2021; Costa jussà et al., 2022), which comprises parallel sentences in over 200 languages³. We use the *devtest* subset from FLORES-200, which contains 1,012 sentences. We test LLMs in a zero-shot inference scenario. We use the following English prompt to request text translations from the specified source and to the specified target language:

“ $\{lang_a\}$: $\{sentence_{lang_a}\}$
Translate the above $\{lang_a\}$ text into $\{lang_b\}$
 $\{lang_b\}$: ”

The translation and evaluation are performed at *sentence-level*; the inference is conducted in a single run for each test sentence. For automatic evaluation of MT quality, we use COMET⁴ (Rei et al., 2020, 2022) as it has been shown to have a higher correlation with human judgments than BLEU (Papineni et al., 2002) and to be more suitable for unrelated system comparison (Kocmi et al., 2024).

Multiple-Choice Question Answering For MCQA experiments, we employ the Belebele dataset, a benchmark in multiple-choice machine reading comprehension (Bandarkar et al., 2024). This dataset pairs each question with a short passage from the FLORES-200 dataset. Each question includes four multiple-choice answers, with one being the correct option. The dataset consists of 900 questions involving 488 distinct passages, each linked to one or two related questions. We use LLMs in a zero-shot inference scenario. We use the following English prompt where “ $\{context\}$ ”, “ $\{question\}$ ” and “ $\{answer_{\#}\}$ ” are in a specific language (Latvian, Estonian, etc.):

“This is the context: ' $\{context\}$ '. This is the question: ' $\{question\}$ '. Here are the 4 candidate answers: '1) $\{answer_1\}$ '; '2) $\{answer_2\}$ '; '3) $\{answer_3\}$ '; '4) $\{answer_4\}$ '. Report only the correct answer's ID (1, 2, 3, 4) using the mandatory JSON format: $\{answer_id : " \}$.”

The prompt explicitly requests the ID (e.g. ‘1’, ‘2’, ‘3’, or ‘4’) of the correct answer formatted in JSON. Our evaluation metric is accuracy.

³<https://github.com/facebookresearch/flores/tree/main/flores200>.

⁴<https://huggingface.co/Unbabel/wmt22-comet-da>

DeepL		GPT		Llama: 3		3.1		3.2	NeMo	Gemma 2		Phi 3	
		3.5-T	4o	8B	70B	8B	70B	3B	12B	9B	27B	3B	14B
EN-LT	0.92	0.88	0.91	0.62	0.83	0.61	0.84	0.46	0.73	0.86	0.89	0.26	0.32
EN-LV	0.92	0.88	0.91	0.59	0.82	0.58	0.83	0.44	0.72	0.83	0.88	0.25	0.27
EN-ET	0.93	0.92	0.92	0.65	0.86	0.63	0.87	0.48	0.75	0.84	0.89	0.27	0.37
EN-CS	0.93	0.91	0.92	0.81	0.90	0.82	0.90	0.66	0.84	0.89	0.91	0.25	0.51
LT-EN	0.87	0.86	0.88	0.77	0.81	0.77	0.82	0.75	0.82	0.87	0.87	0.32	0.34
LV-EN	0.89	0.87	0.89	0.77	0.83	0.78	0.82	0.76	0.84	0.87	0.88	0.33	0.34
ET-EN	0.90	0.90	0.90	0.78	0.83	0.79	0.82	0.76	0.85	0.88	0.89	0.33	0.34
CS-EN	0.89	0.89	0.89	0.86	0.88	0.86	0.87	0.86	0.87	0.89	0.89	0.32	0.33
Avg.	0.91	0.89	0.90	0.73	0.85	0.73	0.85	0.65	0.80	0.87	0.89	0.29	0.35

Table 1: Automatic MT quality evaluation results in COMET scores across models and translation directions. DeepL and OpenAI GPT 3.5-Turbo and 4o are provided for reference. Top results by open-weight models for each translation direction are marked in **bold**.

	GPT		Llama: 3		3.1		3.2	NeMo	Gemma 2		Phi 3
	3.5-T	4o	8B	70B	8B	70B	3B	12B	9B	27B	14B
LT	0.734	0.941	0.607	0.768	0.618	0.834	0.435	0.715	0.861	0.898	0.001
LV	0.756	0.950	0.571	0.710	0.581	0.783	0.410	0.689	0.869	0.914	0.002
EN	0.903	0.962	0.883	0.938	0.872	0.947	0.740	0.898	0.931	0.943	0.886
ET	0.773	0.928	0.576	0.770	0.560	0.821	0.397	0.686	0.859	0.893	0.003
CS	0.818	0.937	0.769	0.888	0.743	0.892	0.676	0.800	0.907	0.910	0.296
Avg.	0.797	0.944	0.681	0.815	0.675	0.855	0.532	0.758	0.885	0.912	0.238

Table 2: Automatic MCQA evaluation results measuring accuracies across models and languages. OpenAI GPT 3.5-Turbo and 4o are provided for reference. Top results by open-weight models for each language are marked in **bold**.

	Llama 3.1		3.2	Gemma 2	
	8B	70B	3B	9B	27B
Δ MT	0.074	0.009	0.015	0.006	0.001
Δ MCQA	0.100	0.058	0.004	0.010	0.000

Table 3: Performance drop (difference between Avg. scores across all languages) for several 4bit models compared to their respective full precision versions on the two tasks – MT (COMET points) and MCQA (accuracy).

Text Quality in Free Form Question Answering

To assess LLMs’ ability to generate answers that adhere to Lithuanian and Latvian conventions and grammatical norms, we prompt models to answer ten free-form questions such as “*When did the Soviet Union collapse, how many new countries appeared, and what are their names?*” and “*Provide a definition of artificial intelligence.*” We conduct human evaluation by two Lithuanian and Latvian native speakers and linguistics experts. We require evaluators to count text errors, mark grammatically incorrect words or incorrect inflexions, mark invented words not existing in the language, and mark words within syntactically incorrect sentence structures (see Table 4).

We also assess whether the provided answers are factually correct. However, the factual accuracy results lack statistical significance due to the small sample size and should be interpreted with caution. For instance, it happened that GPT-4 answered all ten questions correctly for the Latvian language, but this outcome reflects a preliminary observation rather than a deep investigation. The results, therefore, should be viewed as part of a pilot study and not as definitive findings.

3 Results and Discussion

MT evaluation results (see Table 1) demonstrate the Gemma 2 family as the most capable open-weight model family. Gemma 2 27B emerges as the best locally deployable model, yielding COMET scores on par with OpenAI’s GPT-3.5 Turbo and only marginally worse than GPT-4o. Although specialised proprietary MT models like DeepL achieve the highest average score (0.91), freely available Gemma 2 models are not far off, with average COMET scores of 0.89 and 0.87 for the 27B and 9B versions, respectively. In this context, the Llama family has little to offer, with the Llama 3.0 and 3.1 70B param-

		GPT 4o		Llama 3.1 8B 70B	Gemma 2 27B	Lt-Llama 2 7B 13B
LT	Words/Sentences	320/38		724/68 300/24	1273/135	1132/69 1020/58
	Error Rate (%)	3.44		12.98 7.67	4.08	1.15 0.98
	Answer acc.	0.9		0.2 0.9	0.9	0.5 0.4
LV	Words/Sentences	1249/91		362/27 619/39	1171/97	- -
	Error Rate (%)	2.48		18.51 11.31	5.98	- -
	Answer acc.	1		0.3 0.9	0.8	- -

Table 4: Human evaluation results for text generation quality in free form question answering.

ter models surpassing the much smaller Gemma 2 9B only in two out of eight translation directions. Smaller models, like Llama’s 3B and 8B versions and Mistral’s 12B NeMo, show equally meagre results given the high performance of Gemma 2 9B. Lastly, the results of Phi 3 prove that these models have very little support for multilingualism.

Quantisation impact on MT quality analysis (see Table 3) reveals that while Llama models are negatively affected to some degree, the performance of full-precision models does not justify their use either. Increased inference time and memory requirements for the 70B model are too prohibitive unless several top-of-the-shelf enterprise-grade GPUs are available⁵. However, the full precision 8B parameter Llama 3.1 still does not reach the performance of the Gemma 2 9B 4bit version (0.80 vs 0.87). Gemma 2 family models, on the other hand, show a statistically insignificant drop in translation performance when the 4bit version is used, suggesting that their architecture is very robust to quantization.

While the current MT results provide valuable insights into LLM capabilities, future work could benefit from more fine-grained error analysis using frameworks like MQM (Multidimensional Quality Metrics) and ESA (Error Span Annotation). These approaches allow detailed classification of errors-such as those related to accuracy, fluency, and terminology, and help quantify their impact on text usability. Incorporating these methods could provide deeper insights into model limitations and guide targeted improvements, particularly for smaller languages like Lithuanian, Latvian, and Estonian.

MCQA results (see Table 2) show that the Gemma 2 27B parameter model outperforms GPT-3.5 Turbo across all languages, coming second only to OpenAI’s flagship model, GPT-4o. No-

tably, Gemma 2 27B achieves the highest accuracy among the open-weight models, outperforming Llama 3.1 and NeMo models for the Lithuanian, Latvian, Estonian, and Czech. The Phi models, however, perform poorly, particularly in non-English languages, and their results often fail to meet the required JSON output format, providing detailed responses instead of just answer IDs.

Quantization impact on MCQA analysis unveils a similar picture as the analysis above for MT: Llama models are more sensitive to quantization, while Gemma 2 are more robust. As a result, Gemma 2 models show little performance degradation when much more efficient 4bit models are used. It’s worth noting, however, that the accuracy drops because quantization differs depending on the amount of data each language has. Less spoken languages like Lithuanian, Latvian, and Estonian are affected more than English and Czech, for which overall results are better. For example, the Llama 3.1 70B model loses 0.06, 0.12, and 0.08 accuracy for Lithuanian, Latvian, and Estonian, respectively, but only 0.01 and 0.02 for English and Czech.

Text Generation Quality evaluation results (see Table 4) show that most models produce more than one error per 100 words. The Lt-Llama 2 models, specifically fine-tuned for Lithuanian, are the exception, with an error rate of just 0.98% and no invented words. Among multilingual models, OpenAI’s GPT-4o achieves strong performance, with 0.94 grammatically incorrect or incorrectly inflected words per 100 for Lithuanian and 1.52 for Latvian, while generating a very small number of invented words (0.31 and 0.56 per 100 for Lithuanian and Latvian, respectively). In contrast, Llama 3.1 models show significant shortcomings, with the highest frequency of grammatical errors: 8.01 per 100 for Lithuanian and 11.88 for Latvian. Additionally, Llama 3.1 generates a substantial number of invented words: 4.28 per 100 for

⁵<https://huggingface.co/blog/llama31>

Lithuanian and 3.87 for Latvian. Gemma 2 models perform considerably better, with 2.36 grammatical errors per 100 for Lithuanian and 4.10 for Latvian, and fewer invented words: 0.39 per 100 for Lithuanian and 1.45 for Latvian. These findings highlight clear quality differences among models. While Llama 3.1's high error rates make it unsuitable for most commercial applications, Gemma 2 strikes a better balance, approaching GPT-4o's quality but still falling short. Notably, Lt-Llama 2 sets the strongest benchmark with near-perfect output, minimal grammatical errors, and no invented words. On average, users can expect at least one linguistic error in every 2-3 sentences from the best open-weight models like Gemma 2, or every sentence for models like Llama 3.1, unless further multilingual specialization becomes available.

Lesser-spoken Languages like Lithuanian, Latvian, and Estonian have less support in open-weight models compared to more populous languages such as Czech. These differences are more pronounced in smaller and lower-quality models, especially in tasks where models generate text in lesser-spoken languages (e.g., MT from English into Lithuanian). Comparatively good results for Czech suggest that these disparities are related to the amount of data each LLM has seen for each language during training, rather than factors such as the structural complexity of the language.

4 Conclusions

Our findings demonstrate that certain open-weight LLMs, such as the Gemma 2 family, achieve performance comparable to top-tier commercial products, such as general-purpose models like OpenAI's GPT-4o and specialized machine translation services like DeepL. This progress enables local, secure, and private solutions, supporting the development of sovereign AI for many language tasks in governmental, defence, and other data-sensitive private sectors. Nevertheless, unless specifically fine-tuned for languages like Lithuanian, most multilingual models are still surprisingly prone to lexical hallucinations, highlighting the need for 1) high-quality language data for languages of the Baltic states and 2) research on language-specialized LLMs.

Acknowledgments

This research has been supported by the ICT Competence Centre (www.itkc.lv) within the project *2.4 Daudzvalodīgs uzņēmuma informācijas semantiskās meklēšanas un atbilžu gatavošanas risinājums* (2.4 Multilingual Semantic Search and Question-Answering Solution for Enterprises) of EU Structural funds, ID no 5.1.1.2.i.0/1/22/A/CFLA/008.

References

- Marah Abdin, Jyoti Aneja, Hany Awadalla, Ahmed Awadallah, Ammar Ahmad Awan, Nguyen Bach, Amit Bahree, Arash Bakhtiari, Jianmin Bao, Harkirat Behl, Alon Benhaim, Misha Bilenko, Johan Bjorck, Sébastien Bubeck, Martin Cai, Qin Cai, Vishrav Chaudhary, Dong Chen, Dongdong Chen, Weizhu Chen, Yen-Chun Chen, Yi-Ling Chen, Hao Cheng, Parul Chopra, Xiyang Dai, Matthew Dixon, Ronen Eldan, Victor Fragoso, Jianfeng Gao, Mei Gao, Min Gao, Amit Garg, Allie Del Giorno, Abhishek Goswami, Suriya Gunasekar, Emman Haider, Junheng Hao, Russell J. Hewett, Wenxiang Hu, Jamie Huynh, Dan Iter, Sam Ade Jacobs, Mojan Javaheripi, Xin Jin, Nikos Karampatziakis, Piero Kauffmann, Mahoud Khademi, Dongwoo Kim, Young Jin Kim, Lev Kurilenko, James R. Lee, Yin Tat Lee, Yuanzhi Li, Yunsheng Li, Chen Liang, Lars Liden, Xihui Lin, Zeqi Lin, Ce Liu, Liyuan Liu, Mengchen Liu, Weishung Liu, Xiaodong Liu, Chong Luo, Piyush Madan, Ali Mahmoudzadeh, David Majercak, Matt Mazzola, Caio César Teodoro Mendes, Arindam Mitra, Hardik Modi, Anh Nguyen, Brandon Norick, Barun Patra, Daniel Perez-Becker, Thomas Portet, Reid Pryzant, Heyang Qin, Marko Radmilac, Liliang Ren, Gustavo de Rosa, Corby Rosset, Sambudha Roy, Olatunji Ruwase, Olli Saarikivi, Amin Saied, Adil Salim, Michael Santacrose, Shital Shah, Ning Shang, Hiteshi Sharma, Yelong Shen, Swadheen Shukla, Xia Song, Masahiro Tanaka, Andrea Tupini, Praneetha Vaddamanu, Chunyu Wang, Guanhua Wang, Lijuan Wang, Shuohang Wang, Xin Wang, Yu Wang, Rachel Ward, Wen Wen, Philipp Witte, Haiping Wu, Xiaoxia Wu, Michael Wyatt, Bin Xiao, Can Xu, Jiahang Xu, Weijian Xu, Jilong Xue, Sonali Yadav, Fan Yang, Jianwei Yang, Yifan Yang, Ziyi Yang, Donghan Yu, Lu Yuan, Chenruidong Zhang, Cyril Zhang, Jianwen Zhang, Li Lyna Zhang, Yi Zhang, Yue Zhang, Yunan Zhang, and Xiren Zhou. 2024. Phi-3 technical report: A highly capable language model locally on your phone.
- Lucas Bandarkar, Davis Liang, Benjamin Muller, Mikel Artetxe, Satya Narayan Shukla, Donald Husa, Naman Goyal, Abhinandan Krishnan, Luke Zettlemoyer, and Madian Khabsa. 2024. The belebele benchmark: a parallel reading comprehension

dataset in 122 language variants. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 749–775, Bangkok, Thailand and virtual meeting. Association for Computational Linguistics.

Marta R. Costa jussà, James Cross, Onur Çelebi, Maha Elbayad, Kenneth Heafield, Kevin Heffernan, Elahe Kalbassi, Janice Lam, Daniel Licht, Jean Maillard, Anna Y. Sun, Skyler Wang, Guillaume Wenzek, Al Youngblood, Bapi Akula, Loïc Barrault, Gabriel Mejia Gonzalez, Prangthip Hansanti, John Hoffman, Semarley Jarrett, Kaushik Ram Sadagopan, Dirk Rowe, Shannon Spruit, Chau Tran, Pierre Andrews, Necip Fazil Ayan, Shruti Bhosale, Sergey Edunov, Angela Fan, Cynthia Gao, Vedanuj Goswami, Francisco Guzmán, Philipp Koehn, Alexandre Mourachko, Christophe Ropers, Safiyyah Saleem, Holger Schwenk, and Jeff Wang. 2022. No language left behind: Scaling human-centered machine translation. *CoRR*, abs/2207.04672.

Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, Arun Rao, Aston Zhang, Aurelien Rodriguez, Austen Gregerson, Ava Spataru, Baptiste Roziere, Bethany Biron, Binh Tang, Bobbie Chern, Charlotte Caucheteux, Chaya Nayak, Chloe Bi, Chris Marra, Chris McConnell, Christian Keller, Christophe Touret, Chunyang Wu, Corinne Wong, Cristian Canton Ferrer, Cyrus Nikolaidis, Damien Allonsius, Daniel Song, Danielle Pintz, Danny Livshits, David Esobu, Dhruv Choudhary, Dhruv Mahajan, Diego Garcia-Olano, Diego Perino, Dieuwke Hupkes, Egor Lomakin, Ehab AlBadawy, Elina Lobanova, Emily Dinan, Eric Michael Smith, Filip Radenovic, Frank Zhang, Gabriel Synnaeve, Gabrielle Lee, Georgia Lewis Anderson, Graeme Nail, Gregoire Mialon, Guan Pang, Guillem Cucurell, Hailey Nguyen, Hannah Korevaar, Hu Xu, Hugo Touvron, Iliyan Zarov, Imanol Arrieta Ibarra, Isabel Kloumann, Ishan Misra, Ivan Evtimov, Jade Copet, Jaewon Lee, Jan Geffert, Jana Vranes, Jason Park, Jay Mahadeokar, Jeet Shah, Jelmer van der Linde, Jennifer Billock, Jenny Hong, Jenya Lee, Jeremy Fu, Jianfeng Chi, Jianyu Huang, Jiawen Liu, Jie Wang, Jiecao Yu, Joanna Bitton, Joe Spisak, Jongsoo Park, Joseph Rocca, Joshua Johnston, Joshua Saxe, Junteng Jia, Kalyan Vasuden Alwala, Kartikeya Upasani, Kate Plawiak, Ke Li, Kenneth Heafield, Kevin Stone, Khalid El-Arini, Krithika Iyer, Kshitiz Malik, Kuenley Chiu, Kunal Bhalla, Lauren Rantala-Yeary, Laurens van der Maaten, Lawrence Chen, Liang Tan, Liz Jenkins, Louis Martin, Lovish Madaan, Lubo Malo, Lukas Blecher, Lukas Landzaat, Luke de Oliveira, Madeline Muzzi, Mahesh Pasupuleti, Mannat Singh, Manohar Paluri, Marcin Kardas, Mathew Old-

ham, Mathieu Rita, Maya Pavlova, Melanie Kam-badur, Mike Lewis, Min Si, Mitesh Kumar Singh, Mona Hassan, Naman Goyal, Narjes Torabi, Nikolay Bashlykov, Nikolay Bogoychev, Niladri Chatterji, Olivier Duchenne, Onur Çelebi, Patrick Al-rassy, Pengchuan Zhang, Pengwei Li, Petar Vasic, Peter Weng, Prajjwal Bhargava, Pratik Dubal, Praveen Krishnan, Punit Singh Koura, Puxin Xu, Qing He, Qingxiao Dong, Ragavan Srinivasan, Raj Ganapathy, Ramon Calderer, Ricardo Silveira Cabral, Robert Stojnic, Roberta Raileanu, Rohit Girdhar, Rohit Patel, Romain Sauvestre, Ronnie Polidoro, Roshan Sumbaly, Ross Taylor, Ruan Silva, Rui Hou, Rui Wang, Saghar Hosseini, Sahana Chennabasappa, Sanjay Singh, Sean Bell, Seohyun Sonia Kim, Sergey Edunov, Shaoliang Nie, Sharan Narang, Sharath Rapparthi, Sheng Shen, Shengye Wan, Shruti Bhosale, Shun Zhang, Simon Vandenhende, Soumya Batra, Spencer Whitman, Sten Sootla, Stephane Collet, Suchin Gururangan, Sydney Borodinsky, Tamar Herman, Tara Fowler, Tarek Sheasha, Thomas Georgiou, Thomas Scialom, Tobias Speckbacher, Todor Mihaylov, Tong Xiao, Ujjwal Karn, Vedanuj Goswami, Vibhor Gupta, Vignesh Ramanathan, Viktor Kerkez, Vincent Conguet, Virginie Do, Vish Vogeti, Vladan Petrovic, Weiwei Chu, Wenhan Xiong, Wenyin Fu, Whitney Meers, Xavier Martinet, Xiaodong Wang, Xiaoqing Ellen Tan, Xinfeng Xie, Xuchao Jia, Xuewei Wang, Yaelle Goldschlag, Yashesh Gaur, Yasmine Babaei, Yi Wen, Yiwen Song, Yuchen Zhang, Yue Li, Yunying Mao, Zacharie Delpierre Coudert, Zheng Yan, Zhengxing Chen, Zoe Papakipos, Aaditya Singh, Aaron Grattafiori, Abha Jain, Adam Kelsey, Adam Shajnfeld, Adithya Gangidi, Adolfo Victoria, Ahuva Goldstand, Ajay Menon, Ajay Sharma, Alex Boesenberg, Alex Vaughan, Alexei Baevski, Allie Feinstein, Amanda Kallet, Amit Sangani, Anam Yunus, Andrei Lupu, Andres Alvarado, Andrew Caples, Andrew Gu, Andrew Ho, Andrew Poulton, Andrew Ryan, Ankit Ramchandani, Annie Franco, Aparajita Saraf, Arkabandhu Chowdhury, Ashley Gabriel, Ashwin Bharambe, Assaf Eisenman, Azadeh Yazdan, Beau James, Ben Maurer, Benjamin Leonhardi, Bernie Huang, Beth Loyd, Beto De Paola, Bhargavi Paranjape, Bing Liu, Bo Wu, Boyu Ni, Braden Hancock, Bram Wasti, Brandon Spence, Brani Stojkovic, Brian Gamido, Britt Montalvo, Carl Parker, Carly Burton, Catalina Mejia, Changan Wang, Changkyu Kim, Chao Zhou, Chester Hu, Ching-Hsiang Chu, Chris Cai, Chris Tindal, Christoph Feichtenhofer, Damon Civin, Dana Beaty, Daniel Kreymer, Daniel Li, Danny Wyatt, David Adkins, David Xu, Davide Testuggine, Delia David, Devi Parikh, Diana Liskovich, Didem Foss, Dingkan Wang, Duc Le, Dustin Holland, Edward Dowling, Eissa Jamil, Elaine Montgomery, Eleonora Presani, Emily Hahn, Emily Wood, Erik Brinkman, Esteban Arcaute, Evan Dunbar, Evan Smothers, Fei Sun, Felix Kreuk, Feng Tian, Firat Ozgenel, Francesco Caggioni, Francisco Guzmán, Frank Kanayet, Frank Seide, Gabriela Medina Florez, Gabriella Schwarz, Gada Badeer, Georgia Swee, Gil Halpern, Govind

- Thattai, Grant Herman, Grigory Sizov, Guangyi, Zhang, Guna Lakshminarayanan, Hamid Shojanazeri, Han Zou, Hannah Wang, Hanwen Zha, Haroun Habeeb, Harrison Rudolph, Helen Suk, Henry Aspegren, Hunter Goldman, Ibrahim Damla, Igor Molybog, Igor Tufanov, Irina-Elena Veliche, Itai Gat, Jake Weissman, James Geboski, James Kohli, Japhet Asher, Jean-Baptiste Gaya, Jeff Marcus, Jeff Tang, Jennifer Chan, Jenny Zhen, Jeremy Reizenstein, Jeremy Teboul, Jessica Zhong, Jian Jin, Jingyi Yang, Joe Cummings, Jon Carvill, Jon Shepard, Jonathan McPhie, Jonathan Torres, Josh Ginsburg, Junjie Wang, Kai Wu, Kam Hou U, Karan Saxena, Karthik Prasad, Kartikay Khandelwal, Katayoun Zand, Kathy Matosich, Kaushik Veeraraghavan, Kelly Michelena, Keqian Li, Kun Huang, Kunal Chawla, Kushal Lakhotia, Kyle Huang, Lailin Chen, Lakshya Garg, Lavender A, Leandro Silva, Lee Bell, Lei Zhang, Liangpeng Guo, Licheng Yu, Liron Moshkovich, Luca Wehrstedt, Madian Khabisa, Manav Avalani, Manish Bhatt, Maria Tsimpoukelli, Martynas Mankus, Matan Hasson, Matthew Lennie, Matthias Reso, Maxim Groshev, Maxim Naumov, Maya Lathi, Meghan Keneally, Michael L. Seltzer, Michal Valko, Michelle Restrepo, Mihir Patel, Mik Vyatskov, Mikayel Samvelyan, Mike Clark, Mike Macey, Mike Wang, Miquel Jubert Hermoso, Mo Metanat, Mohammad Rastegari, Munish Bansal, Nandhini Santhanam, Natascha Parks, Natasha White, Navyata Bawa, Nayan Singhal, Nick Egebo, Nicolas Usunier, Nikolay Pavlovich Laptev, Ning Dong, Ning Zhang, Norman Cheng, Oleg Chernoguz, Olivia Hart, Omkar Salpekar, Ozlem Kalinli, Parkin Kent, Parth Parekh, Paul Saab, Pavan Balaji, Pedro Rittner, Philip Bontrager, Pierre Roux, Piotr Dollar, Polina Zvyagina, Prashant Ratanchandani, Pritish Yuvraj, Qian Liang, Rachad Alao, Rachel Rodriguez, Rafi Ayub, Raghotham Murthy, Raghu Nayani, Rahul Mitra, Raymond Li, Rebekkah Hogan, Robin Battey, Rocky Wang, Rohan Maheswari, Russ Howes, Rutu Rinott, Sai Jayesh Bondu, Samyak Datta, Sara Chugh, Sara Hunt, Sargun Dhillon, Sasha Sidorov, Satadru Pan, Saurabh Verma, Seiji Yamamoto, Sharadh Ramaswamy, Shaun Lindsay, Shaun Lindsay, Sheng Feng, Shenghao Lin, Shengxin Cindy Zha, Shiva Shankar, Shuqiang Zhang, Shuqiang Zhang, Sinong Wang, Sneha Agarwal, Soji Sajuyigbe, Soumith Chintala, Stephanie Max, Stephen Chen, Steve Kehoe, Steve Satterfield, Sudarshan Govindaprasad, Sumit Gupta, Sungmin Cho, Sunny Virk, Suraj Subramanian, Sy Choudhury, Sydney Goldman, Tal Remez, Tamar Glaser, Tamara Best, Thilo Kohler, Thomas Robinson, Tianhe Li, Tianjun Zhang, Tim Matthews, Timothy Chou, Tzook Shaked, Varun Vontimitta, Victoria Ajayi, Victoria Montanez, Vijai Mohan, Vinay Satish Kumar, Vishal Mangla, Vitor Albiero, Vlad Ionescu, Vlad Poenaru, Vlad Tiberiu Mihalescu, Vladimir Ivanov, Wei Li, Wenchen Wang, Wenwen Jiang, Wes Bouaziz, Will Constable, Xiaocheng Tang, Xiaofang Wang, Xiaojuan Wu, Xiaolan Wang, Xide Xia, Xilun Wu, Xinbo Gao, Yanjun Chen, Ye Hu, Ye Jia, Ye Qi, Yenda Li, Yilin Zhang, Ying Zhang, Yossi Adi, Youngjin Nam, Yu, Wang, Yuchen Hao, Yundi Qian, Yuzi He, Zach Rait, Zachary DeVito, Zef Rosnbrick, Zhaoduo Wen, Zhenyu Yang, and Zhiwei Zhao. 2024. The llama 3 herd of models.
- Naman Goyal, Cynthia Gao, Vishrav Chaudhary, Peng-Jen Chen, Guillaume Wenzek, Da Ju, Sanjana Krishnan, Marc’Aurelio Ranzato, Francisco Guzmán, and Angela Fan. 2021. The flores-101 evaluation benchmark for low-resource and multilingual machine translation.
- Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, Léo Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timothée Lacroix, and William El Sayed. 2023. Mistral 7b.
- Tom Kocmi, Vilém Zouhar, Christian Federmann, and Matt Post. 2024. Navigating the metrics maze: Reconciling score magnitudes and accuracies. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1999–2014, Bangkok, Thailand. Association for Computational Linguistics.
- Gemma Team Thomas Mesnard, Cassidy Hardin, Robert Dadashi, Surya Bhupatiraju, Shreya Pathak, L. Sifre, Morgane Riviere, Mihir Kale, J Christopher Love, Pouya Dehghani Tafti, Léonard Hussenot, Aakanksha Chowdhery, Adam Roberts, Aditya Barua, Alex Botev, Alex Castro-Ros, Ambrose Slone, Am’elie H’elieu, Andrea Tacchetti, Anna Bulanova, Antonia Paterson, Beth Tsai, Bobak Shahriari, Charline Le Lan, Christopher A. Choquette-Choo, Clément Crepy, Daniel Cer, Daphne Ippolito, David Reid, Elena Buchatskaya, Eric Ni, Eric Noland, Geng Yan, George Tucker, George-Christian Muraru, Grigory Rozhdestvenskiy, Henryk Michalewski, Ian Tenney, Ivan Grishchenko, Jacob Austin, James Keeling, Jane Labanowski, Jean-Baptiste Lespiau, Jeff Stanway, Jenny Brennan, Jeremy Chen, Johan Ferret, Justin Chiu, Justin Mao-Jones, Katherine Lee, Kathy Yu, Katie Millican, Lars Lowe Sjoesund, Lisa Lee, Lucas Dixon, Machel Reid, Maciej Mikula, Mateo Wirth, Michael Sharman, Nikolai Chinaev, Nithum Thain, Olivier Bachem, Oscar Chang, Oscar Wahltinez, Paige Bailey, Paul Michel, Petko Yotov, Pier Giuseppe Sessa, Rahma Chaabouni, Ramona Comanescu, Reena Jana, Rohan Anil, Ross McIlroy, Ruibo Liu, Ryan Mullins, Samuel L Smith, Sebastian Borgeaud, Sertan Girgin, Sholto Douglas, Shree Pandya, Siamak Shakeri, Soham De, Ted Klimentko, Tom Hennigan, Vladimir Feinberg, Wojciech Stokowiec, Yu hui Chen, Zafarali Ahmed, Zhitao Gong, Tris Brian Warkentin, Ludovic Peran, Minh Giang, Clément Farabet, Oriol Vinyals, Jeffrey Dean, Koray Kavukcuoglu, Demis Hassabis, Zoubin Ghahramani, Douglas Eck, Joelle Barral, Fernando Pereira,

- Eli Collins, Armand Joulin, Noah Fiedel, Evan Senter, Alek Andreev, and Kathleen Kenealy. 2024. Gemma: Open models based on gemini research and technology. *ArXiv*, abs/2403.08295.
- Yasmin Moslem, Rejwanul Haque, John D. Kelleher, and Andy Way. 2023. Adaptive machine translation with large language models. In *Proceedings of the 24th Annual Conference of the European Association for Machine Translation*, pages 227–237, Tampere, Finland. European Association for Machine Translation.
- Artūras Nakvosas, Povilas Daniušis, and Vytas Mulevičius. 2024. Open llama2 model for the lithuanian language.
- OpenAI, Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, Red Avila, Igor Babuschkin, Suchir Balaji, Valerie Balcom, Paul Baltescu, Haiming Bao, Mohammad Bavarian, Jeff Belgum, Irwan Bello, Jake Berdine, Gabriel Bernadett-Shapiro, Christopher Berner, Lenny Bogdonoff, Oleg Boiko, Madelaine Boyd, Anna-Luisa Brakman, Greg Brockman, Tim Brooks, Miles Brundage, Kevin Button, Trevor Cai, Rosie Campbell, Andrew Cann, Brittany Carey, Chelsea Carlson, Rory Carmichael, Brooke Chan, Che Chang, Fotis Chantzis, Derek Chen, Sully Chen, Ruby Chen, Jason Chen, Mark Chen, Ben Chess, Chester Cho, Casey Chu, Hyung Won Chung, Dave Cummings, Jeremiah Currier, Yunxing Dai, Cory Decareaux, Thomas Degry, Noah Deutsch, Damien Deville, Arka Dhar, David Dohan, Steve Dowling, Sheila Dunning, Adrien Ecoffet, Atty Eleti, Tyna Eloundou, David Farhi, Liam Fedus, Niko Felix, Simón Posada Fishman, Juston Forte, Isabella Fulford, Leo Gao, Elie Georges, Christian Gibson, Vik Goel, Tarun Gogineni, Gabriel Goh, Rapha Gontijo-Lopes, Jonathan Gordon, Morgan Grafstein, Scott Gray, Ryan Greene, Joshua Gross, Shixiang Shane Gu, Yufei Guo, Chris Hallacy, Jesse Han, Jeff Harris, Yuchen He, Mike Heaton, Johannes Heidecke, Chris Hesse, Alan Hickey, Wade Hickey, Peter Hoeschele, Brandon Houghton, Kenny Hsu, Shengli Hu, Xin Hu, Joost Huizinga, Shantanu Jain, Shawn Jain, Joanne Jang, Angela Jiang, Roger Jiang, Haozhun Jin, Denny Jin, Shino Jomoto, Billie Jonn, Heewoo Jun, Tomer Kaftan, Łukasz Kaiser, Ali Kamali, Ingmar Kanitscheider, Nitish Shirish Keskar, Tabarak Khan, Logan Kilpatrick, Jong Wook Kim, Christina Kim, Yongjik Kim, Jan Hendrik Kirchner, Jamie Kiros, Matt Knight, Daniel Kokotajlo, Łukasz Kondraciuk, Andrew Kondrich, Aris Konstantinidis, Kyle Kopic, Gretchen Krueger, Vishal Kuo, Michael Lampe, Ikai Lan, Teddy Lee, Jan Leike, Jade Leung, Daniel Levy, Chak Ming Li, Rachel Lim, Molly Lin, Stephanie Lin, Mateusz Litwin, Theresa Lopez, Ryan Lowe, Patricia Lue, Anna Makanju, Kim Malfacini, Sam Manning, Todor Markov, Yaniv Markovski, Bianca Martin, Katie Mayer, Andrew Mayne, Bob McGrew, Scott Mayer McKinney, Christine McLeavey, Paul McMillan, Jake McNeil, David Medina, Aalok Mehta, Jacob Menick, Luke Metz, Andrey Mishchenko, Pamela Mishkin, Vinnie Monaco, Evan Morikawa, Daniel Mossing, Tong Mu, Mira Murati, Oleg Murk, David Mély, Ashvin Nair, Reiichiro Nakano, Rajeev Nayak, Arvind Neelakantan, Richard Ngo, Hyeonwoo Noh, Long Ouyang, Cullen O’Keefe, Jakub Pachocki, Alex Paino, Joe Palermo, Ashley Pantuliano, Giambattista Parascandolo, Joel Parish, Emy Parparita, Alex Passos, Mikhail Pavlov, Andrew Peng, Adam Perelman, Filipe de Avila Belbute Peres, Michael Petrov, Henrique Ponde de Oliveira Pinto, Michael, Pokorny, Michelle Pokrass, Vitchyr H. Pong, Tolly Powell, Alethea Power, Boris Power, Elizabeth Proehl, Raul Puri, Alec Radford, Jack Rae, Aditya Ramesh, Cameron Raymond, Francis Real, Kendra Rimbach, Carl Ross, Bob Rotsted, Henri Roussez, Nick Ryder, Mario Saltarelli, Ted Sanders, Shibani Santurkar, Girish Sastry, Heather Schmidt, David Schnurr, John Schulman, Daniel Selsam, Kyla Sheppard, Toki Sherbakov, Jessica Shieh, Sarah Shoker, Pranav Shyam, Szymon Sidor, Eric Sigler, Maddie Simens, Jordan Sitkin, Katarina Slama, Ian Sohl, Benjamin Sokolowsky, Yang Song, Natalie Staudacher, Felipe Petroski Such, Natalie Summers, Ilya Sutskever, Jie Tang, Nikolas Tezak, Madeleine B. Thompson, Phil Tillet, Amin Tootoonchian, Elizabeth Tseng, Preston Tuggle, Nick Turley, Jerry Tworek, Juan Felipe Cerón Uribe, Andrea Vallone, Arun Vijayvergiya, Chelsea Voss, Carroll Wainwright, Justin Jay Wang, Alvin Wang, Ben Wang, Jonathan Ward, Jason Wei, CJ Weinmann, Akila Welihinda, Peter Welinder, Jiayi Weng, Lilian Weng, Matt Wiethoff, Dave Willner, Clemens Winter, Samuel Wolrich, Hannah Wong, Lauren Workman, Sherwin Wu, Jeff Wu, Michael Wu, Kai Xiao, Tao Xu, Sarah Yoo, Kevin Yu, Qiming Yuan, Wojciech Zaremba, Rowan Zellers, Chong Zhang, Marvin Zhang, Shengjia Zhao, Tianhao Zheng, Juntang Zhuang, William Zhuk, and Barret Zoph. 2024. Gpt-4 technical report.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, pages 311–318.
- Ricardo Rei, José G. C. de Souza, Duarte Alves, Chrysoula Zerva, Ana C Farinha, Taisiya Glushkova, Alon Lavie, Luisa Coheur, and André F. T. Martins. 2022. COMET-22: Unbabel-IST 2022 submission for the metrics shared task. In *Proceedings of the Seventh Conference on Machine Translation (WMT)*, pages 578–585, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- Ricardo Rei, Craig Stewart, Ana C Farinha, and Alon Lavie. 2020. COMET: A neural framework for MT evaluation. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Pro-*

cessing (EMNLP), pages 2685–2702, Online. Association for Computational Linguistics.

Gemma Team, Morgane Riviere, Shreya Pathak, Pier Giuseppe Sessa, Cassidy Hardin, Surya Bhatiraju, Léonard Hussenot, Thomas Mesnard, Bobak Shahriari, Alexandre Ramé, Johan Ferret, Peter Liu, Pouya Tafti, Abe Friesen, Michelle Casbon, Sabela Ramos, Ravin Kumar, Charline Le Lan, Sammy Jerome, Anton Tsitsulin, Nino Vieillard, Piotr Stanczyk, Sertan Girgin, Nikola Momchev, Matt Hoffman, Shantanu Thakoor, Jean-Bastien Grill, Behnam Neyshabur, Olivier Bachem, Alanna Walton, Aliaksei Severyn, Alicia Parrish, Aliya Ahmad, Allen Hutchison, Alvin Abdagic, Amanda Carl, Amy Shen, Andy Brock, Andy Coenen, Anthony Laforge, Antonia Paterson, Ben Bastian, Bilal Piot, Bo Wu, Brandon Royal, Charlie Chen, Chintu Kumar, Chris Perry, Chris Welty, Christopher A. Choquette-Choo, Danila Sinopalnikov, David Weinberger, Dimple Vijaykumar, Dominika Rogozińska, Dustin Herbison, Elisa Bandy, Emma Wang, Eric Noland, Erica Moreira, Evan Senter, Evgenii Eltyshov, Francesco Visin, Gabriel Rasskin, Gary Wei, Glenn Cameron, Gus Martins, Hadi Hashemi, Hanna Klimczak-Plucińska, Harleen Batra, Harsh Dhand, Ivan Nardini, Jacinda Mein, Jack Zhou, James Svensson, Jeff Stanway, Jetha Chan, Jin Peng Zhou, Joana Carrasqueira, Joana Iljazi, Jocelyn Becker, Joe Fernandez, Joost van Amersfoort, Josh Gordon, Josh Lipschultz, Josh Newlan, Ju yeong Ji, Kareem Mohamed, Kartikeya Badola, Kat Black, Katie Millican, Keelin McDonell, Kelvin Nguyen, Kiranbir Sodhia, Kish Greene, Lars Lowe Sjoesund, Lauren Usui, Laurent Sifre, Lena Heuermann, Leticia Lago, Lilly McNealus, Livio Baldini Soares, Logan Kilpatrick, Lucas Dixon, Luciano Martins, Machel Reid, Manvinder Singh, Mark Iversen, Martin Görner, Mat Velloso, Mateo Wirth, Matt Davidow, Matt Miller, Matthew Rahtz, Matthew Watson, Meg Risdal, Mehran Kazemi, Michael Moynihan, Ming Zhang, Minsuk Kahng, Minwoo Park, Mofi Rahman, Mohit Khatwani, Natalie Dao, Nenshad Bardoliwalla, Nesh Devanathan, Neta Dumai, Nilay Chauhan, Oscar Wahltinez, Pankil Botarda, Parker Barnes, Paul Barham, Paul Michel, Pengchong Jin, Petko Georgiev, Phil Culliton, Pradeep Kuppala, Ramona Comanescu, Ramona Merhej, Reena Jana, Reza Ardeshtir Rokni, Rishabh Agarwal, Ryan Mullins, Samaneh Saadat, Sara Mc Carthy, Sarah Perrin, Sébastien M. R. Arnold, Sebastian Krause, Shengyang Dai, Shruti Garg, Shruti Sheth, Sue Ronstrom, Susan Chan, Timothy Jordan, Ting Yu, Tom Eccles, Tom Hennigan, Tomas Kocisky, Tulsee Doshi, Vihan Jain, Vikas Yadav, Vilobh Meshram, Vishal Dharmadhikari, Warren Barkley, Wei Wei, Wenming Ye, Woohyun Han, Woosuk Kwon, Xiang Xu, Zhe Shen, Zhitao Gong, Zichuan Wei, Victor Cotruta, Phoebe Kirk, Anand Rao, Minh Giang, Ludovic Peran, Tris Warkentin, Eli Collins, Joelle Barral, Zoubin Ghahramani, Raia Hadsell, D. Sculley, Jeanine Banks, Anca Dragan, Slav Petrov, Oriol Vinyals, Jeff Dean, Demis Hassabis, Koray Kavukcuoglu, Clement Farabet, Elena Buchatskaya,

Sebastian Borgeaud, Noah Fiedel, Armand Joulin, Kathleen Kenealy, Robert Dadashi, and Alek Andreev. 2024. Gemma 2: Improving open language models at a practical size.

Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023. Llama 2: Open foundation and fine-tuned chat models.