

## 20 Advanced Topics 3: Sub-word MT

Up until this point, we have treated words as the atomic unit that we are interested in training on. However, this has the problem of being less robust to low-frequency words, which is particularly a problem for neural machine translation systems that have to limit their vocabulary size for efficiency purposes. In this chapter, we first discuss a few of the phenomena that cannot be easily tackled by pure word-based approaches but can be handled if we look at the word’s characters, and then discuss some methods to handle these phenomena.

### 20.1 Sub-word Phenomena

There are a large number of examples in which subword structure can be useful for translation systems, a few of which are outlined in Figure 59.

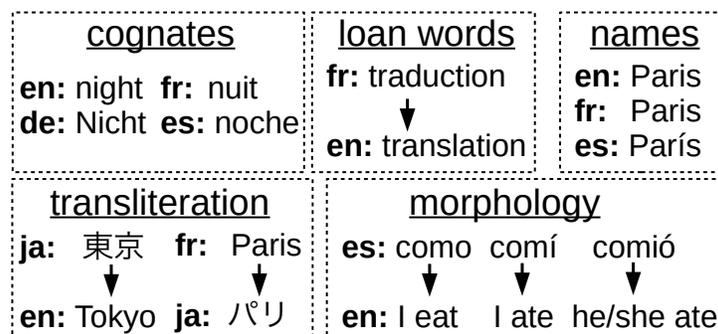


Figure 59: An example of phenomena for which sub-word information is useful.

For these words, the surface form of the words shows some non-random similarity between the source and target languages. In the extreme, we can think of examples where the words are exactly the same between the source and target sentences. For example, this is common when translating proper names, such as the “Paris” in the top-right of the figure. This can be handled by copying words directly from the source to target, as described in previous chapters (Section 19.3, [9])

However, there are many cases where words are similar, but not exactly the same. For example, this is true for **cognates**, words which share a common origin but have diverged at some point in the evolution of respective languages. For example, the word “night” in English is shared in some form with the words “Nacht” in German, “nuit” in French, and “noche” in Spanish. These reflect the fact that “night” in English descended from “nakht” in proto-Germanic (shared with German), which in turn descended from “nekw̥t” in proto-Indo-European (shared with all four languages above) [19, 10]. This is also true for **loan words**, which are not a result of gradual change in language, but are instead borrowed as-is from another language. One example of a loan word is “translation” (as well as most other words that end with “-ion” in English), which was borrowed from its French counterpart. While these words are not exactly the same, precluding the use of a copy mechanism, models that can appropriately handle these similarities could improve accuracy for these phenomena.

Another phenomenon that is worth noting is **transliteration**. Transliteration is the process of converting words with identical or similar pronunciations from one script to another.

For example, Japanese is written in a different script than European languages, and thus words such as “Tokyo” and “Paris”, which are pronounced similarly in both languages, must nevertheless be converted appropriately.

Finally, **morphology** is another notable phenomenon that affects, and requires handling of, subword structure. Morphology is the systematic changing of word forms according to their grammatical properties such as tense, case, gender, part of speech, and others. In the example above, the Spanish verb changes according to the tense (present or past) as well as the person of the subject (first or third). These sorts of systematic changes are not captured by word-based models, but can be captured by models that are aware of some sort of subword structure.

In the following sections, we will see how to design models to handle these phenomena.

## 20.2 Character-based Translation

The first, and simplest, method for moving beyond words as the atomic unit for translation is to perform **character-based translation**, simply using characters to perform translation between the. In other words, instead of treating words as the symbols in  $F$  and  $E$ , we simply treat characters as the symbols in these sequences.

### 20.2.1 Symbolic character-based models

This method was proposed by [23] in the context of phrase-based statistical machine translation. In this method, the entire phrase-based pipeline, from alignment to phrase extraction to decoding, is applied to character strings instead. This showed a certain amount of success for very similar languages that had lots of cognates and strong correspondence between the words in their vocabularies, such as Catalan and Spanish.

However, there are two difficulties with performing translations over characters like this. First, because the correspondence between individual characters in the source and target sentences can be tenuous, it is common that models such as the IBM models, which generally prefer one-to-one alignments fail when there is not a strong one-to-one or temporally consistent alignment. This difficulty can be alleviated somewhat by instead performing many-to-many alignment, aligning strings of characters to other strings of characters [18]. This allows the model to extract phrases that work across multiple granularities.

### 20.2.2 Neural character-based models

Within the framework of neural MT, there are also methods to perform character-based translation. Because neural MT methods inherently capture long-distance context through the use of recurrent neural networks, competitive results can be achieved without explicit segmentation into phrases [4].

There are also a number of methods that attempt to go further, creating models that are character-aware, but nonetheless incorporate the idea that we would like to combine characters into units that are approximately the same size as a word. A first example is the idea of **pyramidal encoders** [2]. The idea behind this method is that we have multiple levels of stacked encoders where each successive level of encoding uses a coarser granularity. For example, the pyramidal encoder shown on the left side of Figure 60 takes in every character at its first layer, but each successive layer only takes the output of the first layer every two

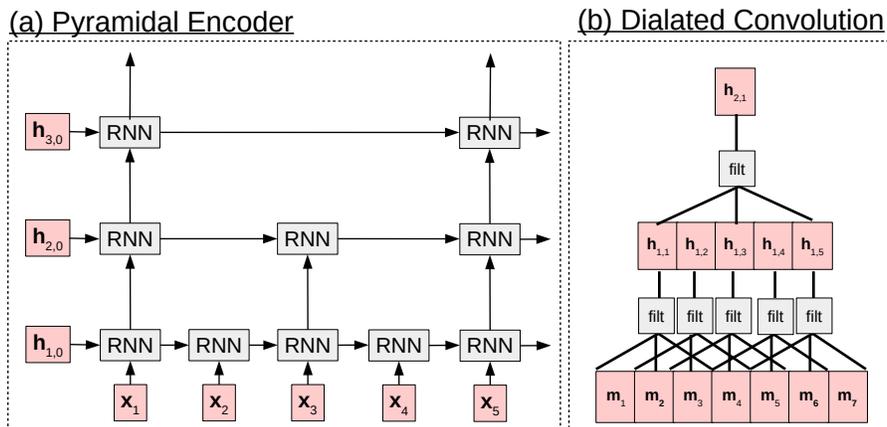


Figure 60: Encoders that reduce the resolution of input.

time steps, reducing the resolution of the output by two. A very similar idea in the context of convolutional networks is **dilated convolutions** [24], which perform convolutions that skip time steps in the middle, as shown in the right side of Figure 60.

One other important consideration for character-based models (both neural and symbolic) is their computational burden. With respect to neural models, one very obvious advantage from the computational point of view is that using characters limits the size of the output vocabulary, reducing the computational bottleneck in calculating large softmaxes over a large vocabulary of words. On the other hand, the length of the source and target sentence will be significantly longer (multiplied by the average length of a word), which means that the number of RNN time steps required for each sentence will increase significantly. In addition, [16] report that a larger hidden layer size is necessary to efficiently capture intra-word dynamics for character-based models, resulting in an increase in a further increase in computation time.

### 20.3 Hybrid Word-character Models

The previous section covered models which simply map characters one-by-one into the target sequence with no concept of “words” or “tokenization”. It is also possible to create models that work on the word level most of the time, but fall back to the character level when appropriate.

One example of this is models for transliteration, where the model decides to translate character-by-character only if it decides that the word should be transliterated. For example, [11] come up with a model that identifies **named entities** (e.g. people or places) in the source text using a standard named entity recognizer, and decides whether and how to transliterate it. This is a difficult problem because some entities may require a mix of translation and transliteration; for example “Carnegie Mellon University” is a named entity, but while “Carnegie” and “Mellon” may be transliterated, “University” will often be translated into the appropriate target language. [6] adapt this to be integrated in a phrase-based translation system.

Word-character hybrid models have also been implemented within the neural MT paradigm. One method, proposed by [14], uses word boundaries to specify the granularity of the encod-

ing, so that each word is encoded by a single vector (c.f. the fixed-length encoding of the previous section). The encoding of the word can be performed using a bi-directional LSTM over its characters, with the final states in each direction being concatenated into the word representation. On the decoding side, an RNN generates word representations one-by-one, then the over-arching word representation is used to generate the target word one character at a time.

Alternatively, it is also possible to use this character-based model only for words that do not exist in the vocabulary [15]. This is beneficial, as it allows the model to directly generate words that are in its vocabulary, presumably more frequent and well-learned words, while falling back to a character-based representation when it is necessary.

## 20.4 Rethinking Tokenization

Another way to improve translation of lower-frequency words is to come up with a new standard for tokenizing characters into words that splits words of lower frequency into smaller units.

One extreme example of this would be the tokenization of sentences in languages that do not explicitly mark words with white space delimiting word boundaries. These languages include Chinese, Japanese, Thai, and several others. In these languages, it is common to create a **word segmenter** trained on data manually annotated with word boundaries, then apply this to the training and testing data for the machine translation system. In these languages, the accuracy of word segmentation has a large impact on results, with poorly segmented words often being translated incorrectly, often as unknown words. In particular, [3] note that it is extremely important to have a consistent word segmentation algorithm that usually segments words into the same units regardless of context. This is due to the fact that any differences in segmentation between the MT training data and the incoming test sentence may result in translation rules or neural net statistics not appropriately covering the mis-segmented word. As a result, it may be preferable to use a less accurate but more consistent segmentation when such a trade-off exists.

One thing to note is that the work in the previous paragraph is entirely *supervised* segmentation, where we have data manually annotated with word boundaries. It is also possible to perform **unsupervised word segmentation**, where original corpora (consisting of character strings)  $\mathcal{F}$  and  $\mathcal{E}$  are provided to a training algorithm, and boundaries splitting these into segmented corpora  $\bar{\mathcal{F}}$  and  $\bar{\mathcal{E}}$  are learned directly from raw text. The most prominent method for unsupervised word segmentation [8, 17] attempts to maximize the probability of this raw text using a word-based language model:

$$\log P_{\text{LM}}(\bar{\mathcal{E}}) = \sum_{\bar{E} \in \bar{\mathcal{E}}} \sum_{t=1}^{|\bar{E}|} \log P(\bar{e}_t | \bar{e}_1^{t-1}), \quad (182)$$

$$\text{s.t. } \forall_{\langle E, \bar{E} \rangle \in \langle \mathcal{E}, \bar{\mathcal{E}} \rangle} E = \text{concat}(\bar{E}). \quad (183)$$

These models additionally add a bias against the vocabulary of the model getting too large, and describe a method to search for this maximum likelihood solution, usually with an iterative procedure that re-samples the segmentations of sentences one-by-one (**Gibbs sampling**).

As a simple, faster, but potentially less accurate method for unsupervised word segmentation, [20] have recently proposed a method based on a technique called “byte pair encoding

(BPE)”, which is fast and relatively effective. The method finds segmentations by starting with an initial segmentation  $E = \bar{E}$ , where each token is its own character, then iteratively combining together the most frequent 2-gram in the corpus. The intuition behind the model is that more frequent strings (with sufficient training data) should be treated as a single unit, while less frequent strings should be split together into their component parts to prevent sparsity.

In unsupervised word segmentation algorithms, it can also be useful to segment both sides of the corpus with a single model at the same time [20]. The reason for this is because many words share character strings, as noted in Section 20.1. This would basically entail running the same unsupervised word segmentation, only over both the source and target corpora  $\mathcal{F}$  and  $\mathcal{E}$  at the same time, instead of running them simultaneously.

## 20.5 Models of Morphology

As mentioned above, morphology is the phenomenon of word forms being transformed in a consistent way corresponding to their syntactic function. Because these transformations are consistent across words, models that can capture these transformations appropriately can be used to more accurately translate from or to inflected word forms. One common way to incorporate morphological information into MT is through a three step process of (1) *analysis*, which calculates various features of the word to be translated such as the lemma or morphological tags, (2) *translation*, which translates this factored representation making various independence assumptions, and (3) *generation*, which creates the surface form from each of these factors [12].

For the analysis and generation steps, the classical way of creating an analyzer is through manually created rules with expressed as finite-state automata that take in the characters of the word one at a time and output possible analyses [1]. This allows us to create a concise list of analysis candidates for words in the dictionary with relatively high precision, and these methods are still widely used for a number of languages for which good morphological analyzers exist. However, these methods rely on a linguist capable of making rules, an extensive dictionary in the language, and lack the ability to disambiguate hypotheses based on context. As a result, there are also methods based on symbolic [7] or neural [21] approaches to perform morphological analysis and contextual disambiguation.

Once we have an analysis of one or both sides of the translation pair, we can proceed to do translation. One easy and effective way to do so when using a morphologically rich language on the source side is simply to analyze the source corpus, split the words into their lemmas and morphological tags, and use this split data as input to a normal MT system [13]. This is very similar to the subword splitting approach in Section 20.4, so it is worth considering the differences. In the case of **concatenative morphology**, where a word can be viewed as the concatenation of its morphemes (e.g. the word “undecided” can be viewed as the concatenation of “un+decide+d”), subword splitting methods may be sufficient. However, there are also more difficult cases such as **infix morphology**, where the inner parts of a word are changed due to morphological processes (e.g. in English “goose” and its plural “geese”). In these more complicated cases, normalizing to the lemma can be an effective way to increase the generalization capabilities of the translation model.

On the other hand, if we have rich morphology on the target side, we can also think of converting our target corpus into a sequence of lemmas and morphological tags and translating

into these [22, 5]. Before showing the result to the end user, we need to generate the surface form from these lemmas and tags (e.g. going from “goose +PLURAL” to “geese”). Similarly, this process can be done with a rule-based generation model created by a linguist, or through data-driven approaches, much like morphological analysis. One thing to note is that in general, translation into morphologically rich languages is considered more difficult than translation from morphologically rich languages to a morphologically poor language such as English. The reason for this is twofold. First, morphologically rich languages tend to require more long-distance agreement between the forms of words; similarly to how the subject being first, second, or third person affects the conjugation of the verb (e.g. “I run” vs. “he runs”), in morphologically rich languages it is not uncommon for the person, case, gender, or other information to be matched between different words in the sentence. Second, by adding an extra generation step, it is common for error propagation to occur, with errors in the first step cascading to errors in the second step.

## 20.6 Further Reading

There are a number of additional topics related to sub-word models that interested readers can examine:

**Factored translation models:** Factored translation models split a word up into several factors and translate them given some independence assumptions on what factors influence others [12]. For example, a word may be split into a lemma factor, a tense factor, and a plural factor, each of which would be translated independently, then combined together in a final generation step. This is different from the simple pre/post-processing approaches described above in that these various factors are tightly integrated within the translation model, and translation is still performed on a word-by-word basis.

## 20.7 Exercise

As an exercise, you could try to either

1. Train a character-based neural machine translation system using your existing code. Note its advantages and disadvantages compared to your existing word-based model, including training speed and accuracy.
2. Train a system using byte-pair encoding. This would entail using a

## References

- [1] Kenneth R. Beesley and Lauri Karttunen. *Finite State Morphology*. CSLI Studies in Computational Linguistics. CSLI Publications, 2003.
- [2] William Chan, Navdeep Jaitly, Quoc Le, and Oriol Vinyals. Listen, attend and spell: A neural network for large vocabulary conversational speech recognition. In *Proceedings of the International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, pages 4960–4964. IEEE, 2016.
- [3] Pi-Chuan Chang, Michel Galley, and Christopher D. Manning. Optimizing Chinese word segmentation for machine translation performance. In *Proceedings of the 3rd Workshop on Statistical Machine Translation (WMT)*, pages 224–232, 2008.

- [4] Junyoung Chung, Kyunghyun Cho, and Yoshua Bengio. A character-level decoder without explicit segmentation for neural machine translation. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 1693–1703, 2016.
- [5] Ann Clifton and Anoop Sarkar. Combining morpheme-based machine translation with post-processing morpheme prediction. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics (ACL)*, 2011.
- [6] Nadir Durrani, Hassan Sajjad, Hieu Hoang, and Philipp Koehn. Integrating an unsupervised transliteration model into statistical machine translation. In *Proceedings of the 14th European Chapter of the Association for Computational Linguistics (EACL)*, pages 148–153, 2014.
- [7] Greg Durrett and John DeNero. Supervised learning of complete morphological paradigms. In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*, pages 1185–1195, 2013.
- [8] Sharon Goldwater, Thomas L. Griffiths, and Mark Johnson. A Bayesian framework for word segmentation: Exploring the effects of context. *Cognition*, 112(1), 2009.
- [9] Jiatao Gu, Zhengdong Lu, Hang Li, and Victor O.K. Li. Incorporating copying mechanism in sequence-to-sequence learning. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 1631–1640, 2016.
- [10] Douglas Harper et al. Online etymology dictionary, 2001.
- [11] Kevin Knight and Jonathan Graehl. Machine transliteration. *Computational Linguistics*, 24(4):599–612, 1998.
- [12] Philipp Koehn and Hieu Hoang. Factored translation models. In *Proceedings of the 2007 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning (EMNLP-CoNLL)*, 2007.
- [13] Young-Suk Lee. Morphological analysis for statistical machine translation. In *Proceedings of the 2004 Human Language Technology Conference of the North American Chapter of the Association for Computational Linguistics (HLT-NAACL)*, pages 57–60, 2004.
- [14] Wang Ling, Isabel Trancoso, Chris Dyer, and Alan W Black. Character-based neural machine translation. *arXiv preprint arXiv:1511.04586*, 2015.
- [15] Minh-Thang Luong and Christopher D. Manning. Achieving open vocabulary neural machine translation with hybrid word-character models. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 1054–1063, Berlin, Germany, August 2016. Association for Computational Linguistics.
- [16] Tomáš Mikolov, Ilya Sutskever, Anoop Deoras, Hai-Son Le, and Stefan Kombrink. Subword language modeling with neural networks.
- [17] Daichi Mochihashi, Takeshi Yamada, and Naonori Ueda. Bayesian unsupervised word segmentation with nested Pitman-Yor modeling. In *Proceedings of the 47th Annual Meeting of the Association for Computational Linguistics (ACL)*, 2009.
- [18] Graham Neubig, Taro Watanabe, Shinsuke Mori, and Tatsuya Kawahara. Machine translation without words through substring alignment. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 165–174, 2012.
- [19] Oxford English Oxford. *Oxford English Dictionary*. Oxford: Oxford University Press, 2009.
- [20] Rico Sennrich, Barry Haddow, and Alexandra Birch. Neural machine translation of rare words with subword units. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 1715–1725, 2016.

- [21] Qinlan Shen, Daniel Clothiaux, Emily Tagtow, Patrick Littell, and Chris Dyer. The role of context in neural morphological disambiguation. In *Proceedings of the 26th International Conference on Computational Linguistics (COLING)*, pages 181–191, 2016.
- [22] Kristina Toutanova, Hisami Suzuki, and Achim Ruopp. Applying morphology generation models to machine translation. In *Proceedings of the 46th Annual Meeting of the Association for Computational Linguistics (ACL)*, pages 514–522, 2008.
- [23] David Vilar, Jan-T. Peter, and Hermann Ney. Can we translate letters? In *Proceedings of the 2nd Workshop on Statistical Machine Translation (WMT)*, pages 33–39, 2007.
- [24] Fisher Yu and Vladlen Koltun. Multi-scale context aggregation by dilated convolutions. 2016.