# Investigating Multi-Pivot Ensembling with Massively Multilingual Machine Translation Models 

Alireza Mohammadshahi ${ }^{* 1,2}$ Jannis Vamvas ${ }^{1}$ Rico Sennrich ${ }^{1}$<br>${ }^{1}$ University of Zurich ${ }^{2}$ EPFL alireza.mohammadshahi@epfl.ch<br>\{vamvas,sennrich\}@cl.uzh.ch


#### Abstract

Massively multilingual machine translation models allow for the translation of a large number of languages with a single model, but have limited performance on low- and very-low-resource translation directions. Pivoting via high-resource languages remains a strong strategy for low-resource directions, and in this paper we revisit ways of pivoting through multiple languages. Previous work has used a simple averaging of probability distributions from multiple paths, but we find that this performs worse than using a single pivot, and exacerbates the hallucination problem because the same hallucinations can be probable across different paths. We also propose MaxEns, a novel combination strategy that makes the output biased towards the most confident predictions, hypothesising that confident predictions are less prone to be hallucinations. We evaluate different strategies on the FLORES benchmark for 20 low-resource language directions, demonstrating that MaxEns improves translation quality for low-resource languages while reducing hallucination in translations, compared to both direct translation and an averaging approach. On average, multi-pivot strategies still lag behind using English as a single pivot language, raising the question of how to identify the best pivoting strategy for a given translation direction. ${ }^{1}$


## 1 Introduction

Early work on multilingual neural machine translation (NMT) has explored combining source segments in different source languages (Zoph and Knight, 2016; Firat et al., 2016a), an idea that is also compatible with pivoting through intermediate languages. For example, one could translate from Dutch to Ukrainian by first translating the Dutch source to English and Russian, and then

[^0]making a combined prediction to Ukrainian. In the simplest case, this combination is achieved by predicting probability distributions for each source language and averaging these predictions in an ensemble-like manner (Firat et al., 2016a).

With massively multilingual NMT models (NLLB Team et al., 2022; Mohammadshahi et al., 2022a; Goyal et al., 2022; Wenzek et al., 2021; Zhang et al., 2020; Fan et al., 2021; Aharoni et al., 2019; Arivazhagan et al., 2019), one can in principle translate directly in any translation direction. While early models relied on zero-shot generalization for many directions, recent improvements include massive data collection efforts (Schwenk et al., 2021; El-Kishky et al., 2020) and synthetic data creation via back-translation (Edunov et al., 2018; Sennrich et al., 2016). However, these models still have low performance on many low-resource translation directions ${ }^{2}$ and pivottranslation via high-resource languages remains a strong baseline. Fan et al. (2021) also investigate the combination of multiple translation paths, which they call multi-source self-ensemble, that slightly improves over the direct translation and a single pivot for zero-shot language pairs.

In this paper, we investigate this multi-source self-ensembling strategy more closely, with a focus on preventing completely defunct translations such as hallucinations. However, we find that simple averaging is sub-optimal and may increase the number of hallucinations in the output, a typical failure case in low-resource settings. We relate this to a recent finding that hallucinations are sticky, meaning that different models trained on the same data and architecture may produce similar hallucinations (Guerreiro et al., 2023a). We also find evidence of such stickiness when combining multiple translation paths, and propose a new ensembling strategy that, instead of averaging probabilities, picks the

[^1]output with the maximum probability across different paths: MaxEns. This is partially inspired by the finding that model confidence is a good heuristic for avoiding hallucinations, which tend to be lowconfidence predictions (Guerreiro et al., 2023b).

We perform experiments on the FLORES benchmark (Goyal et al., 2022) for 20 low-resource translation directions by using two massively multilingual NMT models, SMaLL100 (Mohammadshahi et al., 2022a) and M2M100 (Fan et al., 2021). Our results show that while the average ensemble outperforms the direct translation, it still underperforms using only English as a pivot, both in terms of spBLEU and the number of hallucinations. MaxEns performs significantly better than the averaging strategy for both translation performance and hallucination. Specially, MaxEns has competitive translation performance with English pivoting on average, but still lags behind it on the hallucination performance. To sum up, our contributions are:

- We explore why a naive multi-pivot strategy with massively multilingual models can underperform single-pivot translation. Then, we propose MaxEns, a more robust ensembling technique for multi-pivot translation with multilingual NMT models.
- We evaluate different ensembling strategies on 20 low-resource translation directions of FLORES benchmark, and demonstrate that multi-pivot ensembling still lags behind the English pivoting.


## 2 Related Work

Several approaches exploited different multipivoting methods to improve the performance of NMT models, specifically for low-resource language directions (Macháček et al., 2023; Dabre et al., 2021; Kim et al., 2019; Cheng et al., 2017; Firat et al., 2016b). Macháček et al. (2023) analyzed the robustness of multi-source NMT in transcription errors. Dabre et al. (2021) improved the performance of simultaneous NMT by translating the source language into pivot languages, then applying the multi-source translation method (Zoph and Knight, 2016). Firat et al. (2016b) proposed a novel zero-resource translation approach by exploiting the multi-way multilingual NMT model, introduced by Firat et al. (2016a), and improved the performance over traditional pivot-based translation (Wu and Wang,

2007; Utiyama and Isahara, 2007). Cheng et al. (2017) introduced the pivot-based NMT model by jointly training source-to-pivot and pivot-to-target directions. Currey and Heafield (2019) proposed an alternative method by applying a monolingual pivot-language data for zero-resource NMT via back-translation (Sennrich et al., 2016).

## 3 Ensembling Methods

When performing direct translation, the score of a translation $Y$ given a source sequence $X_{s r c}$ is computed as follows:

$$
\begin{equation*}
s\left(Y ; X_{s r c}\right)=\sum_{i=1}^{|Y|} \log p\left(y_{i} \mid y_{<i}, X_{s r c}\right) \tag{1}
\end{equation*}
$$

where $p\left(y_{i} \mid y_{<i}, X_{\text {src }}\right)$ is the predicted probability of the $i$-th target token $y_{i}$ given the previous tokens $y_{<i}$ and the source sequence $X_{s r c}$.

For multi-pivot ensembling, we select a set of pivot languages $M=\left\{\mu_{1}, \mu_{2}, \ldots, \mu_{K}\right\}$ and generate the corresponding pivot translations $X_{M}=\left\{X_{\mu_{1}}, X_{\mu_{2}}, \ldots, X_{\mu_{K}}\right\}$. The final translation is generated by ensembling predictions, conditioned on the individual pivot translations.

In the following, we describe two approaches for such an ensembling: the multilingual averaging method and our MaxEns approach.

Multilingual Average (MultiAvg). Inspired by Fan et al. (2021); Firat et al. (2016b), we average the predicted probabilities of a token $y_{i}$ across all pivot languages: ${ }^{3}$

$$
\begin{equation*}
s\left(Y ; X_{M}\right)=\sum_{i=1}^{|Y|} \log \frac{1}{|M|} \sum_{k=1}^{|M|} p\left(y_{i} \mid y_{<i}, X_{\mu_{k}}\right) . \tag{2}
\end{equation*}
$$

where $|Y|$ and $|M|$ are the number of target tokens and pivots, respectively.
Maximum Ensemble (MaxEns). As novel combination strategy that biases the prediction towards the more confident pivot, we propose the following approach:

$$
\begin{equation*}
s\left(Y ; X_{M}\right)=\sum_{i=1}^{|Y|} \max _{k=1}^{|M|}\left[\log p\left(y_{i} \mid y_{<i}, X_{\mu_{k}}\right)\right] . \tag{3}
\end{equation*}
$$

[^2]

Figure 1: A sample translation of Afrikaans to Asturian by using SMaLL100. Translations of individual pivots are shown in the middle, output translations of MaxEns and MultiAvg on the right. MaxEns method eliminates the hallucination, as it follows the more confident pivot (here, Spanish). Glosses in English are presented within gray boxes.
where it chooses the maximum score between predictions of pivots for token $y_{i}$. Intuitively, MaxEns selects the most confident pivot language when generating token $y_{i}$.

## 4 Results and Discussion

### 4.1 Experimental Setup

Models. We used M2M100 and SMaLL100 as our massively multilingual NMT models. M2M100 is trained on large-scale multilingual corpora (Schwenk et al., 2021; El-Kishky et al., 2020) with a novel data mining procedure, that uses language similarities. We exploit M2M100 variant with 418M parameters. SMaLL100 (Mohammadshahi et al., 2022a) is a distilled version of M2M100 12B with 330M parameters. It has been trained with uniform sampling across all language pairs on nearly $6 \%$ of M2M100 pre-training dataset, and achieved competitive performance with M2M100 with 1.2B parameters.

Evaluation Setting. Inspired by Fan et al. (2021), we use the FLORES-101 benchmark (Goyal et al., 2022). It contains 3,001 sentences derived from English Wikipedia, and translated into 101 languages by human. We use the devtest subset for the evaluation. To better understand the effect of multilingual pivoting, we chose five low-resource (or very low) languages from different branches of Indo-European, including Germanic, Romance, Slavic, Indo-Aryan, and Iranian. These languages are Afrikaans, Asturian, Croatian, Urdu, and

Pashto. We evaluate on all permutations of these languages, which results in 20 language pairs. As pivot languages, we use English, Spanish and French. English has the largest amount of bitext overall in the training data of M2M100, and Spanish and French have the largest amount of bitext with English (Fan et al., 2021).

Evaluation Metrics. spBLEU ${ }^{4}$ is used to measure the translation performance (Goyal et al., 2022). For the hallucination measurement, we apply a coarse estimation method inspired by Lee et al. (2019); Müller and Sennrich (2021), counting the proportion of sentences with ChrF (Popović, $2015)^{5}$ less than $20 .{ }^{6}$ Additionally, we use top n-gram (TNG) (Guerreiro et al., 2023b; Raunak et al., 2022, 2021) for detecting oscillatory hallucinations. ${ }^{7}$ We apply significance testing with $p=0.05 .{ }^{8}$ Beam size 5 is used for inference.

### 4.2 Results \& Discussion

Figure 1 illustrates an example of multi-pivot translation of Afrikaans to Asturian by using English, French, and Spanish as pivots. Translations via English and French pivots are hallucinations, while the translation via the Spanish pivot is more related to the reference translation. The output of MaxEns

[^3]| Language Pairs | Direct | MultiAvg | MaxEns | EN Pivot |
| :---: | :---: | :---: | :---: | :---: |
| Afrikaans-Asturian | 20.6 | 21.8 | 22.5 | 21.0 |
| Afrikaans-Croatian | 22.1 | 22.5 | 22.4 | 22.8 |
| Afrikaans-Urdu | 14.0 | 13.8 | 13.9 | 14.4 |
| Afrikaans-Pashto | 5.9 | 5.8 | 6.0 | 6.0 |
| Asturian-Afrikaans | 18.5 | 19.9 | 19.8 | 21.0 |
| Asturian-Croatian | 16.1 | 20.4 | 20.0 | 20.4 |
| Asturian-Urdu | 8.4 | 11.8 | 11.8 | 12.6 |
| Asturian-Pashto | 3.6 | 5.0 | 5.0 | 5.4 |
| Croatian-Afrikaans | 20.5 | 20.7 | 20.8 | 21.2 |
| Croatian-Asturian | 19.7 | 20.8 | 21.6 | 19.7 |
| Croatian-Urdu | 13.5 | 13.1 | 12.9 | 13.2 |
| Croatian-Pashto | 5.0 | 5.2 | 5.2 | 5.6 |
| Urdu-Afrikaans | 12.1 | 13.0 | 13.2 | 13.8 |
| Urdu-Asturian | 7.7 | 11.7 | 12.8 | 12.2 |
| Urdu-Croatian | 11.2 | 12.0 | 11.9 | 12.2 |
| Urdu-Pashto | 4.7 | 4.4 | 4.2 | 4.6 |
| Pashto-Afrikaans | 10.2 | 11.0 | 11.0 | 11.4 |
| Pashto-Asturian | 6.9 | 9.9 | 10.9 | 10.4 |
| Pashto-Croatian | 8.7 | 9.9 | 10.0 | 9.8 |
| Pashto-Urdu | 10.0 | 9.2 | 9.2 | 9.5 |
| Average | 12.0 | 13.1 | 13.3 | 13.4 |

Table 1: Average spBLEU (higher is better) of different pivoting methods for M2M100 and SMaLL100 on selected language pairs of FLORES-101. Best systems (not significantly outperformed by any other) in bold.
approach is closer to the translation achieved by using Spanish as the pivot language, since the NMT model is more confident for this pivot (maximum of output probability distributions for the first token of English, French, and Spanish pivots are $0.15,0.14$, and 0.91 , respectively). In contrast, the output of the MultiAvg method is a hallucination. Tables 1 and 2 show translation and hallucination performances on 20 language directions, respectively. ${ }^{9}$ TNG scores for measuring oscillatory hallucinations are provided in Appendix B. MultiAvg approach achieves better translation performance and lower hallucination compared to the direct translation. However, MultiAvg method lags behind the English pivoting approach in terms of both translation quality ( 13.4 vs. 13.1 ) and the occurrence of hallucinations ( $18.8 \%$ vs. $22.5 \%$ ). ${ }^{10}$ Applying the MaxEns method instead for combining pivots tightens this gap, and leads to better translation and reduces the occurrence of hallucinations. Specifically, MaxEns reaches competitive translation quality with English

[^4]| Language Pairs | Direct | MultiAvg | MaxEns | EN Pivot |
| :---: | :---: | :---: | :---: | :---: |
| Afrikaans-Asturian | 5.9 | 6.4 | 4.0 | 4.3 |
| Afrikaans-Croatian | 1.7 | 2.2 | 2.2 | 1.9 |
| Afrikaans-Urdu | 11.3 | 13.9 | 13.4 | 11.1 |
| Afrikaans-Pashto | 49.5 | 54.1 | 53.2 | 49.8 |
| Asturian-Afrikaans | 8.8 | 6.2 | 7.8 | 2.1 |
| Asturian-Croatian | 19.5 | 5.3 | 8.2 | 3.5 |
| Asturian-Urdu | 40.7 | 23.4 | 22.9 | 18.0 |
| Asturian-Pashto | 66.7 | 62.2 | 62.2 | 55.9 |
| Croatian-Afrikaans | 1.1 | 1.5 | 1.2 | 1.3 |
| Croatian-Asturian | 6.1 | 6.3 | 3.8 | 5.0 |
| Croatian-Urdu | 13.2 | 16.5 | 16.3 | 12.7 |
| Croatian-Pashto | 54.7 | 58.4 | 56.6 | 53.1 |
| Urdu-Afrikaans | 9.8 | 9.2 | 9.4 | 4.9 |
| Urdu-Asturian | 31.5 | 19.9 | 14.7 | 12.5 |
| Urdu-Croatian | 14.1 | 15.4 | 15.5 | 11.8 |
| Urdu-Pashto | 56.8 | 66.0 | 66.8 | 60.1 |
| Pashto-Afrikaans | 9.1 | 10.5 | 10.0 | 6.6 |
| Pashto-Asturian | 26.9 | 21.2 | 15.7 | 15.6 |
| Pashto-Croatian | 18.4 | 18.5 | 18.7 | 16.7 |
| Pashto-Urdu | 24.7 | 33.2 | 33.3 | 28.9 |
| Average | 23.5 | 22.5 | 21.8 | 18.8 |

Table 2: Average percentage ( $100 \%$ ) of hallucinations (chrF $<20$; lower is better) of different pivoting methods for M2M100 and SMaLL100 on selected language pairs of FLORES-101.
pivoting on average ( 13.3 vs. 13.4), while still underperforming on hallucinations, as most of the parallel sentences of pre-training data for M2M100 and SMaLL100 are paired with English.

In general, the optimal strategy differs across translation directions, highlighting the potential for future research on determining the most effective translation strategy for each direction without depending on the development data for each.

## 5 Conclusion

We investigate more closely the multi-source selfensembling method of Fan et al. (2021) for combining multiple translation paths to improve translations of low-resource (or very-low) language pairs. Specifically, this approach (named MultiAvg, here) averages the predictions of probability distributions of each source language in an ensemble-like manner. We evaluated it on 20 low-resource language pairs of FLORES-101 benchmark by using two massively multilingual NMT models, SMALL100 and M2M100. The MultiAvg method performs better than direct translation in terms of both translation quality and hallucinations, however it lags behind applying only English as pivot. Then, we proposed MaxEns method, a novel combination
method that chooses the maximum of prediction probabilities of pivots for each designated target token. This approach results in a better translation quality compared to MultiAvg, while reducing hallucinations. On average, it achieves competitive performance with English pivoting with regard to the translation quality metric, but performs worse with regard to the hallucination metric. The most effective translation strategy varies depending on the translation direction, suggesting the need for future research to identify the optimal strategy for each direction independently of the specific development data. We hope our findings are a starting point for the broader integration of ensemble techniques within the context of massively multilingual NMT.

The insights of our experiments, specifically the stickiness of hallucinations with different inputs, have inspired our follow-up work on source-contrastive decoding (Sennrich et al., 2024), which has empirically shown to be an effective strategy to mitigate hallucinations. Future work could revisit multi-pivot ensembling in combination with source-contrastive decoding.

## Limitations

We apply our method to two common massively multilingual NMT models, SMALL100 and M2M100; future work can extend our work to more recent large models e.g. NLLB200 (NLLB Team et al., 2022) and LLMs (Touvron et al., 2023; Workshop et al., 2022). We tested our approach on a subset of 20 low-resource language directions, future research can study the method for further language directions, including medium-resource language pairs.

## Acknowledgement

This work was funded by the Swiss National Science Foundation (project MUTAMUR; no. 176727/213976).

## References

Roee Aharoni, Melvin Johnson, and Orhan Firat. 2019. Massively multilingual neural machine translation. In Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers), pages 3874-3884, Minneapolis, Minnesota. Association for Computational Linguistics.

Naveen Arivazhagan, Ankur Bapna, Orhan Firat, Dmitry Lepikhin, Melvin Johnson, Maxim Krikun, Mia Xu Chen, Yuan Cao, George Foster, Colin Cherry, Wolfgang Macherey, Zhifeng Chen, and Yonghui Wu. 2019. Massively multilingual neural machine translation in the wild: Findings and challenges.

Yong Cheng, Qian Yang, Yang Liu, Maosong Sun, and Wei Xu. 2017. Joint training for pivot-based neural machine translation. In Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI-17, pages 3974-3980.

Anna Currey and Kenneth Heafield. 2019. Zero-resource neural machine translation with monolingual pivot data. In Proceedings of the 3rd Workshop on Neural Generation and Translation, pages 99-107, Hong Kong. Association for Computational Linguistics.

Raj Dabre, Aizhan Imankulova, Masahiro Kaneko, and Abhisek Chakrabarty. 2021. Simultaneous multi-pivot neural machine translation.

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.

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), pages 5960-5969, Online. 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, Naman Goyal, Tom Birch, Vitaliy Liptchinsky, Sergey Edunov, Michael Auli, and Armand Joulin. 2021. Beyond english-centric multilingual machine translation. Journal of Machine Learning Research, 22(107):1-48.

Orhan Firat, Kyunghyun Cho, and Yoshua Bengio. 2016a. Multi-way, multilingual neural machine translation with a shared attention mechanism. In Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, pages 866-875, San Diego, California. Association for Computational Linguistics.

Orhan Firat, Baskaran Sankaran, Yaser Al-onaizan, Fatos T. Yarman Vural, and Kyunghyun Cho. 2016b. Zero-resource translation with multi-lingual neural machine translation. In Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing, pages 268-277, Austin, Texas. Association for Computational Linguistics.

Naman Goyal, Cynthia Gao, Vishrav Chaudhary, Peng-Jen Chen, Guillaume Wenzek, Da Ju, Sanjana Krishnan, Marc' Aurelio Ranzato, Francisco Guzmán, and Angela Fan. 2022. The Flores-101 evaluation benchmark for low-resource and multilingual machine translation. Transactions of the Association for Computational Linguistics, 10:522-538.

Nuno M. Guerreiro, Duarte M. Alves, Jonas Waldendorf, Barry Haddow, Alexandra Birch, Pierre Colombo, and André F. T. Martins. 2023a. Hallucinations in Large Multilingual Translation Models. Transactions of the Association for Computational Linguistics, 11:1500-1517.

Nuno M. Guerreiro, Elena Voita, and André Martins. 2023b. Looking for a needle in a haystack: A comprehensive study of hallucinations in neural machine translation. In Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics, pages 1059-1075, Dubrovnik, Croatia. Association for Computational Linguistics.

Yunsu Kim, Petre Petrov, Pavel Petrushkov, Shahram Khadivi, and Hermann Ney. 2019. Pivot-based transfer learning for neural machine translation between non-English languages. In Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing ( $E M N L P-I J C N L P$ ), pages 866-876, Hong Kong, China. Association for Computational Linguistics.

Philipp Koehn. 2004. Statistical significance tests for machine translation evaluation. In Proceedings of the 2004 Conference on Empirical Methods in Natural Language Processing, pages 388-395, Barcelona, Spain. Association for Computational Linguistics.

Katherine Lee, Orhan Firat, Ashish Agarwal, Clara Fannjiang, and David Sussillo. 2019. Hallucinations in neural machine translation.

Dominik Macháček, Peter Polak, Ondřej Bojar, and Raj Dabre. 2023. Robustness of multi-source MT to transcription errors. In Findings of the Association for Computational Linguistics: ACL 2023, pages 3707-3723, Toronto, Canada. Association for Computational Linguistics.

Alireza Mohammadshahi, Vassilina Nikoulina, Alexandre Berard, Caroline Brun, James Henderson, and Laurent Besacier. 2022a. SMaLL-100: Introducing shallow multilingual machine translation model for low-resource languages. In Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing, pages 8348-8359, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.

Alireza Mohammadshahi, Vassilina Nikoulina, Alexandre Berard, Caroline Brun, James Henderson, and Laurent Besacier. 2022b. What do compressed multilingual machine translation models forget? In Findings of the Association for Computational

Linguistics: EMNLP 2022, pages 4308-4329, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.

Mathias Müller and Rico Sennrich. 2021. Understanding the properties of minimum Bayes risk decoding in neural machine translation. In Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers), pages 259-272, Online. Association for Computational Linguistics.

NLLB Team, Marta R. Costa-jussà, James Cross, Onur Çelebi, Maha Elbayad, Kenneth Heafield, Kevin Heffernan, Elahe Kalbassi, Janice Lam, Daniel Licht, Jean Maillard, Anna Sun, Skyler Wang, Guillaume Wenzek, Al Youngblood, Bapi Akula, Loic Barrault, Gabriel Mejia Gonzalez, Prangthip Hansanti, John Hoffman, Semarley Jarrett, Kaushik Ram Sadagopan, Dirk Rowe, Shannon Spruit, Chau Tran, Pierre Andrews, Necip Fazil Ayan, Shruti Bhosale, Sergey Edunov, Angela Fan, Cynthia Gao, Vedanuj Goswami, Francisco Guzmán, Philipp Koehn, Alexandre Mourachko, Christophe Ropers, Safiyyah Saleem, Holger Schwenk, and Jeff Wang. 2022. No language left behind: Scaling human-centered machine translation.

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.

Matt Post. 2018. A call for clarity in reporting BLEU scores. In Proceedings of the Third Conference on Machine Translation: Research Papers, pages 186-191, Brussels, Belgium. Association for Computational Linguistics.

Vikas Raunak, Arul Menezes, and Marcin JunczysDowmunt. 2021. The curious case of hallucinations in neural machine translation. In Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics. Human Language Technologies, pages 1172-1183, Online. Association for Computational Linguistics.

Vikas Raunak, Matt Post, and Arul Menezes. 2022. SALTED: A framework for SAlient long-tail translation error detection. In Findings of the Association for Computational Linguistics: EMNLP 2022, pages 5163-5179, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.

Holger Schwenk, Guillaume Wenzek, Sergey Edunov, Edouard Grave, Armand Joulin, and Angela Fan. 2021. CCMatrix: Mining billions of high-quality parallel sentences on the web. In Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers), pages 6490-6500, 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.

Rico Sennrich, Jannis Vamvas, and Alireza Mohammadshahi. 2024. Mitigating hallucinations and off-target machine translation with source-contrastive and language-contrastive decoding. In Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 2: Short Papers), pages 21-33, St. Julian's, Malta. Association for Computational Linguistics.

Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurelien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. 2023. Llama: Open and efficient foundation language models.

Masao Utiyama and Hitoshi Isahara. 2007. A comparison of pivot methods for phrase-based statistical machine translation. In Human Language Technologies 2007: The Conference of the North American Chapter of the Association for Computational Linguistics; Proceedings of the Main Conference, pages 484-491, Rochester, New York. Association for Computational Linguistics.

Guillaume Wenzek, Vishrav Chaudhary, Angela Fan, Sahir Gomez, Naman Goyal, Somya Jain, Douwe Kiela, Tristan Thrush, and Francisco Guzmán. 2021. Findings of the WMT 2021 shared task on large-scale multilingual machine translation. In Proceedings of the Sixth Conference on Machine Translation, pages 89-99, Online. Association for Computational Linguistics.

BigScience Workshop, :, Teven Le Scao, Angela Fan, Christopher Akiki, Ellie Pavlick, Suzana Ilić, Daniel Hesslow, Roman Castagné, Alexandra Sasha Luccioni, François Yvon, Matthias Gallé, Jonathan Tow, Alexander M. Rush, Stella Biderman, Albert Webson, Pawan Sasanka Ammanamanchi, Thomas Wang, Benoît Sagot, Niklas Muennighoff, Albert Villanova del Moral, Olatunji Ruwase, Rachel Bawden, Stas Bekman, Angelina McMillan-Major, Iz Beltagy, Huu Nguyen, Lucile Saulnier, Samson Tan, Pedro Ortiz Suarez, Victor Sanh, Hugo Laurençon, Yacine Jernite, Julien Launay, Margaret Mitchell, Colin Raffel, Aaron Gokaslan, Adi Simhi, Aitor Soroa, Alham Fikri Aji, Amit Alfassy, Anna Rogers, Ariel Kreisberg Nitzav, Canwen Xu, Chenghao Mou, Chris Emezue, Christopher Klamm, Colin Leong, Daniel van Strien, David Ifeoluwa Adelani, Dragomir Radev, Eduardo González Ponferrada, Efrat Levkovizh, Ethan Kim, Eyal Bar Natan, Francesco De Toni, Gérard Dupont, Germán Kruszewski, Giada Pistilli, Hady Elsahar, Hamza Benyamina, Hieu Tran, Ian Yu, Idris Abdulmumin, Isaac Johnson, Itziar Gonzalez-Dios, Javier de la Rosa, Jenny Chim, Jesse Dodge, Jian

Zhu, Jonathan Chang, Jörg Frohberg, Joseph Tobing, Joydeep Bhattacharjee, Khalid Almubarak, Kimbo Chen, Kyle Lo, Leandro Von Werra, Leon Weber, Long Phan, Loubna Ben allal, Ludovic Tanguy, Manan Dey, Manuel Romero Muñoz, Maraim Masoud, María Grandury, Mario Šaško, Max Huang, Maximin Coavoux, Mayank Singh, Mike Tian-Jian Jiang, Minh Chien Vu, Mohammad A. Jauhar, Mustafa Ghaleb, Nishant Subramani, Nora Kassner, Nurulaqilla Khamis, Olivier Nguyen, Omar Espejel, Ona de Gibert, Paulo Villegas, Peter Henderson, Pierre Colombo, Priscilla Amuok, Quentin Lhoest, Rheza Harliman, Rishi Bommasani, Roberto Luis López, Rui Ribeiro, Salomey Osei, Sampo Pyysalo, Sebastian Nagel, Shamik Bose, Shamsuddeen Hassan Muhammad, Shanya Sharma, Shayne Longpre, Somaieh Nikpoor, Stanislav Silberberg, Suhas Pai, Sydney Zink, Tiago Timponi Torrent, Timo Schick, Tristan Thrush, Valentin Danchev, Vassilina Nikoulina, Veronika Laippala, Violette Lepercq, Vrinda Prabhu, Zaid Alyafeai, Zeerak Talat, Arun Raja, Benjamin Heinzerling, Chenglei Si, Davut Emre Taşar, Elizabeth Salesky, Sabrina J. Mielke, Wilson Y. Lee, Abheesht Sharma, Andrea Santilli, Antoine Chaffin, Arnaud Stiegler, Debajyoti Datta, Eliza Szczechla, Gunjan Chhablani, Han Wang, Harshit Pandey, Hendrik Strobelt, Jason Alan Fries, Jos Rozen, Leo Gao, Lintang Sutawika, M Saiful Bari, Maged S. Al-shaibani, Matteo Manica, Nihal Nayak, Ryan Teehan, Samuel Albanie, Sheng Shen, Srulik Ben-David, Stephen H. Bach, Taewoon Kim, Tali Bers, Thibault Fevry, Trishala Neeraj, Urmish Thakker, Vikas Raunak, Xiangru Tang, Zheng-Xin Yong, Zhiqing Sun, Shaked Brody, Yallow Uri, Hadar Tojarieh, Adam Roberts, Hyung Won Chung, Jaesung Tae, Jason Phang, Ofir Press, Conglong Li, Deepak Narayanan, Hatim Bourfoune, Jared Casper, Jeff Rasley, Max Ryabinin, Mayank Mishra, Minjia Zhang, Mohammad Shoeybi, Myriam Peyrounette, Nicolas Patry, Nouamane Tazi, Omar Sanseviero, Patrick von Platen, Pierre Cornette, Pierre François Lavallée, Rémi Lacroix, Samyam Rajbhandari, Sanchit Gandhi, Shaden Smith, Stéphane Requena, Suraj Patil, Tim Dettmers, Ahmed Baruwa, Amanpreet Singh, Anastasia Cheveleva, Anne-Laure Ligozat, Arjun Subramonian, Aurélie Névéol, Charles Lovering, Dan Garrette, Deepak Tunuguntla, Ehud Reiter, Ekaterina Taktasheva, Ekaterina Voloshina, Eli Bogdanov, Genta Indra Winata, Hailey Schoelkopf, JanChristoph Kalo, Jekaterina Novikova, Jessica Zosa Forde, Jordan Clive, Jungo Kasai, Ken Kawamura, Liam Hazan, Marine Carpuat, Miruna Clinciu, Najoung Kim, Newton Cheng, Oleg Serikov, Omer Antverg, Oskar van der Wal, Rui Zhang, Ruochen Zhang, Sebastian Gehrmann, Shachar Mirkin, Shani Pais, Tatiana Shavrina, Thomas Scialom, Tian Yun, Tomasz Limisiewicz, Verena Rieser, Vitaly Protasov, Vladislav Mikhailov, Yada Pruksachatkun, Yonatan Belinkov, Zachary Bamberger, Zdeněk Kasner, Alice Rueda, Amanda Pestana, Amir Feizpour, Ammar Khan, Amy Faranak, Ana Santos, Anthony Hevia, Antigona Unldreaj, Arash Aghagol, Arezoo Abdollahi, Aycha Tammour, Azadeh HajiHosseini,

Bahareh Behroozi, Benjamin Ajibade, Bharat Saxena, Carlos Muñoz Ferrandis, Daniel McDuff, Danish Contractor, David Lansky, Davis David, Douwe Kiela, Duong A. Nguyen, Edward Tan, Emi Baylor, Ezinwanne Ozoani, Fatima Mirza, Frankline Ononiwu, Habib Rezanejad, Hessie Jones, Indrani Bhattacharya, Irene Solaiman, Irina Sedenko, Isar Nejadgholi, Jesse Passmore, Josh Seltzer, Julio Bonis Sanz, Livia Dutra, Mairon Samagaio, Maraim Elbadri, Margot Mieskes, Marissa Gerchick, Martha Akinlolu, Michael McKenna, Mike Qiu, Muhammed Ghauri, Mykola Burynok, Nafis Abrar, Nazneen Rajani, Nour Elkott, Nour Fahmy, Olanrewaju Samuel, Ran An, Rasmus Kromann, Ryan Hao, Samira Alizadeh, Sarmad Shubber, Silas Wang, Sourav Roy, Sylvain Viguier, Thanh Le, Tobi Oyebade, Trieu Le, Yoyo Yang, Zach Nguyen, Abhinav Ramesh Kashyap, Alfredo Palasciano, Alison Callahan, Anima Shukla, Antonio Miranda-Escalada, Ayush Singh, Benjamin Beilharz, Bo Wang, Caio Brito, Chenxi Zhou, Chirag Jain, Chuxin Xu, Clémentine Fourrier, Daniel León Periñán, Daniel Molano, Dian Yu, Enrique Manjavacas, Fabio Barth, Florian Fuhrimann, Gabriel Altay, Giyaseddin Bayrak, Gully Burns, Helena U. Vrabec, Imane Bello, Ishani Dash, Jihyun Kang, John Giorgi, Jonas Golde, Jose David Posada, Karthik Rangasai Sivaraman, Lokesh Bulchandani, Lu Liu, Luisa Shinzato, Madeleine Hahn de Bykhovetz, Maiko Takeuchi, Marc Pàmies, Maria A Castillo, Marianna Nezhurina, Mario Sänger, Matthias Samwald, Michael Cullan, Michael Weinberg, Michiel De Wolf, Mina Mihaljcic, Minna Liu, Moritz Freidank, Myungsun Kang, Natasha Seelam, Nathan Dahlberg, Nicholas Michio Broad, Nikolaus Muellner, Pascale Fung, Patrick Haller, Ramya Chandrasekhar, Renata Eisenberg, Robert Martin, Rodrigo Canalli, Rosaline Su, Ruisi Su, Samuel Cahyawijaya, Samuele Garda, Shlok S Deshmukh, Shubhanshu Mishra, Sid Kiblawi, Simon Ott, Sinee Sang-aroonsiri, Srishti Kumar, Stefan Schweter, Sushil Bharati, Tanmay Laud, Théo Gigant, Tomoya Kainuma, Wojciech Kusa, Yanis Labrak, Yash Shailesh Bajaj, Yash Venkatraman, Yifan Xu, Yingxin Xu, Yu Xu, Zhe Tan, Zhongli Xie, Zifan Ye, Mathilde Bras, Younes Belkada, and Thomas Wolf. 2022. Bloom: A 176b-parameter open-access multilingual language model.

Hua Wu and Haifeng Wang. 2007. Pivot language approach for phrase-based statistical machine translation. In Proceedings of the 45th Annual Meeting of the Association of Computational Linguistics, pages 856-863, Prague, Czech Republic. Association for Computational Linguistics.

Biao Zhang, Philip Williams, Ivan Titov, and Rico Sennrich. 2020. Improving massively multilingual neural machine translation and zero-shot translation. In Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, pages 1628-1639, Online. Association for Computational Linguistics.

Barret Zoph and Kevin Knight. 2016. Multi-source neural translation. In Proceedings of the 2016 Conference of the North American Chapter of the Association
for Computational Linguistics: Human Language Technologies, pages 30-34, San Diego, California. Association for Computational Linguistics.

## Appendix A Individual Scores of M2M100 and SMaLL100 Models

## A.A M2M100 Results

| Language Pairs | Direct | MultiAvg | MaxEns | EN Pivot |
| :---: | :---: | :---: | :---: | :---: |
| Afrikaans-Asturian | 19.3 | 20.2 | 21.0 | 20.2 |
| Afrikaans-Croatian | 20.8 | 21.1 | 21.0 | 21.4 |
| Afrikaans-Urdu | 14.0 | 13.6 | 13.8 | 14.4 |
| Afrikaans-Pashto | 5.4 | 5.4 | 5.5 | 5.6 |
| Asturian-Afrikaans | 14.2 | 16.5 | 16.1 | 18.0 |
| Asturian-Croatian | 11.1 | 19.0 | 18.3 | 19.4 |
| Asturian-Urdu | 6.3 | 11.4 | 11.4 | 12.6 |
| Asturian-Pashto | 2.4 | 4.4 | 4.5 | 5.0 |
| Croatian-Afrikaans | 17.6 | 17.9 | 17.9 | 18.2 |
| Croatian-Asturian | 18.8 | 19.5 | 20.6 | 19.1 |
| Croatian-Urdu | 13.6 | 13.3 | 13.1 | 13.4 |
| Croatian-Pashto | 4.4 | 4.9 | 4.9 | 5.4 |
| Urdu-Afrikaans | 9.0 | 9.8 | 10.0 | 10.6 |
| Urdu-Asturian | 7.1 | 9.9 | 10.8 | 10.9 |
| Urdu-Croatian | 8.9 | 10.0 | 9.8 | 10.1 |
| Urdu-Pashto | 4.2 | 3.8 | 3.6 | 4.2 |
| Pashto-Afrikaans | 8.3 | 9.3 | 8.9 | 9.3 |
| Pashto-Asturian | 7.8 | 9.6 | 10.4 | 9.7 |
| Pashto-Croatian | 8.0 | 9.0 | 9.0 | 8.5 |
| Pashto-Urdu | 9.8 | 8.9 | 8.9 | 9.0 |
| Average | 10.6 | 11.9 | 12.0 | 12.2 |

Table 3: spBLEU (higher is better) of different pivoting methods for M2M100 model on selected language pairs of FLORES-101 (Goyal et al., 2022) benchmark. Best systems (not significantly outperformed by any other) in bold.

| Language Pairs | Direct | MultiAvg | MaxEns | EN Pivot |
| :---: | :---: | :---: | :---: | :---: |
| Afrikaans-Asturian | 7.0 | 8.8 | 5.0 | 3.6 |
| Afrikaans-Croatian | 2.6 | 2.6 | 2.4 | 2.3 |
| Afrikaans-Urdu | 12.2 | 15.5 | 14.5 | 11.5 |
| Afrikaans-Pashto | 53.6 | 59.7 | 57.0 | 54.8 |
| Asturian-Afrikaans | 15.7 | 10.0 | 13.0 | 2.8 |
| Asturian-Croatian | 35.0 | 7.3 | 12.6 | 4.4 |
| Asturian-Urdu | 55.8 | 26.7 | 27.3 | 18.4 |
| Asturian-Pashto | 73.4 | 68.0 | 67.8 | 58.6 |
| Croatian-Afrikaans | 1.4 | 1.9 | 1.5 | 1.7 |
| Croatian-Asturian | 6.5 | 8.4 | 4.5 | 3.9 |
| Croatian-Urdu | 13.3 | 17.0 | 16.4 | 13.6 |
| Croatian-Pashto | 59.1 | 62.0 | 58.8 | 57.5 |
| Urdu-Afrikaans | 15.2 | 14.5 | 14.8 | 7.7 |
| Urdu-Asturian | 38.9 | 28.0 | 21.5 | 15.9 |
| Urdu-Croatian | 20.3 | 22.3 | 22.2 | 17.9 |
| Urdu-Pashto | 60.2 | 72.7 | 73.3 | 65.1 |
| Pashto-Afrikaans | 11.5 | 12.8 | 12.7 | 9.0 |
| Pashto-Asturian | 23.1 | 25.1 | 17.9 | 18.0 |
| Pashto-Croatian | 20.5 | 21.6 | 22.5 | 21.1 |
| Pashto-Urdu | 26.6 | 35.9 | 35.5 | 31.9 |
| Average | 27.6 | 26.0 | 25.0 | 21.0 |

Table 4: The percentage of hallucinations ( $\operatorname{chrF}<20$; lower is better) of different pivoting methods for M2M100 model on selected language pairs of FLORES-101 (Goyal et al., 2022) benchmark.

## A.B SMaLL100 Results

| Language Pairs | Direct |  |  |  |  | MultiAvg | MaxEns | EN Pivot |
| :--- | ---: | ---: | ---: | ---: | :---: | :---: | :---: | :---: |
| Afrikaans-Asturian | 22.0 | 23.4 | $\mathbf{2 4 . 0}$ | 21.7 |  |  |  |  |
| Afrikaans-Croatian | 23.5 | $\mathbf{2 3 . 9}$ | $\mathbf{2 3 . 8}$ | $\mathbf{2 4 . 1}$ |  |  |  |  |
| Afrikaans-Urdu | 13.9 | 14.0 | 14.0 | $\mathbf{1 4 . 4}$ |  |  |  |  |
| Afrikaans-Pashto | $\mathbf{6 . 4}$ | 6.1 | $\mathbf{6 . 4}$ | $\mathbf{6 . 4}$ |  |  |  |  |
| Asturian-Afrikaans | 22.8 | $\mathbf{2 3 . 3}$ | $\mathbf{2 3 . 4}$ | $\mathbf{2 3 . 8}$ |  |  |  |  |
| Asturian-Croatian | 21.1 | $\mathbf{2 1 . 8}$ | $\mathbf{2 1 . 6}$ | $\mathbf{2 1 . 4}$ |  |  |  |  |
| Asturian-Urdu | 10.5 | 12.1 | 12.2 | $\mathbf{1 2 . 5}$ |  |  |  |  |
| Asturian-Pashto | 4.8 | $\mathbf{5 . 6}$ | $\mathbf{5 . 5}$ | $\mathbf{5 . 7}$ |  |  |  |  |
| Croatian-Afrikaans | 23.4 | 23.4 | 23.7 | $\mathbf{2 4 . 2}$ |  |  |  |  |
| Croatian-Asturian | 20.6 | $\mathbf{2 2 . 1}$ | $\mathbf{2 2 . 5}$ | 20.2 |  |  |  |  |
| Croatian-Urdu | $\mathbf{1 3 . 3}$ | 12.8 | 12.6 | $\mathbf{1 3 . 0}$ |  |  |  |  |
| Croatian-Pashto | $\mathbf{5 . 6}$ | $\mathbf{5 . 4}$ | $\mathbf{5 . 4}$ | $\mathbf{5 . 8}$ |  |  |  |  |
| Urdu-Afrikaans | 15.1 | 16.1 | 16.3 | $\mathbf{1 7 . 0}$ |  |  |  |  |
| Urdu-Asturian | 8.3 | 13.4 | $\mathbf{1 4 . 8}$ | 13.4 |  |  |  |  |
| Urdu-Croatian | 13.4 | 14.0 | 13.9 | 14.3 |  |  |  |  |
| Urdu-Pashto | $\mathbf{5 . 1}$ | $\mathbf{5 . 0}$ | $\mathbf{4 . 8}$ | $\mathbf{5 . 0}$ |  |  |  |  |
| Pashto-Afrikaans | 12.0 | 12.7 | 12.9 | $\mathbf{1 3 . 5}$ |  |  |  |  |
| Pashto-Asturian | 6.0 | 10.2 | $\mathbf{1 1 . 3}$ | $\mathbf{1 1 . 1}$ |  |  |  |  |
| Pashto-Croatian | 9.4 | $\mathbf{1 0 . 8}$ | $\mathbf{1 1 . 0}$ | $\mathbf{1 1 . 0}$ |  |  |  |  |
| Pashto-Urdu | $\mathbf{1 0 . 2}$ | 9.5 | 9.4 | 9.9 |  |  |  |  |
| Average | 13.4 | 14.3 | $\mathbf{1 4 . 5}$ | $\mathbf{1 4 . 4}$ |  |  |  |  |

Table 5: spBLEU (higher is better) of different pivoting methods for SMaLL100 model on selected language pairs of FLORES-101 (Goyal et al., 2022) benchmark. Best systems (not significantly outperformed by any other) in bold.

| Language Pairs | Direct | MultiAvg | MaxEns | EN Pivot |
| :--- | ---: | ---: | ---: | ---: |
| Afrikaans-Asturian | 4.7 | 4.1 |  | 2.9 |
| Afrikaans-Croatian | 0.9 | 1.9 | 2.1 | 1.5 |
| Afrikaans-Urdu | 10.4 | 12.3 | 12.2 | 10.8 |
| Afrikaans-Pashto | 45.4 | 48.4 | 49.4 | 44.8 |
| Asturian-Afrikaans | 1.9 | 2.5 | 2.6 | 1.5 |
| Asturian-Croatian | 4.1 | 3.3 | 3.9 | 2.7 |
| Asturian-Urdu | 25.6 | 20.1 | 18.5 | 17.5 |
| Asturian-Pashto | 59.9 | 56.5 | 56.7 | 53.3 |
| Croatian-Afrikaans | 0.8 | 1.1 | 0.9 | 0.9 |
| Croatian-Asturian | 5.6 | 4.2 | 3.1 | 6.1 |
| Croatian-Urdu | 13.2 | 16.0 | 16.1 | 11.8 |
| Croatian-Pashto | 50.3 | 54.7 | 54.5 | 48.7 |
| Urdu-Afrikaans | 4.4 | 4.0 | 4.0 | 2.1 |
| Urdu-Asturian | 24.1 | 11.8 | 7.9 | 9.0 |
| Urdu-Croatian | 7.9 | 8.5 | 8.8 | 5.7 |
| Urdu-Pashto | 53.4 | 59.3 | 60.3 | 55.2 |
| Pashto-Afrikaans | 6.7 | 8.3 | 7.3 | 4.2 |
| Pashto-Asturian | 30.7 | 17.3 | 13.6 | 13.2 |
| Pashto-Croatian | 16.4 | 15.4 | 14.9 | 12.4 |
| Pashto-Urdu | 22.7 | 30.5 | 31.1 | 25.9 |
| Average | 19.5 | 19.0 | 18.5 | 16.6 |

Table 6: The percentage of hallucinations ( $\operatorname{chrF}<20$; lower is better) of different pivoting methods for SMaLL100 model on selected language pairs of FLORES-101 (Goyal et al., 2022) benchmark.

## Appendix B Results of oscillatory hallucinations based on TNG metric

| Language Pairs | Direct |  |  | MultiAvg |
| :--- | ---: | :---: | :---: | :---: |
| MaxEns | EN Pivot |  |  |  |
| Afrikaans-Asturian | 2.2 | 2.1 | 0.7 | 1.3 |
| Afrikaans-Croatian | 0.3 | 0.3 | 0.2 | 0.2 |
| Afrikaans-Urdu | 1.2 | 1.5 | 1.1 | 0.6 |
| Afrikaans-Pashto | 6.7 | 6.7 | 5.1 | 4.7 |
| Asturian-Afrikaans | 5.1 | 3.0 | 3.9 | 0.2 |
| Asturian-Croatian | 11.9 | 1.5 | 3.0 | 0.4 |
| Asturian-Urdu | 15.7 | 4.1 | 4.3 | 0.7 |
| Asturian-Pashto | 20.9 | 12.8 | 11.6 | 5.0 |
| Croatian-Afrikaans | 0.1 | 0.2 | 0.2 | 0.0 |
| Croatian-Asturian | 2.2 | 2.0 | 1.0 | 1.2 |
| Croatian-Urdu | 1.5 | 1.3 | 1.1 | 0.5 |
| Croatian-Pashto | 7.2 | 6.4 | 4.9 | 4.3 |
| Urdu-Afrikaans | 8.2 | 3.3 | 3.3 | 0.9 |
| Urdu-Asturian | 15.8 | 6.2 | 4.0 | 2.4 |
| Urdu-Croatian | 2.5 | 2.6 | 2.8 | 0.8 |
| Urdu-Pashto | 9.3 | 11.7 | 9.5 | 4.2 |
| Pashto-Afrikaans | 2.7 | 3.6 | 3.5 | 0.8 |
| Pashto-Asturian | 16.0 | 6.2 | 4.1 | 2.0 |
| Pashto-Croatian | 3.6 | 2.9 | 3.0 | 1.0 |
| Pashto-Urdu | 2.5 | 4.8 | 4.3 | 1.6 |
| Average | 7.3 | 4.5 | 4.0 | 1.9 |

Table 7: Average percentage (100\%) of hallucinations (TNG metric; lower is better) of different pivoting methods for M2M100 and SMaLL100 on selected language pairs of FLORES-101.

| Language Pairs | Direct |  |  |  |
| :--- | ---: | ---: | ---: | ---: |
| MultiAvg | MaxEns | EN Pivot |  |  |
| Afrikaans-Asturian | 1.8 | 2.6 |  | 1.1 |
| Afrikaans-Croatian | 0.5 | 0.49 | 0.2 | 0.5 |
| Afrikaans-Urdu | 1.9 | 2.7 | 1.6 | 0.9 |
| Afrikaans-Pashto | 10.4 | 11.6 | 9.1 | 8.0 |
| Asturian-Afrikaans | 9.8 | 5.6 | 6.8 | 0.2 |
| Asturian-Croatian | 22.6 | 2.5 | 5.2 | 0.4 |
| Asturian-Urdu | 25.3 | 7.6 | 7.7 | 1.2 |
| Asturian-Pashto | 36.8 | 23.6 | 21.3 | 9.2 |
| Croatian-Afrikaans | 0.1 | 0.2 | 0.2 | 0.0 |
| Croatian-Asturian | 2.3 | 2.7 | 1.1 | 0.2 |
| Croatian-Urdu | 2.2 | 2.1 | 1.7 | 0.7 |
| Croatian-Pashto | 11.8 | 11.1 | 8.5 | 7.5 |
| Urdu-Afrikaans | 15.2 | 6.1 | 6.1 | 1.5 |
| Urdu-Asturian | 18.2 | 9.0 | 6.8 | 2.1 |
| Urdu-Croatian | 4.1 | 4.8 | 4.9 | 1.2 |
| Urdu-Pashto | 13.3 | 21.0 | 16.8 | 7.4 |
| Pashto-Afrikaans | 4.2 | 4.3 | 4.7 | 1.1 |
| Pashto-Asturian | 12.7 | 6.9 | 5.1 | 1.2 |
| Pashto-Croatian | 4.2 | 3.7 | 3.7 | 1.6 |
| Pashto-Urdu | 3.3 | 7.5 | 6.2 | 2.8 |
| Average | 10.9 | 7.6 | 6.7 | 2.8 |

Table 8: The percentage of hallucinations (TNG metric; lower is better) of different pivoting methods for M2M100 model on selected language pairs of FLORES-101 (Goyal et al., 2022) benchmark.

| Language Pairs | Direct |  | MultiAvg | MaxEns |
| :--- | ---: | :---: | :---: | :---: |
| EN Pivot |  |  |  |  |
| Afrikaans-Asturian | 2.5 | 1.6 | 0.2 | 2.1 |
| Afrikaans-Croatian | 0.0 | 0.1 | 0.2 | 0.1 |
| Afrikaans-Urdu | 0.4 | 0.3 | 0.6 | 0.2 |
| Afrikaans-Pashto | 2.9 | 1.8 | 1.2 | 1.3 |
| Asturian-Afrikaans | 0.4 | 0.5 | 0.9 | 0.2 |
| Asturian-Croatian | 1.2 | 0.4 | 0.7 | 0.4 |
| Asturian-Urdu | 6.1 | 0.7 | 0.9 | 0.2 |
| Asturian-Pashto | 5.0 | 2.0 | 1.9 | 0.9 |
| Croatian-Afrikaans | 0.0 | 0.1 | 0.1 | 0.0 |
| Croatian-Asturian | 2.1 | 1.3 | 0.8 | 2.2 |
| Croatian-Urdu | 0.8 | 0.4 | 0.4 | 0.3 |
| Croatian-Pashto | 2.6 | 1.7 | 1.3 | 1.2 |
| Urdu-Afrikaans | 1.1 | 0.6 | 0.4 | 0.3 |
| Urdu-Asturian | 13.4 | 3.3 | 1.2 | 2.6 |
| Urdu-Croatian | 0.8 | 0.4 | 0.7 | 0.4 |
| Urdu-Pashto | 5.3 | 2.3 | 2.1 | 1.0 |
| Pashto-Afrikaans | 1.2 | 2.9 | 2.2 | 0.4 |
| Pashto-Asturian | 19.4 | 5.4 | 3.1 | 2.7 |
| Pashto-Croatian | 3.0 | 2.1 | 2.3 | 0.3 |
| Pashto-Urdu | 1.6 | 2.2 | 2.4 | 0.4 |
| Average | 3.6 | 1.5 | 1.2 | 0.9 |

Table 9: The percentage of hallucinations (TNG metric; lower is better) of different pivoting methods for SMALL100 model on selected language pairs of FLORES-101 (Goyal et al., 2022) benchmark.


[^0]:    *Work done while working at University of Zurich. Currently co-founder of Leeroo.
    ${ }^{1}$ The implementation is publicly available at https://github.com/ZurichNLP/MultiPivotNMT.

[^1]:    ${ }^{2}$ SentencePiece BLEU of $63 \%$ translation directions in M2M-100 is lower than 12 (Mohammadshahi et al., 2022b).

[^2]:    ${ }^{3}$ We tried both averaging probabilities and log-probabilities in preliminary experiments, and averaging probabilities worked better in terms of translation performance and hallucination.

[^3]:    ${ }^{4}$ BLEU is computed after tokenization with SentencePiece with 256 K tokens (Goyal et al., 2022).
    ${ }^{5}$ sacrebleu 2.3.1 (Post, 2018) with ChrF3 is used.
    ${ }^{6}$ Threshold based on manual inspection.
    ${ }^{7}$ We follow Guerreiro et al. (2023b) and use $n=4$ and $t=2$.
    ${ }^{8}$ Paired bootstrap resampling (Koehn, 2004) with sacrebleu.

[^4]:    ${ }^{9}$ Average scores of M2M100 and SMaLL100 are shown in Table 1 and 2. Individual scores are provided in Appendix A.
    ${ }^{10} 4.5 \%$ vs. $7.3 \%$ based on the TNG metric, as shown in Table 7.

