

Design of an Open-Source Architecture for Neural Machine Translation

Séamus Lankford
ADAPT Centre,
MTU, Cork, Ireland
seamus.lankford@mtu.ie

Haithem Afli
ADAPT Centre,
MTU, Cork, Ireland
haithem.afli@mtu.ie

Andy Way
ADAPT Centre,
DCU, Dublin, Ireland
andy.way@dcu.ie

Abstract

adaptNMT is an open-source application that offers a streamlined approach to the development and deployment of Recurrent Neural Networks and Transformer models. This application is built upon the widely-adopted OpenNMT ecosystem, and is particularly useful for new entrants to the field, as it simplifies the setup of the development environment and creation of train, validation, and test splits. The application offers a graphing feature that illustrates the progress of model training, and employs SentencePiece for creating subword segmentation models. Furthermore, the application provides an intuitive user interface that facilitates hyperparameter customization. Notably, a single-click model development approach has been implemented, and models developed by adaptNMT can be evaluated using a range of metrics. To encourage eco-friendly research, adaptNMT incorporates a green report that flags the power consumption and kgCO₂ emissions generated during model development. The application is freely available.¹

1 Credits

This research is supported by Science Foundation Ireland through the ADAPT Centre (Grant 13/RC/2106) (www.adaptcentre.ie) at Dublin City University. This research was also funded by the Munster Technological University.

© 2023 The authors. This article is licensed under a Creative Commons 4.0 licence, no derivative works, attribution, CC-BY-ND.

¹<http://github.com/adaptNMT>

2 Introduction

Explainable Artificial Intelligence (XAI) (Arrieta et al., 2020) aims to ensure that the outcomes of AI solutions are easily comprehensible to humans. In light of this goal, adaptNMT has been developed to provide users with a form of Explainable Neural Machine Translation (XNMT). The typical NMT process comprises several independent stages, including setting up the environment, preparing the dataset, training subword models, parameterizing and training the main models, evaluating and deploying them. By adopting a modular approach, this framework has established an effective NMT model development process that caters to both technical and non-technical practitioners in the field. To address the environmental impact of building and running large AI models (Henderson et al., 2020; Jooste et al., 2022b), we have also produced a “green report” that calculates carbon emissions. While primarily intended as an information aid, this report will hopefully encourage the development of reusable and sustainable models.

This research endeavors to create models and applications that address the challenges of language technology, which will be particularly beneficial for those new to the field of Machine Translation (MT) and those seeking to learn more about NMT.

The application is built on OpenNMT² (Klein et al., 2017) and thus inherits all of its features. Unlike many NMT toolkits, a command line interface (CLI) is not used, and the interface is designed and fully implemented in Google Colab.³ For both educational and research purposes, a cloud-hosted solution like Colab is often more user-friendly.

²<https://opennmt.net>

³colab.research.google.com

Additionally, the training of models can be monitored and controlled via a Google Colab mobile app, which is useful for long-run builds. The adaptNMT framework also includes GUI controls that allow for the customization of all crucial parameters needed for NMT model training.

The application can be run in local mode to utilize existing infrastructure or hosted mode for rapid infrastructure scaling. A deploy function is also included to allow for the immediate deployment of trained models.

This paper begins by presenting background information on NMT and NMT tools in Section 3, followed by a detailed description of the adaptNMT architecture and its key features in Section 4. The system is discussed in Section 5 before concluding with a discussion of future work in Section 6. A more in-depth system description, coupled with an empirical evaluation of models developed using the application, is outlined in a separate paper (Lankford et al., 2023a).

3 Related Work

3.1 NMT

In addition to the ongoing research dedicated to developing state-of-the-art (SOTA) NMT models, comprehensive descriptions of this technology are readily available in the literature, making it accessible to individuals who are new to the field or have limited technical expertise (Way, 2019).

NMT has benefitted from the availability of large parallel corpora, leading to the development of high-performing MT models. The field of MT has experienced significant advancements through the application of NMT, particularly after the introduction of the Transformer (Vaswani et al., 2017) architecture, which has resulted in SOTA performance across multiple language pairs (Bojar et al., 2017; Bojar et al., 2018; Lankford et al., 2021a; Lankford et al., 2021b; Lankford et al., 2022a; Lankford et al., 2022b).

3.2 NMT Tools

In essence, adaptNMT is an IPython wrapper built on OpenNMT, enabling it to benefit from OpenNMT’s extensive feature set and continuous code maintenance. However, adaptNMT takes abstraction to a higher level than OpenNMT, with greater focus on usability, particularly for newcomers. As a result, adaptNMT facilitates easy and fast deployment, offering features such as more pre-

processing, as well as GUI control over model creation. Moreover, it incorporates green features in line with current research efforts towards smaller models with reduced carbon footprints, making it suitable for educational and research environments alike.

Other commonly used frameworks for developing NMT systems include FAIRSEQ⁴ (Ott et al., 2019), an open-source sequence modelling toolkit based on PyTorch that allows for training models for translation, summarization, and language modelling. Marian⁵ (Junczys-Dowmunt et al., 2018), on the other hand, is an NMT framework based on dynamic computation graphs and developed using C++. OpenNMT is an open-source NMT framework that has been widely adopted in the research community and covers the entire MT workflow from data preparation to live inference.

4 Architecture of adaptNMT

After providing a general overview of NMT and NMT development systems, we introduce the adaptNMT tool, which enables users to configure the components of the NMT development process. The platform’s system architecture is depicted in Figure 1. The tool is built as an IPython notebook and leverages the Pytorch implementation of OpenNMT for training models. Additionally, SentencePiece is used to train subword models. Using a Jupyter notebook facilitates sharing the application with other members of the MT community, and the application’s setup is simplified since all necessary packages are downloaded dynamically as the application runs.

The system has two deployment options: running it locally or as a Colab instance via Google Cloud. In order to build translation models, the system requires parallel text corpora for both the source and target languages. A Tensorboard visualization allows for real-time monitoring of the model training process. At runtime, users can select to use the system for either model building or translation services, or both. Additionally, as depicted in Figure 1, the system enables the generation of an ensemble output during translation. Finally, trained models can be easily deployed to a pre-configured location.

⁴<https://github.com/facebookresearch/fairseq>

⁵<https://marian-nmt.github.io>




Figure 1: Proposed architecture for adaptNMT: a language-agnostic NMT development environment. The system is designed to run either in the cloud or using local infrastructure. Models are trained using parallel corpora. Visualization and extensive logging enable real-time monitoring. Models are developed using vanilla RNN-based NMT, Transformer-based approaches or transfer learning using a fine-tuning approach. Translation and evaluation can be carried out using either single models or ensembles.

4.1 adaptNMT

The application may be run as an IPython Jupyter notebook or as a Google Colab application. Given the ease of integrating large Google drive storage into Colab, the application has been used exclusively as a Google Colab application for our own experiments.

4.1.1 Initialization and logging

Initialization enables connection to Google Drive to run experiments, automatic installation of Python, OpenNMT,⁶ SentencePiece,⁷ Pytorch and other applications. The visualization section enables real-time graphing of model development. All log files are stored and can be viewed to inspect training convergence, the model’s training and validation accuracy and changes in learning rates.

4.1.2 Modes of operation

There are two modes of operation: local or cloud. In local mode, the application is run so

⁶<https://opennmt.net>

⁷<https://github.com/google/sentencepiece>

that models are built using the user’s local GPU resources. The option to use cloud mode enables users to develop models using Google’s GPU clusters. For shorter training times, the unpaid Colab option is adequate. However, for a small monthly subscription, the Google Colab Pro option is worthwhile since users have access to improved GPU and compute resources. Furthermore, using Google Cloud may be considered as the “green option” since its platform uses 100% renewables (Lacoste et al., 2019).

4.1.3 Customization of models

The system has been developed to allow users to select variations to the underlying model architecture. A vanilla RNN or Transformer approach may be selected to develop the NMT model. The customization mode enables users to specify the exact parameters required for the chosen approach. One of the features, AutoBuild, enables a user to build an NMT model in three simple steps: (i) upload source and target files, (ii) select RNN or Transformer, and (iii) click AutoBuild.

4.1.4 Use of subword segmentation

In the NMT development process, users can specify the type of optimizer for learning and choose from different subword models. The subword model functionality allows for the selection of a subword model type and the choice of vocabulary size, currently offering either a SentencePiece unigram or a SentencePiece BPE model.

A user may upload a dataset which includes the train, validation and test splits for both source and target languages. In cases where a user has not already created the required splits for model training, single source and target files may be uploaded. Automated splitting of the uploaded dataset into train, validation, and test files is then performed based on the user’s chosen split ratio.

Given that building NMT models typically demands long training times, an automatic notification feature is incorporated that informs the user by email when model training has been completed.

4.1.5 Translation and evaluation

The application supports not only the training of models but also the translation and evaluation of model performance. For translation using pre-built models, users can specify the model name as a hyperparameter which is subsequently used to translate and evaluate the test files. The option for creating an ensemble output is also available, with users simply naming the models to be used in generating the ensemble output.

Once the system has been built, the user can select the model to be used for translating the test set. While human evaluation is often considered the most insightful approach for evaluating translation quality, it can be limited by factors such as availability, cost, and subjectivity. Thus, automatic evaluation metrics are frequently employed, particularly by developers monitoring incremental progress of systems. A further discussion on the advantages and disadvantages of human and automatic evaluation is available in the literature (Way, 2018).

Several automatic evaluation metrics provided by SacreBleu⁸ (Post, 2018) are used: BLEU (Papineni et al., 2002), TER (Snover et al., 2006) and ChrF (Popović, 2015). Translation quality can also be evaluated using Meteor (Denkowski and Lavie, 2014) and F1 score (Melamed et al., 2003). Note that BLEU, ChrF, Meteor and F1 are precision-

based metrics, so higher scores are better, whereas TER is an error-based metric and lower scores indicate better translation quality. Evaluation options available include standard (truecase) and lowercase BLEU scores, a sentence-level BLEU score option, ChrF1 and ChrF3.

There are three levels of logging: model development logs for graphing, training console output and experimental results. A references section outlines resources which are relevant to developing, using and understanding adaptNMT. Validation during training is currently conducted using model accuracy and perplexity (PPL).

4.2 Infrastructure

Rapid prototype development is possible through a Google Colab Pro subscription using NVIDIA Tesla P100 PCIe 16GB graphic cards and up to 27GB of memory when available.

5 Discussion

Numerous tools have been developed to assess the carbon footprint of NLP (Bannour et al., 2021). The notion of sustainable NLP has also gained momentum as an independent research track, with high-profile conferences such as the *EACL 2021 Green and Sustainable NLP* track dedicating resources to this area.⁹

Given these developments, we have incorporated a "green report" into adaptNMT that logs the kgCO₂ generated during model development. This aligns with the industry’s increasing focus on quantifying the environmental impact of NLP. In fact, it has been demonstrated that high-performing MT systems can be developed with much lower carbon footprints, leading to significant energy cost savings for a real translation company (Jooste et al., 2022a).

The risks associated with relying on Large Language Models (LLMs) have been well-documented in the literature. The discussion surrounding these models emphasizes not only their environmental impact but also the inherent biases and dangers they pose for low-resource languages (Bender et al., 2021). It is important to note that smaller, in-domain datasets can yield high-performing NMT models, and the adaptNMT framework makes this approach easily accessible and understandable.

⁸<https://github.com/mjpost/sacrebleu>

⁹<https://2021.eacl.org/news/green-and-sustainable-nlp>

6 Conclusion and Future Work

We have introduced adaptNMT, an application that manages the entire NMT model development, evaluation, and deployment workflow.

As for future work, our development efforts will be directed towards incorporating new transfer learning methods and improving our ability to track environmental costs. We will integrate modern zero-shot and few-shot approaches, as seen in the GPT3 (Brown et al., 2020) and Facebook LASER (Artetxe and Schwenk, 2019) frameworks. While the existing adaptNMT application is focused on customizing NMT models, we will also develop a separate application, adaptLLM (Lankford et al., 2023b; Lankford et al., 2023c), for fine-tuning LLMs. In particular, adaptLLM will cater for low-resource language pairs such as NLLB (Costa-jussà et al., 2022).

The green report integrated into the application represents our first implementation of a sustainable NLP feature within adaptNMT. We plan to enhance this feature by improving the user interface and providing recommendations on how to develop greener models. As an open-source project, we invite the community to contribute new ideas and improvements to the development of this feature.

References

- Arrieta, Alejandro Barredo, Natalia Díaz-Rodríguez, Javier Del Ser, Adrien Bannetot, Siham Tabik, Alberto Barbado, Salvador García, Sergio Gil-López, Daniel Molina, Richard Benjamins, et al. 2020. Explainable artificial intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible ai. *Information Fusion*, 58:82–115.
- Artetxe, Mikel and Holger Schwenk. 2019. Massively multilingual sentence embeddings for zero-shot cross-lingual transfer and beyond. *Transactions of the Association for Computational Linguistics*, 7:597–610.
- Bannour, Nesrine, Sahar Ghannay, Aurélie Névéal, and Anne-Laure Ligozat. 2021. Evaluating the carbon footprint of NLP methods: a survey and analysis of existing tools. In *Proceedings of the Second Workshop on Simple and Efficient Natural Language Processing*, pages 11–21, Virtual, November. Association for Computational Linguistics.
- Bender, Emily M., Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. On the dangers of stochastic parrots: Can language models be too big? In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '21, page 610–623, New York, NY, USA. Association for Computing Machinery.
- Bojar, Ondřej, Rajen Chatterjee, Christian Federmann, Yvette Graham, Barry Haddow, Shujian Huang, Matthias Huck, Philipp Koehn, Qun Liu, Varvara Logacheva, Christof Monz, Matteo Negri, Matt Post, Raphael Rubino, Lucia Specia, and Marco Turchi. 2017. Findings of the 2017 conference on machine translation (WMT17). In *Proceedings of the Second Conference on Machine Translation*, pages 169–214, Copenhagen, Denmark, September. Association for Computational Linguistics.
- Bojar, Ondřej, Christian Federmann, Mark Fishel, Yvette Graham, Barry Haddow, 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, October. Association for Computational Linguistics.
- Brown, Tom, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Costa-jussà, Marta R, James Cross, Onur Çelebi, Maha Elbayad, Kenneth Heafield, Kevin Heffernan, Elahe Kalbassi, Janice Lam, Daniel Licht, Jean Maillard, et al. 2022. No language left behind: Scaling human-centered machine translation. *arXiv preprint arXiv:2207.04672*.
- Denkowski, Michael and Alon Lavie. 2014. Meteor universal: Language specific translation evaluation for any target language. In *Proceedings of the Ninth Workshop on Statistical Machine Translation*, pages 376–380, Baltimore, Maryland, USA, June. Association for Computational Linguistics.
- Henderson, Peter, Jieru Hu, Joshua Romoff, Emma Brunskill, Dan Jurafsky, and Joelle Pineau. 2020. Towards the systematic reporting of the energy and carbon footprints of machine learning. *Journal of Machine Learning Research*, 21(248):1–43.
- Jooste, Wandri, Rejwanul Haque, and Andy Way. 2022a. Knowledge distillation: A method for making neural machine translation more efficient. *Information*, 13(2).
- Jooste, Wandri, Andy Way, Rejwanul Haque, and Riccardo Superbo. 2022b. Knowledge distillation for sustainable neural machine translation. In *Proceedings of the 15th Biennial Conference of the Association for Machine Translation in the Americas (Volume 2: Users and Providers Track and Government Track)*, pages 221–230, Orlando, USA, September. Association for Machine Translation in the Americas.

- Junczys-Dowmunt, Marcin, 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, July. Association for Computational Linguistics.
- Klein, Guillaume, Yoon Kim, Yuntian Deng, Jean Senellart, and Alexander Rush. 2017. OpenNMT: Open-source toolkit for neural machine translation. In *Proceedings of ACL 2017, System Demonstrations*, pages 67–72, Vancouver, Canada, July. Association for Computational Linguistics.
- Lacoste, Alexandre, Alexandra Luccioni, Victor Schmidt, and Thomas Dandres. 2019. Quantifying the carbon emissions of machine learning. *arXiv preprint arXiv:1910.09700*.
- Lankford, Séamus, Haithem Afi, and Andy Way. 2021a. Transformers for low-resource languages: Is féidir linn! In *Proceedings of Machine Translation Summit XVIII: Research Track*, pages 48–60.
- Lankford, Séamus, Haithem Afi, and Andy Way. 2021b. Machine translation in the covid domain: an English-Irish case study for LoResMT 2021. In *Proceedings of the 4th Workshop on Technologies for MT of Low Resource Languages (LoResMT2021)*, pages 144–150, Virtual, August. Association for Machine Translation in the Americas.
- Lankford, Séamus, Haithem Afi, Órla Ní Loinsigh, and Andy Way. 2022a. gaHealth: An English-Irish bilingual corpus of health data. In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 6753–6758, Marseille, France, June. European Language Resources Association.
- Lankford, Séamus, Haithem Afi, and Andy Way. 2022b. Human evaluation of English-Irish Transformer-Based NMT. *Information*, 13(7):309. 19pp.
- Lankford, Séamus, Haithem Afi, and Andy Way. 2023a. adaptNMT: an open-source, language-agnostic development environment for neural machine translation. *Language Resources and Evaluation*.
- Lankford, Séamus, Haithem Afi, and Andy Way. 2023b. adaptLLM: Fine-tuning large language models for improved machine translation. an open-source approach with NLLB and GPT-J integration. In *Journal of Artificial Intelligence Research [manuscript submitted]*, El Segundo, CA, United States. AI Access Foundation.
- Lankford, Séamus, Haithem Afi, and Andy Way. 2023c. Fine-tuning LLMs using low-resource language pairs: A comparative study with loresmt2021 shared task baselines. In *Proceedings of Machine Translation Summit XIX: Research Track [manuscript submitted]*, Macau SAR, China, September 4-8. Asia-Pacific Association for Machine Translation (AAMT).
- Melamed, I. Dan, Ryan Green, and Joseph P. Turian. 2003. Precision and recall of machine translation. In *Companion Volume of the Proceedings of HLT-NAACL 2003 - Short Papers*, pages 61–63, Edmonton, Canada.
- Ott, Myle, Sergey Edunov, Alexei Baevski, Angela Fan, Sam Gross, Nathan Ng, David Grangier, and Michael Auli. 2019. fairseq: A fast, extensible toolkit for sequence modeling. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (Demonstrations)*, pages 48–53, Minneapolis, Minnesota, June. Association for Computational Linguistics.
- Papineni, Kishore, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, PA.
- Popović, Maja. 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, September. Association for Computational Linguistics.
- Post, Matt. 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, October. Association for Computational Linguistics.
- Snover, Matthew, Bonnie Dorr, Rich Schwartz, Linnea Micciulla, and John Makhoul. 2006. A study of translation edit rate with targeted human annotation. In *Proceedings of the 7th Conference of the Association for Machine Translation in the Americas: Technical Papers*, pages 223–231, Cambridge, Massachusetts, USA, August 8-12. Association for Machine Translation in the Americas.
- Vaswani, Ashish, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS’17*, page 6000–6010, Red Hook, NY, USA. Curran Associates Inc.
- Way, Andy. 2018. Quality expectations of machine translation. In Moorkens, Joss, Sheila Castilho, Federico Gaspari, and Stephen Doherty, editors, *Translation Quality Assessment: From Principles to Practice*, pages 159–178. Springer, Cham, Switzerland.
- Way, Andy. 2019. Machine translation: Where are we at today? In Angelone, Eric, Gary Massey, and Maureen Ehrensberger-Dow, editors, *The Bloomsbury Companion to Language Industry Studies*, page 311–332. Bloomsbury, London.