

What do Neural Machine Translation Models Learn about Morphology?

Yonatan Belinkov¹ Nadir Durrani² Fahim Dalvi² Hassan Sajjad² James Glass¹

¹MIT Computer Science and Artificial Intelligence Laboratory, Cambridge, MA 02139, USA
{belinkov, glass}@mit.edu

²Qatar Computing Research Institute, HBKU, Doha, Qatar
{ndurrani, faimaduddin, hsajjad}@qf.org.qa

Abstract

Neural machine translation (MT) models obtain state-of-the-art performance while maintaining a simple, end-to-end architecture. However, little is known about what these models learn about source and target languages during the training process. In this work, we analyze the representations learned by neural MT models at various levels of granularity and empirically evaluate the quality of the representations for learning morphology through extrinsic part-of-speech and morphological tagging tasks. We conduct a thorough investigation along several parameters: word-based vs. character-based representations, depth of the encoding layer, the identity of the target language, and encoder vs. decoder representations. Our data-driven, quantitative evaluation sheds light on important aspects in the neural MT system and its ability to capture word structure.¹

1 Introduction

Neural network models are quickly becoming the predominant approach to machine translation (MT). Training neural MT (NMT) models can be done in an end-to-end fashion, which is simpler and more elegant than traditional MT systems. Moreover, NMT systems have become competitive with, or better than, the previous state-of-the-art, especially since the introduction of sequence-to-sequence models and the attention mechanism (Bahdanau et al., 2014; Sutskever et al., 2014). The improved translation quality is often attributed to better handling of non-local dependencies and morphology generation (Luong

and Manning, 2015; Bentivogli et al., 2016; Toral and Sánchez-Cartagena, 2017).

However, little is known about what and how much these models learn about each language and its features. Recent work has started exploring the role of the NMT encoder in learning source syntax (Shi et al., 2016), but research studies are yet to answer important questions such as: (i) what do NMT models learn about word morphology? (ii) what is the effect on learning when translating into/from morphologically-rich languages? (iii) what impact do different representations (character vs. word) have on learning? and (iv) what do different modules learn about the syntactic and semantic structure of a language? Answering such questions is imperative for fully understanding the NMT architecture. In this paper, we strive towards exploring (i), (ii), and (iii) by providing quantitative, data-driven answers to the following specific questions:

- Which parts of the NMT architecture capture word structure?
- What is the division of labor between different components (e.g. different layers or encoder vs. decoder)?
- How do different word representations help learn better morphology and modeling of infrequent words?
- How does the target language affect the learning of word structure?

To achieve this, we follow a simple but effective procedure with three steps: (i) train a neural MT system on a parallel corpus; (ii) use the trained model to extract feature representations for words in a language of interest; and (iii) train a classifier using extracted features to make predictions

¹Our code is available at <https://github.com/boknilev/nmt-repr-analysis>.

for another task. We then evaluate the quality of the trained classifier on the given task as a proxy to the quality of the extracted representations. In this way, we obtain a quantitative measure of how well the original MT system learns features that are relevant to the given task.

We focus on the tasks of part-of-speech (POS) and full morphological tagging. We investigate how different neural MT systems capture POS and morphology through a series of experiments along several parameters. For instance, we contrast word-based and character-based representations, use different encoding layers, vary source and target languages, and compare extracting features from the encoder vs. the decoder.

We experiment with several languages with varying degrees of morphological richness: French, German, Czech, Arabic, and Hebrew. Our analysis reveals interesting insights such as:

- Character-based representations are much better for learning morphology, especially for low-frequency words. This improvement is correlated with better BLEU scores. On the other hand, word-based models are sufficient for learning the structure of common words.
- Lower layers of the encoder are better at capturing word structure, while deeper networks improve translation quality, suggesting that higher layers focus more on word meaning.
- The target language impacts the kind of information learned by the MT system. Translating into morphologically-poorer languages leads to better source-side word representations. This is partly, but not completely, correlated with BLEU scores.
- The neural decoder learns very little about word structure. The attention mechanism removes much of the burden of learning word representations from the decoder.

2 Methodology

Given a source sentence $s = \{w_1, w_2, \dots, w_N\}$ and a target sentence $t = \{u_1, u_2, \dots, u_M\}$, we first generate a vector representation for the source sentence using an encoder (Eqn. 1) and then map this vector to the target sentence using a decoder (Eqn. 2) (Sutskever et al., 2014):

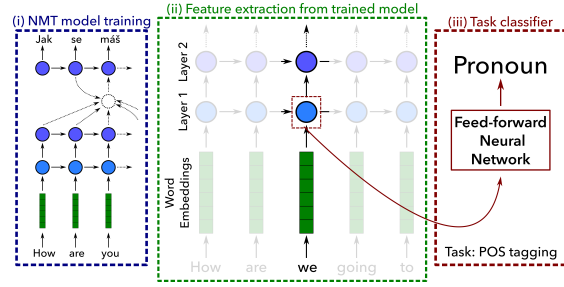


Figure 1: Illustration of our approach: (i) NMT system trained on parallel data; (ii) features extracted from pre-trained model; (iii) classifier trained using the extracted features. Here a POS tagging classifier is trained on features from the first hidden layer.

$$\text{ENC} : s = \{w_1, w_2, \dots, w_N\} \mapsto \mathbf{s} \in \mathbb{R}^k \quad (1)$$

$$\text{DEC} : \mathbf{s} \in \mathbb{R}^k \mapsto t = \{u_1, u_2, \dots, u_M\} \quad (2)$$

In this work, we use long short-term memory (LSTM) (Hochreiter and Schmidhuber, 1997) encoder-decoders with attention (Bahdanau et al., 2014), which we train on parallel data.

After training the NMT system, we freeze the parameters of the encoder and use ENC as a feature extractor to generate vectors representing words in the sentence. Let $\text{ENC}_i(s)$ denote the encoded representation of word w_i . For example, this may be the output of the LSTM after word w_i . We feed $\text{ENC}_i(s)$ to a neural classifier that is trained to predict POS or morphological tags and evaluate the quality of the representation based on our ability to train a good classifier. By comparing the performance of classifiers trained with features from different instantiations of ENC, we can evaluate what MT encoders learn about word structure. Figure 1 illustrates this process. We follow a similar procedure for analyzing representation learning in DEC.

The classifier itself can be modeled in different ways. For example, it may be an LSTM over outputs of the encoder. However, as we are interested in assessing the quality of the representations learned by the MT system, we choose to model the classifier as a simple feed-forward neural network with one hidden layer and a ReLU non-linearity. Arguably, if the learned representations are good, then a non-linear classifier should be able to extract useful information from them.² We empha-

²We also experimented with a linear classifier and observed similar trends to the non-linear case, but overall lower results; Qian et al. (2016b) reported similar findings.

	Ar	De	Fr	Cz
	Gold/Pred	Gold/Pred	Pred	Pred
Train Tokens	0.5M/2.7M	0.9M/4M	5.2M	2M
Dev Tokens	63K/114K	45K/50K	55K	35K
Test Tokens	62K/16K	44K/25K	23K	20K
POS Tags	42	54	33	368
Morph Tags	1969	214	–	–

Table 1: Statistics for annotated corpora in Arabic (Ar), German (De), French (Fr), and Czech (Cz).

size that our goal is not to beat the state-of-the-art on a given task, but rather to analyze what NMT models learn about morphology. The classifier is trained with a cross-entropy loss; more details about its architecture are given in the supplementary material (appendix A.1).

3 Data

Language pairs We experiment with several language pairs, including morphologically-rich languages, that have received relatively significant attention in the MT community. These include Arabic-, German-, French-, and Czech-English pairs. To broaden our analysis and study the effect of having morphologically-rich languages on both source and target sides, we also include Arabic-Hebrew, two languages with rich and similar morphological systems, and Arabic-German, two languages with rich but different morphologies.

MT data Our translation models are trained on the WIT³ corpus of TED talks (Cettolo et al., 2012; Cettolo, 2016) made available for IWSLT 2016. This allows for comparable and cross-linguistic analysis. Statistics about each language pair are given in Table 1 (under Pred). We use official dev and test sets for tuning and testing. Reported figures are the averages over test sets.

Annotated data We use two kinds of datasets to train POS and morphological classifiers: gold-standard and predicted tags. For predicted tags, we simply used freely available taggers to annotate the MT data. For gold tags, we use gold-annotated datasets. Table 1 provides statistics for datasets with gold and predicted tags; see the supplementary material (appendix A.2) for more details about taggers and gold data. We train and test our classifiers on predicted annotations, and similarly on gold annotations, when we have them. We report both results wherever available.

	Gold	Pred	BLEU
	Word/Char	Word/Char	Word/Char
Ar-En	80.31/93.66	89.62/95.35	24.7/28.4
Ar-He	78.20/92.48	88.33/94.66	9.9/10.7
De-En	87.68/94.57	93.54/94.63	29.6/30.4
Fr-En	–	94.61/95.55	37.8/38.8
Cz-En	–	75.71/79.10	23.2/25.4

Table 2: POS accuracy on gold and predicted tags using word-based and character-based representations, as well as corresponding BLEU scores.

4 Encoder Analysis

Recall that after training the NMT system we freeze its parameters and use it only to generate features for the POS/morphology classifier. Given a trained encoder ENC and a sentence s with POS/morphology annotation, we generate word features $ENC_i(s)$ for every word in the sentence. We then train a classifier that uses the features $ENC_i(s)$ to predict POS or morphological tags.

4.1 Effect of word representation

In this section, we compare different word representations extracted with different encoders. Our word-based model uses a word embedding matrix which is initialized randomly and learned with other NMT parameters. For a character-based model we adopt a convolutional neural network (CNN) over character embeddings that is also learned during training (Kim et al., 2015); see appendix A.1 for specific settings. In both cases we run the encoder over these representations and use its output $ENC_i(s)$ as features for the classifier.

Table 2 shows POS tagging accuracy using features from different NMT encoders. Char-based models always generate better representations for POS tagging, especially in the case of morphologically-richer languages like Arabic and Czech. We observed a similar pattern in the full morphological tagging task. For example, we obtain morphological tagging accuracy of 65.2/79.66 and 67.66/81.66 using word/char-based representations from the Arabic-Hebrew and Arabic-English encoders, respectively.³ The superior morphological power of the char-based model also manifests in better translation quality (measured by BLEU), as shown in Table 2.

³The results are not far below dedicated taggers (e.g. 95.1/84.1 on Arabic POS/morphology (Pasha et al., 2014)), indicating that NMT models learn quite good representations.




Figure 2: POS and morphological tagging accuracy of word-based and character-based models per word frequency in the training data. Best viewed in color.




Figure 3: Improvement in POS/morphology accuracy of character-based vs. word-based models for words unseen/seen in training, and for all words.




Figure 4: Increase in POS accuracy with char- vs. word-based representations per tag frequency in the training set; larger bubbles reflect greater gaps.

Impact of word frequency Let us look more closely at an example case: Arabic POS and morphological tagging. Figure 3 shows the effect of using word-based vs. char-based feature representations, obtained from the encoder of the Arabic-Hebrew system (other language pairs exhibit similar trends). Clearly, the char-based model is superior to the word-based one. This is true for the overall accuracy (+14.3% in POS, +14.5% in morphology), but more so in OOV words (+37.6% in POS, +32.7% in morphology). Figure 2 shows that the gap between word-based and char-based representations increases as the frequency of the word in the training data decreases. In other words, the more frequent the word, the less need there is for character information. These findings make intuitive sense: the char-based model is able to learn character n-gram patterns that are important for identifying word structure, but as the word becomes more frequent the word-based model has seen enough examples to make a decision.

Analyzing specific tags In Figure 5 we plot confusion matrices for POS tagging using word-based and char-based representations (from Arabic encoders). While the char-based representations are overall better, the two models still share similar misclassified tags. Much of the confusion comes from wrongly predicting nouns (NN, NNP). In the word-based case, relatively many tags with determiner (DT+NNP, DT+NNPS, DT+NNS, DT+VBG) are wrongly predicted as non-determined nouns (NN, NNP). In the char-based case, this hardly happens. This suggests that the char-based representations are predictive of the presence of a determiner, which in Arabic is expressed as the prefix “Al-” (the definite article), a pattern easily captured by a char-based model.

In Figure 4 we plot the difference in POS accuracy when moving from word-based to char-based representations, per POS tag frequency in the training data. Tags closer to the upper-right corner occur more frequently in the training set and are

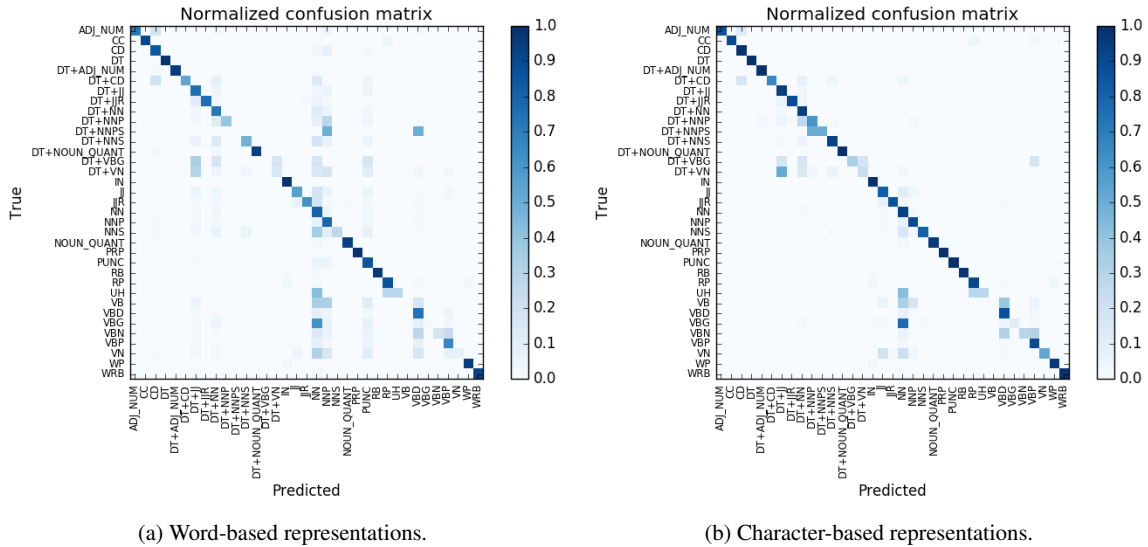


Figure 5: Confusion matrices for POS tagging using word-based and character-based representations.

better predicted by char-based compared to word-based representations. There are a few fairly frequent tags (in the middle-bottom part of the figure) whose accuracy does not improve much when moving from word- to char-based representations: mostly conjunctions, determiners, and certain particles (CC, DT, WP). But there are several very frequent tags (NN, DT+NN, DT+JJ, VBP, and even PUNC) whose accuracy improves quite a lot. Then there are plural nouns (NNS, DT+NNS) where the char-based model really shines, which makes sense linguistically as plurality in Arabic is usually expressed by certain suffixes (“-wn/yn” for masc. plural, “-At” for fem. plural). The char-based model is thus especially good with frequent tags and infrequent words, which is understandable given that infrequent words typically belong to frequent open categories like nouns and verbs.

4.2 Effect of encoder depth

Modern NMT systems use very deep architectures with up to 8 or 16 layers (Wu et al., 2016; Zhou et al., 2016). We would like to understand what kind of information different layers capture. Given a trained NMT model with multiple layers, we extract feature representations from the different layers in the encoder. Let $ENC_l^i(s)$ denote the encoded representation of word w_i after the l -th layer. We can vary l and train different classifiers to predict POS or morphological tags. Here we focus on the case of a 2-layer encoder-decoder model for simplicity ($l \in \{1, 2\}$).




Figure 6: POS tagging accuracy using representations from layers 0 (word vectors), 1, and 2, taken from encoders of different language pairs.

Figure 6 shows POS tagging results using representations from different encoding layers across five language pairs. The general trend is that passing word vectors through the NMT encoder improves POS tagging, which can be explained by the contextual information contained in the representations after one layer. However, it turns out that representations from the 1st layer are better than those from the 2nd layer, at least for the purpose of capturing word structure. Figure 7 demonstrates that the same pattern holds for both word-based and char-based representations, on Arabic POS and morphological tagging. In all cases, layer 1 representations are better than layer 2 representations.⁴ In contrast, BLEU scores ac-

⁴We found this result to be also true in French, German, and Czech experiments; see appendix A.3.




Figure 7: POS and morphological tagging accuracy across layers. Layer 0: word vectors or char-based representations before the encoder; layers 1 and 2: representations after the 1st and 2nd layers.

usually increase when training 2-layer vs. 1-layer models (+1.11/+0.56 BLEU for Arabic-Hebrew word/char-based models). Thus translation quality improves when adding layers but morphology quality degrades. Intuitively, it seems that lower layers of the network learn to represent word structure while higher layers are more focused on word meaning. A similar pattern was recently observed in a joint language-vision deep recurrent network (Gelderloos and Chrupała, 2016).

4.3 Effect of target language

While translating from morphologically-rich languages is challenging, translating into such languages is even harder. For instance, our basic system obtains BLEU scores of 24.69/23.2 on Arabic/Czech to English, but only 13.37/13.9 on English to Arabic/Czech. How does the target language affect the learned source language representations? Does translating into a morphologically-rich language require more knowledge about source language morphology? In order to investigate these questions, we fix the source language and train NMT models using different target languages. For example, given an Arabic source side, we train Arabic-to-English/Hebrew/German systems. These target languages represent a morphologically-poor language (English), a morphologically-rich language with similar morphology to the source language (Hebrew), and a morphologically-rich language with different morphology (German). To make a fair comparison, we train the models on the intersection of the training data based on the source language. In this way the experimental setup is

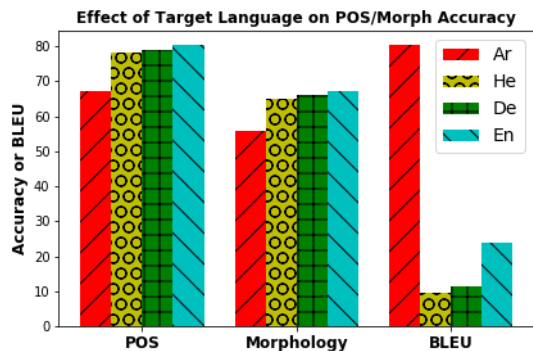


Figure 8: Effect of target language on representation quality of the Arabic source.

completely identical: the models are trained on the same Arabic sentences with different translations.

Figure 8 shows POS and morphological tagging accuracy of word-based representations from the NMT encoders, as well as corresponding BLEU scores. As expected, translating into English is easier than translating into the morphologically-rich Hebrew and German, resulting in higher BLEU scores. Despite their similar morphological systems, translating Arabic to Hebrew is worse than Arabic to German, which can be attributed to the richer Hebrew morphology compared to German. POS and morphology accuracies share an intriguing pattern: the representations that are learned when translating into English are better for predicting POS or morphology than those learned when translating into German, which are in turn better than those learned when translating into Hebrew. This is remarkable given that English is a morphologically-poor language that does not display many of the morphological properties that are found in the Arabic source. In contrast, German and Hebrew have richer morphologies, so one could expect that translating into them would make the model learn more about morphology.

A possible explanation for this phenomenon is that the Arabic-English model is simply better than the Arabic-Hebrew and Arabic-German models, as hinted by the BLEU scores in Table 2. The inherent difficulty in translating Arabic to Hebrew/German may affect the ability to learn good representations of word structure. To probe this more, we trained an Arabic-Arabic autoencoder on the same training data. We found that it learns to recreate the test sentences extremely well, with very high BLEU scores (Figure 8). However, its

word representations are actually inferior for the purpose of POS/morphological tagging. This implies that higher BLEU does not necessarily entail better morphological representations. In other words, a better translation model learns more informative representations, but only when it is actually learning to translate rather than merely memorizing the data as in the autoencoder case. We found this to be consistently true also for character-based experiments, and in other language pairs.

5 Decoder Analysis

So far we only looked at the encoder. However, the decoder `DEC` is a crucial part in an MT system with access to both source and target sentences. In order to examine what the decoder learns about morphology, we first train an NMT system on the parallel corpus. Then, we use the trained model to encode a source sentence and extract features for words in the target sentence. These features are used to train a classifier on POS or morphological tagging on the target side.⁵ Note that in this case the decoder is given the correct target words one-by-one, similar to the usual NMT training regime.

Table 3 (1st row) shows the results of using representations extracted with `ENC` and `DEC` from the Arabic-English and English-Arabic models, respectively. There is clearly a huge drop in representation quality with the decoder.⁶ At first, this drop seems correlated with lower BLEU scores in English to Arabic vs. Arabic to English. However, we observed similar low POS tagging accuracy using decoder representations from high-quality NMT models. For instance, the French-to-English system obtains 37.8 BLEU, but its decoder representations give a mere 54.26% accuracy on English POS tagging.

As an alternative explanation for the poor quality of the decoder representations, consider the fundamental tasks of the two NMT modules: encoder and decoder. The encoder’s task is to create a generic, close to language-independent representation of the source sentence, as shown by recent evidence from multilingual NMT (Johnson et al., 2016). The decoder’s task is to use this representation to generate the target sentence in a specific

⁵In this section we only experiment with predicted tags as there are no parallel data with gold POS/morphological tags that we are aware of.

⁶Note that the decoder results are above a majority baseline of 20%, so the decoder is still learning something about the target language.

Attn	POS Accuracy		BLEU	
	ENC	DEC	Ar-En	En-Ar
✓	89.62	43.93	24.69	13.37
✗	74.10	50.38	11.88	5.04

Table 3: POS tagging accuracy using encoder and decoder representations with/without attention.

language. Presumably, it is sufficient for the decoder to learn a strong language model in order to produce morphologically-correct output, without learning much about morphology, while the encoder needs to learn quite a lot about source language morphology in order to create a good generic representation. In the following section we show that the attention mechanism also plays an important role in the division of labor between encoder and decoder.

5.1 Effect of attention

Consider the role of the attention mechanism in learning useful representations: during decoding, the attention weights are combined with the decoder’s hidden states to generate the current translation. These two sources of information need to jointly point to the most relevant source word(s) and predict the next most likely word. Thus, the decoder puts significant emphasis on mapping back to the source sentence, which may come at the expense of obtaining a meaningful representation of the current word. We hypothesize that the attention mechanism hurts the quality of the target word representations learned by the decoder.

To test this hypothesis, we train NMT models with and without attention and compare the quality of their learned representations. As Table 3 shows (compare 1st and 2nd rows), removing the attention mechanism decreases the quality of the encoder representations, but improves the quality of the decoder representations. Without the attention mechanism, the decoder is forced to learn more informative representations of the target language.

5.2 Effect of word representation

We also conducted experiments to verify our findings regarding word-based versus character-based representations on the decoder side. By character representation we mean a character CNN on the input words. The decoder predictions are still done at the word-level, which enables us to use its hidden states as word representations.

Table 4 shows POS accuracy of word-based vs. char-based representations in the encoder and decoder. While char-based representations improve the encoder, they do not help the decoder. BLEU scores behave similarly: the char-based model leads to better translations in Arabic-to-English, but not in English-to-Arabic. A possible explanation for this phenomenon is that the decoder’s predictions are still done at word level even with the char-based model (which encodes the target input but not the output). In practice, this can lead to generating unknown words. Indeed, in Arabic-to-English the char-based model reduces the number of generated unknown words in the MT test set by 25%, while in English-to-Arabic the number of unknown words remains roughly the same between word-based and char-based models.

6 Related Work

Analysis of neural models The opacity of neural networks has motivated researchers to analyze such models in different ways. One line of work visualizes hidden unit activations in recurrent neural networks that are trained for a given task (Elman, 1991; Karpathy et al., 2015; Kádár et al., 2016; Qian et al., 2016a). While such visualizations illuminate the inner workings of the network, they are often qualitative in nature and somewhat anecdotal. A different approach tries to provide a quantitative analysis by correlating parts of the neural network with linguistic properties, for example by training a classifier to predict features of interest. Different units have been used, from word embeddings (Köhn, 2015; Qian et al., 2016b), through LSTM gates or states (Qian et al., 2016a), to sentence embeddings (Adi et al., 2016). Our work is most similar to Shi et al. (2016), who use hidden vectors from a neural MT encoder to predict syntactic properties on the English source side. In contrast, we focus on representations in morphologically-rich languages and evaluate both source and target sides across several criteria. Vylomova et al. (2016) also analyze different representations for morphologically-rich languages in MT, but do not directly measure the quality of the learned representations.

Word representations in MT Machine translation systems that deal with morphologically-rich languages resort to various techniques for representing morphological knowledge, such as word segmentation (Nieflen and Ney, 2000; Koehn and

	POS Accuracy		BLEU	
	ENC	DEC	Ar-En	En-Ar
Word	89.62	43.93	24.69	13.37
Char	95.35	44.54	28.42	13.00

Table 4: POS tagging accuracy using word-based and char-based encoder/decoder representations.

Knight, 2003; Badr et al., 2008) and factored translation and reordering models (Koehn and Hoang, 2007; Durrani et al., 2014). Characters and other sub-word units have become increasingly popular in neural MT, although they had also been used in phrase-based MT for handling morphologically-rich (Luong et al., 2010) or closely related language pairs (Durrani et al., 2010; Nakov and Tiedemann, 2012). In neural MT, such units are obtained in a pre-processing step—e.g. by byte-pair encoding (Sennrich et al., 2016) or the word-piece model (Wu et al., 2016)—or learned during training using a character-based convolutional/recurrent sub-network (Costa-jussà and Fonollosa, 2016; Luong and Manning, 2016; Vylomova et al., 2016). The latter approach has the advantage of keeping the original word boundaries without requiring pre- and post-processing. Here we focus on a character CNN which has been used in language modeling and machine translation (Kim et al., 2015; Belinkov and Glass, 2016; Costa-jussà and Fonollosa, 2016; Jozefowicz et al., 2016; Sajjad et al., 2017). We evaluate the quality of different representations learned by an MT system augmented with a character CNN in terms of POS and morphological tagging, and contrast them with a purely word-based system.

7 Conclusion

Neural networks have become ubiquitous in machine translation due to their elegant architecture and good performance. The representations they use for linguistic units are crucial for obtaining high-quality translation. In this work, we investigated how neural MT models learn word structure. We evaluated their representation quality on POS and morphological tagging in a number of languages. Our results lead to the following conclusions:

- Character-based representations are better than word-based ones for learning morphology, especially in rare and unseen words.

- Lower layers of the neural network are better at capturing morphology, while deeper networks improve translation performance. We hypothesize that lower layers are more focused on word structure, while higher ones are focused on word meaning.
- Translating into morphologically-poorer languages leads to better source-side representations. This is partly, but not completely, correlated with BLEU scores.
- The attentional decoder learns impoverished representations that do not carry much information about morphology.

These insights can guide further development of neural MT systems. For instance, jointly learning translation and morphology can possibly lead to better representations and improved translation. Our analysis indicates that this kind of approach should take into account factors such as the encoding layer and the type of word representation.

Another area for future work is to extend the analysis to other word representations (e.g. byte-pair encoding), deeper networks, and more semantically-oriented tasks such as semantic role-labeling or semantic parsing.

Acknowledgments

We would like to thank Helmut Schmid for providing the Tiger corpus, members of the MIT Spoken Language Systems group for helpful comments, and the three anonymous reviewers for their useful suggestions. This research was carried out in collaboration between the HBKU Qatar Computing Research Institute (QCRI) and the MIT Computer Science and Artificial Intelligence Laboratory (CSAIL).

References

Yossi Adi, Einat Kermany, Yonatan Belinkov, Ofer Lavi, and Yoav Goldberg. 2016. Fine-grained Analysis of Sentence Embeddings Using Auxiliary Prediction Tasks. *arXiv preprint arXiv:1608.04207*.

Ibrahim Badr, Rabih Zbib, and James Glass. 2008. [Segmentation for English-to-Arabic Statistical Machine Translation](#). In *Proceedings of the 46th Annual Meeting of the Association for Computational Linguistics on Human Language Technologies: Short Papers*. Columbus, Ohio, HLT-Short '08, pages 153–156. <http://dl.acm.org/citation.cfm?id=1557690.1557732>.

Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural Machine Translation by Jointly Learning to Align and Translate. *arXiv preprint arXiv:1409.0473*.

Yonatan Belinkov and James Glass. 2016. Large-Scale Machine Translation between Arabic and Hebrew: Available Corpora and Initial Results. In *Proceedings of the Workshop on Semitic Machine Translation*. Association for Computational Linguistics, Austin, Texas, pages 7–12.

Luisa Bentivogli, Arianna Bisazza, Mauro Cettolo, and Marcello Federico. 2016. [Neural versus Phrase-Based Machine Translation Quality: a Case Study](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, Austin, Texas, pages 257–267. <https://aclweb.org/anthology/D16-1025>.

Mauro Cettolo. 2016. An Arabic-Hebrew parallel corpus of TED talks. In *Proceedings of the Workshop on Semitic Machine Translation*. Association for Computational Linguistics, Austin, Texas, pages 1–6.

Mauro Cettolo, Christian Girardi, and Marcello Federico. 2012. WIT³: Web Inventory of Transcribed and Translated Talks. In *Proceedings of the 16th Conference of the European Association for Machine Translation (EAMT)*. Trento, Italy, pages 261–268.

Marta R. Costa-jussà and José A. R. Fonollosa. 2016. [Character-based Neural Machine Translation](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*. Association for Computational Linguistics, Berlin, Germany, pages 357–361. <http://anthology.aclweb.org/P16-2058>.

Nadir Durrani, Philipp Koehn, Helmut Schmid, and Alexander Fraser. 2014. [Investigating the Usefulness of Generalized Word Representations in SMT](#). In *Proceedings of COLING 2014, the 25th International Conference on Computational Linguistics: Technical Papers*. Dublin City University and Association for Computational Linguistics, Dublin, Ireland, pages 421–432. <http://www.aclweb.org/anthology/C14-1041>.

Nadir Durrani, Hassan Sajjad, Alexander Fraser, and Helmut Schmid. 2010. [Hindi-to-Urdu Machine Translation through Transliteration](#). In *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics, Uppsala, Sweden, pages 465–474. <http://www.aclweb.org/anthology/P10-1048>.

Jeffrey L Elman. 1991. Distributed representations, simple recurrent networks, and grammatical structure. *Machine learning* 7(2-3):195–225.

- Lieke Gelderloos and Grzegorz Chrupała. 2016. From phonemes to images: levels of representation in a recurrent neural model of visually-grounded language learning. In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*. The COLING 2016 Organizing Committee, Osaka, Japan, pages 1309–1319. <http://aclweb.org/anthology/C16-1124>.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural Computation* 9(8):1735–1780.
- Melvin Johnson, Mike Schuster, Quoc V Le, Maxim Krikun, Yonghui Wu, Zhifeng Chen, Nikhil Thorat, Fernanda Viégas, Martin Wattenberg, Greg Corrado, et al. 2016. Google’s Multilingual Neural Machine Translation System: Enabling Zero-Shot Translation. *arXiv preprint arXiv:1611.04558*.
- Rafal Jozefowicz, Oriol Vinyals, Mike Schuster, Noam Shazeer, and Yonghui Wu. 2016. Exploring the Limits of Language Modeling. *arXiv preprint arXiv:1602.02410*.
- Ákos Kádár, Grzegorz Chrupała, and Afra Alishahi. 2016. Representation of linguistic form and function in recurrent neural networks. *arXiv preprint arXiv:1602.08952*.
- Andrej Karpathy, Justin Johnson, and Fei-Fei Li. 2015. Visualizing and Understanding Recurrent Networks. *arXiv preprint arXiv:1506.02078*.
- Yoon Kim. 2016. Seq2seq-attn. <https://github.com/harvardnlp/seq2seq-attn>.
- Yoon Kim, Yacine Jernite, David Sontag, and Alexander M Rush. 2015. Character-aware Neural Language Models. *arXiv preprint arXiv:1508.06615*.
- Diederik Kingma and Jimmy Ba. 2014. Adam: A Method for Stochastic Optimization. *arXiv preprint arXiv:1412.6980*.
- Philipp Koehn and Hieu Hoang. 2007. Factored Translation Models. In *Proceedings of the 2007 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning (EMNLP-CoNLL)*. Association for Computational Linguistics, Prague, Czech Republic, pages 868–876. <http://www.aclweb.org/anthology/D07-1091>.
- Philipp Koehn and Kevin Knight. 2003. Empirical Methods for Compound Splitting. In *10th Conference of the European Chapter of the Association for Computational Linguistics*. pages 187–194. <http://www.aclweb.org/anthology/E03-1076>.
- Arne Köhn. 2015. What’s in an Embedding? Analyzing Word Embeddings through Multilingual Evaluation. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, Lisbon, Portugal, pages 2067–2073. <http://aclweb.org/anthology/D15-1246>.
- Minh-Thang Luong and Christopher D. Manning. 2015. Stanford Neural Machine Translation Systems for Spoken Language Domains. In *Proceedings of the International Workshop on Spoken Language Translation*. Da Nang, Vietnam.
- Minh-Thang Luong and D. Christopher Manning. 2016. Achieving Open Vocabulary Neural Machine Translation with Hybrid Word-Character Models. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics, pages 1054–1063. <https://doi.org/10.18653/v1/P16-1100>.
- Minh-Thang Luong, Preslav Nakov, and Min-Yen Kan. 2010. A Hybrid Morpheme-Word Representation for Machine Translation of Morphologically Rich Languages. In *Proceedings of the 2010 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, pages 148–157. <http://aclweb.org/anthology/D10-1015>.
- Thomas Mueller, Helmut Schmid, and Hinrich Schütze. 2013. Efficient Higher-Order CRFs for Morphological Tagging. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, Seattle, Washington, USA, pages 322–332. <http://www.aclweb.org/anthology/D13-1032>.
- Preslav Nakov and Jörg Tiedemann. 2012. Combining Word-Level and Character-Level Models for Machine Translation Between Closely-Related Languages. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*. Jeju, Korea, ACL ’12, pages 301–305. <http://aclweb.org/anthology/P12-2059>.
- Sonja Nieflen and Hermann Ney. 2000. Improving SMT quality with morpho-syntactic analysis. In *COLING 2000 Volume 2: The 18th International Conference on Computational Linguistics*. <http://www.aclweb.org/anthology/C00-2162>.
- Arfath Pasha, Mohamed Al-Badrashiny, Mona Diab, Ahmed El Kholy, Ramy Eskander, Nizar Habash, Manoj Pooleery, Owen Rambow, and Ryan Roth. 2014. MADAMIRA: A Fast, Comprehensive Tool for Morphological Analysis and Disambiguation of Arabic. In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC’14)*. Reykjavik, Iceland, pages 1094–1101.
- Peng Qian, Xipeng Qiu, and Xuanjing Huang. 2016a. Analyzing Linguistic Knowledge in Sequential Model of Sentence. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, Austin, Texas, pages 826–835. <https://aclweb.org/anthology/D16-1079>.

- Peng Qian, Xipeng Qiu, and Xuanjing Huang. 2016b. Investigating Language Universal and Specific Properties in Word Embeddings. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics, Berlin, Germany, pages 1478–1488. <http://www.aclweb.org/anthology/P16-1140>.
- Adwait Ratnaparkhi. 1998. *Maximum Entropy Models for Natural Language Ambiguity Resolution*. Ph.D. thesis, University of Pennsylvania, Philadelphia, PA.
- Hassan Sajjad, Fahim Dalvi, Nadir Durrani, Ahmed Abdelali, Yonatan Belinkov, and Stephan Vogel. 2017. Challenging Language-Dependent Segmentation for Arabic: An Application to Machine Translation and Part-of-Speech Tagging. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics, Vancouver, Canada.
- Helmut Schmid. 1994. Part-of-Speech Tagging with Neural Networks. In *Proceedings of the 15th International Conference on Computational Linguistics (Coling 1994)*. Coling 1994 Organizing Committee, Kyoto, Japan, pages 172–176.
- Helmut Schmid. 2000. LoPar: Design and Implementation. Bericht des Sonderforschungsbereiches “Sprachtheoretische Grundlagen für die Computerlinguistik” 149, Institute for Computational Linguistics, University of Stuttgart.
- 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)*. Association for Computational Linguistics, Berlin, Germany, pages 1715–1725. <http://www.aclweb.org/anthology/P16-1162>.
- Xing Shi, Inkit Padhi, and Kevin Knight. 2016. Does String-Based Neural MT Learn Source Syntax? In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, Austin, Texas, pages 1526–1534. <https://aclweb.org/anthology/D16-1159>.
- Ilya Sutskever, Oriol Vinyals, and Quoc VV Le. 2014. Sequence to Sequence Learning with Neural Networks. In *Advances in neural information processing systems*. pages 3104–3112.
- Antonio Toral and Víctor M. Sánchez-Cartagena. 2017. A Multifaceted Evaluation of Neural versus Phrase-Based Machine Translation for 9 Language Directions. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*. Association for Computational Linguistics, Valencia, Spain, pages 1063–1073. <http://aclweb.org/anthology/E17-1100>.
- Ekaterina Vylomova, Trevor Cohn, Xuanli He, and Gholamreza Haffari. 2016. Word Representation Models for Morphologically Rich Languages in Neural Machine Translation. *arXiv preprint arXiv:1606.04217*.
- Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, et al. 2016. Google’s Neural Machine Translation System: Bridging the Gap between Human and Machine Translation. *arXiv preprint arXiv:1609.08144*.
- Jie Zhou, Ying Cao, Xuguang Wang, Peng Li, and Wei Xu. 2016. Deep Recurrent Models with Fast-Forward Connections for Neural Machine Translation. *Transactions of the Association for Computational Linguistics* 4:371–383. <https://transacl.org/ojs/index.php/tacl/article/view/863>.

A Supplementary Material

A.1 Training Details

POS/Morphological classifier The classifier used for all prediction tasks is a feed-forward network with one hidden layer, dropout ($\rho = 0.5$), a ReLU non-linearity, and an output layer mapping to the tag set (followed by a Softmax). The size of the hidden layer is set to be identical to the size of the encoder’s hidden state (typically 500 dimensions). We use Adam (Kingma and Ba, 2014) with default parameters to minimize the cross-entropy objective. Training is run with mini-batches of size 16 and stopped once the loss on the dev set stops improving; we allow a patience of 5 epochs.

Neural MT system We train a 2-layer LSTM encoder-decoder with attention. We use the seq2seq-attn implementation (Kim, 2016) with the following default settings: word vectors and LSTM states have 500 dimensions, SGD with initial learning rate of 1.0 and rate decay of 0.5, and dropout rate of 0.3. The character-based model is a CNN with a highway network over characters (Kim et al., 2015) with 1000 feature maps and a kernel width of 6 characters. This model was found to be useful for translating morphologically-rich languages (Costa-jussà and Fonollosa, 2016). The MT system is trained for 20 epochs, and the model with the best dev loss is used for extracting features for the classifier.

A.2 Data and Taggers

Datasets All of the translation models are trained on the Ted talks corpus included in WIT³ (Cettolo et al., 2012; Cettolo, 2016). Statistics about each language pair are available on the WIT³ website: <https://wit3.fbk.eu>. For experiments using gold tags, we used the Arabic Treebank for Arabic (with the versions and splits described in the MADAMIRA manual (Pasha et al., 2014)) and the Tiger corpus for German.⁷

POS and morphological taggers We used the following tools to annotate the MT corpora: MADAMIRA (Pasha et al., 2014) for Arabic POS and morphological tags, Tree-Tagger (Schmid, 1994) for Czech and French POS tags, LoPar (Schmid, 2000) for German POS and morphological tags, and MXPOST (Ratnaparkhi, 1998) for English POS tags. These tools are recommended

⁷<http://www.ims.uni-stuttgart.de/forschung/ressourcen/korpora/tiger.html>

on the Moses website.⁸ As mentioned before, our goal is not to achieve state-of-the-art results, but rather to study what different components of the NMT architecture learn about word morphology. Please refer to Mueller et al. (2013) for representative POS and morphological tagging accuracies.

A.3 Supplementary Results

We report here results that were omitted from the paper due to the space limit. Table 5 shows encoder results using different layers, languages, and representations (word/char-based). As noted in the paper, all the results consistently show that i) layer 1 performs better than layers 0 and 2; and ii) char-based representations are better than word-based for learning morphology. Table 6 shows that translating into a morphologically-poor language (English) leads to better source representations, and Table 7 provides additional decoder results.

	Layer 0	Layer 1	Layer 2
	Word/Char (POS)		
De	91.1/92.0	93.6/95.2	93.5/94.6
Fr	92.1/92.9	95.1/95.9	94.6/95.6
Cz	76.3/78.3	77.0/79.1	75.7/80.6
	Word/Char (Morphology)		
De	87.6/88.8	89.5/91.2	88.7/90.5

Table 5: POS and morphology accuracy on predicted tags using word- and char-based representations from different layers of *-to-En systems.

Source	Target		
	English	Arabic	Self
German	93.5	92.7	89.3
Czech	75.7	75.2	71.8

Table 6: Impact of changing the target language on POS tagging accuracy. Self = German/Czech in rows 1/2 respectively.

	En-De	En-Cz	De-En	Fr-En
POS	53.6	36.3	53.3	54.1
BLEU	23.4	13.9	29.6	37.8

Table 7: POS accuracy and BLEU using decoder representations from different language pairs.

⁸[http://www.statmt.org/ Moses.ExternalTools](http://www.statmt.org/ Moses/?n=Moses.ExternalTools)