







# DAKULTUR: Evaluating the Cultural Awareness of Language Models for Danish with Native Speakers

Max Müller-Eberstein   Mike Zhang  

Elisa Bassignana   Peter Brunsgaard Trolle  Rob van der Goot  

 IT University of Copenhagen, Denmark  Aalborg University, Denmark

 Pioneer Center for Artificial Intelligence, Denmark

mamy@itu.dk jjz@cs.aau.dk

## Abstract

Large Language Models (LLMs) have seen widespread societal adoption. However, while they are able to interact with users in languages beyond English, they have been shown to lack cultural awareness, providing anglocentric or inappropriate responses for underrepresented language communities. To investigate this gap and disentangle linguistic versus cultural proficiency, we conduct the first cultural evaluation study for the mid-resource language of Danish, in which native speakers prompt different models to solve tasks requiring cultural awareness. Our analysis of the resulting 1,038 interactions from 63 demographically diverse participants highlights open challenges to cultural adaptation: Particularly, how currently employed automatically translated data are insufficient to train or measure cultural adaptation, and how training on native-speaker data can more than double response acceptance rates. We release our study data as DAKULTUR—the first native Danish cultural awareness dataset.<sup>1</sup>

## 1 Introduction

Culture encompasses shared beliefs, norms, and worldviews (Tylor, 1871; Braff and Nelson, 2020), and tightly interweaves with language (Kramsch, 1998, 2014). These bidirectional influences affect how LLMs perform on culturally-sensitive tasks (Hovy and Yang, 2021). Contemporary LLMs are predominantly trained on English data, yet their global usage has outpaced their cultural coverage (Shi et al., 2023; Huang et al., 2023). Recent findings suggest that many models fail to adequately represent non-anglophone cultures (Hershcovich et al., 2022; Zhang et al., 2023; Liu et al., 2024),

resulting in culturally misaligned outputs that undermine user trust (Hovy and Yang, 2021; Litschko et al., 2023; Ge et al., 2024).

Addressing cultural misalignment requires assessing linguistic forms, common ground, aboutness, and values (Hershcovich et al., 2022). Prior efforts to improve alignment across these dimensions include the diversification of training data, as well as involving native speakers in evaluations (Tay et al., 2020; Huang and Yang, 2023; Cao et al., 2023; Naous et al., 2024; Wang et al., 2024). However, it remains unclear which LLM adaptation strategies (i.e., data selection, training methodology) lead to the highest linguistic and cultural alignment—especially for smaller languages.

This work investigates these questions by focusing on Danish, a mid-resource language that shares typological features with English, yet differs culturally to a non-trivial degree. By performing controlled experiments with respect to linguistic and cultural adaptation, we further aim to identify guidelines for culturally adapting LLMs to languages with similar characteristics and resource constraints. Our contributions are:

- DAKULTUR: The first native Danish dataset for the cultural evaluation of LLMs.
- A corresponding study, showing that native Danish data is critical to improving human acceptance rates (14%→42%), and accurate automatic cultural evaluation.
- An analysis of the effects of demographic factors (region, age, gender) on model alignment across different cultural topics.

## 2 DAKULTUR

### 2.1 Study and Data Collection Setup

To obtain a holistic perspective on Danish culture, we construct DAKULTUR based on a cultural evaluation study with native speakers in the loop. For

<sup>⊗</sup>These authors contributed equally.

<sup>1</sup>Dataset and code at <https://mxij.me/x/dakultur>. This study was approved by the ethics committee of the IT University of Copenhagen on 24th June 2024.




Figure 1: **Demographic Statistics for Our 63 Study Participants**, who were asked to optionally provide the region, where one grew up (Fig. 1a), age range in decades (Fig. 1b), and gender identity (Fig. 1c). 94% of respondents opted to provide this information.

this purpose, we build an open online interface (Fig. 3a), through which we task participants to compose prompts which require an understanding of Danish culture (Fig. 3c). We then use three different language models (Section 3.1) to generate answers<sup>2</sup>, which participants rate with an accept or reject, plus optional comments (Fig. 3d).

While the study is anonymous, we ask for optional demographic information (Fig. 3b), in order to assess the intra-cultural diversity of the respondents. We aim to collect only the minimal set of demographic features most likely to affect cultural responses, while not discouraging casual participation. Namely, we ask for the *region* where one grew up, (the five regions of Denmark, plus *other* for, e.g., people having grown up abroad), *age range* in decades, and *gender identity* (female, male, other).

After data collection, we manually validated the responses for relevance and correctness, and added topic annotations with a distinct set of five Danish speakers (Section 2.3). The resulting validated study data, in the form of DAKULTUR, not only serves to evaluate the cultural capabilities of the examined LLMs, but also constitutes—to the best of our knowledge—the first native Danish instruction dataset, with culturally-specific instructions, and human preference annotations.

## 2.2 Study Statistics

Our study was conducted over a period of two months, and was mainly advertised across higher educational institutions in Denmark. It attracted 1,038 input-response pairs with human quality judgments, from 63 participants (after validation).

**Demographics.** 94% of study respondents opted to provide demographic information, for which we find coverage of all regions (Fig. 1a) and gender

identities (Fig. 1c), as well as most age ranges except for <20 and >70 (Fig. 1b). We observe a slight skew towards participants who report having grown up in or around the Capital Region, that is 7% above the expected population share, while participants from Mid/Northern Jutland and Southern Denmark are underrepresented by 4–12%.

**Quality.** Generally, participants provided high-quality input, with 94.49% of prompts passing our post-study validation (Section 2.3). They further cover a diverse range of cultural concepts, as shown in the spread of topics in Fig. 2. Prompts are more frequently phrased as questions than as instructions (e.g., “how does a hot-dog stand look?” versus “describe how a hot-dog stand looks like”). Furthermore, the majority of inputs query the models’ cultural knowledge directly instead of via its situational awareness of societal norms (e.g., by prompting models to resolve dilemmatic situations). As prompts in the latter format are much more time-intensive to create, this is likely to be expected. Participants further steered clear of politically and morally-charged topics, despite their anonymity. The resulting collection of cultural prompts therefore contains cultural concepts, that appear to enjoy a more uniform consensus.

## 2.3 Post-study Validation

Post-study, we validate and analyze the resulting data in another round of annotation with a distinct set of five Danish speakers. The study data is split across annotators, and each annotator is tasked to verify whether an input is dependent on a Danish cultural context (i.e., valid for this study), as well as which broader main topic it belongs to. For annotating topics, we employ an open coding strategy (Strauss, 1987), which resulted in the following 12 topics (+ other):

<sup>2</sup>Answer order was shuffled after each trial.

MODEL	LANGUAGE						CULTURE		
	LA	NER	SA	AS	CSR	QA	PE	CT	DK
LLAMA2-7B <sub>base</sub>	33.4	23.7	61.5	65.5	29.8	63.5	38.6	57.7	—
+ INST <sub>da</sub>	36.1	28.5	62.9	66.4	29.0	64.4	49.1	58.5	13.9
LLAMA2-7B <sub>chat</sub>	47.4	24.6	66.2	66.3	<b>32.2</b>	61.3	46.7	55.2	—
+ INST <sub>da</sub>	43.4	29.7	65.9	65.8	31.0	62.5	57.3	55.6	15.0
SNAKMODEL	<b>52.9</b>	<b>29.8</b>	<b>66.7</b>	<b>66.6</b>	29.5	<b>64.7</b>	<b>71.1</b>	<b>71.9</b>	<b>42.4</b>

Table 1: **Results on the ScandEval Benchmark (Test) and DAKULTUR (DK)**. Higher scores are better, with exact metrics depending on the sub-task (Section 3.1). We differentiate between linguistic tasks (under LANGUAGE), and cultural tasks (under CULTURE). Additionally, we include scores for the English LLAMA2-7B<sub>base</sub> and LLAMA2-7B<sub>chat</sub> (Touvron et al., 2023). The best score per sub-task is highlighted in **bold**.

- **arts**: media and their place in society (e.g., “name five popular Danish TV programs”).
- **education**: regarding the education system (e.g., “which university is best to learn about AI in Denmark?”).
- **food**: regarding dishes and culinary traditions (e.g., “can I serve herring on french bread?”).
- **geography**: regions, cities, and climate (e.g., “where can you go on vacation in the south?”).
- **language**: proficiency in appropriate responses and proverbs (e.g., “what does it mean to be a pineapple in its own juice?”).
- **lifestyle**: everyday activities that are not as strict as norms (e.g., “what should I prepare when going to a Danish beach?”).
- **norms**: implicit rules that are followed in Danish society (e.g., “explain the effect of ‘the law of Jante’ on Danish mentality”).
- **politics**: knowledge of the political system, figures, and parties (e.g., “how do I become a member of the regional parliament?”).
- **traditions**: customs and events, observed across multiple generations (e.g., “what do you do with a 25-year-old who’s single?”).
- **transport**: knowledge and customs regarding transportation systems (e.g., “can you turn left on a bicycle at a traffic light?”).
- **trivia**: factual knowledge about people, places, historical events, sports etc., which are not part of the other topics (e.g., “in what year was the reunification of Southern Jutland?”).
- **work**: procedures and behaviors, that are appropriate for a professional context (e.g., “how do I ask my manager for a raise?”).

### 3 Cultural Evaluation

We next investigate the results of our cultural evaluation study, and compare the metrics from DAKULTUR with those of automatic benchmarks.

#### 3.1 Experimental Setup

**Models.** In our study, we use three LLMs based on LLAMA2-7B<sub>base</sub> (Touvron et al., 2023), adapted to Danish using distinctive training strategies: Danish language modeling training (LMT<sub>da</sub>), and instruction tuning on translated data (INST<sub>da</sub>). The corresponding models are LLAMA2-7B<sub>base</sub>+INST<sub>da</sub>, LLAMA2-7B<sub>chat</sub>+INST<sub>da</sub>, and SNAKMODEL (Zhang et al., 2024; LLAMA2-7B<sub>base</sub>+LMT<sub>da</sub>+INST<sub>da</sub>).

**Automatic Evaluation.** To compare the human judgments in DAKULTUR with existing automatic metrics, we employ the Danish part of ScandEval (Nielsen, 2023), across its sub-tasks on linguistic acceptability (LA from ScaLA<sup>3</sup>); named entity recognition (NER from DANSK; Hvingelby et al., 2020); sentiment analysis (SA from AngryTweets; Pauli et al., 2021); abstractive summarization (AS from Nordjylland-News; Kinch, 2023); commonsense reasoning (CSR from HellaSwag; Zellers et al., 2019); and question answering (QA from ScandiQA<sup>4</sup>). ScandEval further includes two culturally-oriented tasks: Danske Talemåder (PE; Nielsen, 2023), which prompts for meanings behind Danish proverbs, and a collection of Danish Citizenship Tests (CT; Nielsen, 2024).

<sup>3</sup>Based on Danish data from the Universal Dependencies dataset from (Kromann and Lyng, 2004).

<sup>4</sup>Note that ScandiQA is a translation of the English MKQA dataset Longpre et al., 2021, and does not strictly focus on Scandinavian knowledge.




Figure 2: **Acceptance/Rejection Rates** across SNAKMODEL, LLAMA2-7B<sub>chat</sub>+INST<sub>da</sub> and LLAMA2-7B<sub>base</sub>+INST<sub>da</sub> as judged by participants in DAKULTUR. Left: overall results; Right: results by topic.

### 3.2 Results

**Automatic Metrics.** Results on ScandEval (Table 1) show that training on native Danish data (i.e., SNAKMODEL) leads to the greatest performance gains across the board. While the unadapted English models perform comparably on some tasks, it is important to note that ScandEval employs constrained generation. When prompted without constraint, both LLAMA2-7B<sub>base</sub> and LLAMA2-7B<sub>chat</sub> generate English responses. Instruction tuning using translated data is already sufficient to enforce Danish responses (even when prompted in English), which is why we employ the +INST<sub>da</sub> variants in our human study. Nonetheless, we observe that translated data is insufficient to induce much cultural knowledge into the model, as only SNAKMODEL improves on the cultural tasks of PE and CT to a substantial degree.

**DAKULTUR Results.** In terms of acceptance rates, SNAKMODEL obtains a rate more than twice as high compared to the other models (Fig. 2). Nonetheless, with a maximum acceptance rate of 42%, none of the models appears to provide particularly well-adapted responses—highlighting the gap between cultural versus linguistic adaptation. Qualitatively, we observe that answers are almost never rejected due to linguistic errors, but rather due to incorrect or incomplete factual content.

Our post-study analysis reveals that the cultural topics of *food* and *traditions* are most popular, and that SNAKMODEL achieves acceptance rates over ten times as high for these topics. While training on native data improves performance across all top-

ics, gains are larger for implicit cultural knowledge (e.g., *lifestyle*, *norms*) than for facts (e.g., *trivia*, *geography*, *politics*). In Appendix C, we further show how topics and acceptance rates vary by demographics. Female-identity participants tend toward *food*, *lifestyle*, *education*, and *norms*, while male-identity participants focus more on *politics*, *trivia*, and *geography*. Additionally, younger participants and those from the capital region report slightly higher acceptance rates.

### 4 Conclusion

In this work, we introduced DAKULTUR—the first native Danish cultural evaluation dataset. By constructing it via a native-speaker-driven evaluation study, and applying a thorough post-study validation, we are able to share 1,038 high-quality input-response pairs for future Danish NLP research. Our cultural evaluation using DAKULTUR highlights that language modeling training using native data is already sufficient to more than double human-judged cultural awareness—especially for popular cultural topics. Simultaneously, the maximum acceptance rate of 42% highlights that more research is needed to fully align anglocentric LLMs to smaller language communities, such as Danish. In terms of evaluation methodologies, the fact that human judgments align more with the smaller, yet culturally-relevant and non-translated sub-tasks of the automatic ScandEval benchmark (PE, CT, as well as LA) is encouraging, since small amounts of high-quality data may already be sufficient to accurately estimate an LLM’s cultural awareness.

## Limitations

While we strive for broad coverage of the Danish cultural landscape, culture itself has a high degree of inherent subjectivity and variability. As such, future work using DAKULTUR should be cognizant of the context in which its data was obtained. Our cultural evaluation study was advertised primarily at higher educational institutions. Although we are aware of word-to-mouth advertisement stretching to demographic groups beyond this initial cluster (as evidenced by the range of represented age groups), the study likely does not capture the full breadth of the Danish cultural landscape. By gathering demographics for intra-cultural differences with regard to topics and user acceptance rates, we nonetheless aim to enable analyses with respect to how much cultural consensus might vary with respect to different topics. We believe this is crucial information for practitioners designing downstream systems, as contemporary models seem to, for instance, align slightly better with male-identity participants under 30 from the capital region.

On the technical side, we hope that future work will be able to validate our findings across more base models and languages. Our choice of Danish and LLAMA2-7B-based models was primarily driven by data and compute resource availability. Similarly, while DAKULTUR can theoretically be used for small-scale instruction tuning or model alignment, its size is far from contemporary, automatically generated datasets. For cultural evaluation purposes, we nonetheless believe that it offers a representative out-of-the-box solution for developers of future Danish LMs.

## Acknowledgments

We thank the NLPnorth group at ITU and the AAU-NLP group at AAU for feedback on earlier version of this draft. Elisa Bassignana is supported by a research grant (VIL59826) from VILLUM FONDEN. Mike Zhang is supported by a research grant (VIL57392) from VILLUM FONDEN.

## References

- Lara Braff and Katie Nelson. 2020. Introduction to anthropology. In *Perspectives: An Open Invitation to Cultural Anthropology*, pages 3–28. Society for Anthropology in Community Colleges.
- Yong Cao, Li Zhou, Seolhwa Lee, Laura Cabello, Min Chen, and Daniel Hershcovich. 2023. [Assessing](#)

[cross-cultural alignment between ChatGPT and human societies: An empirical study](#). In *Proceedings of the First Workshop on Cross-Cultural Considerations in NLP (C3NLP)*, pages 53–67, Dubrovnik, Croatia. Association for Computational Linguistics.

- Xiao Ge, Chunchen Xu, Daigo Misaki, Hazel Rose Markus, and Jeanne L Tsai. 2024. [How culture shapes what people want from ai](#). In *Proceedings of the CHI Conference on Human Factors in Computing Systems*, CHI '24, New York, NY, USA. Association for Computing Machinery.

- Daniel Hershcovich, Stella Frank, Heather Lent, Miryam de Lhoneux, Mostafa Abdou, Stephanie Brandl, Emanuele Bugliarello, Laura Cabello Piqueras, Ilias Chalkidis, Ruixiang Cui, Constanza Fierro, Katerina Margatina, Phillip Rust, and Anders Søgaard. 2022. [Challenges and strategies in cross-cultural NLP](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6997–7013, Dublin, Ireland. Association for Computational Linguistics.

- Dirk Hovy and Diyi Yang. 2021. [The importance of modeling social factors of language: Theory and practice](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 588–602, Online. Association for Computational Linguistics.

- Haoyang Huang, Tianyi Tang, Dongdong Zhang, Xin Zhao, Ting Song, Yan Xia, and Furu Wei. 2023. [Not all languages are created equal in LLMs: Improving multilingual capability by cross-lingual-thought prompting](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 12365–12394, Singapore. Association for Computational Linguistics.

- Jing Huang and Diyi Yang. 2023. [Culturally aware natural language inference](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 7591–7609, Singapore. Association for Computational Linguistics.

- Rasmus Hvingelby, Amalie Brogaard Pauli, Maria Barrett, Christina Rosted, Lasse Malm Lidegaard, and Anders Søgaard. 2020. [DaNE: A named entity resource for Danish](#). In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 4597–4604, Marseille, France. European Language Resources Association.

- Oliver Kinch. 2023. [Nordjylland news summarization](#).

- Claire Kramsch. 1998. *Language and culture*. Oxford introductions to language study. Oxford University Press, Oxford.

- Claire Kramsch. 2014. [Language and culture](#). *AILA Review*, 27(1):30–55.

- Matthias Trautner Kromann and Stine Kern Lynge. 2004. The danish dependency treebank v. 1.0.
- Robert Litschko, Max Müller-Eberstein, Rob van der Goot, Leon Weber-Genzel, and Barbara Plank. 2023. [Establishing trustworthiness: Rethinking tasks and model evaluation](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 193–203, Singapore. Association for Computational Linguistics.
- Chen Liu, Fajri Koto, Timothy Baldwin, and Iryna Gurevych. 2024. [Are multilingual LLMs culturally-diverse reasoners? an investigation into multicultural proverbs and sayings](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 2016–2039, Mexico City, Mexico. Association for Computational Linguistics.
- Fangyu Liu, Emanuele Bugliarello, Edoardo Maria Ponti, Siva Reddy, Nigel Collier, and Desmond Elliott. 2021. [Visually grounded reasoning across languages and cultures](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 10467–10485, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Shayne Longpre, Yi Lu, and Joachim Daiber. 2021. Mkqa: A linguistically diverse benchmark for multilingual open domain question answering. *Transactions of the Association for Computational Linguistics*, 9:1389–1406.
- Tarek Naous, Michael J. Ryan, Alan Ritter, and Wei Xu. 2024. [Having beer after prayer? measuring cultural bias in large language models](#). *Computing Research Repository*, arxiv:2305.14456. Version 4.
- 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.
- Dan Saattrup Nielsen. 2024. [Danish citizen test](#).
- 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)*, pages 460–466, Reykjavik, Iceland (Online). Linköping University Electronic Press, Sweden.
- Freda Shi, Mirac Suzgun, Markus Freitag, Xuezhi Wang, Suraj Srivats, Soroush Vosoughi, Hyung Won Chung, Yi Tay, Sebastian Ruder, Denny Zhou, Dipanjan Das, and Jason Wei. 2023. [Language models are multilingual chain-of-thought reasoners](#). In *The Eleventh International Conference on Learning Representations*.
- Anselm L Strauss. 1987. *Qualitative analysis for social scientists*. Cambridge university press.
- Yi Tay, Donovan Ong, Jie Fu, Alvin Chan, Nancy Chen, Anh Tuan Luu, and Chris Pal. 2020. [Would you rather? a new benchmark for learning machine alignment with cultural values and social preferences](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5369–5373, Online. Association for Computational Linguistics.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shrutu Bhosale, et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Edward Burnett Tylor. 1871. *Primitive culture: researches into the development of mythology, philosophy, religion, art, and custom*, volume 2. J. Murray.
- Wenxuan Wang, Wenxiang Jiao, Jingyuan Huang, Ruyi Dai, Jen-tse Huang, Zhaopeng Tu, and Michael Lyu. 2024. [Not all countries celebrate thanksgiving: On the cultural dominance in large language models](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6349–6384, Bangkok, Thailand. Association for Computational Linguistics.
- 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.
- Mike Zhang, Max Müller-Eberstein, Elisa Bassignana, and Rob van der Goot. 2024. [Snakmodel: Lessons learned from training an open danish large language model](#). *Preprint*, arXiv:2412.12956.
- Xiang Zhang, Senyu Li, Bradley Hauer, Ning Shi, and Grzegorz Kondrak. 2023. [Don't trust ChatGPT when your question is not in English: A study of multilingual abilities and types of LLMs](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 7915–7927, Singapore. Association for Computational Linguistics.

## Appendix

### A Study Interface

We build a web-based evaluation interface (study flow shown in Fig. 3), which allows study participants to prompt the three LLMs simultaneously<sup>5</sup> with tasks and questions, that require cultural awareness (Fig. 3c), and to rate (accept/reject) and comment on the models’ responses (Fig. 3d). The study guidelines (Figs. 3a and 3c) broadly lay out which dimensions of cultural awareness the study aims to investigate—i.e., common ground, aboutness, objectives and values, in addition to linguistic form and style (Hershcovich et al., 2022), which is implicit, given the study’s monolingual nature. Following prior work on culturally diverse dataset creation (Liu et al., 2021), we opted for an elicitation setup in order to avoid biasing responses towards a limited set of cultural concepts and topics.

While the study is conducted anonymously, we ask for optional demographic information (Fig. 3b), in order to assess the intra-cultural diversity of the respondents. For this purpose, we aimed to collect only the minimal set of demographic features, that we deemed most likely to affect cultural responses, while not discouraging casual participation. Namely, we ask for the region, where one grew up, (the five regions of Denmark, plus *other* for, e.g., people having grown up abroad), age range in decades, and gender identity (female, male, other).

In test trials, we noticed that, while participants intuitively prompted for a wide variety of culturally-relevant topics, they typically did so in a multi-turn conversational manner, which our single-turn, instruction-tuned models often fail to answer. For instance, the prompt “Hello! Could you tell me about [...]?”, frequently produces the response, “Yes, I can.”, with no further relevant content. To encourage single-turn instruction-style inputs, we iterated over multiple guideline formulations, of which we found, “Ask one question or give one task about Danish culture [...] to the three virtual assistants below”, to produce the most compatible results (see full translations in Appendix B).

<sup>5</sup>Note that the order in which responses are displayed is randomized with each prompt.

## B Translations

### B.1 Landing Page with Guidelines

Thanks for your interest in our research project!

**Purpose** We examine cultural skills/competencies with artificial intelligence and would like you to assess our three virtual assistants’ knowledge of Danish culture. This includes, for example, norms, art and laws in Danish society, as well as Danes’ knowledge, beliefs, customs and habits.

**Task** On the following pages, you should ask the virtual assistants to perform tasks and assess their answers one-by-one. You can ask them questions or ask them for descriptions, e.g., “tell me how to change the back tire of my bike”.

**Data policy** As this is a research project, input, feedback and optional demographic data are stored in a dataset. All data is collected anonymously.

If you agree to the above terms, continue by clicking Accept below.

### B.2 Demographic Information

Your anonymous session ID is:  
SESSION\_ID

Save it in a safe place since it is required if you would like to get your answers removed from our dataset.

**Demographics** Enter your demographic information below (one or more can be omitted). This helps us to ensure that we get a more diverse data set.

**Region** (where you grew up)  
*do not wish to disclose*

**Age**  
*do not wish to disclose*

**Gender Identity**  
*do not wish to disclose*

Click *Start* to get started!

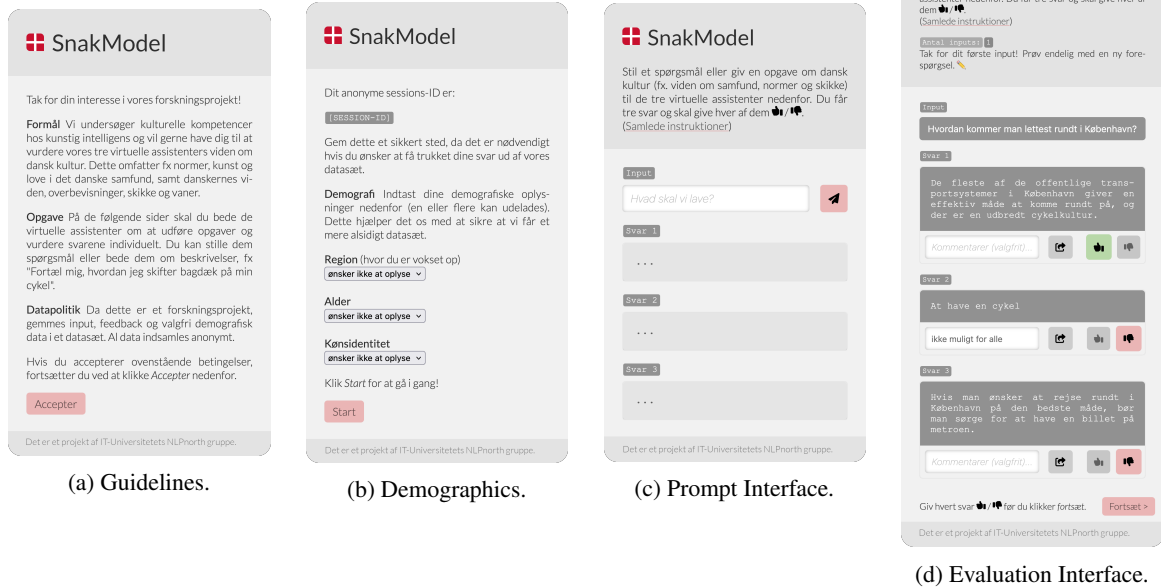


Figure 3: **Study Interface for Human Cultural Evaluation.** Participants are guided through the guidelines (Fig. 3a), optional demographic registration (Fig. 3b), before being asked to prompt the three LLMs simultaneously (Fig. 3c), and to evaluate the model responses (Fig. 3d). Translations of the guidelines, interface, and examples can be found in Appendix B.

### B.3 Prompt Interface

Ask one question or give one task about Danish culture (e.g., knowledge of society, norms and customs) to the three virtual assistants below. You will receive three answers, which you can each rate with a thumbs-up/down.

#### Input

*What shall we do?*

#### Answer 1

#### Answer 2

#### Answer 3

### B.4 Response Evaluation Interface

Thanks for your first input! Go ahead, and try another request!

#### Input

*What's the easiest way to get around in Copenhagen?*

#### Answer 1

Most of the public transport systems in Copenhagen provide an effective way to get around, and there is a widespread

cycling-culture.

#### Answer 2

To have a bike

*Comment: not possible for everyone*

#### Answer 3

If you want to travel around Copenhagen, you should make sure to have a ticket for the subway.

Give each answer a thumbs-up/down before clicking *Continue*.

## C Topics and Acceptance Rates per Demographic

For each demographic dimension, we merge the available categories into two groups, in order to have a large enough amount of information to compare. This leads to splits along male/female,  $< 30/\geq 30$ , capital region/other regions. The distribution as well as acceptance rates are shown in Fig. 4 for gender, Fig. 5 for age, and Fig. 6 for region.






Figure 4: **Acceptance/Rejection Rates and Distribution across Topics** for the female/male gender identity demographic groups.




Figure 5: **Acceptance/Rejection Rates and Distribution across Topics** for the age ranges  $\geq 29$  and  $\leq 30$ .




Figure 6: **Acceptance/Rejection Rates and Distribution across Topics** for participants from the capital versus other regions.