



Assessing the Tolerance of Neural Machine Translation Systems Against Speech Recognition Errors

Nicholas Ruiz^{1,2}, Mattia Antonino Di Gangi¹, Nicola Bertoldi¹, Marcello Federico¹

¹Fondazione Bruno Kessler, Italy

²Interactions, LLC, USA

nruiz@interactions.com, {digangi,bertoldi,federico}@fbk.eu

Abstract

Machine translation systems are conventionally trained on textual resources that do not model phenomena that occur in spoken language. While the evaluation of neural machine translation systems on textual inputs is actively researched in the literature, little has been discovered about the complexities of translating spoken language data with neural models. We introduce and motivate interesting problems one faces when considering the translation of automatic speech recognition (ASR) outputs on neural machine translation (NMT) systems. We test the robustness of sentence encoding approaches for NMT encoder-decoder modeling, focusing on word-based over byte-pair encoding. We compare the translation of utterances containing ASR errors in state-of-the-art NMT encoder-decoder systems against a strong phrase-based machine translation baseline in order to better understand which phenomena present in ASR outputs are better represented under the NMT framework than approaches that represent translation as a linear model.

Index Terms: speech translation, machine translation, evaluation, neural machine translation

1. Introduction

A substantial amount of progress has been made in Neural Machine Translation (NMT) for text documents. Research has shown that the encoder-decoder model with an attention mechanism generates high quality translations that exploit long range dependencies in an input sentence [1]. While NMT has proven to yield significant improvements for text translation over log-linear approaches to MT such as phrase-based machine translation (PBMT), it has yet to be shown the extent to which gains purported in the literature generalize to the scenario of spoken language translation (SLT), where the input sequence may be corrupted by noise in the audio signal and uncertainties during automatic speech recognition (ASR) decoding. Are NMT models implicitly better at modeling and mitigating ASR errors than the former state-of-the-art approaches to machine translation? As a preliminary work, we analyze the impact of ASR errors on neural machine translation quality by studying the properties of the translations provided by an encoder-decoder NMT system with an attention mechanism, against a strong baseline PBMT system that rivals the translation quality of Google TranslateTM on TED talks.

We address the following questions regarding NMT:

1. How do NMT systems react when ASR transcripts are provided as input?
2. Do ASR error types in word error alignments impact SLT quality the same for NMT as PBMT? Or is NMT implicitly more tolerant against ASR errors?

3. Which types of sentences does NMT handle better than PBMT, and vice-versa?

To address these questions, we explore the impact of feeding ASR hypotheses, which may contain noise, disfluencies, and different representations on the surface text, to a NMT system that has been trained on TED talk transcripts that do not reflect the noisy conditions of ASR. Our experimental framework is similar to that of [2, 3], with the addition of a ranking experiment to evaluate the quality of NMT against our PBMT baseline. These experiments are intended as an initial analysis with the purpose to suggesting directions to focus on in the future.

2. Neural versus Statistical MT

Before beginning our analysis, we summarize some of the biggest differences between NMT and other forms of statistical machine translation, such as PBMT.

[4] compare neural machine translation against three top-performing statistical machine translation systems in the TED talk machine translation track from IWSLT 2015.¹ The evaluation set consists of 600 sentences and 10,000 words, post-edited by five professional translators. In addition to reporting a 26% relative improvement in multi-reference TER (mTER), [5]’s encoder-decoder attention-based NMT system trained on full words outperformed state of the art statistical machine translation (SMT) systems on English-German, a language pair known to have issues with morphology and whose syntax differs significantly from English in subordinate clauses. [4]’s analysis yields the following observations:

- *Precision versus Sentence length:* Although NMT outperformed every comparable log-linear MT system, they confirmed [6]’s findings that translation quality deteriorates rapidly as the sentence length approaches 35 words.
- *Morphology:* NMT translations have better case, gender and number agreement than PBMT systems.
- *Lexical choice:* NMT made 17% fewer lexical errors than any PBMT system.
- *Word order:* NMT yielded fewer *shift* errors in TER alignments than any SMT system. NMT yielded significantly higher Kendall Reordering Score (KRS) [7] values than any PBMT system. NMT generated 70% fewer verb order errors than the next-best hybrid phrase and syntax-based system.

Several SMT modeling challenges are exacerbated in NMT. While log-linear SMT translation models can handle large word

¹The International Workshop of Spoken Language Translation.

vocabularies, NMT systems require careful modeling to balance vocabulary coverage and network size, since each token introduced increases the size of its hidden layers. Because of this constraint, [8] observe that only 69% of German nouns are covered with 30,000 words on English-German WMT 2014 system.² Although noun compound splitting works well for German→English, English→German model performance not improve significantly. In particular named entities (e.g. persons, organizations, and locations) are underrepresented.

On the other hand, NMT has the ability to model subword units such as characters [9] or coarser grained segmentations of low frequency words [10] without substantial changes to the system architecture, unlike other SMT approaches.

[11] have additionally demonstrated NMT’s ability to translate between multiple language pairs with a neural translation model trained with a single attention mechanism.

Although NMT models translate with higher precision, models are slow to train even with the most powerful GPUs – often taking weeks for the strongest systems to complete training. On the other hand, large order PBMT systems trained in the ModernMT framework³ may be trained within a few hours and can be adapted in near real-time with translation memories containing post-editions by professional translators.

3. Research Methodology

Similar to our experimental framework in [2, 3], we collect English ASR hypotheses from the eight submissions on the tst_{2012} test set in the IWSLT 2013 TED talk ASR track [12]. Coupled with reference translations from the MT track, we construct a dataset consisting of the eight English ASR hypotheses for 1,124 utterances, a single unpunctuated reference transcript from the ASR track, and the reference translations from the English-French MT track. The English ASR hypotheses and reference transcript are normalized and punctuated according to the same approach as described in [3]. We use both BLEU [13] and Translation Edit Rate (TER) [14] as global evaluation metrics. TER and Δ TER over gold standard ASR outputs are used to observe sentence-level trends. We compute automatic translation scores, sentence-level system ranking, and take a closer look at the types of errors observed in the data. Below, we briefly describe the MT systems used in this experiment.

3.1. Neural MT system

Our NMT system is based on FBK’s primary MT submission to the IWSLT 2016 evaluation for English-French TED talk translation [15]. The system is based on the sequence-to-sequence encoder-decoder architecture proposed in [1] and further developed by [5, 10]. The system is trained on full-word text units to allow a direct comparison with our PBMT counterpart. We refer to this system as NEURAL for the remainder of our experiments.

3.2. Phrase-based MT system

Our phrase-based MT system (which we refer to as MMT) is built upon the ModernMT framework: an extension of the phrase-based MT framework introduced in [16] that enables context-aware translation for fast and incremental training. Context-aware translation is achieved by the partitioning

²WMT 2014 training data consists primarily of news texts, European parliament proceedings, and web crawled data. <http://www.statmt.org/wmt14/translation-task.html>

³<http://www.modernmt.eu>

Table 1: A comparison of Neural MT versus Phrase-based MT on the SLT evaluation of TED talks (tst_{2012}) from the IWSLT 2013 evaluation campaign [12]. Evaluation results are compared to a gold standard that assumes that no ASR errors occur.

ASR system	WER ↓	MMT		NEURAL	
		BLEU ↑	TER ↓	BLEU ↑	TER ↓
gold	0.0	43.4	39.5	47.9	35.4
fbk	16.5	35.6	48.4	38.5	45.6
kit	10.1	38.1	45.0	41.8	41.7
mitll	11.4	37.7	45.8	41.4	42.4
naist	10.6	38.1	45.0	41.8	41.5
nict	9.2	38.7	44.7	42.5	41.1
prke	16.6	34.9	48.7	38.1	45.8
rwth	11.7	37.2	46.1	41.4	42.3
uedin	12.3	37.3	46.1	40.8	42.9

Table 2: Average utterance-level translation TER and Δ TER scores for the MMT and Neural MT systems. The average Neural MT TER scores are 3% better than the PBMT counterpart.

SysID	MMT		NEURAL		DIFFERENCE	
	TER ↓	Δ TER ↓	TER	Δ TER	TER	Δ TER
gold	39.6	0.0	35.6	0.0	-4.0	0.0
fbk	49.3	9.7	46.6	11.0	-2.7	1.3
kit	45.9	6.3	42.7	7.1	-3.2	0.8
mitll	46.8	7.2	43.7	8.1	-3.1	0.9
naist	45.6	6.1	42.1	6.5	-3.5	0.5
nict	45.1	5.5	41.9	6.3	-3.1	0.9
prke	49.4	9.8	46.5	10.9	-2.9	1.1
rwth	47.0	7.4	43.2	7.6	-3.8	0.2
uedin	46.7	7.1	43.8	8.2	-2.9	1.1

of the training data into homogeneous domains by a *context analyzer* (CA), which permits the rapid construction and interpolation of domain-specific translation, language, and reordering models based on the context received during a decoding run. It also permits the underlying models to be modified online. The decoder also exploits other standard features (phrase, word, distortion, and unknown word penalties) and performs cube-pruning search. A detailed description of the ModernMT project can be found in [17].

4. SLT Evaluation

We first report the translation results on the evaluation task in Table 1. NMT outperforms our best PBMT system by 4.5 BLEU in the absence of ASR errors (*gold*) and by approximately 3 BLEU across all ASR hypothesis inputs. Overall, the introduction of ASR errors results in decreases in BLEU by 5.5(\pm 0.8) and 5.4(\pm 0.8) and TER increases of 6.0(\pm 0.9) and 6.2(\pm 0.9) for MMT and NEURAL, respectively.

Table 2 provides the average sentence-level TER and Δ TER scores, which report the degradation of SLT quality by the presence of ASR errors. Although the average TER scores from the MMT outputs are higher, the Δ TER scores are lower than their NEURAL counterparts, implying that the MMT SLT outputs are closer to their gold standard MT outputs. This may suggest that NMT is more sensitive to local changes to an input caused by minor ASR errors.

4.1. MT system ranking

Are there ASR error conditions in which PBMT remains a better solution than NMT, and if so, what are the properties of these utterances that makes them difficult for NMT? We take a closer look at the sentence-level translation scores by ranking the performance of each MT system on the utterances where ASR errors exist, in order to understand how each MT system handles

Table 3: Ranked evaluation of the SLT utterances containing ASR errors in *tst*₂₀₁₂. (Left) Counts of the winner decisions and the percentage of all of the decisions that were influenced by ASR errors. (Right) Mean TER scores across each sentence in the ranked set. The remainder of winner decisions are made on error-free ASR transcripts.

Lab	SLT Rank by sentence			TER (avg) ↓	
	Winner	Count	Percentage	MMT	NEURAL
fbk	MMT	257	32.4	51.1	64.0
	NEURAL	373	47.1	58.9	44.5
	Tie	162	20.5	54.1	54.1
kit	MMT	213	30.6	46.3	59.1
	NEURAL	347	49.9	55.0	41.1
	Tie	135	19.4	52.9	52.9
mitll	MMT	194	27.6	48.4	61.4
	NEURAL	351	49.9	55.5	41.5
	Tie	159	22.6	52.2	52.2
naist	MMT	189	28.3	43.9	56.5
	NEURAL	342	51.2	54.8	41.2
	Tie	137	20.5	52.5	52.5
nict	MMT	184	31.8	46.3	58.4
	NEURAL	286	49.4	54.0	40.8
	Tie	109	18.8	56.0	56.0
prke	MMT	256	31.6	48.0	60.3
	NEURAL	378	46.7	57.7	44.1
	Tie	175	21.6	55.8	55.8
rwth	MMT	221	29.9	47.0	59.2
	NEURAL	383	51.8	55.5	41.3
	Tie	135	18.3	55.0	55.0
uedin	MMT	219	30.6	47.9	59.2
	NEURAL	348	48.7	56.5	42.8
	Tie	148	20.7	52.2	52.2

noisy input. For each utterance, we rank the systems based on their the sentence-level TER scores computed on their translation outputs over each ASR hypothesis. We also mark ties, in which both systems yield the same TER score. Results containing the counts and percentage of wins by MT system are provided in Table 3.

The NEURAL and MMT scores are tied on over 20% of the utterances. For the better performing ASR systems (e.g. NICT, KIT), we observe a slightly higher proportion of utterances with better NMT translations and a reduced number of ties. On the right-hand side of Table 3 we report the average TER scores within each ranking partition of the data. For example, for the utterances that are translated better by MMT, we observe that the average TER scores for NEURAL have an absolute average improvement of 10% in TER over MMT. The converse is also true, suggesting that there is a subset of utterances that MMT translates better than NEURAL.

We look into the translation errors caused by ASR errors by plotting the changes in MT system ranking as we shift from the perfect ASR scenario to the actual ASR results from the evaluation (Table 4). Across all ASR outputs, 70.2% of the MT evaluation ranking decisions remain the same when ASR errors create noisy input. The NEURAL model retains a higher rank 7.5% more often than MMT as ASR errors are introduced. Ranking ties remain 55.5% of the time. Of the remaining, the NEURAL model outperforms MMT 5.5% more often in the presence of ASR errors (25.0% versus 19.5%). These results confirm that at the corpus level NMT produces higher scoring translations in the presence of ASR errors.

4.2. Translation examples

Although NMT may outperform phrase-based SMT, our experiment shows that MMT still outperforms NEURAL 30.1% of the time. In order to understand this behavior, we provide three ex-

Table 4: Changes in sentence-level TER rankings as ASR errors are introduced.

Gold Winner	ASR Winner	Absolute %	Relative %
MMT	MMT	19.8	69.0
	NEURAL	4.6	16.1
	Tie	4.3	14.9
NEURAL	MMT	6.4	12.4
	NEURAL	39.4	76.5
	Tie	5.7	11.1
Tie	MMT	3.9	19.5
	NEURAL	5.0	25.0
	Tie	11.0	55.5

amples of key differences between in how NEURAL and MMT mitigate FBK’s ASR errors (Fig. 1). In utterance U4, NEURAL is missing the translation of two content words from its vocabulary. In the absence of errors NEURAL_{gold} passes the source word “embody” through to its output without translating it. During ASR, “embody” is misrecognized as “body”, which is also passed through without a translation. We find it strange that “body” was not translated as “corps”, given that other utterances containing “body” receive that translation. After investigating further, we came across other cases of gold transcripts where “body” was not translated at all. Utterances U212, U214, and U242 have the phrase “body architect”, but only U212 has a translation for the word “body”:

I call myself a body architect.
je m’appelle un corps architecte.

As a body architect, I’m fascinated with the human body
en tant qu’architecte, je me suis retrouvé avec le corps humain

As a body architect, I’ve created
en tant qu’architecte, j’ai créé

It is likely that NMT may not be able to translate contextual patterns it hasn’t observed before. MMT on the other hand provides valid translation for both words; although the meaning of the sentence is lost due to the translation of ASR errors. A PBMT system will translate phrases consistently, as long as there is not another overlapping phrase pair in the translation model that leads to a path in the search graph with a higher score.

Utterance U85 in the TED talk test set shows longer range effects of ASR errors on translation in NMT. FBK’s ASR recognized the utterance as “But when I **step** back, I felt myself at the cold, hard center of a perfect storm.” In the translation of ASR_{gold}, both MT systems translate the expression “stepped back” in the sense of “returned”. MMT_{gold} reorders “centre” incorrectly. ASR_{hyp} has a single error where the past tense suffix “-ed” on “step” was lost. NEURAL_{ASR} provides an adequate translation as “je recule”, but in the process, the attention mechanism seems to have taken the incorrect source word and translation as context that corrupts the remainder of the translation. While MMT_{ASR} makes a translation error at the beginning of the sense, the remainder of the translated sentence remains the same as its gold translation. This suggests that ASR errors may have longer range effects on NMT systems in languages that are even observable in sentences that lack long distance dependencies.

Utterance U296 demonstrates an example where misrecognitions of short function words can cause the duplication of content words in NMT. While MMT handles the misrecognition “and” ⇒ “an” by backing off by translating it independently from other phrases in the sentence, NEURAL, attaches “photo” both to the article “an” and additionally outputs “photo” at its usual position. As innocuous closed-class word errors that occur often in ASR, this could yield a significant problem in NMT.

		TER	Δ TER
ASR _{gold}	I embody the central paradox.		(U4)
ASR _{hyp}	I body the central paradox.		
Translation	j' incarne le paradoxe central .		
MMT _{ASR}	je corps au paradoxe central .	50.0	50.0
NEURAL _{ASR}	je body le paradoxe central .	33.33	16.66
MMT _{gold}	j' incarne le paradoxe central .	0.0	
NEURAL _{gold}	j' embody le paradoxe central .	16.67	
ASR _{gold}	But when I stepped back, I felt myself at the cold, hard center of a perfect storm.		(U85)
ASR _{hyp}	But when I step back, I felt myself at the cold, hard center of a perfect storm.		
Translation	mais quand j' ai pris du recul , je me suis sentie au centre froid , et dur d' une tempête parfaite .		
NEURAL _{ASR}	mais quand je recule , je me sentais dans le froid et le centre d' une tempête parfaite .	47.83	13.05
NEURAL _{gold}	mais quand je suis revenu , je me sentais au centre froid et dur d' une tempête parfaite .	34.78	
ASR _{gold}	And he emailed me this picture.		(U296)
ASR _{hyp}	An emailed me this picture.		
Translation	il m' a envoyé cette photo .		
MMT _{ASR}	un m' a envoyé cette photo .	14.29	0.0
NEURAL _{ASR}	une photo m' a envoyé cette photo .	28.57	14.28
MMT _{gold}	et il m' a envoyé cette photo .	14.29	
NEURAL _{gold}	et il m' a envoyé cette photo .	14.29	

Figure 1: Three examples of changes in NMT errors (NEURAL) caused by ASR errors: (1) the effects of unobserved context; (2) long distance effects of local ASR errors; (3) duplication of content words caused by substitution errors on short function words. Alternative translations are provided by MMT.

5. Mixed-effects analysis and error distribution

In order to quantify the effects of ASR errors on each system, we build linear mixed-effects models [18] in a similar manner to our mixed-effects analysis in [2, 3]. We construct two sets of mixed-effects models, using the word error rate scores of the 8 ASR hypotheses as independent variables and the resulting increase in translation errors Δ TER as the response variable. The models contain random effect intercepts that account for the variance associated with the ASR system (SysID), the intrinsic difficulty of translating a given utterance (UttID), and a random effects slope accounting for the variability of word error rate scores (WER) across systems. Instead of treating each different MT system as a random effect in a joint mixed-effect model, we construct a mixed-effects model for each MT system with the purpose of comparing the degree to which each ASR error type explains the increase in translation difficulty. The models are built using R [19] and the *lme4* library [20]. The fixed-effects coefficients and the variance of the random effects for each model are shown in Table 5.

Our first models (*WER-only*) focus on the effects of the global WER score on translation quality (Δ TER). Our fitted models claim that each point of WER yields approximately the same change in Δ TER for NEURAL (0.61 ± 0.020) and MMT (0.56 ± 0.019).

Our second models (*WER_{basic}*) break WER into its **S**ubstitution, **D**eletion, and **I**nsertion error alignments, each being normalized by the length of the reference transcript. According to the fixed effects of the model, insertion errors have a greater impact on translation quality in NMT than deletions. More importantly, substitution errors have a significantly stronger impact in NMT on translation quality, which reflects the behavior we observe in the translation examples from Fig. 1. MMT appears to be affected by insertion and deletion error

Table 5: *Mixed-effects summary, comparing Neural MT (NEURAL) compared to Phrase-based MT (MMT) for SLT. Top: WER score as a single predictor of translation Δ TER. Bottom: Decomposing WER into the basic alignment error operations. Statistical significance at $p < 10^{-4}$ is marked with \bullet .*

WER-only (null model)				
Fixed effects	NEURAL		MMT	
	β	Std. Error	β	Std. Error
(Intercept)	4.35e-3	2.68e-3	-2.08e-5	1.92e-3
WER	6.09e-1	1.98e-2 \bullet	5.58e-1	1.85e-2 \bullet
Random effects	Variance	Std. Dev.	Variance	Std. Dev.
UttID (Intercept)	0.01	0.08	0.00	0.05
WER	0.23	0.48	0.22	0.47
SysID (Intercept)	0	0	0	0
Residual	0.01	0.07	0.00	0.06
WER _{basic} (Levenshtein alignment errors)				
Fixed effects	NEURAL		MMT	
	β	Std. Error	β	Std. Error
(Intercept)	4.87e-3	2.69e-3	-5.76e-5	1.93e-3
Sub	6.80e-1	2.10e-2 \bullet	5.35e-1	1.96e-2 \bullet
Del	4.28e-1	2.41e-2 \bullet	5.94e-1	2.20e-2 \bullet
Ins	5.59e-1	3.01e-2 \bullet	5.98e-1	2.68e-2 \bullet

Table 6: *Average ASR error counts for utterances translated best with NEURAL, MMT, or a tie. Translation TER is compared between the best MT system and the inferior MT system. Computed on utterances with reference (gold) length between 10-20 words.*

	NEURAL	MMT	Tie
Length _{gold}	14.75 \pm 0.27	14.78 \pm 0.37	14.92 \pm 0.44
WER	17.74 \pm 1.69	19.39 \pm 2.54	17.28 \pm 2.77
Sub	1.19 \pm 0.14	1.15 \pm 0.17	0.20 \pm 0.12
Del	0.22 \pm 0.05	0.32 \pm 0.07	0.33 \pm 0.12
Ins	0.22 \pm 0.05	0.27 \pm 0.08	0.35 \pm 0.12
Sub-span	1.00 \pm 0.15	0.92 \pm 0.19	0.75 \pm 0.21
TER	42.00 vs 57.66	47.71 vs 64.02	49.69

types equally.

We compare the average ASR error type frequencies in the FBK ASR utterances where NEURAL or MMT yield a better TER score. We introduce the “phonetic substitution span” error type from [3] to cover multi-word substitution errors (e.g. “anatomy” \Rightarrow “and that to me”). Focusing on utterances between 10 and 20 words, we observe in Table 6 that the cases where NEURAL scores highest consist of utterances with fewer deletion errors (0.22 versus 0.32). Although further investigation is needed to understand the interplay between substitution and deletion ASR errors in NMT, it is interesting to note that mmt seems to be more adept to handle error-prone ASR outputs, given the higher average WER (19.4% vs 17.7%).

6. Conclusion

We have introduced a preliminary analysis of the impact of ASR errors on SLT for models trained by neural machine translation systems. In particular, we identify the following as areas to focus on in new research in evaluating NMT for spoken language translation scenarios: (1) contextual patterns not observed during training – SMT systems usually can back off to shorter sized entries in their translation table; NMT behavior can be erratic. (2) localized and minor ASR errors can cause long distance errors in translation. (3) NMT duplicates content words when minor ASR errors cause the modification of function words. Most of the observable errors above are caused by minor substitution errors caused by noisy ASR. We will expand this analysis further by evaluating NMT architectures that model coverage as well as the representation of inputs with subword units.

7. References

- [1] D. Bahdanau, K. Cho, and Y. Bengio, “Neural Machine Translation by Jointly Learning to Align and Translate,” in *5th International Conference on Learning Representations*. San Diego, USA: ICLR, 2015.
- [2] N. Ruiz and M. Federico, “Assessing the Impact of Speech Recognition Errors on Machine Translation Quality,” in *Association for Machine Translation in the Americas (AMTA)*, Vancouver, Canada, 2014, pp. 261–274.
- [3] —, “Phonetically-Oriented Word Error Alignment for Speech Recognition Error Analysis in Speech Translation,” in *IEEE Workshop on Automatic Speech Recognition and Understanding (ASRU)*. Scottsdale, Arizona: IEEE, December 2015.
- [4] L. Bentivogli, A. Bisazza, M. Cettolo, and M. Federico, “Neural versus phrase-based machine translation quality: a case study,” in *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing, EMNLP 2016, Austin, Texas, USA, November 1-4, 2016*, 2016, pp. 257–267. [Online]. Available: <http://aclweb.org/anthology/D/D16/D16-1025.pdf>
- [5] M.-T. Luong and C. D. Manning, “Stanford neural machine translation systems for spoken language domain,” in *International Workshop on Spoken Language Translation*, Da Nang, Vietnam, 2015.
- [6] K. Cho, B. van Merriënboer, D. Bahdanau, and Y. Bengio, “On the properties of neural machine translation: Encoder-decoder approaches,” in *Proceedings of the Eighth Workshop on Syntax, Semantics and Structure in Statistical Translation*, 2014, pp. 103–111. [Online]. Available: <http://aclweb.org/anthology/W/W14/W14-4012.pdf>
- [7] A. Birch, P. Blunsom, and M. Osborne, “A quantitative analysis of reordering phenomena,” in *StatMT '09: Proceedings of the Fourth Workshop on Statistical Machine Translation*. Morristown, NJ, USA: Association for Computational Linguistics, 2009, pp. 197–205.
- [8] F. Hirschmann, J. Nam, and J. Fürnkranz, “What Makes Word-level Neural Machine Translation Hard: A Case Study on English-German Translation,” in *Proceedings of the 25th International Conference on Computational Linguistics*, Osaka, Japan, December 2016.
- [9] J. Chung, K. Cho, and Y. Bengio, “A character-level decoder without explicit segmentation for neural machine translation,” *arXiv preprint arXiv:1603.06147*, 2016.
- [10] R. Sennrich, B. Haddow, and A. Birch, “Neural machine translation of rare words with subword units,” in *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics, ACL 2016, August 7-12, 2016, Berlin, Germany, Volume 1: Long Papers*, 2016. [Online]. Available: <http://aclweb.org/anthology/P/P16/P16-1162.pdf>
- [11] O. Firat, K. Cho, B. Sankaran, F. T. Y. Vural, and Y. Bengio, “Multi-way, multilingual neural machine translation,” *Computer Speech & Language*, 2016.
- [12] M. Cettolo, J. Niehues, S. Stüker, L. Bentivogli, and M. Federico, “Report on the 10th IWSLT Evaluation Campaign,” in *Proc. of the International Workshop on Spoken Language Translation*, December 2013.
- [13] K. Papineni, S. Roukos, T. Ward, and W.-J. Zhu, “BLEU: a method for automatic evaluation of machine translation,” in *Proceedings of the 40th Annual Meeting of the Association of Computational Linguistics (ACL)*, Philadelphia, PA, 2002, pp. 311–318. [Online]. Available: <http://aclweb.org/anthology-new/P/P02/P02-1040.pdf>
- [14] M. Snover, B. Dorr, R. Schwartz, L. Micciulla, and J. Makhoul, “A study of translation edit rate with targeted human annotation,” in *5th Conference of the Association for Machine Translation in the Americas (AMTA)*, Boston, Massachusetts, August 2006.
- [15] M. A. Farajian, R. Chatterjee, C. Conforti, S. Jalalvand, V. Balaraman, M. A. Di Gangi, D. Ataman, M. Turchi, M. Negri, and M. Federico, “FBK’s Neural Machine Translation Systems for IWSLT 2016,” in *Proceedings of the 9th International Workshop on Spoken Language Translation (IWSLT)*, Seattle, WA, USA, December 2016.
- [16] P. Koehn and J. Schroeder, “Experiments in Domain Adaptation for Statistical Machine Translation,” in *Proceedings of the Second Workshop on Statistical Machine Translation*. Prague, Czech Republic: Association for Computational Linguistics, June 2007, pp. 224–227. [Online]. Available: <http://www.aclweb.org/anthology/W/W07/W07-0233>
- [17] N. Bertoldi, D. Caroselli, D. Madl, M. Cettolo, and M. Federico, “ModernMT – Second Report on Database and MT Infrastructure,” European Union Horizon 2020 research and innovation programme, Tech. Rep. D.32, December 2016.
- [18] S. R. Searle, “Prediction, mixed models, and variance components,” Biometrics Unit, Cornell University, Tech. Rep. BU-468-M, June 1973. [Online]. Available: <http://hdl.handle.net/1813/32559>
- [19] R Core Team, *R: A Language and Environment for Statistical Computing*, R Foundation for Statistical Computing, Vienna, Austria, 2013.
- [20] D. Bates, M. Maechler, B. Bolker, and S. Walker, *lme4: Linear mixed-effects models using Eigen and S4*, 2014, r package version 1.1-6. [Online]. Available: <http://CRAN.R-project.org/package=lme4>