

Who Benchmarks the Benchmarks? A Case Study of LLM Evaluation in Icelandic

Finnur Ágúst Ingimundarson¹, Steinunn Rut Friðriksdóttir², Bjarki Ármannsson^{2,3},
Iris Edda Nowenstein², Steinþór Steingrímsson³

¹University of Zürich, ²University of Iceland, ³The Árni Magnússon Institute for Icelandic Studies
finnuragust.ingimundarson@uzh.ch, {srf, irisen}@hi.is, {bjarki.armannsson,
steinthor.steingrimsson}@arnastofnun.is

Abstract

This paper evaluates current Large Language Model (LLM) benchmarking for Icelandic, identifies problems, and calls for improved evaluation methods in low/medium-resource languages in particular. We show that benchmarks that include synthetic or machine-translated data that have not been verified in any way, commonly contain severely flawed test examples that are likely to skew the results and undermine the tests' validity. We warn against the use of such methods without verification in low/medium-resource settings as the translation quality can, at best, only be as good as MT quality for a given language at any given time. Indeed, the results of our quantitative error analysis on existing benchmarks for Icelandic show clear differences between human-authored/-translated benchmarks vs. synthetic or machine-translated benchmarks.

Keywords: benchmarking, machine translation, low-resource languages, Icelandic

1. Introduction

Benchmarking has been an integral part of development in NLP as a way of assessing performance on various tasks and traditionally on tasks with a single-label output, such as POS-tagging, NER, sentiment analysis, etc. With the advent of LLMs, that are not only capable of performing all of these traditional tasks and yet more (Brown et al., 2020) but perhaps even of creating such data, the need for robust benchmarks is even greater than before. In their wake also follows a demand for creating new benchmarks that are better adapted to the capabilities of LLMs, i.e. beyond evaluating single-label output, and importantly, beyond tasks which purely target the scope of expertise of NLP researchers (Bowman and Dahl, 2021).

This means that benchmark set creators increasingly find themselves creating benchmarks for natural language production and understanding in languages they do not speak and within subjects of which they are not experts. For example, there is a clear difference between constructing a NER evaluation metric in a language the researcher knows and generating a machine-translated benchmarking set for the evaluation of an LLMs' capabilities in medical text summarization (Alaa et al., 2025). Furthermore, the field – and therefore benchmarking – has largely been dominated by English (and a few other high-resource languages (Vayani et al., 2024; Wang et al., 2024b)), but given the multilingual capabilities of LLMs, it is important to put them to the test in a valid manner. For one thing, it can be informative to see how different scripts or more morphologically rich languages compare. This is of particular importance for low-resource languages,

as they run the risk of being submerged in the digital ocean of English data.

For low- and medium-resource languages, it is often hard to keep up with the pace and establish what already exists for high-resource languages. Benchmarking is no exception. As benchmark production with human input can be time-consuming and costly, one common approach in such settings is to machine-translate existing benchmarks for other languages or use LLMs to create synthetic benchmarks, for instance by generating question-answer pairs from pre-existing data such as Wikipedia. But what are the drawbacks to this approach and do they affect low-resource languages disproportionately?

In this paper we look at the current state of benchmarking for Icelandic as a case study of a low/medium-resource language, discuss the effects and results including machine-translated and synthetic data, and present the results of a quantitative error analysis on existing benchmarks for Icelandic. They show clear differences between human-authored or human-translated benchmarks vs. synthetic or machine-translated benchmarks, especially where the output has not been verified and/or corrected.¹

2. Related Work

Discussion on the pitfalls and problems of language model benchmarking, and the question of how to best compare the capabilities of different models,

¹All code and data available at https://github.com/finaing/who_benchmarks_the_benchmarks

has been ongoing in the NLP community for years. At least as early as [Bowman and Dahl \(2021\)](#), it had become apparent that the best-performing NLU models of the day, though far from perfect in real-world applications, had almost completely saturated existing benchmarks, begging the question of how best to measure actual progress. Four years later, this is still an open and important question as evidenced by [Eriksson et al. \(2025\)](#) which identifies ‘key shortcomings’ in benchmarking practices over the last decade, including ‘issues in the design and application’, ‘misaligned incentives’, ‘construct validity issues,’ and other shortcomings directly related to the points made in this paper.

As previously discussed, it has become commonplace practice in recent years to automate the construction of benchmarks and evaluation sets either fully or partially. Here, statistical power (which is certainly important, see e.g. [Card et al. \(2020\)](#)) is seemingly prioritized over precision and overview, despite the fact that machine learning models are frequently ‘right for the wrong reasons’ ([McCoy et al., 2019](#)), capable of achieving high scores through learning unintended shallow heuristics. Indeed, there is considerable evidence that smaller, carefully curated test sets are often equally good or better indicators of performance than larger, synthetic sets (see e.g. [Gardner et al. \(2020\)](#); [Shaib et al. \(2024\)](#); [Yadav et al. \(2024\)](#)). This is in line with the importance of *validity* in any kind of testing. Present already in e.g. [Cronbach and Meehl \(1955\)](#), validity is a key concept in psychometrics and addresses whether a test or evaluation instruments truly measures what it intends to measure (the construct). In the context of LLM testing and particularly in clinical applications, calls have been made to prioritize construct validity and employ a Benchmark-Validation-First Approach to (medical) LLM evaluation ([Alaa et al., 2025](#)). To paraphrase, benchmarks need to be right for the right reasons.

Even automatically constructing test items using a structured template, now the norm for grammatical evaluation sets ([Warstadt et al., 2020](#); [Huebner et al., 2021](#); [Song et al., 2022](#); [Jumelet et al., 2025](#)), can easily result in semantically bizarre examples and items solvable through superficial heuristics rather than genuine linguistic competence ([Vázquez Martínez et al., 2023](#)). Similarly, the widespread practice of evaluating systems for different languages by automatically translating existing benchmarks (usually from English) has repeatedly been shown to result in translation artifacts which affect accuracy and omit language- and culture-specific context ([Clark et al., 2020](#); [Park et al., 2024](#); [Chen et al., 2024](#); [Liu et al., 2025](#)). For instance, [Umutlu et al. \(2025\)](#) examine 17 Turkish datasets (a “mid-resource” language), many of which are translations/adaptations

from English. They identify major shortcomings related to coherence, fluency, cultural alignment, and increased risk of bias and error. Similarly, [Yao et al. \(2024\)](#) note major shortcomings in culturally specific machine translations where exact translation pairs might not exist. Furthermore, [Semenov and Sennrich \(2025\)](#) demonstrate shortcomings of template-based translations in multilingual benchmarks, showing that sentence-level translations yield more reliable results.

As noted by [Vayani et al. \(2024\)](#), multilingual benchmarks remain limited in scope and disproportionately represent high-resource languages, neglecting the cultural and linguistic diversity of lower-resource languages. Moreover, benchmarks are often absorbed into training corpora, resulting in contamination and inflated evaluation scores ([Yang et al., 2023](#)).

When evaluation datasets rely on translations of English benchmarks without cultural adaptation, models may optimize for unnatural constructions rather than authentic usage. This issue is amplified in small-language settings, where limited corpora increase the influence of benchmarks on model optimization. Indigenous languages in particular face severe digital resource scarcity, both in corpus size and representativeness. As [Wiechetek et al. \(2024\)](#) argue, when developers lack proficiency or active engagement with the target language, systematic errors may go unnoticed. Scarcity increases reliance on synthetic data, which can recursively dominate future training cycles, amplifying distortions and reinforcing inaccuracies. In extreme cases, synthetic outputs may even invert the meaning of source texts, causing direct harm to the language community.

Consequently, models may perform well on translated test sets while failing to capture culturally grounded semantics or authentic discourse patterns, leading to systematic overestimation of model capabilities ([Sainz et al., 2023](#)). This concern is supported by [Wang et al. \(2024a\)](#), who show that LLM performance is inconsistent across languages, with degradation in lower-resource settings and stronger results on English cultural reasoning tasks. Similarly, [Wang et al. \(2024b\)](#) demonstrate that LLMs tend to reproduce English-centric cultural assumptions even when prompted in other languages. Direct translation of culturally specific English benchmarks may therefore reward English-oriented reasoning rather than culturally situated understanding in the target language.

3. Leaderboards for Icelandic

As of October 2025, there are two publicly available leaderboards for Icelandic. One is maintained by the Icelandic LT company, Miðeind, on Hugging

Face² and the other is a part of the EuroEval project, previously ScandEval (Nielsen, 2023).³

3.1. Miðeind’s leaderboard

At the time of writing, Miðeind’s leaderboard evaluates 59 LLMs across 6 different benchmarks:

- *WinoGrande-IS*: a human-translated and localized version of the WinoGrande test set (Snæbjarnarson et al., 2022; Sakaguchi et al., 2021)
- *GED*: binary sentence-level GED with data from the Icelandic Error Corpus (Arnardóttir et al., 2021)
- *Inflection*: inflection of 300 Icelandic adjective-noun pairs in all four cases and both numbers⁴
- *Belebele*: Icelandic subset of Belebele (reading comprehension) (Bandarkar et al., 2024)⁵
- *ARC-Challenge-IS*: machine-translated version of the ARC-Challenge benchmark (Clark et al., 2018)⁶
- *WikiQA-IS*: 1.9k QA pairs created by GPT-4o from Icelandic Wikipedia data, manually verified and corrected (Arnardóttir et al., 2025)

3.2. EuroEval’s leaderboard

EuroEval’s leaderboard contains benchmarks and evaluation for 18 European languages, at the time of writing, and evaluates all types of language models, fine-tuning encoder models and either zero- or few-shot prompting decoder models for different tasks.

Like Miðeind, EuroEval uses WinoGrande-IS and Belebele, but a different machine-translated version of the ARC-Challenge. It additionally uses WikiQA-IS in a different format, where an LLM is used to generate three plausible answer alternatives, turning it into a multiple-choice answering task, but the benchmark has not been published in this format. Other datasets or benchmarks included are:

- *Hotter and Colder Sentiment*: sentiment analysis dataset (Friðriksdóttir et al., 2025)
- *MIM-GOLD-NER*: NER dataset (Ingólfssdóttir et al., 2020, paper), (Ingólfssdóttir et al., 2020, data)
- *ScaLA-IS* (Nielsen, 2023), *IceEC* (sampled data from the Icelandic Error Corpus), subset of *IceLinguistic* (Ármannsson et al., 2025): All used for a binary evaluation of grammaticality or GED

²huggingface.co/spaces/mideind/icelandic-llm-leaderboard

³euroeval.com

⁴huggingface.co/datasets/mideind/icelandic-inflection-all-flat

⁵huggingface.co/datasets/facebook/belebele/viewer/isl_Latn

⁶huggingface.co/datasets/mideind/icelandic-arc-challenge

- *NQil*: “Natural Questions in Icelandic”, dataset for open domain QA with retrieval (Snæbjarnarson and Einarsson, 2022)⁷

- *MultiWikiQA-IS*: created using LLM Gemini-1.5-pro on Wikipedia data (306 languages) (Smart, 2025)⁸

- *MMLU-IS*: machine-translated version of the US MMLU benchmark (Hendrycks et al., 2021)⁹

- *HellaSwag-IS*: machine-translated version of the HellaSwag benchmark (Zellers et al., 2019)¹⁰

- *RRN*: summarization dataset (Sverrisson and Einarsson, 2023)¹¹

It is important to note two things. First, in some cases datasets are being turned into benchmarks, despite not being intended as such originally. This goes for *Hotter and Colder*, *MIM-GOLD-NER*, *IceEC*, *NQil* and *RRN*. Second, EuroEval labels some benchmarks as “unofficial”. As far as we can see, there is no information to be found on the website on what factors are decisive in making this distinction between benchmarks, but the label apparently refers to benchmarks for which the results are not published on the leaderboard. “Official” benchmarks are: *Hotter and Colder Sentiment*, *MIM-GOLD-NER*, *SCaLA-IS*, *NQil*, *IcelandicKnowledge*, *WinoGrande-IS* and *RRN*. However, the published code on Github and the EuroEval Python package include scripts to create all of the benchmarks (and splits), regardless of whether they are unofficial or not, and benchmark models on them.¹²

4. Quantitative Error Analysis

A review of benchmarks that are currently being used for Icelandic revealed various problems in some of the datasets and, to some extent, in the way they are being used. To keep this paper from consisting of seemingly cherry-picked examples, we conducted a quantitative error analysis of subsets of most of the benchmarks described in Section 3. The review did not include *RRN*, *MIM-GOLD-NER*, *IceLinguistic* and *Hotter and Colder Sentiment*. The first two for technical reasons and the other two to be impartial in our judgments, as some of the authors participated in their creation and we acknowledge the drawback of that. The *Inflection* benchmark should have been included,

⁷huggingface.co/datasets/vesteinn/icelandic-qa-NQiI

⁸huggingface.co/datasets/alexandrainst/multi-wiki-qa

⁹huggingface.co/datasets/alexandrainst/m_mmlu

¹⁰huggingface.co/datasets/alexandrainst/m_hellaswag/viewer/isl

¹¹huggingface.co/datasets/thors/RRN

¹²github.com/EuroEval/EuroEval

Benchmark	Samples	Total size	Synth	IRR	CI	p-value
ARC-IS (EuroEval)	112	1,116	x	0.39	(0.28, 0.51)	< .001
ScaLA-IS	102	1,024		0.52	(0.36, 0.67)	< .001
IceEC	250	58,239		0.63	(0.55, 0.71)	< .001
IcelandicQA	200	2,000		0.54	(0.39, 0.68)	< .001
Belebele-IS	90	900		0.49	(0.36, 0.63)	< .001
MultiWikiQA-IS	250	5,004	x	0.71	(0.66, 0.77)	< .001
ARC-challenge-IS (Miðeind)	112	1,119	x	0.61	(0.50, 0.72)	< .001
GED	20	200		0.74	(0.14, 1)	0.018
NQil	213	2,132		0.73	(0.66, 0.79)	< .001
MMLU-IS	28	277	x	0.31	(0.01, 0.62)	0.044
WinoGrande-IS	109	1,088		0.54	(0.36, 0.72)	< .001
HellaSwag-IS	250	9,368	x	0.22	(0.12, 0.32)	< .001

Table 1: Overview of benchmarks, sample size, total size, inter-rater reliability (IRR) scores along with their confidence intervals and p-values

but simply was not at this stage, which is unfortunate. Furthermore, we did not include *Icelandic-Knowledge*, EuroEval’s version of the *IcelandicQA* dataset, that has been turned into a multiple-choice question-answering task and three new answer options are created by an LLM. As mentioned, the benchmark has not been published in this format and we did not make use of EuroEval’s Python package to create our own version.

To keep the sample sizes of the datasets within feasible boundaries, we used the following approach: If the dataset was longer than 250 rows, we took 250 randomly selected examples, else randomly selected 10%. Where different splits were available, we used the train split, and for *NQil* we only extracted rows that had an answer, like EuroEval, as rows with no answer are only meant to serve as negative examples in model training.

The samples were annotated by three of the authors of this paper, all of whom have a background in both Icelandic linguistics and NLP, and we used the following three labels and guidelines:

- IC: one of the following: incorrect (e.g. wrong answer to a question), severely flawed, clearly/poorly machine-translated
- F: flawed example, contains errors (spelling, grammaticality, vocabulary)
- OK: Valid example

In the case of the question-answering benchmarks, *IcelandicQA*, *NQil* and *MultiWikiQA-IS*, we focused on the QA pairs in our error annotation and only considered the context to verify validity, i.e. the context itself is not evaluated. We admit that our labels could have been more fine-grained with regards to error types to perhaps give some kind of estimate of their effect on the evaluation. Although the IC-label was intended to cover all examples that would potentially affect the evaluation validity there is a clear difference between factually wrong answers and very poor machine translations. The former will certainly affect the evaluation whereas an LLM might still be able to “make sense” of the latter and provide the correct answer.

An overview of the sample size, original length and the inter-rater reliability (IRR) with confidence intervals is shown in Table 1, where the x-label in the column ‘synth’ indicates whether the benchmark was either machine-translated or created by an LLM. Note that *IcelandicQA* does not have such a label as the model output was manually verified and corrected where necessary. For calculating the IRR, we use Krippendorff’s alpha and weighted agreement coefficients (ordinal) to account for the ordinal nature of the ratings, i.e. IC > F > OK.

For each of the benchmarks we have three annotations and overall the level of agreement is *good* (following Regier et al. (2013)), *questionable* for ARC-IS (EE), MMLU-IS and HellaSwag-is, and *very good* for MultiWikiQA-is, ARC-Challenge-is, GED and NQil. We note that the low IRR score in HellaSwag-IS is only due to disagreement between the IC and F labels as the annotated sample has no valid examples. This has, however, not been analyzed any further and the IRR values should be taken into account for our interpretation of the results.

4.1. Results

The results of our error annotation are shown in Figure 1, where the mean proportion of each label across all raters is shown for each benchmark with a 95% confidence interval. In Appendix A we show the label proportions for each rater across the benchmarks.

First, we observe a clear difference in benchmarks according to whether or not humans have been involved in their creation, translation or verification. Based on our samples, we find *Hellaswag* to be almost entirely severely flawed and *MMLU* very flawed, with less than 10% of the sample being valid (admittedly, the sample size for *MMLU-IS* could have been bigger and the inter-rater agreement is questionable based on our interpretation). Furthermore, we find that the machine-translated version of ARC that EuroEval uses (ARC-is) has an OK proportion of only 20%, whereas the localized version *Miðeind* uses (ARC-Challenge-is) is better,

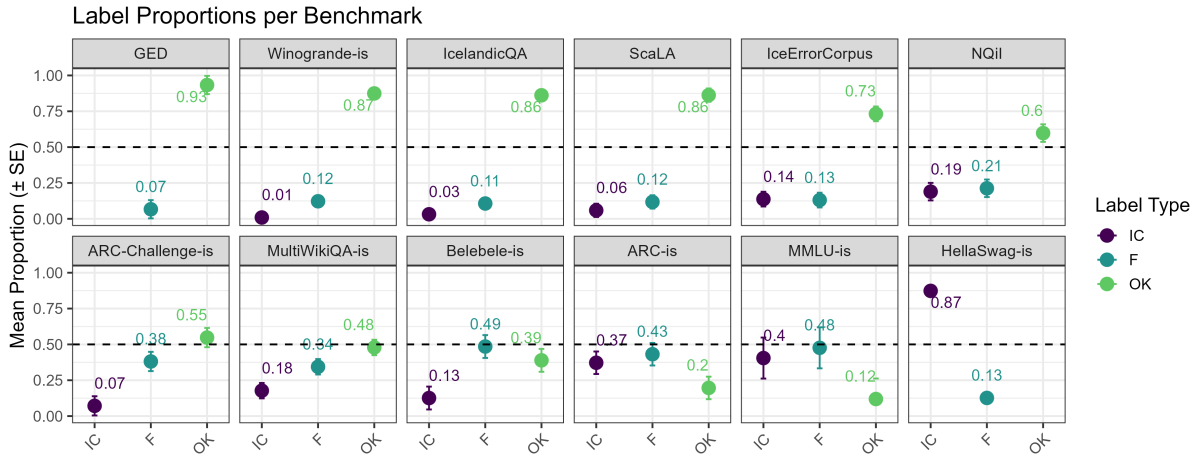


Figure 1: Mean proportion and 95% confidence interval for each of the labels across annotations and benchmarks, arranged by OK proportion in descending order

with a much lower IC-proportion (7% compared to 37%) and a higher OK-proportion (55%).

For *MultiWikiQA-IS*, less than half of the examples in our sample receive the OK-label. Our error labels admittedly make no distinction between factual and linguistic errors, but both types of flaws can affect the benchmark’s validity. In future evaluations, we aim to distinguish between these two types of flaws although they can sometime appear in conjunction. The following example illustrates the type of problematic items in *MultiWikiQA-IS*: An article about the Icelandic philosopher Salvör Nordal is presented as context and the question is (in our translation): “Which married couple had the son Sigurjón Nordal?” The male name, Sigurjón, does not appear in the article and the only Icelander to ever have had that name is completely irrelevant to the context, but the answer to the question is Salvör’s parents. In contrast, IcelandicQA, where QA pairs generated by an LLM were manually verified and corrected, is much better, although it contains validity errors as well.

Surprisingly, the Icelandic version of Belebele has a rather low proportion of OK examples. In the paper by Bandarkar et al. (2024), the authors state that it “was created end-to-end without the use of machine translation technology, relying solely on experts fluent in English and the target language.” We have, however, reasons to believe that the translators involved in translating the data into Icelandic used MT services, at least to some extent. One blatant example is an item where one of the answer options, which only include geopolitical entities, is the Icelandic word *kalkúnn* which translates to the animal turkey, not the country, in English. Here, *Tyrkland* would have been the correct translation and this answer option also happens to be

the correct one. No fluent expert in English and Icelandic would make such a mistake. A good contrast to these machine-translated benchmarks is the Icelandic version of the WinoGrande benchmark (Snæbjarnarson et al., 2022), which was translated by humans and localized. Based on our sample, it contains no severely flawed examples.

Regarding the results for NQil (Snæbjarnarson and Einarsson, 2022), it is important to keep in mind that this is a case of a dataset being turned into a benchmark despite this not being the intended use – it was originally created for open domain QA retrieval. As it is being used as a benchmark, however, we evaluate it as such. Based on our sample, there unfortunately seem to be some validity errors in the QA pairs. In one example, context about Kim Jong-Un is paired with a question about the capital of Bolivia, whereas the answer is related to the Kim Jong-Un (nqii_3570). In another example, a badly-constructed pair consists of a strange question about when tap water in Iceland was first boiled and the answer is that ‘it boils at 100°C’ (nqii_7841), where the context is an article about water in general (we refer to Section 5 for more discussion).

5. Discussion

We would like to begin by stressing that the goal is not to create flawless benchmarks, but to improve current practices, both by revising and re-releasing existing benchmarks where possible and by designing stronger future ones. In light of our findings, we argue that it is important to change the way benchmarking is approached in lower-resource settings and we will start by addressing the use of MT.

5.1. Machine-translated benchmarks

In general, we find it questionable to use machine-translated versions of foreign/non-localized benchmarks, especially in low/medium-resource settings and even more so if the output is not verified or corrected. Such an approach immediately raises several questions: 1) How was the MT model chosen? 2) Is the quality of the output evaluated in any way? 3) How suitable is the original benchmark for benchmarking another language?

To start with the first question, one point of reference could be findings of WMT, the largest conference in the MT community, where Icelandic has been a part of the general shared MT task in the last two years. In 2024 (Kocmi et al.), human translators achieved the highest score for EN→IS with 93.1 on a 0–100 human evaluation scale (with 0 being ‘no meaning preserved’ and 100 being that the meaning and grammar is completely consistent with the source). The best performing ML system, Dubformer (a private translation service), scored 84.9, followed by Claude 3.5 Sonnet at 81.9. Incidentally, two different translations by Claude 3.5 Sonnet of the ARC-Challenge are being used for Icelandic evaluation by Miðeind and EuroEval. This raises the question of whether those translations can, out-of-the-box, be expected to be of better quality than 81.9%, and whether that level is sufficient for benchmark construction.

As for the second question, the only MT quality assessment we identified for these benchmarks is the one conducted by Smart (2025) for MultiWikiQA. In that study, crowd workers rated the fluency of a small sample of translated items on a 1–3 scale, with Icelandic (n=150) scoring ~2.68/3. Note that the fluency was meant to be evaluated *regardless* of whether the question was “unanswerable or require[d] context to be answered”. We refer to Section 4 for our evaluation, but at this point we would like to note that fluency evaluation alone, without any measure of adequacy or accuracy, is unlikely to give a comprehensive assessment of the dataset.

Regarding the third question, we use the ARC-Challenge as an example. The original benchmark consists of human-authored, grade-school science questions from the US. One might ask how pertinent it is to evaluate LLM performance in Icelandic on it in the first place, regardless of whether it is human- or machine-translated, just as it would be questionable to make Icelandic grade-school pupils answer the same questions without accounting for cultural differences and some localization of content. To take a specific example from the original benchmark, in item Mercury_7084648 the question is about the causes of low precipitation in Nevada. It is unclear how culturally or educationally relevant such a question is in Icelandic context, highlighting the broader issue of contextual misalignment

in translated benchmarks.

As far as localization is concerned, in Miðeind’s version, the MT model seems to have translated most names that occur in the benchmark – presumably this was included in instructions to an LLM. At a quick glance, this seems to be a valid localization method, but it has some unintended side-effects as example 1 in Table 2 shows, where the name Louis Pasteur is erroneously changed to *Louis Guðmundsson*, removing all relation to the real person, which is of relevance for the question.

Another example of the shortcomings of this approach are items 2a and 2b in Table 2. In Miðeind’s version, *bi-flagellate* is translated as *tvíflagella*, which is not an Icelandic word (*tvísvipa* is correct), and *autotroph* should be *frumbjarga lífvera*, rendering the translation incorrect. EuroEval’s version is only half-translated and the Icelandic translation of *protist* is wrong. These examples show how the MT can fail on scientific terminology, which is unfortunate for a benchmark based on science questions. In light of this, can it be said that performance on such data meaningfully reflects a model’s Icelandic reading comprehension abilities, if that is what the benchmark is intended to measure? Furthermore, it is concerning that the output has seemingly not been filtered prior to release. In EuroEval’s version, one question (ID Mercury_401011) includes the Icelandic word *gös* ‘gases’ repeated 399 times, nothing else.¹³

Another example of a machine-translated benchmark is *HellaSwag-IS*, excluded by Miðeind and marked “unofficial” by EuroEval. Even a cursory inspection reveals to any native speaker how poor the quality of translation is. Our quantitative error analysis (Section 4) confirms that the Icelandic version is severely flawed, and we argue that it should be discarded. The IRR score is admittedly low, but not a single item was labeled as valid by any of the annotators, which is telling for its quality.

5.2. No native speaker involvement

Another concern arises when benchmarks are developed without the involvement of native speakers of the target language. One such example is the ScaLA benchmark (Nielsen, 2023), which covers the Nordic languages as well as German, Dutch and English. ScaLA measures linguistic acceptability through binary judgments of sentences that are either in their original form (“correct” sentences) or artificially corrupted, either by deleting a word or by swapping two adjacent words. The data is from

¹³In the EuroEval code there are automatic checks in place to not include “overly repetitive samples” (and samples with overly short or long texts) but we are referring to the published data.

Ref	Original	MT / Our translation
1	Louis Pasteur created a process that reduced the amount of bacteria in milk. How does this process most likely benefit people?	M: Louis Guðmundsson fann upp aðferð sem dró úr fjölda baktería í mjólk. Hvernig nýtist þessi aðferð fólki líklega best?
2a	Which protist is a bi-flagellate autotroph?	M: Hvaða frumvera er tvíflagella sjálfnærandi lífvera?
2b	-	EE: Hvaða mótefni er bi-flagellate autotroph?
3	Hvers vegna erum við ekki með lögin þannig að hver sem opinberlega hæðist að, rógber, smánar og ógnar manni og svo framvegis án tillits Forseti hringir. til nokkurs skapaðar hlutar sæti refsingu?	<i>Why are our laws not written such that anyone who publicly mocks, defames, humiliates and threatens someone etc., with no regard to Speaker rings. anything at all, is punished?</i>
4	Varðandi smábátahlutfallið í þessu frumvarpi verð ég að segja ég hef ekki skoðað skiptinguna á því mjög „grundlgt“.	<i>Regarding the proportion of small fishing boats in this bill, I have to say that I have not looked into the distribution very thoroughly.</i>

Table 2: Examples of flaws in Icelandic benchmarks. For space reasons we use the same column for either machine-translated examples from the benchmarks or, in italics, our translation of Icelandic examples for the foreign reader. (**M:** Miðeind, **EE:** EuroEval.)

the Universal Dependencies framework¹⁴ and the authors assume the original sentences are grammatically correct. Although they describe steps taken to ensure that the corrupted sentences are indeed ungrammatical, this verification was conducted only on a random sample of Danish data.

The problems with the Icelandic subset of this benchmark are twofold. The former is related to peculiarities in the data, of which the creators of ScaLA were probably unaware. Most of the data consists of transcribed speeches from the Icelandic parliament. While the speeches are proofread and produced in a formal setting, it is still a word-for-word transcription of spoken language and not written text. While not necessarily ungrammatical, they often include sentences that are perhaps not ideal for such a test.

Moreover, the data includes procedural insertions such as a) *Forseti hringir.* ‘Speaker rings’ and b) *Gripíð fram í:* ‘Interruption’, which mark, respectively, a) the moment when the speaker of the house rings a bell to indicate that the allocated time to an MP is up, and b) when someone interrupts a speech *ex auditorio*. An example of a) is shown in 3 in Table 2. Although the sentence is labeled as correct, such insertions disrupt the syntactic structure of the surrounding speech. There are unfortunately also two instances in the benchmarks where these two phrases (a and b) are themselves subjected to artificial corruption through word-order reversal (*hringir Forseti.* and *fram Gripíð í:*), which makes little sense as the phrases themselves are already corruptions of the sentence.

The second problem is that in some cases the Icelandic sentence is still grammatically correct after corruption. One such example is number 4 in Table 2. In the corrupted version, the noun in boldface is deleted, but the resulting sentence is just as correct and grammatical as the original. Lastly, there are examples of original sentences being grammatically wrong, resulting in erroneous labels.

A related issue concerns EuroEval’s use of the

¹⁴<https://universaldependencies.org>

Icelandic Error Corpus (Ingason et al., 2021) for binary grammatical error detection (GED). The corpus consists of proofread texts from various sources, with corrections made independently by different proofreaders, unfortunately resulting in some inconsistencies. Errors were classified in five different categories: *coherence, grammar, orthography, style, vocabulary*. EuroEval, however, samples the data without regard to error type and treats the ‘has_error’ label as an indicator of the grammaticality of a sentence. This is problematic. A stylistic correction – of which there 39 examples in our sample – hardly changes an ungrammatical sentence to a grammatical one or vice versa. For GED purposes, filtering the data by relevant categories such as grammar or orthography would be more appropriate, which Miðeind seems to have done in their use of the same data. Similarly, there are instances where there is an error in a sentence that is supposed to be error-free, which again leads to erroneous labels, as well as errors in corrected sentences, but the latter are not used by EuroEval so they have no bearing on the benchmarking.

These deficiencies in both benchmarks do not mean that they are “bad”, but simply show how their creation or use as benchmarks could have been improved with native speaker involvement and proper validation of the data for these purposes. We would like to add that in both cases, the data that makes up these benchmarks was collected or created by native speakers but it is their use as benchmarks that is under scrutiny.

5.3. MT without native speaker involvement

Although it is not a part of public leaderboards for benchmarking in Icelandic, a recent benchmark illustrates similar concerns, namely *GAM-BIT+*: *A Challenge Set for Evaluating Gender Bias in Machine Translation Quality Estimation Metrics* by Filandrianos et al. (2025). There, Icelandic is used as one example of a grammatically gendered language (along with 10 others). For each

English source text with gender-ambiguous occupation terms, two target language translations are created, intended to differ “only in the grammatical gender of the occupation and the dependent grammatical elements”. If, for instance, the English source sentence includes the word *professor*, the two target sentences should differ only in the gender of the translation of that word (and its dependent elements). The assumption is therefore that this is possible in the languages involved, but that is simply not the case for Icelandic, as such minimal gender pairs in occupation terms are few (see e.g. [Friðriksdóttir and Einarsson \(2024\)](#) where 381 out of 394 occupation terms examined were masculine, eight feminine and five neuter). The word *professor* is used in the example taken in the paper to demonstrate the approach for all the languages included. In Italian they use *il professore* : *la professoressa*, in Serbian *profesor* : *profesorica*, etc. But in the Icelandic example, the sole difference lies in the use of the titles *Sir* (icel. *Herra*) and *Mrs.* (icel. *Frú*) with the noun (“*Herra/Frú prófessorinn*”), as no feminine version of the word exists. These titles are, however, very rarely used in modern Icelandic and of course do not constitute a modification of the grammatical gender of the target noun.

In Icelandic, masculine occupation nouns are generally used irrespective of the referent’s gender, though some feminine counterparts exist in compounds, such as *forstöðumaður* : *forstöðukona* ‘director man/woman (of e.g. an institute)’. In the dataset, however, many minimal pairs rely on wrong or hallucinated feminine forms. Examples of this include *rithöfundarins* : *rithöfundarkonunnar*, hallucinated feminine term for *writer*; *puntari* : *puntara*, hallucinated translation of *prompter* (in a theater) regardless of gender; and *kjarnorkuverskjörinn* : *kjarnorkuverskjörinn (fem.)*, where the translation of the occupation itself (*nuclear power plant operator*) is hallucinated, the words are identical, and the feminine gender is added in parenthesis (*fem.*) to indicate a gender difference that does not exist.

Third, it is important to note that masculine agreement in pronouns, adjectives etc. with masculine nouns is grammatically correct in Icelandic regardless of the referent’s gender. Feminine agreement may, however, be used to specifically express the feminine gender of the person (and similarly, the neuter, for a non-binary person). There are some Icelandic examples in the benchmark of this form, i.e. where the only difference lies in agreement with the noun but not the noun itself. However, as the authors say, the target translations were “constructed to be semantically and lexically identical, differing only in the gender marking of the occupation and the dependent grammatical elements.” It thus seems that examples with agreement only in dependent grammatical elements and not in the

gender marking of the occupation do not align with the authors’ intention, further illustrating how ill-suited Icelandic is for this approach.

The translations were generated with Claude Sonnet 3.5 and verified with Claude Sonnet 3.7 (same family as the translation model) to ensure that the gender of the occupation was the only difference between the two target sentences and Icelandic achieved the lowest accuracy among the included languages (86.75%) in this validation. On the face of the arguments presented here, we believe that the quality of the Icelandic data is actually lower than this percentage indicates. A more rigorous and quantitative evaluation of the benchmark is needed to verify this, but it is already clear that the methodology can lead to serious flaws in the benchmark. We would lastly like to note that this dataset is intended to evaluate gender bias in MT and the authors claim to reveal “statistically significant score shifts driven solely by grammatical gender [...] with a clear overall tendency to favor masculine forms”. As far as Icelandic is concerned, we wonder to which extent these results are themselves biased by enforced and hallucinated feminine forms as in the examples above.

5.4. Typos, errors and whether they matter

LLMs are robust to linguistic errors, so why be finicky about that? The answer depends on the purpose and nature of the benchmark. Consider the NQil dataset, where *natural questions* are meant to indicate naturally occurring questions, such as a user might ask an LLM or search engines. In practice, such questions can and do often contain errors that do not necessarily prevent the system from correctly interpreting or answering the question. This can also be useful as training data for models, to expose them to a wider variety of language than only the standard. However, our focus is not on training utility but on the use of such datasets as benchmarks. In cases like this, with QA benchmarks, we are more concerned with validity errors, as in the examples in Section 4.1, and not natural linguistic variation by humans. On the other hand, the machine-translated versions of ARC originate from carefully constructed US grade-school science questions, presumably well-formed and appropriate for the task in English. This raises a different concern, whether it is acceptable to evaluate models on translated versions that introduce linguistic distortions through “translationese” or outright errors which no native speaker would accept.

A broader question, which we leave for future work, is to what extent “translationese”, LLM hallucinations or flaws in benchmarks matter. Someone might contend that the most of the benchmarks

involved in this paper, regardless of their flaws, are sufficient to show differences in model performance. But to what extent do the flaws affect the outcome? And does it matter if they do? As an example, we can think of a benchmark that has 80% valid examples and 20% severely flawed. Given a large enough sample size for sufficient statistical power, are the 20% “acceptable” if the rest shows clear difference in model performance? We argue that, regardless of short-term performance comparisons, flawed benchmarks can be harmful if incorporated into optimization pipelines, thereby influencing model development (see Section 2). In general, we believe the goal should always be to design benchmarks which are right for the right reasons.

With increased use of automated methods, ever larger benchmarks can be built and with increasing size, it becomes increasingly unlikely that any human will ever read through all of the examples for verification. The resulting loss of oversight makes systematic quality control even more essential, even if it is only via stratified random sampling.

5.5. Beyond single-label output, towards more diverse benchmarks

As mentioned in Section 1, there is demand for creating new benchmarks that are better adapted to the capabilities of LLMs and go beyond evaluating single-label output. Most of the benchmarks involved in our evaluation are indeed such benchmarks, where the answer is either True/False or an answer option in a multiple-choice task, with only the QA benchmarks as an exception.

We also note that, overall, the benchmarks included in our evaluation are quite similar. *GED*, *ScaLA* and *IceErrorCorpus* all measure very similar phenomena with binary outcomes, raising questions about the added value of reporting them separately. Similarly, the three QA benchmarks are all Wikipedia-based, although there is not necessarily overlap between them, and MultiWikiQA-IS was created despite the existence of *IcelandicQA*, that was created in a similar way but manually verified and corrected, which MultiWikiQA was not. There are two different versions of ARC, which are reading comprehension benchmarks just like *Belebele* and *MMLU*. *WinoGrande-IS* targets commonsense reasoning, while *HellaSwag-IS*, as we argued above, should likely be discarded. It would therefore be of benefit to try and broaden the spectrum of benchmarks for Icelandic and include more diverse tests, for instance with respect to the text generation capabilities of LLMs. This is to some extent done in three of the benchmarks that were not included in our evaluation, *Inflection*, the *Icelandic Linguistic Benchmark* and *RRN* (see Section 3).

6. Conclusion

In light of our findings, we would like to summarize a few key points to which attention should be paid in our opinion when benchmarking low/medium-resource languages:

- Avoid use of machine-translated benchmarks, especially when MT quality is still not sufficiently fluent, and/or when the output (or at least a sample of it) has not been verified and corrected
- Using an LLM to generate data can be tantamount to machine-translating data in terms of text quality and the same caveats apply
- Involve native speakers and prioritize the use of native-authored data

As already outlined, our results show clear differences with regards to native speaker involvement compared to machine-translated data where the output has not been verified or corrected, and it is clear that the involvement of native speakers and native data makes a difference for the better. In the Icelandic context, further opportunities exist to develop native benchmarks, for example by drawing on human-authored and validated materials from the school system (*à la* the ARC Challenge and many more).

As some of the authors of this paper have themselves been involved in the creation of benchmarks for Icelandic, we acknowledge that building good benchmarks is a complicated task and that we are not exempt from criticism of the problems outlined in the overall context of this paper. It is our hope that this can be a contribution to a healthy debate on how to proceed and what to avoid. We emphasize the need for native-authored, community-vetted, and contamination-audited benchmarks, ideally developed through participatory design with native speakers. This approach ensures that evaluation metrics genuinely reflect linguistic and cultural specificities of the language, preventing models from optimizing for superficial patterns introduced by translation artifacts or data leakage.

Lastly, we return to the question posed in the title of this paper: *who benchmarks the benchmarks?* The responsibility must lie with their creators. Benchmark construction cannot be treated as a neutral or one-time contribution, it requires continuous scrutiny, documentation, and critical evaluation. The work presented here originated from a simple review of existing Icelandic datasets, where repeated irregularities and inconsistencies prompted the systematic analysis reported in this paper. Such examination should not be incidental, it should be standard practice.

7. Limitations

We first recognize that the annotation may have been biased by the overall topic of this paper, i.e. flaws in benchmarking, as well as for benchmarks where the annotators knew that they were machine-translated or synthetic. Secondly, the sample size could have been bigger, especially for the 10% samples of benchmarks. Third, in three cases, the level of agreement was only questionable which must be taken into account. Fourth, the interpretation of the border between the error labels ‘IC’ (severely flawed) and ‘F’ (flawed) may have varied from one annotator to the other. Furthermore, the labels do not necessarily reflect whether the error is on the factual or linguistic level (although ‘IC’ was meant for examples that are flawed to such an extent that it potentially affects the results validity). Moreover, as mentioned in Section 4, two benchmarks that are included in leaderboards for Icelandic were left out in the evaluation for impartiality reasons. We concede that this is a drawback and they should preferably be evaluated as well. Finally, our study evaluates benchmark quality rather than directly quantifying the downstream impact of these flaws on model rankings or comparative conclusions. Future work should investigate to what extent the identified issues materially alter evaluation outcomes.

8. Acknowledgements

We would like to thank our reviewers for useful comments and feedback as well as those who gave their opinion on the paper in earlier versions. This project was in part funded by the Language Technology Programme for Icelandic 2024—2026. The programme, which is managed and coordinated by Almennarómur, is funded by the Icelandic Ministry of Education, Science and Culture. Steinunn Rut Friðriksdóttir was supported by the Eimskip Fund of The University of Iceland no. HEI2025-96428 and the Ludvig Storr Trust no. LSTORR2023-93030.

9. Bibliographical References

Ahmed Alaa, Thomas Hartvigsen, Niloufar Golchini, Shiladitya Dutta, Frances Dean, Inioluwa Deborah Raji, and Travis Zack. 2025. [Position: Medical Large Language Model Benchmarks Should Prioritize Construct Validity](#). In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pages 80991–81004, Vancouver, Canada. PMLR.

Bjarki Ármannsson, Finnur Ágúst Ingimundarson, and Einar Freyr Sigurðsson. 2025. [An Icelandic](#)

[Linguistic Benchmark for Large Language Models](#). In *Proceedings of the Joint 25th Nordic Conference on Computational Linguistics and 11th Baltic Conference on Human Language Technologies (NoDaLiDa/Baltic-HLT 2025)*, pages 37–47, Tallinn, Estonia. University of Tartu Library.

Pórunn Arnardóttir, Elías Bjartur Einarsson, Garðar Ingvarsson Juto, Þorvaldur Páll Helgason, and Hafsteinn Einarsson. 2025. [WikiQA-IS: Assisted Benchmark Generation and Automated Evaluation of Icelandic Cultural Knowledge in LLMs](#). In *Proceedings of the Third Workshop on Resources and Representations for Under-Resourced Languages and Domains (RESOURCEFUL-2025)*, pages 64–73, Tallinn, Estonia. University of Tartu Library.

Pórunn Arnardóttir, Xindan Xu, Dagbjört Guðmundsdóttir, Lilja Björk Stefánsdóttir, and Anton Karl Ingason. 2021. [Creating an error corpus: Annotation and applicability](#). In *Proceedings of CLARIN 2021*, pages 59–63, Virtual Edition.

Lucas Bandarkar, Davis Liang, Benjamin Muller, Mikel Artetxe, Satya Narayan Shukla, Donald Husa, Naman Goyal, Abhinandan Krishnan, Luke Zettlemoyer, and Madian Khabsa. 2024. [The Belebele Benchmark: a Parallel Reading Comprehension Dataset in 122 Language Variants](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, page 749–775, Bangkok, Thailand. Association for Computational Linguistics.

Samuel R. Bowman and George Dahl. 2021. [What Will it Take to Fix Benchmarking in Natural Language Understanding?](#) In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4843–4855, Online. Association for Computational Linguistics.

Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language Models are Few-Shot Learners](#). In *Proceedings of the 34th International Conference on Neural Information Processing Systems, NIPS '20*, Red Hook, NY, USA. Curran Associates Inc.

- Dallas Card, Peter Henderson, Urvashi Khandelwal, Robin Jia, Kyle Mahowald, and Dan Jurafsky. 2020. [With Little Power Comes Great Responsibility](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9263–9274, Online. Association for Computational Linguistics.
- Pinzhen Chen, Simon Yu, Zhicheng Guo, and Barry Haddow. 2024. [Is It Good Data for Multilingual Instruction Tuning or Just Bad Multilingual Evaluation for Large Language Models?](#) In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 9706–9726, Miami, Florida, USA. Association for Computational Linguistics.
- Jonathan H. Clark, Eunsol Choi, Michael Collins, Dan Garrette, Tom Kwiatkowski, Vitaly Nikolaev, and Jennimaria Palomaki. 2020. [TyDi QA: A Benchmark for Information-Seeking Question Answering in Typologically Diverse Languages](#). *Transactions of the Association for Computational Linguistics*, 8:454–470.
- Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and Oyvind Tafjord. 2018. [Think you have Solved Question Answering? Try ARC, the AI2 Reasoning Challenge](#). *ArXiv*, abs/1803.05457.
- Lee J. Cronbach and Paul E. Meehl. 1955. [Construct validity in psychological tests](#). *Psychological Bulletin*, 52(4):281–302.
- Maria Eriksson, Erasmo Purificato, Arman Noroozian, João Vinagre, Guillaume Chaslot, Emilia Gomez, and David Fernández-Llorca. 2025. [AI Benchmarks: Interdisciplinary Issues and Policy Considerations](#). In *ICML Workshop on Technical AI Governance (TAIG)*.
- George Filandrianos, Orfeas Menis Mastromichalakis, Wafaa Mohammed, Giuseppe Attanasio, and Chrysoula Zerva. 2025. [GAMBIT: A Challenge Set for Evaluating Gender Bias in Machine Translation Quality Estimation Metrics](#). In *Proceedings of the Tenth Conference on Machine Translation*, pages 314–326, Suzhou, China. Association for Computational Linguistics.
- Steinunn Rut Friðriksdóttir and Hafsteinn Einarsson. 2024. [Gendered Grammar or Ingrained Bias? Exploring Gender Bias in Icelandic Language Models](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 7596–7610, Torino, Italia. ELRA and ICCL.
- Steinunn Rut Friðriksdóttir, Dan Saattrup Nielsen, and Hafsteinn Einarsson. 2025. [Hotter and Colder: A New Approach to Annotating Sentiment, Emotions, and Bias in Icelandic Blog Comments](#). In *Proceedings of the Joint 25th Nordic Conference on Computational Linguistics and 11th Baltic Conference on Human Language Technologies (NoDaLiDa/Baltic-HLT 2025)*, pages 181–191, Tallinn, Estonia. University of Tartu Library.
- Matt Gardner, Yoav Artzi, Victoria Basmov, Jonathan Berant, Ben Bogin, Sihao Chen, Pradeep Dasigi, Dheeru Dua, Yanai Elazar, Ananth Gottumukkala, Nitish Gupta, Hannaneh Hajishirzi, Gabriel Ilharco, Daniel Khoshdel, Kevin Lin, Jiangming Liu, Nelson F. Liu, Phoebe Mulcaire, Qiang Ning, Sameer Singh, Noah A. Smith, Sanjay Subramanian, Reut Tsarfay, Eric Wallace, Ally Zhang, and Ben Zhou. 2020. [Evaluating Models’ Local Decision Boundaries via Contrast Sets](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1307–1323, Online. Association for Computational Linguistics.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2021. [Measuring Massive Multitask Language Understanding](#). In *Proceedings of the International Conference on Learning Representations (ICLR)*, Vienna, Austria.
- Philip A. Huebner, Elinor Sulem, Fisher Cynthia, and Dan Roth. 2021. [BabyBERTa: Learning More Grammar With Small-Scale Child-Directed Language](#). In *Proceedings of the 25th Conference on Computational Natural Language Learning*, pages 624–646, Online. Association for Computational Linguistics.
- Svanhvít Lilja Ingólfssdóttir, Ásmundur Guðjónsson, and Hrafn Loftsson. 2020. [Named Entity Recognition for Icelandic: Annotated Corpus and Models](#). In *Statistical Language and Speech Processing, 8th International Conference*, pages 46–57, Cardiff, UK. Springer International Publishing.
- Jaap Jumelet, Leonie Weissweiler, Joakim Nivre, and Arianna Bisazza. 2025. [MultiBLiMP 1.0: A Massively Multilingual Benchmark of Linguistic Minimal Pairs](#). arXiv: 2504.02768.
- Tom Kocmi, Eleftherios Avramidis, Rachel Bawden, Ondřej Bojar, Anton Dvorkovich, Christian Federmann, Mark Fishel, Markus Freitag, Thamme Gowda, Roman Grundkiewicz, Barry Haddow, Marzena Karpinska, Philipp Koehn, Benjamin Marie, Christof Monz, Kenton Murray, Masaaki Nagata, Martin Popel, Maja Popović, Mariya Shmatova, Steinþór Steingrímsson, and Vilém Zouhar. 2024. [Findings of the WMT24 General Machine Translation Shared Task: The LLM Era](#)

- Is Here but MT Is Not Solved Yet. In *Proceedings of the Ninth Conference on Machine Translation*, pages 1–46, Miami, Florida, USA. Association for Computational Linguistics.
- Chaoqun Liu, Wenxuan Zhang, Jiahao Ying, Mahani Aljunied, Anh Tuan Luu, and Lidong Bing. 2025. [SeaExam and SeaBench: Benchmarking LLMs with Local Multilingual Questions in Southeast Asia](#). In *Findings of the Association for Computational Linguistics: NAACL 2025*, pages 6119–6136, Albuquerque, New Mexico. Association for Computational Linguistics.
- R. Thomas McCoy, Ellie Pavlick, and Tal Linzen. 2019. [Right for the Wrong Reasons: Diagnosing Syntactic Heuristics in Natural Language Inference](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3428–3448, Florence, Italy. Association for Computational Linguistics.
- 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.
- ChaeHun Park, Koanho Lee, Hyesu Lim, Jae-seok Kim, Junmo Park, Yu-Jung Heo, Du-Seong Chang, and Jaegul Choo. 2024. [Translation Deserves Better: Analyzing Translation Artifacts in Cross-lingual Visual Question Answering](#). In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 5193–5221, Bangkok, Thailand. Association for Computational Linguistics.
- Darrel A. Regier, William E. Narrow, Diana E. Clarke, Helena C. Kraemer, S. Janet Kuramoto, Emily A. Kuhl, and David J. Kupfer. 2013. [DSM-5 Field Trials in the United States and Canada, Part II: Test-Retest Reliability of Selected Categorical Diagnoses](#). *The American Journal of Psychiatry*, 170(1):59–70.
- Oscar Sainz, Jon Campos, Iker García-Ferrero, Julen Etxaniz, Oier Lopez de Lacalle, and Eneko Agirre. 2023. NLP evaluation in trouble: On the need to measure LLM data contamination for each benchmark. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 10776–10787.
- Keisuke Sakaguchi, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. 2021. [WinoGrande: an adversarial winograd schema challenge at scale](#). *Communications of the ACM*, 64(9):99–106.
- Kirill Semenov and Rico Sennrich. 2025. [Measuring the Effect of Disfluency in Multilingual Knowledge Probing Benchmarks](#). In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*.
- Chantal Shaib, Joe Barrow, Alexa Siu, Byron Wallace, and Ani Nenkova. 2024. [How Much Annotation is Needed to Compare Summarization Models?](#) In *Proceedings of the Third Workshop on Bridging Human–Computer Interaction and Natural Language Processing*, pages 51–59, Mexico City, Mexico. Association for Computational Linguistics.
- Dan Saattrup Smart. 2025. [MultiWikiQA: A Reading Comprehension Benchmark in 300+ Languages](#). arXiv:2509.04111.
- Vésteinn Snæbjarnarson and Hafsteinn Einarsson. 2022. [Natural Questions in Icelandic](#). In *Proceedings of the Thirteenth Language Resources and Evaluation Conference (LREC)*, pages 4488–4496, Marseille, France. European Language Resources Association.
- Vésteinn Snæbjarnarson, Haukur Barri Símonarson, Pétur Orri Ragnarsson, Svanhvít Lilja Ingólfssdóttir, Haukur Jónsson, Vilhjalmur Thorsteinsson, and Hafsteinn Einarsson. 2022. [A Warm Start and a Clean Crawled Corpus - A Recipe for Good Language Models](#). In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 4356–4366, Marseille, France. European Language Resources Association.
- Yixiao Song, Kalpesh Krishna, Rajesh Bhatt, and Mohit Iyer. 2022. [SLING: Sino Linguistic Evaluation of Large Language Models](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 4606–4634, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Pór Sverrisson and Hafsteinn Einarsson. 2023. [Abstractive Text Summarization for Icelandic](#). In *Proceedings of the 24th Nordic Conference on Computational Linguistics (NoDaLiDa)*, pages 17–31, Tórshavn, Faroe Islands. University of Tartu Library.
- Elif Ecem Umutlu, Ayse Aysu Cengiz, Ahmet Kaan Sever, Seyma Erdem, Burak Aytan, Busra Tufan, Abdullah Topraksoy, Esra Darıcı, and Cagri Toraman. 2025. [Evaluating the Quality of Benchmark Datasets for Low-Resource Languages: A Case Study on Turkish](#). In *Proceedings of the Fourth Workshop on Generation, Evaluation and Metrics (GEM²)*, pages 471–487, Vienna, Austria and virtual meeting. Association for Computational Linguistics.

- Ashmal Vayani, Dinura Dissanayake, Hasindri Watawana, Noor Ahsan, Nevasini Sasikumar, Omkar Thawakar, Henok Biadgign Ademteu, Yahya Hmaiti, Amandeep Kumar, Kartik Kuckreja, Mykola Maslych, Wafa Al Ghallabi, Mikhail Mihaylov, Chao Qin, Abdelrahman M. Shaker, Mike Zhang, Mahardika Krisna Ihsani, Amiel Esplana, Monil Gokani, Shachar Mirkin, Harsh Singh, Ashay Srivastava, Endre Hamerlik, Fathinah Asma Izzati, Fadillah Adamsyah Maani, Sebastian Cavada, Jenny Chim, Rohit Gupta, Sanjay Manjunath, Kamila Zhumakhanova, Feno Heriniaina Rabevohitra, Azril Amirudin, Muhammad Ridzuan, Daniya Kareem, Ketan More, Kunyang Li, Pramesh Shakya, Muhammad Saad, Amirpouya Ghasemaghaei, Amirbek Djanibekov, Dilshod Azizov, Branislava Jankovic, Naman Bhatia, Alvaro Cabrera, Johan S. Obando-Ceron, Olympiah Otieno, Fabian Farestam, Muztoba Rabhani, Sanoojan Baliah, Santosh Sanjeev, Abduragim Shtanchaev, Maheen Fatima, Thao Nguyen, Amrin Kareem, Toluwani Aremu, Nathan Xavier, Amit Bhatkal, Hawau Toyin, Aman Chadha, Hisham Cholakkal, Rao Muhammad Anwer, Michael Felsberg, Jorma Laaksonen, Tamar Solorio, Monojit Choudhury, Ivan Laptev, Mubarak Shah, Salman Khan, and Fahad Khan. 2024. [All languages matter: Evaluating llms on culturally diverse 100 languages](#). *CoRR*, abs/2411.16508.
- Héctor Javier Vázquez Martínez, Annika Heuser, Charles Yang, and Jordan Kodner. 2023. [Evaluating Neural Language Models as Cognitive Models of Language Acquisition](#). In *Proceedings of the 1st GenBench Workshop on (Benchmarking) Generalisation in NLP*, pages 48–64, Singapore. Association for Computational Linguistics.
- Bin Wang, Zhengyuan Liu, Xin Huang, Fangkai Jiao, Yang Ding, AiTi Aw, and Nancy Chen. 2024a. [SeaEval for Multilingual Foundation Models: From Cross-Lingual Alignment to Cultural Reasoning](#). 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 370–390.
- Wenxuan Wang, Wenxiang Jiao, Jingyuan Huang, Ruyi Dai, Jen-tse Huang, Zhaopeng Tu, and Michael Lyu. 2024b. [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.
- Alex Warstadt, Alicia Parrish, Haokun Liu, Anhad Mohananey, Wei Peng, Sheng-Fu Wang, and Samuel R. Bowman. 2020. [BLiMP: The Benchmark of Linguistic Minimal Pairs for English](#). *Transactions of the Association for Computational Linguistics*, 8:377–392.
- Linda Wiecheteck, Flammie Pirinen, Maja Lisa Kappfjell, Trond Trosterud, Børre Gaup, and Sjur Nørstebø Moshagen. 2024. [The Ethical Question – Use of Indigenous Corpora for Large Language Models](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 15922–15931, Torino, Italia. ELRA and ICCL.
- Ankit Yadav, Himanshu Beniwal, and Mayank Singh. 2024. [PythonSaga: Redefining the benchmark to evaluate code generating LLMs](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 17113–17126, Miami, Florida, USA. Association for Computational Linguistics.
- Shuo Yang, Wei-Lin Chiang, Lianmin Zheng, Joseph E. Gonzalez, and Ion Stoica. 2023. [Rethinking Benchmark and Contamination for Language Models with Rephrased Samples](#). arXiv:2311.04850.
- Binwei Yao, Ming Jiang, Tara Bobinac, Diyi Yang, and Junjie Hu. 2024. [Benchmarking Machine Translation with Cultural Awareness](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 13078–13096, Miami, Florida, USA. 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.

10. Language Resource References

- Anton Karl Ingason, Lilja Björk Stefánsdóttir, Þórunn Arnardóttir, and Xindan Xu. 2021. [Ice-landic error corpus \(IceEC\) version 1.1](#). CLARIN-IS.
- Svanhvít Lilja Ingólfssdóttir, Ásmundur Alma Guðjónsson, and Hrafn Loftsson. 2020. [MIM-GOLD-NER – named entity recognition corpus \(21.09\)](#). CLARIN-IS.

A. Appendix: Label Proportions per Rater

In Figure 2 we show the distribution of labels per rater across the benchmarks.



Figure 2: Label proportions per rater across the evaluated benchmarks