

# Danoliteracy of Generative Large Language Models

Søren Vejlgård Holm<sup>1,2</sup>, Lars Kai Hansen<sup>1</sup>, Martin Carsten Nielsen<sup>2</sup>

<sup>1</sup>Technical University of Denmark, Anker Engelunds Vej 1, 2800 Kongens Lyngby, Denmark,

<sup>2</sup>Alvenir, Applebys Plads 7, 1411 København K, Denmark

Correspondence: swiho@dtu.dk

## Abstract

The language technology moonshot moment of Generative Large Language Models (GLLMs) was not limited to English: These models brought a surge of technological applications, investments, and hype to low-resource languages as well. However, the capabilities of these models in languages such as Danish were, until recently, difficult to verify beyond qualitative demonstrations due to a lack of applicable evaluation corpora. We present a GLLM benchmark to evaluate *Danoliteracy*, a measure of Danish language and cultural competency across eight diverse scenarios such as Danish citizenship tests and abstractive social media question answering. This limited-size benchmark was found to produce a robust ranking that correlates to human feedback at  $\rho \sim 0.8$  with GPT-4 and Claude Opus models achieving the highest rankings. Analyzing these model results across scenarios, we find one strong underlying factor explaining 95% of scenario performance variance for GLLMs in Danish, suggesting a  $g$  factor of model consistency in language adaptation.

## 1 Introduction

Benchmarks shape technologies. By acting as normative guidelines for technology applications, benchmarks imply directions of research and development that ultimately impact users (Liang et al., 2022). GLLMs specifically have emerged as a technology with near-universal impact, including lower-resource languages such as Danish (Olsen, 2023). If the challenging task of general GLLM evaluation is not extended to low-resource languages, practitioners start from scratch for each model use case, inhibiting practical adoption or possibly resulting in risky, undertested implementations.

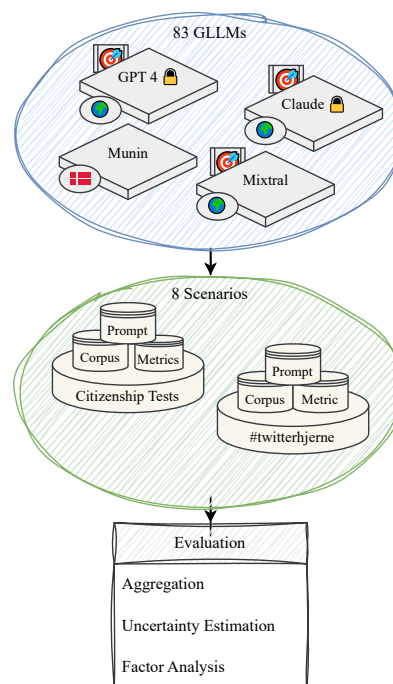


Figure 1: The overall evaluation setup: A collection of GLLMs, including closed-source (lock symbol) instruct-tuned (bulls-eye) and multilingual (globe) ones, were evaluated in Danish across diverse use-case scenarios.

We take up the challenge of creating a GLLM evaluation benchmark for Danish, a North Germanic language spoken by 6 million people, primarily in the Nordic country of Denmark. As depicted in Figure 1, our approach is to create a compilation of small-scale, diverse evaluation scenarios combined into one general benchmark to reveal model Danoliteracy. By Danoliteracy we refer to the level of GLLM real-world knowledge (RWK), natural language understanding (NLU), and natural language generation (NLG) in Danish.

This paper presents the resulting *Danoliterate Benchmark*, describing evaluation methods, datasets and model results. We analyze these results with the goals of validating and exploring evaluation methodology. An important part of this analysis is to investigate the feasibility of such evaluation: Does this small-scale, language-specific approach achieve a significant ranking of GLLMs? Even if a non-spurious leaderboard can be discerned from the result, it is not enough to validate the benchmark which might actually show something orthogonal to Danoliteracy. Thus, as a benchmark validation tool, we additionally present a user survey, collecting the preferences of Danish speakers when interacting with hidden pairs of GLLMs in an arena A/B test setup.

The availability of a suite of meaningful benchmark results allows us to investigate GLLM behavior: Initially, we explore which specific models are most Danoliterate and how different types of GLLMs compare. Beyond that, we are particularly interested in capability consistency across tasks: If a GLLM performs strongly in one Danish use-case scenario, does this performance generalize to other Danish scenarios across different domains and objectives?

We hope so. If the answer is no, practitioners are without general results to trust, requiring a full model re-evaluation for each downstream use. However, if capability consistency is present, we should be able to find a single underlying axis that correlates with performance across diverse scenarios. Such a general dimension of Danoliteracy can be compared to the  $g$  factor of general human intelligence (Spearman, 1904). If one significant, main factor is found, it implies a level of stability that can help guide the expectations of practitioners across GLLM implementations in varying and even novel Danish tasks.

The contributions presented in this paper can be summarized as follows:

- An open-source benchmark for GLLMs in low-resource languages with an evaluation framework and a live leaderboard site.
- The release of a set of novel evaluation datasets for Danish.
- Evidence that GPT-4 and Claude Opus models are currently uniquely capable in Danish, outperforming other closed models which in turn overcome open-weights models.

- Evidence suggesting the existence of a Danoliteracy  $g$  factor in GLLMs supported by preliminary results from our open-source human feedback study.

## 2 Related Work

### 2.1 GLLM Evaluation

The hard task of evaluating free-generation, multitask models has been attempted in many ways. Liang et al. define an empirical approach for revealing model behaviour: Evaluate each model on a compilation of many scenarios and use-cases of interest, spanning different languages, domains and task categories – ideally across multiple performance dimensions in addition to raw model capability such as efficiency, bias, and robustness (Liang et al., 2022).

This scenario compilation approach has been applied in many ways to GLLMs: The HELM Lite benchmark presents evaluations of GLLMs on question answering (QA) and translation tasks (Liang et al., 2023). Influential benchmarks include the Huggingface OpenLLM Leaderboard (Beeching et al., 2023) and other implementations of the knowledge-based scenarios MMLU (Hendrycks et al., 2021) and HellaSwag (Zellers et al., 2019).

These benchmarks mainly use comparison or similarity algorithms to parse model answers e.g. for finding a chosen option for multiple-choice QA. Other approaches include applying other GLLMs to grade generations (Zheng et al., 2023) (OpenAI, 2023) or using human feedback (Chiang et al., 2024).

### 2.2 Low-resource NLP Evaluation

Most broadly reported GLLM evaluations are only or primarily performed on examples in English. Approaches to evaluate lesser-resourced languages include both attempts to compile massively multilingual benchmarks either by automatic translation (Lai et al., 2023) or dataset curation (Ahuja et al., 2023).

Other approaches focus on one language exclusively in attempts to evaluate GLLM language performance beyond surface-level lexical or syntactical literacy. Using this method, practitioners can align scenario domain, cultural content, and real-world facts with the setting of the language, though a lack of relevant data can be problematic (Liu et al., 2023).

Specifically in Danish, the comprehensive Scan-dEval benchmark, which packages scenarios across eight languages divided into NLU and NLG leaderboards, implements evaluation on GLLMs in Danish on eight NLG scenarios with some overlap in dataset sources with this work (Nielsen, 2023).

### 2.3 GLLM $g$ Factor

The idea that GLLM performance is strongly correlated across tasks has been noted previously by for example Ilić who carried out factor analysis on the Open LLM Leaderboard and GLUE, (Wang et al., 2018) obtaining results similar to ours (Ilić, 2023).

## 3 Methods

### 3.1 Datasets

The eight scenario datasets are divided into three broad categories: Scenarios testing RWK, scenarios requiring models to perform free NLG and those that imply solving classical NLU tasks.

#### Real-world Knowledge

1. **Citizenship Test** is a novel dataset of 605 multiple-choice questions acquired from governmental tests that require applicants for Danish citizenship to demonstrate familiarity with national societal structure, culture, and history (siri.dk, 2023).
2. **HyggeSwag** is a novel manual translation<sup>1</sup> of 125 HellaSwag (Zellers et al., 2019) ActivityNet (Caba Heilbron et al., 2015) questions testing commonsense natural language inference as a multiple-choice task to pick the only completion consistent with real-world logic.
3. **Gym 2000** is a small, novel extraction of 50 literature comprehension multiple-choice questions from the Danish Centre for Reading Research (CRR) aimed at high-schoolers (Arnbak and Elbro, 2000).

#### Free NLG

- 4 **#twitterhjerne** is a novel abstractive question-answering dataset containing 78 anonymized question tweets from the Danish hashtag of that name, translated to *Twitter Brain*, where users ask the social media hive mind for help, input or recommendations. For each question

<sup>1</sup>The text was translated by the authors with each translation being validated by another author completing the inference task.




Figure 2: Selected model normalized results across the eight scenarios divided into three categories as described in Section 3.1. Claude Opus is overtaken by GPT-4 on the NER task but wins on an NLG task. LLaMa 3 70B, the SOTA open-weights model, lags behind on NLU and knowledge-based tasks. A Danish-specialized model with only 1.1B parameters, DanskGPT-tiny Chat, benchmarks well in NLG but fails on knowledge and understanding.

tweet, 2-9 reference answer tweets were extracted making it possible to use the  $\text{score}_{o_{10}}$  metric (1).

- 5 **Nordjylland News** is an existing news summarization dataset (Kinch, 2023) from which a subset of 300 short news articles with corresponding summaries were used.

#### NLU Tasks

- 6 **Cloze Self Test** is another small, novel extraction from CRR materials (Jensen et al., 2015), this one containing 33 cloze-style questions evaluated as multiple-choice selection.
- 7 **DaNE** is an existing canonical Danish NER dataset with four entity categories (Hvingelby et al., 2020) from which a subset of 256 examples were used.
- 8 **Angry Tweets** is an existing sentiment classification dataset (Brogaard Pauli et al., 2021, Sec. 4) with three sentiment categories from which 256 examples were used for multiple-choice prompts.




Figure 3: The non-normalized metric scores across evaluation scenarios for two models that were judged highly according to human feedback. Uncertainties are 95% confidence intervals according to the bootstrapping procedure and the micro-average is displayed for each model.

All datasets are released on the Huggingface Datasets Hub with dataset cards<sup>2</sup> except for the two small datasets extracted from CRR which require practitioners to re-run data collection for personal use. More details on dataset licensing and collection as well as data examples can be found in Appendix B.

### 3.2 Evaluation

Each evaluation scenario consisted of a dataset, a prompt template, and a chosen metric.

Most of the available datasets allowed primarily for testing discriminative RWK and NLU of the GLLM by requiring it to select between multiple-choice answers. For these multiple-choice scenarios, frequency of generating the correct option number was reported as model accuracy.

Two metrics were used for NLG. First, summarization was implemented using a similarity score between model summary and a reference summary  $s(\mathcal{T}_{\mathcal{M}}, \mathcal{T}_{\text{ref}})$ . Secondly, we implemented abstractive question answering tasks for the specific type of dataset  $D$  where each question has not just one correct answer but a corresponding set of reference, human-generated answers. This was done by scoring GLLMs using the frequency with which generated answers were the odd-one-out, defined by the lowest total similarity to all possible answers

<sup>2</sup>Datasets can be found on danoliterate.compute.dtu.dk/Scenarios

$T = \{\mathcal{T}_{\mathcal{M}}\} \cup \{\mathcal{T}_{\text{ref}, i}\}_{i=1..k}$  as shown in Eq. 1. For similarity scores  $s$ , the BERT score algorithm (Zhang\* et al., 2020) based on the DFM Encoder Large (Enevoldsen et al., 2022) was used.

$$\text{score}_{\text{o1o}} = \mathbb{P}_D \left[ \mathcal{T}_{\mathcal{M}} = \underset{t_1 \in T}{\text{argmin}} \sum_{t_2 \in T} s(t_1, t_2) \right] \quad (1)$$

Finally, few-shot named entity recognition (NER) was implemented for GLLMs using 3-shot prompting and the GPT-NER multiple queries idea (Wang et al., 2023). Here, word-level entity class predictions were aligned and the standard NER micro-average F1 scores were calculated using the SeqEval framework (Nakayama, 2018).

Scenarios were operationalized by prompting GLLMs in the scenario language, Danish, and structuring prompts with headers marked with the # character as in Markdown. In order to use the same prompts for instruct-tuned and base GLLMs, prompts started with the instruction and ended with a text leading towards an answer in the continuation as shown in Figure 4. Prompting and metric implementation details are covered in Appendix A.

- 1 Write a one-sentence summary of the text.
- 2 # TEXT
- 3 Lorem ipsum dolor sit amet ...
- 4 # SUMMARY
- 5 A summary of the text could be:

Figure 4: The general prompting approach translated to English.

### 3.3 Models

Both local prediction of open-weights models and API access to externally hosted GLLMs were implemented. 54 autoregressive, decoder language models trained for general text generation were included. Models were tested if we saw any reason to suspect a degree of Danoliteracy, thus including multilingual models with possibly small amounts of Danish training data as well as other Mainland Scandinavian monolingual models but excluding strictly English-only models. Both base and instruct-tuned models were evaluated. All model generation was performed with greedy decoding and with a maximum number of generated model tokens of 256. OpenAI’s o1 model was allowed to generate internal tokens freely. The models run

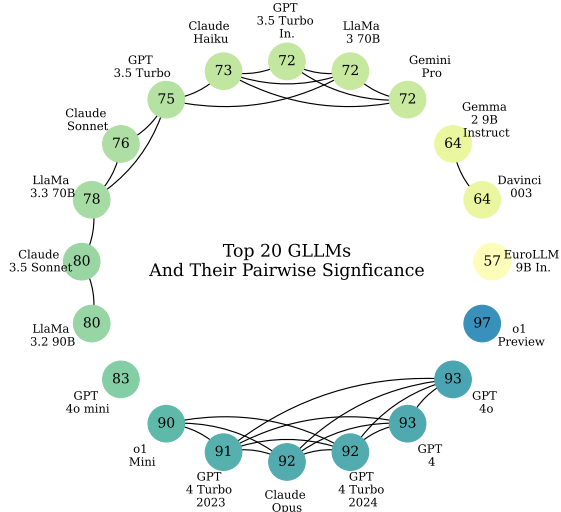


Figure 5: Model Danoliteracy Index across all scenarios for top performers. Two model nodes are connected iff the bootstrapping procedure could not reveal significant benchmark performance difference at  $\alpha = 0.05$ . Together with the special o1 model, Claude Opus and the GPT 4 family models are consistent winners.

locally ranged in sizes from 124M parameters to 13B parameters, resulting in a total project GPU use of  $\sim 100$  hours on a single Nvidia H100.

### 3.4 The Danoliterate Framework

A modular, open-source evaluation framework was implemented in Python, using Huggingface Transformers (Wolf et al., 2019) and Datasets (Lhoest et al., 2021) as central tools as well as Hydra (Yadan, 2019) and a Weights and Biases-integration (Biewald, 2020) for structuring experimentation. This framework, danoliterate, is released on GitHub<sup>3</sup> under the MIT License.

Furthermore, an interactive site displaying the leaderboard as well as other benchmark results and examples was produced using the Streamlit framework. See Figure 6 for a screenshot of this frontend and Appendix A.4 for versions of software dependencies.

### 3.5 Human Feedback

For a subset of 18 instruct-tuned models, we have set up a parallel study to collect human judgment on model performance. Volunteers were presented with a anonymized pair of models and were asked

<sup>3</sup>github.com/sorenmulli/danoliterate

Model	Check 1	Check 2	Check 3	Score	Citizenship Test
OpenAI Davinci 003	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>		65	69±2
Mixtral (@ Groq)	<input checked="" type="checkbox"/>	<input type="checkbox"/>	46.7	52	61±2
SOLAR 10.7B Instruct	<input checked="" type="checkbox"/>	<input type="checkbox"/>	10.7	51	66±2
Qwen1.5 7B chat	<input checked="" type="checkbox"/>	<input type="checkbox"/>	7.7	50	55±2
Starling-LM-7B-beta	<input checked="" type="checkbox"/>	<input type="checkbox"/>	7.2	43	45±2
Heidrun Mistral 7B Chat	<input checked="" type="checkbox"/>	<input type="checkbox"/>	7.2	42	54±2
LLaMa 2 (@ Groq)	<input checked="" type="checkbox"/>	<input type="checkbox"/>	69.0	41	54±2

Figure 6: A screenshot from the leaderboard frontend allowing users to explore how model results change with different metric choices as well as inspecting model output examples and reading further details on evaluation scenarios.

to report their preferred model. This was done based on side-by-side model answers on at least three prompts selected by the volunteer from a pool of 100 prompt examples. Prompt selection was chosen independently of the Danoliterate Benchmark by creating one Danish prompt for each of 100 popular generative AI use-cases according to Zao-Sanders (Zao-Sanders, 2024). The study is ongoing: At the time of writing, 477 responses were analyzed. More details on data collection and analysis can be seen in Appendix C.

## 4 Results

### 4.1 Benchmark Feasibility

Benchmarks must have a sufficiently clear signal to be useful. The ranking in the final leaderboard should be determined by meaningful model differences and influenced minimally by sampling noise.

To quantify benchmark noise, we implemented blocked bootstrapping, resampling all examples with replacement and aggregating all  $N = 8$  scenario scores for each of the  $M = 83$  models. For each  $K = 10,000$  bootstrap samples, the  $M \times N$  model scenario results were aggregated into one overall Danoliteracy Index for each model  $d_M$ . This index was computed by considering one scenario at a time, assigning to the winner index 100 and to the lowest-scoring model index 0 with a linear scaling between the two. The micro-average across the  $N$  scenarios is reported as the resulting Danoliteracy Index of the sample.

The median index is presented as the main leaderboard of this report<sup>4</sup> with top 20 shown in

<sup>4</sup>The full  $M \times N$  version can be seen at the live leaderboard site: danoliterate.compute.dtu.dk/Leaderboard.






Figure 7: Model Danish capabilities based on human feedback. Values are normalized Bradley-Terry coefficients where two models are connected if the coefficients are not significantly different in the ranking model described in Appendix C.3.

Figure 5. For these, pairwise model index comparisons were performed using the bootstrap samples, correcting  $p$  values to control the false discovery rate across  $\frac{1}{2}M \times (M - 1)$  comparisons (Benjamini and Hochberg, 1995), presenting significant differences at  $\alpha = 5\%$ .

The results show groups of similarly performing models whose Danoliteracy cannot be distinguished. This increases the lower you go: Many mediocre models, especially non instruct-tuned models, get  $d_M \sim 20$ : As an example, this small-sample size, curated benchmark cannot reveal a difference between the base models LLaMa 2 7B,  $d_{L2} = 20$ , and LLaMa 3 8B,  $d_{L3} = 23$ .

However, robust separation is visible for some models, providing basis for statements like "The different GPT-4 models benchmark at the same level but clearly perform better than GPT-3 models" or "The bigger the Claude 3, the better the performance – but even the cheap Haiku version performs at GPT-3.5 level" or "Small, Danish-specialized models like Heidrun can perform at LLaMa 2 70B level but LLaMa 3 has moved the SOTA for open-weights models in Danish". This signal allows us to learn more about reasons for model performance which we explore in the next section.

First, we turn to the important question of validity: We see a robust benchmark signal resulting in

a significant ranking but must question the meaning of the signal. One superficial indication of a meaningful signal is that, as expected, the ranking correlates significantly with model parameter counts<sup>5</sup>  $\rho \sim 0.6$ . However, more importantly: We find that it does correlate with the preliminary results of our Danish human judgement survey.

Ranking human judgement using the Bradley-Terry model as in (Chiang et al., 2024, Sec. 4), we achieve a ranking shown in Figure 7. We observe meaningful differences compared to the Danoliteracy Index: For example, Claude models are more competitive against GPT 4 and the title as best included open-weights model is taken by Nexusflow Starling (Zhu et al., 2023) from LLaMa 3 70B. Crucially, however, the general ranking is similar, resulting in a correlation<sup>6</sup> of  $\rho \sim 0.8$  with the Danoliteracy Index for these 18 judged models. We note this as a high value. As a comparison, the Danoliteracy index has a weaker correlation with English benchmarks like HELM Lite and the Open LLM Leaderboard<sup>7</sup>,  $\rho \sim 0.5$ .

Thus, the results from our monolingual scenario compilation approach differ from those from English benchmarks while importantly, showing high correspondence to judgments made by Danish speakers.

## 4.2 Model Outcomes

The leading models are familiar, proprietary top products. Though LLaMa 3 reaches GPT 3.5 level, Figure 8 shows that most models capable in Danish do not have openly available weights. Furthermore, these top performers are generally also large and instruct-tuned: Quantitatively, models get about  $\sim 0.5$  further Danoliteracy Index points per additional billion parameters and around  $\sim 15$  from instruct-tuning<sup>8</sup>.

The substantial requirements of dataset and model scale as well as creation of instruct datasets might explain why nationally anchored organizations have not been able to come up with Danish-first models competing with the multilingual behemoths. Such multilingual models have the advan-

<sup>5</sup>Here, only 38 open models with known parameter counts were considered.  $\alpha = 5\%$  confidence interval: [0.37; 0.78].

<sup>6</sup> $\alpha = 5\%$  confidence interval: [0.6; 0.9]

<sup>7</sup>This is based on 12 models that overlap between this benchmark and HELM Lite ( $\alpha = 5\%$  confidence interval: [-0.2; 0.8]) and the 15 that overlap with the Open LLM Leaderboard ( $\alpha = 5\%$  confidence interval: [0.; 0.8]).

<sup>8</sup>From fitting a naïve linear model on the results including only models with known parameter counts, see Appendix D.1




Figure 8: Model Danoliteracy Index for groups of models. For the models tested in Danish, closed weights dominate open-weights in a remarkably clear way. Bigger models are better and on this benchmark, instruct-tuning is necessary to achieve high benchmark scores.

tage of linguistic and factual knowledge enhancement across training languages which also benefit them in the monolingual setting.

### 4.3 Capability Dimensionality

The previous analysis primarily considered the aggregated benchmark results across scenarios. What is going on at the scenario level? While different model capability profiles can be seen, exemplified in Figure 2 and Figure 3, the main first impression is that model performance at one benchmark scenario strongly predicts performance at other scenarios: One principal component explains 75% of the model result variance across the eight scenarios.

This finding leads us to the conclusion that a "general factor of Danoliteracy" exists. We investigate this further using Exploratory Factor Analysis (EFA) on the  $M \times N$  scenario result matrix, analyzing the underlying result dimensionality: How many factors are needed to explain the variance induced by model results over the  $N$  scenarios?

This analysis, further detailed in Appendix D.2, shows a sharp drop in factor eigenvalue when moving from one to two factors as shown in Figure 9.




Figure 9: Factor Analysis on model results across eight scenarios reveal one underlying dimension of Danoliteracy deemed significant by Horn's Parallel Analysis.




Figure 10: The scenarios most contributing to the underlying signal of Danoliteracy: The factual multiple-choice evaluation scenario containing Danish citizenship tests most strongly explains benchmark performance.

According to both Horn's Parallel Analysis (Horn, 1965) and the Kaiser Criterion requiring relevant factors to have  $\lambda > 1$  (Kaiser, 1960), the resulting number of significant factors is 1. Loadings for this factor is shown in Figure 10 suggesting that the RWK scenarios of HyggeSwag and, importantly, the Danish culturally aligned Citizenship Test scenario explain the largest part of the dynamics of the model results.

GLLM capability being consistent across different tasks is not just suggested by this benchmark: Carrying out the same analysis for Scandinavian and English scenario compilations shows an underlying benchmark dimensionality much lower than scenario count, with significant factor count close to 1 as seen in Table 1.

Benchmark	D	$\sigma_{F1}$	$\sigma_{F2}$	$F_K$	$F_{PA}$
Danoliterate	$83 \times 8$	93%	7%	1	1
ScandEval Da.	$199 \times 8$	88%	12%	1	2
ScandEval Full	$199 \times 24$	77%	16%	2	2
HELM Lite	$90 \times 10$	78%	19%	2	2
OpenLLM	$2859 \times 6$	97%	3%	1	2

Table 1: How much variance did the first and second factors in EFA explain for the Danoliterate Benchmark as well as the ScandEval benchmark, both full and Danish subset (Nielsen, 2023), English benchmarks HELM Lite (Liang et al., 2022) and OpenLLM (Beeching et al., 2023) as of January 2025. Leaderboard dimensionality, model count  $\times$  scenario count, is presented along with suggested significant factor count by the Kaiser criterion and Parallel Analysis: All benchmarks have an important first component.

## 5 Conclusions

Based on the ability to robustly discover model groupings at different Danish capability levels and correlate these rankings with human feedback, we conclude that a scenario compilation approach can meaningfully reveal GLLM capabilities. We show that, in Danish, open-weights GLLMs currently lag behind large, closed, multilingual, instruct-tuned models, such as GPT-4 and Claude Opus.

For our evaluation setup, we observe one underlying factor in model capability across the diverse test scenarios. This observation is supported by similar structures in other Danish, English and multilingual scenario compilations which we consider a positive result for low-resource evaluation: By using curated and language-specific scenarios, the general landscape of GLLM capabilities for a given low-resource language can be meaningfully inferred even if resources limit the scale.

## 6 Concerns of the Ethical Impacts

This work releases a benchmark and leaderboard with the hope of a positive outcome of increased understanding of potentials and limitations of GLLMs in Danish. However, we note some risks in the use of such leaderboards.

The results presented here only focus on model capability but, on the leaderboard site, versions of other important dimensions for model applicability are presented; such as model efficiency, model likelihood calibration and model generated output toxicity. However, these are presented with a disclaimer as preliminary results and our work on

other crucial dimensions such as GLLM performance fairness across gender and nationalities or robustness to input noise have not been released due to limitations to current datasets to robustly carry out these analyses.

There is an increased risk of bias, fairness and toxicity violations in low-resource languages to which models are less tuned. Problematically, when the evaluation situation is also low-resource, these risks might be undiscovered for practitioners that only focus on a model capability. Further work is crucially needed but for now, the leaderboard site displays a disclaimer against blindly trusting that high benchmark numbers mean predictable downstream performance or applying GLLMs with unchecked assumptions about robustness, fairness, bias, and toxicity.

## 7 Limitations

The study only focuses on one language, Danish, with limited comparisons to other language results.

The presented benchmark consists of eight specific scenarios: Although we find high correlation between scenario results, all our statements about model performance on Danish in general are evidently biased by the scenario selection. A similar statement can be made about prompt and metric design decisions though these seem robust in ablation studies in Appendix A.

We stress the importance of the uncertainty quantification for this benchmark where all scenarios are small-scale,  $n < 1,000$ : The bootstrap analysis revealed some model result differences, such as Mixtral 8x7B ( $d_M = 54$ ) and Qwen1.5 7B Chat ( $d_M = 50$ ), are not significant at the desired level and might be spurious. Other differences such as that between GPT-4 and Claude Opus might be obscured by the important Citizenship Test scenario, where these models achieve close to 100% accuracy (Figure 3), being saturated by SOTA GLLMs. Though most other scenarios still show far-from-perfect accuracy, more difficult scenarios are needed to accommodate future developments.

As all evaluation data is publicly available, unintentional or malignant dataset contamination is possible. This issue requires attention but might, in the short-term, be less of a risk for low-resource language evaluation with smaller and less widely published corpora.



## Acknowledgments

We would like to thank all the anonymous reviewers for the insightful and helpful comments. This work was supported by the Pioneer Centre for AI, DNRF grant number P1.

## References

- Kabir Ahuja, Harshita Diddee, Rishav Hada, Millicent Ochieng, Krithika Ramesh, Prachi Jain, Akshay Nambi, Tanuja Ganu, Sameer Segal, Mohamed Ahmed, Kalika Bali, and Sunayana Sitaram. 2023. MEGA: Multilingual evaluation of generative AI. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 4232–4267, Singapore. Association for Computational Linguistics.
- Elisabeth Arnbak and Carsten Elbro. 2000. Læsetekster for gymnasium, hf mv. Uddannelsesstyrelsens Internetpublikationer. Adgang 2000.
- Maurice S. Bartlett. 1951. The effect of standardization on a  $\chi^2$  approximation in factor analysis. *Biometrika*, 38(3/4):337–344.
- Edward Beeching, Clémentine Fourier, Nathan Habib, Sheon Han, Nathan Lambert, Nazneen Rajani, Omar Sanseviero, Lewis Tunstall, and Thomas Wolf. 2023. Open llm leaderboard. [https://huggingface.co/spaces/HuggingFaceH4/open\\_llm\\_leaderboard](https://huggingface.co/spaces/HuggingFaceH4/open_llm_leaderboard).
- Yoav Benjamini and Yosef Hochberg. 1995. Controlling the false discovery rate: A practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society: Series B (Methodological)*, 57(1):289–300.
- Lukas Biewald. 2020. Experiment tracking with weights and biases. Software available from wandb.com.
- Amalie Brogaard Pauli, Maria Barrett, Ophélie Lacroix, and Rasmus Hvingelby. 2021. DaNLP: An open-source toolkit for danish natural language processing. In *Proceedings of the 23rd Nordic Conference on Computational Linguistics (NoDaLiDa 2021)*.
- Fabian Caba Heilbron, Victor Escorcia, Bernard Ghanem, and Juan Carlos Niebles. 2015. Activitynet: A large-scale video benchmark for human activity understanding. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 961–970.
- 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.
- Kenneth Enevoldsen et al. 2022. dfm-encoder-large-v1. A Transformer encoder model, part of the BERT family, intended for Danish natural language tasks.
- Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. 2021. Measuring massive multitask language understanding. In *Proceedings of the 9th International Conference on Learning Representations (ICLR)*.
- Matthew Honnibal, Ines Montani, Sofie Van Landeghem, and Adriane Boyd. 2020. spaCy: Industrial-strength Natural Language Processing in Python. Software available from spacy.io.
- John L. Horn. 1965. A rationale and test for the number of factors in factor analysis. *Psychometrika*, 30(2):179–185.
- Rasmus Hvingelby, Amalie Brogaard Pauli, Maria Barrett, Christina Rosted, Lasse Malm Lidegaard, and Anders Sjøgaard. 2020. DaNE: A named entity resource for Danish. In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 4597–4604, Marseille, France. European Language Resources Association.
- David Ilić. 2023. Unveiling the general intelligence factor in language models: A psychometric approach.
- Katrine Lyskov Jensen, Anna Steenberg Gellert, and Carsten Elbro. 2015. Rapport om udvikling og afprøvning af selvtest af læsning – en selvtest af voksnes læsefærdigheder på nettet.
- Henry F. Kaiser. 1960. The application of electronic computers to factor analysis. *Educational and Psychological Measurement*, 20(1):141–151.
- Henry F. Kaiser. 1970. A second generation little jiffy. *Psychometrika*, 35:401–415.
- Oliver Kinch. 2023. Nordjylland news summarization. Hugging Face Datasets.
- Viet Lai, Chien Nguyen, Nghia Ngo, Thuat Nguyen, Franck Dernoncourt, Ryan Rossi, and Thien Nguyen. 2023. Okapi: Instruction-tuned large language models in multiple languages with reinforcement learning from human feedback. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 318–327, Singapore. Association for Computational Linguistics.
- Vladimir I. Levenshtein. 1966. Binary codes capable of correcting deletions, insertions, and reversals. *Soviet Physics Doklady*, 10(8):707–710.
- Quentin Lhoest, Albert Villanova del Moral, Yacine Jernite, Abhishek Thakur, Patrick von Platen, Suraj Patil, Julien Chaumond, Mariama Drame, Julien Plu, Lewis Tunstall, Joe Davison, Mario Šaško, Guntjan Chhablani, Bhavitvya Malik, Simon Brandeis, Teven Le Scao, Victor Sanh, Canwen Xu, Nicolas

- Patry, Angelina McMillan-Major, Philipp Schmid, Sylvain Gugger, Clément Delangue, Théo Matussière, Lysandre Debut, Stas Bekman, Pierric Cistac, Thibault Goehringer, Victor Mustar, François Lagunas, Alexander Rush, and Thomas Wolf. 2021. Datasets: A community library for natural language processing. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 175–184, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Percy Liang, Rishi Bommasani, Tony Lee, Dimitris Tsipras, Dilara Soylu, Michihiro Yasunaga, Yian Zhang, Deepak Narayanan, Yuhuai Wu, Ananya Kumar, Benjamin Newman, Binhang Yuan, Bobby Yan, Ce Zhang, Christian Cosgrove, Christopher D. Manning, Christopher Ré, Diana Acosta-Navas, Drew A. Hudson, Eric Zelikman, Esin Durmus, Faisal Ladhak, Frieda Rong, Hongyu Ren, Huaxiu Yao, Jue Wang, Keshav Santhanam, Laurel Orr, Lucia Zheng, Mert Yuksekgonul, Mirac Suzgun, Nathan Kim, Neel Guha, Niladri Chatterji, Omar Khattab, Peter Henderson, Qian Huang, Ryan Chi, Sang Michael Xie, Shibani Santurkar, Surya Ganguli, Tatsunori Hashimoto, Thomas Icard, Tianyi Zhang, Vishrav Chaudhary, William Wang, Xuechen Li, Yifan Mai, Yuhui Zhang, and Yuta Koreeda. 2022. Holistic evaluation of language models.
- Percy Liang, Yifan Mai, Josselin Somerville, Farzaan Kaiyom, Tony Lee, and Rishi Bommasani. 2023. HELM Lite: Lightweight and broad capabilities evaluation.
- Chin-Yew Lin. 2004. ROUGE: A package for automatic evaluation of summaries. In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.
- Peng Liu, Lemei Zhang, Terje Nissen Farup, Even W. Lauvrak, Jon Espen Ingvaldsen, Simen Eide, Jon Atle Gulla, and Zhirong Yang. 2023. Nlebench+norglm: A comprehensive empirical analysis and benchmark dataset for generative language models in norwegian.
- Samuel R. Mathias. 2024. Horns: Horn’s parallel analysis in python. <https://github.com/sammosummo/Horns>.
- Hiroki Nakayama. 2018. seqeval: A python framework for sequence labeling evaluation. Software available from <https://github.com/chakki-works/seqeval>.
- 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.
- Jesper Olsen. 2023. Hvor taler du flot dansk! selv skaberne er forbløffede over chatbottens sprogøre.
- OpenAI. 2023. evals. <https://github.com/openai/evals>. Accessed: 2023-11-22.
- siri.dk. 2023. Danskundervisning og prøver for udlændinge.
- Charles Spearman. 1904. "General Intelligence," Objectively Determined and Measured. *The American Journal of Psychology*, 15(2):201–292.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. 2018. GLUE: A multi-task benchmark and analysis platform for natural language understanding. In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 353–355, Brussels, Belgium. Association for Computational Linguistics.
- Shuhe Wang, Xiaofei Sun, Xiaoya Li, Rongbin Ouyang, Fei Wu, Tianwei Zhang, Jiwei Li, and Guoyin Wang. 2023. Gpt-ner: Named entity recognition via large language models.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2019. Transformers: State-of-the-art natural language processing. <https://github.com/huggingface/transformers>. Hugging Face, Brooklyn, USA.
- Omry Yadan. 2019. Hydra - a framework for elegantly configuring complex applications. Github.
- Marc Zao-Sanders. 2024. How people are really using genai. *Harvard Business Review*. Technology and analytics.
- Rowan Zellers, Ari Holtzman, Yonatan Bisk, Ali Farhadi, and Yejin Choi. 2019. HellaSwag: Can a machine really finish your sentence? In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4791–4800, Florence, Italy. Association for Computational Linguistics.
- Tianyi Zhang\*, Varsha Kishore\*, Felix Wu\*, Kilian Q. Weinberger, and Yoav Artzi. 2020. Bertscore: Evaluating text generation with bert. In *International Conference on Learning Representations*.
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric. P Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. 2023. Judging llm-as-a-judge with mt-bench and chatbot arena.
- Banghua Zhu, Evan Frick, Tianhao Wu, Hanlin Zhu, Karthik Ganesan, Wei-Lin Chiang, Jian Zhang, and Jiantao Jiao. 2023. Starling-7b: Improving llm helpfulness and harmlessness with rlaiif.

## A Evaluation Methodology

### A.1 Prompting

An example of a prompt following the structure shown in Figure 4 is the Citizenship Test example shown in Figure 11.

```

1 Svar kun med tallet for den rigtige
  mulighed.
2 # SPØRGSMÅL
3 Hvilket af følgende lande har flest
  indvandrere og efterkommere i
  Danmark oprindelse i?
4 # SVARMULIGHEDER
5 1. Pakistan
6 2. Iran
7 3. Tyrkiet
8 # SVAR
9 Svaret er mulighed nummer

```

Figure 11: Prompting of an example from the Citizenship Test scenario. Translated, this question prompt reads: *Answer only with the number corresponding to the correct answer. # QUESTION From which of the following countries do the highest number of immigrants and descendants in Denmark have their roots? # OPTIONS 1. Pakistan 2. Iran 3. Turkey # ANSWER The answer is option number.*

**Alternatives** Some prompt alternatives to the approach shown in Figures 4 and 11 were run for a subset of GLLMs: For the Citizenship Test scenario, Table 2, results were similar when not presenting options to the models, instead parsing their output choice by selecting the option with highest similarity to their generation. The same table suggests that changing the Citizenship Test scenario to a simpler prompt without markdown headers lowered results minimally..

For the summarization task, Nordjylland News, an alternative prompt with longer and more detailed instructions had no effect on model results shown in Table 3.

Translating all instruction text in the prompt format while keeping data content in Danish maintained or improved non-Danish model results, Table 4 on the Gym 2000 scenario.

### A.2 GLLM Output Parsing

**Multiple-choice** A numbered option was considered selected if it was the only number generated by the model. In the case of multiple generated option numbers, the most frequent number was

	Std.	Simple Q	No opt.
Gemini Pro	$85 \pm 1$	$78 \pm 1$	$79 \pm 1$
GPT 3.5 Turbo	$82 \pm 1$	$77 \pm 1$	$82 \pm 1$
Mistral 7B Instruct	$47 \pm 2$	$50 \pm 2$	$49 \pm 2$
Mistral 7B	$45 \pm 2$	$44 \pm 2$	$41 \pm 2$
Dano. Mistral 7B	$43 \pm 2$	$45 \pm 2$	$59 \pm 2$
LlaMa 2 7B	$39 \pm 2$	$42 \pm 2$	$36 \pm 2$
Dano. LlaMa 2 7B	$37 \pm 2$	$40 \pm 2$	
Dummy Baseline	$36 \pm 2$	$36 \pm 2$	$36 \pm 2$

Table 2: Alternative prompting and scoring approaches to the Citizenship Test run for a subset of models including a baseline outputting a fixed, random string and Danish-tuned versions of base GLLMs. Std. is the prompt version presented in the benchmark, Figure 11, Simple Q removes the first instruction and the markdown headers, simply presenting the question, the options and the final text. No opt. asks the question openly without multiple-choice options, choosing argmax similarity score, as in Section 3.2, as model choice. Presented with 95% Wald confidence interval.

	Std.	Detailed
Gemini Pro	$74 \pm 2$	$74 \pm 2$
GPT 3.5 Turbo	$73 \pm 2$	$74 \pm 2$
Mistral 7B Instruct (v0.2)	$70 \pm 2$	$71 \pm 2$
Mistral 7B	$62 \pm 3$	$58 \pm 3$
Dano. Mistral 7B	$57 \pm 3$	$59 \pm 3$
Dano. LlaMa 2 7B	$52 \pm 3$	$58 \pm 3$
LlaMa 2 7B	$54 \pm 3$	$56 \pm 3$
Dummy Baseline	$43 \pm 3$	$43 \pm 3$

Table 3: Impact of Nordjylland News alternative prompting.

chosen. Maximum generation likelihood-based selection was also implemented and is available for open-weights models on the frontend leaderboard but is not presented here.

**GPT-NER** Following (Wang et al., 2023), for each example, the GLLM was prompted four times, once for each entity category in the DaNE dataset (Hvingelby et al., 2020). The model was instructed to mark all words belonging to this entity category with @. These were parsed to one, multi-class prediction, handling overlap by selecting the generation with highest likelihood for models exposing probabilities. To mitigate small errors resulting in catastrophic results, model output annotated words were aligned using Levenshtein matching to the input example word list (Levenshtein, 1966).

	Std.	English
Gemini Pro	$61 \pm 8$	$64 \pm 8$
GPT 3.5 Turbo	$45 \pm 9$	$52 \pm 9$
Danoliterate Mistral 7B	$48 \pm 9$	$36 \pm 8$
Mistral 7B	$39 \pm 8$	$45 \pm 9$
Mistral 7B Instruct (v0.2)	$36 \pm 8$	$42 \pm 9$
LLama 2 7B	$33 \pm 8$	$33 \pm 8$
Danoliterate LLaMa 2 7B	$27 \pm 7$	$30 \pm 7$
Dummy Baseline	$21 \pm 6$	$21 \pm 6$

Table 4: How the Gym 2000 results change if the prompt instructions are in English. The instructed models handle this strongly while the Danoliterate Mistral model fails to perform under English prompting.

### A.3 Metrics

Differences in results for the similarity-based metrics used for #twitterhjerne and Nordjylland News summarization are presented in Tables 5 and 6. Model rank is minimally changed.

	Odd-one-out	Avg. sim.	Min. sim.	Max. sim.
GPT 3.5 Turbo	$29 \pm 5$	$64 \pm 5$	$61 \pm 5$	$66 \pm 5$
Gemini Pro	$31 \pm 5$	$63 \pm 5$	$61 \pm 5$	$66 \pm 5$
GPT 4	$35 \pm 5$	$63 \pm 5$	$61 \pm 5$	$66 \pm 5$
SOLAR 10.7B Instruct	$50 \pm 6$	$62 \pm 5$	$60 \pm 5$	$65 \pm 5$
LLaMa 2 13B Chat	$60 \pm 5$	$61 \pm 5$	$58 \pm 5$	$63 \pm 5$
Mistral 7B Instruct	$64 \pm 5$	$62 \pm 5$	$59 \pm 5$	$64 \pm 5$
Dano. Mistral 7B	$96 \pm 1$	$54 \pm 6$	$52 \pm 6$	$57 \pm 6$
Dano. LLaMa 2 7B	$97 \pm 1$	$53 \pm 6$	$50 \pm 6$	$55 \pm 6$
OpenAI Davinci 002	$99 \pm 0.3$	$52 \pm 6$	$49 \pm 6$	$54 \pm 6$
LLaMa 2 7B	$100$	$51 \pm 6$	$49 \pm 6$	$53 \pm 6$
Mistral 7B	$99 \pm 0.3$	$50 \pm 6$	$48 \pm 6$	$52 \pm 6$
Dummy Baseline	$100$	$44 \pm 6$	$43 \pm 6$	$46 \pm 6$

Table 5: Standard version of #twitterhjerne using the odd-one-out metric (1) compared to a simpler metric just reporting average similarity score.

### A.4 Software Dependencies

The relevant Python packages and their versions are presented in Table 7. Python version 3.11.8 was used.

	BERT similarity	ROUGE-1	ROUGE-L
Gemini Pro	$74 \pm 2$	$35 \pm 3$	$28 \pm 2$
GPT 3.5 Turbo	$73 \pm 2$	$32 \pm 2$	$25 \pm 2$
GPT 4	$73 \pm 2$	$32 \pm 2$	$23 \pm 2$
SOLAR 10.7B Instruct	$71 \pm 2$	$28 \pm 2$	$20 \pm 2$
Mistral 7B Instruct	$70 \pm 2$	$25 \pm 2$	$18 \pm 2$
LLaMa 2 13B Chat	$69 \pm 2$	$17 \pm 2$	$12 \pm 1$
Mistral 7B	$62 \pm 3$	$16 \pm 1$	$12 \pm 1$
Dano. Mistral 7B	$57 \pm 3$	$11 \pm 1$	$8 \pm 1$
OpenAI Davinci 002	$55 \pm 3$	$9 \pm 1$	$7 \pm 1$
Dummy Baseline	$43 \pm 3$	$11 \pm 1$	$8 \pm 1$
LLaMa 2 7B	$54 \pm 3$	$6 \pm 1$	$5 \pm 1$
Dano. LLaMa 2 7B	$52 \pm 3$	$5 \pm 1$	$4 \pm 0.5$

Table 6: Nordjylland News summarization results presented with an alternative lemma-based similarity score. The score is computed by lemmatizing text using the SpaCy framework (Honnibal et al., 2020) and then computing the ROUGE score (Lin, 2004)

Library	Version
google-cloud-aiplatform	1.38.1
openai	0.28.1
anthropic	0.21.3
groq	0.4.2
pandas	1.5.3
datasets	2.14.5
transformers	4.36.1
torch	2.1.1
evaluate	0.4.0
rouge_score	0.1.2
bert_score	0.3.13
huggingface_hub	0.19.4
hydra-core	1.3.2

Table 7: Evaluation framework Python dependencies and used versions.

## B Evaluation Corpora Details

### B.1 Data Permissions

- Citizenship Test:** All rights reserved "Styrelsen for International Rekruttering og Integration". Written permission was given for the data to be re-released as an appendix to Academic work.
- HyggeSwag:** MIT.
- Gym 2000:** Unreleased. Written permission was given by CRR for Academic use but not for re-releasing the dataset.
- #twitterhjerne:** CC-BY-4.0.
- Nordjylland News:** CC-0-1.0.

6. **Cloze Self Test:** Unreleased. Written permission was given by CRR for Academic use but not for re-releasing the dataset.

7. **DaNE:** CC-By-Sa-4.0-

8. **Angry Tweets:** CC-By-4.0.

## B.2 Data Content

All novel datasets were manually inspected for offensive content. Some crime-related and sexual themes were found in Nordjylland News examples but deemed unproblematic. The #twitterhjerne dataset was manually anonymized, removing all examples with personally identifiable content.

## B.3 Examples

Below, one prompted example per evaluation corpus is presented.

1. Citizenship Test: See Figure 4.

2. HyggeSwag

```
1 Svar kun med tallet for den
  rigtige fortsættelse af
  sætningen
2 # SÆTNING
3 En gruppe venner sidder på
  slæder på toppen af bakken.
  De to venner
4 # SVARMULIGHEDER
5 1. er udstyr kørende ned ad
  bakken med en udstyrsrem på.
6 2. presser deres rygge op mod en
  klippe.
7 3. skubber en slæde med et reb,
  da hele bakken er dækket
  med sne.
8 4. skubbes ned ad bakken, og de
  glider til bunden.
9 # SVAR
10 Den rigtige fortsættelse er
  mulighed nummer
```

3. #twitterhjerne

```
1 Skriv et kort tweet på dansk,
  der besvarer nedenstående
  spørgsmål. Svar kun med
  tweetet.
2 # TWEET MED SPØRGSMÅL
3 Sønnen vil gerne lave
  #pebernødder. De par gange
  jeg har prøvet det, blev de
  kun OK. Er der nogen, der
  kan anbefale en opskrift?
  #twitterhjerne
4 # TWEET MED SVAR
5 Et svar kunne være:
```

4. Gym 2000

```
1 "Selv før jeg lærte Max Kelada
  at kende, var jeg
  indstillet på ikke at kunne
  lide ham. Krigen var lige
  blevet afsluttet, og
  passagertrafikken på de
  store oceandampere var
  livlig. Det var meget
  vanskeligt at få plads, og
  man måtte finde sig i at
  tage, hvad skibsagenterne
  tilbød én. Man kunne ikke
  vente at få en kahyt for
  sig selv, og jeg var
  temmelig taknemmelig over
  at få en, hvor der kun var
  to køjer. Men da jeg
  erfarede navnet på min
  medpassager, sank mit
  humør. Det betød lukkede
  køjer, så det ikke ville
  være muligt at få den
  mindste smule frisk luft om
  natten. Det var ubehageligt
  nok at dele kahyt i fjorten
  dage med hvem som helst
  (jeg rejste fra San
  Francisco til Yokohama),
  men jeg ville have været
  mindre bekymret ved tanken,
  hvis min medpassagers navn
  havde været Smith eller
  Brown."
```

```
2
3 Svar kun med tallet for den
  rigtige mulighed
4 # SPØRGSMÅL
5 Hvorfor var det svært at få en
  kahyt for sig selv?
6 # SVARMULIGHEDER
7 1. Det var moderne at tage på
  krydstogt.
8 2. Det var midt i ferieperioden
9 3. Mange mennesker flyttede til
  USA
10 4. Krigen var lige forbi.
11 # SVAR
12 Svaret er mulighed nummer
```

5. Cloze Self Test

```
1 "Henrik bladede frem til siderne
  med boligannoncer. Deres
  lejlighed var <MASK> for
  lille til dem, så nu ledte
  de efter noget større. De
  ville gerne flytte lidt
  tættere på kysten. De ledte
  efter en lille gård , hvor
  der var plads til at holde
  et par heste ."
2 Erstat det maskerede ord i
  ovenstående tekst (markeret
  med '<MASK>') med et af
  følgende ord: indrettet,
  solgt, annonceret, blevet.
  Svar *kun* med det rigtige
  ord:
```

6. Nordjylland News

```
1 Skriv et kort dansk resumé på én
  enkelt sætning af følgende
  tekst.
```



2 # TEKST  
3 Manden kom kørende på Sønder  
Havnevej ved kiosken på  
havnen i Aalbæk, da han  
påkørte flere afmærkninger  
på stedet og fortsatte  
direkte ind i den bygning,  
hvor kiosken holder til.  
Der skete i forbindelse med  
påkørslen skade på  
bygningen. Uden for sad en  
mand, og han blev i lav  
fart påkørt af bilen ført  
af 53-årig mand. Den  
uheldige kiosk-gæst blev  
kørt til sygehuset med  
lettere skader.  
Nordjyllands Politi  
oplyser, at den 53-årige  
blev anholdt og sigtet for  
at køre i spirituspåvirket  
tilstand. Han er efter endt  
afhøring løsladt igen.

4 # RESUMÉ  
5 Et resumé på en sætning er:

## 7. Angry Tweets

1 Vurdér, om sentimentet i  
følgende tweet er  
'positiv', 'neutral' eller  
'negativ'. Svar kun med et  
enkelt ord.

2 # TWEET  
3 @USER Klæk det æg!  
4 # SENTIMENT:  
5 Sentimentet var

## 8. DaNE (prompting for location)

1 Fuldfør annotering af sidste  
eksempel i opgaven.

2 Her er en lingvists arbejde med  
at annotere entiteter af  
typen 'lokation'.

3 # TEKST  
4 Det blev naboens store, sorte  
hund også, "siger  
Københavns politidirektør,  
Poul Eefsen,  
galgenhumoristisk til B.T.  
efter et stort smykkekup i  
hans Holte-villa og en  
række tilsvarende kup i  
området.

5 # ANNOTERING  
6 Det blev naboens store , sorte  
hund også , " siger  
@@København##  
politidirektør , Poul  
Eefsen , galgenhumoristisk  
til B.T. efter et stort  
smykkekup i hans  
Holte-villa og en række  
tilsvarende kup i området .

7 # TEKST  
8 Diskussionen om forklaringen på  
det "japanske økonomiske  
mirakel" har især drejet  
sig om, hvorvidt man kunne  
nøjes med økonomiske  
faktorer i sin forklaring,  
eller om det også er  
nødvendigt at inddrage

særlige kulturelle og  
historiske forhold for at  
finde en rimelig forklaring.

9 # ANNOTERING  
10 Diskussionen om forklaringen på  
det " japanske økonomiske  
mirakel " har især drejet  
sig om , hvorvidt man kunne  
nøjes med økonomiske  
faktorer i sin forklaring ,  
eller om det også er  
nødvendigt at inddrage  
særlige kulturelle og  
historiske forhold for at  
finde en rimelig forklaring  
.

11 # TEKST  
12 De lyssky fremmede elementer af  
enhver art, der har sneget  
sig til landet, er fjenden.

13 # ANNOTERING  
14 De lyssky fremmede elementer af  
enhver art , der har sneget  
sig til landet , er fjenden  
.

15  
16 # TEKST  
17 "Vi tar'en tysker frem, vi  
tar'en tysker tilbage, vi  
tar'en tysker frem, åååårrr  
så ryster vi ham lidt!"

18 # ANNOTERING

## C Survey

### C.1 Survey Design and Instructions

For 100 GLLM use-cases divided into six categories (Zao-Sanders, 2024), we translated use-cases and categories into Danish and crafted an example prompt in Danish corresponding to that theme. We saved model answers from 18 models and used them in the survey to allow interactivity without requiring infrastructure for true dynamic model responses.

The survey front-end allowed volunteers to pick between the 100 prompts separated into categories, seeing model outputs from "Model A" and "Model B" side-by-side, streamed with a delay of 0.1 seconds between each word to simulate model generation. The volunteer was then instructed to try out at least a total of three prompts before answering. The answer consists of a question of preference, with optional additional Likert scales for each model and a text field for more details. The user instruction was in Danish meant *Two models have now been secretly selected for you: Model A and Model B. Test them out by choosing a prompt under a category that interests you. Look at the models' responses and get an impression of both A and B. Now choose a new prompt and please provide your*

## Hemmeligt par af sprogmodeller #1 ud af 4

Nu er der i hemmelighed valgt to modeller for dig: **A** og **B**.

Afprøv dem ved at vælge en prompt under en kategori, der interesserer dig. Se modellernes svar og få et indtryk af både A og B. Vælg nu en ny prompt og giv endelig din vurdering efter mindst 3 prompts.

### 1. Vælg prompts

Udforsk de seks kategorier og vælg en prompt, der interesserer dig.

**Læring og undervisning**   **Personlig støtte**   **Teknisk problemløsning og programmering**

Prompt 1: Følg medicinske symptomer

Prompt 2: Rådgivning om iværksætteri

Prompt 3: Forretningsmæssig rådgivning

Prompt 4: Hjælpe med at fokusere

Prompt 5: Terapi

### 3. Giv din vurdering

Du har set 3 forskellige eksempler for disse modeller.

Hvilken model foretrak du?

A    B    Ved ikke

Hvad synes du om **A**?

😄    😊    😐    😞    😡    Ved ikke

Hvad synes du om **B**?

😄    😊    😐    😞    😡    Ved ikke

Evt. uddybende tekst:

Tilbage   **Næste**

Afslut modellerne (læser dine svar)

### 2. Se modellernes svar

Jeg føler mig ensom på efter at være startet i l.g. Det føles, som om alle bliver bedste venner omkring mig, men jeg har det svært med så mange nye mennesker. Svar gerne kort: Er det normalt, hvad jeg føler?

**Model A**

Ja, det er normalt! Det kan være overvældende at starte i en ny skole eller klasse, og det kan tage tid at bygge relationer

**Model B**

Ja, det er normalt at føle ensomhed i begyndelsen af et nyt studieprogram, hvor man møder mange nye mennesker. Det kan

Figure 12: A screenshot of the survey UI presented to users. Under header 1, users select prompts of their interest. Header 2 contains model answers side-by-side and in header 3, volunteers fill in their preference and can get model identities revealed and move on to another pair.

assessment after at least 3 prompts.<sup>9</sup> See Figure 12 for an overview of the A/B test user interface.

## C.2 Volunteers

The survey is openly available online, inviting users to voluntarily try out the A/B tests, filling out their preferences. The survey was promoted on social media networks and newsletters. Most of the promotions were made on channels for AI enthusiasts or professionals. Volunteers were made aware that the data would contribute to studies into GLLM evaluation in Danish.

Volunteers could optionally fill in demographic details before carrying out A/B test which is shown in Figure 13 suggesting a bias towards young, male AI professionals.

<sup>9</sup>Da.: Nu er der i hemmelighed valgt to modeller for dig: Model A og Model B. Afprøv dem ved at vælge en prompt under en kategori, der interesserer dig. Se modellernes svar og få et indtryk af både A og B. Vælg nu en ny prompt og giv endelig din vurdering efter mindst 3 prompts.




Figure 13: Demographic details filled out by survey volunteers.

Work is ongoing to increase scale and diversity of respondents.

## C.3 Ranking Model

All 18 models included in the survey were sampled uniformly and the user model preference was used for ranking. An initial version ranking is the model win frequency presented in Table 8. As a model for the human preferences, we follow Chiang et al. to employ the Bradley-Terry model in a non-parametric fashion, using the sandwich robust standard errors (Chiang et al., 2024, Sec. 4, Sec. 5, Appendix B). The approach produces a linear model coefficient per model with estimated standard errors. These can be used for a paired Wilk's test to present significance of differences at  $\alpha = 0.05$  level.

## D Analysis Methodology Details

### D.1 Model Outcomes Linear Model

The model

$$d_m = \beta_0 + \beta_1 p_{\mathcal{M}} + \beta_2 \mathbb{I}(\mathcal{M} \in \text{instruct}) + \varepsilon, \quad (2)$$

where  $\varepsilon \sim \mathcal{N}(0, \sigma^2)$  and  $p_m$  is the number of model parameters in billions, was fitted with results shown in Table 9.

### D.2 Factor Analysis

For the EFA on the  $83 \times 8$  scenario results, the Bartlett Sphericity (Bartlett, 1951)  $p$  value is

	Lower	Estimate	Upper
Claude Opus	87%	94%	100%
Claude Sonnet	65%	76%	89%
GPT 4o	65%	76%	88%
Gemini Pro	59%	71%	85%
Claude Haiku	59%	71%	83%
GPT 4 Turbo 2024-04-09	55%	67%	80%
Starling 7B	53%	65%	81%
GPT 3.5 Turbo	48%	61%	72%
Heidrun 7B Chat	35%	48%	61%
Gemma 7B In.	29%	44%	60%
Mixtral 8x7B	24%	40%	56%
LlaMa 3 70B	24%	40%	53%
LlaMa 3 8B In.	19%	32%	44%
SOLAR 10.7B In.	17%	29%	41%
Munin NeuralBeagle	13%	26%	39%
Qwen1.5 7B Chat	11%	22%	33%
Mistral 7B In. v0.2	9%	20%	30%
DanskGPT-tiny Chat	6%	18%	29%

Table 8: How frequently each model wins their A/B tests with uncertainty estimation to a 95% confidence interval from bootstrapping blocked per volunteer.

Parameter	Value	Std. error	t-value
$\hat{\beta}_0$	17.2	8	
$\hat{\beta}_1$	15	3	5
$\hat{\beta}_2$	0.4	0.01	4
$\hat{\sigma}$	9		

Table 9: Linear model fitted with  $\hat{R}^2 = 0.6$  to the Danoliteracy Index for 38 open-weights models with known parameter counts.

$< 2 \cdot 10^{-16}$  and the Kaiser-Meyer-Olkin Test (Kaiser, 1970) yields a variance proportion of 90%, both suggesting that the data is usable for EFA. Fitting an EFA using the Scikit-learn Factor Analysis model yields  $\lambda_1 = 5.9$ ,  $\lambda_2 = 0.3$ . Explained factor variance is calculated as eigenvalue proportion of summed eigenvalues, and the analysis is repeated for scenario results acquired from the open API at [scandeval.com/danish-nlg/](https://scandeval.com/danish-nlg/), at [crfm.stanford.edu/helm/lite/v1.3.0/#/leaderboard](https://crfm.stanford.edu/helm/lite/v1.3.0/#/leaderboard), and using the OpenLLM Leaderboard Scraper GitHub project<sup>10</sup>. The datasets updated to most recent versions on January 13th, 2025.

All datasets were subjected Horn’s Parallel Analysis (Horn, 1965) simulating 1000 datasets of same shape but without correlation structure: This was implemented using the Python package `horns` (Mathias, 2024).

<sup>10</sup>[github.com/Weyaxi/scrape-open-llm-leaderboard](https://github.com/Weyaxi/scrape-open-llm-leaderboard)