

In-Domain African Languages Translation Using LLMs and Multi-armed Bandits

Pratik Rakesh Singh, Kritarth Prasad, Mohammadi Zaki and Pankaj Wasnik

Media Analysis Group, Sony Research India

{pratik.singh, kritarth.prasad, mohammadi.zaki, pankaj.wasnik}@sony.com

Abstract

Neural Machine Translation (NMT) systems face significant challenges when working with low-resource languages, particularly in domain adaptation tasks. These difficulties arise due to limited training data and suboptimal model generalization. As a result, selecting an optimal model for translation is crucial for achieving strong performance on in-domain data, particularly in scenarios where fine-tuning is not feasible or practical. In this paper, we investigate strategies for selecting the most suitable NMT model for a given domain using bandit-based algorithms, including Upper Confidence Bound, Linear UCB, Neural Linear Bandit, and Thompson Sampling. Our method effectively addresses the resource constraints by facilitating optimal model selection with high confidence. We evaluate the approach across three African languages and domains, demonstrating its robustness and effectiveness in both scenarios where target data is available and where it is absent.

1 Introduction

Advancements in multilingual machine translation models have significantly expanded language coverage, enabling translations even for low-resource languages. These models have also demonstrated strong performance in general domains, such as News, Movies, and more (Barrault et al., 2020) (Saunders, 2022a). Additionally, with the rise of large language models, methods like few-shot learning and in-context learning have shown notable improvements in domain adaptation tasks (Garcia et al., 2023) (Aycock and Bawden, 2024). Despite these advancements, the performance of these models remains highly dependent on the quality and scope of pre-training data as well as the model size, particularly for low-resource languages. It is common for a model to perform well in one domain but struggle in another, which presents a

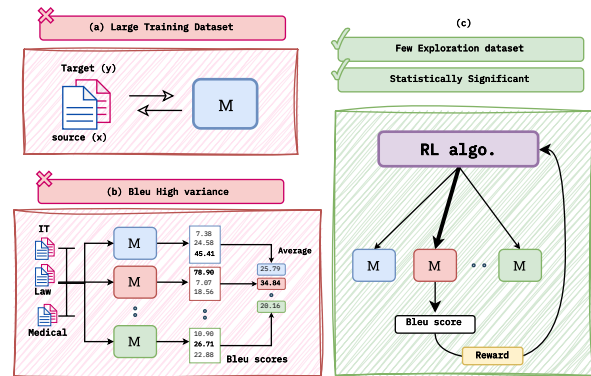


Figure 1: Motivation for Reinforcement Learning for model selection in machine translation: (a) Using a large dataset for training may be inefficient or impractical for low-resource settings, (b) BLEU scores vary significantly across domains, making model selection unreliable, (c) Reinforcement learning enables efficient model selection with fewer data and statistical significance.

significant challenge in the selection of the most suitable NMT system for a given task.

A commonly used approach for domain adaptation in Neural Machine Translation (NMT) tasks is fine-tuning NMT models on in-domain data using various strategies (Chu and Wang, 2018). As shown in Figure 1 (a), this approach faces significant challenges, particularly in low-resource settings where in-domain data is scarce. Moreover, fine-tuning often leads to a degradation in performance on general-domain data due to the issue of catastrophic forgetting (Thompson et al., 2019), further complicating the task of maintaining robust model performance across different domains.

Selection-based approaches have gained significant attention in recent Neural Machine Translation (NMT) systems, where the task is to identify the best possible model from a given set. A widely adopted method for this is the use of a Selection Block (SB) (Salazar et al., 2020) (Liu and Liu, 2021), which reranks models based on the spe-

cific task at hand. In recent years, reinforcement learning (RL)-based approaches have emerged as promising techniques for optimizing the selection of these models (Prasad et al., 2025). However, a limitation of many of these approaches is their reliance on large amounts of data to demonstrate performance gains over individual NMT systems.

One potential approach is to select the best model, which is trained on a general-domain dataset, and assume it will perform well on the in-domain dataset without the need for fine-tuning. Typically, one would evaluate models on the test dataset using common machine translation metrics, such as BLEU, to determine the best-performing model. However, in resource-constrained settings, these metrics can exhibit high variance, and there is often limited control over the statistical significance of the observed differences. This issue is illustrated in Figure 1 (b), where the mean of these metrics might provide a misleading impression of the best NMT system. In such cases, a few outlier examples could skew the BLEU score, leading to the wrong selection of the model. Moreover, a single evaluation does not capture the full variability in system performance, particularly when working with a small validation set. This underscores the need for model selection methods that not only choose the best NMT systems but also provide a statistical basis for the selection process, thereby mitigating the risks of misleading conclusions based on limited data.

To address the above challenges of model selection for Domain adaptation in resource-constrained settings, one possible approach could be to estimate the most optimal NMT system using fewer data samples, thereby reducing the reliance on large datasets. This can be achieved through the use of bandit-based algorithms (Zhou, 2016) (Bouneffouf et al., 2020), which allow for efficient exploration and exploitation of model performance, facilitating the identification of the best-performing system for the given domain with minimal data. As shown in Figure 1 (c), by leveraging these techniques, it is possible to make more informed decisions about model selection, even when In-domain data availability is limited, ensuring effective performance in low-resource scenarios. Our key contributions are summarized as follows:

- We propose a bandit-based approach to estimate optimal systems for a domain in a resource-constraint setting.

- We evaluate our approach on English to multiple African languages in multiple domains and report the performance of the popular bandit algorithms when applied to domain-specific model selection task.

2 Related works and Motivation

Domain Adaptation in Neural Machine Translation (NMT) refers to methods aimed at adjusting translation models trained on general-domain data to perform effectively in specific target domains with distinctive characteristics (Saunders, 2022b). Effective domain adaptation typically addresses data scarcity and domain mismatch (Pang et al., 2024) issues through data-centric and model-centric approaches. Data-centric strategies include back-translation using monolingual target data (Poncelas et al., 2019; Jin et al., 2020), forward-translation and self-learning (Chinea-Ríos et al., 2017), and synthetic data generation via noise introduction or lexicon-based methods (Vaibhav et al., 2019; Hu et al., 2019; Peng et al., 2020; Zhang et al., 2022). Model-centric approaches introduce domain-specific parameters or modules like domain tagging, embedding manipulation, adapter-based methods, and pointer-generators leveraging dictionaries (Kobus et al., 2017; Stergiadis et al., 2021; Pham et al., 2019; Bapna and Firat, 2019; Chen et al., 2021).

In low-resource scenarios, approaches such as data augmentation through bilingual lexicon-based replacements, transfer learning, and pretrained multilingual models have been employed (Nag et al., 2020; Liu et al., 2021). However, despite significant progress in both DA and low-resource NMT, domain adaptation techniques remain underexplored and challenging specifically for low-resource languages, where often the only available parallel data are very limited (Siddhant et al., 2022; Ranathunga et al., 2023).

Selection-Based Approach Recently, various selection methods have been introduced prior to the fusion step in multi-agent candidate selection. Significant research has focused on summarization tasks, including training reranking models based on evaluation metrics (Ravaut et al., 2023), employing contrastive learning for effective candidate ranking (Liu and Liu, 2021), and utilizing pairwise ranking methods to directly compare candidate summaries (Jiang et al., 2023). In the field of

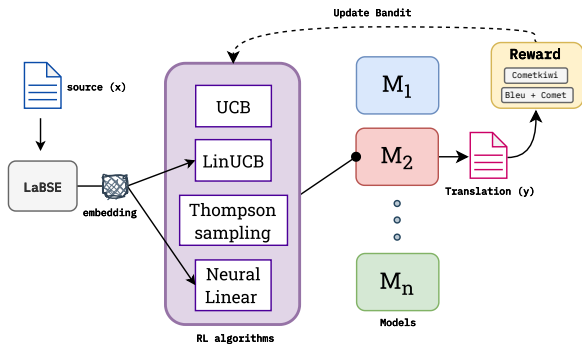


Figure 2: Block diagram of the proposed bandit-based model selection strategy.

neural machine translation (NMT), recent studies by (Prasad et al., 2025) have explored model selection strategies using a DQN-based approach. However, all these selection methods require substantial amounts of parallel data for effective training. Notably, limited research has investigated contextual bandit (Lu et al., 2010) approaches, which require significantly less data, to generalize agent selection based on provided context in low-resource machine translation scenarios.

Motivation for Bandit based approach A major challenge in Low-Resource Machine Translation (LRMT) is the scarcity of high-quality training datasets. This issue is further compounded in domain-specific translation, where the data becomes even more limited. While general-domain NMT systems exhibit reasonable performance across a broad range of tasks, their efficacy significantly fluctuates across different domains and languages. To mitigate this variability and identify the most effective NMT model for a given task with limited data, an optimal selection strategy is essential. This strategy must not only consider the available training data but also provide statistically-backed confidence in the model’s selection.

Bandit-based approaches have been widely explored in recommendation systems, where recommendations are generated based on past interactions with users (Silva et al., 2022). This methodology is well-suited for selecting optimal NMT systems in scenarios where only a small in-domain dataset is available, utilizing an appropriate reward function for NMT performance (Boursier and Perchet, 2024) (Nguyen et al., 2017). Furthermore, reference-less reward mechanisms offer a promising avenue for applying these bandit-based methods in target-free domain-specific machine transla-

tion tasks, as demonstrated by recent works (Obuchowski et al., 2024).

3 Methodology

As previously discussed, selecting the best model from a pool by evaluating a subset of data and then applying it to the entire test set is both computationally expensive and unreliable. Determining the necessary sample size to ensure the optimality of the chosen model becomes extremely important in such cases. Hence we take a more principled way to dynamically choose the machine translation model on-the-fly by treating the model selection process as a multi-armed bandit problem. We explore popular bandit algorithms designed for regret minimization, which, under mild theoretical assumptions, are proven to achieve (near-)optimal cumulative rewards over time. Below, we provide a brief overview of our methodology, as illustrated in Figure 2.

Each source sentence x is passed through a Language-agnostic BERT Sentence Encoder (LaBSE) to obtain a feature vector which we denote by overloading $x \in \mathbb{R}^d$. This vector x acts as the context vector in the contextual bandit algorithms considered in this work. The MT system pool act as the arms $\{M_1, M_2, \dots, M_n\}$ in our multi-armed bandit setup. Once the arm is chosen by the MAB algorithm, the corresponding MT system is chosen to translate the source sentence x to obtain y in the target language. Next, a reward is generated depending on x, y and whether a reference gold translation is available (see the next section for detail on ‘reward’) to obtain a scalar r .

Next, we provide a brief explanation of the arm selection strategy and update rules for each of the bandit algorithms we explore.

Upper Confidence Bound (UCB): UCB (Auer et al., 2002) relies on the principle of Optimism in the Face of Uncertainty (OFU). It selects the arm that maximizes an ‘upper confidence bound’ of its estimated reward.

Select arm:

$$a_t := \arg \max_{a \in \{M_1, M_2, \dots, M_n\}} \left(\hat{\mu}_a(t) + \alpha \sqrt{\frac{\log t}{N_a(t)}} \right)$$

Update empirical means of all arms.

where $\hat{\mu}_a(t)$ is the empirical reward obtained from pulling arm a till round t , $N_a(t)$ is the number of times arm a is pulled till round t , and α is a

confidence parameter.

Thompson Sampling (TS): Thompson Sampling (Thompson, 1933) is a Bayesian approach where we maintain a posterior distribution over each arm’s expected reward and sample from it. In particular, in our case we maintain a Beta distribution over each arm’s reward which has two parameters α, β which are initially set to 0. The arm selection and parameter update rules are as follows:

Select arm:

$$a_t := \arg \max_{a \in \{M_1, M_2, \dots, M_n\}} \theta_a \sim P(\cdot | \alpha_a, \beta_a)$$

Update:

$$\alpha_a \leftarrow \alpha_a + r, \beta_a \leftarrow \beta_a + 1 - r$$

Linear UCB (LinUCB): LinUCB (Li et al., 2010; Abbasi-yadkori et al., 2011) extends UCB to contextual bandits, assuming that rewards follow a linear function of the context/feature vector $x_t \in \mathbb{R}^d$ as explained before. In particular, we make the following assumption on the reward function that $\forall t \geq 1, r := x_t^T \theta_a + noise$ for all arms a .

Select arm:

$$a_t := \arg \max_{a \in \{M_1, M_2, \dots, M_n\}} \left(x_t^T \hat{\theta}_a + \alpha \sqrt{x_t^T A_a^{-1} x_t} \right)$$

where, $A_a := \sum_{s=1}^t \mathbb{1}\{a_s == a\} x_s x_s^T$, $b_a := \sum_{s=1}^t \mathbb{1}\{a_s == a\} r_s x_s$ and $\hat{\theta}_a := A_a^{-1} b_a$ is the Least Squares estimate of the true parameter θ_a of arm a .

Neural LinUCB (NL): Neural LinUCB (Xu et al., 2020) is a deep-learning extension of LinUCB, replacing the linear model with a neural network that maps features to a latent representation before applying LinUCB. In particular we replace the context vector x by a neural network $f(x; w)$ parameterized by w . The arm selection strategy and the update rule remain the same as in LinUCB with x replaced with $f(x; w)$.

Rewards: The rewards serve as the primary signal in bandit-based settings, guiding both the learning process and decision-making of the algorithms. The main objective in a Multi-Armed Bandit (MAB) problem is to maximize rewards by balancing exploration and exploitation.

In Neural Machine Translation (NMT), model performance is typically evaluated using standard

metrics such as BLEU and COMET. These metrics are particularly crucial for assessing how well a model translates within a specific domain. BLEU measures how accurately the model translates domain-specific vocabulary, while COMET evaluates the semantic similarity of the model’s output to the reference translation within the given domain. Using these metrics, we consider two types of reward signals as follows:

- **When parallel data is present:** When we have source along with the reference (gold) translation, we consider a combination of BLEU (Post, 2018) and a reference-based comet as shown below. Note that both the BLEU and comet scores have been normalized to lie between $[0,1]$.

$$\text{Reward} = \lambda \cdot \text{BLEU} + (1 - \lambda) \cdot \text{COMET}.$$

Here λ is a hyperparameter in $[0,1]$. In our experiments, we find that $\lambda = 0.4$ achieves the best results in our case.

- **Target-free scenario:** When only the source sentence is present, and the target is absent, which is typically the case in low-resource languages, especially in domain translation task, we use a reference-less MT metric like CometKiwi (Rei, 2022) (normalized between $[0,1]$), as the reward signal.

The λ is a controllability parameter that enables us to control the influence of metrics on reward, and reward_norm is the normalization value to normalize the metrics.

4 Experimental Setup

Datasets and Evaluation metrics: For our experiments, we utilize parallel datasets for English-to-African language translation, focusing on three African languages: English-to-Yoruba (en-yo), English-to-Swahili (en-sw) and English-to-Igbo (en-ig). We sample 1,000 parallel samples for validation (seed data for model convergence) and testing each. The datasets span three domains: News, Movies, and Religious texts.

- **News Domain:** We use the Lafand-MT dataset (Adelani et al., 2022), which contains parallel data for English-to-16 African languages, gathered from various news corpora.

Datasets	News Igbo	News Yoruba	News Swahili	Movies Igbo	Movies Yoruba	Movies Swahili	Religious Igbo	Religious Yoruba	Religious Swahili
Aya101	12.98	5.20	23.53	7.48	3.847	23.98	19.03	11.53	5.04
Gemma2 9B	8.17	3.55	24.24	5.21	1.97	25.07	9.48	2.12	3.9
Llama 3.1 8B	4.59	3.28	17.24	2.85	1.09	10.90	6.79	3.43	2.25
Madlad	6.91	1.11	8.92	7.087	1.19	24.47	3.19	1.36	32.36
NLLB	19.73	9.67	27.57	9.60	12.90	30.78	34.72	14.96	28.01
UCB	19.83	9.539	28.275	9.6	12.90	30.78	34.72	14.96	32.36
TS	19.48	9.74	26.95	9.36	12.88	27.29	34.34	14.34	32.54
LinUCB	19.73	9.54	27.80	9.80	12.90	29.9	34.8	13.7	32.45
NL	19.74	9.67	27.57	9.24	13.12	24.37	34.72	15.67	32.39

Table 1: Performance on BLEU metrics when Parallel data is present.

Datasets	News Igbo	News Yoruba	News Swahili	Movies Igbo	Movies Yoruba	Movies Swahili	Religious Igbo	Religious Yoruba	Religious Swahili
Aya101	12.98	5.20	23.53	7.48	3.847	23.98	19.03	11.53	5.04
Gemma2 9B	8.17	3.55	24.24	5.21	1.97	25.07	9.48	2.12	3.9
Llama 3.1 8B	4.59	3.28	17.24	2.85	1.09	10.90	6.79	3.43	2.25
Madlad	6.91	1.11	8.92	7.087	1.19	24.47	3.19	1.36	32.36
NLLB	19.73	9.67	27.57	9.60	12.90	30.78	34.72	14.96	28.01
UCB	19.83	9.53	28.27	7.48	12.90	30.78	34.72	14.96	28.01
TS	19.48	9.74	26.95	9.33	12.89	27.95	34.34	14.34	32.54
LinUCB	18.5	9.67	27.5	8.9	12.90	29.7	34.54	11.51	32.67
NL	19.73	9.67	27.57	7.48	12.90	23.98	34.72	11.53	32.36

Table 2: Performance on BLEU metrics in Target-Free scenario

- **Movies Domain:** We leverage the OpenSubtitles dataset (Lison and Tiedemann, 2016), which includes parallel translations of dialogues from various movies and TV shows. This domain is essential for capturing informal language usage and conversational nuances in translations.
- **Religious Texts Domain:** We compile a dataset from various sources, including CCAIghed (El-Kishky et al., 2020) and Tanzil (available at <https://tanzil.net/>), which contains Quran translations in multiple languages. This domain is particularly valuable for translating formal, religious content.

To assess the performance of the models, we rely on BLEU (Papineni et al., 2002), a widely accepted metric in machine translation. BLEU effectively measures the degree of overlap between the model-generated translations and reference translations, capturing the adequacy of domain-specific vocabulary translation, which is specifically effective for In-Domain Translations

Models: We test the effectiveness of our approach by using baselines, which also act as arms for bandits. The models used are a Mixture of LLMs and Foundational models like Aya101 (Üstün et al., 2024), NLLB200 3.3B (NLLB) (Team, 2022), Madlad400 10B (Madlad) (Kudugunta et al., 2023), Gemma2 9B (et al., 2024b) and Llama3.1 8B (et al., 2024a) all the model used for experiments are in its based or pre-trained state.

Choice of Hyper-parameters: For evaluation of our proposed bandit-based strategies, we use four

Algorithm	Parameter	Value
UCB	λ	[0.4-0.6]
Thompson Sampling	Prior Distribution	Beta(0,0)
LinUCB	α	1.5
	λ	0.4
Neural LinUCB	Arm model network	2 layer MLP with 50 neurons each

Table 3: Values of hyper-parameters used in our experiments.

popular arm-selection bandit algorithms that are UCB, LinUCB, Neural Bandit, Thompson Sampling as discussed in detail in Sec. 3. The values of the hyper-parameters specific to each bandit strategy are given in Table 3. The selection of Hyper-parameters was done based on hit and trail method, where initial few sentences of validations were used for convergence of algorithm and the algorithm was tested on rest of the remaining sentences.

5 Results

When parallel data is present: The hyper-parameters are tuned on the validation set and freed for all the algorithms. The Bandit-based models explored the reward that is a weighted summation of BLEU and reference-based Comet on the exploration set (Section 3 Rewards), the performance is evaluated using 1000 test samples per domain and language. As shown in Table 1, most of the Bandit-based approaches, when selecting the optimal arm, either perform on par with or surpass the best possible NMT model, demonstrating the effectiveness of our proposed method. On average, UCB (Upper Confidence

Bound) achieves superior performance across all languages and domains, outperforming the best model, NLLB, by an average improvement of 2.68% in BLEU score. Additionally, in specific cases—such as Neural Linear Bandit (NL) in the Movies domain for Yoruba, UCB in the News domain for Igbo and Swahili, and Thompson Sampling (TS) in the News domain for Yoruba, Religious, and Swahili the Bandit-based algorithms surpass the performance of the best possible NMT model. This suggests that these algorithms can occasionally select alternative NMT systems, resulting in slight but notable improvements in translation quality. In summary, the results indicate that Bandit-based approaches can effectively identify the best-performing NMT models, even with very small training sets, highlighting the robustness and utility of our proposed training strategy.

Results on Target-free Scenario: In this experiment, we explore the scenario where target translations for the in-domain dataset are unavailable, a common challenge in low-resource language settings. Such cases can be addressed using reference-less rewards (Section 3 Rewards), specifically leveraging CometKiwi-based metrics for NMT evaluation. The exploration of the Bandit-based models follows the same setup as discussed previously, with testing performed on 1000 samples. As shown in Table 2, the Bandit-based approaches successfully identify the best arms for translations even in the absence of target translations for reward generation. Among the various Bandit algorithms, UCB performs the best, followed by LinUCB and NL. Notably, in some instances, the Bandit-based selection slightly outperforms the individual best models, underscoring the flexibility of our approach. This demonstrates that our method can effectively be applied to model selection in target-free domain translation task where reference translations are not available.

6 Conclusion

In this paper, we presented a bandit-based approach for selecting the most suitable NMT model for domain adaptation, particularly in low-resource settings. Our method effectively balances exploration and exploitation by leveraging strategies such as Upper Confidence Bound, Linear UCB, Neural Linear Bandit, and Thompson Sampling, enabling op-

timal model selection with high confidence. Experimental results across multiple African languages and domains confirm the robustness of our approach, demonstrating its effectiveness both in the presence and absence of target domain data. Our findings highlight the potential of bandit-based methods to improve NMT performance in resource-constrained environments, paving the way for a more efficient and adaptive model selection process.

References

- Yasin Abbasi-yadkori, Dávid Pál, and Csaba Szepesvári. 2011. [Improved algorithms for linear stochastic bandits](#). In *Advances in Neural Information Processing Systems*, volume 24. Curran Associates, Inc.
- David Adelani *et al.* 2022. [A few thousand translations go a long way! leveraging pre-trained models for African news translation](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3053–3070, Seattle, United States. Association for Computational Linguistics.
- Peter Auer, Nicolò Cesa-Bianchi, and Paul Fischer. 2002. [Finite-time analysis of the multiarmed bandit problem](#). *Machine Learning*, 47(2):235–256.
- Seth Aycok and Rachel Bawden. 2024. [Topic-guided example selection for domain adaptation in LLM-based machine translation](#). In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics: Student Research Workshop*, pages 175–195, St. Julian’s, Malta. Association for Computational Linguistics.
- Ankur Bapna and Orhan Firat. 2019. [Simple, scalable adaptation 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 1538–1548, Hong Kong, China. Association for Computational Linguistics.
- Loïc Barrault, Magdalena Biesialska, Marta R. Costajussà, Fethi Bougares, and Olivier Galibert. 2020. [Findings of the first shared task on lifelong learning machine translation](#). In *Proceedings of the Fifth Conference on Machine Translation*, pages 56–64, Online. Association for Computational Linguistics.
- Djallel Bouneffouf, Irina Rish, and Charu Aggarwal. 2020. [Survey on applications of multi-armed and contextual bandits](#). In *2020 IEEE Congress on Evolutionary Computation (CEC)*, pages 1–8.
- Etienne Boursier and Vianney Perchet. 2024. [A survey on multi-player bandits](#). *Preprint*, arXiv:2211.16275.

- Guanhua Chen, Yun Chen, Yong Wang, and Victor O. K. Li. 2021. Lexical-constraint-aware neural machine translation via data augmentation. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI'20*.
- Mara Chinea-Ríos, Álvaro Peris, and Francisco Casacuberta. 2017. Adapting neural machine translation with parallel synthetic data. In *Proceedings of the Second Conference on Machine Translation*, pages 138–147, Copenhagen, Denmark. Association for Computational Linguistics.
- Chenhui Chu and Rui Wang. 2018. A survey of domain adaptation for neural machine translation. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 1304–1319, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Ahmed El-Kishky, Vishrav Chaudhary, Francisco Guzman, and Philipp Koehn. 2020. Ccaligned: A massive collection of cross-lingual web-document pairs. *Preprint*, arXiv:1911.06154.
- Xavier Garcia, Yamini Bansal, Colin Cherry, George Foster, Maxim Krikun, Fangxiaoyu Feng, Melvin Johnson, and Orhan Firat. 2023. The unreasonable effectiveness of few-shot learning for machine translation. *Preprint*, arXiv:2302.01398.
- Junjie Hu, Mengzhou Xia, Graham Neubig, and Jaime Carbonell. 2019. Domain adaptation of neural machine translation by lexicon induction. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2989–3001, Florence, Italy. Association for Computational Linguistics.
- 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.
- Di Jin, Zhijing Jin, Joey Tianyi Zhou, and Peter Szolovits. 2020. A simple baseline to semi-supervised domain adaptation for machine translation. *Preprint*, arXiv:2001.08140.
- Catherine Kobus, Josep Crego, and Jean Senellart. 2017. Domain control for neural machine translation. *Preprint*, arXiv:1612.06140.
- Sneha Kudugunta, Isaac Caswell, Biao Zhang, Xavier Garcia, Christopher A. Choquette-Choo, Katherine Lee, Derrick Xin, Aditya Kusupati, Romi Stella, Ankur Bapna, and Orhan Firat. 2023. Madlad-400: A multilingual and document-level large audited dataset. *Preprint*, arXiv:2309.04662.
- Lihong Li, Wei Chu, John Langford, and Robert E. Schapire. 2010. A contextual-bandit approach to personalized news article recommendation. In *Proceedings of the 19th International Conference on World Wide Web, WWW '10*, page 661–670, New York, NY, USA. Association for Computing Machinery.
- Pierre Lison and Jörg Tiedemann. 2016. OpenSubtitles2016: Extracting large parallel corpora from movie and TV subtitles. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC'16)*, pages 923–929, Portorož, Slovenia. European Language Resources Association (ELRA).
- Yixin Liu and Pengfei Liu. 2021. SimCLS: A simple framework for contrastive learning of abstractive summarization. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 1065–1072, Online. Association for Computational Linguistics.
- Zihan Liu, Genta Indra Winata, and Pascale Fung. 2021. Continual mixed-language pre-training for extremely low-resource neural machine translation. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 2706–2718, Online. Association for Computational Linguistics.
- Tyler Lu, David Pal, and Martin Pal. 2010. Contextual multi-armed bandits. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, volume 9 of *Proceedings of Machine Learning Research*, pages 485–492, Chia Laguna Resort, Sardinia, Italy. PMLR.
- Sreyashi Nag, Mihir Kale, Varun Lakshminarasimhan, and Swapnil Singhavi. 2020. Incorporating bilingual dictionaries for low resource semi-supervised neural machine translation. *Preprint*, arXiv:2004.02071.
- Khanh Nguyen, Hal Daumé III, and Jordan Boyd-Graber. 2017. Reinforcement learning for bandit neural machine translation with simulated human feedback. *Preprint*, arXiv:1707.07402.
- Aleksander Obuchowski, Barbara Kludel, Piotr Frąckowski, Sebastian Krajna, Wasyl Badyra, Michał Czubenko, and Zdzisław Kowalczyk. 2024. Target-free domain adaptation through cross-adaptation (student abstract). *Proceedings of the AAAI Conference on Artificial Intelligence*, 38(21):23602–23603.
- Jianhui Pang, Fanghua Ye, Longyue Wang, Dian Yu, Derek F. Wong, Shuming Shi, and Zhaopeng Tu. 2024. Salute the classic: Revisiting challenges of machine translation in the age of large language models. *Preprint*, arXiv:2401.08350.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th Annual Meeting on Association for Computational Linguistics, ACL '02*, page 311–318, USA. Association for Computational Linguistics.

- Wei Peng, Chongxuan Huang, Tianhao Li, Yun Chen, and Qun Liu. 2020. [Dictionary-based data augmentation for cross-domain neural machine translation](#). *Preprint*, arXiv:2004.02577.
- MinhQuang Pham, Josep Crego, François Yvon, and Jean Senellart. 2019. [Generic and specialized word embeddings for multi-domain machine translation](#). In *Proceedings of the 16th International Conference on Spoken Language Translation*, Hong Kong. Association for Computational Linguistics.
- Alberto Poncelas, Gideon Maillette de Buy Wenniger, and Andy Way. 2019. [Adaptation of machine translation models with back-translated data using transductive data selection methods](#). *Preprint*, arXiv:1906.07808.
- 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.
- Kritarth Prasad, Mohammadi Zaki, Pratik Singh, and Pankaj Wasnik. 2025. [Faster machine translation ensembling with reinforcement learning and competitive correction](#). *Preprint*, arXiv:2501.15219.
- Surangika Ranathunga, En-Shiun Annie Lee, Marjana Prifti Skenduli, Ravi Shekhar, Mehreen Alam, and Rishemjit Kaur. 2023. [Neural machine translation for low-resource languages: A survey](#). *ACM Comput. Surv.*, 55(11).
- Mathieu Ravaut, Shafiq Joty, and Nancy F. Chen. 2023. [Summareranker: A multi-task mixture-of-experts re-ranking framework for abstractive summarization](#). *Preprint*, arXiv:2203.06569.
- Ricardo *et al.* Rei. 2022. [CometKiwi: IST-unbabel 2022 submission for the quality estimation shared task](#). In *Proceedings of the Seventh Conference on Machine Translation (WMT)*, pages 634–645, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- Julian Salazar, Davis Liang, Toan Q. Nguyen, and Katrin Kirchhoff. 2020. [Masked language model scoring](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2699–2712, Online. Association for Computational Linguistics.
- Danielle Saunders. 2022a. [Domain adaptation and multi-domain adaptation for neural machine translation: A survey](#). *Preprint*, arXiv:2104.06951.
- Danielle Saunders. 2022b. [Domain adaptation and multi-domain adaptation for neural machine translation: A survey](#). *Preprint*, arXiv:2104.06951.
- Aditya Siddhant, Ankur Bapna, Orhan Firat, Yuan Cao, Mia Xu Chen, Isaac Caswell, and Xavier Garcia. 2022. [Towards the next 1000 languages in multilingual machine translation: Exploring the synergy between supervised and self-supervised learning](#). *Preprint*, arXiv:2201.03110.
- Nícollas Silva, Heitor Werneck, Thiago Silva, Adriano C.M. Pereira, and Leonardo Rocha. 2022. [Multi-armed bandits in recommendation systems: A survey of the state-of-the-art and future directions](#). *Expert Systems with Applications*, 197:116669.
- Emmanouil Stergiadis, Satendra Kumar, Fedor Kovalev, and Pavel Levin. 2021. [Multi-domain adaptation in neural machine translation through multidimensional tagging](#). In *Proceedings of Machine Translation Summit XVIII: Users and Providers Track*, pages 396–420, Virtual. Association for Machine Translation in the Americas.
- NLLB Team. 2022. [No language left behind: Scaling human-centered machine translation](#). *Preprint*, arXiv:2207.04672.
- Aaron Grattafiori *et al.* 2024a. [The llama 3 herd of models](#). *Preprint*, arXiv:2407.21783.
- Gemma Team *et al.* 2024b. [Gemma 2: Improving open language models at a practical size](#). *Preprint*, arXiv:2408.00118.
- Brian Thompson, Jeremy Gwinnup, Huda Khayrallah, Kevin Duh, and Philipp Koehn. 2019. [Overcoming catastrophic forgetting during domain adaptation of 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 2062–2068, Minneapolis, Minnesota. Association for Computational Linguistics.
- William R Thompson. 1933. [On the likelihood that one unknown probability exceeds another in view of the evidence of two samples](#). *Biometrika*, 25(3-4):285–294.
- Vaibhav Vaibhav, Sumeet Singh, Craig Stewart, and Graham Neubig. 2019. [Improving robustness of machine translation with synthetic noise](#). 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 1916–1920, Minneapolis, Minnesota. Association for Computational Linguistics.
- Pan Xu, Zheng Wen, Handong Zhao, and Quanquan Gu. 2020. [Neural contextual bandits with deep representation and shallow exploration](#). *CoRR*, abs/2012.01780.
- Hongxiao Zhang, Hui Huang, Jiale Gao, Yufeng Chen, Jinan Xu, and Jian Liu. 2022. [Iterative constrained back-translation for unsupervised domain adaptation of machine translation](#). In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 5054–5065, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.

Li Zhou. 2016. [A survey on contextual multi-armed bandits](#). *Preprint*, arXiv:1508.03326.

Ahmet Üstün, Viraat Aryabumi, Zheng-Xin Yong, Wei-Yin Ko, Daniel D'souza, Gbemileke Onilude, Neel Bhandari, Shivalika Singh, Hui-Lee Ooi, Amr Kayid, Freddie Vargus, Phil Blunsom, Shayne Longpre, Niklas Muennighoff, Marzieh Fadaee, Julia Kreutzer, and Sara Hooker. 2024. [Aya model: An instruction finetuned open-access multilingual language model](#). *Preprint*, arXiv:2402.07827.