

Lingua Custodia at WMT'19: Attempts to Control Terminology

Franck Burlot

Lingua Custodia

Montigny-le-Bretonneux, France

franck.burlot@linguacustodia.com

Abstract

This paper describes Lingua Custodia's submission to the WMT'19 news shared task for German-to-French on the topic of the EU elections. We report experiments on the adaptation of the terminology of a machine translation system to a specific topic, aimed at providing more accurate translations of specific entities like political parties and person names, given that the shared task provided no in-domain training parallel data dealing with the restricted topic. Our primary submission to the shared task uses backtranslation generated with a type of decoding allowing the insertion of constraints in the output in order to guarantee the correct translation of specific terms that are not necessarily observed in the data.

1 Introduction

A sub-task of the WMT'19 News Translation shared task has been jointly organized by the University of Le Mans and Lingua Custodia: the translation of news articles dealing with the topic of the 2019 European Parliament elections for the French-German language pair. This brings back French, a language absent from the News Translation task since 2015, and pairs it with German, a morphologically richer language than English. Finally, the EU election topic brings new challenges to the task.

Such a restriction of the domain to a single topic makes the task very different from the translation of any news data. We propose to roughly define a domain according to two major dimensions:

- **Syntactic structure.** The European election topic probably has no or few syntactic and stylistic differences with the general news domain, since we are in both cases dealing with news articles with the same characteristics. On the other hand, sentences in newspapers are generally longer than in casual discourse.

- **Terminology.** A specific topic implies a specific terminology. For instance, the system should not attempt a literal translation of the German politician's name *Wagenknecht*. It should also be aware of the specific translations of political party names in the press of the target language: the French party *France Insoumise* should not be translated into German. Furthermore, the French movement *gilets jaunes* (yellow vests) is referred to in the German press as *Gelbwesten*, and a literal translation, such as *gelbe Westen*, is inaccurate.

There exist efficient methods for domain adaptation in neural MT (Luong and Manning, 2015; Chu and Wang, 2018). The experiments introduced in this paper attempt to explore techniques that help to specifically adapt the terminology of a system to a restricted topic. However, a serious difficulty stands in the way: among the parallel data provided for the task, only 1,701 sentence pairs deal with the EU elections (development set). Recent monolingual data in German and French is available and contains several sentences using the required terminology, but we then lack the correct translations of the terms of interest.

This paper describes Lingua Custodia's attempts to specifically control the terminology generated by a Machine Translation (MT) system, using only the data provided at the Conference. The resulting German-to-French system was submitted at WMT'19.

In the first section, we provide an overview of our baselines and point out several terminology issues. We then describe our experiments with constrained decoding to control terminology. The last section introduces an attempt to relax the hard constraints applied to the decoder.

2 Baseline

The training parallel data provided for the task consisted of nearly 10M sentences, including *Europarl* (Koehn, 2005), *Common-crawl*, *News-commentary* and *Bicleaner07*. The former was the biggest (over 7M sentences) and also the noisiest corpus, containing bad characters, short phrases with only numbers, lists of products, sentences in the wrong language, obviously machine translated sentences, etc.

2.1 Data selection

We have performed a filtering of the *Bicleaner07* corpus in order to reduce the impact of noisy samples on the MT system, using *LC-Pruner*, a in-house system that was submitted at the First Automatic Translation Memory Cleaning Shared Task (Barbu et al., 2016). The system extracts several monolingual and bilingual features that are fed to a random forest classifier aimed at predicting if a sentence pair is a good translation and whether each sentence is well formed. It is based on the following features:

- Total sentence pair length
- Source/target length ratio
- Average token length
- Uppercase token count comparison
- Source/target punctuation comparison
- Source/target number comparison
- Language identification using *langid.py* (Lui and Baldwin, 2012)
- Cognates
- Source and target language model scores
- Hunalign scores (Varga et al., 2007)
- Zipporah adequacy scores (Xu and Koehn, 2017), using a probabilistic bilingual dictionary computed on *Europarl*.

Random forest parameters are optimized using expert feedback on a set of parallel sentences automatically selected by the model across several iterations. We have run 3 iterations, assessing the quality of 20 sentence pairs each time. The result is a binary classification of each sentence pair

based on a score between 0 and 1. We have experimented with two selection criteria, keeping sentence pairs scoring above 0.5 and above 0.8, which led to respectively nearly 4M and 2M finally accepted sentences. The results are introduced in Section 2.3.

2.2 System setup

German and French pre-processing was performed using in-house normalization and tokenization tools. Trucasing models were learnt, using Moses scripts (Koehn et al., 2007), on the monolingual news data provided at the Conference, on all 2017-2018 data for French and 10M sentences from 2018 for German. A shared French-German BPE vocabulary (Sennrich et al., 2016b) was built with 30k merge operations on all the parallel data available for the task, except *Bicleaner07*.

We have trained baseline systems for French-German in both directions. Transformer *base* (Vaswani et al., 2017) models were trained using the Sockeye toolkit (Hieber et al., 2017) on two Nvidia 1080Ti GPU cards. Most of the standard hyper-parameters have been used. The model dimension included 512 units. The initial learning rate was set to 0.0003 with a warmup on for 30k updates. Due to the small quantity of training data available, we decided to slightly increase dropout between layers (0.2) and label smoothing (0.2). Validations were performed every 20k updates and patience was set to 15. Since this setup contained no training data relevant to the EU election topic, we decided to hold out the provided development set for another purpose, and used a general news domain test set: *Newstest-2012*. We finally wished to sample more sentence pairs from news-related corpora during training. Since no such method is implemented in the Sockeye toolkit for minibatch generation, we simply trained the baselines on a single copy of *Bicleaner07* and *Common-crawl*, and took two copies of *Europarl* and 6 of *News-commentary*.

2.3 Results and terminology issues

The systems were tested on the official development set, *Euelections-dev-2019*, as well as *Newstest-2013* and the official test set *Newstest-2019*. BLEU scores were computed with *SacreBLEU* (Post, 2018) and are shown in Table 1.

Experiments with different data filtering criteria for the *Bicleaner07* corpus were introduced in subsection 2.1. We observe that keeping a bigger

French-to-German			
	Euelections-dev-2019	Newstest-2013	Newstest-2019
Baseline	25.98	23.48	26.94
German-to-French			
	Euelections-dev-2019	Newstest-2013	Newstest-2019
LC.Pruner 2M	31.07	27.49	33.04
LC.Pruner 4M	30.96	27.29	33.16

Table 1: BLEU scores for French-German baselines

set of data does not lead to any clear improvements, at least in terms of BLEU. Thus we have kept *LC.Pruner 2M* as the main baseline for further training in Section 4.1.

The translation from English into German of *Euelections-dev-2019* by our baseline shows consistent terminology issues. The systems has difficulties translating the name of the movement *gilets jaunes* (yellow vests). Out of the 19 occurrences of the expression in the French source, only 4 are correctly translated as the compound *Gelbwesten*. We noted several translations as *gelbe Westen*, the translation of the adjective *jaunes* only, as well as full omissions. We also noted that the French party *France Insoumise* was translated literally as *unbeugsame Frankreich*, instead of simply being copied, the name of the politician *Nicolas Dupont-Aignan* was translated as *Nicolas Dumont-Aignan*, etc. Our best baseline translates the German side of this test into French with the same kind of difficulties: *Gelbwesten* is sometimes translated as *la veste jaune*, etc.

3 Terminology control

We argue that a system specialized in a specific topic should be able to provide the right translations for terms that are relevant to this topic. The baselines we have just introduced fail to translate important terminology. We now seek to adapt these baselines to the EU election terminology.

3.1 Constrained decoding

One way to integrate such knowledge of a specific terminology into the MT system is by using *constrained decoding* (Hokamp and Liu, 2017). The *Grid Beam Search* algorithm guarantees the presence of one or several given phrases in the MT output. This method does not require any change in the model or its parameters, thus the algorithm does not model any sort of token-level source-to-target relation, but simply forces the beam search to go through the target constraint. The challenge

for the decoder is then to correctly insert the constrained phrase in the rest of the sentence.

Post and Vilar (2018) proposed a variant of this algorithm with a significant lower computational complexity. We used their implementation available in the Sockeye toolkit.

3.2 Lexicon extraction

We have extracted bilingual lexicons from two sources: the official development set provided for the task (*Euelections-dev-2019*), and the monolingual French and German data made available at WMT.

3.2.1 Parallel EU election data

We have decided to use the official development set (*Euelections-dev-2019*) as the main source of terminology, for the simple reason that it is the only parallel data available containing the specific terminology of the EU elections with reliable human translations.

Alignments were learnt using Fastalign (Dyer et al., 2013) on a concatenation of *News-commentary* and *Euelections-dev-2019*, and we used them to extract a phrase table from the former with the Moses toolkit. We removed a phrase pair whenever the probability of the German side, given the French side, was below 0.5. This ensured that we never keep more than one translation for a French phrase¹.

The resulting phrases were furthermore filtered according to their domain. We computed Moore-Lewis (Moore and Lewis, 2010) scores of the source French phrases. The out-of-domain language model was computed on the French side of the parallel data (section 2), and the in-domain model on the French monolingual news data 2018 available at WMT. Although this corpus does not contain exclusively articles about the EU elections, we believe its terminology distribution may

¹Since there can be several French translations for one German phrase, the current terminology can only be used for translation into German.

be closer to what is observed in *Euelections-dev-2019*, because the corpus relates more recent news. We kept the best 2000 phrase pairs according to their Moore-Lewis score.

Finally, we kept the phrase pairs for which the German side appeared at least once in the German monolingual news 2018 corpus, in order to filter out obviously bad expressions that remained. We ended up with 773 phrase pairs, among which could be found the correct translation of *gilets jaunes* (yellow vests).

3.2.2 Monolingual news data

As an attempt to address the issue of person name mistranslations, we extracted named entities from the French monolingual news 2018 corpus. First, we tagged the corpus with an in-house French named entity recognizer. We then computed the tagged named entity occurrence counts over the same corpus and removed the ones occurring less than 9 times. The translations of the extracted expressions into German are unknown, so we looked for the named entities that are not translated, but copied into German. We therefore kept the entries that had an occurrence count higher than 9 in the German news monolingual 2018 corpus. As a result, the name *Poutine* in French would be removed because it translates into a different word in German (*Putin*), whereas *Dupont-Aignan* would be kept, as it stays the same in both languages. This procedure produced nearly 20k phrase pairs.

Prior to inference, constraints extracted from the development set are applied every time a source-side constraint is found in the source sentence to be translated. Named Entity constraints extracted from monolingual data are applied in a different way. The same named entity classifier as above is used to tag the source sentence and a constraint is applied when: 1. the source constraint matches a part of the sentence ; 2. the matched sentence part has been tagged as a named entity.

We are well aware that bilingual terminology extraction is a complex task and that more sophisticated models need to be investigated. We chose to employ these simple heuristics only because we lacked time. We did run experiments with tools, allowing us to extract bilingual lexicons from monolingual data, namely *Muse* (Conneau et al., 2017) and *BiLex* (Zhang et al., 2017). However, we found them not suited for our requirements, because 1. the global quality of the lexicons was too low to be inserted in a MT de-

coder as hard constraints, and 2. only single-word phrases were extracted and we wished to extract multi-word expressions as well. Future work should include methods for phrase pair extraction from monolingual data (Marie and Fujita, 2018; Artetxe et al., 2019).

3.3 Constrained French-to-German baseline

The scores of the French-to-German baseline with and without constraints are shown in Table 2. We used a beam size of 20 for constrained decoding, as recommended in the Sockeye documentation², and a default beam size of 5 for the unconstrained decoding. The final models are averages of the 4 best checkpoints in terms of BLEU on the validation set. Applying constraints to *Euelections-dev-2019* adds 2 BLEU points to the baseline, but this should not be considered as an improvement, since parts of the reference translations were inserted as constraints. We observe that constrained decoding has nearly no impact on the BLEU score for *Newstest-2013*, and that it even slightly degrades the score for *Newstest-2019*.

The low impact of the constraints on *Newstest-2013* may be explained by the fact that this set is irrelevant with regard to the EU election topic, leading to the insertion of few constraints: 465 constraints were inserted in 3000 sentences. As a comparison, 751 constraints were inserted in the 1701 sentences of *Newstest-2019*. Looking more closely at the outputs of the different systems, we observed several cases where : 1. the constraint was erroneously inserted in the sentence; 2. the insertion of a constraint seemed to disturb the decoder, which resulted in broken sentences. Table 3 illustrates a case where the constraint helped to correct a mistranslation, but both issues occurred. The French party *France Insoumise* was translated literally by the baseline into *Ununterwürfiges Frankreich*, and one of our constraints successfully forced the right translation of this expression. First, the subject of the first clause (*les populistes de gauche*) has been replaced by the constraints, which should have been inserted in the end of the sentence, like in the baseline. Second, the constrained output ignores the whole section about the raise of classical populist parties.

Although several constraints may potentially

²<https://awsllabs.github.io/sockeye/inference.html>

	Euelections-dev-2019	Newstest-2013	Newstest-2019
Baseline	25.98	23.48	26.94
+ Constraints	27.87	23.42	26.66

Table 2: BLEU scores for French-to-German with constrained decoding

Source	Même si les populistes de gauche ont bien moins de succès en Europe que les acteurs d’extrême-droite, ils peuvent encore s’imposer, comme le montre l’ascension de partis classiques d’opposition tels que Podemos en Espagne et La France Insoumise en France.
Constraints	Podemos, France Insoumise
English	Even if left-wing populists have far less success in Europe than right-wing actors, they can still prevail, as evidenced by the rise of classic opposition parties such as Podemos in Spain and France Insoumise in France.
Baseline	Obwohl die Linkspopulisten in Europa deutlich weniger erfolgreich sind als die Rechtsextremen, können sie sich immer noch durchsetzen, wie der Aufstieg klassischer Oppositionsparteien wie Podemos in Spanien und Frankreichs Ununterwürfiges Frankreich zeigt.
+ Constraints	Podemos in Spanien und France Insoumise in Frankreich haben zwar deutlich weniger Erfolg als rechtsextreme Populisten, aber sie können sich noch immer durchsetzen.
Reference	Auch wenn die Linkspopulisten in Europa weitaus weniger erfolgreich sind als die Rechts-außen-Player, können sie sich durchaus Geltung verschaffen, wie der Aufstieg klassischer Herausforderer-Parteien wie Podemos in Spanien und La France Insoumise in Frankreich zeigt.

Table 3: Example of French-to-German translation with and without constrained decoding (*Newstest-2019*)

help the adaptation of a MT system to the specific terminology of the EU elections, it may be possible that the positive impact it could have on BLEU is mitigated by the broken translations the constraints tend to produce.

4 Relaxed use of constraints

We assume that the strict insertion of terminology through constrained decoding sometimes breaks output sentences, partly because the decoder would have never generated such an expression by itself. More specifically, the decoder assigns a low probability to the constrained phrase, which leads to a harmful disruption during the beam search.

Using parallel data containing the required terminology to fine-tune a system is an obvious good way to adapt a system, and it has the advantage to leave the decoder unchanged. Although we have no such data available for training, we do have monolingual French data that contains at least a big part of the EU election terminology we wish to acquire: the monolingual news 2018 corpus released within the shared task. We could use our French-to-German baseline to backtranslate these sentences (Sennrich et al., 2016a), but this would have the effect of introducing mistranslations in the source, which would break the strict source-target mapping we need to learn. For instance, if the French phrase *gilets jaunes* is backtranslated as *gelbe Westen*, the final German-to-French system would learn to translate *gelbe Westen* into French, but could very well still produce erroneous translations of the correct source expression *Gelbwesten*.

To address this issue, we propose to apply the strict constraints (section 3.1) to the French-to-German baseline used for backtranslation. Although we condemn ourselves to certain broken translated outputs, we have the guarantee that the extracted constraints will be learnt by the system. Another advantage of this strategy is that the constraints are inserted in different contexts, which should help the decoder learn to insert constrained terms in the output sentences more correctly.

4.1 Synthetic parallel datasets

The French news monolingual corpus 2018 comes under the general news domain. We attempted to extract the sentences dealing with the EU election topic using Moore-Lewis data selection strategy (Moore and Lewis, 2010). We chose the French side of *Euelections-dev-2019* as our in-domain corpus, with the hope that it will favor sentences containing the constraints we have extracted from it, in order to maximize the presence of constraint pairs in the backtranslated data. We finally selected the best 2M sentences in terms of Moore-Lewis score.

We provide both constrained and unconstrained translations for the resulting French sentences, using the same beam sizes as in Section 3.3. The constrained setup inserted 673,670 phrases in 2M German sentences.

4.2 Results

We used the German-to-French baseline trained on 2M sentences from *Bicleaner07* (section 2.3) as a starting point for fine-tuning using the constrained

	Euelections-dev-2019	Newstest-2013	Newstest-2019
Baseline	31.07	27.49	33.04
Unconstrained	34.06	28.07	35.64
Constrained	34.04	27.99	35.45
Ensemble	34.31	28.10	35.62

Table 4: BLEU scores for German-to-French systems fine-tuned on backtranslated data

and unconstrained versions of the backtranslation. The backtranslated data was mixed with *Europarl* and *News-commentary* corpora. We first tried to use *Newstest-2012* for validation, but only a slight improvement was observed throughout the training in terms of BLEU. In order to avoid stopping the training too early, we finally decided to run validation on *Euelections-dev-2019*. This most certainly led to overestimated BLEU scores, since the backtranslation data has been selected according to its proximity to this development set (section 4.1). However, it allowed the stopping criterion to fire later during training.

The final models we introduce are averages of the 4 best checkpoints in terms of BLEU on *Euelections-dev-2019*. We also provide results for an ensemble of 8 checkpoints (4 best constrained and 4 best unconstrained). We kept the same hyper-parameters as described in Section 2.2, except we lowered the learning rate from 0.0003 to 0.0001, used no warmup, and ran more frequent validations (every 10k updates).

The result of these fine-tuning procedures are shown in Table 4. Both backtranslation setups provide the best improvements we observed on *Newstest-2019* (+2.5). However, we see no significant difference between the constrained and unconstrained setups. This could be expected, since our experiment was focused on a small set of terms we wished the systems to generate, which can only lead to local improvements with low impact on the BLEU score. The ensemble of 8 models combining both setups is our primary submission to the shared task.

We have run a small analysis of the outputs given by both setups for *Newstest-2019*. We observed that the constrained system correctly copied the German name *Alexander Gauland*³, whereas the unconstrained system erroneously translated the first name into *Alexandre*. The constrained system also translated *europäischen Vermögenssteuer* (European wealth tax) into the

acronym *ISF européen*⁴, which seems more usual in the press about the EU elections, compared to the literal translation of the unconstrained system as *impôt européen sur la fortune*. Several phrases that were in our extracted constraints were correctly translated by the unconstrained system as well. Unconstrained backtranslation (Sennrich et al., 2016a) thus seems to be sufficient to adapt the terminology of a system to a specific system, at least in our setup with few low-quality automatically extracted lexical constraints. However, both systems produce consistent errors on terms that we failed to capture in constraints, which leads us to think that higher quality constraints should have a bigger positive impact on terminology adaptation.

5 Conclusions

We have described Lingua Custodia’s submission to WMT’19 News Translation shared task. We attempted to adapt the terminology of a MT system to the EU election topic without relevant parallel training data. Forcing the decoder to generate specific terms can help, although it disturbs the decoder, which may lead to broken output sentences. Using hard constraint insertion to generate backtranslated target monolingual data showed no improvement in terms of BLEU scores, but we have observed local improvements in the generated terminology. The system that has been submitted to the shared task is an ensemble of both constrained and unconstrained models.

Lexically constrained decoding is highly dependent on the quality of the bilingual constraints available. In future work, we plan to search for other techniques for automatic lexical constraint extraction in order to improve recall and reach a better terminology coverage. We also plan to investigate new techniques to relax the hard constraints applied to the decoder, in order to impose less disturbance to the beam search and avoid broken output sentences.

³Constraint: *Alexander Gauland* → *Alexander Gauland*

⁴Constraint: *Vermögenssteuer* → *ISF*

References

- Mikel Artetxe, Gorka Labaka, and Eneko Agirre. 2019. An effective approach to unsupervised machine translation. *CoRR*, abs/1902.01313.
- Eduard Barbu, Carla Parra Escartín, Luisa Bentivogli, Matteo Negri, Marco Turchi, Constantin Orasan, and Marcello Federico. 2016. The first automatic translation memory cleaning shared task. *Machine Translation*, 30(3-4):145–166.
- Chenhui Chu and Rui Wang. 2018. A survey of domain adaptation for neural machine translation. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 1304–1319, Santa Fe, New Mexico, USA. Association for Computational Linguistics.
- Alexis Conneau, Guillaume Lample, Marc’Aurelio Ranzato, Ludovic Denoyer, and Hervé Jégou. 2017. Word translation without parallel data. *CoRR*, abs/1710.04087.
- Chris Dyer, Victor Chahuneau, and Noah A. Smith. 2013. A simple, fast, and effective reparameterization of IBM model 2. In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 644–648, Atlanta, Georgia. Association for Computational Linguistics.
- Felix Hieber, Tobias Domhan, Michael Denkowski, David Vilar, Artem Sokolov, Ann Clifton, and Matt Post. 2017. Sockeye: A toolkit for neural machine translation. *CoRR*, abs/1712.05690.
- 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.
- Philipp Koehn. 2005. A parallel corpus for statistical machine translation. In *Proc. MT-Summit*, Phuket, Thailand.
- Philipp Koehn, Hieu Hoang, Alexandra Birch, Chris Callison-Burch, Marcello Federico, Nicola Bertoldi, Brooke Cowan, Wade Shen, Christine Moran, Richard Zens, Chris Dyer, Ondřej Bojar, Alexandra Constantin, and Evan Herbst. 2007. Moses: Open source toolkit for statistical MT. In *Proc. ACL: Systems Demos*, pages 177–180, Prague, Czech Republic.
- Marco Lui and Timothy Baldwin. 2012. *langid.py: An off-the-shelf language identification tool*. In *Proceedings of the ACL 2012 System Demonstrations*, pages 25–30, Jeju Island, Korea. Association for Computational Linguistics.
- Minh-Thang Luong and Christopher D. Manning. 2015. Stanford neural machine translation systems for spoken language domain. In *International Workshop on Spoken Language Translation*, Da Nang, Vietnam.
- Benjamin Marie and Atsushi Fujita. 2018. Phrase table induction using monolingual data for low-resource statistical machine translation. *ACM Trans. Asian Low-Resour. Lang. Inf. Process.*, 17(3):16:1–16:25.
- Robert C. Moore and William Lewis. 2010. Intelligent selection of language model training data. In *Proceedings of the ACL 2010 Conference Short Papers*, ACLShort ’10, pages 220–224, Stroudsburg, PA, USA. 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, Belgium, Brussels. 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.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016a. Improving neural machine translation models with monolingual data. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 86–96, Berlin, Germany. Association for Computational Linguistics.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016b. 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. Association for Computational Linguistics.
- Dániel Varga, Péter Halácsy, Andras Kornai, Viktor Nagy, László Németh, and Viktor Trón. 2007. *Parallel corpora for medium density languages*, pages 247–258. John Benjamins.
- 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.
- Hainan Xu and Philipp Koehn. 2017. Zipporah: a fast and scalable data cleaning system for noisy web-crawled parallel corpora. In *Proceedings of the 2017*

Conference on Empirical Methods in Natural Language Processing, pages 2945–2950, Copenhagen, Denmark. Association for Computational Linguistics.

Meng Zhang, Haoruo Peng, Yang Liu, Huan-Bo Luan, and Maosong Sun. 2017. [Bilingual lexicon induction from non-parallel data with minimal supervision](#). In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, February 4-9, 2017, San Francisco, California, USA.*, pages 3379–3385.