

Bilingual Synchronization: Restoring Translational Relationships with Editing Operations

Jitao Xu[†]

Josep Crego[‡]

François Yvon[†]

[†]Université Paris-Saclay, CNRS, LISN, 91400, Orsay, France

[‡]SYSTRAN, 5 rue Feydeau, 75002, Paris, France

{jitao.xu, francois.yvon}@lmsi.fr, josep.crego@systrangroup.com

Abstract

Machine Translation (MT) is usually viewed as a one-shot process that generates the target language equivalent of some source text from scratch. We consider here a more general setting which assumes an initial target sequence, that must be transformed into a valid translation of the source, thereby restoring parallelism between source and target. For this *bilingual synchronization task*, we consider several architectures (both autoregressive and non-autoregressive) and training regimes, and experiment with multiple practical settings such as simulated interactive MT, translating with Translation Memory (TM) and TM cleaning. Our results suggest that *one single generic edit-based system*, once fine-tuned, can compare with, or even outperform, dedicated systems specifically trained for these tasks.

1 Introduction

Neural Machine Translation (NMT) systems have made tangible progress in recent years (Bahdanau et al., 2015; Vaswani et al., 2017), as they started to produce usable translations in production environments. NMT is generally viewed as a one-shot activity process in autoregressive approaches, which generates the target translation based on the sole source side input. Recently, Non-autoregressive Machine Translation (NAT) models have proposed to perform iterative refinement decoding (Lee et al., 2018; Ghazvininejad et al., 2019; Gu et al., 2019), where translations are generated through an iterative revision process, starting with a possibly empty initial hypothesis.

This paper focuses on the revision part of the machine translation (MT) process and consider *bilingual synchronization* (Bi-sync), which we define as follows: given a pair of a source (f) and a target (\tilde{e}) sentences, *which may or may not be mutual translations*, the task is to compute a revised version e of \tilde{e} , such that e is an actual translation of

f . This is necessary when the source side of an existing translation is edited, requiring to update the target and keep both sides synchronized. Bi-sync subsumes standard MT, where the synchronization starts with an empty target ($\tilde{e} = []$). Other interesting cases occur when parts of the initial target can be reused, so that the synchronization only requires a few changes.

Bi-sync encompasses several tasks: synchronization is needed in interactive MT (IMT, Knowles and Koehn, 2016) and bilingual editing (Bronner et al., 2012), with \tilde{e} the translation of a previous version of f ; in MT with lexical constraints (Hokamp and Liu, 2017), where \tilde{e} contains target-side constraints (Susanto et al., 2020; Xu and Carpuat, 2021); in Translation Memory (TM) based approaches (Bulte and Tezcan, 2019), where \tilde{e} is a TM match for a similar example; in automatic post-editing (APE) (do Carmo et al., 2021), where \tilde{e} is an MT output.

We consider here several implementations of sequence-to-sequence models dedicated to these situations, contrasting an autoregressive model with a non-autoregressive approach. The former is similar to Bulte and Tezcan (2019), where the source sentence and the initial translation are concatenated as one input sequence; the latter uses the Levenshtein Transformer (LevT) of Gu et al. (2019). We also study various ways to generate appropriate training samples (f , \tilde{e} , e). Our experiments consider several tasks, including TM cleaning, which attempts to fix and synchronize noisy segments in a parallel corpus. This setting is more difficult than Bi-sync, as many initial translations are already correct and need to be left unchanged. Our results suggest that one single AR system, once fine-tuned, can favorably compare with dedicated systems for each of these tasks. To recap, our main contributions are (a) the generalization of several tasks subsumed by a generic synchronization objective, allowing us to develop a unified perspective about otherwise unrelated subdomains of MT; (b) the design of a

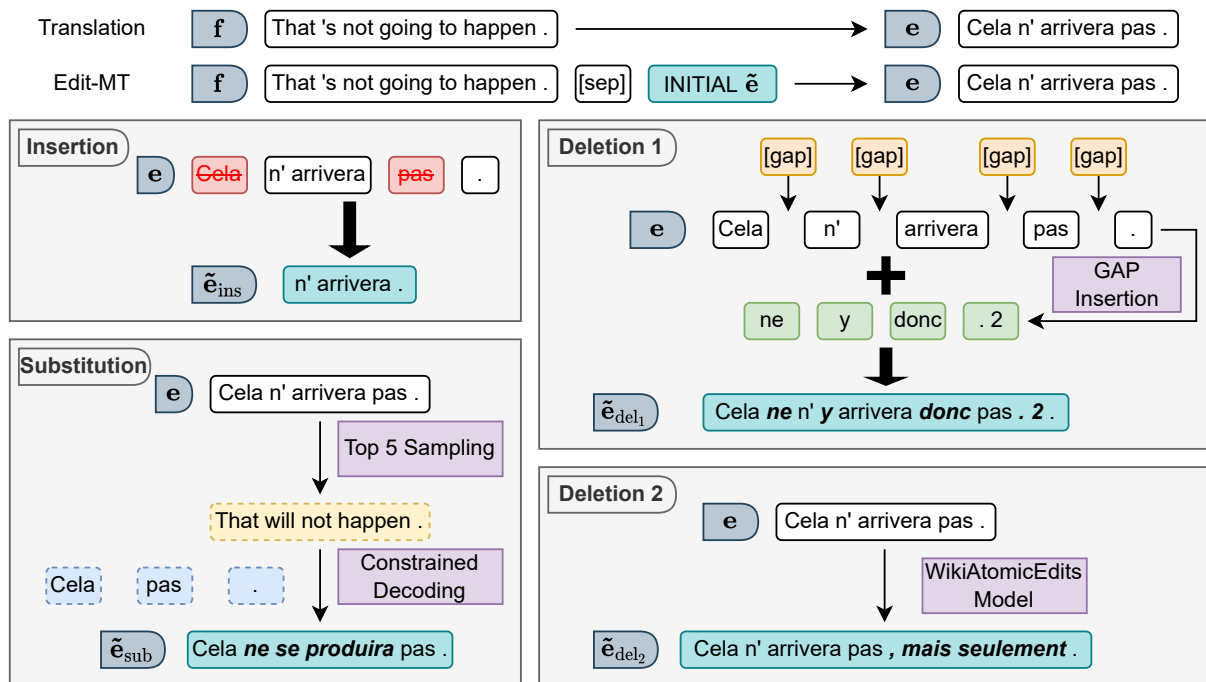


Figure 1: Methods for generating synthetic initial translations \tilde{e} for each edit type. Rectangle purple boxes refer to separate models used to generate the desired operations. Differences in artificial initial translations (in blue boxes) are marked in bold. Initial translations \tilde{e}_{ins} for insertion are generated by randomly removing segments in the reference sentence e . For \tilde{e}_{sub} , e is first back-translated into an intermediate sentence f^* using top-5 sampling, then translated back to \tilde{e}_{sub} with LCD. The first method to generate \tilde{e}_{del_1} randomly inserts [gap] tokens into e and decodes with a GAP insertion model (Xiao et al., 2022). \tilde{e}_{del_1} is obtained by replacing [gap] with the predicted segments. The second method automatically edits e with a model trained on WikiAtomicEdits data.

training procedure for a generic edit-based model; (c) an empirical validation on five settings and domains.

2 Methods

2.1 Generating Editing Data

We consider a general scenario where, given a pair of sentences f and \tilde{e} , assumed to be related, but not necessarily parallel, we aim to generate a target sentence e that is parallel to f . We would also like \tilde{e} and e to be close, as \tilde{e} is often a valid translation of a sentence \tilde{f} that is close to f . Training such models requires triplets (f, \tilde{e}, e) . While large amounts of parallel bilingual data are available for many language pairs, they are hardly ever associated to related translations \tilde{e} (except for APE). We therefore study ways to simulate synthetic \tilde{e} from e , while preserving large portions of e in \tilde{e} . String edits can be decomposed into a sequence of three basic edits (insertions, substitutions and deletions), we design our artificial samples so that edits from \tilde{e} to e only involve one type of operation (Figure 1).

Insertions We mainly follow Xiao et al. (2022) to generate initial translations \tilde{e}_{ins} for insertion by randomly deleting segments from e . For each e , we first randomly sample an integer $k \in [1, 5]$, then randomly remove k non-overlapping segments from e . The length of each removed segment is also randomly sampled with a maximum of 5 tokens. We also impose that the overall ratio of removed segments does not exceed 0.5 of e . Different from Xiao et al. (2022), \tilde{e}_{ins} does not include any placeholders to locate the positions of removed segments. This makes \tilde{e}_{ins} a more realistic starting point as the insertion positions are rarely known in practical settings. Our preliminary experiments also show that identifying insertion positions makes the infilling task easier than when they are unknown.

Substitutions To simulate substitutions, we apply round-trip translation with lexically constrained decoding (LCD, Post and Vilar, 2018) to generate initial translations for substitution \tilde{e}_{sub} . Round-trip translation is already used for the APE task in Junczys-Dowmunt and Grundkiewicz (2016). This requires two standard NMT models separately trained on parallel data, one for each direction. For

each training example (f, e) , we first (a) translate e into an intermediate source sentence f^* using top-5 sampling (Edunov et al., 2018);¹ (b) generate an abbreviated version \tilde{e}'_{ins} using the method described above for insertions. We then translate f^* using LCD, with \tilde{e}'_{ins} as constraints, to obtain \tilde{e}_{sub} . In this way, we ensure that at least half of e remains unchanged in \tilde{e}_{sub} , while the other parts have been substituted. To increase diversity, \tilde{e}'_{ins} (used to create \tilde{e}_{sub}) is sampled with a different random seed than \tilde{e}_{ins} (used for the insertion task).

Deletions Simulating deletions requires the initial translation \tilde{e}_{del} to be an extension of e . We propose two strategies to generate \tilde{e}_{del} . The first uses a GAP insertion model as in Xiao et al. (2022), in which word segments are randomly replaced with a placeholder $[\text{gap}]$ to generate \tilde{e}_{gap} . The task is then to predict the missing segments based on the concatenation of f and \tilde{e}_{gap} as input. This differs from our own insertion task, as (a) insertion positions are identified as a $[\text{gap}]$ symbol in \tilde{e}_{gap} and (b) generation only computes the sequence of missing segments e_{seg} , rather than a complete sentence.

We use GAP to generate extra segments for a pair of parallel sentences as follows. We randomly insert $k \in [1, 5]$ $[\text{gap}]$ tokens into e , concatenate it with f and use GAP to predict extra segments, yielding the synthetic target sentence \tilde{e}_{del_1} . This method always extends parallel sentences with additional segments on the target side. However, these segments are arbitrary and may not contain any valid semantic information, nor be syntactically correct.

We thus consider a second strategy, based on actual edit operations collected in the WikiAtomicEdits dataset² (Faruqui et al., 2018), which contains edits of an original segment x and the resulting segment x' , with exactly one insertion or deletion operation for each example, collected from Wikipedia edit history. This notably ensures that both versions of each utterance are syntactically correct. We treat the deletion data of WikiAtomicEdits as “reversed” insertions, and use both of them to train a seq-to-seq wiki model ($x_{\text{short}} \rightarrow x_{\text{long}}$), generating longer sentences from shorter ones. The wiki model is then used to expand e into an \tilde{e}_{del_2} . Compared to \tilde{e}_{del_1} , \tilde{e}_{del_2} is syntactically more correct. However, it is also by design very close (one edit away) to e .

¹Early experiments showed that using sampling instead of beam search increases the diversity of the generated \tilde{e}_{sub} .

²<https://github.com/google-research-datasets/wiki-atomic-edits>

As both simulation methods have merits and flaws, we randomly select examples from \tilde{e}_{del_1} and \tilde{e}_{del_2} to build the final synthetic initial translation samples for the deletion operation \tilde{e}_{del} .

Copy and Translate Operations To handle parallel sentences that do not require any changes, we add a fourth copy operation, where the initial translation \tilde{e}_{cp} is equal to the target sentence ($\tilde{e}_{\text{cp}}=e$). Hence, the data used to learn edit operations is built with triplets (f, \tilde{e}, e) where \tilde{e} is uniformly randomly selected from \tilde{e}_{ins} , \tilde{e}_{sub} , \tilde{e}_{del} and \tilde{e}_{cp} . Finally, to maintain the capacity to perform standard MT from scratch, we also consider samples where \tilde{e} is empty. The implementation of standard MT varies slightly upon approaches, as we explain below.

2.2 Model Architectures

We implement Bi-sync with Transformer-based (Vaswani et al., 2017) autoregressive and non-autoregressive models. The former (Edit-MT) is a regular Transformer with a combined input made of the concatenation of f and \tilde{e} ; the latter (Edit-LevT) is the LevT of Gu et al. (2019).

Edit-MT In this model, \tilde{e} is simply concatenated to f , with a special token to separate the two sentences. This technique has been used *e.g.* in Dabre et al. (2017) for multi-source MT or in Bulte and Tezcan (2019) for translating with a similar example. The input side of the editing training data is thus f $[\text{sep}]$ \tilde{e} , as shown in Figure 1 (top).

On the target side, we add a categorical prefix to indicate the type of edit(s) associated with a given training sample, as is commonly done for multi-domain or multilingual MT. For each basic edit (insertion, substitution and deletion), we use a binary tag to indicate if the operation is required. For instance, an \tilde{e}_{ins} needing insertions would have tags $[\text{ins}]$ $[\text{!sub}]$ $[\text{!del}]$ prepended to e . Copy corresponds to all three tags set to negative as $[\text{!ins}]$ $[\text{!sub}]$ $[\text{!del}]$. The tagging scheme provides us with various ways to perform edit-based MT: (a) we can perform inference without knowing the required edit type of \tilde{e} by generating tags then translations; (b) when the edits are known, we can generate translations with desired edits by using the corresponding tags as a forced prefix; (c) inference can also only output the edit tags and predict the relation between f and \tilde{e} . The ability to perform standard MT is preserved by training with a balanced mixture of editing data and parallel data.

The latter corresponds to an empty \tilde{e} . For these examples, the target side does not contain any tags.

Edit-LevT LevT is trained with a randomly noised version of the reference as initial target, and decodes with empty sentences. In Bi-sync, we instead initialize the target side with the given \tilde{e} for training and inference. To perform standard MT with the same model, we train with a tunable mixture of these two strategies, where p controls the proportion of each type of samples. Taking $p = 0$ is equivalent to train a LevT model with only parallel data. We use $p = 0.5$ in our experiments, making it equivalent to mixing editing and parallel data for the Edit-MT model. The value of p can be carefully designed with a schedule or curriculum to optimize the behavior of Edit-LevT, which we leave for future work. For Edit-LevT, we do not use any tags, as Edit-LevT already includes an internal mechanism to predict the edit operation(s).

3 Bilingual Re-synchronization

3.1 Datasets and Experimental Settings

We first evaluate Edit-MT and Edit-LevT on a basic resynchronization task where \tilde{e} is assumed to be the translation of a former version of f , and only a limited amount of edits is sufficient to restore parallelism. We conduct experiments on WMT14 English-French data³ in both directions (En-Fr & Fr-En) and evaluate on two test sets. The first is an artificial derivation of the standard newstest2014 set, and the second is the small parallel sentence compression dataset⁴ of [Ive and Yvon \(2016\)](#).

For the original newstest2014, we generate an initial version \tilde{e} for each test sentence and each edit operation according to the methods of Section 2.1. For deletion, we test the performance of both generation methods, resulting in four versions (Ins, Sub, Del₁, Del₂) of newstest2014 with 3,003 sentences each. The sentence compression dataset contains a subset of documents also selected from newstest2014, where sentences are compressed by human annotators while remaining parallel in the two languages. We only retain utterances for which the compressed and original versions actually differ on both sides, resulting in 526 test sentences.

Both experiments use the same training data, where we discard examples with invalid language tag as computed by `fasttext` language identifica-

tion model⁵ ([Bojanowski et al., 2017](#)), yielding a training corpus of 33.9M examples. We tokenize all data using Moses⁶ and build a shared source-target vocabulary with 32K Byte Pair Encoding (BPE) units ([Sennrich et al., 2016](#)) learned with `subword-nmt`.⁷ Since we use both parallel and artificial editing data to train edit-based models, the total training data contains about 68M utterances.

We conduct experiments using `fairseq`⁸ ([Ott et al., 2019](#)). Edit-MT relies on the Transformer-base model of [Vaswani et al. \(2017\)](#). Model and training configurations are in Appendix A. Performance is computed with `SacreBLEU` ([Post, 2018](#)).

3.2 Results

We first separately evaluate the learnability of each edit operation on our synthetic newstest2014 sets. We also derive two tasks from the compression dataset: parallel sentence compression (comp) and extension (ext). For compression, the task consists of producing a compressed target sentence e_{comp} given the compressed source f_{comp} and the original target e . For extension, the model should produce e with f and the compressed target e_{comp} . These two tasks are respectively similar to the deletion and insertion tasks. There are slight differences, though, as (a) \tilde{e} for these settings is always syntactically correct, (b) segments that are removed or inserted are selected for their lower informativeness. Therefore, these tasks are more about restoring an adequate, rather than a fluid, translation.

As mentioned in Section 2.2, the generation of translations in Edit-MT models is conditioned on predicted or oracle editing tags that are prefixed to the output: these two situations are contrasted using forced-prefix decoding with the correct tags. For the compression and extension tasks, we use the deletion and insertion tags, respectively.

Tables 2 and 3 report results for both directions, to be compared with a “do-nothing” baseline which simply copies \tilde{e} as the output. Edit-MT is able to edit the given \tilde{e} for all types of required edits much better than Edit-LevT. It obtains large gains over the *copy* baseline for insertion,⁹ substitution and deletion for both directions. When tested on

⁵<https://fasttext.cc/>

⁶<https://github.com/moses-smt/mosesdecoder>

⁷<https://github.com/rsennrich/subword-nmt>

⁸<https://github.com/pytorch/fairseq>

⁹BLEU gains for the insertion operation are artificially high. This is because the baseline is hindered by a high brevity penalty.

³<https://www.statmt.org/wmt14>

⁴<https://github.com/fyvo/ParallelCompression>

En-Fr	0	1	2	3	4	5	6	7	8-10	>10	All
N	458	2076	1133	1152	1117	1033	880	776	1852	1535	12012
copy	100.0	88.9	84.7	81.5	77.8	75.2	70.3	68.0	65.3	52.9	75.9
Edit-MT	96.4	92.6	89.3	87.3	85.8	83.3	82.3	80.4	79.4	71.7	84.1
+ tag	95.5	92.6	90.4	88.3	86.9	84.4	83.0	81.3	80.1	72.1	84.7
Edit-LevT	95.9	88.3	84.3	80.6	77.2	74.2	69.7	67.9	65.3	57.8	74.8

Fr-En	0	1	2	3	4	5	6	7	8-10	>10	All
N	341	2050	1439	1178	1125	1054	945	832	1643	1405	12012
copy	100.0	90.1	85.8	80.6	76.4	71.8	68.7	67.0	63.8	51.7	75.6
Edit-MT	97.6	92.8	91.1	87.8	83.8	82.5	80.1	80.2	77.3	69.2	83.6
+ tag	98.3	93.2	91.6	88.6	84.8	83.3	80.8	80.5	77.9	69.9	84.2
Edit-LevT	97.9	89.0	84.4	79.6	75.0	71.0	67.3	66.2	64.8	54.6	73.6

Table 1: BLEU scores for all edit types (Ins, Sub, Del₁, Del₂) broken down by the edit distance Δ between \tilde{e} and e for both En-Fr and Fr-En. Each column represents a range of distances. N denotes the number of sentences in each group. *All* is computed by concatenating all test sentences.

Model	Ins	Sub	Del ₁	Del ₂	Comp	Ext
copy	54.0	71.5	71.0	78.7	68.4	66.7
Edit-MT	75.9	77.0	86.9	94.7	73.1	67.9
+ tag	76.9	78.5	88.6	94.7	74.0	72.7
Edit-LevT	65.3	73.9	72.5	78.7	67.8	67.7

Table 2: BLEU scores for Edit-MT and Edit-LevT on resynchronization tasks for En-Fr. Deletions are evaluated separately for two generation methods (Del₁ and Del₂). + tag refers to decoding with the oracle tag as a forced-prefix. Best performance is in bold.

Model	Ins	Sub	Del ₁	Del ₂	Comp	Ext
copy	51.8	70.9	71.0	78.7	63.4	61.4
Edit-MT	73.6	74.6	87.5	95.8	65.8	69.3
+ tag	74.6	76.2	89.1	96.2	67.0	71.6
Edit-LevT	66.5	72.4	72.3	78.4	62.5	64.3

Table 3: BLEU scores for Edit-MT and Edit-LevT models on resynchronization tasks for Fr-En.

the compression and extension tasks, which have different edit distributions to the artificial editing data, Edit-MT still improves \tilde{e} by 1.2-4.7 BLEU for En-Fr and 2.4-7.9 BLEU for Fr-En. By prefixing Edit-MT with the oracle editing type tags, we can further boost the performance on almost every task for both directions. Edit-LevT can also improve \tilde{e} in most test situations, even though the gains are lower than Edit-MT models. However, due to the non-autoregressive nature of Edit-LevT, it obtains a decoding speedup of 2.3–3 \times with respect to Edit-MT when tested with the same inference batch size on the same hardware, as recommended by Helcl et al. (2022).

To better understand Edit-MT and Edit-LevT on the resynchronization task, we further analyze their performance with respect to the edit distance Δ

between \tilde{e} and e . For the results in Table 1, we merge test sentences of all edit types (Ins, Sub, Del₁, Del₂) into one test set, then break them down by the value of Δ . For both directions, prefixing the oracle editing tags for Edit-MT yields a stable improvement for almost all Δ . Edit-LevT performs similar to Edit-MT when no editions are needed, but only starts to improve the *copy* baseline when more edits are required ($\Delta \geq 8$).

4 Translating with Translation Memories

As explained above, Bi-sync encompasses example-based MT, whereby an existing similar translation retrieved from a TM is turned into an adequate translation of the source. Edit-MT actually uses the same architecture as the retrieval-based models of Bulte and Tezcan (2019); Xu et al. (2020). In this section, we study the performance of our synchronization models in this practical scenario.

4.1 Datasets

We use the same multi-domain corpus as Xu et al. (2020), which contains 11 different domains for the En-Fr direction, collected from OPUS¹⁰ (Tiedemann, 2012). We search for similar translations using Fuzzy Match.¹¹ The similarity between two English source sentences is computed as:

$$\text{sim}(\mathbf{f}, \tilde{\mathbf{f}}) = 1 - \frac{\text{ED}(\mathbf{f}, \tilde{\mathbf{f}})}{\max(|\mathbf{f}|, |\tilde{\mathbf{f}}|)}, \quad (1)$$

where $\text{ED}(\mathbf{f}, \tilde{\mathbf{f}})$ is the edit distance between \mathbf{f} and $\tilde{\mathbf{f}}$, and $|\mathbf{f}|$ is the length of \mathbf{f} . The intuition is that the closer \mathbf{f} and $\tilde{\mathbf{f}}$ is, the more suitable \tilde{e} will be. To

¹⁰<https://opus.nlpl.eu/>

¹¹<https://github.com/SYSTRAN/fuzzy-match>

Model	ECB	EMEA	Epps	GNOME	JRC	KDE	News	PHP	TED	Ubuntu	Wiki	All
copy	59.8	64.5	34.4	70.3	67.6	55.3	12.0	38.6	30.8	51.6	47.4	52.6
FM	72.1	72.3	58.3	80.6	83.2	66.9	28.0	47.2	62.9	69.3	68.8	67.3
FM [#]	69.3	68.1	58.2	74.2	80.1	65.2	28.6	44.3	62.6	68.1	69.0	65.0
Edit-MT	59.3	62.5	34.7	69.8	68.0	50.6	12.1	38.0	31.2	52.3	45.6	51.8
+ tag	60.3	63.0	35.7	70.3	68.4	51.9	12.9	38.8	32.6	52.1	45.6	52.6
+ R + tag	56.0	53.9	45.9	64.9	68.5	50.0	17.7	39.7	44.9	59.9	52.8	53.3
+ FT + tag	70.6	71.5	57.8	78.2	82.0	66.2	28.0	45.1	61.1	67.7	66.8	66.0
+ FT + R + tag	66.4	63.6	57.3	71.3	77.6	60.5	28.0	42.1	60.9	66.7	65.0	62.3
Edit-LevT	59.5	62.6	34.8	69.0	66.9	54.2	12.2	37.8	31.0	49.0	44.8	52.0
+ R	49.6	49.1	43.6	57.8	61.0	42.9	16.5	34.3	41.0	49.2	50.6	47.6
+ FT	60.5	63.2	35.2	68.3	66.5	54.5	12.1	39.0	32.9	50.7	45.9	52.4

Table 4: BLEU scores for the multi-domain test sets. *All* is computed by concatenating test sets from all domains, with 11k sentences in total. *Copy* refers to copying the similar translation in the output. *+R* implies using the related segments instead of a full initial sentence for inference. Best performance in each block are in bold.

study the ability of our models to actually make use of TMs instead of memorizing training examples, we also test on two unseen domains: OpenOffice from OPUS and the PANACEA environment corpus¹² (ENV). Detailed description and statistics about these corpora are in Appendix B. Xu et al. (2020) proposed an alternative where, for each similar translation, they masked out segments that were not aligned with the source. This means that the initial similar translation only contains segments that are directly related to the source input. We also reproduce this *related* setting, which is very similar to the insertion task.

As we are mostly interested in the edit behavior, we split the data by keeping 1,000 sentences with a sufficiently similar translation ($\text{sim} > 0.6$) in the TM as the test set for each domain. The remaining data is used for training. We use all found similar translations for training and only the best similar translation for testing. This results in 4.4M parallel sentences (*para*), in which 2.6M examples are also associated to a similar translation (*similar*) or just the related segments (*related*). Data preprocessing is similar as in Section 3.1.

4.2 Experimental Settings

Our baselines reproduce two settings for TM-based MT: the FM setting of Bulte and Tezcan (2019) and the FM[#] setting of Xu et al. (2020). The former is trained using *para + similar* data, and the latter uses *para + related*. These two baselines are trained with the same configuration as in Section 3.1. We also report scores obtained by simply copying the retrieved similar translations, as in Section 3.2.

¹²<http://catalog.elda.org/en-us/repository/browse/ELRA-W0057/>

Edit-MT and Edit-LevT from Section 3 differ from FM and FM[#] both in task, and also in training domains. Hence, we consider fine-tuning (FT) our models. For Edit-MT, we use *para+similar+ related* data and fine-tune for only 1 epoch with a learning rate of $8e^{-5}$. As we do not have information about the edit operations required to change a similar translation into the reference, we set all editing tags as on for *similar* data, and prefix the output with [ins] [sub] [del]. For the *related* data, we conjecture that mostly insertions are needed as the irrelevant segments have already been removed. We thus only activate the insertion tag. For Edit-LevT, we only use *similar+related* and fine-tune for 1 epoch with a learning rate of $9e^{-5}$. Our fine-tuned models can perform both translation with a similar sentence and with related segments.

4.3 Results and Analysis

We evaluate with BLEU, and also show TER scores in Appendix C. We reproduce in Table 4 the overall good performance of FM and FM[#]. Both significantly improve the initial similar translations. The generic Edit-MT performs much worse, and does not even match the copy results. When prefixed with the editing tag¹³ (+tag), we observe small improvements (+0.8 BLEU on average), that are further increased in the *related* scenario (+R). FT yields a much larger boost in performance (+13.4 BLEU). This highlights the effect of the task and domain mismatches on our initial results with Edit-MT. The *related* setting also benefits from FT, albeit by a smaller margin (+9 BLEU).

¹³We conjecture with a substitution tag for zero-shot inference as we have no information about the required edit type. Fine-tuned model uses the same tag as FT data.

Model $\Delta(\tilde{e}, e)$	0	1	2	3	4	5	6	7	8-10	>10	All
N	540	2096	1107	882	827	782	689	607	1193	2277	11000
copy	100.0	82.3	74.2	67.2	62.1	51.7	50.5	48.2	40.7	33.8	52.6
FM	91.6	93.3	86.5	82.3	79.2	70.5	69.0	68.0	60.9	49.5	67.3
Edit-MT+ tag	95.3	80.8	72.9	68.0	62.7	52.4	50.7	49.6	41.5	34.3	52.6
+ FT + tag	91.6	91.1	85.8	80.9	77.7	68.8	68.4	66.6	59.0	48.0	66.0
Edit-LevT	94.5	79.5	72.1	66.3	61.0	49.9	49.6	47.5	40.5	33.7	52.0
+ FT	96.6	78.9	72.0	65.8	61.1	50.4	50.4	48.0	40.8	34.3	52.4

Table 5: BLEU scores for the multi-domain test sets broken down by the edit distance Δ between \tilde{e} and e . Each column represents a range of distances. N denotes the number of sentences in each group.

Model	Office	ENV
copy	54.7	59.6
FM	66.8	75.4
FM [#]	64.0	70.6
Edit-MT+ tag	56.2	60.3
+ FT + tag	68.6	78.6
Edit-LevT	53.0	60.0
+ FT	51.9	59.6

Table 6: BLEU scores on unseen domains.

Our best results overall, using FT, are superior to FM[#] and close to that of FM. This has practical implications, since FM[#] and FM are specifically trained to transform a retrieved translation, whereas the generic Edit-MT is initially trained with artificial edits then only slightly fine-tuned on the in-domain data. To appreciate this difference, we evaluate our models on two unseen domains (Office and ENV), neither of which are used to train FM and FM[#], nor to fine-tune Edit-MT. Results in Table 6 unambiguously show that in this setting, the fine-tuned Edit-MT outperforms FM, suggesting that our edit-based model has not only adapted to the domain, but also to the task, as it can effectively perform zero-shot TM-based translations. Results obtained with the Edit-LevT model on these test sets lag far behind: even with FT on in-domain data, Edit-LevT still struggles to improve over the *copy* baseline.

We also perform the analysis with a breakdown by the edit distance Δ as in Section 3.2 on the multi-domain test sets. For the results in Table 5, all 11k test sentences are merged into one test set, then broken down by values of Δ . The generic Edit-MT starts to improve the similar translation for $\Delta \geq 3$. However, it is difficult for Edit-MT to detect very small changes ($\Delta < 3$) without FT. Once fine-tuned, Edit-MT performs similar to FM for small changes, which further confirms that Edit-MT adapts to the TM-based translation task. Edit-LevT models, however, only slightly improve the copy baseline for

similar translations requiring large changes ($\Delta \geq 8$). However, it is better than other models at detecting \tilde{e} s that do not need edits ($\Delta = 0$). We also provide results broken down the merged test set by different edit operations in Appendix C.

5 Parallel Corpus Cleaning

Our model restores synchronization between a pair of sentences. This is also useful for parallel TM cleaning tasks. Given a source sentence f and a possibly incorrect translation \tilde{e} , we want to detect non-parallelism and *to perform appropriate fixes*. We study how Edit-MT fares with this new problem on two publicly available datasets: first on the SemEval 2012&3 Task 8: Cross-lingual Textual Entailment (CLTE, Negri et al., 2011, 2012, 2013),¹⁴ then with the OpenSubtitles corpus (Lison and Tiedemann, 2016).

5.1 Cross-lingual Textual Entailment

The CLTE task aims to identify multi-directional entailment relationships between two sentences x_1 and x_2 written in different languages. We evaluate on the Fr-En direction, where x_1 is in French and x_2 is in English. The tagging mechanism of Edit-MT (see Section 2.2) can readily be used for this classification task. Data descriptions and slight adjustments of tagging scheme are in Appendix D.

We treat x_1 as f and x_2 as \tilde{e} to match the input format of Edit-MT, and perform zero-shot inference reusing the same Edit-MT model as in Section 3. We concatenate x_1 and x_2 as input, and truncate the target sequence by only taking the first three edit tags as the predicted label for the corresponding input pair, treating Edit-MT as a mere classification (CLF) model. We also slightly fine-tune Edit-MT with the 500 examples of CLTE training data for 5 epochs with a learning rate of $8e-5$.

¹⁴<https://ict.fbk.eu/clte-benchmark/>

Results are in Table 7, together with the best scores reported in Negri et al. (2013) for years 2012 and 2013 and the scores reported in Carpuat et al. (2017), which are the best reported performance we could find. Note that these scores may be quite weak, as pre-trained language models did not exist at that time. We have tried to apply a pre-trained XLM model (Conneau and Lample, 2019) to report better baselines for CLTE. However, fine-tuning XLM with only 500 sentences for 5 epochs did not outperform the reported baselines, as the fine-tuned XLM model needs to train new parameters for the linear output layer. Our Edit-MT, on the contrary, does not require any additional parameters during fine-tuning.

Methods	2012	2013
Best SemEval13	0.570	0.458
Carpuat et al. (2017)	0.604	0.436
Edit-MT En-Fr	0.350	0.284
+ FT	0.716	0.466
Edit-MT Fr-En	0.376	0.288
+ FT	0.710	0.530

Table 7: Accuracy scores on the SemEval CLTE tasks. FT denotes Edit-MT fine-tuned for classification.

As can be seen in Table 7, out-of-the-box Edit-MT fails to clearly detect the entailment relationships. This is not surprising, as there is a significant difference between our editing data and the CLTE test sets. For instance, the insertion initial translation \tilde{e}_{ins} is always grammatically incorrect, while all sentences in CLTE are syntactically correct. However, after slight fine-tuning on the CLTE data, Edit-MT for both directions can quickly adapt to the task, achieving state-of-the-art performance. This again hints that Edit-MT actually learns to identify various cases of non-parallelism.

5.2 Fixing OpenSubtitles Corpus

We further evaluate the ability of Edit-MT to detect parallel sentences and fix noisy data. We experiment with the OpenSubtitles¹⁵ data (Lison and Tiedemann, 2016) for the En-Fr direction. In this corpus, the French side is translated from English, but with noisy segments. A standard approach is to filter out noisy sentences from the training data when building systems. We aim to study whether Edit-MT can automatically identify and edit, rather than discard, noisy sentence pairs, so that training can use the full set of parallel data. We measure

¹⁵<https://opus.nlpl.eu/OpenSubtitles-v2018.php>

performance on the 10,159 segments of the En-Fr Microsoft Spoken Language Translation (MSLT) task (Federmann and Lewis, 2016), which simulates a real world MT scenario.

The OpenSubtitles data is processed similarly as in Section 3.1. We first use the fine-tuned CLF model in the CLTE task to predict the relation for all sentence pairs in OpenSubtitles data. About 60% of the data is classified as parallel, indicating that no edit operation is predicted for these segments. Models trained on the 60% clean data are denoted as *filtered*. For the other 40% presumably noisy data, we reuse the Edit-MT En-Fr model of Section 3 to fix the translations, using the predicted edit tag as a prefix on the target side (see Sections 3 and 4). Models trained on the edited data are noted as *fixed*. We train NMT models with either all data (*full*) or just the 40% noisy data as baselines. For comparison, we also train a model using the same data size (15.8M) as the noisy subset, randomly selected from the *filtered* subset.

As shown in Table 8, aggressively filtering the noisy data improves over using the full training corpus (+2 BLEU for *filtered*) more than revising it (+1 BLEU for *filtered* + *fixed*). The second set of results yields similar conclusions with smaller datasets: here, the effect of automatically fixing a set of initially noisy data improves the BLEU score by 7.2 points and closes half of the gap with a clean corpus of the same size. Note that these results are obtained without adaptation, simply reusing the pre-trained Edit-MT model of Section 3. This suggests that in situations where the training data is small and noisy, editing-based strategies may provide an effective alternative to filtering.

Cleaning Method	BLEU	Corpus size
full	44.7	41.6M
filtered	46.7	25.8M
filtered + fixed	45.7	41.6M
noisy	32.2	15.8M
fixed	39.4	15.8M
filtered (15.8M)	46.7	15.8M

Table 8: BLEU scores on MSLT tasks of models trained with different subsets of OpenSubtitles.

6 Related Work

The prediction of translations based on a source sentence and an initial translation is first explored in the context of IMT, using a left-to-right sentence completing framework proposed by Langlais

et al. (2000). The proposals of Green et al. (2014); Knowles and Koehn (2016); Santy et al. (2019) explore ways to generate translations based on given prefix hints. A more general setting, enabling arbitrary insertions thanks to LCD, is studied for online IMT systems in Huang et al. (2021). Note that LCD was initially developed for other purposes, namely enforcing lexical or terminological constraints (Hokamp and Liu, 2017; Post and Villar, 2018; Hu et al., 2019). As this approach induces large decoding overheads, recent works in this thread explore NAT techniques: Susanto et al. (2020) propose to inject lexical constraints into an edit-based LevT (Gu et al., 2019), an approach improved by Xu and Carpuat (2021) with an additional repositioning operator.

Recent attempts to revise initial translations are explored by Marie and Max (2015), who propose a touch-based scenario where users select usable translations segments, while the more questionable ones are automatically retranslated iteratively. This idea is revisited by Grangier and Auli (2018), where undesired words in the initial translation are crossed-out. The authors use a dual source encoder to represent the initial translation along with the source sentence, which is also explored by Wang et al. (2020) in a touch-editing case. The text infilling task is also considered by Xiao et al. (2022), based on a single source encoder; see also Yang et al. (2021) and Lee et al. (2021) for related proposals. These studies consider a slightly different task than ours, as they only predict the missing part of the initial translation. Nevertheless, they can all be adapted to our generic Bi-sync scenario.

Similar approaches have also been studied in APE. Multi-source architectures have been explored in *e.g.* (Junczys-Dowmunt and Grundkiewicz, 2018; Tebbifakhr et al., 2018; Shin and Lee, 2018; Pal et al., 2018), whereas Hokamp (2017); Lopes et al. (2019) jointly encode source and translation as one input. (Wisniewski et al., 2015; Libovický et al., 2016; Bérard et al., 2017) focus on learning edit operations. Junczys-Dowmunt and Grundkiewicz (2016) also propose to generate APE training data with round-trip translation.

Bi-sync also encompasses TM-based methods. Gu et al. (2018) use a second encoder to represent TM matches, an idea extended with a more compact representation of TM matches in Xia et al. (2019). As explained above, Bulte and Tezcan (2019) use a single encoder, concatenating TM seg-

ments with the source. Xu et al. (2020) further add a second embedding feature indicating related segments in TM matches, and Pham et al. (2020) propose to simultaneously consider the source and target sides of retrieved TMs. Retrieval-based MT is also explored in He et al. (2021); Khandelwal et al. (2021); Cai et al. (2021), trying to make the performance gain less dependent on the quality of the retrieved TM matches, or to enforce a tighter coupling between TM matches and translations.

7 Conclusion

This work introduced Bi-sync, the task of generating translations of a source sentence by editing a related target sentence. We have proposed various ways to create artificial initial translations for different editing types, that are needed for training. We have explored both autoregressive and non-autoregressive architectures, observing experimentally that our autoregressive Edit-MT model trained with artificial triplets performs bilingual resynchronization tasks in several real world scenarios. Edit-MT can also be quickly adapted to retrieval-based MT tasks, where we compared favorably to dedicated models. Finally, Edit-MT can also fix TMs by detecting parallel sentences and correct imperfect translations without adaptation. Another application that we wish to explore is APE.

Our NAT approach Edit-LevT is lagging behind Edit-MT. In the future, we would like to explore more NAT systems, which are computationally faster, and improve their performance. We intend to consider training curriculums and to modify the LevT model to better fit the Bi-sync task. We would also like to study ways to reduce the load of fully re-decoding the input sequence, especially when small changes, that need to be reproduced in the target, are iteratively applied to the source sentence.

Limitations

The generation of editing data for each type of edit operations requires lots of efforts and resources. This requires two separately trained NMT models to generate the data for substitution via round-trip translation with LCD. The computational cost of LCD is very high compared to regular beam search, therefore consuming many computational resources. The generation of editing data for deletion also requires one separately trained model for each method with a complete decoding of the entire training corpora. Even though our data generation

procedure is effective, the generation process may not be environmentally friendly. Due to computational limits, we were not able to conduct experiments on other languages pairs and tasks such as APE, in which large public datasets are available for other language pairs other than En-Fr.

As we decomposed the edits from one initial translation to the reference by basic edits (insertion, substitution and deletion), the generic Edit-MT and Edit-LevT models can mostly perform one type of edits at a time. It might be worth studying to combine several edit types into one single generated example, in order to approach more realistic scenarios for the generic models.

We mainly measured our results with BLEU and some additional scores in TER. However, other metrics like COMET (Rei et al., 2020) can also be interesting: as pointed out by Helcl et al. (2022), BLEU may be less appropriate to measure valid translations than COMET for NAT models.

Acknowledgement

We would like to thank the anonymous reviewers for their valuable suggestions and comments. This work was granted access to the HPC resources of IDRIS under the allocation 2022-[AD011011580R2] made by GENCI. The first author is partly funded by SYSTRAN and by a grant (Transwrite) from Région Ile-de-France.

References

- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. [Neural machine translation by jointly learning to align and translate](#). In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Alexandre Bérard, Laurent Besacier, and Olivier Pietquin. 2017. [LIG-CRISAL submission for the WMT 2017 automatic post-editing task](#). In *Proceedings of the Second Conference on Machine Translation*, pages 623–629, Copenhagen, Denmark. Association for Computational Linguistics.
- Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. 2017. [Enriching word vectors with subword information](#). *Transactions of the Association for Computational Linguistics*, 5:135–146.
- Amit Bronner, Matteo Negri, Yashar Mehdad, Angela Fahrni, and Christof Monz. 2012. [Cosyne: Synchronizing multilingual wiki content](#). In *Proceedings of the Eighth Annual International Symposium on Wikis and Open Collaboration, WikiSym '12*, New York, NY, USA. Association for Computing Machinery.
- Bram Bulte and Arda Tezcan. 2019. [Neural fuzzy repair: Integrating fuzzy matches into neural machine translation](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1800–1809, Florence, Italy. Association for Computational Linguistics.
- Deng Cai, Yan Wang, Huayang Li, Wai Lam, and Lemao Liu. 2021. [Neural machine translation with monolingual translation memory](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 7307–7318, Online. Association for Computational Linguistics.
- Marine Carpuat, Yogarshi Vyas, and Xing Niu. 2017. [Detecting cross-lingual semantic divergence for neural machine translation](#). In *Proceedings of the First Workshop on Neural Machine Translation*, pages 69–79, Vancouver. Association for Computational Linguistics.
- Alexis Conneau and Guillaume Lample. 2019. [Cross-lingual language model pretraining](#). In *Advances in Neural Information Processing Systems*, volume 32, pages 7059–7069. Curran Associates, Inc.
- Raj Dabre, Fabien Cromieres, and Sadao Kurohashi. 2017. [Enabling multi-source neural machine translation by concatenating source sentences in multiple languages](#). *CoRR*, abs/1702.06135.
- Félix do Carmo, Dimitar Shterionov, Joss Moorkens, Joachim Wagner, Murhaf Hossari, Eric Paquin, Dag Schmidtke, Declan Groves, and Andy Way. 2021. [A review of the state-of-the-art in automatic post-editing](#). *Machine Translation*, 35(2):101–143.
- Sergey Edunov, Myle Ott, Michael Auli, and David Grangier. 2018. [Understanding back-translation at scale](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 489–500, Brussels, Belgium. Association for Computational Linguistics.
- Manaal Faruqui, Ellie Pavlick, Ian Tenney, and Dipanjan Das. 2018. [WikiAtomicEdits: A multilingual corpus of Wikipedia edits for modeling language and discourse](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 305–315, Brussels, Belgium. Association for Computational Linguistics.
- Christian Federmann and William D. Lewis. 2016. [Microsoft speech language translation \(MSLT\) corpus: The IWSLT 2016 release for English, French and German](#). In *Proceedings of the 13th International Conference on Spoken Language Translation*, Seattle, Washington D.C. International Workshop on Spoken Language Translation.

- Marjan Ghazvininejad, Omer Levy, Yinhan Liu, and Luke Zettlemoyer. 2019. [Mask-predict: Parallel decoding of conditional masked language models](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 6112–6121, Hong Kong, China. Association for Computational Linguistics.
- David Grangier and Michael Auli. 2018. [QuickEdit: Editing text & translations by crossing words out](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 272–282, New Orleans, Louisiana. Association for Computational Linguistics.
- Spence Green, Sida I. Wang, Jason Chuang, Jeffrey Heer, Sebastian Schuster, and Christopher D. Manning. 2014. [Human effort and machine learnability in computer aided translation](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1225–1236, Doha, Qatar. Association for Computational Linguistics.
- Jiatao Gu, Changhan Wang, and Junbo Zhao. 2019. [Levenshtein transformer](#). In *Advances in Neural Information Processing Systems*, volume 32, pages 11181–11191. Curran Associates, Inc.
- Jiatao Gu, Yong Wang, Kyunghyun Cho, and Victor O.K. Li. 2018. [Search engine guided neural machine translation](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 32(1).
- Qiuxiang He, Guoping Huang, Qu Cui, Li Li, and Lemao Liu. 2021. [Fast and accurate neural machine translation with translation memory](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 3170–3180, Online. Association for Computational Linguistics.
- Jindřich Helcl, Barry Haddow, and Alexandra Birch. 2022. [Non-autoregressive machine translation: It’s not as fast as it seems](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1780–1790, Seattle, United States. Association for Computational Linguistics.
- Chris Hokamp. 2017. [Ensembling factored neural machine translation models for automatic post-editing and quality estimation](#). In *Proceedings of the Second Conference on Machine Translation*, pages 647–654, Copenhagen, Denmark. Association for Computational Linguistics.
- Chris Hokamp and Qun Liu. 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)*, pages 1535–1546, Vancouver, Canada. Association for Computational Linguistics.
- J. Edward Hu, Huda Khayrallah, Ryan Culkin, Patrick Xia, Tongfei Chen, Matt Post, and Benjamin Van Durme. 2019. [Improved lexically constrained decoding for translation and monolingual rewriting](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 839–850, Minneapolis, Minnesota. Association for Computational Linguistics.
- Guoping Huang, Lemao Liu, Xing Wang, Longyue Wang, Huayang Li, Zhaopeng Tu, Chengyan Huang, and Shuming Shi. 2021. [Transmart: A practical interactive machine translation system](#). *CoRR*, abs/2105.13072.
- Hakan Inan, Khashayar Khosravi, and Richard Socher. 2017. [Tying word vectors and word classifiers: A loss framework for language modeling](#). In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24–26, 2017, Conference Track Proceedings*. OpenReview.net.
- Julia Ive and François Yvon. 2016. [Parallel sentence compression](#). In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 1503–1513, Osaka, Japan. The COLING 2016 Organizing Committee.
- Marcin Junczys-Dowmunt and Roman Grundkiewicz. 2016. [Log-linear combinations of monolingual and bilingual neural machine translation models for automatic post-editing](#). In *Proceedings of the First Conference on Machine Translation: Volume 2, Shared Task Papers*, pages 751–758, Berlin, Germany. Association for Computational Linguistics.
- Marcin Junczys-Dowmunt and Roman Grundkiewicz. 2018. [MS-UEdin submission to the WMT2018 APE shared task: Dual-source transformer for automatic post-editing](#). In *Proceedings of the Third Conference on Machine Translation: Shared Task Papers*, pages 822–826, Belgium, Brussels. Association for Computational Linguistics.
- Urvashi Khandelwal, Angela Fan, Dan Jurafsky, Luke Zettlemoyer, and Mike Lewis. 2021. [Nearest neighbor machine translation](#). In *International Conference on Learning Representations*.
- Rebecca Knowles and Philipp Koehn. 2016. [Neural interactive translation prediction](#). In *Conferences of the Association for Machine Translation in the Americas: MT Researchers’ Track*, pages 107–120, Austin, TX, USA. The Association for Machine Translation in the Americas.

- Philippe Langlais, George Foster, and Guy Lapalme. 2000. [TransType: a computer-aided translation typing system](#). In *ANLP-NAACL 2000 Workshop: Embedded Machine Translation Systems*.
- Dongjun Lee, Junhyeong Ahn, Heesoo Park, and Jaemin Jo. 2021. [IntelliCAT: Intelligent machine translation post-editing with quality estimation and translation suggestion](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing: System Demonstrations*, pages 11–19, Online. Association for Computational Linguistics.
- Jason Lee, Elman Mansimov, and Kyunghyun Cho. 2018. [Deterministic non-autoregressive neural sequence modeling by iterative refinement](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1173–1182, Brussels, Belgium. Association for Computational Linguistics.
- Jindřich Libovický, Jindřich Helcl, Marek Tlustý, Ondřej Bojar, and Pavel Pecina. 2016. [CUNI system for WMT16 automatic post-editing and multimodal translation tasks](#). In *Proceedings of the First Conference on Machine Translation: Volume 2, Shared Task Papers*, pages 646–654, Berlin, Germany. Association for Computational Linguistics.
- Pierre Lison and Jörg Tiedemann. 2016. [OpenSubtitles2016: Extracting large parallel corpora from movie and TV subtitles](#). In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC'16)*, pages 923–929, Portorož, Slovenia. European Language Resources Association (ELRA).
- António V. Lopes, M. Amin Farajian, Gonçalo M. Correia, Jonay Trénous, and André F. T. Martins. 2019. [Unbabel's submission to the WMT2019 APE shared task: BERT-based encoder-decoder for automatic post-editing](#). In *Proceedings of the Fourth Conference on Machine Translation (Volume 3: Shared Task Papers, Day 2)*, pages 118–123, Florence, Italy. Association for Computational Linguistics.
- Benjamin Marie and Aurélien Max. 2015. [Touch-based pre-post-editing of machine translation output](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1040–1045, Lisbon, Portugal. Association for Computational Linguistics.
- Matteo Negri, Luisa Bentivogli, Yashar Mehdad, Danilo Giampiccolo, and Alessandro Marchetti. 2011. [Divide and conquer: Crowdsourcing the creation of cross-lingual textual entailment corpora](#). In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing*, pages 670–679, Edinburgh, Scotland, UK. Association for Computational Linguistics.
- Matteo Negri, Alessandro Marchetti, Yashar Mehdad, Luisa Bentivogli, and Danilo Giampiccolo. 2012. [Semeval-2012 task 8: Cross-lingual textual entailment for content synchronization](#). In **SEM 2012: The First Joint Conference on Lexical and Computational Semantics – Volume 1: Proceedings of the main conference and the shared task, and Volume 2: Proceedings of the Sixth International Workshop on Semantic Evaluation (SemEval 2012)*, pages 399–407, Montréal, Canada. Association for Computational Linguistics.
- Matteo Negri, Alessandro Marchetti, Yashar Mehdad, Luisa Bentivogli, and Danilo Giampiccolo. 2013. [Semeval-2013 task 8: Cross-lingual textual entailment for content synchronization](#). In *Second Joint Conference on Lexical and Computational Semantics (*SEM), Volume 2: Proceedings of the Seventh International Workshop on Semantic Evaluation (SemEval 2013)*, pages 25–33, Atlanta, Georgia, USA. Association for Computational Linguistics.
- Myle Ott, Sergey Edunov, Alexei Baevski, Angela Fan, Sam Gross, Nathan Ng, David Grangier, and Michael Auli. 2019. [fairseq: A fast, extensible toolkit for sequence modeling](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (Demonstrations)*, pages 48–53, Minneapolis, Minnesota. Association for Computational Linguistics.
- Santanu Pal, Nico Herbig, Antonio Krüger, and Josef van Genabith. 2018. [A transformer-based multi-source automatic post-editing system](#). In *Proceedings of the Third Conference on Machine Translation: Shared Task Papers*, pages 827–835, Belgium, Brussels. Association for Computational Linguistics.
- Minh Quang Pham, Jitao Xu, Josep Crego, François Yvon, and Jean Senellart. 2020. [Priming neural machine translation](#). In *Proceedings of the Fifth Conference on Machine Translation*, pages 462–473, Online. Association for Computational Linguistics.
- Matt Post. 2018. [A call for clarity in reporting BLEU scores](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 186–191, Brussels, Belgium. Association for Computational Linguistics.
- Matt Post and David Vilar. 2018. [Fast lexically constrained decoding with dynamic beam allocation for 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)*, pages 1314–1324, New Orleans, Louisiana. Association for Computational Linguistics.
- Ofir Press and Lior Wolf. 2017. [Using the output embedding to improve language models](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 157–163, Valencia, Spain. Association for Computational Linguistics.

- Ricardo Rei, Craig Stewart, Ana C Farinha, and Alon Lavie. 2020. [COMET: A neural framework for MT evaluation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2685–2702, Online. Association for Computational Linguistics.
- Sebastin Santy, Sandipan Dandapat, Monojit Choudhury, and Kalika Bali. 2019. [INMT: Interactive neural machine translation prediction](#). 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): System Demonstrations*, pages 103–108, Hong Kong, China. Association for Computational Linguistics.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. [Neural machine translation of rare words with subword units](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1715–1725, Berlin, Germany. Association for Computational Linguistics.
- Jaehun Shin and Jong-Hyeok Lee. 2018. [Multi-encoder transformer network for automatic post-editing](#). In *Proceedings of the Third Conference on Machine Translation: Shared Task Papers*, pages 840–845, Belgium, Brussels. Association for Computational Linguistics.
- Raymond Hendy Susanto, Shamil Chollampatt, and Liling Tan. 2020. [Lexically constrained neural machine translation with Levenshtein transformer](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3536–3543, Online. Association for Computational Linguistics.
- Amirhossein Tebbifakhr, Ruchit Agrawal, Matteo Negri, and Marco Turchi. 2018. [Multi-source transformer with combined losses for automatic post editing](#). In *Proceedings of the Third Conference on Machine Translation: Shared Task Papers*, pages 846–852, Belgium, Brussels. Association for Computational Linguistics.
- Jörg Tiedemann. 2012. [Parallel data, tools and interfaces in OPUS](#). In *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC’12)*, pages 2214–2218, Istanbul, Turkey. European Language Resources Association (ELRA).
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 5998–6008. Curran Associates, Inc.
- Qian Wang, Jiajun Zhang, Lemao Liu, Guoping Huang, and Chengqing Zong. 2020. [Touch editing: A flexible one-time interaction approach for translation](#). In *Proceedings of the 1st Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 10th International Joint Conference on Natural Language Processing*, pages 1–11, Suzhou, China. Association for Computational Linguistics.
- Guillaume Wisniewski, Nicolas Pécheux, and François Yvon. 2015. [Why predicting post-edition is so hard? failure analysis of LIMS submission to the APE shared task](#). In *Proceedings of the Tenth Workshop on Statistical Machine Translation*, pages 222–227, Lisbon, Portugal. Association for Computational Linguistics.
- Mengzhou Xia, Guoping Huang, Lemao Liu, and Shuming Shi. 2019. [Graph based translation memory for neural machine translation](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 33(01):7297–7304.
- Yanling Xiao, Lemao Liu, Guoping Huang, Qu Cui, Shujian Huang, Shuming Shi, and Jiajun Chen. 2022. [BiTIIMT: A bilingual text-infilling method for interactive machine translation](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1958–1969, Dublin, Ireland. Association for Computational Linguistics.
- Jitao Xu, Josep Crego, and Jean Senellart. 2020. [Boosting neural machine translation with similar translations](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1580–1590, Online. Association for Computational Linguistics.
- Weijia Xu and Marine Carpuat. 2021. [EDITOR: An Edit-Based Transformer with Repositioning for Neural Machine Translation with Soft Lexical Constraints](#). *Transactions of the Association for Computational Linguistics*, 9:311–328.
- Zhen Yang, Yingxue Zhang, Ernan Li, Fandong Meng, and Jie Zhou. 2021. [Wets: A benchmark for translation suggestion](#). *CoRR*, abs/2110.05151.

A Edit-based Model Configurations

We conduct experiments using fairseq¹⁶ (Ott et al., 2019). Edit-MT relies on the Transformer-base model of Vaswani et al. (2017). We use a hidden size of 512 and a feedforward size of 2,048. We optimize with Adam with a maximum learning rate of 0.0007, an inverse square root decay schedule, and 4,000 warmup steps. We also tie all input and output embedding matrices (Press and Wolf, 2017; Inan et al., 2017). Edit-MT is trained with mixed precision and a batch size of 8,192 tokens on 4 V100 GPUs for 300k iterations. We save checkpoints for every 3,000 iterations and average the last 10 saved checkpoints for inference. For Edit-LevT, we follow Gu et al. (2019), using a maximum learning rate of 0.0005 with 10,000 warmup steps and a larger batch size of 16,384. For inference, we set a maximum decoding round of 10.

B Details for MT with TMs

B.1 Datasets and Processing Details

Domain	Train	FM ratio	FM train
ECB	195,956	51.73%	234,943
EMEA	373,235	65.68%	624,109
Epps	2,009,489	10.12%	465,228
GNOME	55,391	39.31%	42,697
JRC	503,437	50.87%	587,859
KDE	180,254	36.00%	136,456
News	151,423	2.12%	4,048
PHP	16,020	34.93%	10,350
TED	159,248	11.90%	39,895
Ubuntu	9,314	20.32%	1,738
Wiki	803,704	19.87%	409,755
Total	4,457,471	24.27%	2,557,078
Office	49,845	43.76%	-
ENV	13,632	6.81%	-

Table 9: Data used for experiments of Section 4.2. *FM ratio* is the ratio of sentences with at least one matched similar translations, *FM train* is the actual number of examples augmented with a similar translation used for training, after setting aside 1,000 test sentences for each domain. Each training sentence is matched with up to 3 similar translations.

Our experiments use the same multi-domain corpus as Xu et al. (2020). This corpus contains 11 different domains for the En-Fr direction, collected from OPUS¹⁷ (Tiedemann, 2012): documents from

the European Central Bank (ECB); from the European Medicines Agency (EMA); Proceedings of the European Parliament (Epps); legislative texts of the European Union (JRC); News Commentaries (News); TED talk subtitles (TED); parallel sentences extracted from Wikipedia (Wiki); localization files (GNOME, KDE and Ubuntu) and manuals (PHP). All these data were deduplicated prior to training. To evaluate the ability of our models to actually make use of TMs instead of memorizing training examples, we also test on two unseen domains: OpenOffice from OPUS and the PANACEA environment corpus¹⁸ (ENV). We follow Xu et al. (2020) and search for the top 3 similar translations based on Fuzzy Match with a similarity score greater than 0.6 as computed by Equation (1) on the source side and without an exact match. Note that the ratio of sentences with at least one similar translation greatly varies across domains, as shown in Table 9. When reproducing the *related* setting, Xu et al. (2020) used a placeholder token to mark the positions where segments are deleted. As discussed in Section 2.1, our models do not include such information. Therefore, we do not use placeholders for the *related* data.

C Additional Results

Table 10 reports results on multi-domain test set for the task of translating with TMs measured on TER using SacreBLEU (Post, 2018). TER results show that even the generic Edit-MT model actually identifies useful edits, as we see improvements with respect to the copy baseline.

Table 11 shows results on BLEU broken down the aggregate test set by edit operations. We observe that the generic Edit-MT struggles to perform well when substitutions are needed. FT vastly improves the ability to substitute and delete from \tilde{e} , separately or even in combination. Fine-tuned Edit-MT even outperforms FM when only deletions are required.

D Details for Parallel Corpus Cleaning

In the CLTE task, the goal is to identify multi-directional entailment relationships between two sentences x_1 and x_2 , written in different languages. Each (x_1, x_2) pair in the dataset is annotated with one of the following relations: *Bidirectional* ($x_1 \leftrightarrow x_2$): the two fragments entail each

¹⁶<https://github.com/pytorch/fairseq>

¹⁷<https://opus.nlpl.eu/>

¹⁸<http://catalog.elda.org/en-us/repository/browse/ELRA-W0057/>

TER ↓	ECB	EMEA	Epps	GNOME	JRC	KDE	News	PHP	TED	Ubuntu	Wiki	All
copy	0.435	0.377	0.659	0.263	0.294	0.457	0.999	0.703	0.653	0.372	0.502	0.488
FM	0.286	0.301	0.374	0.160	0.138	0.312	0.642	0.572	0.328	0.227	0.306	0.314
FM#	0.299	0.332	0.365	0.202	0.152	0.309	0.634	0.605	0.327	0.234	0.309	0.327
Edit-MT	0.418	0.390	0.641	0.260	0.285	0.455	0.904	0.646	0.643	0.359	0.520	0.472
+ tag	0.412	0.382	0.638	0.256	0.283	0.452	0.899	0.643	0.635	0.360	0.520	0.468
+ R + tag	0.397	0.422	0.508	0.268	0.237	0.428	0.825	0.636	0.520	0.282	0.457	0.430
+ FT + tag	0.300	0.308	0.376	0.182	0.146	0.314	0.643	0.595	0.351	0.237	0.329	0.326
+ FT + R + tag	0.322	0.357	0.375	0.219	0.168	0.342	0.640	0.623	0.350	0.238	0.338	0.346
Edit-LevT	0.437	0.390	0.655	0.267	0.296	0.463	0.996	0.720	0.656	0.399	0.528	0.496
+ R	0.434	0.446	0.503	0.314	0.294	0.471	0.780	0.649	0.515	0.370	0.478	0.458
+ FT	0.443	0.390	0.663	0.276	0.306	0.467	1.008	0.722	0.677	0.389	0.531	0.502

Table 10: TER scores on multi-domain test sets. *All* is computed by concatenating test sets from all domains, with 11k sentences in total. *Copy* refers to copying the similar translation in the output. *+R* implies using the related segments instead of a full initial sentence for inference. Best performance in each block are in bold.

BLEU↑	=	Ins	Sub	Del	Ins+Sub	Ins+Del	Sub+Del	Ins+Sub+Del
<i>N</i>	540	316	3260	316	2865	58	2634	1011
copy	100.0	72.0	67.9	75.4	32.5	69.8	34.0	47.3
FM	91.6	80.6	86.6	82.9	50.0	67.4	58.4	63.0
Edit-MT+ tag	95.3	75.7	67.0	77.2	34.2	68.5	37.4	48.0
+ FT + tag	91.6	79.7	84.6	85.8	48.3	69.9	57.6	60.8
Edit-LevT	94.5	75.3	65.9	73.3	33.1	69.4	33.7	46.9
+ FT	96.6	75.5	65.6	74.1	34.3	70.1	33.9	47.0

Table 11: BLEU scores for the multi-domain test sets broken down by the edit operations between \tilde{e} and e . Each column represents a combination of edits. *N* denotes the number of sentences in each group.

other (semantic equivalence); *Forward* ($x_1 \Rightarrow x_2$ & $x_1 \not\Leftarrow x_2$): unidirectional entailment from x_1 to x_2 ; *Backward* ($x_1 \not\Rightarrow x_2$ & $x_1 \Leftarrow x_2$): unidirectional entailment from x_2 to x_1 ; *No Entailment* ($x_1 \not\Leftarrow x_2$): no entailment between x_1 and x_2 . The dataset contains a training set of 500 pairs, and two test sets of the same size (test-2012 and test-2013).

CLTE Fr-En	Edit-MT En-Fr	Edit-MT Fr-En
Bidirectional	Copy	Copy
Forward	Deletion	Insertion
Backward	Insertion	Deletion
No Entailment	Substitution	Substitution

Table 12: Label conversion scheme between CLTE task and Edit-MT editing tags.

The tagging mechanism of Edit-MT described in Section 2.2 can readily be used for this classification task with slight adjustments represented in Table 12.

We have not performed hyperparameter searching for FT, even though carefully fine-tuned models may achieve even better performance.