

Understanding Model Robustness to User-generated Noisy Texts

Jakub Náplava and Martin Popel and Milan Straka and Jana Straková

Charles University,

Faculty of Mathematics and Physics,

Institute of Formal and Applied Linguistics

{naplava, popel, straka, strakova}@ufal.mff.cuni.cz

Abstract

Sensitivity of deep-neural models to input noise is known to be a challenging problem. In NLP, model performance often deteriorates with naturally occurring noise, such as spelling errors. To mitigate this issue, models may leverage artificially noised data. However, the amount and type of generated noise has so far been determined arbitrarily. We therefore propose to model the errors statistically from grammatical-error-correction corpora. We present a thorough evaluation of several state-of-the-art NLP systems' robustness in multiple languages, with tasks including morpho-syntactic analysis, named entity recognition, neural machine translation, a subset of the GLUE benchmark and reading comprehension. We also compare two approaches to address the performance drop: a) training the NLP models with noised data generated by our framework; and b) reducing the input noise with external system for natural language correction. The code is released at <https://github.com/ufal/kazitext>.

1 Introduction

Although there has recently been an amazing progress in variety of NLP tasks (Vaswani et al., 2017; Devlin et al., 2019) with some models even reaching performance comparable to humans on certain domains (Ge et al., 2018; Popel et al., 2020), it has been shown that the models are very sensitive to noise in data (Belinkov and Bisk, 2017; Rychalska et al., 2019).

Multiple areas of NLP have been studied to evaluate the effect of noise in data (Belinkov and Bisk, 2017; Heigold et al., 2018; Ribeiro et al., 2018; Glockner et al., 2018) and a framework for text corruption to test NLP models robustness is also available (Rychalska et al., 2019). However, all these systems introduce noise in a custom-defined, arbitrary level and typically for a single language.

We suggest modeling natural noise statistically from corpora and we propose a framework with the following distinctive features:

- The error probabilities are estimated on real-world grammatical-error-correction corpora.
- The intended noisiness can be scaled to a desired level and various aspects (types) of errors can be turned on/off to test the NLP systems robustness to specific error types.

Furthermore, we also present a thorough evaluation of several current state-of-the-art NLP systems' with varying level of data noisiness and a selection of error aspects in multiple languages. The NLP tasks include morpho-syntactic analysis, named entity recognition, neural machine translation, a subset of GLUE benchmark and reading comprehension. We conclude that:

- The amount of noise is far more important than the distribution of error types.
- Sensitivity to noise differs greatly among NLP tasks. While tasks such as lemmatization require correcting the input text, only an approximate understanding is sufficient for others.

We also compare two approaches for increasing models robustness to noise: training with noise and external grammatical-error-correction (GEC) preprocessing. Our findings suggest that training with noise is beneficial for models with large capacity and large training data (neural machine translation), while the preprocessing with grammatical-error-correction is more suitable for limited-data classification tasks, such as morpho-syntactic analysis.

Finally, we also offer an evaluation on authentic noise: We assembled a new dataset with authentic Czech noisy sentences translated into English and we evaluate the noise-mitigating strategies in the neural machine translation task on this dataset.

2 Related Work

Many empirical findings have shown the fact that data with natural noise deteriorate NLP systems performance. [Belinkov and Bisk \(2017\)](#) found that natural noise such as misspellings and typos cause significant drops in BLEU scores of character-level machine translation models. To increase the model’s robustness, they trained the model on a mixture of original and noisy input and found out that it learnt to address certain amount of errors. Similar findings were observed by [Heigold et al. \(2018\)](#) who tested machine translation and morphological tagging under three types of word-level errors.

[Ribeiro et al. \(2018\)](#) defined a set of substitution rules that produce semantically equivalent text variants. They used them to test systems in machine comprehension, visual question answering and sentiment analysis. [Glockner et al. \(2018\)](#) created a new test set for natural language inference and showed that current systems do not generalize well even for a single-word replacements by synonyms and antonyms.

[Rychalska et al. \(2019\)](#) implemented a framework for introducing multiple noise types into text such as removing or swapping articles, rewriting digit numbers into words or introducing errors in spelling. They found out in four NLP tasks that even recent state-of-the-art systems based on contextualized word embeddings are not completely robust against such natural noise. They also retrained the systems on noisy data and observed improvements for certain error types.

Similarly to [Rychalska et al. \(2019\)](#), we also developed a general framework that allows to test a variety of NLP tasks. The difference is that we estimate the probabilities of individual error types from real-world error corpora. This makes the generated sentences more similar to what humans would do. Moreover, since we defined the individual error types with no language-specific rules, we can apply it to multiple languages with an available annotated grammatical-error corpus.

Grammatical-error corpora are typically used as training data for estimating error statistics in GEC systems. In a setting similar to ours, [Choe et al. \(2019\)](#); [Rozovskaya et al. \(2017\)](#) also estimated error statistics and used them to generate additional training data for GEC systems. However, compared to our approach, they defined only a small set of predefined error categories and used it specifically

for training GEC systems whereas we also use it to assess model performance in noisy scenarios.

Authentic Noise Evaluation The growing interest in developing production-ready machine translation models that are robust to natural noise resulted in the First Shared Task on Machine Translation Robustness ([Li et al., 2019](#)). The shared task used the MTNT dataset ([Michel and Neubig, 2018](#)), which consists of noisy texts collected from Reddit and their translations between English and French and English and Japanese.

Improving Model Robustness Using Noisy Data Majority of research on improving model robustness is dedicated to training on a mixture of original and noisy data. The same procedure is usually used for generating both the test corpus and training data ([Belinkov and Bisk, 2017](#); [Heigold et al., 2018](#); [Ribeiro et al., 2018](#); [Rychalska et al., 2019](#)).

To generate synthetic training data, researchers in machine translation and GEC often use so called back-translation ([Sennrich et al., 2016](#)). A reverse model translating in the opposite direction (i.e. from the target language to the source language or from the clean sentence into noisy sentence, respectively) is trained ([Rei et al., 2017](#); [Náplava, 2017](#); [Kasewa et al., 2018](#); [Xie et al., 2018](#)). It is then used on the corpus of clean sentences to generate noisy input data. While this approach might generate high-quality synthetic data, it requires large volumes of training data.

We evaluate two approaches to alleviate performance drop on noisy data: We either train the system on a mixture of synthetic (generated statistically from real error corpora) and original authentic data; or we use an external grammatical-error-correction system to correct the noisy data before inputting them to the system itself. We are not aware of any other work that compares these two approaches and we believe that both approaches may be beneficial under certain conditions.

3 Modeling Natural Noise from Corpora

Robustness of NLP models to natural noise would ideally be evaluated on texts with authentic noise, with error corrections annotated by humans. (We present such authentic data evaluation in Section 7.) This perfect-world setting, however, requires an immense annotation effort, as multiple target domains have to be covered by well-educated human annotators for multiple NLP tasks in a range of languages. To ease the annotation burden, we propose

a new framework, named KaziText, for introducing natural-like errors in a text.

The core of KaziText is a set of several common error type classes, *aspects* (following naming convention of Rychalska et al., 2019). The *aspects* are composable (can be combined) and the probability of the aspect manifestation as well as the aspect’s internal probabilities are estimated from grammatical-error-correction corpora.

3.1 Noising Aspects

One of the main objectives of our error aspects’ design was to avoid manually designed rules, especially those derived from a single language. An ideal approach, automatically inferring the aspects themselves, is however limited by the amount of available data. Therefore, we defined a rich set of aspects which can be estimated from the data:

1. **Diacritics** Strip diacritics either from a whole sentence or randomly from individual characters.
2. **Casing** Change casing of a word, distinguishing between changing the first letter and other.
3. **Spelling** Insert, remove, replace or swap individual characters (*wrong* → *worng*) or use ASpell¹ to transform a word to other existing word (*break* → *brake*).
4. **Suffix/Prefix** Replace common suffix (*do* → *doing*) and prefix (*bid* → *forbid*).
5. **Punctuation** Insert, remove or replace punctuation.
6. **Whitespace** Remove or insert spaces in text.
7. **Word Order** Reorder several adjacent words.
8. **Common Other** Insert, replace or substitute common phrases as seen in data (*the* → *a*, *a lot of* → *many*). This is the aspect which should learn language specific rules.

The natural errors found in real-world texts rarely fall into mutually exclusive categories. Casing errors are also spelling errors; common other aspect covers all other aspects. Therefore, some of the aspects naturally overlap. We therefore opted for evaluating the aspects in a *cumulative* manner in the designed order.

When designing the order of the aspects, our goal was to respect the natural inclusion of aspects and also error severity. We therefore start with diacritical-only changes, given that for example in

¹<http://aspell.net/>

Czech, users may deliberately write without diacritics. We then add casing changes, spelling errors and then suffix/prefix changes (the latter being morphologically motivated spelling errors). The first four aspects do not modify tokenization, making them suitable for tokenization-dependent tasks like POS tagging or lemmatization.

The remaining aspects change the number of tokens or token boundaries. The punctuation, whitespace and word order aspects are relatively independent, with the common other aspect covering all of them and thus being the last one.

3.2 Estimating Noising Aspects Probabilities

We use grammatical-error-correction (GEC) datasets to estimate probabilities of individual aspects. The GEC datasets are distributed in M2 format,² which for a tokenized input noisy sentence contains a set of correcting edits. Each correcting edit contains the corresponding input sentence span, the correction itself and the error type. The noising aspect probabilities are estimated by frequency analysis.³ To accurately model the distribution of amount of errors in different sentences, we also measure the standard deviation of the token edit probability per sentence.

We collected M2 files from various grammatical-error-correction corpora in 4 languages: English, Czech, Russian and German. The majority of annotated content comes from Second Learners of the particular language and in addition, more speaker groups are available in English and Czech:

- English
 - Natives: LOCNESS v2.1 (Granger, 1998)
 - Second Learners: NUCLE (Dahlmeier et al., 2013), FCE (Yannakoudakis et al., 2011), Write & Improve (Yannakoudakis et al., 2018)
- Czech
 - Natives: essays of Czech primary schools students, in submission process
 - Natives Informal: web discussions data, in submission process
 - Second Learners: AKCES-GEC (Šebesta et al., 2019)
 - Romani: AKCES-GEC (Šebesta et al., 2019) – Romani ethnic minority children and teenagers using Czech

²GEC file format since the CoNLL-2013 shared task

³We refer to the published source code for details.

Language	Corpus	Sentences	Error rate	Domain
English	<i>NUCLE</i> (Dahlmeier et al., 2013)	57 151	6.6%	SL
	<i>FCE</i> (Yannakoudakis et al., 2011)	33 236	11.5%	SL
	<i>W&I</i> (Yannakoudakis et al., 2018)	37 704	11.7%	SL
	<i>LOCNESS</i> (Granger, 1998)	988	4.7%	native students
Czech	Romani part of <i>AKCES-GEC</i> (Šebesta et al., 2019)	16 030	20.3%	Romani heritage speakers
	SL part of <i>AKCES-GEC</i> (Šebesta et al., 2019)	31 341	22.1%	SL essays
	<i>Natives Informal</i>	11 608	15.6%	web discussions
	<i>Natives</i>	7 696	5.8%	native students
German	<i>Falko-MERLIN</i> (Boyd, 2018)	24 077	16.8%	SL essays
Russian	<i>RULEC-GEC</i> (Rozovskaya and Roth, 2019)	12 480	6.4%	SL, heritage speakers

Table 1: Comparison of used GEC corpora in size, token error rate and domain. SL = second language learners.

- German (Second Learners): Falko-MERLIN GEC Corpus (Boyd, 2018)
- Russian (Second Learners): RULEC-GEC (Rozovskaya and Roth, 2019)

An overview of the sizes and error rates of the datasets above is presented in Table 1.

We call the resulting single file containing all aspect probabilities for one group of speakers a *profile*. The *profile* therefore describes the grammatical style of a particular given group of users, derived from M2 file annotations.

Each profile has a development and test version originating from the respective M2 development and test files. The test profiles are used for synthesising data intended directly for assessing models’ performance in noisy setting while the development profile is intended for creating data for training the models.

3.3 Adjusting the Percentage of Token Edits

In order to reach an intended percentage of *token edits*, which directly corresponds to the amount of noise in the generated data, we correspondingly scale the aspects’ probabilities. We refer to the percentage of token edits in the original corpus as a *corpus error level*.

3.4 Noising the Data

When noising an input sentence, we first sample a token edit probability from the error amount distribution, scaled according to the required number of token edits. We then introduce the desired aspects with the chosen error level.

We allowed the framework to generate any noising aspect, including adding new tokens, in test sets without token-level gold annotations: neural machine translation, GLUE benchmark, tokens outside named entities in NER and to-

kens outside the answer in reading comprehension.

When introducing errors into classification test sets with token-level gold annotations, we need to maintain the original tokenization. For this reason, we allowed only the first 4 aspects for the following data: morpho-syntactic analysis, tokens inside named entity spans in NER and tokens inside answers in the reading comprehension task.

All experiments are repeated with 5 different random seeds and we report means with standard deviations.

4 Evaluated Tasks

4.1 Morpho-syntactic Analysis

Model We employed UDPipe (Straka et al., 2019), a tool for morpho-syntactic analysis.

Dataset We used the Universal Dependencies 2.3 (Nivre et al., 2018) corpus (UD 2.3).⁴

Metrics We utilized the following metrics (Zeman et al., 2018) – **UPOS**: coarse POS tags accuracy, **UFeats**: fine-grained morphological features accuracy, **Lemmas**: lemmatization accuracy, **LAS**: labeled attachment score and **MLAS**: combination of morphological tags and syntactic relations.

4.2 Named Entity Recognition

Model Recently published architecture (Straková et al., 2019) was used for NER evaluation.

Dataset For English and German, we evaluated on the standard CoNLL-2003 shared task data (Tjong

⁴Many English UD test set tokens contain casing or spelling errors, propagated into lemmas, rendering such data unsuitable for analysis. We try to use only error-free English test documents and we therefore drop all test documents containing a sentence starting with a lowercase character, keeping more than half of the data. Apart from the lemmatization accuracy, the results for full test set are nearly identical.

Kim Sang and De Meulder, 2003); for Czech, we used a fine-grained Czech Named Entity Corpus 2.0 (Ševčíková et al., 2007) with 46 types of nested entities.

Metric The evaluation metric is F1 score.

4.3 Neural Machine Translation

Model We chose a state-of-the-art Czech-to-English NMT system CUBBITT (Popel et al., 2020), but we trained it on the newest version (2.0) of the CzEng parallel corpus (Kocmi et al., 2020). We trained with batch size of ca. 23k tokens for 550k steps, saved a checkpoint each hour (ca. 4600 steps) and selected the checkpoint with the highest dev-set BLEU (which was at 547k steps).

Dataset We use WMT17 (newstest2017, 3005 sentences)⁵ as our development set. Our test set is a concatenation of WMT13, WMT16 and WMT18 (8982 sentences in total).

Metric We evaluate the translation quality with case-insensitive BLEU score.⁶

4.4 GLUE Benchmark

We select a subset of GLUE (Wang et al., 2018) tasks, namely Microsoft Research Paraphrase Corpus (MRPC), Semantic Textual Similarity Benchmark (STS-B), Quora Question Pairs (QQP) and The Stanford Sentiment Treebank (SST-2). We finetune BERT on each of these tasks and evaluate them on various levels of noise.

Model We finetune pretrained BERT with an additional feed-forward neural network with one hidden layer predicting score on particular task’s data. We use *bert-base-cased* configuration and Hugging-Face’s Transformers (Wolf et al., 2019) implementation.

Dataset We use official GLUE datasets as provided by <https://gluebenchmark.com/tasks>.

Metric We report following metrics: F1 for MRPC and QQP, Pearson-Spearman Corr for STS-B and accuracy for SST-2.

⁵<http://statmt.org/wmt17>

⁶We use SacreBLEU (Post, 2018) with signature BLEU+case.lc+numrefs.l+smooth.exp+tok.intl+version.1.4.14. When using the case-sensitive version, the results show similar trends, except for the Casing aspect, which causes more harm to the score, as could be expected. However, it is questionable if copying the “wrong” casing to the translation (e.g. not capitalizing the first word in a sentence or using all-uppercase) should be considered a translation error. We thus opted for case-insensitive BLEU as our primary metric.




Figure 1: Proportional distribution of the first 4 aspects (diacritics, casing, spelling, affixes) in Czech and English.

4.5 Reading Comprehension

Model We utilize a BERT base architecture with a standard SQuAD classifier on top (Devlin et al., 2019).

Dataset We employ English SQuAD 2 (Rajpurkar et al., 2018) and its Czech translation (Macková and Straka, 2020).


Metric Our experiments are evaluated using F1 score.

5 Robustness to Noise

We evaluated the models robustness both to the amount of noise (Figure 2) and to error types (Figures 3 and 4).


A unifying trend can be observed in models performance with respect to increasing percentage of token edits. Solid lines in Figure 2 display the morpho-syntactic MLAS, NER F1 and NMT BLEU on texts with up to 30% of token edits. The relative performance decreases roughly linearly with the amount of token edits, in accordance with previous findings (Rychalska et al., 2019). The tendency is consistent across tasks, languages and profiles: For example, compare the Czech and English Second Learners profiles in morpho-syntactic analysis (Figure 2a) or Czech Native Speakers and Czech Second Learners profiles in the NMT clean model (Figure 2c), which exhibit similar behaviour despite their differing distributions of aspects (Figure 1). This consistency implies that it is the sheer amount of noise rather than the distribution of aspects, that contributes to the model performance deterioration. More results are available in Supplementary Material (Figures S2 and S3).

Estimating the amount of noise is important, as the *corpus error level* differs greatly across languages and profiles. For example, compare the Second Learners profile in English (11.3% token edits) and Czech (27.1%) in Figure 2a, or in Czech, see Native Speakers (6.4%) and Second Learners



(a) Morpho-syntactic analysis (relative MLAS), Second Learners

(b) NER (relative F1), Second Learners



(c) NMT (BLEU), Czech Native Speakers and Second Learners

(d) GLUE (relative task score), English Native Speakers and Second Learners

Figure 2: Increasing percentage of token edits with clean model, noise-trained model and grammatical-error-correction. Numbers near lines are absolute values.

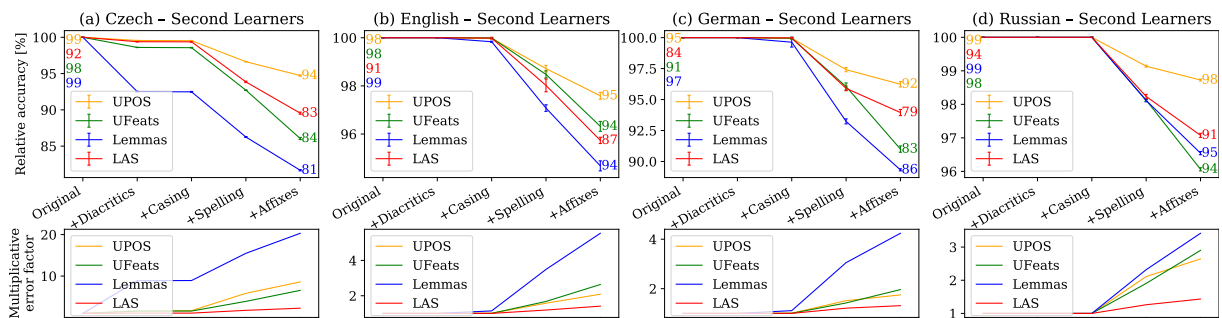

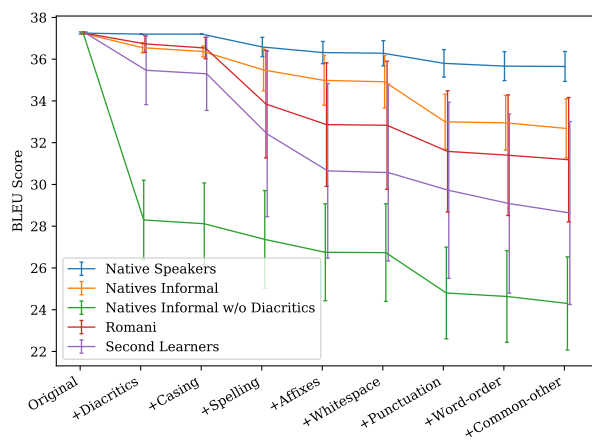


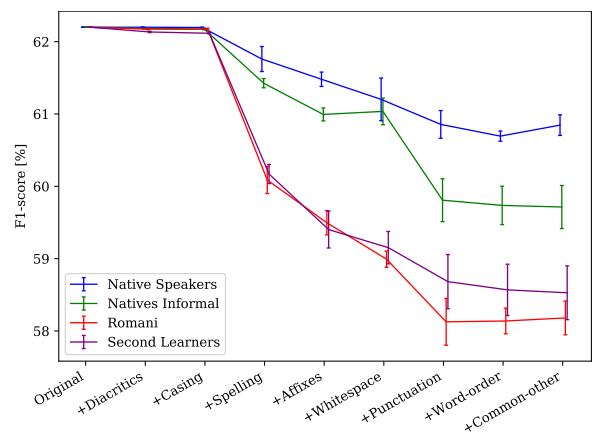
Figure 3: Morpho-syntactic tasks with additive noising aspects in Second Learners profile across (a) Czech, (b) English, (c) German and (d) Russian. The amount of introduced errors is the corpus error level for each aspect. **Upper row** Accuracy relative to original data accuracy, numbers near lines are absolute values. **Lower row** Multiplicative factor of errors to original data error rate.



(a) NER (relative F1), selected profiles in CS, EN and DE



(b) NMT (BLEU), all Czech profiles



(c) Reading comprehension (relative F1), all Czech profiles

Figure 4: Evaluation with additive noising aspects. The amount of introduced errors is the corpus error level for each aspect. Numbers near lines are absolute values.

(27.1%) in Figure 2c. Testing near the estimated noisiness level provides more accurate evaluation of the models' performance.

From a qualitative point of view, spelling and

affixes make for the major performance drop in morpho-syntactic analysis (Figure 3), NER (Figure 4a) and NMT (Figure 4b).

Some tasks are more sensitive to certain aspects: Casing is a crucial aspect for NER. This is clearly shown in the Czech Natives Informal profile, which contains text scraped from the internet discussions and contains nontrivial amount of casing errors (Figure 4a). We further elaborate the casing aspect effect on NER in Section S2 in Supplementary Material. In NMT and reading comprehension, errors in punctuation seem to decrease the model performance consistently across all profiles (Figures 4b and 4c, respectively).

For Czech as a language with diacritic marks, diacritics is an interesting aspect. We can see that when it is introduced at a *corpus error level*, the Czech model's performance on Lemmas drops by circa 7 percent. Figure S1 in Supplementary Material further illustrates that performance significantly deteriorates when all diacritics is stripped, which is quite common in informal Web texts. Similarly, to emphasize the effect of the diacritization aspect on NMT, we created a new profile *Natives Informal w/o Diacritics* from the Natives Informal profile by stripping all diacritization. Figure 4b shows that not using diacritics at all results in a performance drop of ca. 10 BLEU points.

Some tasks are more sensitive to noise than others. Lemmatization is the most sensitive to errors (20 times more errors when processing Czech Second Learners texts with a clean model, see Figure 3), which is understandable, given that all lemma characters must be generated correctly from a corrupted surface token. The effect on POS tagging is the least pronounced (Figure 3), although 8 times as many errors in Czech (when processing noisy texts with a clean model) makes the POS tags much less reliable.

6 Noise-coping Strategies

We implemented and evaluated two strategies to alleviate the performance drop on noisy inputs: *external* and *internal* correction. In the *external* correction approach, we use a separately trained grammatical-error-correction model to de-noise texts before inputting them to the model itself. In the *internal* correction approach, we instead directly train the model on a combination of noisy and authentic texts.

We hypothesise that the external approach may

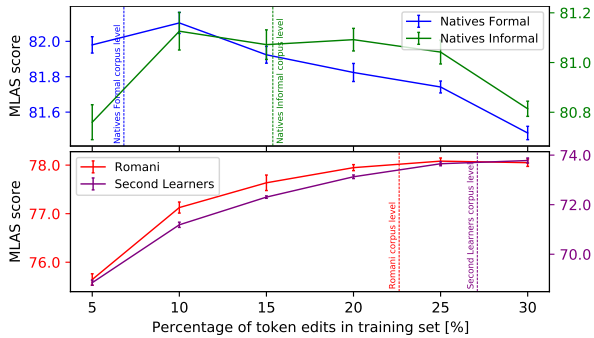


Figure 5: Morpho-syntactic analysis: Training data increasingly noised with each single profile, evaluation with the corresponding profile corpus error level.

be better in scenarios with small amount of annotated data. In such cases, only few iterations over training data are typically performed to prevent overfitting, and we suppose that learning the task itself and denoising at the same time would harm its performance a lot. Contrarily, with enough data and appropriate model capacity, learning the denoising and the task jointly may reduce the amount of potential false positives that might be otherwise proposed by the external language corrector.

6.1 External Correction Model

We use the grammatical-error-correction system of Náplava and Straka (2019) in our experiments. Their models trained on Czech, German and Russian achieve state-of-the-art results and slightly below state-of-the-art results on English. We use their “pretrained” version.

We modified the pipeline of Náplava and Straka (2019) to train on detokenized text. Furthermore, we also trained new grammatical-error-correction models which only make corrections that strictly keep the given tokenization (important in morpho-syntactic annotations). To sum up, we trained two types of grammatical-error-correction models: 1. detokenized error correction model (for NMT) 2. tokenization-preserving error-correction model (for morpho-syntactic tasks and NER).

6.2 Training on Noisy Data

In the *internal* approach to increase model robustness, we train the systems on a mixture of original and noisy data, while keeping the number of training steps unchanged. The noisy data are generated using the KaziText framework operating on development profiles and concatenated to original data.

We noise the training data with appropriately estimated corpus error levels in all our experiments.

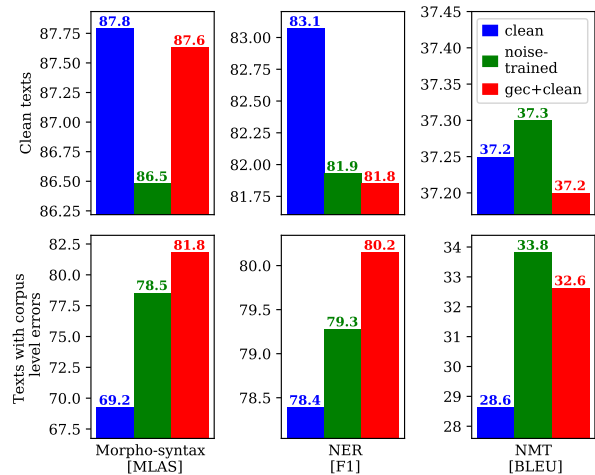


Figure 6: Comparison of three models (clean, noise-trained and GEC-preprocessed) on three tasks in Czech Second Learners profile. **Upper row** Original clean test data. **Lower row** Test data with corpus error level noise.

To illustrate the effect of noise level introduced into training data, we trained the UDPipe on variably noised morpho-syntactic data for all four Czech profiles. In each single profile, we increasingly noised the morpho-syntactic training data and evaluated on the testing data noised with the corresponding profile *corpus error level*. In all cases, the best performance is found near the corpus error level (Figure 5).

When training the NMT model, the best checkpoint on a development set consisting of concatenated standard WMT17 and WMT17 noised with our framework is selected.

We train a single model for each language on a concatenation of noisy data generated by all profiles of the particular language. This makes the final model generalize well across all profiles, although training a single model for each profile could make sense for other scenarios.

6.3 Evaluation

We present the effect of both the *internal* and *external* noise-coping strategies in Figure 2. There are two main points of interest in the graphs: the first one showing performance of models on clean texts and the second one showing model performance on texts with corpus level errors. Additionally, an excerpt showing performance of Czech Second Learners on these two levels is presented in Figure 6.

It is not a surprise that the model trained on clean training data surpasses the noise-coping models on

System	BLEU on Faust	
	Noisy	Cleaned
clean	43.3	50.9
noise-trained	47.0	50.5
gec+clean	44.1	50.4

Table 2: NMT results on authentic user noisy texts. We report BLEU on the Faust-Noisy test set with noisy input sentences and also on Faust-Cleaned that has manually corrected sentences on input.

the clean test data. Adapting to noise clearly comes with a cost. Surprisingly though, the clean model head start is only marginal in the NMT task.

The clean models perform substantially worse than either of the two proposed methods in all three tasks when errors are introduced in the same amount as the corpus error level (marked with vertical lines in Figure 2). Therefore, whenever noisy inputs of particular domain are expected, it is beneficial to adapt to noise using either of the two methods.

With increasing noise, the gap between the clean model and the *external* and *internal* model grows in all three tasks (Figure 2). There is a threshold at which the noise-coping models surpass the clean model for each task. Interestingly, the threshold oscillates around relatively low noise levels up to 5% of token edits.

Finally, we confirm our initial hypothesis that *external* approach with GEC model works better than *internal* approach on low resource tasks: morpho-syntactic analysis and named-entity recognition. The *internal* approach then outperforms *external* approach on machine translation task for which there is a large amount of training data and a model with greater capacity.

7 Evaluating on Authentic User Text

We assembled a new dataset for MT evaluation consisting of 2223 authentic Czech noisy input sentences translated into English, which we release at <http://hdl.handle.net/11234/1-3775>. The sentences originate from the project FAUST⁷ where they were collected from various users of reverso.net. The advantage of this dataset is that in addition to the original Czech noisy sentences, there are manually corrected Czech sentences and manual translations to English.

⁷<https://ufal.mff.cuni.cz/grants/faust>

On this dataset, we evaluate our neural machine translation models from Section 4.3 and Section 6, specifically the *clean* model trained on clean data, *noise-trained* model trained on a mixture of authentic and noised data and their combination with external grammatical-error-correction system. The results of these systems on authentic noisy texts are presented in Table 2. It is evident that noise-trained model outperforms clean model by a large margin on Faust-Noisy data while not losing much precision on Faust-Cleaned data. Similarly to our conclusions in Section 6, the external grammatical-error-correction system helps the clean model on noisy data, however is inferior to noise-trained model.

8 Conclusions

We estimated natural error probabilities statistically from real-world grammatical-error-correction corpora in order to model and generate noisy inputs for machine learning tasks. We extensively evaluated several state-of-the-art NLP downstream systems with respect to their robustness to input noise, both in increasing level of text noisiness and in variations of error types. We confirmed that the noise hurts the model performance substantially and we compared two coping strategies: training with noise and preprocessing with GEC, concluding that each strategy is beneficial in different scenarios. Finally, we also presented authentic noisy data evaluation using a newly assembled dataset for machine translation with authentic Czech noisy sentences translated to English. We release both the new framework (under MPL 2.0) at <https://github.com/ufal/kazitext> and the newly assembled dataset (under CC BY-NC-SA license) at <http://hdl.handle.net/11234/1-3775>.

Acknowledgements

This work has been supported by the Grant Agency of the Czech Republic, project EXPRO LUSyD (GX20-16819X). This research was also partially supported by SVV project number 260 575 and GAUK 578218 of the Charles University. It has also been using data provided by the LINDAT/CLARIAH-CZ Research Infrastructure (<https://lindat.cz>), supported by the Ministry of Education, Youth and Sports of the Czech Republic (Project No. LM2018101).

References

- Yonatan Belinkov and Yonatan Bisk. 2017. Synthetic and natural noise both break neural machine translation. *arXiv preprint arXiv:1711.02173*.
- Adriane Boyd. 2018. Using wikipedia edits in low resource grammatical error correction. In *Proceedings of the 2018 EMNLP Workshop W-NUT: The 4th Workshop on Noisy User-generated Text*, pages 79–84.
- Yo Joong Choe, Jiyeon Ham, Kyubyong Park, and Yeol Yoon. 2019. A neural grammatical error correction system built on better pre-training and sequential transfer learning. In *Proceedings of the Fourteenth Workshop on Innovative Use of NLP for Building Educational Applications*, pages 213–227.
- Daniel Dahlmeier, Hwee Tou Ng, and Siew Mei Wu. 2013. [Building a large annotated corpus of learner english: The nus corpus of learner english](#). In *Proceedings of the eighth workshop on innovative use of NLP for building educational applications*, pages 22–31.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186.
- Tao Ge, Furu Wei, and Ming Zhou. 2018. Reaching human-level performance in automatic grammatical error correction: An empirical study. *arXiv preprint arXiv:1807.01270*.
- Max Glockner, Vered Shwartz, and Yoav Goldberg. 2018. Breaking nli systems with sentences that require simple lexical inferences. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 650–655.
- Sylviane Granger. 1998. The computer learner corpus: A versatile new source of data for SLA research. In Sylviane Granger, editor, *Learner English on Computer*, pages 3–18. Addison Wesley Longman, London and New York.
- Georg Heigold, Stalin Varanasi, Günter Neumann, and Josef van Genabith. 2018. How robust are character-based word embeddings in tagging and mt against word scrambling or random noise? In *Proceedings of the 13th Conference of the Association for Machine Translation in the Americas (Volume 1: Research Papers)*, pages 68–80.
- Sudhanshu Kasewa, Pontus Stenetorp, and Sebastian Riedel. 2018. Wronging a right: Generating better errors to improve grammatical error detection. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4977–4983.
- Tom Kocmi, Martin Popel, and Ondrej Bojar. 2020. Announcing CzEng 2.0 parallel corpus with over 2 gigawords. *arXiv preprint arXiv:2007.03006*.
- Xian Li, Paul Michel, Antonios Anastasopoulos, Yonatan Belinkov, Nadir Durrani, Orhan Firat, Philipp Koehn, Graham Neubig, Juan Pino, and Hassan Sajjad. 2019. Findings of the first shared task on machine translation robustness. In *Proceedings of the Fourth Conference on Machine Translation (Volume 2: Shared Task Papers, Day 1)*, pages 91–102.
- Kateřina Macková and Milan Straka. 2020. Reading comprehension in czech via machine translation and cross-lingual transfer. In *Text, Speech, and Dialogue*, pages 171–179, Cham. Springer International Publishing.
- Paul Michel and Graham Neubig. 2018. Mntnt: A testbed for machine translation of noisy text. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 543–553.
- Jakub Náplava. 2017. Natural language correction. *Diploma Thesis*.
- Jakub Náplava and Milan Straka. 2019. Grammatical error correction in low-resource scenarios. In *Proceedings of the 5th Workshop on Noisy User-generated Text (W-NUT 2019)*, pages 346–356.
- Joakim Nivre et al. 2018. [Universal dependencies 2.3](#). LINDAT/CLARIN digital library at the Institute of Formal and Applied Linguistics (ÚFAL), Faculty of Mathematics and Physics, Charles University.
- Martin Popel, Marketa Tomkova, Jakub Tomek, Łukasz Kaiser, Jakob Uszkoreit, Ondřej Bojar, and Zdeněk Žabokrtský. 2020. Transforming machine translation: a deep learning system reaches news translation quality comparable to human professionals. *Nature Communications*, 11(1):1–15.
- Matt Post. 2018. [A call for clarity in reporting BLEU scores](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 186–191, Belgium, Brussels. Association for Computational Linguistics.
- Pranav Rajpurkar, Robin Jia, and Percy Liang. 2018. [Know what you don’t know: Unanswerable questions for SQuAD](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 784–789, Melbourne, Australia. Association for Computational Linguistics.
- Marek Rei, Mariano Felice, Zheng Yuan, and Ted Briscoe. 2017. Artificial error generation with machine translation and syntactic patterns. In *Proceedings of the 12th Workshop on Innovative Use of NLP for Building Educational Applications*, pages 287–292.

- Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 2018. Semantically equivalent adversarial rules for debugging nlp models. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 856–865.
- Alla Rozovskaya and Dan Roth. 2019. Grammar error correction in morphologically rich languages: The case of russian. *Transactions of the Association for Computational Linguistics*, 7:1–17.
- Alla Rozovskaya, Dan Roth, and Mark Sammons. 2017. Adapting to learner errors with minimal supervision. *Computational Linguistics*, 43(4):723–760.
- Barbara Rychalska, Dominika Basaj, Alicja Gosiewska, and Przemysław Biecek. 2019. Models in the wild: On corruption robustness of neural nlp systems. In *International Conference on Neural Information Processing*, pages 235–247. Springer.
- Karel Šebesta, Zuzanna Bedřichová, Kateřina Šormová, Barbora Štindlová, Milan Hrdlička, Tereza Hrdličková, Jiří Hana, Vladimír Petkevič, Tomáš Jelínek, Svatava Škodová, Petr Janeš, Kateřina Lundáková, Hana Skoumalová, Šimon Sládek, Piotr Pierscieniak, Dagmar Toufarová, Milan Straka, Alexandr Rosen, Jakub Náplava, and Marie Poláčková. 2019. [AKCES-GEC grammatical error correction dataset for czech](#). LINDAT/CLARIAH-CZ digital library at the Institute of Formal and Applied Linguistics (ÚFAL), Faculty of Mathematics and Physics, Charles University.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. Improving neural machine translation models with monolingual data. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 86–96.
- Magda Ševčíková, Zdeněk Žabokrtský, and Oldřich Krůza. 2007. Named entities in czech: Annotating data and developing ne tagger. In *Text, Speech and Dialogue*, pages 188–195, Berlin, Heidelberg. Springer Berlin Heidelberg.
- Milan Straka, Jana Straková, and Jan Hajič. 2019. [Evaluating contextualized embeddings on 54 languages in POS tagging, lemmatization and dependency parsing](#).
- Jana Straková, Milan Straka, and Jan Hajic. 2019. [Neural architectures for nested NER through linearization](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5326–5331, Florence, Italy. Association for Computational Linguistics.
- Erik F. Tjong Kim Sang and Fien De Meulder. 2003. [Introduction to the CoNLL-2003 shared task: Language-independent named entity recognition](#). In *Proceedings of the Seventh Conference on Natural Language Learning at HLT-NAACL 2003*, pages 142–147.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. 2018. Glue: A multi-task benchmark and analysis platform for natural language understanding. In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 353–355.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2019. Huggingface’s transformers: State-of-the-art natural language processing. *ArXiv*, abs/1910.03771.
- Ziang Xie, Guillaume Genthial, Stanley Xie, Andrew Y Ng, and Dan Jurafsky. 2018. Noising and denoising natural language: Diverse backtranslation for grammar correction. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 619–628.
- Helen Yannakoudakis, Ted Briscoe, and Ben Medlock. 2011. [A new dataset and method for automatically grading esol texts](#). In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies-Volume 1*, pages 180–189. Association for Computational Linguistics.
- Helen Yannakoudakis, Øistein E Andersen, Ardeshir Geranpayeh, Ted Briscoe, and Diane Nicholls. 2018. [Developing an automated writing placement system for esl learners](#). *Applied Measurement in Education*, 31(3):251–267.
- Daniel Zeman, Jan Hajič, Martin Popel, Martin Potthast, Milan Straka, Filip Ginter, Joakim Nivre, and Slav Petrov. 2018. [CoNLL 2018 shared task: Multilingual parsing from raw text to universal dependencies](#). In *Proceedings of the CoNLL 2018 Shared Task: Multilingual Parsing from Raw Text to Universal Dependencies*, pages 1–21, Brussels, Belgium. Association for Computational Linguistics.