

Evaluating Large Language Models with Human Feedback: Establishing a Swedish Benchmark

Birger Moëll¹, Gustav Eje Henter¹, Jonas Beskow¹

¹KTH Royal Institute of Technology
{bmoell, ghe, beskow}@kth.se

Abstract

In the rapidly evolving field of artificial intelligence, large language models (LLMs) have demonstrated significant capabilities across numerous applications. However, the performance of these models in languages with fewer resources, such as Swedish, remains underexplored. This study introduces a comprehensive human benchmark to assess the efficacy of prominent LLMs in understanding and generating Swedish language texts using forced choice ranking. We employ a modified version of the ChatbotArena benchmark, incorporating human feedback to evaluate twelve different models, including GPT-4, GPT-4o, GPT-3.5, various Claude and Llama models, and bespoke models like Dolphin-2.9-llama3b-8b-flashback and BeagleCatMunin. These models were chosen based on their performance on LMSYS chatbot arena and the Scandeval benchmarks. We release the chatbotarena.se benchmark as a tool to improve our understanding of language model performance in Swedish with the hopes that it will be widely used. We aim to create a leaderboard once sufficient data has been collected and analysed.

1 Introduction

Large language models (LLMs) have demonstrated exceptional capabilities across various applications, significantly advancing the field of natural language processing. However, their effectiveness in low-resource languages remains underexplored. The majority of these models are optimized for English or other high-resource languages, leading to a notable performance disparity when applied to less commonly used languages. This gap in model performance has significant implications:

- **Accessibility:** Individuals who are native speakers of low-resource languages but not proficient in English are less able to benefit from the advancements in AI and machine

learning. This creates a barrier to accessing technology that could otherwise support education, business, and communication.

- **Research Limitations:** The lack of scientific investigation into LLM performance across diverse linguistic landscapes means that improvements and innovations are often not tailored to the needs of all potential users. This oversight can perpetuate inequities in technology access and effectiveness.
- **Potential model improvements:** Since models are not primarily tailored for the language, improved understanding in combination with fine-tuning techniques has the potential to significantly advance model performance for the target language.

1.1 Motivation for creating Swedish Chatbot Arena

Benchmarks for Swedish exist, with the primary ones being *Scandeval* (Nielsen, 2023) and *Superlim* (Berdicevskis et al., 2023). These tools give a good indication of model performance on specific language tasks, but the reliance on automatic benchmarking has some drawbacks, e.g. optimizing model performance with respect to the benchmarks can skew results and good performing models might not be what is preferred by humans. To address these issues we introduce the Swedish Chatbot Arena, a Swedish human preference benchmark, where the goal is to assess what models humans prefer. Humans can be seen as the gold standard for evaluating LLMs with LMSYS chatbot arena being the defacto standard for rating models reliably (Chiang et al., 2024). Our benchmark draws heavily on the LMSYS benchmark and we have [adapted the source code](#) to work for our use case. Apart from updating language on the interface and the set of models used, we also added several measures to minimize bot activity and spam.

1.2 Evaluation strategies

1.2.1 Subjective evaluation / Human feedback ranking

Human feedback plays a crucial role in the democratic evaluation of large language models (LLMs). As these models increasingly influence various sectors such as media, education, and even governance, it is essential to ensure that they operate in a manner that aligns with the values and needs of diverse user groups. By involving humans in the voting process on the quality of LLMs, we achieve several democratic objectives:

- **Representation:** Including a broad spectrum of individuals in the evaluation process ensures that the models serve the needs of a diverse population rather than a select few. This inclusivity helps prevent biases that may otherwise emerge if only a limited demographic is considered.
- **Accountability:** Human feedback mechanisms make developers and stakeholders accountable to the end-users. By actively soliciting feedback, model developers are encouraged to make improvements that are transparent and in the best interest of the public.
- **Adaptability:** Through continuous human engagement, LLMs can be dynamically adjusted to meet changing societal norms and expectations. This adaptability is crucial in maintaining the relevance and appropriateness of automated systems in a democratic society.
- **Trust:** When end-users contribute to the evaluation and enhancement of LLMs, it builds trust in the technology. Trust is essential for the wider adoption and acceptance of AI technologies in society.
- **Ethical considerations:** Human judgment is indispensable in navigating the complex ethical landscapes that AI systems, like LLMs, often encounter. Voting on model quality allows for ethical deliberations that purely data-driven approaches may overlook.

Swedish Chatbot Arena uses forced ranking between two models to create a benchmark of model performance using Elo ratings. Elo rating is a method used to calculate the relative skill levels of players in zero-sum games based on their performance against each other. This ensures a fair

evaluation of models while ensuring that a broad aspect of the public can be a part in democratically evaluating LLM performance.

1.2.2 Objective evaluation / Automatic evaluation tools

Automatic evaluation tools such as Scandeval (Nielsen, 2023) play an indispensable role in the development and refinement of large language models (LLMs). These tools provide a systematic and standardized approach to assessing model performance across multiple benchmarks, offering several distinct advantages:

- **Efficiency:** Automatic tools can evaluate models much more quickly than human-based assessments. This speed is crucial for iterating over model designs, allowing developers to refine and test models continuously without significant delays.
- **Consistency:** These tools apply the same standards and methodologies across different models, ensuring that the evaluations are consistent. This uniformity is essential for fair comparisons between models, facilitating a clearer understanding of each model's strengths and weaknesses.
- **Scalability:** Automated tools can handle large volumes of evaluation tasks simultaneously, making them scalable to the needs of rapidly developing AI technologies. This capability is particularly valuable in the context of large-scale LLMs, which require extensive testing across diverse datasets.
- **Repeatability:** Automatic evaluations can be repeated under the same conditions to verify results. This repeatability helps in ensuring the reliability of the evaluations and allows researchers to systematically track improvements over time.
- **Benchmarking:** Objective evaluation provide benchmarks that are specifically designed to test various aspects of LLM performance, such as linguistic understanding, contextual awareness, and the ability to generate coherent text. This targeted assessment helps pinpoint specific areas where a model may need further development.

The use of automated evaluation tools is crucial in the LLM development lifecycle as they provide

rapid feedback that is essential for timely and effective model training and refinement. By leveraging these tools, developers can accelerate the development process and enhance the overall quality and effectiveness of their models.

1.3 Objective and subjective evaluation

We believe that combining objective evaluation with a tool such as Scandeval with subjective evaluation with a tool such as Swedish Chatbot Arena is the most efficient way to get an accurate assessment of model performance.

2 Method

2.1 Evaluation platform

Our software evaluation platform is a fork of the LMSYS project which currently runs on a RTX-4090 GPU with enough memory to serve two open source models.

2.2 Model Selection

Twelve language models were initially selected for inclusion based on evaluation on LMSYS and Scandeval on Swedish-language tasks. Limitations such as availability in the EU / availability through an API limited model selection. A majority of the models were served through an API, while two models Dolphin-2.9-llama3b-8b-flashback and BeagleCatMunin were served through a local GPU.

The models currently included in the benchmark can be viewed in table 1.

Table 1: Language Models Selected

| Model | Model Size | Open License | LMSYS | Scandeval | Training |
|--------------------|------------|--------------|-------|-----------|--------------------|
| GPT-4o | - | No | 1 | - | - |
| GPT-4-turbo | - | No | 2 | - | - |
| GPT-4 | - | No | 2 | 1.07 | - |
| GPT-3.5 | - | No | 29 | 1.88 | - |
| Claude Opus | - | No | 2 | - | - |
| Claude Sonnet | - | No | 8 | - | - |
| Claude Haiku | - | No | 13 | - | - |
| Llama70b-Instruct | 70b | Yes | 7 | 1.82 | base + instruction |
| Llama-8b-Instruct | 8b | Yes | 17 | 2.52 | base + instruction |
| AI-Sweden-20b | 20b | No | - | 3.28 | base + fine-tune |
| Dolphin-flashback* | 8b | Yes | - | 2.36 | fine-tune + merge |
| BeagleCatMunin | 8b | Yes | - | 2.37 | fine-tune + merge |

*Full name, Dolphin-2.9-llama3b-8b-flashback

2.3 Open AI models

The models GPT-3.5 and GPT-4, GPT-4-turbo and GPT-4o were natural to include since they are commonly used. GPT-3.5 is the original ChatGPT model, GPT-4 has been a highly capable model in many evaluations and GPT-4o is ranking number one on the LMSYS Chatbot Arena (Chiang et al., 2024).

2.4 Anthropic Models

Anthropic released Anthropic version 3 on March 4th, 2024 with the models Haiku, Sonnet and Opus. Training techniques includes pretraining on large data to acquire language capabilities through word prediction, as well as human feedback techniques. (Anthropic, 2023) Anthropic used a technique called Constitutional AI that uses RLAIIF for feedback for alignment. (Bai et al., 2022).

The Opus overtook GPT-4 to become the best performing model on the LMSYS benchmark before the release of GPT-4.5. The models are highly capable and although they are not reviewed on Scandeval benchmark, they deserve to be reviewed further in this evaluation.

2.5 Llama Models

On the 18th of April 2024 Meta open sourced Llama 3. Llama 3 was trained on 15 trillion tokens collected from public sources. Over 5% of the data consists of high quality non English data. (Meta, 2023) Both the Llama 70b and the Llama 8b models are capable with high scores on both LMSYS chatbot arena and Scandeval.

2.5.1 Dolphin-2.9-llama3b-8b-flashback

Several fine-tunes and merges have been done on Llama3 models for Swedish. Dolphin-2.9-llama3b-8b-flashback is a high performing example of a llama-3 fine tune. The model is a merge of timpal0l/Llama-3-8B-flashback-v1 and cognitivecomputations/dolphin-2.9-llama3-8b which in turn are fine-tunes on Swedish language data (Flashback) and English language data. Notably the underlying dolphin model is uncensored to remove alignment and bias. This model is part of the the benchmark for research purpose but can generate unethical responses.

2.6 Open models

2.6.1 BeagleCatMunin

BeagleCatMunin is similar to Dolphin-2.9-llama3b-8b-flashback by being a merge of fine-tuned models. The model is based on the Mistral-7b architecture. The underlying Swedish model has been trained on flashback data and it is merged with a highly performant Scandinavian model, RJuro/munin-neuralbeagle-7b.

2.6.2 AI-Sweden-20B

AI-Sweden 20b is a model created by AI Sweden and trained on the Nordic Pile dataset.

3 Types of models evaluated

Since the benchmark contains both open source and closed source, not all knowledge regarding training of the models is available. However, based on information regarding the open models some conclusions regarding models can be drawn.

3.1 Instruction tuned models

Instruction tuning involves creating a dataset of instructions (prompts) and answer for a base trained model to follow. Instruction tuning has been shown to increase model performance and Llama 8b Instruct and Llama 70b Instruct are examples of instruction tuned models in this benchmark.

3.2 Reinforcement Learning from Human Feedback

Reinforcement Learning from Human Feedback is a step in training where a model uses human feedback to improve model performance. Most models included in this benchmark has at some point been trained with RLHF. Notably Dolphin-2.9-llama3b-8b-flashback is the only model where an attempt has been made to remove preferences made from RLFH training.

3.3 Merged models

Merging(Yadav et al., 2023) is a technique where models are combined to form a new model by combining the weights inside the neural network directly. This data free training method has several benefits with perhaps the biggest one being the ability to merge a model without the need for a GPU. This in combination with the high quality performance of merge models have made them a common technique for working with models.

4 Next steps

We hope that everyone will help out in our work on improving Swedish Language models by evaluating models on chatbotarena.se, the home of our benchmark.

Once evaluations are done and we have gathered enough data, we will present results of the most performant models for Swedish. We will also expand the scope of the benchmark to involve the rest of the Scandinavian languages.

4.1 How can you help?

The goal of our project is to improve access to high quality language models in Swedish. The easiest

way to help is to simply evaluate models. We also welcome pull requests and feedback on our repo¹.

References

- Anthropic. 2023. Model card for claude 3. https://www-cdn.anthropic.com/de8ba9b01c9ab7cbabf5c33b80b7bbc618857627/Model_Card_Claude_3.pdf. Accessed: 2024-05-02.
- Yuntao Bai, Saurav Kadavath, Sandipan Kundu, Amanda Askell, Jackson Kernion, Andy Jones, Anna Chen, Anna Goldie, Azalia Mirhoseini, Cameron McKinnon, Carol Chen, Catherine Olsson, Christopher Olah, Danny Hernandez, Dawn Drain, Deep Ganguli, Dustin Li, Eli Tran-Johnson, Ethan Perez, Jamie Kerr, Jared Mueller, Jeffrey Ladish, Joshua Landau, Kamal Ndousse, Kamile Lukosuite, Liane Lovitt, Michael Sellitto, Nelson Elhage, Nicholas Schiefer, Noemi Mercado, Nova DasSarma, Robert Lasenby, Robin Larson, Sam Ringer, Scott Johnston, Shauna Kravec, Sheer El Showk, Stanislav Fort, Tamera Lanham, Timothy Telleen-Lawton, Tom Conerly, Tom Henighan, Tristan Hume, Samuel R. Bowman, Zac Hatfield-Dodds, Ben Mann, Dario Amodei, Nicholas Joseph, Sam McCandlish, Tom Brown, and Jared Kaplan. 2022. *Constitutional ai: Harmlessness from ai feedback*. Preprint, arXiv:2212.08073.
- Aleksandrs Berdicevskis, Gerlof Bouma, Robin Kurtz, Felix Morger, Joey Öhman, Yvonne Adesam, Lars Borin, Dana Dannélls, Markus Forsberg, Tim Isbister, Anna Lindahl, Martin Malmsten, Faton Rekathati, Magnus Sahlgren, Elena Volodina, Love Börjeson, Simon Hengchen, and Nina Tahmasebi. 2023. *Superlim: A Swedish language understanding evaluation benchmark*. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 8137–8153, Singapore. Association for Computational Linguistics.
- Wei-Lin Chiang, Lianmin Zheng, Ying Sheng, Anastasios Nikolas Angelopoulos, Tianle Li, Dacheng Li, Hao Zhang, Banghua Zhu, Michael Jordan, Joseph E. Gonzalez, and Ion Stoica. 2024. *Chatbot arena: An open platform for evaluating llms by human preference*. Preprint, arXiv:2403.04132.
- Meta. 2023. Llama 3 release blog. <https://ai.meta.com/blog/meta-llama-3/>. Accessed: 2024-05-02.
- Dan Nielsen. 2023. *ScandEval: A benchmark for Scandinavian natural language processing*. In *Proceedings of the 24th Nordic Conference on Computational Linguistics (NoDaLiDa)*, pages 185–201, Tórshavn, Faroe Islands. University of Tartu Library.
- Prateek Yadav, Derek Tam, Leshem Choshen, Colin Raffel, and Mohit Bansal. 2023. *Ties-merging: Resolving interference when merging models*. Preprint, arXiv:2306.01708.

¹<https://github.com/BirgerMoell/SwedishLLMBenchmark>