

Vuyko Mistral: Adapting LLMs for Low-Resource Dialectal Translation

Roman Kyslyi¹, Yuliia Maksymiuk², Ihor Pysmennyi¹

¹National Technical University of Ukraine "Igor Sikorsky Kyiv Polytechnic Institute"

²Ukrainian Catholic University

Correspondence: kyslyi.roman@lil.kpi.ua, yuliia.maksymyuk@ucu.edu.ua, pysmennyi.ihor@lil.kpi.ua

Abstract

In this paper we introduce the first effort to adapt large language models (LLMs) to the Ukrainian dialect (in our case Hutsul), a low-resource and morphologically complex dialect spoken in the Carpathian Highlands. We created a parallel corpus of 9852 dialect-to-standard Ukrainian sentence pairs and a dictionary of 7320 dialectal word mappings. We also addressed data shortage by proposing an advanced Retrieval-Augmented Generation (RAG) pipeline to generate synthetic parallel translation pairs, expanding the corpus with 52142 examples. We have fine-tuned multiple open-source LLMs using LoRA and evaluated them on a standard-to-dialect translation task, also comparing with few-shot GPT-4o translation. In the absence of human annotators, we adopt a multi-metric evaluation strategy combining BLEU, chrF++, TER, and LLM-based judgment (GPT-4o). The results show that even small(7B) finetuned models outperform zero-shot baselines such as GPT-4o across both automatic and LLM-evaluated metrics. All data, models, and code are publicly released at: <https://github.com/woters/vuyko-hutsul>.

1 Introduction

Despite recent advances in large language models (LLMs), most research and applications remain centered on high-resource languages and their standard variants (Li et al., 2024). This imbalance has significant consequences for linguistic diversity, particularly for underrepresented dialects that lack sufficient textual resources and standardized orthographies (Zhong et al., 2024). Despite being an integral part of the linguistic identity of many communities, dialects are often excluded from NLP tools and research, limiting their accessibility and risking further marginalization and extinction (Syed et al., 2023).

Language technologies and especially LLMs are playing a growing role in the preservation of en-

dangered and underrepresented languages. While much attention has focused on major indigenous languages (e.g., Māori, Quechua, Inuktitut) (Trudgill, 2003; Cooper et al., 2024), dialects of national languages are often overlooked despite facing similar pressures of attrition and assimilation. Dialectal variants, particularly in post-Soviet contexts, often carry suppressed cultural identities that are not reflected in the standard language. These dialects are not only linguistically rich but also culturally vital and deserve computational attention.

Ukrainian, a language low in resources according to global standards itself (Kiulian et al., 2024), exhibits rich internal variation, with dialects such as Hutsul, Boyko and Lemko¹ preserving unique phonetic, lexical and grammatical characteristics. Among these, the Hutsul dialect, spoken in the Carpathian Mountains, is one of the most linguistically distinct and has the most written sources. From the culture standpoint, Hutsul dialect has a great significance as it encapsulates traditions, folklore, and a unique worldview, playing a central role in community identity.

However, the lack of digitized corpora, dictionaries, and processing tools makes it practically invisible to modern LLMs.

Here are some of the linguistic Characteristics of Hutsul dialect:

- *Phonetics*: vowel transformations, such as changing vowels "e" instead of "a", "я"(ya) (example: "як" → "єк", "ягода" → "єгода" ("yak" → "yek", "yahoda" → "yehoda")).
- *Morphology*: unique case endings (-єдь, -сі) ('-yed', '-si') and preserved dual forms дві яблуці ("two apples", with dual form "yablutsi" instead of plural "yabluka").
- *Lexicon*: Romanian, Polish and German

¹<https://en.wikipedia.org/wiki/Hutsuls>

borrowings such as "бринза" (cheese) and "шпацірувати" (go for a walk).²

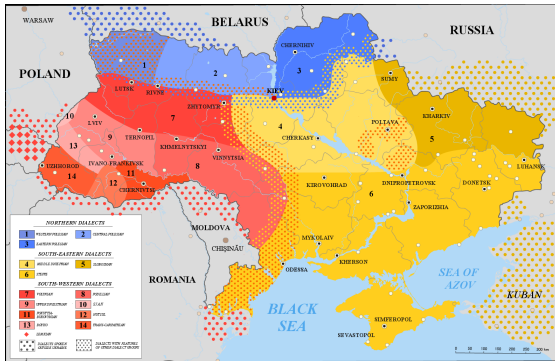


Figure 1: Map of Ukrainian dialects. The Hutsul dialect is located in the southwestern Carpathian region. Source: Wikipedia

In this work, we present an effort to adapt LLMs to the Hutsul dialect of Ukrainian, addressing both data shortage and modeling challenges. Our contributions are:

- A new parallel corpus of original Hutsul-Ukrainian (9852 sentence pairs), dictionary of 7320 dialectal word mappings and also synthetically extended corpus (52142 sentence pairs), using an advanced RAG approach (detailed described below).
- Fine-tuning of several open-source LLMs for Ukrainian to Hutsul dialect translation task.

We frame our task as standard-to-dialect translation, in which model has to take standard Ukrainian as input and produce grammatically correct (or as close as possible) Hutsul dialect. Our models show that it is feasible to address such translation with limited parallel data and targeted augmentation strategies.

To our knowledge, this is the first work that tries to adapt LLMs to a Ukrainian dialect and among the few globally addressing dialect-to-standard generation using synthetic augmentation.

2 Related Work

2.1 Dialectal NLP and Language Variation

In recent years we can see growing interest in dialect modeling, particularly for Arabic (Zampieri et al., 2017), German (Hollenstein et al., 2020), and

²https://en.wikipedia.org/wiki/Eastern_Romance_influence_on_Slavic_languages

Romance languages (Ramponi and Plank, 2021). These efforts mainly focus on classification, generation, and translation between dialects and their standard variants. However, most research remains concentrated on high-resource languages and dialects with pre-existing NLP resources. At the same time, within Ukrainian language, dialectal NLP remains underexplored. The VarDial workshop series (Zampieri et al., 2024) has supported work for different Slavic languages on related tasks such as cross-dialect machine translation and morphological modeling (Blokland et al., 2024; Kinn and Åfarli, 2024). For example, Kinn and Åfarli (2024) explore MT between Bokmål and Nynorsk, while Blokland et al. (2024) tackle dialectal variation in North Sámi. The SIGMORPHON 2023 shared task (Kirov et al., 2023) highlighted the importance of lexicon-based inflection modeling for low-resource morphological variants.

2.2 Dialect-to-Standard Normalization

The task of normalizing dialectal language to its standard form has been explored using various alignment techniques. Scherrer (2023) evaluated character alignment methods for sentence-level standardization of dialect transcriptions across Finnish, Norwegian, and Swiss German. The study compared approaches from dialectometry, speech processing, and machine translation, finding that trained alignment methods offered only small benefits over simple Levenshtein distance. This suggests that simple yet robust statistical methods may still provide strong baselines in resource-constrained dialectal settings. Moreover, the study underlines the need for tailored preprocessing and alignment tools when working with highly variable and phonetically rich dialect data.

2.3 LLMs and Dialect Adaptation

Several recent studies investigate adapting LLMs to dialectal data. Held and Klakow (2024) propose task-agnostic adapters for dialect adaptation, while Liu et al. (2024) introduce dynamic adapter aggregation based on linguistic distance. Tokenizer retrofitting for morphologically rich dialects is explored by Cs’aki et al. (2023). These works demonstrate that both architectural and data-centric interventions are necessary for effective adaptation. However, these approaches are primarily evaluated on English dialects (e.g., African American English, Indian English) using curated corpora such as Multi-VALUE (Lin et al., 2021), and rely on an-

notated dialect-to-standard pairs, which are rarely available for under-resourced dialects.

2.4 Low-Resource and Synthetic Data Techniques

Our work also benefits from previous research in low-resource translation and text generation with synthetic data. [Gudibande et al. \(2023\)](#) and [Garcia et al. \(2024\)](#) propose retrieval-based or prompt-based augmentation techniques to bootstrap performance in limited-data settings. At the same time we propose our own approach for generating synthetic data using advanced RAG techniques.

3 Dataset Creation

3.1 Parallel Corpus Collection

We constructed the first parallel corpus for the Hutsul dialect and standard Ukrainian by combining multiple sources and annotation strategies. The dataset includes 9852 sentence pairs, manually aligned at the sentence level. Source texts in Hutsul were collected from publicly available books, ethnographic transcripts, folklore websites, and dialect blogs. A significant portion of the dataset is based on the novel "Дідо Іванчик" (Dido Yvanchik) by Petro Shekeryk-Donykiv³, a foundational literary work written in authentic Hutsul. We are especially grateful to the publishing house Дискурс and translator Іван Андруссяк, who kindly approved the use of their modern standard Ukrainian translation for academic purposes.

Standard Ukrainian references in the dataset were either manually translated or sourced from bilingual editions where available. To ensure linguistic diversity, we tried to include examples from both everyday conversation and stylized narrative texts (e.g., folk tales, songs, etc.), but due to data shortage some topics remain uncovered.

3.2 Lexical Resource

We compiled a Hutsul-to-Ukrainian dictionary that now contains about 7 300 word pairs. The work started from the vocabulary that appears in the book "Дідо Іванчик" (Dido Yvanchik), but we soon enlarged it with data taken from websites that explain Hutsul dialect words. Among the most useful web sources were:

- "Dictionary of Hutsul Words"⁴.

³https://pl.wikipedia.org/wiki/Petro_Szekeryk-Donykiw

⁴<https://karnauhova.at.ua/publ/1-1-0-3>

- "Hutsul Hovir"⁵.
- "Dictionary of Ukrainian Dialects of the Carpathian Region"⁶.
- "Explanatory Dictionary of Hutsul Dialects" by Petro Havuka⁷.
- "Hutsul dictionary".⁸

All these pages were automatically scraped. The raw text contained a lot of noise: strange characters, extra commentary, uneven tabulation, and inconsistent separators between the Hutsul entry and its Ukrainian translation. We wrote simple cleaning scripts, converted everything to a single CSV file, and then manually checked the list to remove the last errors. The final result is a clean lexicon with 7 320 Hutsul–Ukrainian pairs. Each entry includes standard and dialectal word forms.

Despite this effort, the lexicon remains biased toward the vocabulary found in literature and folkloric domains. Due to the shortage of Hutsul texts on topics like news, science, or politics, our dataset lacks sufficient lexical diversity in those domains.

3.3 Synthetic Data via Advanced RAG

To overcome shortage of written sources in Hutsul dialect, we developed an advanced RAG pipeline to generate additional Hutsul-standard sentence pairs. The foundation of this pipeline was the dialectal novel "Дідо Іванчик" (Dido Yvanchik), which served as both the primary corpus for retrieval and the source of linguistic examples. We used GPT-4o to build a RAG module capable of retrieving semantically related Hutsul sentences. For each generation step, a prompt was created containing linguistic transformation rules representative of Hutsul phonological and lexical variation.

The construction of the RAG pipeline involved several steps:

1. **Grammar Rule Extraction:** Using "Дідо Іванчик" (Dido Yvanchik) as input, we prompted GPT-4o to extract and structure grammatical transformation rules characteristic of the Hutsul dialect. These included phonological shifts, morphological alternations, and syntactic reordering. We augmented these rules with material from Wikipedia and

⁵<https://rakhiv-mr.gov.ua/hutsulskyj-hovir/>

⁶<https://evrika.if.ua/88/>

⁷<https://evrika.if.ua/1565/>

⁸<http://www.webteka.com/hutsul-language/>

Models of Word Formation in Hutsul Dialects Greshchuk (2016) to create a comprehensive prompt template (see Figure 2).

2. **Indexing via RAG:** We indexed the "Дідо ІВАНЧІК"(Dido Yvanchik) corpus into a retrieval system to serve as a reference base for generating dialectal outputs using text-embedding-3-large⁹.
3. **Candidate Sentence Selection:** Standard Ukrainian sentences were sampled from the UberText corpus (Chaplynskyi (2023)). For each such sentence, we used the RAG module to retrieve the top-3 semantically similar Hutsul-like sentences from "Дідо ІВАНЧІК"(Dido Yvanchik).
4. **Prompt Construction:** The retrieved examples were inserted into the prompt template along with the standard Ukrainian sentence as the source for translation.
5. **Dialect Generation:** GPT-4o was instructed to produce a Hutsul translation of the input sentence using the provided grammar rules and examples as context (see Figure 3).

Below is a main part of our rule-based prompt (Full prompt can be found here: https://github.com/woters/vuyko-hutsul/blob/main/prompts/hutsul_rules_prompt.txt):

```
Here are Grammatical Rules for
Converting Ukrainian Text into the
Hutsul Dialect:
1. Vowel Shifts:
- "як" → "єк" ("yak" → "yek")
- "яблуко" → "єблуко" ("yabluko" →
"yeblyuko")
- "їдеш" → "єдеш" ("yidesh" → "yedesh")
2. Consonant Transformations:
- "дівка" → "гівка" ("divka" → "givka")
- "чого" → "цьо" ("choho" → "cho")
- "ти" → "ци" ("ty" → "tsy")

3. Word Order and Syntax:
- "Я тебе люблю" → "Люблю я тебе" ("I
love you" → "Love I you")
- "Він сміється" → "Він смієтси" ("He is
laughing" → "He laugh-reflexive")
- "Ти знаєш?" → "Ци ти знаєш?" ("Do you
know?" → "Do you know?" with dialectal
marker "tsy")

Apply only contextually appropriate
transformations.
```

⁹<https://platform.openai.com/docs/models/text-embedding-3-large>

This process have created some data alignment challenges in the generated dataset. To address these challenges and also to clean generated dataset we have developed a hybrid alignment strategy. First we leveraged the expected textual similarity between a language and its dialect using difflib's SequenceMatcher¹⁰. This approach directly compares character sequences, effectively identifying pairs even with minor dialectal variations. Pairs falling below a similarity threshold of 0.45 was removed from the dataset. To measure quality of remained sentence pairs we have used several statistical metrics as described by Scherrer (2023):

- **U-src** – proportion of unaligned source characters,
- **U-tgt** – proportion of unaligned target characters,
- **X** - proportion of crossing alignment pairs (swaps)

These metrics were calculated over symmetrized alignment pairs obtained with fast align(Dyer et al., 2013). We have compared alignment metrics across three datasets: the original manually annotated dataset (mainly from "Дідо ІВАНЧІК"(Dido Yvanchik)), raw synthetically generated dataset, and the filtered synthetic dataset.

Before filtering, the synthetic data already exhibited lower proportions of unaligned source and target words (U-src=0.139, U-tgt=0.136) compared to the original data (U-src=0.260, U-tgt=0.265). However, it presented a higher proportion of crossing alignments (X=0.033 vs. 0.022 original), indicating increased structural variability.

To improve the quality of our generated dataset, we applied alignment-based filtering - for each sentence pair, we have used previously calculated statistics(U-src, U-tgt and X) and we empirically defined a thresholds for them: $U\text{-src} < 0.1$, $U\text{-tgt} < 0.1$, and $X < 0.2$.

Any sentence pair that exceeded one or more of these thresholds was excluded from the final data set. This procedure removed inconsistent examples, reducing the number of reorderings, and improving alignment. As the result we got a better quality synthetic dataset with better structural alignment, as demonstrated by the comparative metrics in Table 1.

¹⁰<https://docs.python.org/3/library/difflib.html>

Metric	Original Dataset	Synthetic (Raw)	Synthetic (Filtered)
U-src	0.260	0.139	0.005
U-tgt	0.265	0.136	0.005
X	0.022	0.033	0.019

Table 1: Alignment quality metrics comparison between the original dataset, raw synthetic dataset, and synthetic dataset after alignment-based filtering.

Although we acknowledge that the obtained synthetic data has some variation and lack of certain lexical phrases present in authentic dialectal speech, its inclusion is justified by shortage of Hutsul textual resources. This filtering step effectively improved the consistency and reliability of the synthetic dataset and added additional 52142 phrase pairs to our training dataset.




Figure 2: Overview of the rules generation pipeline based on "Дідо Іванчiк"(Dido Yvanchik), Wikipedia, and Greshchuk (2016).

Although this approach enabled us to significantly enlarge the dataset, it also introduced certain limitations. Specifically, the synthetic data reflects the lexical and topical range of the source corpus, which lacks modern domains such as aviation, technology, news and politics.

As a result, lexical coverage in these areas remains quite sparse or absent (even after generation, words still remain the same as they are in standard Ukrainian). To avoid introducing hallucinated vocabulary, we deliberately excluded modern news and web-based corpora from the generation process.

3.4 Data Splits and Availability

The final corpus was split into 80% training, 10% validation, and 10% test sets. Test and validation sets contain only human-annotated sentence pairs from "Дідо Іванчiк"(Dido Yvanchik).

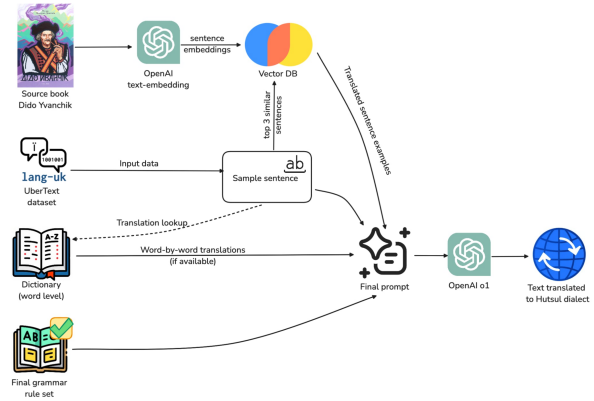


Figure 3: Overview of the synthetic data generation pipeline: A RAG system using "Дідо Іванчiк"(Dido Yvanchik) and UberCorpus retrieves and prompts GPT-4o to generate high-quality Hutsul-Ukrainian pairs.

4 Fine-Tuning

To adapt large language models (LLMs) to the Ukrainian-to-Hutsul translation task, we used parameter-efficient fine-tuning using LoRA (Hu et al., 2021).

We fine-tuned two state-of-the-art open-source models in the 7B–13B parameter range (as we considered our training resources and that the model should be not too big to be able to run locally):

- **Mistral-7B-Instruct v0.3¹¹** – Chosen for its performance-to-size ratio. It outperforms some larger models on many benchmarks, supports multilingual instructions, and includes explicit support for Ukrainian (AI, 2023).
- **LLaMA-3.1 8B Instruct¹²** – The instruction-tuned version of LLaMA 3.1 8B. This model has a strong multilingual support and improved instruction-following ability, making it a good candidate for low-resource translation (Touvron et al., 2024).

Models were selected based on the following criteria:

- **Tokenizer support** – Both models use tokenizers with fallback strategies for rare or out-of-vocabulary tokens, enabling good handling of Cyrillic-based dialects.

¹¹<https://huggingface.co/mistralai/Mistral-7B-Instruct-v0.3>

¹²<https://huggingface.co/meta-llama/Llama-3.1-8B-Instruct>

- Multilingual capabilities – Mistral-7B-Instruct v0.3 explicitly lists Ukrainian among supported languages. LLaMA-3.1 8B Instruct has shown strong generalization capabilities¹³.
- Open licensing and reproducibility – Both models are publicly available under open-source licenses.
- Feasibility on a single GPU – Using LoRA or QLoRA, all selected models can be fine-tuned within the a single not very big GPU.

We also considered other multilingual models such as BLOOMZ (7.1B)¹⁴ and NLLB-200 (3.3B)¹⁵, which offer extensive language coverage. However, these models either underperformed on general language modeling tasks or lacked strong generation quality compared to selected models. Recent benchmarks demonstrate that Mistral-7B-Instruct-v0.3¹⁶ matches or surpasses larger models in translation tasks, particularly in low-resource and instruction-tuned settings (Wu et al., 2023).

4.1 Fine-Tuning Setup

Each model was trained for 3 epochs using LoRa on two dataset variants (complete setup can be found in the Github¹⁷): (1) a manually created Hutsul–Ukrainian parallel corpus, and (2) an extended version that included combined manual and filtered synthetic data.

5 Evaluation

5.1 Metrics

Evaluating dialectal machine translation is not a simple task, as standard reference-based metrics may penalize correct lexical variation. To insure translation quality we calculated the following widely used metrics:

- **BLEU** (Papineni et al., 2002) - a precision-based metric measuring n-gram overlap between hypothesis and reference. While widely used, it may penalize valid lexical and syntactic variations common in dialects.

¹³<https://huggingface.co/blog/akjinda153244/llama31-storm8b>

¹⁴<https://huggingface.co/bigscience/bloomz-7b1>

¹⁵<https://huggingface.co/facebook/nllb-200-3.3B>

¹⁶<https://huggingface.co/mistralai/Mistral-7B-Instruct-v0.3>

¹⁷<https://github.com/woters/vuyko-hutsul>

- **chrF++** (Popović, 2015) computes character n-gram F-scores and has been shown to outperform BLEU on morphologically rich and non-standard languages. It is more robust to minor spelling or inflectional differences, making it particularly suitable for dialectal text.
- **TER** (Translation Edit Rate) (Snover et al., 2006) quantifies the number of edits required to convert the system output into the reference. It captures structural divergence and penalizes reordering errors.

Each metric emphasizes different aspects of translation quality:

- *BLEU* reflects n-gram precision,
- *chrF++* captures morphological similarity and recall,
- *TER* penalizes structural mismatches

We apply these metrics to test set of manually translated Ukrainian–Hutsul sentence pairs (1900 pairs). Rather than aggregating them into a single score, we interpret them jointly to understand different behavioral aspects of each model. For instance, a high chrF++ score alongside a low BLEU score may indicate valid variation in surface realization.

As mentioned before, while these metrics provide a useful baseline, they often struggle to evaluate dialectal outputs. So, following the framework of Aepli et al. (2023), we incorporate LLMs as evaluators.

We prompt GPT-4o model to rate model outputs along three axes:

- **Fluency**: grammaticality and naturalness in the Hutsul dialect.
- **Adequacy**: preservation of the source sentence’s meaning.
- **Dialectal Quality**: consistency with known lexical, phonological, and morphosyntactic properties of Hutsul.

Each evaluation is performed in a zero-shot setting. We thought about including some grammatical rules into the prompt, but to avoid creation of potential bias through this rules decided to use zero-shot instead.

LLM receives the source, model output, and a reference translation and returns scores from 1

(poor) to 5 (excellent). The prompt is structured as follows:

You are a linguistic expert evaluating machine-translated dialectal text. Rate the translation on the following dimensions:

1. Fluency (1–5): Is the output grammatically correct and natural in the target dialect?
2. Adequacy (1–5): Does the output preserve the meaning of the original source?
3. Dialectal Quality (1–5): Does the output reflect the expected phonological, lexical, and grammatical properties of the Hutsul dialect?

Return your answer in this exact JSON format:

```
{ "fluency": x, "adequacy": y, "dialect": z }
```

Do not explain your ratings.

Source (Standard Ukrainian): <source sentence>

Model Output (Hutsul): <model prediction>

Reference (Hutsul): <reference sentence>

As we didn't have an opportunity to perform a human evaluation for our translation, and considering that standard reference-based metrics may not be a good fit for the dialect translation (Aeppli et al., 2023), we have used the LLM-based adequacy and dialect scores as our primary evaluation metrics. These are better aligned with human intuition and more tolerant of different variations than BLEU or TER.

Automatic metrics are used in a supporting role to identify trends such as reordering or character-level similarity. We report all metrics side-by-side. This multi-metric approach enables a more holistic interpretation of model behavior, especially in the absence of human raters.

5.2 Baselines

We compared our fine-tuned models against the GPT-4o baseline. Queried via the OpenAI API, prompted to translate standard Ukrainian into the Hutsul dialect. To ensure consistency and lexical coverage, we used the same RAG context and

dictionary entries as in our synthetic generation pipeline.

We did not include non fine-tuned Mistral or LLaMA models as baselines, since their performance in dialect generation tasks was much worse. Due to their small size, their instruct tuning is insufficient for zero-shot generation in underrepresented languages or dialects.

5.3 Results

As mentioned earlier, we evaluate our models using both automatic metrics and LLM-based judgments. Table 2 presents the BLEU, chrF++, TER scores and GPT-4o as an LLM-based judge, rating each output on a 1–5 scale for fluency, adequacy, and dialectal quality scores computed on a held-out test set of 1900 sentences.

From the results we can see that all fine-tuned models outperform the GPT-4o baseline for every metric. Mistral fine-tuned on combined manually collected and synthetic data performs best overall, with the highest BLEU (74.35), chrF++ (81.89), and dialect rating (3.60). While adequacy scores remain stable across all models (≈ 4.7), dialectal accuracy varies more substantially and proves most sensitive to the source of training data. Also we can see that both, LLaMA and Mistral trained on combined synthetic and manually annotated data show strong scores on automatic metrics but slightly underperform on dialectal quality, highlighting the limitations of our method of generating synthetic data.

5.4 Qualitative Examples

Below we show an example depicting LLM-calculated scores over real data along with respective BLEU, chrF++, and TER metrics. This demonstrates that even small fine-tuned models are slightly better at preserving dialect-specific meaning and lexicon than zero-shot commercial models, but still far from perfect.

Reference (Hutsul): "Прошумавси у вечер, ек зробивси в діда в обох хатах гармідер."
(Eng: *He came to his senses in the evening, after a mess started in both of the grandfather's houses.*)

GPT-4o: "Отетавсі аж увечеру, ек зчинивсі в гіда в обидвох оседочьках гармідер."
(Eng: *He snapped out of it only in the evening, after a chaos started in both of the old man's*

Model	BLEU	chrF++	TER	Fluency	Adequacy	Dialect
GPT-4o	56.64	65.90	34.34	3.76	4.30	3.22
LLaMA (manual annotated + synthetic)	69.02	74.92	22.90	4.11	4.72	3.33
LLaMA (manual annotated only)	59.98	72.61	28.62	4.13	4.72	3.38
Mistral (manual annotated only)	62.36	75.65	28.62	4.14	4.74	3.35
Mistral (manual annotated + synthetic)	74.35	81.89	22.90	4.18	4.72	3.60

Table 2: Automatic and LLM-based evaluation results. BLEU, chrF++, and TER are computed with sacreBLEU. Fluency, adequacy, and dialect quality are rated by GPT-4o (1–5 scale).

huts.)

BLEU: 7.77, chrF++: 32.8, TER: 72.73, Fluency: 4, Adequacy: 4, Dialect: 3

Mistral (manual + synthetic): "Отямивси аж увечері, коли зчинився в діда в обох хатах гармидер." (Eng: *He came to his senses in the evening, when a chaos broke out in both grandfather's houses.*)

BLEU: 31.46, chrF++: 46.62, TER: 54.55, Fluency: 4, Adequacy: 5, Dialect: 3

Mistral (manual only): "Отємивси аж увечері, ек зченивси в діда в обох хатах гармидер." (Eng: *He came to his senses only in the evening, after stirring up a mess in both of the grandfather's houses.*)

BLEU: 34.39, chrF++: 56.92, TER: 45.45, Fluency: 4, Adequacy: 5, Dialect: 4

LLAMA (manual + synthetic): "Отямивси аж у вечерє, коли зчинивси в діда в обох хатах гармидер." (Eng: *He came to his senses in the evening, when a mess started in both grandfather's houses.*)

BLEU: 20.45, chrF++: 46.42, TER: 63.64, Fluency: 4, Adequacy: 4, Dialect: 3

LLAMA (manual only): "Отємивси аж увечері, ек зчинивси у діда в обох хатах гармидер." (Eng: *He came to his senses only in the evening, when a mess in the grandfather's both houses started.*)

BLEU: 24.71, chrF++: 49.59, TER: 54.55, Fluency: 4, Adequacy: 5, Dialect: 3

Limitations

Our work makes first step in Ukrainian dialect adaptation for LLMs, a lot of limitations remain open.

An important limitation is that although we introduced a synthetic data generation pipeline to mitigate limited data availability problem, synthetic

translations may lack native fluency or have stylistic inconsistencies, especially for underrepresented topics. This is particularly can be seen in domains not covered by the original corpus, such as politics, technology, etc. where Hutsul lexicon is either very limited or absent. Despite filtering low-quality generations, automatic evaluation metrics still may overestimate linguistic validity.

In addition, evaluation remains challenging. Automatic metrics such as BLEU and chrF++ often penalize valid dialectal variation (Garcia et al., 2024; Held and Klakow, 2024). To better capture stylistic and synthetic diversity, we use GPT-4o as an LLM-based judge following recent work on LLM-based evaluation frameworks (Wang, 2023; Liu, 2023). However, we note that GPT-4o is not explicitly fine-tuned for dialectal assessment, and its preferences may still align with standard Ukrainian and human evaluation would provide much more reliable assessments.

Also we need to mention that our current methods are tailored to Hutsul, a relatively well-documented dialect within the Ukrainian language. Extension to other dialects or usage of the same approach for other low-resource languages will require adaptation of both the data pipeline and prompting strategies.

Acknowledgments

We express our sincere gratitude to Іван Андрусяк (Ivan Andrusiak) for providing access to his Ukrainian translation of "Дідо Иванчік" (Dido Yvanchik), which served as a cornerstone of our dataset. We also thank the publishing house "Дискурс" (Dyskurs) and its director Василь Карп'юк (Vasyl Karpiuk) for their kind permission to use the text and for their continued support of linguistic and cultural preservation initiatives. Their generosity made this research possible.

References

- Noëmi Aeppli, Sarah Ebling, and Rico Sennrich. 2023. A benchmark for evaluating machine translation metrics on dialects without standard orthography. *arXiv preprint arXiv:2311.16865*.
- Mistral AI. 2023. Introducing mistral 7b and mixtral. <https://mistral.ai/news/mistral-7b/>.
- Rogier Blokland, Trond Trosterud, and Jack Rueter. 2024. Morphological variants in north s’ami dialects. In *Proceedings of the Tenth Workshop on NLP for Similar Languages, Varieties and Dialects (VarDial)*.
- Dmytro Chaplynskyi. 2023. Introducing UberText 2.0: A corpus of modern Ukrainian at scale. In *Proceedings of the Second Ukrainian Natural Language Processing Workshop*, pages 1–10, Dubrovnik, Croatia. Association for Computational Linguistics.
- Ned Cooper, Courtney Heldreth, and Ben Hutchinson. 2024. “it’s how you do things that matters”: Attending to process to better serve indigenous communities with language technologies. In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (EACL 2024)*, pages 204–211.
- G’abor Cs’aki and 1 others. 2023. Tokenizer retrofitting for morphologically rich languages. In *Findings of the Association for Computational Linguistics: ACL 2023*.
- Chris Dyer, Victor Chahuneau, and Noah A. Smith. 2013. A simple, fast, and effective reparameterization of IBM model 2. In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 644–648, Atlanta, Georgia. Association for Computational Linguistics.
- Xavier Garcia and 1 others. 2024. Don’t hallucinate, retrieve! a survey on retrieval-augmented text generation. *arXiv preprint arXiv:2402.07836*.
- Vasyl Greshchuk. 2016. Models of word formation in hutsul dialects (based on the dictionary “hutsul dialectal vocabulary in ukrainian belletristic language”). *Gramatychni Studii*, 6:272–286.
- Aditya Gudibande and 1 others. 2023. Synthetic data scaling for low-resource nlp. *arXiv preprint arXiv:2305.11864*.
- Wolfgang Held and Dietrich Klakow. 2024. Tada: Task-agnostic dialect adapters for multilingual transformers. In *Proceedings of the First Workshop on Domain Adaptation for NLP*.
- Nora Hollenstein, C’edrick Fairon, and Julie Snyers. 2020. German’s many voices: A corpus of regional variation in german. In *Proceedings of the 12th Language Resources and Evaluation Conference*.
- Edward J Hu and 1 others. 2021. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*.
- Torodd Kinn and Tor A. Åfarli. 2024. Exploring parallel machine translation for norwegian nynorsk and bokmål. In *Proceedings of the Tenth Workshop on NLP for Similar Languages, Varieties and Dialects (VarDial)*.
- Christo Kirov, Ryan Cotterell, and 1 others. 2023. Sigmorphon 2023 shared task: Morphological inflection in context. In *Proceedings of the 20th SIGMORPHON Workshop*.
- Artur Kiulian, Anton Polishko, Mykola Khandoga, Oryna Chubych, Jack Connor, Raghav Ravishankar, and Adarsh Arunkumar Shirawalmath. 2024. From bytes to borsch: Fine-tuning gemma and mistral for the ukrainian language representation. *arXiv preprint arXiv:2404.09138*.
- Zihao Li, Yucheng Shi, Zirui Liu, Fan Yang, Ninghao Liu, and Mengnan Du. 2024. Language ranker: A metric for quantifying llm performance across high and low-resource languages. *arXiv preprint arXiv:2404.11553*.
- Zi Lin, Joel Tetreault, and 1 others. 2021. Multi-value: A multilingual, multi-dialect, and multi-task benchmark for language understanding. In *Proceedings of EMNLP 2021*.
- Xiang Liu and 1 others. 2024. Dada: Dynamic adapter aggregation for dialectal adaptation. *arXiv preprint arXiv:2409.11404*.
- Ziwei et al. Liu. 2023. Gpt-4 as an automatic grader: An evaluation of zero-shot and few-shot prompting for text scoring tasks. *arXiv preprint arXiv:2304.02329*.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318.
- Maja Popović. 2015. chrF: character n-gram f-score for automatic mt evaluation. In *Proceedings of the Tenth Workshop on Statistical Machine Translation*, pages 392–395.
- Alan Ramponi and Barbara Plank. 2021. Neural multi-dialect language models for zero-shot cross-dialect transfer. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*.
- Yves Scherrer. 2023. Character alignment for dialect standardization: A comparative evaluation. In *Proceedings of the 1st Workshop on NLP for Less-Resourced Languages*.
- Matthew Snover, Bonnie Dorr, Richard Schwartz, Linnea Micciulla, and John Makhoul. 2006. A study of translation edit rate with targeted human annotation.

- In *Proceedings of the 7th Conference of the Association for Machine Translation in the Americas*, pages 223–231.
- Shahbaz Syed, Ahmad Dawar Hakimi, Khalid Al-Khatib, and Martin Potthast. 2023. [Quantifying the dialect gap and its correlates across languages](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 5196–5210.
- Hugo Touvron, Thibaut Lavril, Alp Yurtsever, and 1 others. 2024. [Llama 3: Open foundation and instruction models](#). *arXiv preprint arXiv:2404.14219*.
- Peter Trudgill. 2003. [Dialect contact and new-dialect formation: The inevitability of colonial englishes](#). In *Proceedings of the 14th International Congress of Phonetic Sciences*, pages 2193–2196.
- Yizhong et al. Wang. 2023. [Llm-eval: Unified, automatic and robust evaluation of large language models with gpt-4](#). *arXiv preprint arXiv:2305.03045*.
- Shijie Wu, Yuxuan Li, Cheng Li, Hao Zhu, and 1 others. 2023. [Benchmarking public large language models in low-resource settings](#). In *Proceedings of the 2023 EMNLP*.
- Marcos Zampieri, Tommi Jauhiainen, Nikola Ljubešić, Noëmi Aepli, Simon Clematide, and Jörg Tiedemann. 2024. [Overview of the vardial evaluation campaign 2024](#). In *Proceedings of the Tenth Workshop on NLP for Similar Languages, Varieties and Dialects (VarDial)*.
- Marcos Zampieri, Shervin Malmasi, Preslav Nakov, Ahmed Ali, and Stephan Vogel. 2017. [Arabic dialect identification for the dsl 2017 shared task](#). In *Proceedings of the Fourth Workshop on NLP for Similar Languages, Varieties and Dialects (VarDial)*.
- Tianyu Zhong, Ziqi Yang, Zhen Liu, Rui Zhang, Yiheng Liu, Hanqi Sun, Yujia Pan, Yiming Li, and Yifan Zhou. 2024. [Opportunities and challenges of large language models for low-resource languages in humanities research](#). *arXiv preprint arXiv:2412.04497*.