

Domain Adaptation and Multi-Domain Adaptation for Neural Machine Translation: A Survey

Danielle Saunders

*Cambridge University Engineering Department
Cambridge, United Kingdom*

DS636@CANTAB.AC.UK

Abstract

The development of deep learning techniques has allowed Neural Machine Translation (NMT) models to become extremely powerful, given sufficient training data and training time. However, systems struggle when translating text from a new domain with a distinct style or vocabulary. Fine-tuning on in-domain data allows good domain adaptation, but requires sufficient relevant bilingual data. Even if this is available, simple fine-tuning can cause overfitting to new data and catastrophic forgetting of previously learned behaviour.

We survey approaches to domain adaptation for NMT, particularly where a system may need to translate across multiple domains. We divide techniques into those revolving around data selection or generation, model architecture, parameter adaptation procedure, and inference procedure. We finally highlight the benefits of domain adaptation and multi-domain adaptation techniques to other lines of NMT research.

1. Introduction

Neural Machine Translation (NMT) has seen impressive advances for some translation tasks in recent years. News and biomedical translation shared tasks from the Conference on Machine Translation (WMT) as early as 2019 identified several systems as performing on par with a human translator for some high-resource language pairs according to human judgements (Barrault et al., 2019; Bawden et al., 2019). Indeed, these tasks involve not only high-resource language pairs but also relatively high-resource domains, with millions of relevant sentence pairs available for training. However, NMT models perform less well on out-of-domain data. A model trained on exclusively news data is unlikely to achieve good performance on the biomedical domain, let alone human parity.

Models trained from scratch on a multi-domain dataset may reach good performance on domains that are sufficiently well-represented in that dataset. However, new domains of interest may arise. For example, generic news translation systems of 2019 do not perform optimally on news from 2020 reporting on the coronavirus pandemic (Anastasopoulos et al., 2020). While it is possible to train a new model from scratch for every new domain of interest, it is not generally practical: high-resource training sets contain millions of sentence pairs and require multiple days for training (Junczys-Dowmunt, 2019). General overviews of NMT (Stahlberg, 2020; Koehn, 2017; Neubig, 2017) tend to focus on the many design decisions involved in this from-scratch NMT system development. However, a ready-trained system can be repurposed to translate new types of language far faster than training a new model, and often performs better on the language of interest as a result (Luong et al., 2015; Freitag & Al-Onaizan, 2016). This repurposing is generally known as domain adaptation.

We interpret domain adaptation as any scheme intended to improve translations from an existing system for a certain topic or genre of language. In the definition of a domain we primarily follow Koehn and Knowles (2017), who state that a domain ‘may differ from other domains in topic, genre, style, level of formality, etc.’. As well as these more ambiguous qualities, they define a domain as ‘a corpus from a specific source’, a fixed, discrete criteria often termed domain provenance.

Existing work on domain-specific translation does frequently treat the domain of ‘test’ language as discrete, known, or completely distinct from other domains (Chu & Wang, 2018). Known-domain translation may be relevant in limited scenarios, such as WMT shared tasks where the topic and genre of text is pre-specified, or bespoke translation systems adapted to customer data. In this scenario, added domain-specific neural network layers (Bapna & Firat, 2019b) can be tuned briefly for good performance on a fixed task with relatively low cost. Such approaches have recently become popular throughout natural language processing (NLP) using large pre-trained language models: BERT (Devlin et al., 2019) and GPT-3 (Brown et al., 2020) among others.

However, even in known-domain settings, discrete domain labels or domain-specific parameters may not always be beneficial. A corpus domain label may not fit all sentences in that corpus, and sentences can belong to more than one domain. In a very general case, sentences supplied to a free online translation system could come from any source and contain features of any domain. This article does explore single-domain adaptation, but also the development of systems that can translate text from multiple domains: *multi-domain* translation. We will also emphasise techniques for adaptation that can incorporate the benefits of a new domain without succumbing to forgetting, brittleness, overfitting, or general failure to successfully translate anything other than a chosen set of adaptation sentences.

The survey begins in Section 2 by exploring what we mean by a domain or multi-domain adaptation problem. Section 3 gives a brief overview of NMT to provide context for later adaptation approaches. In Section 4 we summarize fine-tuning, often used as a default or baseline approach to adaptation, and highlight difficulties that motivate alternative approaches. We then review domain adaptation methods that intervene in various stages of NMT system design, training and use. These can be data-centric (Section 5), change the neural network architecture (Section 6) or adaptation procedure (Section 7), or take place during inference (Section 8). We explore the possibilities and pitfalls present at each step of system development. Section 9 presents case studies on three areas of NMT research that have benefited from framing as domain adaptation problems: low-resource language translation, demographic bias in translation, and document-level translation. Section 10 concludes with a view of some open challenges in domain adaptation for NMT.

The structure of this survey aims to acknowledge that researchers and practitioners often only have control over certain elements of an NMT system. A typical scenario is commercial NMT providers allowing customers access to a pre-trained model with a fixed inference procedure which the providers adapt to customer data (Savenkov, 2018). In this case, the end user only has control over the data used for adaptation. By contrast, a researcher participating in a shared task might have a fixed set of adaptation data to make best use of, and focus on adjusting the network architecture or adaptation procedure itself. Even a user with access only to a black-box commercial tool may be able to ‘adapt’ the output with pre- and post-processing terminology control.

Chu and Wang (2018) previously performed a brief survey on domain adaptation for NMT, later extended with more recent references (Chu & Wang, 2020). While their survey focuses on a handful of domain adaptation methods, primarily mixed fine-tuning, instance weighting and architectural domain control, this article offers a more substantive overview of adaptation techniques. In contrast to prior work we also emphasise the relative advantages and disadvantages of different methods in specific scenarios, describe multi-domain adaptation, and discuss how the described techniques can be applied to NMT challenges not typically thought of as domains.

Specific contributions of this article are as follows:

- We provide a thorough and up-to-date survey of domain adaptation techniques in popular use for machine translation at time of writing.
- We provide a taxonomy of adaptation approaches corresponding to the process of developing an NMT system, letting practitioners identify the most relevant techniques.
- We distinguish between domain adaptation and multi-domain adaptation, and describe relevant techniques for each.
- We focus throughout on various pitfalls in adaptation as they apply to the surveyed techniques, particularly forgetting and overfitting.
- We explore the potential of framing other NMT research areas as adaptation problems.

2. What Constitutes a Domain Adaptation Problem?

Adapting to a ‘domain’ for machine translation has come to refer to a number of disparate concepts. In van der Wees (2017) various aspects combine to form a domain for statistical MT. In this section, we briefly review their findings and the surrounding literature. We then clarify what is meant by a domain, and by a domain adaptation problem, for the purposes of this survey.

2.1 Exploring the Definition of Domain

Many possible categories and sub-categories may describe a language domain for purposes as disparate as education, text classification or data retrieval. These may not be well defined, especially across fields of research (Sinclair & Ball, 1996). However, we concentrate on three primary elements of a domain identified by van der Wees (2017) in the context of machine translation research: provenance, topic and genre.

Provenance is the source of the text, usually a single discrete label. This may be a narrow description, such as one news article by a single author, or an extremely broad description, such as the 37M English-German web-crawled sentence pairs in the cleaned Paracrawl corpus (Bañón et al., 2020). Importantly, the provenance of a test sentence may not be known outside well-defined tasks in the research community, such as the WMT shared tasks. Even if the provenance is known, it may not be useful for adapting translation performance unless it also corresponds to topic or genre. To take an example from commercial MT, it may not be helpful to know that a sentence is generated by a particular customer unless their writing has specific features that should be reflected in a translation. However, it is possible,

if less common, to require domain-specific translations of sentences of a given provenance instead of using sentence content to determine domain. For example, two different customers may have identical source sentences but different preferred translations for certain terms.

Topic is the subject of text, for example news, software, biomedical. Topic often reveals itself in terms of distribution over vocabulary items. Notably, topic domains can overlap: each word in the vocabulary may have different topic-conditional probabilities, and a document (or sentence) may be classified as a mixture of topics (Blei et al., 2003). The topic(s) of a given sentence can also be determined as a mix of latent topics determined over a large dataset. Discrete topics can however be defined, and may be used to resolve lexical or domain ambiguity for translation, as in the case of domain labels (Kobus et al., 2017).

Genre may be interpreted as a concept that is orthogonal to topic, consisting of function, register, syntax and style (Santini, 2004). For example, multiple documents about a company may share topics and use similar vocabulary, such as the company name or specific products. However, a recruitment document, product specification, or product advertisement would all constitute different genres (Lee & Myaeng, 2002). In some languages it is possible to convey the same information about the same topic in multiple genres, for example by changing formality register. Various cues exist for automatic genre detection in text, including relative frequency of different syntactic categories, use of particular characters such as punctuation marks, and sentence length distribution (Kessler et al., 1997).

2.2 Domain Adaptation and Multi-Domain Adaptation

For this article, a domain adaptation problem fulfills the following broad requirements:

1. We wish to *improve translation performance* on some set of sentences with identifiable characteristics. The characteristic may be a distinct vocabulary distribution suggesting a topic, a stylometric feature such as sentence length distribution, other features of language, or otherwise meta-information such as provenance.
2. We wish to *avoid retraining* the system from scratch. Retraining is generally not desirable since it can take a great deal of time and computation, with correspondingly high financial and energy costs. Retraining may also be impossible, practically speaking, if we do not have access to the original training data or computational resources.

We are also interested in *multi-domain* adaptation. For the purposes of this article a multi-domain adaptation problem requires that in addition to the previous circumstances:

3. We wish to achieve good translation performance on text from *more than one domain* using the same system.

We find the multi-domain scenario interpreted in a number of ways in the NMT literature. Some common settings for multi-domain NMT are as follows:

- Domain labels are known for all sentences. This is often assumed when adapting to multiple domains (Pham et al., 2021). In this case multi-domain control or domain-specific subnetworks can be used (Section 6).

- Some sentences have unknown domain labels. This may also apply when domains of interest are considered to be fuzzy and potentially overlapping. The closest ‘seen’ domain could be inferred by a domain classifier (Section 6.2) or multiple possible labels could be accounted for at inference time (Section 8).
- Text from a labeled but unseen domain may appear at inference time. The system may classify it as the closest ‘seen’ domain, as for unknown-domain text in the previous point. This scenario also lends itself to few-shot unsupervised adaptation (Section 5.3), or zero-shot adaptation using priming techniques (Section 8.3).
- The domain of the pre-training data is one of the multiple domains we wish to be able to translate well. This is sometimes referred to as ‘continual learning’ (Cao et al., 2021). Maintaining performance on the pre-training domain involves avoiding catastrophic forgetting during adaptation, a problem we discuss in Section 4.2. While there are many potential solutions to this problem, most popular approaches involve adjusting the training algorithm (Section 7).

Clearly it is possible for some of these multi-domain settings to co-occur. For example, we may want to avoid catastrophic forgetting on the original domain and also account for new domains at inference time. We may have domain labels during training, but need to infer domain at inference time. However, they are simpler to consider in isolation, and solutions addressing each individually can also often be combined.

In Section 9 we provide case studies on separate lines of machine translation research: low-resource language translation, gender handling in translation, and document-level translation. These are well-established NMT research topics in their own right. Elements of at least the latter two may be viewed as relevant to genre or topic in terms of gender- and document-consistent grammar and lexical choice. However, they are not always treated as relevant to domain adaptation. We will demonstrate that these lines of translation research can also be treated as domain adaptation problems.

3. A Brief Overview of Neural Machine Translation

Machine Translation (MT) aims to translate written text from a source natural language to a target language. Originally accomplished with statistical MT (SMT) using phrase-based frequency models, the state-of-the-art in high-resource language pairs like English-German has been achieved by Neural Machine Translation (NMT) since 2016 (Bojar et al., 2016). More recently NMT has also outperformed SMT on low-resource language pairs on generic domains (Sennrich & Zhang, 2019) or those with sufficient domain-specific data (Ahmadnia & Dorr, 2020). Here we summarize NMT, focusing on elements relevant to the domain adaptation techniques in this survey. For an in-depth discussion of general approaches to NMT, we direct the reader to broader surveys, e.g. Koehn (2020), Stahlberg (2020).

3.1 Representing Language for Neural Machine Translation

NMT models usually take as a single training example two token sequences: \mathbf{x} representing a source language sentence and \mathbf{y} representing its target language translation. Tokens are represented as integer IDs $\in V$ where V is the source or target vocabulary, which may

represent words, subwords, or other features such as domain tags. $|V|$ is usually limited to tens of thousands of tokens due to the computational complexity of the softmax that maps embeddings to discrete tokens.

The vocabulary should convey information that is useful for translation, while remaining computationally tractable. Unknown or out-of-vocabulary (OOV) tokens are undesirable, as are large vocabulary sizes and very long sentence representations. This causes trade-offs – for example, a character vocabulary might be very small with no OOVs but very long sequences. Here we describe ways to address these trade-offs when representing language for NMT. The chosen approach often affects domain adaptation, since vocabulary and terminology are strongly indicative of a domain.

3.1.1 WORD VOCABULARIES

Early approaches to NMT focused on word vocabularies, usually either the top $|V|$ words by frequency (Cho et al., 2014) or all training words with frequencies over a threshold (Kalchbrenner & Blunsom, 2013). Out-of-vocabulary (OOV) words are represented by a special UNK token. Word vocabularies must account for sparsity. According to Zipf (1949) a word’s occurrence probability is inversely related to its vocabulary rank: a large proportion of words in a language are rare in a given corpus. Representing every rare word as a unique vocabulary item is inefficient, and may not allow good learned representations.

3.1.2 SUBWORD VOCABULARIES

Subword vocabularies address the rare word problem by representing words as sequences of more frequent subwords. Sennrich et al. (2016c) first propose NMT on subword sequences using Byte Pair Encoding (BPE) (Gage, 1994). A BPE vocabulary initializes with the set of all character tokens, and with all words represented as character sequences. BPE then iteratively merges the most frequent token pair. Frequent words become single tokens, while rare words can be represented as character sequences. The most extreme subword segmentation is a character vocabulary (Ling et al., 2015; Kim et al., 2016; Costa-jussà & Fonollosa, 2016; Cherry et al., 2018). Subword vocabularies can be constructed using syllables (Assylbekov et al., 2017), language model scores (Kudo, 2018) or linguistic information (Ataman et al., 2017; Huck et al., 2017; Macháček et al., 2018). However, BPE is currently widely accepted as a default vocabulary scheme for NMT. BPE granularity can also affect performance: Ding et al. (2019) find that low-resource NMT benefits from fewer BPE merges to maintain frequency for individual subwords. Gallé (2019) and Salesky et al. (2020) similarly find best performance when balancing high-frequency subwords with short sequence lengths overall.

3.1.3 TAGS

NMT language representations primarily vary in granularity, from words to characters. However, sentence representations can also be augmented with elements not present in the original text. One example relevant to domain adaptation uses externally-defined tags to indicate a particular feature of the sentence. Previous work has used sentence tags to convey domain (Kobus et al., 2017; Britz et al., 2017), speaker gender (Vanmassenhove et al., 2018), or target language formality (Sennrich et al., 2016a; Feely et al., 2019). Multiple

tags can be used throughout source and target sentences, for example indicating linguistic features (Sennrich & Haddow, 2016; García-Martínez et al., 2016; Aharoni & Goldberg, 2017; Saunders et al., 2018) or custom terminology use (Dinu et al., 2019; Michon et al., 2020). Tags are usually incorporated throughout training, implicitly requiring the availability of reliable tags for large training datasets, although new tags can also be introduced during model adaptation (Saunders et al., 2020). Tags can be treated as any other token in a sentence, or handled in special ways by the neural network (Section 6).

3.1.4 REPRESENTING EXTRA-SENTENCE CONTEXT

NMT typically translates between sentence pairs but can incorporate additional context, from the previous sentence (Tiedemann et al., 2017) to the whole document (Junczys-Dowmunt, 2019; Macé & Servan, 2019). Early work on context for recurrent NMT showed improvements in BLEU score (Wang et al., 2017) and context-based metrics like pronoun accuracy (Jean et al., 2017). For self-attention models, context improves lexical cohesion and anaphora resolution (Voita et al., 2019b, 2018), but has inconsistent effects on overall quality (Stahlberg et al., 2019). Notably, Kim et al. (2019b) suggest that context use may simply regularize NMT, and that context-specific information for e.g. lexical cohesion can be retained in very minimal forms, such as terminology tags or syntactic annotation – much like domain tags. While we focus on sentence-level NMT in this article, in Section 9.3 we describe approaches to document-level NMT that take a domain adaptation approach.

3.2 Neural Translation Model Architecture

To the NMT architecture, all inputs are sequences of generic ‘tokens’, represented as integer values. The NMT model input is $\mathbf{x} = x_1, \dots, x_I, x_i \in V_{src}$, and it must produce another integer sequence: $\mathbf{y} = y_1, \dots, y_J, y_j \in V_{trg}$. We refer to the phase of updating model parameters given \mathbf{x} and \mathbf{y} as training, to the phase where parameters are fixed and we generate hypotheses $\hat{\mathbf{y}}$ without reference tokens \mathbf{y} as inference, and to the process of producing any sequence with the decoder as decoding. While many NMT networks can learn mappings between \mathbf{x} and \mathbf{y} that generalize to unseen \mathbf{x} during inference, we focus on the Transformer (Vaswani et al., 2017), the de facto standard for NMT in recent years.

3.2.1 CONTINUOUS TOKEN EMBEDDINGS

A Transformer first maps a sequence of integers $\mathbf{x} \in V_{src}$ to embeddings. A $|V| > 10\text{K}$ -dimension token representation is intractable when operating on thousands of tokens per batch. It is also discrete: the effect of changing one token may not generalize. Bengio et al. (2003) proposed instead mapping each vocabulary item to a continuous feature vector or ‘embedding’. These are a more tractable size: typically $d \leq 1024$. Embeddings trained with context-related objectives – as in NMT – may exhibit local smoothness and therefore generalize: words with similar contexts tend to have similar embeddings (Collobert & Weston, 2008; Turian et al., 2010). In particular embeddings are often similar for tokens that belong to the same category – that is, the same domain (Collobert et al., 2011; Mikolov et al., 2013). These properties can aid data selection for adaptation (Section 5).

3.2.2 TRANSFORMERS: ATTENTION-BASED ENCODER-DECODER NETWORKS

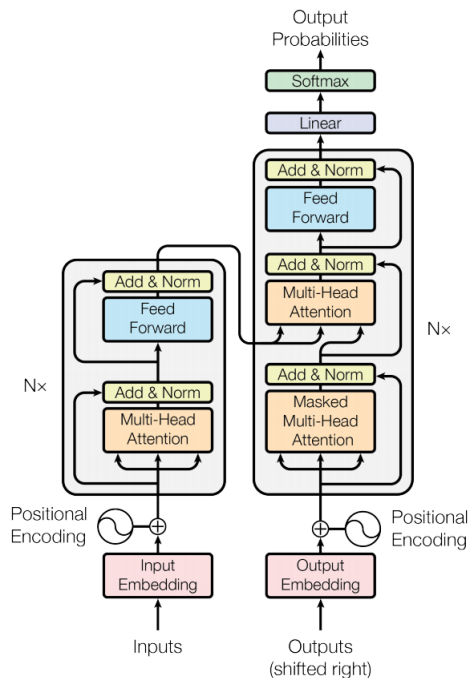


Figure 1: Illustration of the Transformer model architecture, Figure 1 from Vaswani et al. (2017). The model is made up of encoder (left) and decoder (right) subnetworks.

The Transformer network, proposed by Vaswani et al. (2017) and illustrated in Figure 1, is a self-attention-based encoder-decoder model. It remains state-of-the-art for NMT at time of writing (Akhbardeh et al., 2021), improving parallelizability and quality over earlier RNN-based NMT models. The Transformer encoder produces a source sentence embedding from the source token embeddings. Its decoder produces a target embedding from the target tokens during training or the partial translation hypothesis during inference. An encoder-decoder attention network relates the encoder output and decoder state. During inference the decoder translates based on the target embedding and the encoder-decoder attention.

Recurrent model attention mechanisms such as the earlier state-of-the-art described in Bahdanau et al. (2015) relate different positions in different sequences in order to learn a ‘context representation’. The Transformer’s self-attention embedding for a sequence is instead calculated to relate different positions in the same sequence. Self-attention is on source embeddings x at the encoder input and target embeddings y at the decoder input. For a multi-layer encoder or decoder self-attention is calculated on the output of the previous layer. Any of these elements - encoder, decoder, and attention network - can be made domain-specific, or duplicated for a specific domain (Section 6.2).

3.2.3 INCREASING MODEL DEPTH

The Transformer encoder and decoder are usually formed of several identical layers, each operating on the output of the layer before. Increasing the number of layers can cause training difficulties, as loss gradients must propagate through more layers. Difficulties can be mitigated by adding residual networks, changing each layer output $f(z)$ to $f(z) + z$ (He et al., 2016). This gives each encoder or decoder layer access to the first layer’s input. Residual networks are also needed for some domain-specific subnetworks like adapter layers (Bapna & Firat, 2019b) (Section 6.3). Deep models can significantly improve NMT (Wang et al., 2019), but shallower models still often perform better for low-resource NMT (Sennrich & Zhang, 2019; Nguyen & Chiang, 2018). The potential pitfalls of deeper models are worth considering in relation to architectural approaches to domain adaptation (Section 6).

3.3 Training Neural Machine Translation Models

Once an NMT model architecture is determined, its parameters must be adjusted to map from source \mathbf{x} to target \mathbf{y} . NMT model parameters are usually trained by backpropagation with a gradient descent optimizer (Rumelhart et al., 1986), which requires some objective on the training set. Standard training objectives vary weight in the gradient direction of maximising log likelihood, or equivalently minimising cross-entropy loss, for training examples (Baum & Wilczek, 1988; Levin & Fleisher, 1988). Such objectives condition the model on real reference sentences. However, during inference, the translation must be conditioned on the prefix of the model’s own hypothesis $\hat{\mathbf{y}}$. This difference in conditioning can cause poor behaviour during inference (Bengio et al., 2015; Ranzato et al., 2016). Avoiding such over-exposure to training examples has motivated parameter and objective regularization methods during training. Here we focus on objective functions which are relevant to domain adaptation.

3.3.1 CROSS-ENTROPY LOSS

The most common NMT training objective is varying weights θ to maximise log likelihood of training examples, known as Maximum Likelihood Estimation (MLE) (Baum & Wilczek, 1988; Levin & Fleisher, 1988).

$$\hat{\theta} = \operatorname{argmax}_{\theta} \log P(\mathbf{y}|\mathbf{x}; \theta) \quad (1)$$

MLE is equivalent to minimizing cross-entropy loss L_{CE} between the generated output distribution and the references if each token has one reference label $q(y' = y_j|\mathbf{x}; \theta) = \delta(y_j)$:

$$\hat{\theta} = \operatorname{argmin}_{\theta} \sum_{j=1}^{|\mathbf{y}|} -\log P(y_j|\mathbf{y}_{1:j-1}, \mathbf{x}; \theta) = \operatorname{argmin}_{\theta} L_{CE}(\mathbf{x}, \mathbf{y}; \theta) \quad (2)$$

This survey will describe several variations on the MLE loss for domain adaptation (Section 7.1), as well as adaptation-specific applications of non-MLE objectives which seek to address some of the downsides of MLE (Section 7.4).

3.3.2 OBJECTIVE FUNCTION REGULARIZATION

We can change Eq. 2 by adding a regularization term L_{Reg} to the loss function itself:

$$\hat{\theta} = \underset{\theta}{\operatorname{argmin}} [L_{CE}(\mathbf{x}, \mathbf{y}; \theta) + \lambda L_{\text{Reg}}(\theta)] \quad (3)$$

One option is an L2 penalty, $L_{\text{Reg}} = \sum_i \theta_i^2$ – variations on L2 regularization are popular in regularizing adaptation (Section 7.1.2). Alternatively L_{Reg} can be an objective from another task, known as multi-task learning. Translation-specific multi-task terms include a coverage term (Tu et al., 2016), a right-to-left translation objective (Zhang et al., 2019b), the ‘future cost’ of a partial translation (Duan et al., 2020), or a target language modelling objective (Gülcehre et al., 2015; Sriram et al., 2018; Stahlberg et al., 2018a). Another approach is dropout: randomly omitting a subset of parameters θ_{dropout} from optimization for a training batch (Hinton et al., 2012). This corresponds to regularization with $L_{\text{Reg}}(\theta) = \infty$ for $\theta \in \theta_{\text{dropout}}$, 0 otherwise.

Objective function regularization is often used to avoid catastrophic forgetting during NMT domain adaptation. We discuss methods relating to dropout in Section 7.1.1, and other objective regularization methods in Section 7.1.2.

3.3.3 OUTPUT DISTRIBUTION REGULARIZATION

MLE assumes that a reference sentence is far more likely than any other translation. This encourages large differences in likelihood between training examples and language not seen in training: overfitting. Overfitting can reduce the model’s ability to cope with novel data during inference. This is especially relevant for adaptation to small domains with limited training examples. Overfitting can be mitigated by regularizing the distribution over output labels during training, for example using label smoothing (Szegedy et al., 2016). This replaces the single target label $q(y'|\mathbf{x}; \theta) = \delta(y_j)$ used to derive Eq. 2, instead smoothing the label distribution by hyperparameter ϵ towards a uniform distribution over the vocabulary:

$$q(y'|\mathbf{x}; \theta) = (1 - \epsilon)\delta(y_j) + \frac{\epsilon}{|V_{\text{trg}}|} \quad (4)$$

$$\hat{\theta} = \underset{\theta}{\operatorname{argmin}} \sum_{j=1}^{|\mathbf{y}|} \sum_{y' \in V_{\text{trg}}} -q(y'|\mathbf{x}; \theta) \log P(y'|\mathbf{y}_{1:j-1}, \mathbf{x}; \theta) \quad (5)$$

Instead of smoothing with a uniform distribution $\frac{1}{|V_{\text{trg}}|}$, label smoothing can incorporate prior information about the target language. For example, the smoothing distribution can come from a larger ‘teacher’ model (Hinton et al., 2015). Pereyra et al. (2017) explore smoothing towards a unigram distribution over the vocabulary, and address over-confidence directly by penalizing peaky, low-entropy output distributions. Related distribution regularization schemes have been applied to domain adaptation (Section 7.1.3).

3.4 Inference with Neural Machine Translation Models

An NMT model trains on source and target language sentences \mathbf{x} and \mathbf{y} , and learns to model $P(\mathbf{y}|\mathbf{x})$. During inference, the model has access only to \mathbf{x} , and must produce a translation:

$$\hat{\mathbf{y}} = \operatorname{argmax}_{\mathbf{y}} P(\mathbf{y}|\mathbf{x}) \quad (6)$$

In this survey we focus on autoregressive NMT inference, by far the most common approach. For autoregressive inference the model produces one output token at each inference step j , which is conditioned on the source sentence and all previously output tokens. Ideally:

$$\hat{y}_j = \operatorname{argmax}_{y_j} P(y_j|\mathbf{y}_{1:j-1}, \mathbf{x}) \quad (7)$$

There are $|V_{trg}|^j$ possible partial translations ending in the j^{th} output token. Exploring all of them is impractically slow even using likelihood pruning to reduce the total (Stahlberg & Byrne, 2019). Nevertheless, approximations to this inference objective work well in practice. Here we focus on those which have variants used for domain-specific translation.

3.4.1 BEAM SEARCH

Beam search is the most common approximation for NMT inference. It tracks the top N partial hypothesis ‘beams’ by log likelihood. At each inference step all possible single-token expansions of all beams are ranked, and the new top N selected. Search continues until all beams terminate with an end-of-sentence token or exceed a maximum length. ‘Greedy’ search is the special case where $N = 1$, which produces the most likely next token at each step. Beam search variations include optimizing for diverse hypotheses (Vijayakumar et al., 2016; Li & Jurafsky, 2016) or adequacy (Wu et al., 2016). In Section 8.2 we describe beam search variants intended to produce domain-specific translations.

3.4.2 ENSEMBLING

Conducting inference with an ensemble of MT models allows consensus on which tokens to expand and track in beams at each inference step, usually outperforming individual models (Sim et al., 2007; Rosti et al., 2007). Ensembles can integrate scores from different NMT architectures (Stahlberg et al., 2018b) or language models (Vaswani et al., 2013; Gülcehre et al., 2015). Many schemes for ensembling exist, from majority vote to minimizing Bayes risk under some metric (Rokach, 2010). Ensembling by static weighting is commonly used in NMT. For a K -model ensemble with a weight W_k for each model, the ensemble translates:

$$\hat{\mathbf{y}} = \operatorname{argmax}_{\mathbf{y}} P(\mathbf{y}|\mathbf{x}) = \operatorname{argmax}_{\mathbf{y}} \sum_{k=1}^K W_k P_k(\mathbf{y}|\mathbf{x}) \quad (8)$$

A downside of ensembling is reduced efficiency. The expensive softmax calculation takes place K times, and often all ensemble models are stored in memory simultaneously. Efficiency can be improved if one model can learn to reproduce the ensemble’s behaviour, using simplification schemes like ensemble knowledge distillation (Freitag et al., 2017; Fukuda et al., 2017) or ensemble unfolding (Stahlberg & Byrne, 2017). Many approaches to ensembling assume that all ensembled models generate in the same way, e.g. lexical items from the same vocabulary, and that fixed ensemble weights can be used regardless of input content or model domain. Section 8.1 reviews domain-specific ensembling approaches which relax these assumptions.

4. Fine-Tuning as a Domain Adaptation Baseline, and its Difficulties

While the previous section covered ‘default’ approaches to NMT, e.g. a Transformer with a BPE vocabulary trained with cross-entropy loss, this section covers fine-tuning, a ‘default’ approach to domain adaptation for NMT. It also describes some difficulties with this approach which motivate many variations in the remainder of this survey.

Given an in-domain dataset and a pre-trained neural model, domain adaptation can often be achieved by continuing training the model on that dataset: ‘fine-tuning’. Fine-tuning initially involved monolingual data only for SMT (Lavergne et al., 2011). By contrast fine-tuning end-to-end NMT models, to the best of our knowledge first proposed in Luong and Manning (2015), usually requires bilingual parallel data. In both cases the pre-trained model trains for relatively few additional iterations with the original loss function applied to the in-domain data.

Fine-tuning is a straightforward and efficient approach to adaptation: adaptation sets might have thousands of sentence pairs, while training the original model could use millions. Luong and Manning (2015) find that in-domain NMT can improve dramatically under fine-tuning despite the low computational requirements. Servan et al. (2016) observe that fine-tuning on in-domain data may not achieve the same translation quality as retraining the entire system on combined generic and in-domain data, but report a high proportion of the improvement in less than one percent of the time needed for complete retraining.

Variations and alternatives to fine-tuning for domain adaptation are discussed in the remainder of this survey. They include changing the data (Section 5), the network architecture (Section 6), or the tuning procedure itself (Section 7). Alternatively, domain-specificity can be achieved with domain-adaptive inference (Section 8). These variations on simple fine-tuning can be combined. For example, a variation on fine-tuning could hypothetically tune on a synthetically expanded in-domain set (Wei et al., 2020a), using domain-specific parameters (Bapna & Firat, 2019b), and inference-time domain classification to determine whether to use the domain-specific parameters (Huck et al., 2015). However, in the literature it is common to compare proposed domain adaptation interventions in isolation to a baseline of simple fine-tuning on a pre-defined dataset.

Simple fine-tuning is associated with a number of difficulties. We present three here – insufficient tuning data, forgetting, and overfitting – to motivate the variations on fine-tuning data, architecture, tuning and inference presented in the remainder of this survey.

4.1 Difficulty: Not Enough In-Domain Data

Fine-tuning strong pre-trained models on small, trusted sets has become a popular approach for machine translation fixed-domain tasks. A version of this commonly appearing in WMT shared task submissions is tuning on test sets released for the same task in previous years (Schamper et al., 2018; Koehn et al., 2018a; Stahlberg et al., 2019). Shared task adaptation sets are likely to contain many sentences stylistically very similar to those in the test set.

However, we cannot always assume the existence of a sufficiently large and high-quality in-domain parallel dataset. Outside well-defined domain adaptation shared tasks, there may be little or no bilingual in-domain data available for a particular adaptation scenario. In Section 5, we describe data-centric adaptation methods which aim to expand the in-domain parallel corpus available for tuning.

4.2 Difficulty: Forgetting Previous Domains

The difficulty of ‘catastrophic forgetting’ occurs when translating domains other than the fine-tuning domain. If a neural model with strong performance on domain A is fine-tuned on domain B , it often improves performance on B at the expense of extreme performance degradation on A (McCloskey & Cohen, 1989; Ratcliff, 1990). This is especially problematic if we intended to train a model capable of translating text from multiple domains, regardless of when during training a given domain was learned.

Some forgetting may be permissible if the model is intended to translate only a small amount of highly specific data. Examples include adapting a new model to translate each individual test sentence (Farajian et al., 2017; Li et al., 2018; Mueller & Lal, 2019) or document (Kothur et al., 2018). Even in this scenario, forgetting all data that is unlike the tuning data may cause poor performance on even very specific target-domain data – this is the related scenario of overfitting (Section 4.3).

A straightforward approach to avoiding forgetting under fine-tuning is to simply tune for fewer steps (Xu et al., 2019), although this introduces an inherent trade-off between better performance on the new domain and worse performance on the old domain. Other approaches to good performance on both new and old domains with a single model might introduce additional domain-specific parameters or subnetworks for the new domain, as described in Section 6. Alternatively they may involve changing the training approach, for example to add regularization relative to the original domain, as described in Section 7.

4.3 Difficulty: Overfitting to New Domains

Overfitting or ‘exposure bias’ is common when the fine-tuning dataset is very small or repetitive. The model may be capable of achieving excellent performance on the precise adaptation domain B , but any small variations from it – say, B' – cause difficulties. Farajian et al. (2017) and Li et al. (2018) observe this degradation after adapting a model to just tens of training sentences with certain learning settings. Overfitting is particularly relevant in cases of domain mismatch between the test sentence domain and the fine-tuning domain (Wang & Sennrich, 2020).

The effect can also occur when the tuning data contains small irregularities. B' becomes a subset of B containing misaligned sentence pairs, or sentences with idiosyncratic and irrelevant terminology use. In this case test sentences may match topic and genre with overall tuning domain B , but we do not wish to translate them as in B' – the domain with irregularities. Ott et al. (2018) observe one example of this effect: a model trained on data with a high noise proportion may fall back on language model behaviour and generate some common target sentence from the dataset, or simply copy the input sentence.

Overfitting can be mitigated by expanding the in-domain corpus so that it is harder to overfit (Section 5), or by adjusting the adaptation procedure to adapt less aggressively to the new domain (Section 7). An alternative which sidesteps fine-tuning, and therefore much of the potential for overfitting or forgetting, is adjusting the test-time inference procedure without changing the model at all (Section 8).

5. Data-Centric Adaptation Methods

A domain is often identifiable by features of its data. Topic and genre as described in Section 2.1 are defined in terms of vocabulary and syntactic style for in-domain text. However, the existence of a sufficiently large in-domain parallel dataset is often taken for granted by fine-tuning (Section 4), as well as many model adaptation methods described in Sections 6 and 7 of this survey. Here we consider how an in-domain dataset can be created or expanded.

By data-centric adaptation methods, we refer to methods that select or generate appropriate in-domain data. Natural data – produced by a human – may be selected from some large generic corpus or corpora according to some domain-specific criteria (Section 5.1). A special case of natural data selection is filtering an existing in-domain corpus for cleaner or more relevant data (Section 5.2). If only monolingual natural data is available, corresponding source or target sentences can instead be generated by neural models. Data generation may also diversify an easily over-fitted bilingual corpus, or text that is mismatched with the domain of interest (Sennrich et al., 2017). Semi-synthetic bilingual data can be produced by forward or back translation of existing natural data (Section 5.3). Additional in-domain bilingual data can be partially synthesised from an in-domain natural dataset using noising or simplification (Section 5.4). Finally, purely synthetic data - neither source nor target produced by a human - can be generated for adaptation (Section 5.5).

5.1 Selecting Additional Natural Data for Adaptation

Grangier and Iyer (2022) note that selecting a small dataset for domain adaptation involves an inherent trade-off. A smaller dataset may be more domain relevant, but the impact of any deviation from the target domain will be magnified. Nevertheless, data selection from generic parallel text is a popular adaptation approach, and many approaches can be found in the literature. We divide these into methods using discrete token-level measures for retrieval, those using continuous sentence representations for retrieval, or those that use some external model for scoring.

Sentences can be selected by discrete token overlap with in-domain data. Many such methods have been applied to statistical MT, which typically does not involve continuous sentence representations. For example, Eck et al. (2004) retrieve sentences for SMT language model adaptation by TF-IDF and n-gram-overlap relevance measures. The NMT adaptation scenario where few in-domain sentences are available has likewise seen fine-tuning improvements when selecting additional data by straightforward n-gram matching, as demonstrated for single-sentence adaptation by Farajian et al. (2017) and Li et al. (2018). Koehn and Senellart (2010) instead use a fuzzy match score based on word-level edit distance to retrieve similar source sentences for SMT. Xu et al. (2019) apply fuzzy matching retrieval to NMT, using retrieved sentences as in-domain adaptation data for each test sentence. However, they obtain slightly better results using n-gram matching. Other applications of fuzzy matching for NMT use the retrieved sentences directly at inference time as for the Koehn and Senellart (2010) approach – we discuss these in Section 8.3.2.

It can be beneficial to emphasise selected data diversity, avoiding creation of a repetitive adaptation set. Poncelas et al. (2019b) take this approach. They select either by overall n-gram rarity, or by using Feature Decay Algorithms which adjust the ‘value’ of an n-gram

based on how often it has been sampled. They find both approaches outperform TF-IDF similarity selection, which does not account for retrieved dataset diversity.

Tuning on data selected by discrete lexical overlap measures can be highly effective and also relatively efficient. For example, Farajian et al. (2018) and Li et al. (2018) both note that tuning on sentences that are highly similar to test sentences can improve terminology use after adaptation, while requiring only a few adaptation examples per test sentence. However, fine-grained adaptation can result in overfitting if matching is imperfect, with Chen et al. (2020b) finding better performance tuning on sentences selected by n-gram match over the entire test set compared to matching per-sentence.

In-domain data can also be retrieved based on continuous representations. For example, Wang et al. (2017a) expand an in-domain corpus by selecting sentences with embeddings similar to in-domain sentence embeddings. Bapna and Firat (2019a) combine n-gram-based retrieval with dense vector representations to improve candidate retrieval across multiple domains. Lohar et al. (2022) combine continuous and discrete approaches, using an embedding comparison model to determine key phrases from each test input document, then selecting data for adaptation by discrete key phrase matching.

Whether using a discrete or continuous representation, matching usually occurs in the source language, as this allows direct comparison between source and test input sentences. Notable exceptions are the recent approaches of Cai et al. (2021) and Vu et al. (2021), both of whom use a cross-lingual network to retrieve target language sentences with similar embeddings to an input source sentence. While this requires a cross-lingual model, it has the advantage that ‘real’ relevant target sentences are retrieved. Many previously-described approaches assume the retrieved source sentences will have good reference translations, or retrieve only source sentences and then generate synthetic targets (Section 5.3.2).

Sentences can be selected by scoring for domain-relevance using external models. Moore and Lewis (2010) select in-domain data to train language models by scoring the data under in-domain and general domain language models and taking the cross-entropy difference as in-domain relevance. Axelrod et al. (2011) add a bilingual cross-entropy difference term to the Moore and Lewis (2010) approach to select parallel data for SMT adaptation. More recently, Vu and Moschitti (2021) select documents based on support vector machines trained to classify between the target domain data and random mixed-domain data. However, van der Wees et al. (2017) note that such static discriminative filtering schemes can struggle with very similar in-domain and generic corpora. Axelrod (2017) instead suggests ‘cynical data selection’: repeatedly selecting the sentence that most reduces the relative entropy for modelling the domain of interest. Importantly given the ambiguity over the meaning of ‘domain’ mentioned in Section 2.1, this does not involve actually defining a domain. Indeed, Santamaría and Axelrod (2017) note that using classifiers for data selection is not necessarily a good conceptual approach as a sentence pair may easily appear in multiple corpora. They instead reframe the ‘in-domain’ and ‘general domain’ corpora as data that we know we are interested in or do not yet have an opinion about.

Aharoni and Goldberg (2020) dispense with assigned corpus labels and show that adapting to data identified by unsupervised domain clustering using large language models matches or out-performs tuning on the ‘correct’ domain-labelled data. Del et al. (2021) extend this idea to show that NMT models themselves can cluster sentences by embedding

into domains that allow for better adaptation than pre-defined corpora. We regard these results as an indication that arbitrary labels should not be relied upon as domain identifiers.

5.2 Filtering Existing Natural Data

Taking an existing corpus label as indicative of domain is a simple approach to natural data selection for adaptation, despite the deficiencies mentioned at the end of Section 5.1. If we do start with a pre-labeled corpus, data filtering can be applied to ensure that the selected data is actually domain-representative. For example, Taghipour et al. (2011) map sentences in a parallel corpus to a feature space, and mark the most novel pairs in the feature space as noise to be removed. Care must be taken not to diminish the training space too far: for example, Lewis and Eetemadi (2013) filter sentences for SMT while attempting to maintain maximum n-gram coverage.

A special case of data filtering is targeted to remove ‘noisy’ training examples, as for example in the WMT parallel corpus filtering task (Koehn et al., 2018b). Data cleaning may involve ensuring source and target sentences in the training sentence are well-aligned, contain the languages of interest, are not too long or contain too many non-words, such as HTML tags (Khayrallah & Koehn, 2018; Berard et al., 2019). Cleaning methods can be very similar to domain data selection, for example with extensions of previously-mentioned cross-entropy difference filtering (Moore & Lewis, 2010) to bilingual training examples, where the ‘in-domain’ models are trained on verifiably clean data only (Junczys-Dowmunt, 2018a, 2018b). While data cleaning has become a standard step in data preprocessing when training any NMT system, we note it is particularly crucial to domain adaptation if only a small in-domain dataset is available, as discussed in Section 4.3, and is widely mentioned in the submissions for domain-specific shared tasks (Bawden et al., 2019).

5.3 Generating Synthetic Bilingual Adaptation Data from Monolingual Data

Bilingual training data that is relevant to the domain of interest may not be available. However, source or target language monolingual data in the domain of interest is often much easier to acquire. In-domain monolingual data can be used to construct partially synthetic bilingual training corpora by forward- or back translation. This is a case of bilingual data *generation* for adaptation rather than data selection.

Source or target language sentences in the domain of interest may be extracted from monolingual corpora using techniques described previously in this section. A sufficiently strong source-to-target translation model can forward-translate these sentences, or a target-to-source model can back-translate them. The result is aligned in-domain source and target language training sentences. This allows ‘unsupervised’ MT domain adaptation where parallel data is unavailable. Statistical MT has benefited from both forward translation (Schwenk, 2008) and back-translation (Bertoldi & Federico, 2009; Lambert et al., 2011) in the context of domain adaptation, with back translation performing better in direct comparisons. Here we discuss more recent work on adaptation data generation for NMT.

5.3.1 BACK TRANSLATION

Back translation uses natural monolingual data as target sentences, and requires a target-to-source NMT model to generate synthetic source sentences. Back translations are com-

monly used to augment general domain translation corpora, with strong improvements over models not trained on back-translated data (Sennrich et al., 2016b). Even models trained exclusively on back translations can perform similarly to models trained on natural data (Poncelas et al., 2018). More relevant to this survey, back translation is also used to generate adaptation data for domains where in-domain target sentences are available (Sennrich et al., 2017; Jin et al., 2020). In-domain target sentences may be identified with known topic and genre, for example biomedical papers (Abdul Rauf et al., 2020) or coronavirus-related texts (Mahdiah et al., 2020). Parthasarathy et al. (2020) extract target sentences by finding target language terms that are likely to be in the test set, while Vu et al. (2021) use a cross-lingual domain classification model to extract target sentences that are relevant to test source sentences. These target sentences are then back-translated. The synthetic-source sentence pairs are typically used directly for fine-tuning the model, but can also be used as candidates for a domain-specific data selection scheme (Poncelas & Way, 2019).

Notably, Poncelas et al. (2019a) show that back-translated sentences can improve domain adaptation even when the original examples were used to train the general domain model, suggesting it is not simply the in-domain target sentences but the novel synthetic source sentences that aid adaptation. However, the quality and domain-relevance of the synthetic source sentences is also important for adaptation. Wei et al. (2020a) show that using a general domain model to back-translate in-domain sentences gives poor synthetic source sentences, harming adaptation. Instead, they suggest ‘domain repairing’ the synthetic source sentences before adaptation. Kumari et al. (2021) similarly show that domain-filtering back-translated data improves NMT adaptation.

A related approach proposed by Currey et al. (2017) does not back-translate in-domain target sentences at all, but instead copies them to the source side. This can teach the model to produce rare words present in target language in-domain monolingual data without requiring an expensive back-translation model. Burlot and Yvon (2018) show that training on such copied sentence pairs with language-tags on the source side can give similar results to using true back translation.

Zhang and Toral (2019) and Graham et al. (2020) have shown that training on high proportions of back-translations gives a false sense of model quality on back-translated test sets, while the apparent improvements may not be seen on natural test sets. Effectively, back translations can constitute their own domain. Tagging back translated data, as for a domain (Section 6.1) appears to mitigate this effect (Caswell et al., 2019), and is particularly successful in low-resource scenarios (Marie et al., 2020), further emphasising the domain-like nature of back translation. Adapting exclusively to back-translated data risks adapting to this ‘translationese’ domain instead of the true domain of interest. Jiao et al. (2021) propose mitigating this effect by interposing authentic natural parallel data, if available, when tuning on back-translated data.

5.3.2 FORWARD TRANSLATION

Forward translation generates synthetic target sentences with an existing source-to-target NMT model. A sub-type, self-learning, trains a model with its *own* synthetic translations. Forward translation is less common than back translation, perhaps because the synthetic data is on the target side and so any translation errors may be reinforced and produced dur-

ing inference. However, self-learning can mean more efficient domain adaptation compared to back translation, as it requires no additional NMT model. Chinea-Ríos et al. (2017) demonstrate self-learning for domain adaptation, producing data in the target domain by forward translation and then tuning on that data. Zhang and Zong (2016b) find self-learning particularly beneficial for low-resource target languages. As in previous sections, relevant monolingual data can be selected using domain-specific terminology or n-gram overlap (Haque et al., 2020) or continuous representations (Chinea-Ríos et al., 2017).

A variation of forward translation uses one or more much larger or otherwise stronger ‘teacher’ models to generate in-domain forward translations which are then used to train or tune a ‘student’ model. Gordon and Duh (2020) demonstrate that even strong generic models can benefit from training on in-domain forward translations from a teacher. Currey et al. (2020) extend this idea to train a multi-domain system using forward translations from multiple separate in-domain teachers.

Unlike back translation, forward translation can also be applied to precisely the in-domain source text we actually wish to translate, letting us generate synthetic test target sentences. These can be used for direct fine-tuning, or as a seed to retrieve more natural or synthetic parallel sentences for adaptation as in Poncelas et al. (2018). In this way forward translation permits few-shot translation to previously unseen or unknown domains in a multi-domain scenario.

Adapting to forward translated data does not risk adaptation to a ‘domain’ of synthetic inputs as for back translation, since the source sentences remain natural. However, the presence of target-side synthetic sentences still requires caution, since the model may learn to generate translationese for the domain of interest. Any forward translation errors will also be reinforced by adaptation. The outcome of adapting to either forward or back translations depends heavily on the quality of the data-generating system. Notably, Bogoychev and Sennrich (2019) suggest that forward translation is more sensitive to the quality of the generating system, and that humans typically prefer the fluency of systems trained on back translation.

5.4 Artificially Noising and Simplifying Natural Data

In some cases only a small in-domain dataset is available for adaptation, with no additional bilingual nor monolingual corpora large enough to provide additional relevant sentences. In this scenario we can generate new data from the available in-domain set by changing its source or target sentences in some way. Including these variations in tuning can reduce the likelihood of overfitting or over-exposure to a small set of one-to-one adaptation examples (Bishop, 1995), thus addressing one of the fine-tuning pitfalls of Section 4.

A common example adds artificial noise to source sentences, for example by deleting, substituting or permuting characters or words. As well as mitigating overfitting during adaptation, tuning on such sentences can improve NMT robustness. Tuning on noised data can improve performance on new sentences containing natural mistakes (Vaibhav et al., 2019; Karpukhin et al., 2019) or noise introduced synthetically, for example by a poorly-performing automatic speech recognition (ASR) system (Sperber et al., 2017; Gangi et al., 2019). Some ‘noised’ datasets may themselves be domains of interest, as for the work of Tan et al. (2020) adapting NMT to handle a wider range of natural linguistic variation. Kim

et al. (2021) highlight another benefit of performing domain adaptation with noisy data, in their case heavily segmented sentence fragments. They suggest that sufficient noising may serve to redact any confidential information in adaptation data, mitigating privacy concerns while still benefiting in-domain translation performance. Hu and Neubig (2021) also fine-tune on fragmented phrase translations, but note that fragmentation-noising eventually limits performance on non-noisy longer sentences – a case of over-fitting.

New in-domain examples can also be constructed by simplifying some natural source (Li et al., 2020; Hasler et al., 2017) or target (Agrawal & Carpuat, 2019) sentences. Simplifying sources can make translation easier and is applicable to unseen in-domain sentences at inference time. Simplifying targets during training allows specification of output language complexity, which may itself be a domain feature. However, simplification approaches to adaptation are rare compared to noising approaches. This may be due to the difficulty in obtaining training data and models for simplification – these models themselves may be domain-sensitive or domain-mismatched – while noising text is intuitively often easier. Interestingly, Mehta et al. (2020) find that back-translation can function as a simplification system, and that training on the resulting simplified sentences benefits NMT in low resource settings. This may suggest that adapting to back-translations (Section 5.3.1) is beneficial in part due to the simplified synthetic source sentences.

5.5 Synthetic Bilingual Adaptation Data

A final type of data used for adaptation is purely synthetic. Purely synthetic data may be beneficial when even monolingual in-domain natural data is unavailable, either due to lack of resources or an extremely narrow target domain. Synthetic data may be obtained from an external or induced lexicon. Alternatively synthetic sentences may be produced by a template or an external model, perhaps in conjunction with forward or back translation.

Lexicons have been used effectively when dealing with rare words, OOV words or ambiguous words with multiple senses in the training data (Zhao et al., 2018). For SMT lexicons can be used to mine in-domain translation model probabilities directly (Daumé III & Jagarlamudi, 2011). In NMT domain-specific lexicon probabilities may be incorporated into the NMT loss function (Arthur et al., 2016) or bilingual lexicon entries may be used to construct partially synthetic sentence pairs (Zhang & Zong, 2016a).

Another application of synthetic lexicon data is covering words or phrases for which there is an easily obtainable translation, and which the model is likely to be required to translate. This type of data is usually domain specific: NMT for use on social media may require translations for common greetings, while a biomedical NMT system may have a fixed set of terminology to translate. Hu et al. (2019) adapt to this form of lexicon for a pre-defined domain, and find that it is complementary to tuning on back translated data, with lexicons aiding translation of isolated words and back translations aiding overall in-domain fluency. Kothur et al. (2018) adapt to a lexicon containing novel words in a test document, and find that careful hyperparameter tuning is necessary to mitigate overfitting behaviour when adapting to dictionaries for too many epochs. Peng et al. (2020) mitigate this form of overfitting by combining generic domain parallel data with in-domain dictionaries to synthesize in-domain parallel sentences that are similar to generic domain data.

A final approach generates synthetic in-domain monolingual sentences, then translates them. With known vocabulary of interest, Saunders and Byrne (2020b) fill generic templates with in-domain source terminology, then forward translates. Given pre-trained language models for the target language, Moslem et al. (2022) use existing in-domain target sentences as language model prompts to generate additional examples, then back translates. In both cases the resulting synthetic bilingual data is used for domain adaptation.

6. Architecture-Centric Adaptation

Architecture-centric approaches to domain adaptation typically add trainable parameters to the NMT model itself. This may be a single new layer, a domain discriminator, or a new subnetwork. If domains of interest are known ahead of model training, domain-specific parameters may be learned jointly with NMT model training. However, in a domain adaptation scenario where we wish to avoid retraining, domain-specific parameters can be added after pre-training but before any fine-tuning on in-domain data.

Many of the papers surveyed in this section describe learning domain-specific labels or parameters when training NMT models from scratch. Their focus is thus on domain adaptation as in adapting the domain of a specific translation, but *not* as in the goal in Section 2.2 of achieving good in-domain performance without retraining from scratch. We survey these papers regardless for two reasons. One is completeness: many approaches described here are popular ways to control or improve domain-specific translation. The second is that the choice to train from scratch with domain-specific features is often based only on relative simplicity of implementation. Editing the architecture of a pre-trained model can be more complex than simply retraining a new model with a different architecture. We highlight work where domain-specific parameters or other modeling changes are made after pre-training, notably adapter layers (Section 6.3) as indicative that similar ‘from scratch’ approaches described here could also adapt without requiring retraining.

6.1 Domain Labels and Control




Figure 2: Domain tagging schemes mentioned in this section. a) Prepending an inline domain tag (D) to the source (S) (Kobus et al., 2017; Tars & Fishel, 2018). b) Prepending multiple inline domain tokens to the source (Stergiadis et al., 2021). c) Prepending an inline domain tag to the target (T) (Britz et al., 2017). d) Embedding domain tags separately from plain-text tokens, and concatenating the embeddings so the embedded sequence contains per-token domain features (Kobus et al., 2017; Tars & Fishel, 2018).

Where data is domain-labelled, the labels can be used to signal domain for a multi-domain system. Domain labels may come from a human-curated source or from a trained domain-classifier. Tags can be applied to either the source or target sentence. While tags may be treated as simply another token in the source or target sequence – ‘inline’ tagging – they also frequently involve a separate embedding or subnetwork (Figure 2). We therefore discuss tags here in conjunction with architectural considerations.

Tagging the source allows external control of the target domain. One motivation for source tagging is that some sentences can conceivably be translated into multiple domains, particularly given that domain can encompass genre-related concepts like formality register. Another motivation is that contemporary NMT systems may not successfully infer domain from the source sentence. Instead, there is some evidence that they struggle with aspects of language that would indicate domain to a human, such as lexical ambiguity (Emelin et al., 2020) or formality (Hovy et al., 2020). Consequently any available domain information may help to guide the NMT system during adaptation.

Kobus et al. (2017) explore domain-tagging the source sentence for NMT with either a single inline token or as an embedded feature combined with each token embedding (Figure 2). Tars and Fishel (2018) also source-tag domains as either a single tag or as a tag feature, using real known-domain labels as well as those determined by external supervised classifiers and unsupervised clustering schemes. Both Kobus et al. (2017) and Tars and Fishel (2018) find slightly better performance for domain features than for discrete tags. Both also find that tags improve in-domain translation over simple fine-tuning on the target domain, supporting the view that untagged NMT may not successfully infer domain.

Chu et al. (2017) introduce in-domain and generic domain source sentence tags during fine-tuning. They find this improves performance on the target domain, but show improved performance on the generic domain without tags, and suggest that tags discourage the model from taking advantage of any overlap between the generic and target domains. Stergiadis et al. (2021) account for this problem of overlapping domains by using multiple domain tags for each source sentence. Mino et al. (2020) likewise use multiple inline tags indicating domain and whether data is noised, and show benefits over using single tags. Wang et al. (2021) fine-tune on mixed domains, but improve over simple domain tagging by tuning with counterfactual tagged examples, where some generic domain source sentences are also labelled with a best-fit in-domain tag.

While most of the above work focuses on tagging the source sentence, Britz et al. (2017) instead prepend inline domain tags to the target sentence. They suggest that although NMT models may not learn to infer domain without intervention, target tagging can provide that intervention and teach the model to classify the source sentence’s domain. Importantly this does not require the domain to be explicitly determined for new, user-generated source sentences. However, it also complicates the process of controlling the desired domain.

Although many of the papers proposing domain control apply labels to all sentences when training from scratch, this is not strictly necessary. Indeed Pham et al. (2021), comparing several previously proposed tagging approaches, note that introducing new domains by specifying new labels at fine-tuning time is straightforward.

6.2 Domain-Specific Subnetworks

Another architecture-based approach introduces new subnetworks for domains of interest. This could be a domain-specific element of the original NMT architecture (e.g. vocabulary embedding, encoder, decoder), or a domain embedding determined by a subnetwork not typically found in the Transformer.

Vocabulary embeddings can be made wholly or partially domain-specific. This is related to the domain-embedding feature discussed in Section 6.1, but differs in that the whole embedding relates to the vocabulary item, not to a domain tag. For example, Zeng et al. (2018) determine both domain-specific and domain-shared source sentence representations. Pham et al. (2019) change existing word embeddings to have domain-specific features which are activated or deactivated for a given input, improving lexical choice. Lin et al. (2021) cache keywords for individual users with dedicated subnetworks, effectively tracking user ‘domain’-specific vocabulary embeddings to combine with the generic embeddings. Sato et al. (2020) replace generic translation word embeddings with domain-specific vocabulary embeddings learned by domain-specific language models from monolingual data. Dou et al. (2019a) also learn domain-specific embeddings using monolingual target data, but via an auxiliary language modelling task learned jointly with model tuning.

While most work mentioned in Section 6.1 assumes known domain labels or determines them using an external classification model, it is also possible to learn a domain classifier jointly with the NMT model, as in Britz et al. (2017). Similarly Lachaux et al. (2020) represent the target domain as a latent variable input to a shared decoder. In these cases target domain labels are assumed to be available for training, but at inference time domain labels can be input to the model to control the output domain. Even if the desired domain label is unknown at inference time, multiple sets of outputs can be generated by performing inference with different domain label control.

The domain can be incorporated into the model architecture in ways other than simply changing the word embeddings. For example, Michel and Neubig (2018) learn domain embeddings for many individual speaker ‘domains’ during fine-tuning. They use these to bias the softmax directly, resulting in small improvements over tuning with domain tags prepended to the target sentence. Chu and Dabre (2019) likewise learn domain-specific softmax biases for multiple domains, although they do not see improvements over simply appending a domain token to the source sentence. A domain embedding does not need to be determined directly from the source sentence: Wu et al. (2019) and Wang et al. (2020) both learn domain-specific representations across multiple domains from the output of the encoder, to feed to a shared decoder. These approaches encourage domain-specificity in parts of the model other than the source sentence representation.

An alternative approach is to duplicate encoders or decoders for each domain of interest, although this quickly becomes expensive in terms of added model size and therefore computational and energy cost for training and inference. Gu et al. (2019) use both a shared encoder and decoder for all sentences, as well as domain-specific encoders and decoders for each input domain as determined by a classifier. Jiang et al. (2020) define multi-head attention networks for each domain. The overall attention in each layer is a weighting function of the different domain attention networks, determined by domain classifiers for each layer.

6.3 Lightweight Added Parameters

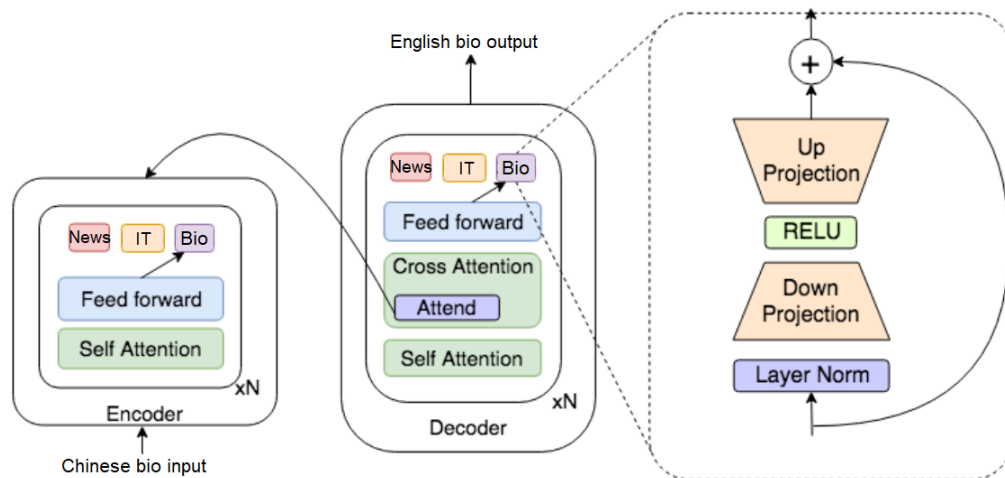


Figure 3: Illustration of adapter layers used for domain adaptation, based on Figure 1 from Bapna and Firat (2019b). A separate set of adapters can be defined for each domain of interest, here News, IT and Bio. The adapter (right) consists of relatively few parameters, and is incorporated into both encoder (left) and decoder (middle) layers.

A lightweight architectural approach to domain adaptation for NMT adds only a limited number of parameters after pre-training. The added parameters are usually adapted on in-domain data while pre-trained parameters are ‘frozen’ – held at their pre-trained values. For example, Vilar (2018) introduces a new component to each hidden unit in the model which can amplify or decrease the contribution of that unit. Bapna and Firat (2019b) instead add new layers, called ‘adapters’. In each case the domain-specific multiplicative unit or adapter layer is tuned on the in-domain data while the rest of the model is frozen.

Adapter layers (Figure 3) have achieved particular popularity as they require very simple and lightweight model modification, and inherently involve no forgetting since the original model parameters are unchanged. They are typically inserted between layers in the encoder and decoder, and may be used alongside a domain discriminator that determines which adapter to use (Pham et al., 2020a). Notably, Abdul Rauf et al. (2020) find that adapters can outperform full model fine-tuning when translating a ‘noisy’ domain, to which the full model overfits. However, many of the above papers highlight that using adapters for NMT domain adaptation requires careful choice of adapter size and number of tuning steps.

6.4 Architecture-Centric Multi-Domain Adaptation

Architectural approaches as described in this section are capable of good performance over multiple domains. In particular, schemes that leave original parameters unchanged and only adapt a small added set of parameters can avoid any performance degradation from forgetting or overfitting by simply using the original parameters. Adapter-like architectural

approaches may therefore have a natural application to continual learning, and can also be a lightweight approach for other multi-domain scenarios.

It is worth noting that using a certain set of parameters for a certain domain implicitly assumes that language domains are discrete, distinct entities. New architecture may be either ‘activated’ if the test set is in-domain, or ‘deactivated’ for better general domain performance as in Vilar (2018). A sentence may be assigned to a single domain, and a label added for that domain as in Tars and Fishel (2018). However, multiple text domains may overlap, and training domains may be mutually beneficial for translation performance – using discrete tags may interfere with this, as found by Chu et al. (2017), and likewise discrete domain-specific subnetworks may not make optimal use of multi-domain data.

One option proposed by Vu et al. (2022) explicitly accounts for the future presence of new, unknown domains when tuning adapter layers. They add a ‘fusion’ adapter layer whose purpose is to generalise from other, domain-specific layers. They find that learning to combine results from domain-specific parameters improves performance on unseen domains. Their finding is corroborated by Pham et al. (2021), who compare various architecture-centric approaches to NMT domain-specificity mentioned in this section, focusing on their performance in a multi-domain setting. They highlight that systems relying on domain tags are hampered by wrong or unknown tags, for example for new domains or unfamiliar examples of existing domains. Among the compared multi-domain systems, they find performance generally underperforms simple fine-tuning for a single given domain, but that when domains are close or overlapping, multi-domain systems that share parameters across different domains are most effective.

7. Training Schemes for Adaptation

Once data is selected or generated for adaptation and a neural architecture is determined and pre-trained, the model can be adapted to the in-domain data. One straightforward approach described in Section 4 is fine-tuning the neural model with the same MLE objective function used in pre-training. However, as mentioned there, simple MLE fine-tuning can cause catastrophic forgetting of old domains and overfitting to new domains. In this section we discuss training schemes intended to mitigate these problems: regularized training, curriculum learning, instance weighting and non-MLE fine-tuning.

7.1 Objective Function Regularization

A straightforward way to mitigate forgetting is to minimize changes to the model parameters. Intuitively, if parameters stay close enough to their pre-trained values they will give similar performance on the pre-training domain. Here we discuss regularization and knowledge distillation techniques which may be applied during adaptation for this purpose.

7.1.1 FREEZING PARAMETERS

One way to ensure model parameters do not change is to simply not update them, effectively dropping them out of adaptation (Section 3.3.3). Research has shown that effective adaptation results can be obtained by varying just a small subset of the original model

parameters¹. For example, Thompson et al. (2018) choose subnetworks of recurrent NMT models to hold at their pre-trained values when fine-tuning on a new domain. Gu and Feng (2020) extend this work to the Transformer architecture. Both note a decrease in forgetting dependent on which subnetwork is frozen, although performance on the new domain is reduced relative to fine-tuning. Wuebker et al. (2018) likewise adapt only a manually defined subset of model parameters, encouraging sparsity in the adapted parameters with L1 regularization to improve efficiency. Deng et al. (2020) factorize all model components into shared and domain-specific, freezing the shared factors when tuning the domain-specific components on multiple domains. In these two cases performance on the new domains is shown to be similar to full fine-tuning, even though only a fraction of the parameters are adapted.

Parameters to be frozen do not necessarily have to be defined manually. Liang et al. (2021) identify low-magnitude parameters as sparse or underused areas of the network that can be easily adapted to new domains. They prune the sparse areas, then tune the original network on the generic domain so that it recovers, after which the sparse areas can be reintroduced and tuned on the new domain while the rest of the network is frozen. Gu et al. (2021) take a very similar approach, but recover the pruned model using knowledge distillation with in-domain data, removing the need for access to the generic dataset. In both cases the authors demonstrate adaptation to multiple domains while significantly reducing forgetting on the generic domain.

As for the approaches described in Section 6.3, these schemes have the advantage that only a small number of parameters need be saved for each adapted model. This is particularly useful for saving memory if ‘personalizing’ many models for fine-grained domains, such as individual users.

7.1.2 REGULARIZING PARAMETERS




Figure 4: Illustration of parameter regularization during fine-tuning from a pre-training (PT) domain to a new domain, based on Figure 1 from Kirkpatrick et al. (2017). If the pre-trained parameters θ^{PT} are adapted with no regularization (NR), good performance on the new domain corresponds to catastrophic forgetting on PT. Applying the same regularization to all parameters (L2) encourages minimal overall change from θ^{PT} . EWC regularization (Kirkpatrick et al., 2017) aims to vary the parameters that are unimportant for PT.

1. Different to adapters (Section 6.3), which introduce *new* parameters and freeze all original parameters.

If adapting all parameters, regularizing the amount of adaptation can mitigate forgetting. Barone et al. (2017) allow all NMT model parameters to vary under L2 regularization relative to their pre-trained values θ^{PT} , while Kirkpatrick et al. (2017) introduce the related approach of Elastic Weight Consolidation (EWC) for computer vision domain adaptation, which effectively scales the L2 regularization applied to each parameter θ_j by a value F_j . If Λ is a scalar weight, a general-form L2 regularized loss function is:

$$\hat{\theta} = \underset{\theta}{\operatorname{argmin}} [L_{CE}(\mathbf{x}, \mathbf{y}; \theta) + \Lambda \sum_j F_j (\theta_j - \theta_j^{PT})^2] \quad (9)$$

We illustrate these approaches in Figure 4. Relative domain importance can be controlled or tuned with Λ , which is larger if the old domain is more important and smaller if the new domain is more important. L2 regularization occurs where $F_j = 1$ for all j . For EWC Kirkpatrick et al. (2017) define F_j as the Fisher information for the pre-training domain estimated over a sample of data from that domain, $(\mathbf{x}^{PT}, \mathbf{y}^{PT})$.

$$F_j = \mathbb{E}[\nabla^2 L_{CE}(\mathbf{x}^{PT}, \mathbf{y}^{PT}; \theta_j^{PT})] \quad (10)$$

EWC has been broadly applied to NMT domain adaptation in recent years. Thompson et al. (2019) demonstrate that it mitigates catastrophic forgetting adapting from a generic to a specific domain. Saunders et al. (2019) demonstrate that this extends to sequential adaptation across multiple domains, and Stahlberg et al. (2019) show that EWC can improve translation on the generic domain when adapting between related domains.

Cao et al. (2021) observe that catastrophic forgetting corresponds to the model reducing the probability of words from previous domains by shrinking their corresponding parameters. They address this by normalizing target word parameters, an approach previously proposed by Nguyen and Chiang (2018) to improve rare word translation by mitigating the bias towards common words. While not quite the same as regularizing parameters via the objective function, the immediate goal of avoiding excessive changes in parameter magnitude during tuning is the same, as is the effect of reducing forgetting.

Parameter regularisation methods involve a higher computational load than parameter freezing, since all parameters in the adapted model are likely to be different from the baseline, and the baseline parameter values are often needed for regularisation. However, regularisation may also function more predictably across domains, language pairs and models than selecting parameters to freeze by subnetwork.

7.1.3 KNOWLEDGE DISTILLATION

Knowledge distillation and similar ‘teacher-student’ model compression schemes effectively use one teacher model to regularize training or tuning of a separate student model (Buciluă et al., 2006; Hinton et al., 2015). Typically the teacher model is a large, pre-trained model, and the student is required to emulate its behaviour with far fewer parameters. The student model is tuned so its *output distribution* remains similar to the teacher’s output distribution over the pre-trained data, often via an addition to the loss function (Kim & Rush, 2016). We note that ‘knowledge distillation’ is sometimes used in the literature to mean training a student model on forward translations produced by a teacher model – this approach is described in Section 5.3.2.

In a domain adaptation context, knowledge distillation encourages similar performance on the pre-training domain using a regularization function between general and in-domain model output distributions. The teacher is another NMT model. Dakwale and Monz (2017) focus on weighting the teacher distribution weighting to reduce catastrophic forgetting, while Khayrallah et al. (2018) aim for better in-domain performance on a single domain and Mghabbar and Ratnamogan (2020) use knowledge distillation to adapt to multiple domains simultaneously. Cao et al. (2021) extend knowledge distillation for domain adaptation to a continual learning scenario where the model is adapted sequentially multiple times. They propose dynamic knowledge distillation, in which older versions of the model carry less weight compared to newer versions.

We see knowledge distillation as similar in spirit to parameter regularization. However, it has a higher computational cost, since two models must actually operate on the data, instead of some parameter values simply being stored in memory. This can be effective when the aim is to compress the teacher model, since in this case the in-domain student model is likely to be much smaller than the other. For models remaining the same size, parameter regularization may be more practical.

7.2 Curriculum Learning




Figure 5: Curricula discussed in this section, moving from dataset A to dataset B. Typically, A consists of ‘easy’ examples and B of ‘difficult’ examples. In a domain adaptation context, A is the generic domain and B the target domain. a): Gradually mixing more of the target domain into the curriculum (Zhang et al., 2019a). b) Simple fine-tuning – pre-training on the generic domain for a long time, then switching to the target domain for a short time – as discussed in section 4. c) Mixed fine-tuning, where the target-domain portion of the curriculum incorporates a fixed proportion of the generic domain (Chu et al., 2017).

Bengio et al. (2009) recognize that humans learn best when concepts are presented in a meaningful order, or a ‘curriculum’. They hypothesize that neural model training can benefit from the same strategy of curriculum learning in terms of learning speed or quality. In broad terms, a curriculum ranks the training examples, and that ranking guides the order in which examples are presented to the model during training or fine-tuning.

Figure 5 illustrates some typical curricula used in domain adaptation for NMT. ‘Curriculum learning’ often refers to Figure 5a, where examples from a new dataset are gradually over-sampled during training. Often simple examples are shown first and difficult examples introduced later (Zhang et al., 2017; Weinshall et al., 2018). Difficulty can be determined in terms of objective data features like sentence length, linguistic complexity or word rarity

(Kocmi & Bojar, 2017; Platanios et al., 2019). Difficulty can also depend on model ‘competence’ for a given example at a certain state (Platanios et al., 2019). A training example may be considered difficult for an NMT model at a given point in training if the example’s embedding norm is large (Liu et al., 2020b), if the training loss of the example is changing significantly between training iterations (Wang et al., 2018), or if the model simply does not translate it well (Dou et al., 2020).

Although an easiest-first ranking is common, other rankings are possible. In fact, Zhang et al. (2018) find that both easy-first and hardest-first schedules can give similar speed improvements. We note that curricula involving translating all training examples during training or tracking competence across epochs may increase the computational cost to an extent that outweighs any benefits of better data ordering.

More relevant to this survey, a curriculum can be constructed from least-domain-relevant examples to most-domain-relevant. In fact, simple fine-tuning (Section 4, Figure 5b) is effectively curriculum learning where the final part of the curriculum contains only in-domain data. However, curriculum-based approaches to domain adaptation generally involve a gradual transition to in-domain data. For example, Wang et al. (2018) use an incremental denoising curriculum to fine-tune a pre-trained NMT model on increasingly clean data from its existing training corpus. Similar ‘cleaning’ fine-tuning curricula data can be learned via reinforcement learning methods (Kumar et al., 2019; Zhao et al., 2020).

A curriculum can be determined in terms of training example similarity with the target or generic domain. Moving smoothly between the generic and target domains may mitigate forgetting and overfitting. One way to move smoothly from the generic to the target domain is to extract target-domain-relevant data from generic data already seen by the model. Zhang et al. (2019a) identify such ‘pseudo in-domain’ data from general corpora and use similarity scores to rearrange training sample order. This effectively treats distance from in-domain data as a curriculum difficulty score. van der Wees et al. (2017) likewise perform domain adaptation while gradually emphasising more in-domain data from the generic dataset each training epoch.

The idea of a data sampling curriculum can be extended to multi-domain models. Sajjad et al. (2017) carry out ‘model stacking’, training on incrementally closer domains to the target domain, spending several epochs on each domain with the aim of maintaining good performance across multiple domains. Wang et al. (2020) use a difficulty-based domain adaptation curriculum across multiple domains simultaneously, using multi-dimensional domain features as difficulty scores. Wu et al. (2021) use model uncertainty on a small, trusted multi-domain dataset to determine a curriculum across corpora for a multi-domain model. Pham et al. (2022) determine multi-domain curricula with a reinforcement learning approach. They show that this outperforms the simple domain curriculum of Zhang et al. (2019a) for single- and two-domain adaptation, while also allowing n-domain adaptation.

A different perspective on on curricula for domain adaptation ‘reminds’ the model of the generic domain during adaptation, if generic data is available at adaptation time. Mixed fine-tuning, proposed by Chu et al. (2017) mixes domain-tagged generic data into the adaptation data for the target domain, effectively incorporating previously-seen data into the later half of a static curriculum (Figure 5c). They demonstrate strongly reduced catastrophic forgetting compared to simple fine-tuning, since the model is tuned on a more diverse dataset. They also show complementary behaviour to domain tagging (Section 6.1).

Song et al. (2020) extend mixed fine-tuning to a multi-stage multi-domain curriculum, sequentially learning new domains while discarding less relevant prior-domain data from the data mix, and show that this improves over single-stage mixed fine-tuning. Hasler et al. (2021) demonstrate that mixed fine-tuning without domain tags is complementary to directly regularizing parameters using EWC (Section 7.1.2). ‘Reminder’ techniques like mixed fine-tuning can be effective even if domains are not treated as distinct entities: Aljundi et al. (2019) demonstrate that maintaining a representative ‘replay buffer’ of past training examples avoids forgetting even without hard domain boundaries available.

7.3 Instance Weighting

Instance weighting adjusts the loss function to weight training examples according to their target domain relevance (Foster et al., 2010). For NMT, an instance weight $W_{x,y}$ for each source-target training example can easily be integrated into a cross-entropy loss function:

$$L(\mathbf{x}, \mathbf{y}; \theta) = \sum_{(\mathbf{x}, \mathbf{y})} -W_{\mathbf{x}, \mathbf{y}} \log P(\mathbf{y}|\mathbf{x}; \theta) \quad (11)$$

A higher weight indicates that a sentence pair is important for training towards the target domain, while a low (or zero) weight will lead to sentences effectively being ignored during tuning. While this involves adjusting the loss function rather than data reordering, the effect is to simulate over- or under-sampling particular examples during adaptation.

The weight may be determined in various ways. It may be the same for all sentences marked as from a given domain, or defined for each sentence using a domain measure like n-gram similarity (Joty et al., 2015) or cross-entropy difference (Wang et al., 2017b). If changes can be made to the model architecture, the instance weight may be determined by a domain classifier (Chen et al., 2017), or an architecture-dependent approach like sentence embedding similarity (Zhang & Xiong, 2018). The same effect is achieved by Farajian et al. (2018) by adjusting the learning rate and number of adaptation epochs dedicated to tiny adaptation sets based on their domain similarity. In each case the outcome is to place more or less emphasis on particular adaptation examples during training.

We view instance weighting as fundamentally the same idea as curriculum learning (Section 7.2). Both schemes bias the model to place more importance on certain training examples, allowing some control over how the model fits, forgets, or overfits certain sentences. Some forms of curriculum learning are implemented in a similar way to instance weighting, with a higher weight applied to examples that fall into the current section of the curriculum, or a zero weight applied to examples that should not yet be shown to the model (Bengio et al., 2009; Dou et al., 2020). One difference is that instance weights for domain adaptation do not usually change as training progresses or model competence changes, but bias the model towards in-domain data in a constant manner.

7.4 Non-MLE Training

As discussed in Section 3.3, MLE training is particularly susceptible to exposure bias, since it tunes for high likelihood only on the sentences available in the training or adaptation corpus. MLE also experiences loss-evaluation metric mismatch, since it optimizes the log

likelihood of training data while machine translation is usually evaluated with translation-specific metrics. Tuning an NMT system with a loss function other than the simple MLE on pairs of training sentences may therefore improve domain adaptation. Here we describe two other training objectives that have been demonstrated as useful precursors or follow-ups to domain-adaptive MLE fine-tuning: minimum risk training and meta-learning.

7.4.1 MINIMUM RISK TRAINING

Discriminative training for MT was introduced for phrase-based machine translation, minimizing the expected cost of model hypotheses with respect to an evaluation metric like document-level BLEU (Papineni et al., 2002). Shen et al. (2016) extend these ideas to Minimum Risk Training (MRT) for NMT, using expected minimum risk at the sequence level with a sentence-level BLEU (sBLEU) cost for end-to-end NMT training. Given N sampled target sequences \mathbf{y}_n and the corresponding reference sequences \mathbf{y}^* , for each sentence pair in a minibatch the MRT objective is:

$$\hat{\theta} = \operatorname{argmin}_{\theta} \sum_{n=1}^N \Delta(\mathbf{y}_n, \mathbf{y}^*) \frac{P(\mathbf{y}_n | \mathbf{x}; \theta)^\alpha}{\sum_{n'=1}^N P(\mathbf{y}_{n'} | \mathbf{x}; \theta)^\alpha} \quad (12)$$

Hyperparameter α controls the smoothness of the sample probability distribution. Function $\Delta(\cdot)$ measures hypothesis cost $\in [0, 1]$, often related to sentence BLEU. MRT has been applied to NMT in various forms since its introduction. Edunov et al. (2018) explore variations on MRT, using samples produced by beam search and an sBLEU-based cost. Wieting et al. (2019) use MRT with sBLEU and sentence similarity metrics. Saunders et al. (2020) use MRT for translation with BLEU calculated over minibatch-level ‘documents’.

MRT is of particular relevance to domain adaptation for two reasons. Firstly, in the NMT literature we find that MRT is exclusively applied to fine-tune a model that has already converged under a maximum likelihood objective. MRT therefore fits naturally into a discussion of improvements to pre-trained NMT models via parameter adaptation.

Secondly, there is some indication that MRT may be effective at reducing the effects of exposure bias. Exposure bias can be a particular difficulty where there is a risk of overfitting a small dataset, which is often the case for domain adaptation, especially if there is a domain mismatch between adaptation and test data (Müller et al., 2020). However, MRT optimizes with respect to model samples rather than tuning set target sentences – the latter may be less diverse and thus easier to overfit. Moreover, a high-quality NMT baseline may produce synthetic samples that are better aligned than available natural data. The ability of MRT to mitigate overfitting is highlighted by Neubig (2016), who notes that MRT tends to produce sentences of the correct length without needing length penalty decoding. When applied to adaptation, MRT has been shown to mitigate exposure bias when the test data domain is very different from the training domain (Wang & Sennrich, 2020), and when adaptation data is in-domain but noisy, poorly aligned or unreliable (Saunders & Byrne, 2020a).

7.4.2 META-LEARNING

Meta-learning was proposed by Finn et al. (2017) as a means of tuning a neural model such that it can easily learn new tasks in few additional steps: ‘learning to learn’. The aim is to find model parameters that can be adapted to a range of domains easily, in few training

steps and with few training examples. This is achieved by introducing a meta-learning phase after pre-training and before fine-tuning on any particular target domain.

During the meta-learning phase the model f_θ sees a number of few-shot adaptation tasks T , where each task consists of training and test examples. In a meta-training step, model parameters θ are first temporarily updated using the training examples, then the updated model’s loss L_T on the task test examples is used as the real meta-training objective. Effectively, test losses for the meta-training tasks are used as the model’s training loss during meta-training. For a given learning rate α :

$$\operatorname{argmin}_\theta \sum_T L_T(f_{\hat{\theta}}) = \sum_T L_T(f_{\theta - \alpha \nabla_\theta L_T(f_\theta)}) \quad (13)$$

Sharaf et al. (2020) apply meta-learning to NMT domain adaptation. They specifically meta-tune the parameters of a single adapter layer (Section 6.3) and simulate small domains for meta-training by sampling sets of just hundreds of parallel sentences. They find that this outperforms simply fine-tuning on the same samples when the number of simulated domains is high, but not when fewer, larger meta-training domains are sampled, potentially due to overfitting the meta-training sets. They also note that meta-training all Transformer parameters instead of an adapter leads to catastrophic forgetting of the original domain. Song et al. (2021) also meta-learn adapter layers, but adjust meta-learning learning rate α dynamically to make the process sensitive to domain differences. They focus on domain differences in terms of model confidence in modelling a particular domain, and how representative each sentence pair is of a particular domain. They find dynamic learning rate adjustment particularly improves meta-learning performance when target domains of interest are very different sizes. Park et al. (2021) apply meta-learning to unsupervised domain adaptation with only monolingual data. They do so by incorporating back-translation and language modelling objectives into the meta-learning phase, and demonstrate strong improvements over both simple fine-tuning and mixed fine-tuning.

Li et al. (2020) meta-learn the encoder and vocabulary embedding spaces, and find this outperforms fine-tuning except for domains with many infrequent words. Zhan et al. (2021) effectively apply a curriculum (Section 7.2) to the meta-learning process. They initiate the meta-learning phase by sampling similar meta-training domains, and gradually introduce more divergent domains. They find small improvements for the domains seen during meta-learning. These two works together suggest that meta-learning may out-perform fine-tuning for domains where test data is well-represented by the tuning data.

8. Inference Schemes for Adaptation

One way to side-step the problem of translating multiple domains in a single model is to simply assign a separate NMT model to each domain and combine them at inference time. Such models can be obtained using the techniques discussed in the previous sections, for example by fine-tuning a single pre-trained model on data from each domain of interest. While this approach is simple, if not memory-efficient, it begs the question of how best to perform translation on an unseen source sentence from an unknown domain. Possible approaches are multi-domain ensembling, and reranking or rescoreing an existing set of translation hypotheses.

We also describe ways to encourage a model to produce domain-specific terminology or phrasing via a pre- or post-processing step at inference time. This approach bypasses even domain-specific tuning of an NMT model.

8.1 Multi-Domain Ensembling

At inference time an NMT ensemble can use predictions from multiple models to produce a translation in a single pass, as described in Section 3.4. Here we discuss two forms of domain-specific ensembling. Domain adaptive ensembling seeks to find a good set of interpolation weights for a traditional ensemble of NMT models when translating a sentence of unknown domain. Retrieval-based ensembling is a recently-proposed technique that interpolates NMT predictions with a distribution over tokens that are likely in context, extracted from a static in-domain datastore.

8.1.1 DOMAIN ADAPTIVE ENSEMBLING

Certain models in an ensemble may be more useful than others for certain inputs. For example, if ensembling a travel-domain model with a science-domain model, we might expect the science model to be more useful for translating medical abstracts. This idea of varying utility across an ensemble is particularly relevant when the domain of a test sentence is unknown and therefore the best ensemble weighting must be determined at inference time. We make an important distinction between the domain of the training data and model – typically known, often with an available development set – and the test data, for which we do not always know the domain or have access to a development set.

With a source sentence of unknown domain, Freitag and Al-Onaizan (2016) demonstrate that reasonable performance can be achieved via a uniform ensemble of general models and in-domain translation models. For models over K domains:

$$\hat{\mathbf{y}} = \operatorname{argmax}_{\mathbf{y}} \sum_{k=1}^K \frac{1}{K} P_k(\mathbf{y}|\mathbf{x}) \quad (14)$$

However, intuition would suggest that the in-domain model should be prioritised when inputs are closer to in-domain data, and the general model if inputs are more similar to generic training data. Uniform ensembling is straightforward, but does not emphasise performance on any particular domain.

Sajjad et al. (2017) also use multi-domain ensembles, but determine an ensemble weight W_k for each model and test set t , tuned on development sets:

$$\hat{\mathbf{y}} = \operatorname{argmax}_{\mathbf{y}} \sum_{k=1}^K W_k(t) P_k(\mathbf{y}|\mathbf{x}) \quad (15)$$

This static weighting approach can be targeted to a particular domain, but tuning can be slow, and requires either that the domain of the test set is known or that a development set representative of the test set is available.

Domain adaptive approaches to inference can instead determine the ensemble weights conditioned only on the source sentence or the partial translation hypothesis, without the need to tune weights manually. For SMT, Huck et al. (2015) use a language model to classify

the domain of an unseen source sentence x when determining which set of parameters to use when generating the translation hypothesis. Liu et al. (2020a) use text similarity between x and the in-domain corpus to determine an ensemble interpolation weight between generic and in-domain NMT models.

$$\hat{\mathbf{y}} = \operatorname{argmax}_{\mathbf{y}} \sum_{k=1}^K W_k(x) P_k(\mathbf{y}|\mathbf{x}) \quad (16)$$

An alternative approach, Bayesian Interpolation, is introduced by Allauzen and Riley (2011) for language model ensembling. Bayesian Interpolation allows adaptive weighting without necessarily specifying that the domain of the test sentence is fixed for the entire sentence. A set of initial approximated ensemble weights $W_k(t)$ defines a task-conditional ensemble as in Equation 15. This becomes a fixed weight ensemble if t is known for a test set or sentence, much like the approaches described in Huck et al. (2015) and Sajjad et al. (2017). However if t is not known, ensemble weights can be estimated at each inference step i , where h_i is history $\mathbf{y}_{1:i-1}$:

$$p(y_i|h_i) = \sum_{t=1}^T p(t, y_i|h_i) = \sum_{k=1}^K p_k(y_i|h_i) \sum_{t=1}^T p(t|h_i) W_k(t)$$

That is, a weighted ensemble with state-dependent mixture weights computable from priors $W_k(t)$ priors and the updated language model posterior. Saunders et al. (2019) extend this formalism to include conditioning on a source sentence and estimating $W_k(t)$ from the source sentence, $W_k(x)$. This permits domain adaptive NMT with multi-domain ensembles when the test sentence domain is unknown and may vary within the sentence. If a new, unseen domain is present, adaptive weighting will favour models for which that new domain is likely. Saunders et al. (2019) demonstrate that adaptive ensemble weighting is most beneficial when ensembling models specializing in different domains. An ensemble of models specializing in similar domains may indeed perform well with a uniform weighting.

8.1.2 RETRIEVAL-BASED ENSEMBLING

An ensemble does not need to consist purely of NMT models, or even involve an additional neural network. Khandelwal et al. (2021) propose k -nearest-neighbour (k NN) machine translation using a domain-specific data-store. The data-store maps NMT decoder states to in-domain target language tokens. At each inference step the NMT model’s predictions are interpolated with a distribution over the k retrieved nearest neighbour tokens given the current model state. This is effectively an ensemble with a distribution from a static context-to-token map.

Zheng et al. (2021a) demonstrate that this technique can benefit from determining k itself adaptively for a given context: retrieving many tokens from rare contexts is likely to result in noisy retrieved neighbours. Zheng et al. (2021b) further show that the k NN data-store itself can be adapted for a given domain in the absence of parallel data using copying and backtranslation (Section 5.3) to create pseudo-parallel data, and adapters (Section 6.3) to learn states from the pseudo-parallel data. Martins et al. (2022) improve efficiency of domain-specific k NN by caching retrieved translations for a given domain. All these variants

on k NN have been shown to improve in-domain translation significantly without requiring an additional NMT model.

8.2 Constrained Inference and Rescoring

Ensembling uses multiple models to produce a translation simultaneously. An alternative is a multi-pass approach to inference, in which one model produces an initial translation which is adjusted or corrected using another model. While this approach involves multiple models performing their own separate inference passes, it can be more efficient than ensembling. Multi-pass approaches do not involve holding multiple models in memory at once, and the second translation pass is commonly held close to the initial translation in some way, reducing the number of translations to be scored.

The initial model may produce multiple translations – for example, the highest scoring N translations following beam search (Section 3.4.1). In this case a second inference pass can rescore this N-best list using a different model or loss function. For example, Minimum Bayes Risk (MBR) decoding rescoring N-best lists or lattices to improve single system performance for SMT (Kumar & Byrne, 2004; Tromble et al., 2008; de Gispert et al., 2010), or for NMT if a sufficiently diverse lattice can be defined, for example, from SMT n-gram posteriors (Stahlberg et al., 2017). In the context of domain adaptation, Dou et al. (2019b) propose rescoring generic NMT model translations with domain-specific language models. Zhang et al. (2018) retrieve in-domain source sentences that are similar to test inputs, and up-score initial NMT hypotheses which contain fragments of their translations.

A related idea which has been applied to domain-specific translation is constrained inference, where the NMT model is only permitted to produce certain translations. The constraints may be generated by another translation model. For example, Stahlberg et al. (2016) generate translation hypotheses with an SMT model, which are then represented as a lattice that constrains NMT decoding. Khayrallah et al. (2017) constrain NMT output to a domain-specific lattice to improve adequacy in domain adaptation scenarios. The constraints may allow only certain domain-relevant words to change: Saunders and Byrne (2020b) use lattice constraints when rescoring translations with an adapted model to correct mistakes in a gender ‘domain’. Constrained decoding can mitigate forgetting under a domain-adapted model, as the adapted model only changes the original translation in pre-determined permitted ways.

The initial beam search may itself be constrained to domain-specific language. Hokamp and Liu (2017) and Hasler et al. (2018) both adjust the beam search algorithm to force domain-specific terminology into the output of a generic model, although both approaches are slower than straightforward beam search. An extension of the Hokamp and Liu (2017) approach by Xie and Way (2020) finds that using alignment information can significantly improve the efficiency of constrained beam search, with a similar BLEU score.

We note that the initial hypotheses used for rescoring or with constraints do not have to be generated by a domain-specific model, or even a model which the user can tune. For example, a user might generate initial translations with a generic commercial tool, then rescore or retranslate the translations with constraints relating to the desired output, as in Saunders and Byrne (2020b). It may be more practical for a user to edit or repair small

aspects of generic translations from commercial systems in domain-specific ways, manually or with their own monolingual tool, than to access a custom NMT system.

8.3 Domain Terminology Via Pre- and Post-Processing

A similar idea to domain terminology ‘correction’ during inference treats the presence of domain-specific terminology as a pre- and post-processing problem. We discuss two such approaches here. The first incorporates domain-specific terminology into the source sentence in pre- and post-processing steps so as to ensure they are correctly translated. The second preprocesses the input with similar source sentences or translations as a cue for the model.

8.3.1 TERMINOLOGY TAGGING

There are many possible ways to integrate terminology into the system at inference time. Dinu et al. (2019) add inline tags to terminology in the source sentence as a pre-processing step, and then replace the translated tags with corresponding target language terms. Song et al. (2019) simply replace certain source phrases with pre-specified domain-specific target translations as a preprocessing step, encouraging copy behaviour. Michon et al. (2020) compare several variations on inline terminology tags, and find that the approach of Dinu et al. (2019) may give the best scores where terminology is already reasonably well-handled, but that it is less effective at introducing poorly-handled terminology. By contrast with the above approaches Dougal and Lonsdale (2020) inject terminology after inference as a post-processing step using source-target alignments. This has the advantage of not requiring a model to handle tags, and could in principle be used to introduce terminology to a generic or commercial system output. It is, however, reliant on an effective alignment model.

Chen et al. (2020a) take a similar approach to Dinu et al. (2019) and Song et al. (2019), but do not require alignments and only require bilingual dictionaries during inference - they specify the reference terminology in a fixed source position, encouraging the model to learn correct alignments. Niehues (2021) also specify the reference terminology in a fixed source position, but additionally use only the lemma, encouraging the model to learn correct inflections for provided terminology. Lee et al. (2021) add an objective to predict the span of masked source terms during NMT training, allowing them to insert multi-word domain-specific terms during inference. They find similar unigram performance to Chen et al. (2020a), but better reproduction of longer terms.

We note that many of these approaches do not actually require the baseline model to see or produce the domain-specific terminology. This means terminology injection can take place even for new, unseen domains whose terminology which would not otherwise be generated by the model (Michon et al., 2020).

8.3.2 IN-DOMAIN PRIMING

An extension to priming the model with individual terminology incorporates entire related sentences. For a given sentence, similar source sentences and their translations can be extracted from a parallel dataset using fuzzy matching as described in Section 5.1. The similar sentence or its translation can be used as a domain-specific input prompt.

Bulte and Tezcan (2019) propose priming with Neural Fuzzy Repair (NFR) based on token edit distance: they incorporate known translations of similar sentences into the model

input to act as context cues at inference. Xu et al. (2020) perform sentence priming experiments with a wider range of similarity metrics, including continuous representations. They distinguish between related and unrelated target words, bringing their work closer to terminology tagging approaches of Section 8.3.1. Tezcan et al. (2021) develop NFR towards use of subword representations for identifying matches, and also mark related and unrelated target terms. By contrast to these approaches Pham et al. (2020b) ‘prime’ using both source and target sides of related sentence pairs, forcing the model to produce an in-domain target-language ‘cue’ before translating the sentence of interest. Interestingly, Moryossef et al. (2019) successfully prime commercial MT systems with short phrases relating to gender or number, suggesting custom content injection may even be possible in some black-box scenarios.

Pre- and post-processing approaches may work best if the NMT model is trained from scratch such that it ‘sees’ examples of generic terminology tags or priming cues during training. However, importantly, this training is not domain-specific, and the model does not need to be *retrained* for performance on a particular domain. At inference time, these approaches effectively adapt a translation’s domain-specificity without the need to retrain or further change the model itself.

9. Case Studies for (Multi-)Domain Adaptation

We conclude this survey by briefly describing three ongoing lines of NMT research. At first glance they may seem unrelated to the domain adaptation scenarios described above. However, all have seen successful application of domain adaptation techniques in recent years. In each case we briefly describe the problem, and assess whether it meets the criteria for a single-domain or multi-domain adaptation problem using the points summarized in Section 2. We then summarize recent approaches to each problem, focusing on those that use a domain adaptation framing. Some of these approaches have been mentioned elsewhere in this survey: here we put them in the context of their broader line of research. We do not intend this to be a comprehensive survey of work on each problem, and refer to prior surveys where available in each case. Instead we include this section as a demonstration of the power and wide applicability of domain adaptation techniques, as well as guidance for applying the (multi-)domain adaptation framing to new problems.

In this section we demonstrate the general utility of techniques in this survey for notions other than changing text topic or genre. The ‘domains’ referred to in this section are not necessarily implicit given the source sentence in the same way that topic or genre might be. However, we note that even traditional domain adaptation often handles sentences where domain is not implicit. Indeed, identical sentences may be translated in multiple different ways depending on target domain, whether due to lexical ambiguity, customer-specific terminology or genre-appropriate formality. Any of these could be resolved by controlling the target domain. Likewise, we may wish to control an output language, entity gender or document ‘domain’ despite not being able to infer it from the source.

9.1 Low Resource Language Translation as a Domain Adaptation Problem

Neural machine translation systems have achieved impressive performance on language pairs with millions of well-aligned bilingual training sentence pairs available. For many language

pairs, such resources are not available. A survey on natural language processing in such low resource cases is given in Hedderich et al. (2020).

Faced with machine translation between a low resource language pair, we may wish to improve the available resources, for example by generating synthetic bilingual data. Alternatively we may decide to minimize the need for additional resources, for example by adapting an existing NMT model trained to translate from the same source language. These approaches are closely related to those described for domain adaptation in this article: in effect, they treat the low resource language pair as its own domain.

A related approach to low resource language translation is multilingual translation. In multilingual NMT, the goal is to translate to or from more than one language with a single model (Dong et al., 2015; Aharoni et al., 2019). Multilingual NMT is a broad field of research which does not necessarily involve low resource translation. However we are particularly interested in the application of multilingual NMT to low resource languages (Mueller et al., 2020). This can be seen as analogous to a multi-domain adaptation scenario: we wish to achieve good performance on multiple distinct sets of text without compromising translation on any, and ideally with the distinct ‘domains’ benefiting each other.

9.1.1 CAN LOW RESOURCE LANGUAGE TRANSLATION BE TREATED AS A DOMAIN ADAPTATION PROBLEM?

Low resource language translation can be seen as a single-domain adaptation problem in most cases, although for multilingual NMT systems it may become a multi-domain adaptation problem:

- We wish to improve translation performance on sentences from a low-resource language pair. These should be very identifiable.
- We wish to avoid retraining, since we know that training from scratch on the low-resource language pair alone is unlikely to be effective.
- Training a system to translate between only a single pair of languages will not be a multi-‘domain’ problem. However, if we are training a one- or many-to-many multilingual NMT model, we may want to adapt so as to translate into several languages simultaneously. This may mean avoiding forgetting previous language pairs, learning multiple languages simultaneously, or learning new languages in a few-shot manner.

9.1.2 IMPROVING LOW RESOURCE LANGUAGE TRANSLATION BY ADAPTATION

Early and straightforward data-centric approaches to low resource language translation, like domain adaptation, center around tuning an existing model on data for the new language pair, usually sharing either the target language (Zoph et al., 2016) or source language (Kocmi & Bojar, 2018). A distinction is that in this case catastrophic forgetting of the original model’s abilities is less of a concern, since it is likely that a user can pre-determine which languages will be translated with which model. However, many techniques are applicable to both single-language-pair domain adaptation and cross-lingual transfer learning.

If the architecture is to be kept the same, lack of vocabulary overlap can be a difficulty, although less so for related languages or those that can leverage similar BPE vocabularies

(Nguyen & Chiang, 2017). Alternatively, both the data and architecture can change: for example, Kim et al. (2019a) explore use of adding or sharing network components for transfer learning new languages.

Monolingual data is often more freely available than bilingual data for low resource translation, as for small-domain translation. Data-centric approaches to low-resource language translation may therefore adapt to semi-synthetic datasets created by forward- or back-translation (Rubino et al., 2020; Karakanta et al., 2018). Even if this pseudo-parallel data is generated using a relatively weak low resource translation system, it still may be beneficial for further tuning that system (Currey & Heafield, 2019). An alternative option is to generate purely synthetic data: for example Fadaee et al. (2017) target rare words for low resource NMT by situating them in new, synthetic contexts.

If tuning a previously-trained model on a new language pair, it is not usually necessary to maintain performance on the original language pair. However, if e.g. the original source language is shared with the new language pair, it may be better to avoid overfitting the original encoder on the new dataset. To this end, previous work has experimented with freezing some parameters (Ji et al., 2020), or jointly training on the low and high resource language for ‘similar language regularization’ (Neubig & Hu, 2018). It may also be possible to leverage monolingual data from the low resource language during tuning, for example to regularize adaptation (Baziotis et al., 2020). Such approaches are clearly reminiscent of domain adaptive tuning techniques described in Section 7.

A separate approach develops multilingual NMT models, which may translate between very many language pairs. As for single-low-resource-language adaptation, most ‘domain’-adjacent techniques used here are naturally data-centric. For example, if the language set includes both high and low resource language pairs, these may act analogously to high and low resource domains, with lower resource language pairs benefiting from the higher resource lexical representations (Gu et al., 2018). Multilingual models may also benefit from pseudo-parallel datasets, even for language pairs with no parallel data (Firat et al., 2016). New language pairs may be introduced at different points during training, with language or language pair tags, or subsets of the model parameters reserved for individual language pairs (Blackwood et al., 2018; Zhang et al., 2020). Parameter regularization can avoid forgetting previously learned language pairs. (Carrión-Ponz & Casacuberta, 2022). In terms of architecture-centric approaches, the adapter scheme described in Section 6.3 was proposed simultaneously for multi-domain adaptation and multilingual NMT (Bapna & Firat, 2019b; Philip et al., 2020), highlighting the conceptual connection between these lines of research.

9.2 Gender Bias in Machine Translation as a Domain Adaptation Problem

Translation into languages with grammatical gender involves correctly inferring the grammatical gender of all entities in a sentence. In some languages this grammatical gender is dependent on the social gender of human referents. For example, in German, translation of the entity ‘the doctor’ would be feminine for a female doctor – *Die Ärztin* – or masculine for a male doctor – *Der Arzt*. In practice, however, many NMT models struggle to generate such inflections correctly (Prates et al., 2020). Stanovsky et al. (2019) explore these mistakes and demonstrate that they tend to reflect social gender bias: machine translation

tends to translate based on profession-based gender stereotypes instead of using the context of e.g. referential pronouns. A full survey of this behaviour in NMT can be found in Savoldi et al. (2021).

English source	The doctor helps the patient, although she is busy
German reference	Die Ärztin hilft dem Patienten, obwohl sie beschäftigt ist
MT with a bias-related mistake	<i>The [female] doctor helps the [male] patient, although she is busy</i> Der Arzt hilft der Patientin , obwohl sie beschäftigt ist <i>The [male] doctor helps the [female] patient, although she is busy</i>
English source	The nurse helps the patient, although he is busy
German reference	Der Krankenpfleger hilft dem Patienten, obwohl er beschäftigt ist
MT with a bias-related mistake	<i>The [male] nurse helps the [male] patient, although he is busy</i> Die Krankenschwester hilft dem Patienten , obwohl er beschäftigt ist <i>The [female] nurse helps the [male] patient, although he is busy</i>

Table 1: Examples of mistranslation relating to gender bias effects. Bolded words are entities inflected to correspond to the pronoun. MT from Microsoft Translator 05/22.

Gender-based errors are particularly common when translating sentences involving coreference resolution. Table 1 gives two typical examples. These effects may occur because the systems are influenced by the higher frequency of masculine-inflected doctors and feminine-inflected nurses in training data, resulting from historical or cultural imbalances in the society that produces this data.

Such effects are commonly referred to as systems exhibiting gender bias. Interestingly, the effects can be interpreted in much the same way as a traditional text domain: certain gendered terms have become associated with masculine or feminine grammar in the target language. The model effectively derives a ‘gender’ domain implicitly from even gender-ambiguous source sentences (e.g. ‘the doctor helps the patient’), even though this is not desirable. Mitigation techniques typically revolve around controlling this ‘domain’ more sensibly.

9.2.1 CAN GENDER BIAS BE TREATED AS A DOMAIN ADAPTATION PROBLEM?

The problem of mitigating the effects of gender bias in NMT can be cast as a multi-domain adaptation problem:

- We wish to improve translation for sentences with a distinct vocabulary distribution, which can be interpreted as a domain: sentences containing gendered terms which do not match existing social biases, such as female doctors and male nurses.
- We wish to avoid retraining: biases may be reinforced by the overall training set.
- We want to translate at least two domains. One consists of sentences affected by bias, the other of sentences that are unaffected – we do not wish to compromise translation of sentences without gendered terms. We may also want to continually adapt to new ‘domains’ to mitigate newly identified biases, since we are unlikely to successfully pre-determine all relevant biases and their sources.

9.2.2 MITIGATING THE EFFECTS OF GENDER BIAS BY ADAPTATION

While there has been interest in mitigating undesirable data bias prior to training, this may not be practical in terms of computation or in terms of ability to adjust the data (Tomalin et al., 2021). Such approaches may also not be conceptually sensible, as there are countless ways in which biases could conceivably manifest in generated natural language, relating to gender or otherwise (Hovy et al., 2020; Shah et al., 2020), so speaking in terms of simple biases or imbalances that can be addressed is not clearly meaningful. Instead, a range of approaches benefit from treating gender handling as a domain.

Improving gender translation by adapting to new data has been demonstrated in several ways, although as with domain adaptation these approaches generally involve some mitigation of forgetting and/or overfitting during training or inference. Jwalapuram et al. (2020) adapt, using a discriminative loss function to mitigate overfitting, to sentences containing pronouns that were previously mistranslated. Saunders and Byrne (2020b) adapt to semi-synthetic forward-translated gendered sentences and fully synthetic template-generated sentences, exploring adaptation with EWC and constrained retranslation to avoid forgetting. Costa-jussà and de Jorge (2020) fine-tune on natural gender-balanced data, and mitigate forgetting by mixing in a proportion of general-domain data.

The idea of controlling machine translation gender inflections with some form of tag, reminiscent of domain tagging, has been proposed in several forms. Vanmassenhove et al. (2018) incorporate a ‘speaker gender’ tag into all training data, allowing gender to be conveyed at the sentence level. More relevant to domain adaptation, explicit gender tags have been introduced during model adaptation (Saunders et al., 2020). Interestingly, implicit gender tags in the form of gendered ‘prompts’ such as prepending the phrase ‘he said’ or ‘she said’ have been demonstrated by Moryossef et al. (2019) to control translation gender at inference time, reminiscent of ‘priming’ techniques described in Section 8.3. This approach effectively provides gender domain tags without requiring any changes to the model.

9.3 Document-Level Translation as a Domain Adaptation Problem

Document-level machine translation, summarized in Section 3.1.4, has come to refer to two connected ideas. The first is translating one specific document, accounting for its terminology or writing style. The second is translation of documents in general, accounting for context beyond the individual sentence - for example to resolve pronoun coreference when handling gender, as in Section 9.2. Document-specific terminology can benefit from extra-sentential context, as shown in Table 2, so the ideas can be considered linked, and are often explored together in the literature – a thorough survey of which can be found in Maruf et al. (2021). Translating a specific document can clearly be cast as an adaptation problem. We find that incorporating document context into an existing model often also involves adaptation-adjacent approaches, since document-aligned bilingual datasets are usually much smaller than sentence-aligned bilingual datasets (Voita et al., 2019b).

9.3.1 CAN DOCUMENT NMT BE TREATED AS A DOMAIN ADAPTATION PROBLEM?

Improving document-level translation can be seen as a domain adaptation problem, and may be seen as a multi-domain adaptation problem depending on efficiency requirements:

English	[X] is a portrait photographer ₁₂ . She ₁ is known for shooting ₂ in the woods.
German translations of individual sentences	[X] ist Porträtfotograf . Sie ist bekannt dafür, im Wald zu schießen . <i>X is a [male] portrait photographer. She is known for shooting [a weapon] in the woods.</i>
German translations with context	[X] ist Porträtfotografin . Sie ist bekannt dafür, im Wald zu fotografieren . <i>X is a [female] portrait photographer. She is known for taking photographs in the woods.</i>

Table 2: Example of English-to-German translation improved by document context. Two terms with their cross-sentence resolution are numbered in the English sentence: 1 involves an anaphoric pronoun and 2 is a case of lexical ambiguity. Ambiguous German translations (bolded) are resolvable given document context.

- We wish to improve translation performance on (at least) one document, which may use specific terminology or a distinct writing style.
- We wish to avoid retraining from scratch: an individual document or set of document-aligned data is very small compared to typical sentence-aligned generic training sets, so retraining would be especially inefficient.
- If we are not willing to tune one model per document, we may wish to use the same model to translate well across multiple documents. This would be a multi-domain adaptation problem for the sake of efficiency. However, it is possible that only one document is of interest at a time.

9.3.2 IMPROVING DOCUMENT-LEVEL TRANSLATION BY ADAPTATION

Many of the data-centric adaptation approaches described in Section 5 use some ‘seed’ text to extract related bilingual or monolingual data. These can be used for document-level translation with the document to be translated as a seed. An example of this approach is Kothur et al. (2018), who adapt to a lexicon containing novel words in a test document.

Attempts to use extra-sentential context in document translation have primarily incorporated additional sentences into the model input as described in Section 3.1.4. However, an alternative approach reminiscent of domain adaptation is to condense document information into tag or label form (6.1). Jehl and Riezler (2018) use inline document-content tags for patent translation, integrating information at the word and sentence-level. Stojanovski and Fraser (2020) show that incorporating document context can improve domain-specific translation, suggesting that domain tags and document tags may function similarly. Kim et al. (2019b) demonstrate that contextual and lexical information can be incorporated into a very minimal form, reminiscent of tagging approaches, rather than necessitating encoding entire additional sentences or documents.

The neural network architecture may need to be adjusted in order to incorporate context beyond the individual sentence. A typical change is simply adding additional context source

sentences alongside additional encoder parameters. Stojanovski and Fraser (2019) take this approach, initialising the new context parameters randomly in a pre-trained model before tuning on smaller document-level datasets. Voita et al. (2019b) initialise and tune a context-aware decoder on a relatively small amount of document-sensitive data. Ul Haq et al. (2020) first train a generic NMT model, then add and adapt context-sensitive hierarchical attention networks on document-specific data. A final set of domain-adaptation-adjacent approaches to document-level translation incorporates document information at inference time (Section 8): Voita et al. (2019a) carry out monolingual document repair with a specialised system trained on ‘in-domain’ – context-aware – data. Stahlberg et al. (2019) similarly rerank translations with a language model that incorporates document context.

10. Conclusions and Future Directions

Domain adaptation lets NMT models achieve good performance on language of interest with limited training data, and without the cost of retraining the model from scratch. Adaptation may even allow better performance than from-scratch training on a given domain. Some ongoing challenges for machine translation can be framed as domain adaptation problems, inviting the application of the well-tested techniques reviewed in the earlier sections of this article. As our case studies in Section 9 have shown, there is already a trend towards using adaptation for pseudo-domains other than straightforward provenance, topic or genre. It remains to be seen which other lines of NMT research might benefit from domain adaptation techniques.

We conclude this survey with a view to the future, highlighting four areas that are of particular interest for future work on NMT domain adaptation: extremely fine-grained adaptation, unsupervised adaptation, efficiency, and intentional forgetting. We have touched on these at various points in this survey; here we emphasise their relevance to future research.

As observed in Section 2.1, text domains can be difficult to interpret, with topic and style not necessarily consistent even within one document. Fine-grained adaptation of models to very small amounts of data is one way to account for this domain fuzziness and variability (Michel & Neubig, 2018). It has also become a popular way for industrial MT providers to differentiate their products or target them to specific clients (Savenkov, 2018). Models may be tuned to translate text by a particular author (Michel & Neubig, 2018; Buj et al., 2020), or for every sentence (Farajian et al., 2017; Li et al., 2018; Mueller & Lal, 2019). While fine-grained adaptation may improve lexical choice, it is often slow, and risks overfitting to irrelevant data (Li et al., 2018). Future work in fine-grained or ‘extreme’ adaptation may therefore focus on extraction of relevant data (Section 5) or ways to adapt at a per-sentence level without changing the model at all (Section 8.3).

Unsupervised domain adaptation, which tunes a system without access to parallel data, is also growing in popularity (Ramponi & Plank, 2020). A low resource domain – especially a very fine-grained domain – may not have parallel corpora for a particular language pair. However, in-domain *monolingual* data is more widely available, and can let us produce pseudo-parallel data by forward or back translation (Section 5.3), allowing adaptation without genuine ‘supervised’ parallel data. Open questions involve the best ways of extracting large amounts of in-domain data from multiple corpora given small amounts of in-domain data, or of producing pseudo-parallel in-domain data from monolingual in-domain data.

With interest in adapting to ever-finer domains, sufficient in-domain parallel data is increasingly scarce: we predict interest in unsupervised adaptation will grow.

Efficiency concerns stem from the tension between high-quality, fine-grained adaptation and practical use of limited resources. As typical NMT models grow larger (Wei et al., 2020b) and more work emphasises extreme adaptation at the user or sentence level, there is growing awareness that training and storing large numbers of deep neural models is unsustainable from both a financial and an environmental perspective (Strubell et al., 2019; Bender et al., 2021). Energy use and carbon footprint are certainly greater when training from scratch on hundreds of millions of sentences as compared to fine-tuning models only on relevant data. However, it may be that the aggregate effect of ‘personalizing’ large NMT systems for many users outweighs the reduced per-model training time. In general, we predict that efficiency in terms of tuning and storing models will continue to grow in importance. This convergence of requirements may draw attention to approaches which adapt using fewer tuning steps, or do not adapt the entire model. These may be variations on adapter layers (Section 6.3), ways to find sparse or underused subsets of parameters for minimal model adaptation (Section 7.1.1), or means of adapting the output that avoid parameter tuning completely (Section 8.3).

Finally, we note an interesting trend of machine learning research into neural *unlearning* (Kwak et al., 2017; Cao et al., 2018; Neel & Sharifi-Malvajerdi, 2021). We have previously described catastrophic forgetting as a pitfall for adaptation, but forgetting can be viewed as a feature, not a bug. As we saw in case studies on low resource and gender translation, domain adaptation techniques can find application in intentionally forgetting or abandoning certain behaviour which is no longer desirable, rather than specifically learning new behaviour. We might want to unlearn domains we know we will not translate in the future, or abandon a tendency to default to outdated translations, or mitigate more subtle effects resulting from word associations relating to demographic biases. Adaptation for the purposes of intentional forgetting may become more relevant with new risks connected to NLP systems. For example, recent research has suggested that neural language models can be ‘attacked’ to retrieve their training data (Wallace et al., 2020; Gong et al., 2020). Given this risk the ability to redact specific translation examples via unlearning might be desirable, especially in privacy-sensitive cases such as biomedical translation. As a relatively new area of research, the effect of unlearning on NMT remains an open question.

Natural language is both complex and evolving, as are the AI systems that interact with natural language. With this survey we hope to draw attention to the possible benefits and drawbacks of different approaches to domain adaptive translation, as well as their possible applications. We hope that future work on adapting neural machine translation will focus not only on individual domains of immediate interest, but on the range of machine translation abilities that we wish to maintain or abandon.

Acknowledgements

Many thanks to Bill Byrne for helpful comments on earlier versions of this survey, as well as associate editor Marcello Federico and the anonymous reviewers for their feedback. The survey was finalized while working at RWS Language Weaver, and draws from the author’s PhD thesis which was supported by EPSRC grants EP/M508007/1 and EP/N509620/1.

References

- Abdul Rauf, S., Rosales Núñez, J. C., Pham, M. Q., & Yvon, F. (2020). LIMSIS @ WMT 2020. In *Proceedings of the Fifth Conference on Machine Translation*, pp. 803–812, Online. Association for Computational Linguistics.
- Agrawal, S., & Carpuat, M. (2019). Controlling text complexity in 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)*, pp. 1549–1564, Hong Kong, China. Association for Computational Linguistics.
- Aharoni, R., & Goldberg, Y. (2017). Towards string-to-tree neural machine translation. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, Vol. 2, pp. 132–140.
- Aharoni, R., & Goldberg, Y. (2020). Unsupervised domain clusters in pretrained language models. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 7747–7763, Online. Association for Computational Linguistics.
- Aharoni, R., Johnson, M., & Firat, O. (2019). Massively multilingual neural machine translation. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pp. 3874–3884, Minneapolis, Minnesota. Association for Computational Linguistics.
- Ahmadnia, B., & Dorr, B. J. (2020). Low-resource multi-domain machine translation for spanish-farsi: Neural or statistical?. In Shakshuki, E. M., & Yasar, A. (Eds.), *The 11th International Conference on Emerging Ubiquitous Systems and Pervasive Networks (EUSPN 2020) / The 10th International Conference on Current and Future Trends of Information and Communication Technologies in Healthcare (ICTH-2020) / Affiliated Workshops, Madeira, Portugal, November 2-5, 2020*, Vol. 177 of *Procedia Computer Science*, pp. 575–580. Elsevier.
- Akhbardeh, F., Arkhangorodsky, A., Biesialska, M., Bojar, O., Chatterjee, R., Chaudhary, V., Costa-jussa, M. R., España-Bonet, C., Fan, A., Federmann, C., Freitag, M., Graham, Y., Grundkiewicz, R., Haddow, B., Harter, L., Heafield, K., Homan, C., Huck, M., Amponsah-Kaakyire, K., Kasai, J., Khashabi, D., Knight, K., Koemi, T., Koehn, P., Lourie, N., Monz, C., Morishita, M., Nagata, M., Nagesh, A., Nakazawa, T., Negri, M., Pal, S., Tapo, A. A., Turchi, M., Vydrin, V., & Zampieri, M. (2021). Findings of the 2021 conference on machine translation (WMT21). In *Proceedings of the Sixth Conference on Machine Translation*, pp. 1–88, Online. Association for Computational Linguistics.
- Aljundi, R., Lin, M., Goujaud, B., & Bengio, Y. (2019). Gradient based sample selection for online continual learning. In *Advances in Neural Information Processing Systems*, pp. 11816–11825.
- Allauzen, C., & Riley, M. (2011). Bayesian Language Model Interpolation for Mobile Speech Input. In *Proceedings of the Twelfth Annual Conference of the International Speech Communication Association*.

- Anastasopoulos, A., Cattelan, A., Dou, Z.-Y., Federico, M., Federmann, C., Genzel, D., Guzmán, F., Hu, J., Hughes, M., Koehn, P., Lazar, R., Lewis, W., Neubig, G., Niu, M., Öktem, A., Paquin, E., Tang, G., & Tur, S. (2020). TICO-19: the translation initiative for COvid-19. In *Proceedings of the 1st Workshop on NLP for COVID-19 (Part 2) at EMNLP 2020*, Online. Association for Computational Linguistics.
- Arthur, P., Neubig, G., & Nakamura, S. (2016). Incorporating discrete translation lexicons into neural machine translation. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pp. 1557–1567, Austin, Texas. Association for Computational Linguistics.
- Assylbekov, Z., Takhanov, R., Myrzakhmetov, B., & Washington, J. N. (2017). Syllable-aware neural language models: A failure to beat character-aware ones. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pp. 1866–1872, Copenhagen, Denmark. Association for Computational Linguistics.
- Ataman, D., Negri, M., Turchi, M., & Federico, M. (2017). Linguistically motivated vocabulary reduction for neural machine translation from turkish to english. *The Prague Bulletin of Mathematical Linguistics*, 108(1), 331–342.
- Axelrod, A. (2017). Cynical selection of language model training data. *CoRR*, abs/1709.02279.
- Axelrod, A., He, X., & Gao, J. (2011). Domain adaptation via pseudo in-domain data selection. In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing*, pp. 355–362, Edinburgh, Scotland, UK. Association for Computational Linguistics.
- Bahdanau, D., Cho, K., & Bengio, Y. (2015). Neural machine translation by jointly learning to align and translate. In *ICLR*.
- Bañón, M., Chen, P., Haddow, B., Heafield, K., Hoang, H., Esplà-Gomis, M., Forcada, M. L., Kamran, A., Kirefu, F., Koehn, P., Ortiz Rojas, S., Pla Sempere, L., Ramírez-Sánchez, G., Sarrías, E., Strelec, M., Thompson, B., Waites, W., Wiggins, D., & Zaragoza, J. (2020). ParaCrawl: Web-scale acquisition of parallel corpora. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 4555–4567, Online. Association for Computational Linguistics.
- Bapna, A., & Firat, O. (2019a). Non-parametric adaptation for 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)*, pp. 1921–1931, Minneapolis, Minnesota. Association for Computational Linguistics.
- Bapna, A., & Firat, O. (2019b). 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)*, pp. 1538–1548, Hong Kong, China. Association for Computational Linguistics.

- Barone, A. V. M., Haddow, B., Germann, U., & Sennrich, R. (2017). Regularization techniques for fine-tuning in Neural Machine Translation. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pp. 1489–1494.
- Barrault, L., Bojar, O., Costa-jussà, M. R., Federmann, C., Fishel, M., Graham, Y., Haddow, B., Huck, M., Koehn, P., Malmasi, S., Monz, C., Müller, M., Pal, S., Post, M., & Zampieri, M. (2019). Findings of the 2019 conference on machine translation (WMT19). In *Proceedings of the Fourth Conference on Machine Translation (Volume 2: Shared Task Papers, Day 1)*, pp. 1–61, Florence, Italy. Association for Computational Linguistics.
- Baum, E. B., & Wilczek, F. (1988). Supervised learning of probability distributions by neural networks. In *Neural information processing systems*, pp. 52–61.
- Bawden, R., Bretonnel Cohen, K., Grozea, C., Jimeno Yepes, A., Kittner, M., Krallinger, M., Mah, N., Neveol, A., Neves, M., Soares, F., Siu, A., Verspoor, K., & Vicente Navarro, M. (2019). Findings of the WMT 2019 biomedical translation shared task: Evaluation for MEDLINE abstracts and biomedical terminologies. In *Proceedings of the Fourth Conference on Machine Translation (Volume 3: Shared Task Papers, Day 2)*, pp. 29–53, Florence, Italy. Association for Computational Linguistics.
- Baziotis, C., Haddow, B., & Birch, A. (2020). Language model prior for low-resource neural machine translation. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 7622–7634, Online. Association for Computational Linguistics.
- Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (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*, pp. 610–623.
- Bengio, S., Vinyals, O., Jaitly, N., & Shazeer, N. (2015). Scheduled sampling for sequence prediction with recurrent neural networks. In *Advances in Neural Information Processing Systems*, pp. 1171–1179.
- Bengio, Y., Ducharme, R., Vincent, P., & Jauvin, C. (2003). A neural probabilistic language model. *Journal of machine learning research*, 3(Feb), 1137–1155.
- Bengio, Y., Louradour, J., Collobert, R., & Weston, J. (2009). Curriculum learning. In *Proceedings of the 26th annual international conference on machine learning*, pp. 41–48.
- Berard, A., Calapodescu, I., & Roux, C. (2019). Naver labs Europe’s systems for the WMT19 machine translation robustness task. In *Proceedings of the Fourth Conference on Machine Translation (Volume 2: Shared Task Papers, Day 1)*, pp. 526–532, Florence, Italy. Association for Computational Linguistics.
- Bertoldi, N., & Federico, M. (2009). Domain adaptation for statistical machine translation with monolingual resources. In *Proceedings of the Fourth Workshop on Statistical Machine Translation*, pp. 182–189, Athens, Greece. Association for Computational Linguistics.
- Bishop, C. M. (1995). Training with noise is equivalent to tikhonov regularization. *Neural computation*, 7(1), 108–116.

- Blackwood, G., Ballesteros, M., & Ward, T. (2018). Multilingual neural machine translation with task-specific attention. In *Proceedings of the 27th International Conference on Computational Linguistics*, pp. 3112–3122, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Blei, D. M., Ng, A. Y., & Jordan, M. I. (2003). Latent dirichlet allocation. *Journal of machine Learning research*, 3(Jan), 993–1022.
- Bogoychev, N., & Sennrich, R. (2019). Domain, translationese and noise in synthetic data for neural machine translation. *CoRR*, *abs/1911.03362*.
- Bojar, O., Chatterjee, R., Federmann, C., Graham, Y., Haddow, B., Huck, M., Jimeno Yepes, A., Koehn, P., Logacheva, V., Monz, C., Negri, M., N ev ol, A., Neves, M., Popel, M., Post, M., Rubino, R., Scarton, C., Specia, L., Turchi, M., Verspoor, K., & Zampieri, M. (2016). Findings of the 2016 conference on machine translation. In *Proceedings of the First Conference on Machine Translation: Volume 2, Shared Task Papers*, pp. 131–198, Berlin, Germany. Association for Computational Linguistics.
- Britz, D., Le, Q., & Pryzant, R. (2017). Effective domain mixing for neural machine translation. In *Proceedings of the Second Conference on Machine Translation*, pp. 118–126, Copenhagen, Denmark. Association for Computational Linguistics.
- Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D. M., Wu, J., Winter, C., Hesse, C., Chen, M., Sigler, E., Litwin, M., Gray, S., Chess, B., Clark, J., Berner, C., McCandlish, S., Radford, A., Sutskever, I., & Amodei, D. (2020). Language models are few-shot learners. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., & Lin, H. (Eds.), *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Bucilu a, C., Caruana, R., & Niculescu-Mizil, A. (2006). Model compression. In *Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 535–541.
- Buj, D. M., Ib a nez Garc a, D., Parcheta, Z., & Casacuberta, F. (2020). NICE: Neural integrated custom engines. In *Proceedings of the 22nd Annual Conference of the European Association for Machine Translation*, pp. 329–338, Lisboa, Portugal. European Association for Machine Translation.
- Bulte, B., & Tezcan, A. (2019). Neural fuzzy repair: Integrating fuzzy matches into neural machine translation. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 1800–1809, Florence, Italy. Association for Computational Linguistics.
- Burlot, F., & Yvon, F. (2018). Using monolingual data in neural machine translation: a systematic study. In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pp. 144–155, Brussels, Belgium. Association for Computational Linguistics.
- Cai, D., Wang, Y., Li, H., Lam, W., & Liu, L. (2021). Neural machine translation with monolingual translation memory. In Zong, C., Xia, F., Li, W., & Navigli, R. (Eds.),

- Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing, ACL/IJCNLP 2021, (Volume 1: Long Papers), Virtual Event, August 1-6, 2021*, pp. 7307–7318. Association for Computational Linguistics.
- Cao, Y., Yu, A. F., Aday, A., Stahl, E., Merwine, J., & Yang, J. (2018). Efficient repair of polluted machine learning systems via causal unlearning. In Kim, J., Ahn, G., Kim, S., Kim, Y., López, J., & Kim, T. (Eds.), *Proceedings of the 2018 on Asia Conference on Computer and Communications Security, AsiaCCS 2018, Incheon, Republic of Korea, June 04-08, 2018*, pp. 735–747. ACM.
- Cao, Y., Wei, H.-R., Chen, B., & Wan, X. (2021). Continual learning for neural machine translation. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pp. 3964–3974, Online. Association for Computational Linguistics.
- Carrión-Ponz, S., & Casacuberta, F. (2022). Few-shot regularization to tackle catastrophic forgetting in multilingual machine translation. In *Proceedings of the 15th biennial conference of the Association for Machine Translation in the Americas (Volume 1: Research Track)*, pp. 188–199, Orlando, USA. Association for Machine Translation in the Americas.
- Caswell, I., Chelba, C., & Grangier, D. (2019). Tagged back-translation. In *Proceedings of the Fourth Conference on Machine Translation (Volume 1: Research Papers)*, pp. 53–63, Florence, Italy. Association for Computational Linguistics.
- Chen, B., Cherry, C., Foster, G., & Larkin, S. (2017). Cost weighting for neural machine translation domain adaptation. In *Proceedings of the First Workshop on Neural Machine Translation*, pp. 40–46, Vancouver. Association for Computational Linguistics.
- Chen, G., Chen, Y., Wang, Y., & Li, V. O. K. (2020a). Lexical-constraint-aware neural machine translation via data augmentation. In Bessiere, C. (Ed.), *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI 2020*, pp. 3587–3593. ijcai.org.
- Chen, P., Bogoychev, N., & Germann, U. (2020b). Character mapping and ad-hoc adaptation: Edinburgh’s IWSLT 2020 open domain translation system. In *Proceedings of the 17th International Conference on Spoken Language Translation*, pp. 122–129, Online. Association for Computational Linguistics.
- Cherry, C., Foster, G., Bapna, A., Firat, O., & Macherey, W. (2018). Revisiting character-based neural machine translation with capacity and compression. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pp. 4295–4305, Brussels, Belgium. Association for Computational Linguistics.
- Chinea-Ríos, M., Peris, A., & Casacuberta, F. (2017). Adapting neural machine translation with parallel synthetic data. In *Proceedings of the Second Conference on Machine Translation*, pp. 138–147, Copenhagen, Denmark. Association for Computational Linguistics.
- Cho, K., van Merriënboer, B., Gulcehre, C., Bahdanau, D., Bougares, F., Schwenk, H., & Bengio, Y. (2014). Learning phrase representations using RNN encoder–decoder for

- statistical machine translation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 1724–1734, Doha, Qatar. Association for Computational Linguistics.
- Chu, C., & Dabre, R. (2019). Multilingual multi-domain adaptation approaches for neural machine translation. *CoRR*, *abs/1906.07978*.
- Chu, C., Dabre, R., & Kurohashi, S. (2017). An empirical comparison of domain adaptation methods for neural machine translation. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pp. 385–391.
- Chu, C., & Wang, R. (2018). A Survey of Domain Adaptation for Neural Machine Translation. In *Proceedings of COLING 2018, the 27 International Conference on Computational Linguistics: Technical Papers*, pp. 1304–1319.
- Chu, C., & Wang, R. (2020). A survey of domain adaptation for machine translation. *Journal of information processing*, *28*, 413–426.
- Collobert, R., & Weston, J. (2008). A unified architecture for natural language processing: Deep neural networks with multitask learning. In *Proceedings of the 25th International Conference on Machine Learning, ICML '08*, p. 160–167, New York, NY, USA. Association for Computing Machinery.
- Collobert, R., Weston, J., Bottou, L., Karlen, M., Kavukcuoglu, K., & Kuksa, P. (2011). Natural language processing (almost) from scratch. *J. Mach. Learn. Res.*, *12*(null), 2493–2537.
- Costa-jussà, M. R., & de Jorje, A. (2020). Fine-tuning neural machine translation on gender-balanced datasets. In *Proceedings of the Second Workshop on Gender Bias in Natural Language Processing*, pp. 26–34, Barcelona, Spain (Online). Association for Computational Linguistics.
- Costa-jussà, M. R., & Fonollosa, J. A. R. (2016). Character-based neural machine translation. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pp. 357–361, Berlin, Germany. Association for Computational Linguistics.
- Currey, A., & Heafield, K. (2019). Zero-resource neural machine translation with monolingual pivot data. In *Proceedings of the 3rd Workshop on Neural Generation and Translation*, pp. 99–107, Hong Kong. Association for Computational Linguistics.
- Currey, A., Mathur, P., & Dinu, G. (2020). Distilling multiple domains for neural machine translation. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 4500–4511, Online. Association for Computational Linguistics.
- Currey, A., Miceli Barone, A. V., & Heafield, K. (2017). Copied monolingual data improves low-resource neural machine translation. In *Proceedings of the Second Conference on Machine Translation*, pp. 148–156, Copenhagen, Denmark. Association for Computational Linguistics.
- Dakwale, P., & Monz, C. (2017). Fine-tuning for Neural Machine Translation with limited degradation across in-and out-of-domain data. In *Proceedings of the 16th Machine Translation Summit (MT-Summit 2017)*, pp. 156–169.

- Daumé III, H., & Jagarlamudi, J. (2011). Domain adaptation for machine translation by mining unseen words. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pp. 407–412, Portland, Oregon, USA. Association for Computational Linguistics.
- de Gispert, A., Iglesias, G., Blackwood, G., Banga, E. R., & Byrne, W. (2010). Hierarchical phrase-based translation with weighted finite-state transducers and shallow-n grammars. *Computational Linguistics*, 36(3), 505–533.
- Del, M., Korotkova, E., & Fishel, M. (2021). Translation transformers rediscover inherent data domains. *CoRR*, abs/2109.07864.
- Deng, Y., Yu, H., Yu, H., Duan, X., & Luo, W. (2020). Factorized transformer for multi-domain neural machine translation. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pp. 4221–4230, Online. Association for Computational Linguistics.
- Devlin, J., Chang, M.-W., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. 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)*, pp. 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Ding, S., Renduchintala, A., & Duh, K. (2019). A call for prudent choice of subword merge operations in neural machine translation. In *Proceedings of Machine Translation Summit XVII Volume 1: Research Track*, pp. 204–213, Dublin, Ireland. European Association for Machine Translation.
- Dinu, G., Mathur, P., Federico, M., & Al-Onaizan, Y. (2019). Training neural machine translation to apply terminology constraints. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 3063–3068, Florence, Italy. Association for Computational Linguistics.
- Dong, D., Wu, H., He, W., Yu, D., & Wang, H. (2015). Multi-task learning for multiple language translation. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pp. 1723–1732, Beijing, China. Association for Computational Linguistics.
- Dou, Z.-Y., Anastasopoulos, A., & Neubig, G. (2020). Dynamic data selection and weighting for iterative back-translation. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 5894–5904. Association for Computational Linguistics.
- Dou, Z.-Y., Hu, J., Anastasopoulos, A., & Neubig, G. (2019a). Unsupervised domain adaptation for neural machine translation with domain-aware feature embeddings. 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)*, pp. 1417–1422, Hong Kong, China. Association for Computational Linguistics.

- Dou, Z.-Y., Wang, X., Hu, J., & Neubig, G. (2019b). Domain differential adaptation for neural machine translation. In *Proceedings of the 3rd Workshop on Neural Generation and Translation*, pp. 59–69, Hong Kong. Association for Computational Linguistics.
- Dougal, D. K., & Lonsdale, D. (2020). Improving NMT quality using terminology injection. In *Proceedings of the 12th Language Resources and Evaluation Conference*, pp. 4820–4827, Marseille, France. European Language Resources Association.
- Duan, C., Chen, K., Wang, R., Utiyama, M., Sumita, E., Zhu, C., & Zhao, T. (2020). Modeling future cost for neural machine translation. *CoRR*, *abs/2002.12558*.
- Eck, M., Vogel, S., & Waibel, A. (2004). Language model adaptation for statistical machine translation based on information retrieval. In *Proceedings of the Fourth International Conference on Language Resources and Evaluation (LREC'04)*, Lisbon, Portugal. European Language Resources Association (ELRA).
- Edunov, S., Ott, M., Auli, M., Grangier, D., et al. (2018). Classical structured prediction losses for sequence to sequence learning. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, Vol. 1, pp. 355–364.
- Emelin, D., Titov, I., & Sennrich, R. (2020). Detecting word sense disambiguation biases in machine translation for model-agnostic adversarial attacks. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 7635–7653, Online. Association for Computational Linguistics.
- Fadaee, M., Bisazza, A., & Monz, C. (2017). Data augmentation for low-resource neural machine translation. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pp. 567–573, Vancouver, Canada. Association for Computational Linguistics.
- Farajian, M. A., Bertoldi, N., Negri, M., Turchi, M., & Federico, M. (2018). Evaluation of terminology translation in instance-based neural MT adaptation. In *Proceedings of the 21st Annual Conference of the European Association for Machine Translation*, pp. 149–158. European Association for Machine Translation.
- Farajian, M. A., Turchi, M., Negri, M., & Federico, M. (2017). Multi-domain neural machine translation through unsupervised adaptation. In *Proceedings of the Second Conference on Machine Translation*, pp. 127–137, Copenhagen, Denmark. Association for Computational Linguistics.
- Feely, W., Hasler, E., & de Gispert, A. (2019). Controlling Japanese honorifics in English-to-Japanese neural machine translation. In *Proceedings of the 6th Workshop on Asian Translation*, pp. 45–53, Hong Kong, China. Association for Computational Linguistics.
- Finn, C., Abbeel, P., & Levine, S. (2017). Model-agnostic meta-learning for fast adaptation of deep networks. In Precup, D., & Teh, Y. W. (Eds.), *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, Vol. 70 of *Proceedings of Machine Learning Research*, pp. 1126–1135. PMLR.
- Firat, O., Sankaran, B., Al-onaizan, Y., Yarman Vural, F. T., & Cho, K. (2016). Zero-resource translation with multi-lingual neural machine translation. In *Proceedings*

- of the 2016 Conference on Empirical Methods in Natural Language Processing, pp. 268–277, Austin, Texas. Association for Computational Linguistics.
- Foster, G., Goutte, C., & Kuhn, R. (2010). Discriminative instance weighting for domain adaptation in statistical machine translation. In *Proceedings of the 2010 Conference on Empirical Methods in Natural Language Processing*, pp. 451–459, Cambridge, MA. Association for Computational Linguistics.
- Freitag, M., & Al-Onaizan, Y. (2016). Fast domain adaptation for Neural Machine Translation. *CoRR*, *abs/1612.06897*.
- Freitag, M., Al-Onaizan, Y., & Sankaran, B. (2017). Ensemble distillation for neural machine translation..
- Fukuda, T., Suzuki, M., Kurata, G., Thomas, S., Cui, J., & Ramabhadran, B. (2017). Efficient knowledge distillation from an ensemble of teachers. In *INTERSPEECH*.
- Gage, P. (1994). A new algorithm for data compression. *The C Users Journal*, *12*(2), 23–38.
- Gallé, M. (2019). Investigating the effectiveness of BPE: The power of shorter sequences. 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)*, pp. 1375–1381, Hong Kong, China. Association for Computational Linguistics.
- Gangi, M. A. D., Enyedi, R., Brusadin, A., & Federico, M. (2019). Robust neural machine translation for clean and noisy speech transcripts. In *IWSLT 2019, International Workshop on Spoken Language Translation*.
- García-Martínez, M., Barrault, L., & Bougares, F. (2016). Factored neural machine translation. *CoRR*, *abs/1609.04621*.
- Gong, M., Pan, K., Xie, Y., Qin, A. K., & Tang, Z. (2020). Preserving differential privacy in deep neural networks with relevance-based adaptive noise imposition. *Neural Networks*, *125*, 131–141.
- Gordon, M., & Duh, K. (2020). Distill, adapt, distill: Training small, in-domain models for neural machine translation. In *Proceedings of the Fourth Workshop on Neural Generation and Translation*, pp. 110–118, Online. Association for Computational Linguistics.
- Graham, Y., Haddow, B., & Koehn, P. (2020). Statistical power and translationese in machine translation evaluation. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 72–81, Online. Association for Computational Linguistics.
- Grangier, D., & Iyer, D. (2022). The trade-offs of domain adaptation for neural language models. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 3802–3813, Dublin, Ireland. Association for Computational Linguistics.
- Gu, J., Hassan, H., Devlin, J., & Li, V. O. (2018). Universal neural machine translation for extremely low resource languages. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pp. 344–354, New Orleans, Louisiana. Association for Computational Linguistics.

- Gu, S., & Feng, Y. (2020). Investigating catastrophic forgetting during continual training for neural machine translation. In *Proceedings of the 28th International Conference on Computational Linguistics*, pp. 4315–4326, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Gu, S., Feng, Y., & Liu, Q. (2019). Improving domain adaptation translation with domain invariant and specific information. 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)*, pp. 3081–3091, Minneapolis, Minnesota. Association for Computational Linguistics.
- Gu, S., Feng, Y., & Xie, W. (2021). Pruning-then-expanding model for domain adaptation of neural machine translation. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2021, Online, June 6-11, 2021*, pp. 3942–3952. Association for Computational Linguistics.
- Gülcehre, C., Firat, O., Xu, K., Cho, K., Barrault, L., Lin, H.-C., Bougares, F., Schwenk, H., & Bengio, Y. (2015). On using monolingual corpora in Neural Machine Translation. *CoRR*, *abs/1503.03535*.
- Haque, R., Moslem, Y., & Way, A. (2020). Terminology-aware sentence mining for NMT domain adaptation: ADAPT’s submission to the adap-MT 2020 English-to-Hindi AI translation shared task. In *Proceedings of the 17th International Conference on Natural Language Processing (ICON): Adap-MT 2020 Shared Task*, pp. 17–23, Patna, India. NLP Association of India (NLP AI).
- Hasler, E., de Gispert, A., Iglesias, G., & Byrne, B. (2018). Neural machine translation decoding with terminology constraints. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pp. 506–512, New Orleans, Louisiana. Association for Computational Linguistics.
- Hasler, E., de Gispert, A., Stahlberg, F., Waite, A., & Byrne, B. (2017). Source sentence simplification for statistical machine translation. *Computer Speech & Language*, *45*, 221–235.
- Hasler, E., Domhan, T., Trenous, J., Tran, K., Byrne, B., & Hieber, F. (2021). Improving the quality trade-off for neural machine translation multi-domain adaptation. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, EMNLP 2021, Virtual Event / Punta Cana, Dominican Republic, 7-11 November, 2021*, pp. 8470–8477. Association for Computational Linguistics.
- He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778.
- Hedderich, M. A., Lange, L., Adel, H., Strötgen, J., & Klakow, D. (2020). A survey on recent approaches for natural language processing in low-resource scenarios. *CoRR*, *abs/2010.12309*.

- Hinton, G. E., Srivastava, N., Krizhevsky, A., Sutskever, I., & Salakhutdinov, R. R. (2012). Improving neural networks by preventing co-adaptation of feature detectors. *CoRR*, *abs/1207.0580*.
- Hinton, G. E., Vinyals, O., & Dean, J. (2015). Distilling the knowledge in a neural network..
- Hokamp, C., & Liu, Q. (2017). Lexically constrained decoding for sequence generation using grid beam search. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 1535–1546, Vancouver, Canada. Association for Computational Linguistics.
- Hovy, D., Bianchi, F., & Fornaciari, T. (2020). “You Sound Just Like Your Father” Commercial Machine Translation Systems Include Stylistic Biases. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 1686–1690, Online. Association for Computational Linguistics.
- Hu, J., & Neubig, G. (2021). Phrase-level active learning for neural machine translation. *CoRR*, *abs/2106.11375*.
- Hu, J., Xia, M., Neubig, G., & Carbonell, J. (2019). Domain adaptation of neural machine translation by lexicon induction. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 2989–3001, Florence, Italy. Association for Computational Linguistics.
- Huck, M., Birch, A., & Haddow, B. (2015). Mixed-Domain vs. Multi-Domain Statistical Machine Translation. In *Proceedings of the 15th Machine Translation Summit (MT-Summit 2015)*, pp. 240–255.
- Huck, M., Riess, S., & Fraser, A. (2017). Target-side word segmentation strategies for neural machine translation. In *Proceedings of the Second Conference on Machine Translation*, pp. 56–67, Copenhagen, Denmark. Association for Computational Linguistics.
- Jean, S., Lauly, S., Firat, O., & Cho, K. (2017). Does neural machine translation benefit from larger context?. *CoRR*, *abs/1704.05135*.
- Jehl, L., & Riezler, S. (2018). Document-level information as side constraints for improved neural patent translation. In *Proceedings of the 13th Conference of the Association for Machine Translation in the Americas (Volume 1: Research Track)*, pp. 1–12, Boston, MA. Association for Machine Translation in the Americas.
- Ji, B., Zhang, Z., Duan, X., Zhang, M., Chen, B., & Luo, W. (2020). Cross-lingual pre-training based transfer for zero-shot neural machine translation. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*, pp. 115–122. AAAI Press.
- Jiang, H., Liang, C., Wang, C., & Zhao, T. (2020). Multi-domain neural machine translation with word-level adaptive layer-wise domain mixing. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 1823–1834, Online. Association for Computational Linguistics.
- Jiao, R., Yang, Z., Sun, M., & Liu, Y. (2021). Alternated training with synthetic and authentic data for neural machine translation. In Zong, C., Xia, F., Li, W., & Navigli,

- R. (Eds.), *Findings of the Association for Computational Linguistics: ACL/IJCNLP 2021, Online Event, August 1-6, 2021*, Vol. ACL/IJCNLP 2021 of *Findings of ACL*, pp. 1828–1834. Association for Computational Linguistics.
- Jin, D., Jin, Z., Zhou, J. T., & Szolovits, P. (2020). A simple baseline to semi-supervised domain adaptation for machine translation. *CoRR*, *abs/2001.08140*.
- Joty, S., Sajjad, H., Durrani, N., Al-Mannai, K., Abdelali, A., & Vogel, S. (2015). How to avoid unwanted pregnancies: Domain adaptation using neural network models. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pp. 1259–1270, Lisbon, Portugal. Association for Computational Linguistics.
- Junczys-Dowmunt, M. (2018a). Dual conditional cross-entropy filtering of noisy parallel corpora. In *Proceedings of the Third Conference on Machine Translation: Shared Task Papers*, pp. 888–895, Belgium, Brussels. Association for Computational Linguistics.
- Junczys-Dowmunt, M. (2018b). Microsoft’s submission to the WMT2018 news translation task: How I learned to stop worrying and love the data. In *Proceedings of the Third Conference on Machine Translation: Shared Task Papers*, pp. 425–430, Belgium, Brussels. Association for Computational Linguistics.
- Junczys-Dowmunt, M. (2019). Microsoft translator at WMT 2019: Towards large-scale document-level neural machine translation. In *Proceedings of the Fourth Conference on Machine Translation (Volume 2: Shared Task Papers, Day 1)*, pp. 225–233, Florence, Italy. Association for Computational Linguistics.
- Jwalapuram, P., Joty, S., & Shen, Y. (2020). Pronoun-targeted fine-tuning for NMT with hybrid losses. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 2267–2279, Online. Association for Computational Linguistics.
- Kalchbrenner, N., & Blunsom, P. (2013). Recurrent continuous translation models. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pp. 1700–1709, Seattle, Washington, USA. Association for Computational Linguistics.
- Karakanta, A., Dehdari, J., & van Genabith, J. (2018). Neural machine translation for low-resource languages without parallel corpora. *Machine Translation*, *32*(1-2), 167–189.
- Karpukhin, V., Levy, O., Eisenstein, J., & Ghazvininejad, M. (2019). Training on synthetic noise improves robustness to natural noise in machine translation. In *Proceedings of the 5th Workshop on Noisy User-generated Text (W-NUT 2019)*, pp. 42–47, Hong Kong, China. Association for Computational Linguistics.
- Kessler, B., Nunberg, G., & Schutze, H. (1997). Automatic detection of text genre. In *35th Annual Meeting of the Association for Computational Linguistics and 8th Conference of the European Chapter of the Association for Computational Linguistics*, pp. 32–38, Madrid, Spain. Association for Computational Linguistics.
- Khandelwal, U., Fan, A., Jurafsky, D., Zettlemoyer, L., & Lewis, M. (2021). Nearest neighbor machine translation. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.

- Khayrallah, H., & Koehn, P. (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*, pp. 74–83, Melbourne, Australia. Association for Computational Linguistics.
- Khayrallah, H., Kumar, G., Duh, K., Post, M., & Koehn, P. (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)*, pp. 20–25, Taipei, Taiwan. Asian Federation of Natural Language Processing.
- Khayrallah, H., Thompson, B., Duh, K., & Koehn, P. (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*, pp. 36–44.
- Kim, S., Bisazza, A., & Turkmen, F. (2021). Using confidential data for domain adaptation of neural machine translation. In *Proceedings of the Third Workshop on Privacy in Natural Language Processing*, pp. 46–52, Online. Association for Computational Linguistics.
- Kim, Y., Jernite, Y., Sontag, D., & Rush, A. M. (2016). Character-aware neural language models. In *30th AAAI Conference on Artificial Intelligence, AAAI 2016*, pp. 2741–2749. AAAI press.
- Kim, Y., & Rush, A. M. (2016). Sequence-level knowledge distillation. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pp. 1317–1327, Austin, Texas. Association for Computational Linguistics.
- Kim, Y., Petrov, P., Petrushkov, P., Khadivi, S., & Ney, H. (2019a). Pivot-based transfer learning for neural machine translation between non-English languages. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pp. 866–876, Hong Kong, China. Association for Computational Linguistics.
- Kim, Y., Tran, D. T., & Ney, H. (2019b). When and why is document-level context useful in neural machine translation?. In *Proceedings of the Fourth Workshop on Discourse in Machine Translation (DiscoMT 2019)*, pp. 24–34, Hong Kong, China. Association for Computational Linguistics.
- Kirkpatrick, J., Pascanu, R., Rabinowitz, N., Veness, J., Desjardins, G., Rusu, A. A., Milan, K., Quan, J., Ramalho, T., Grabska-Barwinska, A., et al. (2017). Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences of the United States of America*, 114(13), 3521–3526.
- Kobus, C., Crego, J., & Senellart, J. (2017). Domain control for neural machine translation. In *Proceedings of the International Conference Recent Advances in Natural Language Processing, RANLP 2017*, pp. 372–378.
- Kocmi, T., & Bojar, O. (2017). Curriculum learning and minibatch bucketing in neural machine translation. In *Proceedings of the International Conference Recent Advances in Natural Language Processing, RANLP 2017*, pp. 379–386, Varna, Bulgaria. INCOMA Ltd.

- Kocmi, T., & Bojar, O. (2018). Trivial transfer learning for low-resource neural machine translation. In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pp. 244–252, Brussels, Belgium. Association for Computational Linguistics.
- Koehn, P. (2017). Neural machine translation. *CoRR*, *abs/1709.07809*.
- Koehn, P. (2020). *Neural machine translation*. Cambridge University Press.
- Koehn, P., Duh, K., & Thompson, B. (2018a). The JHU machine translation systems for WMT 2018. In *Proceedings of the Third Conference on Machine Translation: Shared Task Papers*, pp. 438–444, Belgium, Brussels. Association for Computational Linguistics.
- Koehn, P., Khayrallah, H., Heafield, K., & Forcada, M. L. (2018b). Findings of the WMT 2018 shared task on parallel corpus filtering. In *Proceedings of the Third Conference on Machine Translation: Shared Task Papers*, pp. 726–739, Belgium, Brussels. Association for Computational Linguistics.
- Koehn, P., & Knowles, R. (2017). Six challenges for neural machine translation. In *Proceedings of the First Workshop on Neural Machine Translation*, pp. 28–39, Vancouver. Association for Computational Linguistics.
- Koehn, P., & Senellart, J. (2010). Convergence of translation memory and statistical machine translation. In *Proceedings of the Second Joint EM+/CNGL Workshop: Bringing MT to the User: Research on Integrating MT in the Translation Industry*, pp. 21–32, Denver, Colorado, USA. Association for Machine Translation in the Americas.
- Kothur, S. S. R., Knowles, R., & Koehn, P. (2018). Document-level adaptation for neural machine translation. In *Proceedings of the 2nd Workshop on Neural Machine Translation and Generation*, pp. 64–73, Melbourne, Australia. Association for Computational Linguistics.
- Kudo, T. (2018). Subword regularization: Improving neural network translation models with multiple subword candidates. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 66–75, Melbourne, Australia. Association for Computational Linguistics.
- Kumar, G., Foster, G., Cherry, C., & Krikun, M. (2019). Reinforcement learning based curriculum optimization for 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)*, pp. 2054–2061, Minneapolis, Minnesota. Association for Computational Linguistics.
- Kumar, S., & Byrne, W. (2004). Minimum Bayes-risk decoding for statistical machine translation. In *Proceedings of the Human Language Technology Conference of the North American Chapter of the Association for Computational Linguistics: HLT-NAACL 2004*, pp. 169–176, Boston, Massachusetts, USA. Association for Computational Linguistics.
- Kumari, S., Jaiswal, N., Patidar, M., Patwardhan, M., Karande, S., Agarwal, P., & Vig, L. (2021). Domain adaptation for NMT via filtered iterative back-translation. In *Proceedings of the Second Workshop on Domain Adaptation for NLP*, pp. 263–271, Kyiv, Ukraine. Association for Computational Linguistics.

- Kwak, C., Lee, J., Park, K., & Lee, H. (2017). Let machines unlearn - machine unlearning and the right to be forgotten. In *23rd Americas Conference on Information Systems, AMCIS 2017, Boston, MA, USA, August 10-12, 2017*. Association for Information Systems.
- Lachaux, M.-A., Joulin, A., & Lample, G. (2020). Target conditioning for one-to-many generation. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pp. 2853–2862, Online. Association for Computational Linguistics.
- Lambert, P., Schwenk, H., Servan, C., & Abdul-Rauf, S. (2011). Investigations on translation model adaptation using monolingual data. In *Proceedings of the Sixth Workshop on Statistical Machine Translation*, pp. 284–293, Edinburgh, Scotland. Association for Computational Linguistics.
- Lavergne, T., Allauzen, A., Le, H.-S., & Yvon, F. (2011). LIMSI’s Experiments in Domain Adaptation for IWSLT11. In *IWSLT 2011, International Workshop on Spoken Language Translation*.
- Lee, G., Yang, S., & Choi, E. (2021). Improving lexically constrained neural machine translation with source-conditioned masked span prediction. In Zong, C., Xia, F., Li, W., & Navigli, R. (Eds.), *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing, ACL/IJCNLP 2021, (Volume 2: Short Papers), Virtual Event, August 1-6, 2021*, pp. 743–753. Association for Computational Linguistics.
- Lee, Y.-B., & Myaeng, S. H. (2002). Text genre classification with genre-revealing and subject-revealing features. In *Proceedings of the 25th annual international ACM SIGIR conference on Research and development in information retrieval*, pp. 145–150.
- Levin, E., & Fleisher, M. (1988). Accelerated learning in layered neural networks. *Complex Systems*, 2, 625–640.
- Lewis, W., & Eetemadi, S. (2013). Dramatically reducing training data size through vocabulary saturation. In *Proceedings of the Eighth Workshop on Statistical Machine Translation*, pp. 281–291, Sofia, Bulgaria. Association for Computational Linguistics.
- Li, J., & Jurafsky, D. (2016). Mutual information and diverse decoding improve neural machine translation. *CoRR*, *abs/1601.00372*.
- Li, R., Wang, X., & Yu, H. (2020). Metamt, a meta learning method leveraging multiple domain data for low resource machine translation. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*, pp. 8245–8252. AAAI Press.
- Li, X., Zhang, J., & Zong, C. (2018). One sentence one model for neural machine translation. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).

- Li, Z., Wang, R., Chen, K., Utiyama, M., Sumita, E., Zhang, Z., & Zhao, H. (2020). Explicit sentence compression for neural machine translation. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI*, pp. 8311–8318. AAAI Press.
- Liang, J., Zhao, C., Wang, M., Qiu, X., & Li, L. (2021). Finding sparse structures for domain specific neural machine translation. In *Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2021, Thirty-Third Conference on Innovative Applications of Artificial Intelligence, IAAI 2021, The Eleventh Symposium on Educational Advances in Artificial Intelligence, EAAI 2021, Virtual Event, February 2-9, 2021*, pp. 13333–13342. AAAI Press.
- Lin, H., Yao, L., Yang, B., Liu, D., Zhang, H., Luo, W., Huang, D., & Su, J. (2021). Towards user-driven neural machine translation. In Zong, C., Xia, F., Li, W., & Navigli, R. (Eds.), *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing, ACL/IJCNLP 2021, (Volume 1: Long Papers), Virtual Event, August 1-6, 2021*, pp. 4008–4018. Association for Computational Linguistics.
- Ling, W., Trancoso, I., Dyer, C., & Black, A. W. (2015). Character-based neural machine translation..
- Liu, J., Zhang, X., Tian, X., Wang, J., & Sangaiah, A. K. (2020a). A novel domain adaption approach for neural machine translation. *Int. J. Comput. Sci. Eng.*, 22(4), 445–453.
- Liu, X., Lai, H., Wong, D. F., & Chao, L. S. (2020b). Norm-based curriculum learning for neural machine translation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 427–436, Online. Association for Computational Linguistics.
- Lohar, P., Madden, S., O’Connor, E., Popovic, M., & Habruseva, T. (2022). Building machine translation system for software product descriptions using domain-specific sub-corpora extraction. In *Proceedings of the 15th biennial conference of the Association for Machine Translation in the Americas (Volume 1: Research Track)*, pp. 1–13, Orlando, USA. Association for Machine Translation in the Americas.
- Luong, M.-T., & Manning, C. D. (2015). Stanford Neural Machine Translation systems for spoken language domains. In *Proceedings of the International Workshop on Spoken Language Translation*, pp. 76–79.
- Luong, T., Sutskever, I., Le, Q., Vinyals, O., & Zaremba, W. (2015). Addressing the rare word problem in neural machine translation. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pp. 11–19, Beijing, China. Association for Computational Linguistics.
- Macé, V., & Servan, C. (2019). Using whole document context in neural machine translation..
- Macháček, D., Vidra, J., & Bojar, O. (2018). Morphological and language-agnostic word segmentation for nmt. In *International Conference on Text, Speech, and Dialogue*, pp. 277–284. Springer.

- Mahdieh, M., Chen, M. X., Cao, Y., & Firat, O. (2020). Rapid domain adaptation for machine translation with monolingual data. *CoRR*, *abs/2010.12652*.
- Marie, B., Rubino, R., & Fujita, A. (2020). Tagged back-translation revisited: Why does it really work?. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 5990–5997, Online. Association for Computational Linguistics.
- Martins, P., Marinho, Z., & Martins, A. (2022). Efficient machine translation domain adaptation. In *Proceedings of the 1st Workshop on Semiparametric Methods in NLP: Decoupling Logic from Knowledge*, pp. 23–29, Dublin, Ireland and Online. Association for Computational Linguistics.
- Maruf, S., Saleh, F., & Haffari, G. (2021). A survey on document-level neural machine translation: Methods and evaluation. *ACM Computing Surveys (CSUR)*, *54*(2), 1–36.
- McCloskey, M., & Cohen, N. J. (1989). Catastrophic interference in connectionist networks: The sequential learning problem. In *Psychology of learning and motivation*, Vol. 24, pp. 109–165. Elsevier.
- Mehta, S., Azarnoush, B., Chen, B., Saluja, A., Misra, V., Bihani, B., & Kumar, R. (2020). Simplify-then-translate: Automatic preprocessing for black-box machine translation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 34.
- Mghabbar, I., & Ratnamogan, P. (2020). Building a multi-domain neural machine translation model using knowledge distillation. In *ECAI 2020 - 24th European Conference on Artificial Intelligence*, Vol. 325, pp. 2116–2123. IOS Press.
- Michel, P., & Neubig, G. (2018). Extreme adaptation for personalized neural machine translation. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pp. 312–318, Melbourne, Australia. Association for Computational Linguistics.
- Michon, E., Crego, J., & Senellart, J. (2020). Integrating domain terminology into neural machine translation. In *Proceedings of the 28th International Conference on Computational Linguistics*, pp. 3925–3937, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., & Dean, J. (2013). Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, pp. 3111–3119.
- Mino, H., Tanaka, H., Ito, H., Goto, I., Yamada, I., & Tokunaga, T. (2020). Content-equivalent translated parallel news corpus and extension of domain adaptation for NMT. In *Proceedings of the 12th Language Resources and Evaluation Conference*, pp. 3616–3622, Marseille, France. European Language Resources Association.
- Moore, R. C., & Lewis, W. (2010). Intelligent selection of language model training data. In *Proceedings of the ACL 2010 Conference Short Papers*, pp. 220–224, Uppsala, Sweden. Association for Computational Linguistics.
- Moryossef, A., Aharoni, R., & Goldberg, Y. (2019). Filling gender & number gaps in neural machine translation with black-box context injection. In *Proceedings of the*

- First Workshop on Gender Bias in Natural Language Processing*, pp. 49–54, Florence, Italy. Association for Computational Linguistics.
- Moslem, Y., Haque, R., Kelleher, J., & Way, A. (2022). Domain-specific text generation for machine translation. In *Proceedings of the 15th biennial conference of the Association for Machine Translation in the Americas (Volume 1: Research Track)*, pp. 14–30, Orlando, USA. Association for Machine Translation in the Americas.
- Mueller, A., & Lal, Y. K. (2019). Sentence-level adaptation for low-resource neural machine translation. In *Proceedings of the 2nd Workshop on Technologies for MT of Low Resource Languages*, pp. 39–47, Dublin, Ireland. European Association for Machine Translation.
- Mueller, A., Nicolai, G., McCarthy, A. D., Lewis, D., Wu, W., & Yarowsky, D. (2020). An analysis of massively multilingual neural machine translation for low-resource languages. In *Proceedings of the 12th Language Resources and Evaluation Conference*, pp. 3710–3718, Marseille, France. European Language Resources Association.
- Müller, M., Rios, A., & Sennrich, R. (2020). Domain robustness in neural machine translation. In *Proceedings of the 14th Conference of the Association for Machine Translation in the Americas (AMTA 2020)*, pp. 151–164, Virtual. Association for Machine Translation in the Americas.
- Neel, S., & Sharifi-Malvajerdi, S. (2021). Descent-to-delete: Gradient-based methods for machine unlearning. In Feldman, V., Ligett, K., & Sabato, S. (Eds.), *Algorithmic Learning Theory, 16-19 March 2021, Virtual Conference, Worldwide*, Vol. 132 of *Proceedings of Machine Learning Research*, pp. 931–962. PMLR.
- Neubig, G. (2016). Lexicons and minimum risk training for neural machine translation: NAIST-CMU at WAT2016. In *Proceedings of the 3rd Workshop on Asian Translation (WAT2016)*, pp. 119–125, Osaka, Japan. The COLING 2016 Organizing Committee.
- Neubig, G. (2017). Neural machine translation and sequence-to-sequence models: A tutorial. *CoRR*, *abs/1703.01619*.
- Neubig, G., & Hu, J. (2018). Rapid adaptation of neural machine translation to new languages. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pp. 875–880, Brussels, Belgium. Association for Computational Linguistics.
- Nguyen, T., & Chiang, D. (2018). Improving lexical choice in neural machine translation. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pp. 334–343, New Orleans, Louisiana. Association for Computational Linguistics.
- Nguyen, T. Q., & Chiang, D. (2017). Transfer learning across low-resource, related languages for neural machine translation. In *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pp. 296–301, Taipei, Taiwan. Asian Federation of Natural Language Processing.
- Niehues, J. (2021). Continuous learning in neural machine translation using bilingual dictionaries. In *Proceedings of the 16th Conference of the European Chapter of the Associ-*

- ation for *Computational Linguistics: Main Volume*, pp. 830–840, Online. Association for Computational Linguistics.
- Ott, M., Auli, M., Grangier, D., & Ranzato, M. (2018). Analyzing uncertainty in neural machine translation. In Dy, J. G., & Krause, A. (Eds.), *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018*, Vol. 80 of *Proceedings of Machine Learning Research*, pp. 3953–3962. PMLR.
- Papineni, K., Roukos, S., Ward, T., & Zhu, W.-J. (2002). Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pp. 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Park, C., Tae, Y., Kim, T., Yang, S., Khan, M. A., Park, L., & Choo, J. (2021). Unsupervised neural machine translation for low-resource domains via meta-learning. In Zong, C., Xia, F., Li, W., & Navigli, R. (Eds.), *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing, ACL/IJCNLP 2021, (Volume 1: Long Papers), Virtual Event, August 1-6, 2021*, pp. 2888–2901. Association for Computational Linguistics.
- Parthasarathy, V., Ramesh, A., Haque, R., & Way, A. (2020). The ADAPT system description for the WMT20 news translation task. In *Proceedings of the Fifth Conference on Machine Translation*, pp. 262–268, Online. Association for Computational Linguistics.
- Peng, W., Huang, C., Li, T., Chen, Y., & Liu, Q. (2020). Dictionary-based data augmentation for cross-domain neural machine translation. *CoRR*, *abs/2004.02577*.
- Pereyra, G., Tucker, G., Chorowski, J., Kaiser, L., & Hinton, G. E. (2017). Regularizing neural networks by penalizing confident output distributions. In *5th International Conference on Learning Representations, ICLR*. OpenReview.net.
- Pham, M. Q., Crego, J.-M., Yvon, F., & Senellart, J. (2019). Generic and specialized word embeddings for multi-domain machine translation. In *International Workshop on Spoken Language Translation*.
- Pham, M. Q., Crego, J. M., Yvon, F., & Senellart, J. (2020a). A study of residual adapters for multi-domain neural machine translation. In *Proceedings of the Fifth Conference on Machine Translation*, pp. 617–628, Online. Association for Computational Linguistics.
- Pham, M. Q., Xu, J., Crego, J., Yvon, F., & Senellart, J. (2020b). Priming neural machine translation. In *Proceedings of the Fifth Conference on Machine Translation*, pp. 516–527, Online. Association for Computational Linguistics.
- Pham, M. Q., Yvon, F., & Crego, J.-M. (2022). Multi-domain adaptation in neural machine translation with dynamic sampling strategies. In *European Association of Machine Translation*, pp. 13–22.
- Pham, M., Crego, J. M., & Yvon, F. (2021). Revisiting multi-domain machine translation. *Transactions of the Association for Computational Linguistics*, *9*, 17–35.

- Philip, J., Berard, A., Gallé, M., & Besacier, L. (2020). Monolingual adapters for zero-shot neural machine translation. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 4465–4470, Online. Association for Computational Linguistics.
- Platanios, E. A., Stretcu, O., Neubig, G., Poczos, B., & Mitchell, T. (2019). Competence-based curriculum learning for 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)*, pp. 1162–1172, Minneapolis, Minnesota. Association for Computational Linguistics.
- Poncelas, A., de Buy Wenniger, G. M., & Way, A. (2019a). Adaptation of machine translation models with back-translated data using transductive data selection methods. In *CICLing 2019 - 20th International Conference on Computational Linguistics and Intelligent Text Processing*.
- Poncelas, A., Maillette de Buy Wenniger, G., & Way, A. (2019b). Transductive data-selection algorithms for fine-tuning neural machine translation. In *Proceedings of The 8th Workshop on Patent and Scientific Literature Translation*, pp. 13–23, Dublin, Ireland. European Association for Machine Translation.
- Poncelas, A., Shterionov, D., Way, A., de Buy Wenniger, G. M., & Passban, P. (2018). Investigating backtranslation in neural machine translation. In *Proceedings of the 21st Annual Conference of the European Association for Machine Translation: 28-30 May 2018, Universitat d’Alacant, Alacant, Spain*, pp. 249–258. European Association for Machine Translation.
- Poncelas, A., & Way, A. (2019). Selecting artificially-generated sentences for fine-tuning neural machine translation. In *Proceedings of the 12th International Conference on Natural Language Generation*, pp. 219–228, Tokyo, Japan. Association for Computational Linguistics.
- Poncelas, A., Wenniger, G. M. d. B., & Way, A. (2018). Data selection with feature decay algorithms using an approximated target side. In *IWSLT 2018, International Workshop on Spoken Language Translation*.
- Prates, M. O., Avelar, P. H., & Lamb, L. C. (2020). Assessing gender bias in machine translation: a case study with google translate. *Neural Computing and Applications*, 32(10), 6363–6381.
- Ramponi, A., & Plank, B. (2020). Neural unsupervised domain adaptation in NLP—A survey. In *Proceedings of the 28th International Conference on Computational Linguistics*, pp. 6838–6855, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Ranzato, M., Chopra, S., Auli, M., & Zaremba, W. (2016). Sequence level training with recurrent neural networks. In *ICLR*.
- Ratcliff, R. (1990). Connectionist models of recognition memory: constraints imposed by learning and forgetting functions.. *Psychological review*, 97(2), 285.
- Rokach, L. (2010). Ensemble-based classifiers. *Artificial Intelligence Review*, 33(1-2), 1–39.

- Rosti, A.-V., Ayan, N. F., Xiang, B., Matsoukas, S., Schwartz, R., & Dorr, B. (2007). Combining outputs from multiple machine translation systems. In *Human Language Technologies 2007: The Conference of the North American Chapter of the Association for Computational Linguistics; Proceedings of the Main Conference*, pp. 228–235, Rochester, New York. Association for Computational Linguistics.
- Rubino, R., Marie, B., Dabre, R., Fujita, A., Utiyama, M., & Sumita, E. (2020). Extremely low-resource neural machine translation for asian languages. *Machine Translation*, 34(4), 347–382.
- Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1986). Learning representations by back-propagating errors. *Nature*, 323(6088), 533–536.
- Sajjad, H., Durrani, N., Dalvi, F., Belinkov, Y., & Vogel, S. (2017). Neural Machine Translation training in a multi-domain scenario. In *IWSLT 2017, International Workshop on Spoken Language Translation*.
- Salesky, E., Runge, A., Coda, A., Niehues, J., & Neubig, G. (2020). Optimizing segmentation granularity for neural machine translation. *Machine Translation*, 34(1), 41–59.
- Santamaría, L., & Axelrod, A. (2017). Data selection with cluster-based language difference models and cynical selection. In *IWSLT 2017, International Workshop on Spoken Language Translation*.
- Santini, M. (2004). State-of-the-art on automatic genre identification. Tech. rep., University of Brighton.
- Sato, S., Sakuma, J., Yoshinaga, N., Toyoda, M., & Kitsuregawa, M. (2020). Vocabulary adaptation for domain adaptation in neural machine translation. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pp. 4269–4279, Online. Association for Computational Linguistics.
- Saunders, D., & Byrne, B. (2020a). Addressing Exposure Bias With Document Minimum Risk Training: Cambridge at the WMT20 Biomedical Translation Task. In *Proceedings of the Fifth Conference on Machine Translation*, pp. 862–869. Association for Computational Linguistics.
- Saunders, D., & Byrne, B. (2020b). Reducing gender bias in neural machine translation as a domain adaptation problem. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 7724–7736, Online. Association for Computational Linguistics.
- Saunders, D., Sallis, R., & Byrne, B. (2020). Neural machine translation doesn’t translate gender coreference right unless you make it. In *Proceedings of the Second Workshop on Gender Bias in Natural Language Processing*, pp. 35–43. Association for Computational Linguistics.
- Saunders, D., Stahlberg, F., & Byrne, B. (2019). UCAM biomedical translation at WMT19: Transfer learning multi-domain ensembles. In *Proceedings of the Fourth Conference on Machine Translation (Volume 3: Shared Task Papers, Day 2)*, pp. 169–174, Florence, Italy. Association for Computational Linguistics.
- Saunders, D., Stahlberg, F., & Byrne, B. (2020). Using context in neural machine translation training objectives. In *Proceedings of the 58th Annual Meeting of the Association*

- for *Computational Linguistics*, pp. 7764–7770, Online. Association for Computational Linguistics.
- Saunders, D., Stahlberg, F., de Gispert, A., & Byrne, B. (2018). Multi-representation ensembles and delayed SGD updates improve syntax-based NMT. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pp. 319–325, Melbourne, Australia. Association for Computational Linguistics.
- Saunders, D., Stahlberg, F., de Gispert, A., & Byrne, B. (2019). Domain adaptive inference for neural machine translation. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 222–228, Florence, Italy. Association for Computational Linguistics.
- Savenkov, K. (2018). State of the domain-adaptive machine translation.. (accessed: Nov 2020).
- Savoldi, B., Gaido, M., Bentivogli, L., Negri, M., & Turchi, M. (2021). Gender Bias in Machine Translation. *Transactions of the Association for Computational Linguistics*, 9, 845–874.
- Schamper, J., Rosendahl, J., Bahar, P., Kim, Y., Nix, A., & Ney, H. (2018). The RWTH Aachen University supervised machine translation systems for WMT 2018. In *Proceedings of the Third Conference on Machine Translation: Shared Task Papers*, pp. 496–503, Belgium, Brussels. Association for Computational Linguistics.
- Schwenk, H. (2008). Investigations on large-scale lightly-supervised training for statistical machine translation. In *2008 International Workshop on Spoken Language Translation, IWSLT 2008, Honolulu, Hawaii, USA, October 20-21, 2008*, pp. 182–189. ISCA.
- Sennrich, R., Birch, A., Currey, A., Gehrmann, U., Haddow, B., Heafield, K., Barone, A. V. M., & Williams, P. (2017). The university of Edinburgh’s neural MT systems for WMT17. In *Proceedings of the Second Conference on Machine Translation*, pp. 389–399.
- Sennrich, R., & Haddow, B. (2016). Linguistic input features improve neural machine translation. In *Proceedings of the First Conference on Machine Translation: Volume 1, Research Papers*, pp. 83–91, Berlin, Germany. Association for Computational Linguistics.
- Sennrich, R., Haddow, B., & Birch, A. (2016a). Controlling politeness in neural machine translation via side constraints. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pp. 35–40, San Diego, California. Association for Computational Linguistics.
- Sennrich, R., Haddow, B., & Birch, A. (2016b). Improving Neural Machine Translation Models with Monolingual Data. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*, Vol. 1, pp. 86–96.
- Sennrich, R., Haddow, B., & Birch, A. (2016c). Neural machine translation of rare words with subword units. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*, Vol. 1, pp. 1715–1725.

- Sennrich, R., & Zhang, B. (2019). Revisiting low-resource neural machine translation: A case study. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 211–221, Florence, Italy. Association for Computational Linguistics.
- Servan, C., Crego, J. M., & Senellart, J. (2016). Domain specialization: a post-training domain adaptation for neural machine translation. *CoRR*, *abs/1612.06141*.
- Shah, D. S., Schwartz, H. A., & Hovy, D. (2020). Predictive biases in natural language processing models: A conceptual framework and overview. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 5248–5264, Online. Association for Computational Linguistics.
- Sharaf, A., Hassan, H., & Daumé III, H. (2020). Meta-learning for few-shot NMT adaptation. In *Proceedings of the Fourth Workshop on Neural Generation and Translation*, pp. 43–53, Online. Association for Computational Linguistics.
- Shen, S., Cheng, Y., He, Z., He, W., Wu, H., Sun, M., & Liu, Y. (2016). Minimum Risk Training for Neural Machine Translation. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*, Vol. 1, pp. 1683–1692.
- Sim, K. C., Byrne, W. J., Gales, M. J., Sahbi, H., & Woodland, P. C. (2007). Consensus network decoding for statistical machine translation system combination. In *2007 IEEE International Conference on Acoustics, Speech and Signal Processing-ICASSP'07*, Vol. 4, pp. IV–105. IEEE.
- Sinclair, J., & Ball, J. (1996). Preliminary recommendations on text typology. Tech. rep., EAGLES (Expert Advisory Group on Language Engineering Standards).
- Song, H., Dabre, R., Fujita, A., & Kurohashi, S. (2020). Domain adaptation of neural machine translation through multistage fine-tuning. In *26th Annual Conference of the Association for Natural Language Processing*.
- Song, K., Zhang, Y., Yu, H., Luo, W., Wang, K., & Zhang, M. (2019). Code-switching for enhancing NMT with pre-specified 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)*, pp. 449–459, Minneapolis, Minnesota. Association for Computational Linguistics.
- Song, Z., Ma, Z., Qi, K., & Liu, G. (2021). Dynamic tuning and weighting of meta-learning for NMT domain adaptation. In Farkas, I., Masulli, P., Otte, S., & Wermter, S. (Eds.), *Artificial Neural Networks and Machine Learning - ICANN 2021 - 30th International Conference on Artificial Neural Networks, Bratislava, Slovakia, September 14-17, 2021, Proceedings, Part V*, Vol. 12895 of *Lecture Notes in Computer Science*, pp. 576–587. Springer.
- Sperber, M., Niehues, J., & Waibel, A. (2017). Toward robust neural machine translation for noisy input sequences. In *IWSLT 2017, International Workshop on Spoken Language Translation*.
- Sriram, A., Jun, H., Satheesh, S., & Coates, A. (2018). Cold fusion: Training seq2seq models together with language models. In *Proc. Interspeech 2018*, pp. 387–391.
- Stahlberg, F. (2020). Neural machine translation: A review. *Journal of Artificial Intelligence Research*, *69*, 343–418.

- Stahlberg, F., & Byrne, B. (2017). Unfolding and shrinking neural machine translation ensembles. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pp. 1946–1956, Copenhagen, Denmark. Association for Computational Linguistics.
- Stahlberg, F., & Byrne, B. (2019). On NMT search errors and model errors: Cat got your tongue?. 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)*, pp. 3356–3362, Hong Kong, China. Association for Computational Linguistics.
- Stahlberg, F., Cross, J., & Stoyanov, V. (2018a). Simple fusion: Return of the language model. In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pp. 204–211, Brussels, Belgium. Association for Computational Linguistics.
- Stahlberg, F., de Gispert, A., & Byrne, B. (2018b). The university of Cambridge’s machine translation systems for WMT18. In *Proceedings of the Third Conference on Machine Translation: Shared Task Papers*, pp. 504–512, Belgium, Brussels. Association for Computational Linguistics.
- Stahlberg, F., de Gispert, A., Hasler, E., & Byrne, B. (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*, pp. 362–368, Valencia, Spain. Association for Computational Linguistics.
- Stahlberg, F., Hasler, E., Waite, A., & Byrne, B. (2016). Syntactically guided neural machine translation. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pp. 299–305, Berlin, Germany. Association for Computational Linguistics.
- Stahlberg, F., Saunders, D., de Gispert, A., & Byrne, B. (2019). CUED@WMT19:EWC&LMs. In *Proceedings of the Fourth Conference on Machine Translation (Volume 2: Shared Task Papers, Day 1)*, pp. 364–373, Florence, Italy. Association for Computational Linguistics.
- Stanovsky, G., Smith, N. A., & Zettlemoyer, L. (2019). Evaluating gender bias in machine translation. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 1679–1684, Florence, Italy. Association for Computational Linguistics.
- Stergiadis, E., Kumar, S., Kovalev, F., & Levin, P. (2021). Multi-domain adaptation in neural machine translation through multidimensional tagging. In *Proceedings of Machine Translation Summit XVIII: Users and Providers Track*, pp. 396–420, Virtual. Association for Machine Translation in the Americas.
- Stojanovski, D., & Fraser, A. (2019). Combining local and document-level context: The LMU Munich neural machine translation system at WMT19. In *Proceedings of the Fourth Conference on Machine Translation (Volume 2: Shared Task Papers, Day 1)*, pp. 400–406, Florence, Italy. Association for Computational Linguistics.

- Stojanovski, D., & Fraser, A. M. (2020). Addressing zero-resource domains using document-level context in neural machine translation. *CoRR*, *abs/2004.14927*.
- Strubell, E., Ganesh, A., & McCallum, A. (2019). Energy and policy considerations for deep learning in NLP. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 3645–3650, Florence, Italy. Association for Computational Linguistics.
- Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., & Wojna, Z. (2016). Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 2818–2826.
- Taghipour, K., Khadivi, S., & Xu, J. (2011). Parallel corpus refinement as an outlier detection algorithm. In *Proceedings of the 13th Machine Translation Summit (MT Summit XIII)*, pp. 414–421.
- Tan, S., Joty, S., Kan, M.-Y., & Socher, R. (2020). It’s morphin’ time! Combating linguistic discrimination with inflectional perturbations. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 2920–2935, Online. Association for Computational Linguistics.
- Tars, S., & Fishel, M. (2018). Multi-domain neural machine translation. In *Proceedings of the 21st Annual Conference of the European Association for Machine Translation: 28-30 May 2018, Universitat d’Alacant, Alacant, Spain*, pp. 259–268. European Association for Machine Translation.
- Tezcan, A., Bulté, B., & Vanroy, B. (2021). Towards a better integration of fuzzy matches in neural machine translation through data augmentation. *Informatics*, *8*(1), 7.
- Thompson, B., Gwinnup, J., Khayrallah, H., Duh, K., & Koehn, P. (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*. Association for Computational Linguistics.
- Thompson, B., Khayrallah, H., Anastasopoulos, A., McCarthy, A. D., Duh, K., Marvin, R., McNamee, P., Gwinnup, J., Anderson, T., & Koehn, P. (2018). Freezing subnetworks to analyze domain adaptation in Neural Machine Translation. In *Proceedings of the Third Conference on Machine Translation*, pp. 124–132.
- Tiedemann, J., Scherrer, Y., et al. (2017). Neural machine translation with extended context. In *Proceedings of the Third Workshop on Discourse in Machine Translation*. Association for Computational Linguistics.
- Tomalin, M., Byrne, B., Concannon, S., Saunders, D., & Ullmann, S. (2021). The practical ethics of bias reduction in machine translation: Why domain adaptation is better than data debiasing. *Ethics and Information Technology*, *23*(1).
- Tromble, R., Kumar, S., Och, F., & Macherey, W. (2008). Lattice Minimum Bayes-Risk decoding for statistical machine translation. In *Proceedings of the 2008 Conference on Empirical Methods in Natural Language Processing*, pp. 620–629, Honolulu, Hawaii. Association for Computational Linguistics.

- Tu, Z., Lu, Z., Liu, Y., Liu, X., & Li, H. (2016). Modeling coverage for neural machine translation. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 76–85, Berlin, Germany. Association for Computational Linguistics.
- Turian, J., Ratinov, L.-A., & Bengio, Y. (2010). Word representations: A simple and general method for semi-supervised learning. In *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*, pp. 384–394, Uppsala, Sweden. Association for Computational Linguistics.
- Ul Haq, S., Abdul Rauf, S., Shoukat, A., & Hira, N.-e. (2020). Improving document-level neural machine translation with domain adaptation. In *Proceedings of the Fourth Workshop on Neural Generation and Translation*, pp. 225–231, Online. Association for Computational Linguistics.
- Vaibhav, V., Singh, S., Stewart, C., & Neubig, G. (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)*, pp. 1916–1920, Minneapolis, Minnesota. Association for Computational Linguistics.
- van der Wees, M. (2017). *What’s in a domain?: Towards fine-grained adaptation for machine translation*. Ph.D. thesis, University of Amsterdam.
- van der Wees, M., Bisazza, A., & Monz, C. (2017). Dynamic data selection for neural machine translation. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pp. 1400–1410, Copenhagen, Denmark. Association for Computational Linguistics.
- Vanmassenhove, E., Hardmeier, C., & Way, A. (2018). Getting gender right in neural machine translation. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pp. 3003–3008, Brussels, Belgium. Association for Computational Linguistics.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., & Polosukhin, I. (2017). Attention is all you need. In *Advances in Neural Information Processing Systems*, pp. 6000–6010.
- Vaswani, A., Zhao, Y., Fossum, V., & Chiang, D. (2013). Decoding with large-scale neural language models improves translation. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pp. 1387–1392, Seattle, Washington, USA. Association for Computational Linguistics.
- Vijayakumar, A. K., Cogswell, M., Selvaraju, R. R., Sun, Q., Lee, S., Crandall, D., & Batra, D. (2016). Diverse beam search: Decoding diverse solutions from neural sequence models. *CoRR*, [abs/1610.02424](https://arxiv.org/abs/1610.02424).
- Vilar, D. (2018). Learning hidden unit contribution for adapting neural machine translation models. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, Vol. 2, pp. 500–505.

- Voita, E., Sennrich, R., & Titov, I. (2019a). Context-aware monolingual repair 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)*, pp. 877–886, Hong Kong, China. Association for Computational Linguistics.
- Voita, E., Sennrich, R., & Titov, I. (2019b). When a good translation is wrong in context: Context-aware machine translation improves on deixis, ellipsis, and lexical cohesion. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 1198–1212, Florence, Italy. Association for Computational Linguistics.
- Voita, E., Serdyukov, P., Sennrich, R., & Titov, I. (2018). Context-aware neural machine translation learns anaphora resolution. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 1264–1274, Melbourne, Australia. Association for Computational Linguistics.
- Vu, T., & Moschitti, A. (2021). Machine translation customization via automatic training data selection from the web. In Hiemstra, D., Moens, M., Mothe, J., Perego, R., Potthast, M., & Sebastiani, F. (Eds.), *Advances in Information Retrieval - 43rd European Conference on IR Research, ECIR 2021, Virtual Event, March 28 - April 1, 2021, Proceedings, Part I*, Vol. 12656 of *Lecture Notes in Computer Science*, pp. 666–679. Springer.
- Vu, T.-T., He, X., Phung, D., & Haffari, G. (2021). Generalised unsupervised domain adaptation of neural machine translation with cross-lingual data selection. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pp. 3335–3346, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Vu, T.-T., Khadivi, S., Phung, D., & Haffari, G. (2022). Domain generalisation of NMT: Fusing adapters with leave-one-domain-out training. In *Findings of the Association for Computational Linguistics: ACL 2022*, pp. 582–588, Dublin, Ireland. Association for Computational Linguistics.
- Wallace, E., Stern, M., & Song, D. (2020). Imitation attacks and defenses for black-box machine translation systems. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 5531–5546, Online. Association for Computational Linguistics.
- Wang, C., & Sennrich, R. (2020). On exposure bias, hallucination and domain shift in neural machine translation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 3544–3552, Online. Association for Computational Linguistics.
- Wang, K., Chen, G., Huang, Z., Wan, X., & Huang, F. (2021). Bridging the domain gap: Improve informal language translation via counterfactual domain adaptation. In *Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2021, Thirty-Third Conference on Innovative Applications of Artificial Intelligence, IAAI 2021, The Eleventh Symposium on Educational Advances in Artificial Intelligence, EAAI 2021, Virtual Event, February 2-9, 2021*, pp. 13970–13978. AAAI Press.

- Wang, L., Tu, Z., Way, A., & Liu, Q. (2017). Exploiting cross-sentence context for neural machine translation. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pp. 2826–2831, Copenhagen, Denmark. Association for Computational Linguistics.
- Wang, Q., Li, B., Xiao, T., Zhu, J., Li, C., Wong, D. F., & Chao, L. S. (2019). Learning deep transformer models for machine translation. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 1810–1822, Florence, Italy. Association for Computational Linguistics.
- Wang, R., Finch, A., Utiyama, M., & Sumita, E. (2017a). Sentence embedding for Neural Machine Translation domain adaptation. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, Vol. 2, pp. 560–566.
- Wang, R., Utiyama, M., Liu, L., Chen, K., & Sumita, E. (2017b). Instance weighting for neural machine translation domain adaptation. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pp. 1482–1488, Copenhagen, Denmark. Association for Computational Linguistics.
- Wang, R., Utiyama, M., & Sumita, E. (2018). Dynamic sentence sampling for efficient training of neural machine translation. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pp. 298–304, Melbourne, Australia. Association for Computational Linguistics.
- Wang, W., Tian, Y., Ngiam, J., Yang, Y., Caswell, I., & Parekh, Z. (2020). Learning a multi-domain curriculum for neural machine translation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 7711–7723, Online. Association for Computational Linguistics.
- Wang, W., Watanabe, T., Hughes, M., Nakagawa, T., & Chelba, C. (2018). Denoising neural machine translation training with trusted data and online data selection. In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pp. 133–143, Belgium, Brussels. Association for Computational Linguistics.
- Wang, Y., Wang, L., Shi, S., Li, V. O. K., & Tu, Z. (2020). Go from the general to the particular: Multi-domain translation with domain transformation networks. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI*, pp. 9233–9241. AAAI Press.
- Wei, H.-R., Zhang, Z., Chen, B., & Luo, W. (2020a). Iterative domain-repaired back-translation. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 5884–5893, Online. Association for Computational Linguistics.
- Wei, X., Yu, H., Hu, Y., Zhang, Y., Weng, R., & Luo, W. (2020b). Multiscale collaborative deep models for neural machine translation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 414–426, Online. Association for Computational Linguistics.

- Weinshall, D., Cohen, G., & Amir, D. (2018). Curriculum learning by transfer learning: Theory and experiments with deep networks. In *International Conference on Machine Learning*, pp. 5238–5246.
- Wieting, J., Berg-Kirkpatrick, T., Gimpel, K., & Neubig, G. (2019). Beyond BLEU: training neural machine translation with semantic similarity. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 4344–4355, Florence, Italy. Association for Computational Linguistics.
- Wu, M., Li, Y., Zhang, M., Li, L., Haffari, G., & Liu, Q. (2021). Uncertainty-aware balancing for multilingual and multi-domain neural machine translation training. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pp. 7291–7305, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Wu, S., Zhang, D., & Zhou, M. (2019). Effective soft-adaptation for neural machine translation. In Tang, J., Kan, M., Zhao, D., Li, S., & Zan, H. (Eds.), *Natural Language Processing and Chinese Computing - 8th CCF International Conference, NLPCC 2019, Dunhuang, China, October 9-14, 2019, Proceedings, Part II*, Vol. 11839 of *Lecture Notes in Computer Science*, pp. 254–264. Springer.
- Wu, Y., Schuster, M., Chen, Z., Le, Q. V., Norouzi, M., Macherey, W., Krikun, M., Cao, Y., Gao, Q., Macherey, K., et al. (2016). Google’s neural machine translation system: Bridging the gap between human and machine translation. *CoRR*, *abs/1609.08144*.
- Wuebker, J., Simianer, P., & DeNero, J. (2018). Compact personalized models for neural machine translation. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pp. 881–886, Brussels, Belgium. Association for Computational Linguistics.
- Xie, G., & Way, A. (2020). Constraining the transformer NMT model with heuristic grid beam search. In *Proceedings of the 14th Conference of the Association for Machine Translation in the Americas (Volume 1: Research Track)*, pp. 36–49, Virtual. Association for Machine Translation in the Americas.
- Xu, J., Crego, J., & Senellart, J. (2019). Lexical micro-adaptation for neural machine translation. In *International Workshop on Spoken Language Translation*.
- Xu, J., Crego, J., & Senellart, J. (2020). Boosting neural machine translation with similar translations. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 1580–1590, Online. Association for Computational Linguistics.
- Zeng, J., Su, J., Wen, H., Liu, Y., Xie, J., Yin, Y., & Zhao, J. (2018). Multi-domain neural machine translation with word-level domain context discrimination. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pp. 447–457, Brussels, Belgium. Association for Computational Linguistics.
- Zhan, R., Liu, X., Wong, D. F., & Chao, L. S. (2021). Meta-curriculum learning for domain adaptation in neural machine translation. In *Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2021, Thirty-Third Conference on Innovative Applications of Artificial Intelligence, IAAI 2021, The Eleventh Symposium on Educational*

Advances in Artificial Intelligence, EAAI 2021, Virtual Event, February 2-9, 2021, pp. 14310–14318. AAAI Press.

- Zhang, B., Williams, P., Titov, I., & Sennrich, R. (2020). Improving massively multilingual neural machine translation and zero-shot translation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 1628–1639, Online. Association for Computational Linguistics.
- Zhang, D., Kim, J., Crego, J., & Senellart, J. (2017). Boosting neural machine translation. In *Proceedings of the Eighth International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pp. 271–276, Taipei, Taiwan. Asian Federation of Natural Language Processing.
- Zhang, J., & Zong, C. (2016a). Bridging neural machine translation and bilingual dictionaries. *CoRR*, *abs/1610.07272*.
- Zhang, J., & Zong, C. (2016b). Exploiting source-side monolingual data in neural machine translation. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pp. 1535–1545, Austin, Texas. Association for Computational Linguistics.
- Zhang, J., Utiyama, M., Sumita, E., Neubig, G., & Nakamura, S. (2018). Guiding neural machine translation with retrieved translation pieces. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pp. 1325–1335, New Orleans, Louisiana. Association for Computational Linguistics.
- Zhang, M., & Toral, A. (2019). The effect of translationese in machine translation test sets. In *Proceedings of the Fourth Conference on Machine Translation (Volume 1: Research Papers)*, pp. 73–81, Florence, Italy. Association for Computational Linguistics.
- Zhang, S., & Xiong, D. (2018). Sentence weighting for neural machine translation domain adaptation. In *Proceedings of the 27th International Conference on Computational Linguistics*, pp. 3181–3190, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Zhang, X., Kumar, G., Khayrallah, H., Murray, K., Gwinnup, J., Martindale, M. J., McNamee, P., Duh, K., & Carpuat, M. (2018). An empirical exploration of curriculum learning for neural machine translation. *CoRR*, *abs/1811.00739*.
- Zhang, X., Shapiro, P., Kumar, G., McNamee, P., Carpuat, M., & Duh, K. (2019a). Curriculum learning for domain adaptation in 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)*, pp. 1903–1915, Minneapolis, Minnesota. Association for Computational Linguistics.
- Zhang, Z., Wu, S., Liu, S., Li, M., Zhou, M., & Xu, T. (2019b). Regularizing neural machine translation by target-bidirectional agreement. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 33, pp. 443–450.
- Zhao, M., Wu, H., Niu, D., & Wang, X. (2020). Reinforced curriculum learning on pre-trained neural machine translation models. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of*

Artificial Intelligence Conference, IAAI 2020, The Tenth AAI Symposium on Educational Advances in Artificial Intelligence, EAAI, pp. 9652–9659. AAAI Press.

- Zhao, Y., Zhang, J., He, Z., Zong, C., & Wu, H. (2018). Addressing troublesome words in neural machine translation. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pp. 391–400, Brussels, Belgium. Association for Computational Linguistics.
- Zheng, X., Zhang, Z., Guo, J., Huang, S., Chen, B., Luo, W., & Chen, J. (2021a). Adaptive nearest neighbor machine translation. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pp. 368–374, Online. Association for Computational Linguistics.
- Zheng, X., Zhang, Z., Huang, S., Chen, B., Xie, J., Luo, W., & Chen, J. (2021b). Non-parametric unsupervised domain adaptation for neural machine translation. In *Findings of the Association for Computational Linguistics: EMNLP 2021, Virtual Event / Punta Cana, Dominican Republic, 16-20 November, 2021*, pp. 4234–4241. Association for Computational Linguistics.
- Zipf, G. K. (1949). *Human behavior and the principle of least effort*. Addison-Wesley Press.
- Zoph, B., Yuret, D., May, J., & Knight, K. (2016). Transfer learning for low-resource neural machine translation. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pp. 1568–1575.