

# Dynamic Terminology Integration for COVID-19 and other Emerging Domains

Toms Bergmanis<sup>†‡</sup> and Mārcis Pinnis<sup>†‡</sup>

<sup>†</sup>Tilde / Vienības gatve 75A, Riga, Latvia

<sup>‡</sup>Faculty of Computing, University of Latvia / Raiņa bulv. 19, Riga, Latvia  
{firstname.lastname}@tilde.lv

## Abstract

The majority of language domains require prudent use of terminology to ensure clarity and adequacy of information conveyed. While the correct use of terminology for some languages and domains can be achieved by adapting general-purpose MT systems on large volumes of in-domain parallel data, such quantities of domain-specific data are seldom available for less-resourced languages and niche domains. Furthermore, as exemplified by COVID-19 recently, no domain-specific parallel data is readily available for emerging domains. However, the gravity of this recent calamity created a high demand for reliable translation of critical information regarding pandemic and infection prevention. This work is part of WMT2021 Shared Task: *Machine Translation using Terminologies*, where we describe Tilde MT systems that are capable of dynamic terminology integration at the time of translation. Our systems achieve up to 94% COVID-19 term use accuracy on the test set of the EN-FR language pair without having access to any form of in-domain information during system training. We conclude our work with a broader discussion considering the Shared Task itself and terminology translation in MT.

## 1 Introduction

This work is part of WMT2021 Shared Task: *Machine Translation using Terminologies*, which is concerned with improving machine translation (MT) accuracy and consistency on *newly developed domains* by utilising word and phrase-level terms. We describe Tilde MT systems that are capable of dynamic terminology integration at inference time. Our submissions consist of translations by terminology-enabled general-purpose MT systems for EN-RU, EN-FR, and CS-DE translation

directions. Our systems are deliberately trained without consideration for the test domain to follow the spirit of the Shared Task—MT for emerging domains. Despite term collections being noisy, our MT systems with dynamic terminology integration improve term translation accuracy proving their usefulness in dynamic adaptation for novel domains, where training-time domain adaptation methods are not feasible.

The remainder of this work describes the methods used for dynamic terminology integration (Section 2) by describing the tasks of terminology filtering, term recognition, and dynamic terminology integration in the translation process. The bulk of Section 2 describes problems due to the low-quality term collections and terminology mismanagement and our solutions to them. We hope that the examples provided will not only illustrate the self-imposed problems by the Shared Task but also will motivate reconsidering the purpose and the desired qualities of a term collection in the context of MT. We then briefly describe the experimental setting and results in Section 3 and Section 4 respectively. We conclude our work with a broader discussion considering the Shared Task and terminology translation in MT in Section 5.

## 2 Methods

This section describes the three tasks necessary for successful MT with terminology: terminology filtering, term recognition, and finally, integration of terminology constraints in the translation process.

### 2.1 Terminology Filtering

To guarantee terminology translation correctness and consistency, which are two quality aspects of terminology translation, term collections must provide unambiguous information about the preferred translation equivalent for each source term's type (full form, short form, or acronym) when listing multiple possible translation equivalents in a tar-

\*Both authors have contributed equally.

|       | # | Source Term                 | Target Side of a Term Entry   |
|-------|---|-----------------------------|---|
| EN-RU | 1 | hand-washing                | мытьё рук   <u>мытья рук</u>   <u>мытья</u>   |
|       | 2 | sneeze                      | <u>чихании</u>   чихания   <u>чихнуть</u>   |
|       | 3 | flu epidemic                | эпидемия гриппа   <u>эпидемия свиного гриппа</u>  |
| EN-FR | 4 | World Health Organization   | Organisation mondiale de la Santé   <i>Organisation Mondiale de la Santé</i>   <i>Organisation mondiale de la santé</i>                       |
|       | 5 | Coronavirus outbreak        | épidémie de coronavirus   |
|       | 6 | <i>coronavirus outbreak</i> | <u>épidémies de coronavirus</u>   |
| CS-DE | 7 | zimní paralympijské hry     | <u>Paralympische Spiele</u>   <u>Sommer-Paralympics</u>   Paralympische Winterspiele   <u>Paralympische Sommerspiele</u>   <u>Paralympics</u> |
|       | 8 | letní paralympijské hry     | <u>Paralympische Spiele</u>   <u>Sommer-Paralympics</u>   <u>Paralympische Winterspiele</u>   Paralympische Sommerspiele   <u>Paralympics</u> |
|       | 9 | narcista                    | Narzissmus   <i>Narzißmus</i>   <u>narzisstisch</u>   |

Table 1: Examples of noisy term entries in the provided term collections. Inflected forms are **blue**, wrong translations of the source term are **red**, alternative spelling variants are *italic*, other possible translations are **bold**, other terms within one term entry are underlined.

get language. Typically preparing such term collections that are useful for MT (or the translation process in general) is the task of a professional translator or a terminologist. In this Shared Task, however, term collections sometimes contain entries that do not feature reciprocal translations of terms (e.g., see examples 1, 3, 7, and 8 in Table 1), have multiple translations per term (example 8 in Table 1), have multiple different terms merged into a single term entry (examples 2 and 9 in Table 1), have different spelling forms listed on the target side (examples 4 and 9 in Table 1), and have different inflections or spelling variants of source terms separated in different term entries (examples 5 and 6 in Table 1). Besides, unlike the common custom of providing terms and their translations in their dictionary forms (e.g., see examples of terms in EuroTermBank<sup>1</sup>, the InterActive Terminology for Europe<sup>2</sup>, the United Nations Terminology Database<sup>3</sup>, and other authoritative term banks), the terminologies provided for the Shared Task often contain entries where the source or the target language form is already inflected (examples 1, 2, and 6 in Table 1).

To reduce the noise present in the provided term collections, we performed filtering by discarding:

1. Term pairs that feature terms consisting of symbols other than digits, letters, apostrophes, white-spaces, and hyphen symbols. This filter allows to identify and discard expressions that do not represent terminology (e.g., full sentences, complete clauses, formulas, expressions consisting of terms and their acronyms

<sup>1</sup><https://eurotermbank.com/>

<sup>2</sup><https://iate.europa.eu/>

<sup>3</sup><https://unterm.un.org/>

within one term entry, etc.; see examples 1-6 and 8 in Table 2).

2. Term pairs where the source term is longer than the target term and the source term contains the target term as a sub-string (or vice versa). This filter is intended to discard term entries representing named entities that are written identically in both source and target languages, but for which one of the sides is incomplete (see examples 7 and 9 in Table 2).
3. Term pairs that represent general language (i.e., are too common). General language phrases are typically ambiguous and may require different translations based on surrounding context as well as external knowledge, which may not be available when translating. Therefore, it may be safer to allow the NMT model to handle the translation of general language phrases. We also do not want to burden the MT model too much with excessively annotated input data since longer segments are typically handled worse by NMT models than shorter segments (Neishi and Yoshinaga, 2019). To identify term entries that are too general, we apply an inverse document frequency (IDF) (Jones, 1972) filter (Pinnis, 2015a). As an example, this filter discarded from the term collections all term entries of the English term “spread” as it is a highly ambiguous word and according to the Collins EN-FR dictionary<sup>4</sup> it may have at least 20 distinct translations. Since the term collection

<sup>4</sup><https://www.collinsdictionary.com/dictionary/EN-FR/spread>

features just one possible translation without any added meta-data, it is safer to not use such terms (considering also the limitations of term recognition when working with emerging domains with scarce or no parallel data).

Besides filtering out the noisy term entries, the type and the quality of the term collections provided in the Shared Task also require selecting one among potentially many term translation equivalents. As noted before, this is typically done by a human, possibly, a domain expert. Nevertheless, we opt for two different strategies. If more than one translation equivalent is provided, it is fair to assume that they are all equally applicable. Thus we propose to select the first translation equivalent in the list. We refer to it as **1<sup>st</sup> Trg Term** term selection strategy. After analyzing the term collections, however, we conclude that it is not the case that all translation equivalents provided in the Shared Task are of equal quality. Therefore, we employ a statistical word alignment-based strategy to select the translation equivalent with the highest alignment score. To compute word alignments we use *eflomal*<sup>5</sup> (Östling and Tiedemann, 2016). We refer to it as **Alignment-based** term selection strategy.

Table 3 gives examples of terms selected by either of the term selection strategy. Examples illustrate that some translation equivalents are of equal quality (examples 1, 3, 6, and 7). However, selecting the first translation equivalent can sometimes give long (examples 6 and 8) or inadequate (example 4) translation equivalents. The Alignment-based term selection strategy also tends to select translation equivalents that are dictionary forms (example 2) instead of inflections.

## 2.2 Term Recognition

Having a term collection, the next task in the MT workflow is term recognition in a running text. Term recognition depending on the morphological typology of the source language and the nature of the domain can prove to be a complex task. Recognition involves term identification in its surface form, which for morphologically complex languages may be hindered by many surface forms a single word can take or by the level of form ambiguity in the case of morphologically impoverished languages (Bergmanis and Goldwater, 2018). To overcome issues posed by the morphology of the

natural language, one can use one of the many off-the-shelf morphological taggers to obtain contextually correct part-of-speech and lemma pairs for each token and perform term recognition on lemmatized collections and texts. We, however, opt for an alternative, a more rudimentary method utilizing language-specific stemmers to normalise the surface forms and do the term recognition on stemmed running text and term collections (Pinnis, 2015a,b). We opt for the stemmer-based approach because, in the production setting, stemming is faster than morphological tagging and has broader coverage for low-resource languages. Besides, to take full advantage of the morpho-syntactic information provided by morphological taggers, similar information must be provided by the term collection. However, as term collections of this Shared Task exemplify, expecting any meta-data is naive.

Last but not least, recognised forms must be word-sense disambiguated if more than one translation sense (i.e., term entry per source-side lexical form) is available. Word-sense disambiguation tools typically are lexicalized classifiers that are trained using large amounts of parallel data. However, the spirit of this Shared Task is MT using terminology for emerging domains where "*parallel data are hard to come by*"<sup>6</sup>. Thus we skip word-sense disambiguation and use just one word-sense per word form.

## 2.3 Integration of Terminology Constraints

In the day-to-day work of professional translators, terminologies are glossaries containing source language terms and their corresponding target language translations in their dictionary forms. Some of the previous work on terminology translation assumes that term entries are given in forms already inflected as required by the target morpho-syntactic context. Thus, such work focuses either on morphologically impoverished languages or is concerned with terminology translation in unrealistic scenarios. Either way, such methods are not relevant for the languages of the Shared Task because all of the target languages, with the exception of Chinese, are to some degree inflective languages.

There is, however, another body of work that addresses translation with terminology while accounting for morphological complexity of the target language (Exel et al., 2020; Niehues, 2021;

<sup>5</sup><https://github.com/robertostling/eflomal>

<sup>6</sup><http://www.statmt.org/wmt21/terminology-task.html>

|       | # | Source Term               | Target Term   |
|-------|---|---------------------------|---|
| EN-FR | 1 | shortness of breath       | essoufflé (e)   |
|       | 2 | Do the Five               | "5 gestes" "barrière" ""                                |
|       | 3 | Coronavirus (COVID-19)    | coronavirus (COVID-19)                                  |
| EN-RU | 4 | CDC                       | центры по контролю и профилактике заболеваний США (CDC) |
|       | 5 | active cases              | активные случаи [заболевания]                           |
|       | 6 | contagious                | заразный: передающийся при контакте                     |
| CS-DE | 7 | Lions Clubs International | Lions Clubs   |
|       | 8 | VIII. hlavový nerv        | Nervus vestibulocochlearis                              |
|       | 9 | Eli Lilly                 | Eli Lilly and Company                                   |

Table 2: Examples of terms removed by basic filtering.

Bergmanis and Pinnis, 2021). We base our submission on Bergmanis and Pinnis (2021) and employ target lemma annotations (TLA) to augment MT training data. An example of a sentence fragment annotated with TLA is "*infections*<sub>ls</sub> *инфекция*<sub>lt</sub> *result*<sub>lw</sub> *in*<sub>lw</sub> *mild*<sub>lw</sub> *symptoms*<sub>lw</sub>", where *ls*, *lt*, *lw* are factors indicating whether token is a source language term, a lemma (the dictionary form) of a target language term, or an ordinary source language word respectively. Systems trained on such data are equipped with a mechanism for passing soft terminology constraints at inference time. An essential property of MT systems trained using TLA is that they learn not just to copy but also to inflect the provided terminology constraints according to the target morphosyntactic context. Therefore, the translation of the sentence above, for example, "инфекций приводят к легким симптомам", contains the plural noun "инфекций" and not just the annotated singular form "инфекция".

### 3 Experiments

**Data.** We use all parallel data provided for the Shared Task for training, except for development data which we use to choose the best model for final submission. Although back-translated monolingual data could, in theory, improve the overall translation quality, we do not use it to train our systems because typically, the monolingual target data is selected based on its similarity with the target domain data. However, the scenario proposed for the Shared Task assumes that the domain is novel; thus, we aim to explore the merits of terminology translation and do not look for extra synthetic target domain data.

**MT Model and Training.** For system training, we use the *Marian toolkit* (Junczys-Dowmunt et al., 2018) because of its factored model functionality

developed within the scope of the *User-Focused Marian project*<sup>78</sup>. In this Shared Task, we train standard MT systems that mostly follow the Transformer (Vaswani et al., 2017) *base* model configuration. The only deviations from the standard configuration are 1) the use of source-side factors (we use factor embeddings of dimensionality 8 and concatenate them with word embeddings), 2) increased *-optimizer-dealy* (from 16 to 24), and 3) an increased maximum sequence length (from 128 to 196 tokens). These changes are necessary purely for TLA support during training and inference: increased sequence length accounts for longer input sequences due to TLA and terminology constraints, while increased optimizer delay compensates for fewer sentences fitting in workspace memory-based batch due to their increased maximum length.

### 4 Results

We trained one NMT system per translation direction and evaluated translation quality on the development sets using the terminology translation evaluation tool provided by the Shared Task<sup>9</sup> (ibn Alam et al., 2021). We compare the baseline translation scenario where no terms are annotated in the source text with improved scenarios where terms are annotated using term collections acquired using the different filtering and term translation equivalent selection strategies.

When analysing the lemmatized exact match accuracy, we must bear in mind that the evaluation data similarly to the term collections feature term

<sup>7</sup><https://marian-project.eu/>

<sup>8</sup><https://github.com/marian-cef/marian-examples/blob/forced-translation/forced-translation/docs/Experiments.md>

<sup>9</sup>[https://github.com/mahfuzibnalalam/terminology\\_evaluation](https://github.com/mahfuzibnalalam/terminology_evaluation)

|       | # | Source Term                | 1 <sup>st</sup> Trg Term  | Alignment-based                           |
|-------|---|----------------------------|---|---|
| EN-FR | 1 | disease outbreak           | apparition de maladie   | épidémie                                  |
|       | 2 | epidemiologist             | épidémiologistes  | épidémiologiste                           |
|       | 3 | wash your hands            | lavez-vous les mains  | laver les mains                           |
| EN-RU | 4 | Center for Disease Control | центры по контролю и профилактике заболеваний США                                   | Центр контроля и профилактики заболеваний |
|       | 5 | COVID-19 crisis            | кризиса, связанного с COVID-19  | кризис из-за COVID-19                     |
|       | 6 | health crisis              | кризисной ситуации, которая сложилась сегодня в сфере общественного здравоохранения | кризис здравоохранения                    |
| CS-DE | 6 | benzoylperoxid             | Dibenzoylperoxid  | Benzoylperoxid                            |
|       | 7 | alternativní medicína      | Alternativmedizin   | Alternative Medizin                       |
|       | 8 | AV ČR                      | Akademie der Wissenschaften der Tschechischen Republik                              | AV ČR                                     |

Table 3: Examples of terms selected by either 1<sup>st</sup> Trg Term or Alignment-based term selection strategy for term entries with multiple translation equivalents.

| EN-FR   | BLEU | Accuracy     | Window 2 | Window 3 | 1 - TERM |
|---|------|--------------|----------|----------|----------|
| Baseline                                      | 44.5 | 0.885        | 0.286    | 0.278    | 0.575    |
| 1 <sup>st</sup> Trg Term                      | 44.2 | 0.922        | 0.284    | 0.278    | 0.572    |
| 1 <sup>st</sup> Trg Term, IDF $\geq$ 5        | 44.3 | 0.905        | 0.281    | 0.275    | 0.574    |
| 1 <sup>st</sup> Trg Term, IDF $\geq$ 7        | 44.4 | 0.890        | 0.285    | 0.280    | 0.574    |
| <b>Alignment-based</b>                        | 44.6 | <b>0.936</b> | 0.291    | 0.285    | 0.576    |
| Alignment-based, IDF $\geq$ 5                 | 44.6 | 0.918        | 0.287    | 0.281    | 0.576    |
| Alignment-based, IDF $\geq$ 7                 | 44.4 | 0.896        | 0.284    | 0.278    | 0.574    |
| EN-RU   | BLEU | Accuracy     | Window 2 | Window 3 | 1 - TERM |
| Baseline                                      | 24.9 | 0.760        | 0.163    | 0.164    | 0.398    |
| 1 <sup>st</sup> Trg Term                      | 24.6 | 0.751        | 0.163    | 0.164    | 0.394    |
| 1 <sup>st</sup> Trg Term, IDF $\geq$ 5        | 25.0 | 0.815        | 0.176    | 0.176    | 0.401    |
| 1 <sup>st</sup> Trg Term, IDF $\geq$ 7        | 25.1 | 0.800        | 0.174    | 0.175    | 0.401    |
| <b>Alignment-based</b>                        | 25.0 | <b>0.833</b> | 0.174    | 0.175    | 0.403    |
| Alignment-based, IDF $\geq$ 5                 | 25.1 | 0.821        | 0.176    | 0.177    | 0.402    |
| Alignment-based, IDF $\geq$ 7                 | 25.2 | 0.792        | 0.171    | 0.172    | 0.401    |
| CS-DE   | BLEU | Accuracy     | Window 2 | Window 3 | 1 - TERM |
| Baseline                                      | 29.9 | 0.824        | 0.376    | 0.368    | 0.390    |
| 1 <sup>st</sup> Trg Term                      | 26.3 | 0.851        | 0.338    | 0.328    | 0.340    |
| 1 <sup>st</sup> Trg Term, IDF $\geq$ 5        | 29.5 | 0.852        | 0.371    | 0.361    | 0.378    |
| 1 <sup>st</sup> Trg Term, IDF $\geq$ 7        | 29.6 | 0.837        | 0.373    | 0.364    | 0.381    |
| Alignment-based                               | 26.8 | 0.849        | 0.346    | 0.337    | 0.348    |
| <b>Alignment-based, IDF<math>\geq</math>5</b> | 29.6 | <b>0.853</b> | 0.373    | 0.364    | 0.386    |
| Alignment-based, IDF $\geq$ 7                 | 29.7 | 0.837        | 0.375    | 0.366    | 0.389    |

Table 4: Development set results for baseline MT systems and systems using terminology integration. MT systems using terminology integration are named after the method used for terminology filtering. The numerically highest scores according the Lemmatized Exact-Match Accuracy (Accuracy) in data setting are **bold**. For detailed description of other metrics consult (ibn Alam et al., 2021).

entries 1) with more than one allowed synonymous translation equivalent (not counting different inflected forms), and 2) where different terms are merged into one entry (see examples in Section 2.1). This consequently means that the evaluation procedure 1) allows terminological ambiguity on the target side, 2) does not allow analysing terminology translation consistency, and 3) may depict a rough estimate of terminology translation accuracy. Therefore, we believe that the lemmatized exact match accuracy results should be analysed with a grain of salt.

That being said, the results in Table 4 show that the metric improves when using a term collection in all but one experiment. The fact that in overall the Alignment-based term collections show better overall translation results (in terms of BLEU) and also allow reaching the highest terminology translation accuracy results shows that relying on the first translation equivalent in a term entry is not a good idea.

We also see that the overall terminology translation quality is already relatively high for the baseline systems ranging from 76% for EN-RU till 88.5% for EN-FR. This makes us wonder whether the evaluated domain can be considered emerging as it features few novel terms and the majority are well handled by the baseline systems. To investigate further, we analysed whether the bilingual terminology that can be found in the development sets is also present in the training data of the NMT systems. We found that for CS-DE, 97.9% and 92.5% of such (unique) bilingual terms are featured at least once or 10 times in the training data respectively. The numbers are even higher if we analyse running terms (tokens) – 99.8% and 98.7% respectively. Since the terminology for EN-RU was human-created and not extracted from parallel data, it shows slightly lower results when analysing unique terms – 93.5% and 88.3% respectively. However, the situation is similar to CS-DE when analysing running terms – 99.1% and 97.8% of bilingual terms found in the EN-RU development data are also found in the training data at least once and 10 times respectively. Based on these findings, we believe that the validation data does not depict an emerging domain and does not help analysing terminology translation quality for emerging domains.

When analysing the overall translation quality (in terms of BLEU), we see that term filtering using

the IDF-based filter is crucial when relying on very noisy and automatically acquired term collections (as was the case with the CS-DE term collection). The results show that translation quality drops by 3 BLEU points when using the unfiltered term collections. This shows that too general (and ambiguous) terminology can be harmful and lower translation quality. The overall translation quality change is marginal for the translation directions that featured human-created term collections (EN-FR and EN-RU), however we do see an increase in terminology translation accuracy.

Our **final submission** consists of machine translations of Shared Task test sets provided by general-purpose MT systems that use dynamic terminology integration using TLA (Bergmanis and Pinnis, 2021). To translate our final submissions, term collections are filtered by basic filters (see Section 2.1) for EN-FR and EN-RU language pairs, while for CS-DE language pair, we also use IDF>5 filtering. We use the statistical word alignment term selection strategy for term entries with multiple translation equivalents for all language pairs. The development set results for the corresponding systems are marked bold in Table 4.

## 5 Discussion

**Shared Task.** Results of automatic metrics show that our baseline systems are already well equipped to translate the development and test sets regardless of their seemingly novel domain. Indeed – we found no statistically significant differences in scores measuring general translation quality between the baseline systems and systems with terminology integration. Preliminary test results suggest a similar pattern in other submissions (c.f. results of submissions by Prompt). The only seemingly meaningful differences are in metrics specifically targeting terminology integration. These results are in stark contrast with previous work (Exel et al., 2020; Niehues, 2021; Bergmanis and Pinnis, 2021) which report significant improvements not only on terminology use targeted metrics but also on metrics measuring general translation quality. This disparity suggests that test data is *not* from an emerging or novel domain (at least as far MT systems trained on the training data provided are concerned). Considering this shortcoming, together with the visibility of WMT Shared Tasks, these results pose a risk of misrepresenting the problem the Shared Task was set out to research. The outcome might be

unintended downplaying of the role of terminology translation for technical domains, which could lead to diminishing interest in terminology translation from the MT research community.

**Term Collections in MT.** Tables 1, 2, and 3 of Section 2 provide numerous examples of problems present in the provided term collections. We believe that these examples illustrate the understanding of the purpose and desired qualities of a term collection not just of those individuals involved in preparing term collections for the Shared Task but also to a broader community of translation professionals. Many of the problematic examples suggest that the shift from human-readable to machine-readable term collections is not there yet, or that it has happened rather formally by merely reformatting the for-human-made term collections into neater TSV-formatted files. While having TSV-formatted files helps for the file to be machine-readable, it does not make the content machine-usable. The standards for-machine-made term collections have to be higher than those made for humans. At least as long as there is no sophisticated intelligence in the MT workflow that is on par with humans to recover from the irregularities and noise present in the term collections typically made for humans. Likewise, the encyclopedia-style entries explaining a concisely coined concept of the source language using a whole sentence to define it in the target language are still present in for-human-made term collections, but they are of no use for current MT systems.

If translation with terminology is supposed to improve MT for novel domains, the term collections, being the supposed source of the expected improvement, have to be of a higher quality than the MT systems they are intended to improve.

## Acknowledgements

This research has been supported by the User-focused Marian project which is co-financed by the European Union's Connecting Europe Facility programme (Action number: 2019-EU-IA-0045).

## References

Md Mahfuz ibn Alam, Antonios Anastasopoulos, Laurent Besacier, James Cross, Matthias Gallé, Philipp Koehn, and Vassilina Nikoulina. 2021. [On the evaluation of machine translation for terminology consistency](#).

Toms Bergmanis and Sharon Goldwater. 2018. Context sensitive neural lemmatization with lematus. 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 1391–1400.

- Toms Bergmanis and Mārcis Pinnis. 2021. Facilitating terminology translation with target lemma annotations. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 3105–3111.
- Miriam Exel, Bianka Buschbeck, Lauritz Brandt, and Simona Doneva. 2020. [Terminology-constrained neural machine translation at SAP](#). In *Proceedings of the 22nd Annual Conference of the European Association for Machine Translation*, pages 271–280, Lisboa, Portugal. European Association for Machine Translation.
- Karen Sparck Jones. 1972. A statistical interpretation of term specificity and its application in retrieval. *Journal of documentation*.
- Marcin Junczys-Dowmunt, Roman Grundkiewicz, Tomasz Dwojak, Hieu Hoang, Kenneth Heafield, Tom Neckermann, Frank Seide, Ulrich Germann, Alham Fikri Aji, Nikolay Bogoychev, André F. T. Martins, and Alexandra Birch. 2018. [Marian: Fast Neural Machine Translation in C++](#). In *Proceedings of ACL 2018, System Demonstrations*, pages 116–121, Melbourne, Australia. Association for Computational Linguistics.
- Masato Neishi and Naoki Yoshinaga. 2019. On the relation between position information and sentence length in neural machine translation. In *Proceedings of the 23rd Conference on Computational Natural Language Learning (CoNLL)*, pages 328–338.
- Jan Niehues. 2021. Continuous learning in neural machine translation using bilingual dictionaries. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 830–840.
- Robert Östling and Jörg Tiedemann. 2016. [Efficient word alignment with Markov Chain Monte Carlo](#). *Prague Bulletin of Mathematical Linguistics*, 106:125–146.
- Mārcis Pinnis. 2015a. [Dynamic terminology integration methods in statistical machine translation](#). In *Proceedings of the 18th Annual Conference of the European Association for Machine Translation*, pages 89–96, Antalya, Turkey.
- Mārcis Pinnis. 2015b. *Terminology Integration in Statistical Machine Translation*. Ph.D. thesis, PhD thesis, Faculty of Computing, University of Latvia, Riga, Latvia.
- 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 *Advances in Neural Information Processing Systems*, pages 5998–6008.