

# On-the-Fly Fusion of Large Language Models and Machine Translation

Hieu Hoang Huda Khayrallah Marcin Junczys-Dowmunt

Microsoft, 1 Microsoft Way, Redmond, WA 98052, USA  
{hihoan,hkhayrallah,marcinjd}@microsoft.com

## Abstract

We propose on-the-fly ensembling of a neural machine translation (NMT) model with a large language model (LLM), prompted on the same task and input. Through experiments on 4 language directions with varying data amounts, we find that a slightly weaker-at-translation LLM can improve translations of a NMT model, and such an ensemble can produce better translations than ensembling two stronger NMT models. We demonstrate that our ensemble method can be combined with various techniques from LLM prompting, such as in context learning and translation context.

## 1 Introduction

For many English NLP tasks, LLMs (Brown et al., 2020; Smith et al., 2022; Chowdhery et al., 2022; Touvron et al., 2023a,b) are the clear state-of-the-art—e.g. sentiment analysis (Zhang et al., 2023c), summarization (Zhang et al., 2023b). However, dedicated NMT outperforms all but the largest closed source LLMs (Jiao et al., 2023) and dedicated MT is stronger in low resource settings (Hendy et al., 2023; Robinson et al., 2023).

We propose a novel integration of a LLM and dedicated NMT model via token-level fusion. This ensembling combines strengths of each model, which emerge from their differences. LLMs are trained on more data than NMT models, and have more parameters. While LLMs are exposed to some parallel data (Briakou et al., 2023), they are trained on vastly more monolingual data, which likely gives them different domain coverage and more fluency than dedicated models. NMT models are trained on the translation task. For example, Jiao et al. (2023) found ChatGPT is more likely to hallucinate but is stronger at translating the spoken domain, while dedicated models are stronger for medical domains and social-media-style noisy text. LLMs can easily be prompted with auxiliary information—such as domain and document context—while that is more complicated for NMT.

In this work we:

- propose on-the-fly ensembling of an MT model with a prompted-for-translation LLM,
- combine it with domain and context prompting,
- demonstrate that a weaker-at-translation LLM can improve translations of a MT model,
- and demonstrate our method is better than MT ensembles and ensembles with non-prompted LLMs.

## 2 Method

We review standard inference of encoder-decoder NMT models and decoder only LLMs and then introduce our proposed ensemble of the two.

**Standard Decoding** In encoder-decoder NMT, the probability of token  $t$  at the  $i^{th}$  time step is:

$$p_{MT}(t_i) = p_{MT}(t_i | t_{j < i}, S) \quad (1)$$

This conditions on source sentence  $S$  as the input to the encoder and  $t_{j < i}$  as the previously generated target tokens in the MT model decoder.

When using a decoder only LLM for translation, the probability of token  $t$  at the  $i^{th}$  time step is:

$$p_{LLM}(t_i) = p_{LLM}(t_i | M, S, t_{j < i}) \quad (2)$$

The concatenation of the prompt  $M$ , source sentence  $S$  and the previous generated targets are all decoder outputs. The LLM model is prefix-decoded through the prompt and source, and then allowed to produce the target tokens. The LLM prompt  $M$  can also include additional content.

**Proposed Ensemble** When combining the two for our ensemble, we have:

$$p_{ensemble}(t_i) = \lambda p_{MT}(t_i) + (1 - \lambda) p_{LLM}(t_i) \quad (3)$$

In the ensemble,  $p_{MT}$  and  $p_{LLM}$  condition on the tokens previously generated by the ensemble.  $p_{LLM}$  still conditions on the prompt, which can be used to infuse the model with auxiliary information (e.g. domain or context).  $p_{ensemble}$  reduces to the LLM when  $\lambda = 0$  and to the MT model when  $\lambda = 1$ .

	German	Russian	Turkish	Hausa
Train	290.4m	38m	49.5m	600k
Valid	1000/1002 WMT21	1000/1002 WMT21	3007 newstest2017	2000 newsdev2021
Test	1984/2037 WMT22	2016/2037 WMT22	3000/3602 newstest2018	4456/4459 newstest2021
TED-100	-	1132	-	-
ParaPat	2000	2000	-	-
CTXPro	2000	2000	-	-

Table 1: Size of datasets used in this work. All numbers are in sentences, except for CTXPro, which is reported in paragraphs. For the validation and testsets that are different in each translation direction, numbers listed are for  $* \rightarrow \text{en} / \text{en} \leftarrow *$ .

### 3 Experimental Setup

We aim to understand how our proposed method performs in high and low resource settings with strong models, and design our experimental setup accordingly.

The parallel and monolingual training data for German and Russian is from the WMT22 (Kocmi et al., 2022)<sup>1</sup> shared task. The Hausa data is from WMT21 (Akhbardeh et al., 2021).<sup>2</sup> The Turkish evaluation data was based on WMT18 (Bojar et al., 2018)<sup>3</sup> and training data also includes additional data from OPUS (Tiedemann, 2012), excluding Paracrawl (Bañón et al., 2020), since such noisy data (Khayrallah and Koehn, 2018) would require filtering (Koehn et al., 2018, 2019, 2020; Sloto et al., 2023).

As domain-specific test sets we use TED-100 (Salesky et al., 2021) and ParaPat (Soares et al., 2020). We also use TED-100 and CTXPro (Wicks and Post, 2023) for document-level experiments.<sup>4</sup>

Table 1 summarizes the parallel training, evaluation and test data and Table 6 in the Appendix summarizes the monolingual data.

We use back translation (Sennrich et al., 2016) (with a 1:1 ratio of parallel to synthetic data) for all language pairs. We train Transformer ‘big’ models for German, Russian and Turkish, and ‘base’ for Hausa (Vaswani et al., 2017) in Marian NMT (Junczys-Dowmunt et al., 2018).<sup>5</sup> We use Llama2 (Touvron et al., 2023b) with 7 and 13 billion pa-

<sup>1</sup><https://www.statmt.org/wmt22/>

<sup>2</sup><https://www.statmt.org/wmt21/>

<sup>3</sup><https://www.statmt.org/wmt18/>

<sup>4</sup>For CTXPro, we select a random sample of 2000 paragraphs for our experiments to reduce compute usage.

<sup>5</sup>We convert models from Marian to Hugging Face format.

rameters as LLMs. The Llama2 32k token SentencePiece model (Kudo and Richardson, 2018) is used for source and target MT tokenization.<sup>6</sup>

The optimal mixing ratio is learnt using grid search  $\lambda \in \{0, 0.1, \dots, 1\}$  on the validation set. We use this same value of  $\lambda$  in domain specific experiments in § 5; we do not re-sweep for each domain. Final results are reported on the test sets, translation quality is measured using COMET-22 (Rei et al., 2022). We use greedy search for decoding. See § A for additional experimental details, including prompts.

### 4 Results

Table 2 shows the translation quality of the ensemble using the 7 billion parameters LLM (col. 1-6). When both models are of reasonable quality (de-en, ru-en, en-ru), ensembling (col. 5) results in better quality than either alone (col. 1 & 2).

In all cases, the LLM quality is worse than the MT model but ensembling with it improves most language directions. For de-en, the MT model is 0.9 COMET stronger than the LLM. The ensemble still improves over the MT model by 0.6 COMET.

The improvement is minor for en-de, where the LLM was 21.9 points worse than MT. The LLM translation quality for Turkish in both direction is poor while the MT is good so the ensembles are essentially reduced to the MT model. Both models are bad for Hausa and the ensembles are unusable. § A.4 shows the effect of  $\lambda$  on translation quality.

**In-context learning:** Table 2 (col. 3) shows 5-shot learning tends to improve LLM quality but has little affect on the ensemble (col. 6).

**Larger LLM:** Xu et al. (2023) found that Llama-13B suffers from off-target issues, degrading translation out-of-English compared to the 7B model. We confirm their results—Table 2 (col. 2 vs 7)—and also reproduce their solution of using 5-shot learning, which can recover and sometime improve LLM quality (col. 8). However, ensembling with the MT model does not require the use of in context learning (col. 10 vs 11). In general, the larger language model is better for the ensemble as de-en, en-de and ru-en all improve. It should also be noted that the MT model adds, at most, 3% to the number of parameters of the 7B LLM allowing the

<sup>6</sup>The target side vocabs must match between the LLM and MT model to be able to ensemble; the source could potentially be different. Preliminary experiments, however, found it better to use the same vocab and be able to tie the embeddings.

	MT	LLM 7B		Ensemble w/ LLM 7B			LLM 13B		Ensemble w/ LLM 13B		
		0-shot	5-shot	$\lambda$	0-shot	5-shot	0-shot	5-shot	$\lambda$	0-shot	5-shot
column:	1	2	3	4	5	6	7	8	9	10	11
de-en	83.5	82.6	82.8	0.7	83.9	83.9	82.6	83.4	0.7	<b>84.1</b>	84.0
en-de	85.4	79.4	79.8	0.7	85.5	85.5	63.4	82.4	0.8	<b>85.6</b>	<b>85.6</b>
ru-en	82.8	82.5	82.5	0.5	84.0	84.1	81.4	83.4	0.5	84.2	<b>84.5</b>
en-ru	83.1	80.4	81.1	0.5	83.9	<b>84.2</b>	36.4	81.1	0.8	83.6	83.7
tr-en	87.2	75.2	75.7	0.8	87.2	87.2	78.9	-	0.8	<b>87.3</b>	<b>87.3</b>
en-tr	89.4	57.8	58.2	1.0	89.4	89.4	40.3	69.4	0.9	89.4	<b>89.5</b>
ha-en	<b>60.1</b>	47.0	49.3	0.3	54.7	54.7	46.9	49.7	0.3	54.7	54.5
en-ha	<b>63.1</b>	33.1	37.6	1.0	<b>63.1</b>	<b>63.1</b>	38.2	35.7	1.0	<b>63.1</b>	<b>63.1</b>

Table 2: COMET-22 on WMT test sets. Ensembling MT & LLM can improve scores in high resource settings where the LLM’s COMET is somewhat worse than the MT.  $\lambda$  is the mixing rate; higher  $\lambda$  puts more emphasis on MT.

ensemble to outperform the nearly 2x bigger 13B LLM.

Ensembles for Turkish and Hausa are still not worthwhile due to the poor LLM quality in these lower resource settings. We use the 7B model in all analysis for the remainder of this work.

## 5 Analysis

### 5.1 MT Model Ensembling

Given the compute resource required to use LLMs (not to mention train them), we compare the results of the MT + LLM ensemble to ensembling two MT models. We create ensembles for German and Russian language pairs consisting of two MT models.<sup>7</sup> As Table 3 shows, using the LLM gives stronger translation quality in all cases except en-de, which is where the LLM underperforms the MT model by 6 COMET points. In all the other situations, *it is better to ensemble the MT model with an LLM, even though the 2<sup>nd</sup> MT model has higher translation quality than the LLM by 0.5 to 2.8 COMET*. This suggests that when selecting models for an ensemble, simply choosing the two highest quality models is insufficient. Instead, ensembling takes advantage of the training diversity in the models to improve quality.

### 5.2 Mixing Ratio Interpretation

The learnt mixing ratio,  $\lambda$ , can be loosely interpreted as a relative utility of the underlying models. For ensembles with German and Russian,  $\lambda$  of 0.7 and 0.5 for the 7B LLM ensemble reflect the nearly equal contribution of both models. Due to

<sup>7</sup>The models differ only in the random seed.

	MT	LLM	MT+LLM	MT+MT	
de-en	83.5	83.7	82.6	<b>83.9</b>	83.8
en-de	85.4	85.4	79.4	85.5	<b>85.7</b>
ru-en	82.8	83.0	82.5	<b>84.0</b>	83.1
en-ru	83.1	83.2	80.4	<b>83.9</b>	83.4

Table 3: COMET-22 score for two MT replicas, the LLM, the MT & LLM ensemble, and the ensemble of the two MT models. The ensembling of the LLM with the MT model has the highest COMET score in all but one language pair, even though both the MT models have higher translation quality than the LLM.

	prompt:	MT	LLM		Ensemble	
		none	general	+domain	general	+domain
TED	ru-en	77.3	78.0	78.5	78.7	<b>78.9</b>
	de-en	79.7	77.1	78.0	<b>80.0</b>	<b>80.0</b>
ParaPat	en-de	79.1	73.8	73.8	<b>79.2</b>	<b>79.2</b>
	ru-en	72.2	74.5	73.9	<b>75.1</b>	75.0
	en-ru	78.5	73.7	73.4	<b>79.0</b>	78.7

Table 4: Prompting with domain can improve COMET-22 for the LLM, but is less effective for the ensemble.

off-target issue described above, the 13B LLM are poor at translating into German and Russian so its contribution to the ensemble is reduced.

For Turkish and Hausa, the LLM offer negligible benefit so most weight is given to the MT model. The mixing ratio space for Hausa-English is flat (see Figure 7(g)) as both underlying models are equality poor so no interpretation should be attached to the results.

### 5.3 Domain Prompting

The flexibility of LLM prompting can be used to add more descriptive task-specific instructions to improve quality (Zhang et al., 2023a). Here, we prompt for domain (TED talks and patents).

Table 4 shows that additional domain information does not guarantee better LLM quality. For the TED-100 test set, ensembling has a 0.2 COMET improvement from an 0.5 LLM increase. Ensembling with or without the domain information in the prompt outperforms either the MT and LLM models alone. For TED, the LLM is stronger than the dedicated MT models, in contrast to our main results.<sup>8</sup> While our dedicated MT models were not trained translation for this specific domain, the LLM likely exposed to monolingual data in this domain. This highlights the complementary strengths of each paradigm—the ensemble leverages both.

### 5.4 Document Context

For document or discourse input—such as TED talks—where the previous translated sentences are often relevant to the sentence to be translated, it may be better to provide the previous sentences and their translation. This contrasts with few-shot prompting where sentences pairs are high quality translations written by humans but are drawn from the validation so may not be relevant to the sentence at hand. Using sentence pairs from the same document should allow the LLM to enforce consistency across sentences and allow it to better translate phenomena that requires document-level context such as pronoun disambiguation.

Figure 1 shows the COMET-22 score against the number of sentence pairs in the prompt on the TED-100 test set. Prompting the LLM with document context outperforms few-shot prompting and the ensemble with context (solid orange line) to outperform all variants of ensembling and LLMs with context or few-shots, as well as the MT model. *Conditioning on the model’s own previous outputs from the same document context outperforms few-shot prompting with the human references of less related sentences.*

Prior work found that document level-specific evaluation is required to evaluate document level phenomena (Läubli et al., 2018; Toral et al., 2018; Vernikos et al., 2022). To this end, we use CTXPro (Wicks and Post, 2023), a specialized test suite

<sup>8</sup>In this work, we used the  $\lambda$  set on the general validation set. Re-sweeping for each specific domain could lead to improved performance.




Figure 1: TED-100 translation quality for various number of prompt examples (for few shot learning or past context). Prompting with context outperforms few shot prompting, and it performs best when ensembled.

		MT	LLM		Ensemble	
context:		none	none	10 sent	none	10 sent
en-de	auxiliary	4.5%	7.2%	<b>28.0%</b>	6.2%	13.7%
	formality	41.9%	38.2%	37.6%	42.7%	<b>43.8%</b>
	gender	44.6%	38.5%	39.0%	<b>45.8%</b>	45.5%
en-ru	auxiliary	2.6%	2.3%	<b>24.6%</b>	2.6%	20.9%
	formality	42.5%	42.6%	46.4%	46.4%	<b>50.0%</b>
	gender	27.4%	31.9%	36.4%	31.6%	<b>37.6%</b>
	inflection	28.9%	22.6%	25.6%	29.2%	<b>31.4%</b>

Table 5: CTXPro accuracy. The ensembled models with context perform particularly well in to Russian.

which evaluates the translation accuracy of targeted words, given the document context.

Table 5 shows accuracy for various phenomena on en-ru and en-de. Adding context improves accuracy in all-but-one test set. Ensembling with context has the highest accuracy in 4 of 7 models. See § A.5 for COMET on this data; the ensemble is always best. So, when balancing COMET and CTXPro accuracy, the ensemble is best.

### 5.5 Unprompted LM Ensembling

Yee et al. (2019) and Petrick et al. (2023) improve translation by ensembling with a smaller-scale language model *without* a task-specific prompt.

We test this by ensembling the MT model with an unprompted LLM. Figure 2 shows that this causes quality to drop precipitously. The divergence from prior work may be due to differences in the base models; for example, Petrick et al. (2023) used an MT model trained on small amount of data, and Yee et al. (2019) trained their own LM. In our




Figure 2: Using an LLM with a translation prompt and without any prompting (ru-en). Unprompted the ensemble is strictly worse than the MT baseline (mixing ratio  $\lambda = 1$ ).

scenario with a strong MT model and a general purpose LLM, we do not see any benefit from using the LLM purely as a language model.

## 6 Related work

**LLMs for MT:** Pretrained LLMs can be prompted directly for translation (Brown et al., 2020; Vilar et al., 2023; Hendy et al., 2023; Robinson et al., 2023; Zhang et al., 2023a; Agrawal et al., 2023), or fine-tuned for MT (Li et al., 2023; Chen et al., 2023; Moslem et al., 2023; Zeng et al., 2023; Xu et al., 2023; Yang et al., 2023). Our approach is complimentary—we leverage prompting and in-context learning. We could also ensemble with a fine-tuned model. Since we perform inference-time combination of the LLM, we do not have the same training-compute burden as fine-tuning.

Much work has explored integrating language models and NMT in various ways (Gulcehre et al., 2015, 2017; Stahlberg et al., 2018; Yee et al., 2019; Petrick et al., 2023), mostly by purely conditioning a language model on the target tokens; in contrast we focus on pretrained LLMs and *prompt* the LLM to produce translations.

**Ensembling:** Diverse inputs can be combined to create stronger ensembles (Hansen and Salamon, 1990; Dietterich, 2000). Various model-combination methods have been used in MT.

System combination of outputs was used for statistical machine translation (SMT) (Bangalore et al., 2001; Heafield and Lavie, 2010; Freitag et al., 2014), and averaging model weights (Junczys-Dowmunt et al., 2016) or ensembling (Chung et al., 2016) are used for NMT. We build upon the latter. Jiang et al. (2023) propose a separate model to combine outputs from LLMs. We ensemble on-the-

fly. Ormazabal et al. (2023) ensemble two LLM from the same family where the smaller LM was finetuned for MT. We create a hybrid ensemble of two distinct architectures and training regimes.

Knowledge distillation (Buciluă et al., 2006; Hinton et al., 2015) inspired methods can be a way to incorporate diverse models during training (Dakwale and Monz, 2017; Khayrallah et al., 2018, 2020), as opposed to during inference. Jiang et al. (2023) introduce a separate model that combines outputs from LLMs. We ensemble on-the-fly.

There are various methods proposed for improving translation quality by combining the adequacy and fluency advantages of SMT and NMT (Devlin et al., 2014; Mi et al., 2016; Junczys-Dowmunt et al., 2016; Stahlberg et al., 2017; Wang et al., 2017; Khayrallah et al., 2017; Ding et al., 2017; Zhang et al., 2021). We combine the strengths of NMT and LLMs.

## 7 Conclusion

We propose an on-the-fly ensembling of a dedicated MT model with an LLM, conditioned on the source and prompted for translation. We demonstrate that an LLM can improve translation quality of a NMT model even if the LLM is weaker at translation, provided the LLM is good enough. We prompt the LLM to imbue the sentence-based MT model with document-level ability, improving on sentence-level and context-focused metrics. We find that ensembling with an LLM performs better than ensembling two MT models, even if each MT model is stronger than the LLM.

While this work focuses on MT, the same techniques can be explored for other tasks, and may be especially useful for situations where the LLM and task-specific model have different properties and strengths.

## 8 Limitations

While we covered four languages to and from English, this is nowhere near enough to be a representative sample of languages and translation directions that would be of interest to others. We used Llama2; there are closed-access models that may be stronger at translation (e.g. GPT-4) but API access is insufficient for this method. As open-source new models are released, this method can be applied to them as well.

We used a single value of  $\lambda$ —which was set on the general domain validation set—for all ex-

periments. We did not re-sweep for each domain. While this is a more general scenario that applies when test-time domain is unknown, results might be improved for focused domains by tuning  $\lambda$  on domain-specific validation sets.

In § 5, we explore different domains (TED talks, subtitles, and patents), and use COMET-22 as a metric. Zouhar et al. (2024) recently demonstrated that neural fine-tuned metrics, such as COMET are not robust to domain shift, but noted that COMET still had the highest overall correlation with human judgements in their domain of study.

## References

- Sweta Agrawal, Chunting Zhou, Mike Lewis, Luke Zettlemoyer, and Marjan Ghazvininejad. 2023. [In-context examples selection for machine translation](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 8857–8873, Toronto, Canada. Association for Computational Linguistics.
- Farhad Akhbardeh, Arkady Arkhangorodsky, Magdalena Biesialska, Ondřej Bojar, Rajen Chatterjee, Vishrav Chaudhary, Marta R. Costa-jussa, Cristina España-Bonet, Angela Fan, Christian Federmann, Markus Freitag, Yvette Graham, Roman Grundkiewicz, Barry Haddow, Leonie Harter, Kenneth Heafield, Christopher Homan, Matthias Huck, Kwabena Amponsah-Kaakyire, Jungo Kasai, Daniel Khashabi, Kevin Knight, Tom Koemi, Philipp Koehn, Nicholas Lourie, Christof Monz, Makoto Morishita, Masaaki Nagata, Ajay Nagesh, Toshiaki Nakazawa, Matteo Negri, Santanu Pal, Allahsera Auguste Tapo, Marco Turchi, Valentin Vydin, and Marcos Zampieri. 2021. [Findings of the 2021 conference on machine translation \(WMT21\)](#). In *Proceedings of the Sixth Conference on Machine Translation*, pages 1–88, Online. Association for Computational Linguistics.
- B. Bangalore, G. Bordel, and G. Riccardi. 2001. [Computing consensus translation from multiple machine translation systems](#). In *IEEE Workshop on Automatic Speech Recognition and Understanding, 2001. ASRU '01.*, pages 351–354.
- Marta Bañón, Pinzhen Chen, Barry Haddow, Kenneth Heafield, Hieu Hoang, Miquel Esplà-Gomis, Mikel L. Forcada, Amir Kamran, Faheem Kirefu, Philipp Koehn, Sergio Ortiz Rojas, Leopoldo Pla Sempere, Gema Ramírez-Sánchez, Elsa Sarrías, Marek Strelec, Brian Thompson, William Waites, Dion Wiggins, and Jaume Zaragoza. 2020. [ParaCrawl: Web-scale acquisition of parallel corpora](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4555–4567, Online. Association for Computational Linguistics.
- Ondřej Bojar, Christian Federmann, Mark Fishel, Yvette Graham, Barry Haddow, Matthias Huck, Philipp Koehn, and Christof Monz. 2018. [Findings of the 2018 conference on machine translation \(WMT18\)](#). In *Proceedings of the Third Conference on Machine Translation: Shared Task Papers*, pages 272–303, Belgium, Brussels. Association for Computational Linguistics.
- Eleftheria Briakou, Colin Cherry, and George Foster. 2023. [Searching for needles in a haystack: On the role of incidental bilingualism in PaLM’s translation capability](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 9432–9452, Toronto, Canada. Association for Computational Linguistics.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.
- Cristian Buciluă, Rich Caruana, and Alexandru Niculescu-Mizil. 2006. [Model compression](#). In *Proceedings of the 12th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '06*, page 535–541, New York, NY, USA. Association for Computing Machinery.
- Yijie Chen, Yijin Liu, Fandong Meng, Yufeng Chen, Jinan Xu, and Jie Zhou. 2023. [Improving translation faithfulness of large language models via augmenting instructions](#).
- Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, Parker Schuh, Kensen Shi, Sasha Tsvyashchenko, Joshua Maynez, Abhishek Rao, Parker Barnes, Yi Tay, Noam Shazeer, Vinodkumar Prabhakaran, Emily Reif, Nan Du, Ben Hutchinson, Reiner Pope, James Bradbury, Jacob Austin, Michael Isard, Guy Gur-Ari, Pengcheng Yin, Toju Duke, Anselm Levskaya, Sanjay Ghemawat, Sunipa Dev, Henryk Michalewski, Xavier Garcia, Vedant Misra, Kevin Robinson, Liam Fedus, Denny Zhou, Daphne Ippolito, David Luan, Hyeontaek Lim, Barret Zoph, Alexander Spiridonov, Ryan Sepassi, David Dohan, Shivani Agrawal, Mark Omernick, Andrew M. Dai, Thanumalayan Sankaranarayanan Pilla, Marie Pellat, Aitor Lewkowycz, Erica Moreira, Rewon Child, Oleksandr Polozov, Katherine Lee, Zongwei Zhou, Xuezhi Wang, Brennan Saeta, Mark Diaz, Orhan Firat, Michele Catasta, Jason Wei, Kathy Meier-Hellstern, Douglas Eck, Jeff Dean, Slav Petrov, and Noah Fiedel. 2022. [Palm: Scaling language modeling with pathways](#).

- Junyoung Chung, Kyunghyun Cho, and Yoshua Bengio. 2016. [A character-level decoder without explicit segmentation for neural machine translation](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1693–1703, Berlin, Germany. Association for Computational Linguistics.
- Praveen Dakwale and Christof Monz. 2017. [Fine-tuning for neural machine translation with limited degradation across in- and out-of-domain data](#). In *Proceedings of Machine Translation Summit XVI: Research Track*, pages 156–169, Nagoya Japan.
- Jacob Devlin, Rabih Zbib, Zhongqiang Huang, Thomas Lamar, Richard Schwartz, and John Makhoul. 2014. [Fast and robust neural network joint models for statistical machine translation](#). In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1370–1380, Baltimore, Maryland. Association for Computational Linguistics.
- Thomas G. Dietterich. 2000. Ensemble methods in machine learning. In *Multiple Classifier Systems*, pages 1–15, Berlin, Heidelberg. Springer Berlin Heidelberg.
- Shuoyang Ding, Huda Khayrallah, Philipp Koehn, Matt Post, Gaurav Kumar, and Kevin Duh. 2017. [The JHU machine translation systems for WMT 2017](#). In *Proceedings of the Second Conference on Machine Translation*, pages 276–282, Copenhagen, Denmark. Association for Computational Linguistics.
- Markus Freitag, Matthias Huck, and Hermann Ney. 2014. [Jane: Open source machine translation system combination](#). In *Proceedings of the Demonstrations at the 14th Conference of the European Chapter of the Association for Computational Linguistics*, pages 29–32, Gothenburg, Sweden. Association for Computational Linguistics.
- Caglar Gulcehre, Orhan Firat, Kelvin Xu, Kyunghyun Cho, Loic Barrault, Huei-Chi Lin, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2015. [On using monolingual corpora in neural machine translation](#).
- Caglar Gulcehre, Orhan Firat, Kelvin Xu, Kyunghyun Cho, and Yoshua Bengio. 2017. [On integrating a language model into neural machine translation](#). *Computer Speech & Language*, 45:137–148.
- L.K. Hansen and P. Salamon. 1990. [Neural network ensembles](#). *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(10):993–1001.
- Kenneth Heafield and Alon Lavie. 2010. [Combining machine translation output with open source: The Carnegie Mellon multi-engine machine translation scheme](#). *The Prague Bulletin of Mathematical Linguistics*, 93:27–36.
- Amr Hendy, Mohamed Abdelrehim, Amr Sharaf, Vikas Raunak, Mohamed Gabr, Hitokazu Matsushita, Young Jin Kim, Mohamed Afify, and Hany Hassan Awadalla. 2023. [How good are gpt models at machine translation? a comprehensive evaluation](#).
- Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. 2015. [Distilling the knowledge in a neural network](#).
- Dongfu Jiang, Xiang Ren, and Bill Yuchen Lin. 2023. [LLM-blender: Ensembling large language models with pairwise ranking and generative fusion](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 14165–14178, Toronto, Canada. Association for Computational Linguistics.
- Wenxiang Jiao, Wenxuan Wang, Jen tse Huang, Xing Wang, Shuming Shi, and Zhaopeng Tu. 2023. [Is chatgpt a good translator? yes with gpt-4 as the engine](#).
- Marcin Junczys-Dowmunt, Tomasz Dwojak, and Rico Sennrich. 2016. [The AMU-UEDIN submission to the WMT16 news translation task: Attention-based NMT models as feature functions in phrase-based SMT](#). In *Proceedings of the First Conference on Machine Translation: Volume 2, Shared Task Papers*, pages 319–325, Berlin, Germany. 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.
- Huda Khayrallah and Philipp Koehn. 2018. [On the impact of various types of noise on neural machine translation](#). In *Proceedings of the 2nd Workshop on Neural Machine Translation and Generation*, pages 74–83, Melbourne, Australia. Association for Computational Linguistics.
- Huda Khayrallah, Gaurav Kumar, Kevin Duh, Matt Post, and Philipp Koehn. 2017. [Neural lattice search for domain adaptation in machine translation](#). In *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 20–25, Taipei, Taiwan. Asian Federation of Natural Language Processing.
- Huda Khayrallah, Brian Thompson, Kevin Duh, and Philipp Koehn. 2018. [Regularized training objective for continued training for domain adaptation in neural machine translation](#). In *Proceedings of the 2nd Workshop on Neural Machine Translation and Generation*, pages 36–44, Melbourne, Australia. Association for Computational Linguistics.
- Huda Khayrallah, Brian Thompson, Matt Post, and Philipp Koehn. 2020. [Simulated multiple reference training improves low-resource machine translation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 82–89, Online. 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.
- Philipp Koehn, Vishrav Chaudhary, Ahmed El-Kishky, Naman Goyal, Peng-Jen Chen, and Francisco Guzmán. 2020. [Findings of the WMT 2020 shared task on parallel corpus filtering and alignment](#). In *Proceedings of the Fifth Conference on Machine Translation*, pages 726–742, Online. Association for Computational Linguistics.
- Philipp Koehn, Francisco Guzmán, Vishrav Chaudhary, and Juan Pino. 2019. [Findings of the WMT 2019 shared task on parallel corpus filtering for low-resource conditions](#). In *Proceedings of the Fourth Conference on Machine Translation (Volume 3: Shared Task Papers, Day 2)*, pages 54–72, Florence, Italy. Association for Computational Linguistics.
- Philipp Koehn, Huda Khayrallah, Kenneth Heafield, and Mikel L. Forcada. 2018. [Findings of the WMT 2018 shared task on parallel corpus filtering](#). In *Proceedings of the Third Conference on Machine Translation: Shared Task Papers*, pages 726–739, Belgium, Brussels. Association for Computational Linguistics.
- Taku Kudo and John Richardson. 2018. [Sentencepiece: A simple and language independent subword tokenizer and detokenizer for neural text processing](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, EMNLP 2018: System Demonstrations, Brussels, Belgium, October 31 - November 4, 2018*, pages 66–71. Association for Computational Linguistics.
- Samuel Läubli, Rico Sennrich, and Martin Volk. 2018. [Has machine translation achieved human parity? a case for document-level evaluation](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4791–4796, Brussels, Belgium. Association for Computational Linguistics.
- Jiahuan Li, Hao Zhou, Shujian Huang, Shanbo Cheng, and Jiajun Chen. 2023. [Eliciting the translation ability of large language models via multilingual finetuning with translation instructions](#).
- Haitao Mi, Zhiguo Wang, and Abe Ittycheriah. 2016. [Vocabulary manipulation for neural machine translation](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 124–129, Berlin, Germany. Association for Computational Linguistics.
- Yasmin Moslem, Rejwanul Haque, John D. Kelleher, and Andy Way. 2023. [Adaptive machine translation with large language models](#). In *Proceedings of the 24th Annual Conference of the European Association for Machine Translation*, pages 227–237, Tampere, Finland. European Association for Machine Translation.
- Aitor Ormazabal, Mikel Artetxe, and Eneko Agirre. 2023. [Comblm: Adapting black-box language models through small fine-tuned models](#).
- Frithjof Petrick, Christian Herold, Pavel Petrushkov, Shahram Khadivi, and Hermann Ney. 2023. [Document-level language models for machine translation](#).
- Ofir Press and Lior Wolf. 2017. [Using the output embedding to improve language models](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 157–163, Valencia, Spain. Association for Computational Linguistics.
- Ricardo Rei, José G. C. de Souza, Duarte Alves, Chrysoula Zerva, Ana C Farinha, Taisiya Glushkova, Alon Lavie, Luisa Coheur, and André F. T. Martins. 2022. [COMET-22: Unbabel-IST 2022 submission for the metrics shared task](#). In *Proceedings of the Seventh Conference on Machine Translation (WMT)*, pages 578–585, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- Nathaniel R. Robinson, Perez Ogayo, David R. Mortensen, and Graham Neubig. 2023. [Chatgpt mt: Competitive for high- \(but not low-\) resource languages](#).
- Elizabeth Salesky, Matthew Wiesner, Jacob Bremerman, Roldano Cattoni, Matteo Negri, Marco Turchi, Douglas W. Oard, and Matt Post. 2021. [Multilingual tedx corpus for speech recognition and translation](#). In *Proceedings of Interspeech*.
- 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.
- Steve Sloto, Brian Thompson, Huda Khayrallah, Tobias Domhan, Thamme Gowda, and Philipp Koehn. 2023. [Findings of the WMT 2023 shared task on parallel data curation](#). In *Proceedings of the Eighth Conference on Machine Translation (WMT)*, Singapore, Singapore. Association for Computational Linguistics.
- Shaden Smith, Mostofa Patwary, Brandon Norick, Patrick LeGresley, Samyam Rajbhandari, Jared Casper, Zhun Liu, Shrimai Prabhunoye, George Zerveas, Vijay Korthikanti, Elton Zhang, Rewon Child, Reza Yazdani Aminabadi, Julie Bernauer, Xia Song, Mohammad Shoeybi, Yuxiong He, Michael



- Houston, Saurabh Tiwary, and Bryan Catanzaro. 2022. [Using deepspeed and megatron to train megatron-turing nlG 530b, a large-scale generative language model.](#)
- Felipe Soares, Mark Stevenson, Diego Bartolome, and Anna Zaretskaya. 2020. [ParaPat: The multi-million sentences parallel corpus of patents abstracts.](#) In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 3769–3774, Marseille, France. European Language Resources Association.
- Felix Stahlberg, James Cross, and Veselin Stoyanov. 2018. [Simple fusion: Return of the language model.](#) In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 204–211, Brussels, Belgium. Association for Computational Linguistics.
- Felix Stahlberg, Adrià de Gispert, Eva Hasler, and Bill Byrne. 2017. [Neural machine translation by minimising the Bayes-risk with respect to syntactic translation lattices.](#) In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 362–368, Valencia, Spain. Association for Computational Linguistics.
- Jörg Tiedemann. 2012. Parallel data, tools and interfaces in opus. In *Proceedings of the Eight International Conference on Language Resources and Evaluation (LREC'12)*, Istanbul, Turkey. European Language Resources Association (ELRA).
- Antonio Toral, Sheila Castilho, Ke Hu, and Andy Way. 2018. [Attaining the unattainable? reassessing claims of human parity in neural machine translation.](#) In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 113–123, Brussels, Belgium. 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. 2023a. [Llama: Open and efficient foundation language models.](#)
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023b. [Llama 2: Open foundation and fine-tuned chat models.](#)
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need.](#) In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Giorgos Vernikos, Brian Thompson, Prashant Mathur, and Marcello Federico. 2022. [Embarrassingly easy document-level MT metrics: How to convert any pretrained metric into a document-level metric.](#) In *Proceedings of the Seventh Conference on Machine Translation (WMT)*, pages 118–128, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- David Vilar, Markus Freitag, Colin Cherry, Jiaming Luo, Viresh Ratnakar, and George Foster. 2023. [Prompting PaLM for translation: Assessing strategies and performance.](#) In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15406–15427, Toronto, Canada. Association for Computational Linguistics.
- Xing Wang, Zhengdong Lu, Zhaopeng Tu, Hang Li, Deyi Xiong, and Min Zhang. 2017. [Neural machine translation advised by statistical machine translation.](#) In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, AAAI'17*, page 3330–3336. AAAI Press.
- Rachel Wicks and Matt Post. 2023. [Identifying context-dependent translations for evaluation set production.](#) In *Proceedings of the Eighth Conference on Machine Translation*, pages 452–467, Singapore. Association for Computational Linguistics.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. [Transformers: State-of-the-art natural language processing.](#) In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Haoran Xu, Young Jin Kim, Amr Sharaf, and Hany Hassan Awadalla. 2023. [A paradigm shift in machine translation: Boosting translation performance of large language models.](#)
- Wen Yang, Chong Li, Jiajun Zhang, and Chengqing Zong. 2023. [Bigtranslate: Augmenting large language models with multilingual translation capability over 100 languages.](#)

- Kyra Yee, Yann Dauphin, and Michael Auli. 2019. [Simple and effective noisy channel modeling for neural machine translation](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 5696–5701, Hong Kong, China. Association for Computational Linguistics.
- Jiali Zeng, Fandong Meng, Yongjing Yin, and Jie Zhou. 2023. [Tim: Teaching large language models to translate with comparison](#).
- Biao Zhang, Barry Haddow, and Alexandra Birch. 2023a. [Prompting large language model for machine translation: A case study](#). In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 41092–41110. PMLR.
- Jiacheng Zhang, Huanbo Luan, Maosong Sun, Feifei Zhai, Jingfang Xu, and Yang Liu. 2021. [Neural machine translation with explicit phrase alignment](#). *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 29:1001–1010.
- Tianyi Zhang, Faisal Ladhak, Esin Durmus, Percy Liang, Kathleen McKeown, and Tatsunori B. Hashimoto. 2023b. [Benchmarking large language models for news summarization](#).
- Wenxuan Zhang, Yue Deng, Bing Liu, Sinno Jialin Pan, and Lidong Bing. 2023c. [Sentiment analysis in the era of large language models: A reality check](#). *arXiv preprint arXiv:2305.15005*.
- Vilém Zouhar, Shuoyang Ding, Anna Currey, Tatyana Badeka, Jenyuan Wang, and Brian Thompson. 2024. [Fine-tuned machine translation metrics struggle in unseen domains](#). *arXiv preprint arXiv:2306.07899*.

## Appendix

### A Experimental Details

#### A.1 Hyperparameters

For German, Russian and Turkish, Transformer ‘big’ models were trained (6 layer encoder-decoder, 1024 embedding dimensions, 4096 feed-forward dimensions, 16 heads) (Vaswani et al., 2017). The base Transformer architecture was used for Hausa (6 layer encoder-decoder, 512 embedding dimensions, 2048 feed-forward dimensions, 8 heads). We use weight tying (Press and Wolf, 2017). We train models using Marian NMT (Junczys-Dowmunt et al., 2018).<sup>9</sup> We convert MT models from Marian to Hugging Face format, to allow for inference with Llama2 (Touvron et al., 2023b) in the Hugging Face library (Wolf et al., 2020).

#### A.2 Monolingual Data

#### A.3 Prompting

Figure 3, Figure 4, Figure 5, and Figure 6 describe the various prompts we use.

#### A.4 $\lambda$

Figure 7 shows translation quality as we vary the mixing ratio,  $\lambda$ . Note that  $p_{\text{ensemble}}$  reduces to the LLM when  $\lambda = 0$  and to the MT model when  $\lambda = 1$ .

For our results in the main section, we selected  $\lambda$  on validation set translation quality. Here we see that in cases where both models are reasonably strong (de-en, ru-en, and en-ru) the ensembling provides a quality boost.

#### A.5 COMET-22 CTXPro

Figure 8 shows the COMET-22 scores corresponding to the document translation accuracy show in Table 5. The ensemble is always best on this data, then the MT, and then the LLM.

---

<sup>9</sup><https://marian-nmt.github.io/>

Translate the following sentence from {src-language} to {tgt-language}:  
 {src-language}: {src}  
 {tgt-language}:

Figure 3: Baseline translation prompt.

Translate the following sentence from {src-language} to {tgt-language} in a {style} style:  
 {src-language}: {src}  
 {tgt-language}:

Figure 4: Translation prompt with domain.

Translate the following sentence from {src-language} to {tgt-language}:  
 {src-language}: {src-1}  
 {tgt-language}: {tgt-1}  
 ...  
 {src-language}: {src-n}  
 {tgt-language}: {tgt-n}  
 {src-language}: {src}  
 {tgt-language}:

Figure 5: n-shot translation prompt.

Translate the following sentence from {src-language} to {tgt-language}:  
 {src-language}: {previous-src-n}  
 {tgt-language}: {previous-translation-n}  
 ...  
 {src-language}: {previous-src}  
 {tgt-language}: {previous-translation}  
 {src-language}: {src}  
 {tgt-language}:

Figure 6: Context-aware translation prompt.

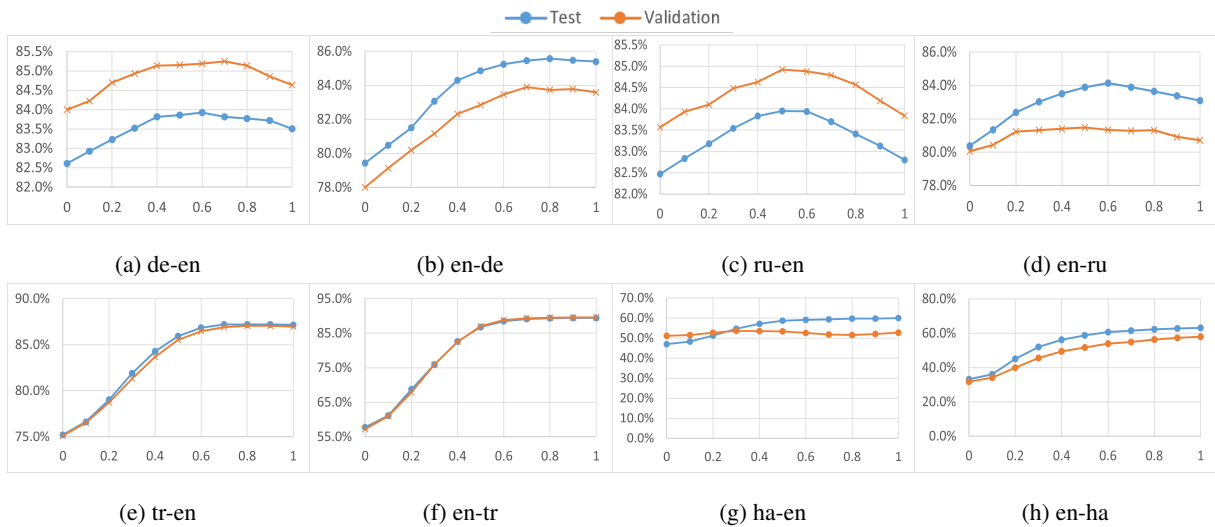


Figure 7: Ensembling MT model with 7B parameter LLM. Graphs shows COMET-22 vs mixing ratio.

	German		Russian		Turkish		Hausa	
	en	de	en	ru	en	tr	en	ha
news-commentary-v18	0.9m	0.5m	0.9m	0.5m				
europarl-v10	2.3m	2.1m						
news (all)	257.2m	468.9m	257.2m	142.7m				
news.2016					18.2m	1.7m		
news.2017					26.8m	3.0m		
news.2018							18.1m	
news.2019							33.6m	
news.2020							41.4m	
CommonCrawl						511.2m		8.5m

Table 6: Monolingual Datasets.

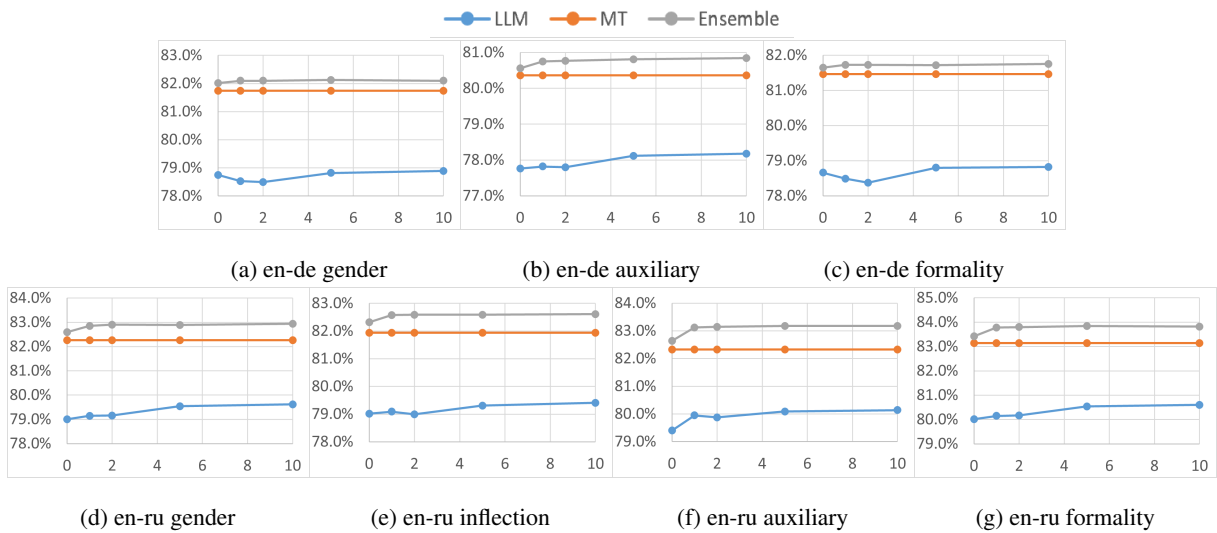


Figure 8: COMET-22 on the data in CTXPro.