

Charles Translator: A Machine Translation System between Ukrainian and Czech

Martin Popel,¹ Lucie Poláková,¹ Michal Novák,¹ Jindřich Helcl,¹ Jindřich Libovický,¹
Pavel Straňák,¹ Tomáš Krabač,¹ Jaroslava Hlaváčová,¹
Mariia Anisimova¹ and Tereza Chlaňová²

¹Charles University, Faculty of Mathematics and Physics, Institute of Formal and Applied Linguistics

²Charles University, Faculty of Arts, Department of East European Studies

Prague, Czech Republic

{surname}@ufal.mff.cuni.cz, mnovak@ufal.mff.cuni.cz, tereza.chlanova@ff.cuni.cz

Abstract

We present Charles Translator, a machine translation system between Ukrainian and Czech, developed as part of a society-wide effort to mitigate the impact of the Russian-Ukrainian war on individuals and society. The system was developed in the spring of 2022 with the help of many language data providers in order to quickly meet the demand for such a service, which was not available at the time in the required quality. The translator was later implemented as an online web interface and as an Android app with speech input, both featuring Cyrillic-Latin script transliteration. The system translates directly, compared to other available systems that use English as a pivot, and thus take advantage of the typological similarity of the two languages. It uses the block back-translation method, which allows for efficient use of monolingual training data. The paper describes the development process, including data collection and implementation, evaluation, mentions several use cases, and outlines possibilities for the further development of the system for educational purposes.

Keywords: machine translation, Ukrainian, Czech

1. Introduction

As a result of the Russian invasion of Ukraine in February 2022, the Czech Republic became one of the main countries to host people forced to flee their homes. According to sources from the UN High Commissioner for Refugees (UNHCR), it is the fourth country with the largest number of Ukrainian refugees. By April 1, 2023, more than 504,000 Ukrainians had been granted temporary protection in the country, of whom more than 325,000 had applied for an extension of their refugee status beyond March 2023.¹ Virtually overnight, there arose the need for a fast and effective means of communication between Czech and Ukrainian speakers, which until then did not have the required quality.

Our motivation to develop such a service, apart from the wish to help reduce the language (and social) barrier between the refugees and the Czech society, is based on several convenient factors: (i) our previous long-term scientific experience in the field of machine translation (MT) and the existence of an appropriate MT method, (ii) the proximity of the two Slavic languages in question, and (iii) the availability of resources: the possibility of obtaining training data from multiple volunteer subjects and the willingness of many researchers to prioritize this line of research, leading to a quick

solution with a quick implementation process.

The translation systems available to the public during the conflict outbreak translated only indirectly between Czech and Ukrainian by pivoting through English. This approach does not take advantage of the typological affinity of the two languages, such as the high inflection with rich morphology enabling great flexibility of word order, pro-drop, partial lexical similarity, e.g. *můj dům – můj дім* (my house), *chladná zima – холодна зима* (cold winter), *krátké vlasy – коротке волосся* (short hair)² and syntactic similarities.

With English as the pivot, some information is inevitably lost. This is most visible in the grammatical categories of gender and politeness: both Ukrainian and Czech distinguish masculine, feminine, and neutral forms of nouns and adjectives, and a formal (*Vy, Bu*) and an informal (*ty, mu*) form of the *you*-pronoun (in singular and plural). Thus, the following Czech example with a female speaker addressing the hearer in a formal (polite) way:

Jsem nemocná. A co Vy?

0-am sick-F.SG. And what you-FORMAL.SG?

'I'm sick. And you?'

²At the same time, there is quite a large number of false friends, e.g.: *квітень – duben* (April, resembling *květen – May* in Czech); *лікарня – nemocnice* (hospital, resembling *lékárna – pharmacy* in Czech), *напад – útok* (attack, resembling *nápad – idea* in Czech).

¹<https://data2.unhcr.org/en/documents/details/104052>

gets incorrectly translated as a male speaker informally addressing the listener:

Я хворий. А ти?
Ja chvoryj. A ty?

I sick-M.SG. And you-INFORMAL.SG?

Similarly, there also can be a fatal shift of meaning when, for instance, the Czech *Jaké léky to jsou?* (*What medicine is that?*) translates as *Що це за наркотики?* (literally *What are these narcotics?*) since English *drugs* can mean both *medicine* and *drugs/narcotics*. Charles Translator translates directly, and thus is not prone to these errors.

In the rest of the paper, we present related work (§ 2); the translator architecture, training and test data, deployment and user interfaces (§ 3); evaluation (§ 4); use cases and usage statistics (§ 5) and we conclude with plans for future (§ 6).

2. Related Work

Current MT methods are largely language independent and rely on the Transformers architecture (Vaswani et al., 2017), making substantial use of back-translation (Sennrich et al., 2016; Edunov et al., 2018) and data filtering in all stages of model training (Junczys-Dowmunt, 2018).

Translation between Czech and Ukrainian was part of the WMT22 evaluation campaign (Kocmi et al., 2022), with most participants relying on language-agnostic methods. The winning system (Nowakowski et al., 2022) enriched the source side of the translation with information about named entities and used complex decoding with a neural model to restore hypotheses. In our submission Popel et al. (2022), we used handcrafted regular expressions to handle errors in named entities during data filtering, which is also part of Charles Translator. Similarly to Alabi et al. (2022), we experimented with romanization, which is not used in the deployed system.

Although the implementation aspects of MT are discussed in the literature (Junczys-Dowmunt et al., 2018; Behnke et al., 2021; Heafield et al., 2022), there is virtually no related work focusing on the deployment of machine translation, which typically remains part of the secret know-how of commercial MT providers.

3. Components of the Translator

Charles Translator consists of the translation service and multiple interfaces for accessing the translator.

3.1. The translation service

Method. We use the Transformer architecture (Vaswani et al., 2017) with iterated block back-translation (Popel et al., 2020), allowing for more efficient monolingual training data use. The system was trained in the same way as the sentence-level English-Czech system of Popel (2018).

Training Data. The collection of training data for the first model took place over a short and intensive period with the help of a wide range of volunteer subjects. Cooperation with Czech-Ukrainian translators, translation agencies, and the authors of the InterCorp parallel corpus (Rosen et al., 2022), a project of the Czech National Corpus, was important for obtaining good quality parallel data.

We also used data available at the OPUS repository (Tiedemann, 2012), namely texts from the Bible (Christodouloupoulos and Steedman, 2015), CCMatrix (Schwenk et al., 2019; Fan et al., 2021), ELRC, EUBookshop, GNOME, KDE4, MultiCCAligned (El-Kishky et al., 2020), MultiParaCrawl,³ OpenSubtitles (Lison and Tiedemann, 2016), QED (Abdelali et al., 2014), Tatoeba, TED2020 (Reimers and Gurevych, 2020), Ubuntu, WikiMatrix (Schwenk et al., 2021), and XLEnt (El-Kishky et al., 2021).

In addition to the (authentic) parallel data, we also used monolingual data for backtranslation: 50M originally Czech sentences from CzEng 2.0 (Kocmi et al., 2020) and 58M originally Ukrainian sentences from WMT NewsCrawl,⁴ the Leipzig Corpora (Biemann et al., 2007), UberText corpus (Khaburska and Tytyk, 2019) and Legal Ukrainian Crawling by ELRC (de Gibert Bonet, 2021).

Test data. To evaluate the system performance in the area for which it was primarily designed, i.e., the daily communication of the refugees with Czech individuals and authorities, we created two test sets: 2,812 sentences for UK→CS (from March and April 2022), and 2,017 sentences for CS→UK (from 2023). We provided these two test sets to the organizers of WMT⁵ and they were published as WMT22 UK→CS (Kocmi et al., 2022) and WMT23 CS→UK (Kocmi et al., 2023), respectively.

Some of the sentences were news crawls provided by the WMT organizers. However, most sentences were selected from the Charles Translator system logs,⁶ anonymized/pseudonymized and

³<https://paracrawl.eu>

⁴<https://data.statmt.org/news-crawl/>

⁵Conference on Machine Translation, <https://www.statmt.org/wmt23/>

⁶From users who agreed to have their data used for further system development, cf. Section 8.

translated by professional translators.

The test sets were also annotated for (i) user type: formal (bureaucracy), news, and other (mostly individual users); and (ii) topic: general personal conversation, work, housing, transportation/travel, school and education, health, and politics. The test sets were designed to be balanced in these respects, but also in sentence length and the number of “noisy” sentences, i.e., user-generated sentences with authentic typos, grammatical and typographic errors, and disfluencies. The 2017 sentences (segments) of the WMT23 CS→UK test set are separated into 5 domains (see Table 3 for evaluation):

- News: 567 segments from WMT news crawl,
- Voice: 533 segments of originally spoken Czech, as recognized by an automatic speech recognition system (ASR), i.e. including some ASR errors,
- Personal: 390 segments of personal conversation,
- Official: 347 segments of formal/public announcements,
- Games: 180 segments of web stories about computer games.

Engine. The backend service is run by the LINDAT/CLARIAH-CZ infrastructure.⁷ It uses one server with 15 CPU cores and 44 GB RAM, with the translation models for UK-CS and CS-UK loaded on 3 GPU cards: Quadro RTX 5000 + 2x GeForce RTX 2080 Ti. The service is accessible via REST API.

The infrastructure was already set up and running for other translation pairs when the Russian invasion started, so we could quickly prepare data, train new models, and provide a robust service.

3.2. Interfaces

We developed multiple interfaces that communicate with the translation service.

Original web interface. We have been providing the translation service for several languages even before we developed the Czech-Ukrainian models, accessible via the original web interface.⁸ Although this web interface is still available, it is more research-oriented. For example, it enables translation via pivoting for pairs without a trained model, regardless of the final translation quality.

New web app. The start of the refugee crisis required a simpler application that would be accessible to end users. To meet this demand,

⁷<https://lindat.cz>

⁸<https://lindat.mff.cuni.cz/services/translation/>

System	BLEU	chrF	COMET
GPT4-5shot	32.8	61.0	90.8
Charles Translator	30.2	57.4	88.0
GTCOM_Peter	29.8	57.6	88.9
CUNI-GA	29.5	57.9	90.9
MUNI-NLP	28.3	57.0	87.0
Lan-BridgeMT	27.5	55.7	86.0
ONLINE-W	26.8	55.0	89.4
ONLINE-B	25.7	54.7	88.8
ONLINE-A	25.4	54.4	88.2
NLLB_MBR_BLEU	25.1	52.3	86.3
NLLB_Greedy	24.9	52.5	86.3
ONLINE-G	24.8	53.7	87.7
ONLINE-Y	24.2	53.4	86.5

Table 1: Results of automatic evaluation on the CS→UK WMT23 test set. Constrained systems (i.e. systems that use only the data provided by the WMT organizers) are marked with a white background.

we developed a new React/Node.js web app that only included Ukrainian-Czech models at the time. Following its creation on an ad-hoc organized hackathon, the development and maintenance of the app was taken over by a single developer. Currently, the app still supports the CS↔UK translation pair only, but we plan to extend it with other pairs soon.

Because Czech and Ukrainian use different scripts, the app supports transliteration in both directions, as illustrated on the left side of Figure 1. Unlike other online translators, we use a transliteration of Ukrainian into the Latin script that is suitable for Czech speakers, e.g. *нашу* is transcribed as *našu* instead of English-oriented *nashu*. When transcribing Czech into Cyrillic script, we keep the acute diacritic marks signaling vowel length, e.g. *článek* is transcribed as *члáneк*.

Android app. After the web frontend was released and publicized, we extended the translation service with an Android app. Besides accessibility-related benefits, an Android app can take advantage of the native speech API, offering both dictation and speech synthesis capabilities in Czech and Ukrainian. Using these features, we also provide the app with a conversation mode, as illustrated on the right side of Figure 1. Recently, the Android app has been extended to include also other translation pairs. The app is implemented in the Kotlin programming language using Jetpack Compose. The app is currently installed by around 2,000 people, of which 1,200 have Czech as their phone language and 600 have Ukrainian.

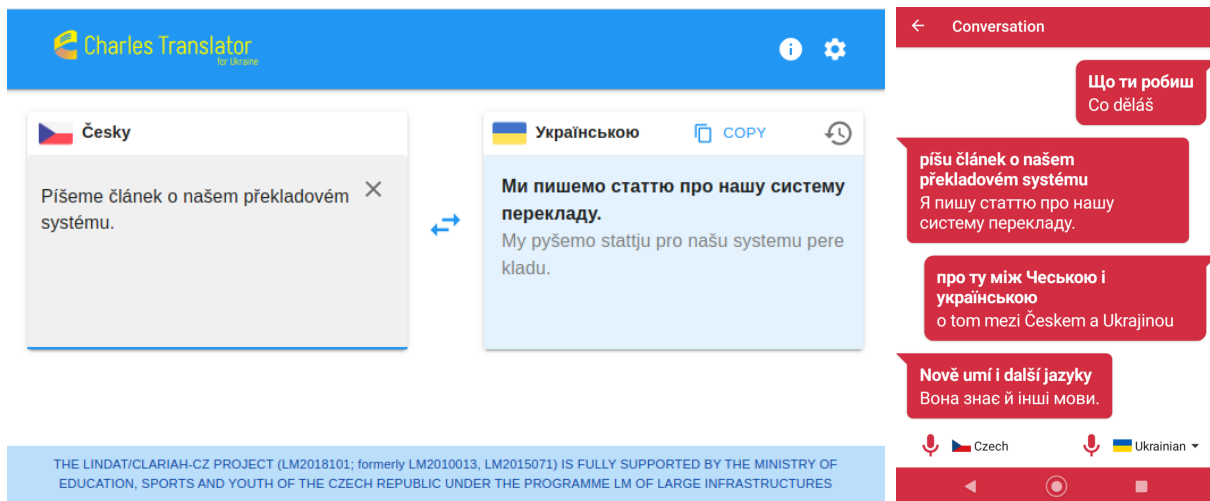


Figure 1: Screenshots of the new web app (left) and the conversation mode in the Android app (right). Note that the Czech translation *mezi Českem a Ukrajinou* (between Czechia and Ukraine) is incorrect. It should be *mezi češtinou a ukrajinštinou* (between Czech and Ukrainian [language]), instead. The error is actually caused by speech recognition that capitalizes the word *чеською*, thus changing its meaning.

System	BLEU	chrF	COMET
AMU	37.0	60.7	104.8
Lan-Bridge	36.5	60.4	94.5
Online-B	36.4	60.3	96.5
HuaweiTSC	36.0	59.6	91.4
Charles Translator-un	35.9	59.0	90.2
Charles Translator	35.8	59.0	88.5
CUNI-JL-JH	35.1	58.7	89.0
Online-A	33.3	57.5	85.4
Online-G	31.5	56.3	84.2
GTCOM	31.3	55.8	80.2
Online-Y	29.6	55.3	78.6
ALMAnaCH-Inria	25.3	50.7	62.4

Table 2: Results of automatic evaluation on the UK→CS WMT22 test set. *Charles Translator-un* is an unconstrained version of our translator, i.e. trained on additional training data from InterCorp.

4. Quality of Translations

We evaluate our system on the Czech-Ukrainian WMT23⁹ General MT test set using three automatic metrics: BLEU (Papineni et al., 2002), ChrF (Popović, 2015) and COMET (Rei et al., 2020).¹⁰ The results in Table 1 show different rankings for each metric, but our system is never more than 8% worse than the best system. Table 2 shows results on the Ukrainian-Czech WMT22 General MT test set. Table 3 shows results, again on Czech-Ukrainian WMT23, but separately for each of the five domains included in the test set,

⁹<https://www.statmt.org/wmt23/>

¹⁰According to the human evaluation (Kocmi et al., 2023), Charles Translator (called CUNI-Transformer) is significantly outperformed only by systems GPT4-5shot and ONLINE-B and the human reference.

including *voice*, which are the ASR outputs.

While our internal manual evaluation shows notable improvements over older versions of our system, there are still occasional errors, especially in the translation of city names and personal names.¹¹ See also the caption of Figure 1 showing a translation error caused by an ASR error.

5. Use Cases

From the very beginning of the project, we have been actively contacting organizations that we anticipated would be assisting Ukrainian refugees, asking about their translation needs. If they were interested, we shared the Charles Translator REST API with them as soon as the first version was ready.

One of the organizations that we partnered with is the Consortium of Migrants Assisting Organizations, an umbrella organization of multiple NGOs dealing with migrants. It maintains the *Pomáhej Ukrajině* (Help Ukraine) web platform,¹² which connects public offers of assistance in various areas (e.g., material aid, housing, education, leisure activities, psychological support) with the needs of individual Ukrainian refugees and organizations involved in aiding migrants. The platform uses our API to automatically translate segments of the offers, particularly those that involve filling in free-form text.

We were contacted by the Police of the Czech Republic, who started gathering information on

¹¹E.g. *Košice* (a Slovak city) translates as *Харьков* (Russian name of the Ukrainian city *Харків*).

¹²<https://www.pomahejukrajine.cz>

System	BLEU for domain						chrF for domain					
	ALL	games	news	official	personal	voice	ALL	games	news	official	personal	voice
GPT4-5shot	32.8	26.7	31.8	36.8	34.5	32.9	61.0	57.4	61.9	63.9	60.1	59.4
Charles Translator	30.2	24.3	30.4	34.2	30.8	29.5	57.4	55.1	59.0	60.6	55.1	54.5
GTCOM_Peter	29.8	24.8	30.7	35.2	29.5	26.0	57.6	54.7	59.6	61.7	55.1	53.2
CUNI-GA	29.5	24.3	30.6	33.4	29.7	26.9	57.9	55.8	60.0	60.6	55.8	54.1
MUNI-NLP	28.3	24.9	27.9	31.9	29.2	27.4	57.0	55.7	58.2	59.8	54.8	54.1
Lan-BridgeMT	27.5	24.0	26.9	31.4	27.8	26.3	55.7	54.1	57.6	58.6	52.7	52.3
ONLINE-W	26.8	20.9	27.3	32.6	26.1	24.0	55.0	51.5	56.9	59.8	51.8	51.4
ONLINE-B	25.7	20.6	25.0	31.5	26.4	24.5	54.7	52.1	56.2	58.6	52.2	51.6
ONLINE-A	25.4	20.5	25.1	30.7	25.8	23.6	54.4	51.1	56.0	58.5	52.3	51.1
NLLB_MBR_BLEU	25.1	22.4	24.4	28.7	25.2	24.9	52.3	50.6	53.5	55.7	50.1	49.0
NLLB_Greedy	24.9	21.7	25.6	28.1	24.2	23.5	52.5	50.8	54.8	55.6	49.0	48.6
ONLINE-G	24.8	20.6	25.1	30.9	24.1	21.1	53.7	51.3	55.8	58.3	50.6	48.5
ONLINE-Y	24.2	20.1	23.4	29.7	23.8	22.7	53.4	51.5	55.3	57.4	49.6	49.8

Table 3: BLEU and chrF results on various domains (subsets) of the CS→UK WMT23 test set.

war crimes committed in Ukraine using an online form,¹³ where the witnesses can report their testimonies in Czech, Ukrainian, Russian, and English. The testimonies are highly sensitive, so they cannot be translated using online translators. We thus provided the Police with the on-premise installation of our translator on their servers.

In September 2023, the translation service showed the following usage statistics:¹⁴ In the Ukrainian→Czech translation direction, there was an average of 30,000 translation requests per day and about two million characters translated per day. These translations were quite concise on average, with an average length of around 60 characters per request, although there are also requests with hundreds of sentences. In the Czech→Ukrainian direction, there were approx. 12,000 requests per day and a total of approx. one million characters translated per day.

After the launch of the service in March 2022 and during its early days, the demand peaked in April 2022 at 223 million characters in the CS→UK direction. Over the rest of 2022, it steadily decreased, with the trend continuing in 2023 from 62 million characters in January to 27 million in September.

In the UK→CS direction, there is no such decrease in usage. Although the highest number is 84 million characters translated in May 2022, the average use until September 2023 has been practically flat, around 65 million characters/month without any trends.

Although both Ukrainian refugees and Czech people communicating with them may need both translation directions, we hypothesize that most of the UK→CS translations are by Ukrainian refugees.

¹³<https://oznameni.policie.cz>

¹⁴Excluding the usage of on-premise installations.

6. Conclusions

We presented Charles Translator, a machine translation system between Ukrainian and Czech based on the block back-translation method. The main motivation behind its rapid development was to facilitate communication between Czech and Ukrainian speakers during the critical period following the migration from Ukraine. The translator is available as API, web app and Android app. Charles Translator’s license allows for free use including the API for non-commercial purposes, so it can be integrated in a wide range of translating activities that are free of charge to the user.

In addition to further work on improving translation quality, we plan to adapt the model for the educational domain in order to create multilingual digital learning materials in the near future. At present, approximately 55,000 Ukrainian children are enrolled in the Czech school system, requiring a high-quality education that seamlessly integrates Ukrainian and Czech in all subjects.

7. Acknowledgements

This work was supported by the project TQ01000458 (EdUKate) financed by the Technology Agency of the Czech Republic (www.tacr.cz) within the Sigma 3 Programme. It was further supported by the Ministry of Education, Youth and Sports of the Czech Republic, Project No. LM2023062 LINDAT/CLARIAH-CZ and Cooperatio AMES of the Faculty of Arts, Charles University.

The authors also express their gratitude to the Institute of the Czech National Corpus (korpus.cz), to České překlady (ceskepreklady.cz) and to the following individuals for their contributions: Zdeněk Kasner, Jiří Mayer, Rudolf Rosa, Ondřej Košarko, Martin Majliš, Oxana Čmelíková, Peter

Polák, Tomáš Musil, David Nápravník and Denys Bojko.

8. Ethical considerations

The sentences in the test sets (described in Section 3.1) were collected with users' opt-in consent and any personal data (except for the names of well-known public figures) were pseudonymized (using random first names and surnames), including all potentially identifying information (addresses, URLs, institution names, phone numbers, names of cities and villages). Sentences where such pseudonymization would not guarantee reasonable anonymity of the users (e.g. describing events uniquely identifying the persons involved) were not included in the test set.

The raw data collected are stored securely on our servers and they are not publicly available. The users can withdraw their consent anytime in the web/mobile app and we immediately stop collecting their data (i.e. texts to be translated). The data already stored in our database can be deleted on request.

9. Bibliographical References

Ahmed Abdelali, Francisco Guzman, Hassan Sajjad, and Stephan Vogel. 2014. [The AMARA corpus: Building parallel language resources for the educational domain](#). In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC'14)*, pages 1856–1862, Reykjavik, Iceland. European Language Resources Association (ELRA).

Jesujoba Alabi, Lydia Nishimwe, Benjamin Muller, Camille Rey, Benoît Sagot, and Rachel Bawden. 2022. [Inria-ALMANaCH at WMT 2022: Does transcription help cross-script machine translation?](#) In *Proceedings of the Seventh Conference on Machine Translation (WMT)*, pages 233–243, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.

Maximiliana Behnke, Nikolay Bogoychev, Alham Fikri Aji, Kenneth Heafield, Graeme Nail, Qianqian Zhu, Svetlana Tchistiakova, Jelmer van der Linde, Pinzhen Chen, Sidharth Kashyap, and Roman Grundkiewicz. 2021. [Efficient machine translation with model pruning and quantization](#). In *Proceedings of the Sixth Conference on Machine Translation*, pages 775–780, Online. Association for Computational Linguistics.

Sergey Edunov, Myle Ott, Michael Auli, and David Grangier. 2018. [Understanding back-translation](#)

[at scale](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 489–500, Brussels, Belgium. Association for Computational Linguistics.

Kenneth Heafield, Biao Zhang, Graeme Nail, Jelmer Van Der Linde, and Nikolay Bogoychev. 2022. [Findings of the WMT 2022 shared task on efficient translation](#). In *Proceedings of the Seventh Conference on Machine Translation (WMT)*, pages 100–108, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.

Marcin Junczys-Dowmunt. 2018. [Dual conditional cross-entropy filtering of noisy parallel corpora](#). In *Proceedings of the Third Conference on Machine Translation: Shared Task Papers*, pages 888–895, Belgium, Brussels. Association for Computational Linguistics.

Marcin Junczys-Dowmunt, Roman Grundkiewicz, Tomasz Dwojak, Hieu Hoang, Kenneth Heafield, Tom Neckermann, Frank Seide, Ulrich Germann, Alham Fikri Aji, Nikolay Bogoychev, André F. T. Martins, and Alexandra Birch. 2018. [Marian: Fast neural machine translation in C++](#). In *Proceedings of ACL 2018, System Demonstrations*, pages 116–121, Melbourne, Australia. Association for Computational Linguistics.

Tom Kocmi, Eleftherios Avramidis, Rachel Bawden, Ondřej Bojar, Anton Dvorkovich, Christian Federmann, Mark Fishel, Markus Freitag, Thamme Gowda, Roman Grundkiewicz, Barry Haddow, Philipp Koehn, Benjamin Marie, Christof Monz, Makoto Morishita, Kenton Murray, Makoto Nagata, Toshiaki Nakazawa, Martin Popel, Maja Popović, and Mariya Shmatova. 2023. [Findings of the 2023 Conference on Machine Translation \(WMT23\): LLMs Are Here but Not Quite There Yet](#). In *Proceedings of the Eighth Conference on Machine Translation*, pages 1–42, Singapore. Association for Computational Linguistics.

Tom Kocmi, Rachel Bawden, Ondřej Bojar, Anton Dvorkovich, Christian Federmann, Mark Fishel, Thamme Gowda, Yvette Graham, Roman Grundkiewicz, Barry Haddow, Rebecca Knowles, Philipp Koehn, Christof Monz, Makoto Morishita, Masaaki Nagata, Toshiaki Nakazawa, Michal Novák, Martin Popel, and Maja Popović. 2022. [Findings of the 2022 conference on machine translation \(WMT22\)](#). In *Proceedings of the Seventh Conference on Machine Translation (WMT)*, pages 1–45, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.

- Artur Nowakowski, Gabriela Pałka, Kamil Guttman, and Mikołaj Pokrywka. 2022. [Adam Mickiewicz University at WMT 2022: NER-assisted and quality-aware neural machine translation](#). In *Proceedings of the Seventh Conference on Machine Translation (WMT)*, pages 326–334, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- 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, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Martin Popel. 2018. [CUNI transformer neural MT system for WMT18](#). In *Proceedings of the Third Conference on Machine Translation: Shared Task Papers*, pages 482–487, Belgium, Brussels. Association for Computational Linguistics.
- Martin Popel, Jindřich Libovický, and Jindřich Helcl. 2022. [CUNI systems for the WMT 22 Czech-Ukrainian translation task](#). In *Proceedings of the Seventh Conference on Machine Translation (WMT)*, pages 352–357, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- Martin Popel, Markéta Tomková, Jakub Tomek, Łukasz Kaiser, Jakob Uszkoreit, Ondřej Bojar, and Zdeněk Žabokrtský. 2020. [Transforming machine translation: a deep learning system reaches news translation quality comparable to human professionals](#). *Nature Communications*, 11(4381):1–15.
- 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, Lisbon, Portugal. Association for Computational Linguistics.
- Ricardo Rei, Craig Stewart, Ana C Farinha, and Alon Lavie. 2020. [COMET: A neural framework for MT evaluation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2685–2702, Online. Association for Computational Linguistics.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. [Improving neural machine translation models with monolingual data](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 86–96, Berlin, Germany. Association for Computational Linguistics.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.

10. Language Resource References

- Chris Biemann, Gerhard Heyer, Uwe Quasthoff, and Matthias Richter. 2007. The leipzig corpora collection-monolingual corpora of standard size. *Proceedings of Corpus Linguistic*, 2007.
- Christos Christodouloupoulos and Mark Steedman. 2015. A massively parallel corpus: the bible in 100 languages. *Language resources and evaluation*, 49:375–395.
- Ona de Gibert Bonet. 2021. [Legal Ukrainian Crawling](#).
- Ahmed El-Kishky, Vishrav Chaudhary, Francisco Guzmán, and Philipp Koehn. 2020. [CCAligned: A massive collection of cross-lingual web-document pairs](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP 2020)*, pages 5960–5969, Online. Association for Computational Linguistics.
- Ahmed El-Kishky, Adithya Renduchintala, James Cross, Francisco Guzmán, and Philipp Koehn. 2021. [XLEnt: Mining a large cross-lingual entity dataset with lexical-semantic-phonetic word alignment](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 10424–10430, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Angela Fan, Shruti Bhosale, Holger Schwenk, Zhiyi Ma, Ahmed El-Kishky, Siddharth Goyal, Mandeep Baines, Onur Celebi, Guillaume Wenzek, Vishrav Chaudhary, et al. 2021. Beyond english-centric multilingual machine translation. *The Journal of Machine Learning Research*, 22(1):4839–4886.
- Anastasiia Khaburska and Igor Tytyk. 2019. Toward language modeling for the ukrainian. *Advances in Data Mining, Machine Learning, and Computer Vision. Proceedings*, pages 71–80.
- Tom Kocmi, Martin Popel, and Ondřej Bojar. 2020. [Announcing CzEng 2.0 Parallel corpus with over 2 gigawords](#). *CoRR*, abs/2007.03006.

- Pierre Lison and Jörg Tiedemann. 2016. [Open-Subtitles2016: Extracting large parallel corpora from movie and TV subtitles](#). In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC'16)*, pages 923–929, Portorož, Slovenia. European Language Resources Association (ELRA).
- Nils Reimers and Iryna Gurevych. 2020. [Making monolingual sentence embeddings multilingual using knowledge distillation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4512–4525, Online. Association for Computational Linguistics.
- Rosen, A. and Vavřín, M. and Zasina, A. J. 2022. [Korpus InterCorp – Czech, Ukrainian](#). Version 14 from 31. 1. 2022.
- Holger Schwenk, Vishrav Chaudhary, Shuo Sun, Hongyu Gong, and Francisco Guzmán. 2021. [WikiMatrix: Mining 135M parallel sentences in 1620 language pairs from Wikipedia](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 1351–1361, Online. Association for Computational Linguistics.
- Holger Schwenk, Guillaume Wenzek, Sergey Edunov, Edouard Grave, and Armand Joulin. 2019. [Ccmatrix: Mining billions of high-quality parallel sentences on the WEB](#). *CoRR*, abs/1911.04944.
- Jörg Tiedemann. 2012. [Parallel data, tools and interfaces in OPUS](#). In *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC'12)*, pages 2214–2218, Istanbul, Turkey. European Language Resources Association (ELRA).