

GECTurk: Grammatical Error Correction and Detection Dataset for Turkish

Atakan Kara, Farrin Marouf Safian, Andrew Bond, Gözde Gül Şahin

Computer Engineering Department
Koç University, Istanbul, Turkey
<https://gglab-ku.github.io/>

Abstract

Grammatical Error Detection and Correction (GEC) tools have proven useful for native speakers and second language learners. Developing such tools requires a large amount of parallel, annotated data, which is unavailable for most languages. Synthetic data generation is a common practice to overcome the scarcity of such data. However, it is not straightforward for morphologically rich languages like Turkish due to complex writing rules that require phonological, morphological, and syntactic information. In this work, we present a flexible and extensible synthetic data generation pipeline for Turkish covering more than 20 expert-curated grammar and spelling rules (a.k.a., writing rules) implemented through complex transformation functions. Using this pipeline, we derive 130,000 high-quality parallel sentences from professionally edited articles. Additionally, we create a more realistic test set by manually annotating a set of movie reviews. We implement three baselines formulating the task as i) neural machine translation, ii) sequence tagging, and iii) prefix tuning with a pretrained decoder-only model, achieving strong results. Furthermore, we perform exhaustive experiments on out-of-domain datasets to gain insights on the transferability and robustness of the proposed approaches. Our results suggest that our corpus, GECTurk, is high-quality and allows knowledge transfer for the out-of-domain setting. To encourage further research on Turkish GEC, we release our datasets, baseline models, and the synthetic data generation pipeline at <https://github.com/GGLAB-KU/gecturk>.

1 Introduction

Grammatical Error Correction (GEC) is among the well-established NLP tasks with dedicated shared tasks (e.g., BEA (Bryant et al., 2019)), benchmarks, and even specific evaluation measures. With increasing interest from the community, the field is

in constant need of novel tools, methods, and more importantly, extensions to other languages.

Recently, there has been an explosion of research about GEC for high-resource languages, especially for English (Rothe et al., 2021; Omelianchuk et al., 2020; Bryant et al., 2019). These recent techniques either formulate the task as neural machine translation, i.e., generation (Rothe et al., 2021), or token classification to detect erroneous tokens (Omelianchuk et al., 2020). The first approach mainly utilizes and engineers vanilla Transformers to generate the corrected text, while the second focuses on engineering a set of errors and transformation rules. Both formulations require a large set of parallel corpora containing grammatically correct and incorrect sentence pairs. Also, the latter approach additionally requires a highly curated dataset with annotations for correcting errors (i.e., location and type of the error). Constructing such a corpus with error annotations is nontrivial—especially for low-resource languages with rich morphology like Turkish. The challenge is due to grammar rules, a.k.a., writing errors being entangled in several layers, such as phonology, morphology, syntax, and semantics. As of today, there are no spelling or grammatical error datasets, as mentioned by Çöltekin et al. (2023), with the exception of the dataset introduced by Arıkan et al. (2019).

To address this, we focus on the Turkish Language and utilize the official writing rules established by the Turkish Language Association¹. We implement corruption, a.k.a. transformation, functions to generate instances that violate a specific rule, which requires challenging analysis of sentences on several linguistic levels, as well as curation of specialized lexicons. Then, we generate a large, synthetic, high-quality annotated corpus by applying transformation functions to professionally edited, modern Turkish articles. In addition to the

¹<https://tdk.gov.tr/>

transformation functions, we implement and share the reverse-transformation functions for validating the generated datasets and developing sequence tagger models, which achieve state-of-the-art in English.

In addition, we compile a corpus of movie reviews and manually annotate 300 sentences with the proposed error types to evaluate the models in a real-life setting. Furthermore, we design and implement several baselines using standard neural machine translation (NMT), sequence tagging and prefix-tuning. The NMT model only generates the correct sentence, but the sequence tagging model is trained to tag the tokens with the error type (if any) and then use our reverse transformation function on this error to generate the correct text. Finally, we perform prefix-tuning (Li and Liang, 2021) on mGPT Shliachko et al. (2022) for both detection and correction to test more recent techniques.

Our findings indicate that our pipeline approach using smaller models perform better than employing larger pretrained models in an end-to-end fashion when it comes to both synthetic and real-world datasets—particularly for the grammatical error detection task. Conversely, we observe that pretraining benefits the models in more realistic cases, despite the larger models still falling behind their simpler counterparts. Our results from out-of-domain tests imply that training on the synthetic dataset gives a strong prior to both smaller and larger models.

Our contributions can be summarized as follows:

- We propose the first comprehensive, expert-curated grammatical error schema for Turkish that covers 25 error types.
- We present a synthetic data generation pipeline that can be used to create arbitrary sized datasets, and can be easily extended to include new grammatical error types or lexicons, and can easily be modified to include custom tools (e.g., morphological analyzer and disambiguator).
- We present the first large-scale, fine-grained public dataset for Turkish grammatical correction and detection, along with a manually annotated realistic test set and strong baseline models.

We make our datasets, baseline models, and synthetic data generation pipeline publicly avail-

able at <https://github.com/GGLAB-KU/gecturk>.

2 Related Work

English GEC Despite having a long history, with the BEA-2019 Shared Task on Grammatical Error Correction (Bryant et al., 2019), the GEC community started using neural models and formulating GEC as a neural machine translation task (i.e., translate from grammatically incorrect to correct sentences), which has become the dominant approach. Another recent approach, GEC-ToR (Omelianchuk et al., 2020), uses the idea of reverse transformations, which can be applied to a list of source tokens $[x_1, \dots, x_n]$, in order to produce the desired correct grammar. They use a sequence tagger with a BERT encoder. Each tag corresponds to a transformation where transformations are applied after the sequence tagging finishes. In contrast, the gT5 model released by Xue et al. (2021) is a multilingual mT5 model fine-tuned on artificially corrupted sentences from the mC4 corpus and uses a span prediction and classification task to fix grammatical errors (Rothe et al., 2021). This does require a lot of additional training time, since the original mT5 model is not initially prepared for a similar task. Their model achieves SOTA results in 4 languages while only training once.

Turkish GEC Previously, Arıkan et al. (2019) proposed a neural sequence tagger model and a synthetically generated dataset to correct “de/da” clitic errors. In Turkish grammar, “de/da” is used both as a locative suffix and a conjunction meaning *also, too* that is written separately. For instance, “-de” is a locative suffix in the sentence “Evde (At home)”; while used as a conjunction here: “Ben de geliyorum (I’m coming too)”. Mistakes in using these clitics are common among native speakers, often due to some contextual subtleties and oral dialect influencing the written language. Öztürk et al. (2020) combined a contextual word embedding model, namely BERT, with a sequence tagger to correct “de/da” clitic errors. Although these errors are common, they constitute only a small portion of grammatical errors made by native speakers. In addition, while there are various forms of this error, the previous work only considers a few. Our data generation strategy considers multiple versions of the “de/da” clitic errors and many more common grammatical errors.

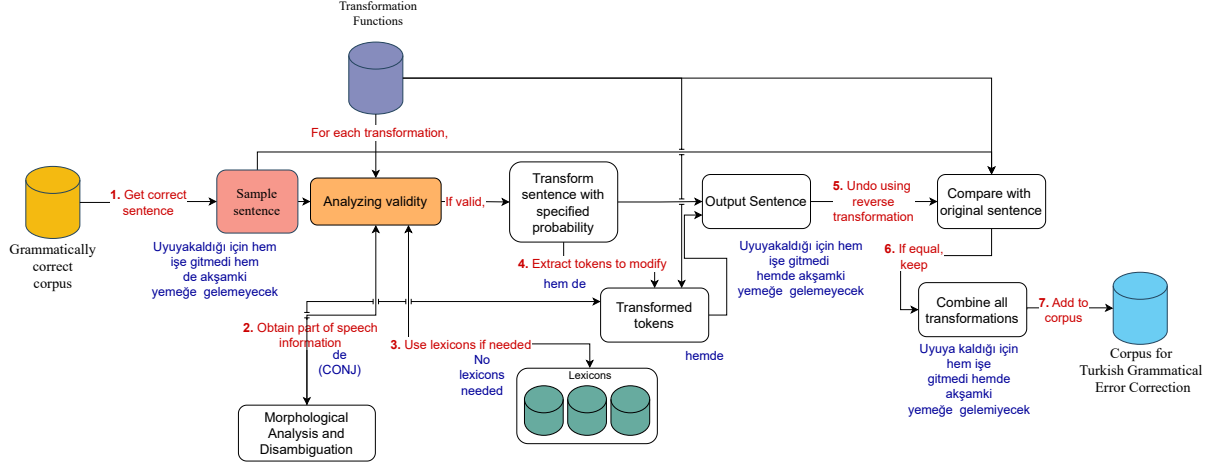


Figure 1: Data generation pipeline. 1) First, a correct sentence is obtained from the grammatically correct corpus. 2) Then, morphological analysis is performed. 3) The validity of the sentence for the transformation function is checked. If the sentence is eligible, the transformation is applied with some probability p . 4) First, selecting tokens to modify, 5) then, checking if the reverse transformations can recover the original tokens. 6) If original cannot be recovered, the sentence is removed. If can be recovered 7) the transformed sentence is added to the corpus.

Parsing-based Approaches While most of the current approaches focus on reverse transformations and sequence tagging, there are several studies that involve the use of parsing techniques. Flickinger and Yu (2013) create a parsing tree, and identify malformed parts of the tree to detect grammatical errors. da Costa (2021) use symbolic parsers and computational grammars for GEC and GED. On the same research line, Flickinger and Packard use bridged analyses combined with parsing to better allow for connecting two phrases in Head-driven Phrase Structure Grammar (HPSG).

3 Synthetic Data Generation

The overall generation process is given in Fig. 1. First, we randomly sample from professionally edited Turkish corpora (§3.1). Then, sentences are corrupted—if possible—following the expert-curated transformation rules explained in §3.2, as well as the use of a morphological analyzer. Finally, pairs of grammatically correct and corrupted sentences are added to the final Turkish GEC corpus following the M^2 scorer (Dahlmeier and Ng, 2012) (MaxMatch) data format.

3.1 Corpus

Our proposed data generation pipeline is built upon the assumption that all input sentences are grammatically correct. Hence, we base our study on previously compiled newspaper corpora (Diri and Amasyali, 2003; Amasyali and Diri, 2006; Can and

Amasyali, 2016; Kemik NLP Group, 2022) that are proofread and went through a professional editing process. The articles are on various topics, including politics, sports, and medicine, and have been written by more than 95 authors for three different newspapers; in total, more than 7000 singly authored documents were collected between 2004-2012. Once we obtained grammatically correct source sentences, we performed several preprocessing steps, such as removing duplicates (2.9% of the combined dataset), ending with 138K unique sentences.

3.2 Transformation

The Turkish Language Association (TDK)², a government agency founded in 1932, is responsible for providing resources to conduct scientific research on written and oral sources of Turkish. Within this scope, they specify and maintain a comprehensive list of publicly available writing rules³. We rely on this expert-curated list to generate forward and backward transformation rules, which we refer to as f and f^{-1} respectively. However, the list is long, and some writing rules are intuitive to native speakers, so that any errors made on these rules sound abnormal to them. We select the grammar rules that are most commonly used incorrectly by native speakers, determined by consultation with Turkish language experts and filtering the list from TDK

²<https://www.tdk.gov.tr/>

³<https://www.tdk.gov.tr/kategori/icerik/yazim-kurallari/>

using their feedback. We do not include any rules that are common for Turkish language learners but rarely made by native speakers. Table 1 provides the full list of the transformation rules produced by this work. The transformations rely on a morphological analyzer, which was essential to get the transformations right for a morphologically rich language like Turkish. For more information on Turkish Morphology, we refer to Oflazer (2014) and Lewis (1985).

Applying f For each sampled sentence, first, we shuffle the list of f s. This ensures that mutually exclusive transformation functions are applied with desired frequencies. Then, we iteratively apply each f on the sentence given with the pseudo-code in Algorithm 1. Here, f gets an input sentence s , morphological analysis of the sentence M_s , an array of indicators for whether any transformation has been applied to the word— $flags$, and parameter $p \in (0,1)$. The algorithm, then, iterates over tokens (or pairs) and checks whether the token has been transformed. If not, it checks whether the token(s) are eligible for f . If eligible, we apply f with the probability p , since not all errors are made with the same frequency by native speakers.⁴

Eligibility Check Some official writing rules require syntactic analysis at the token and sentence levels. For instance, to apply the transformation function CONJ_DE_SEP, one must perform morphological analysis and disambiguation to analyze the part-of-speech tags at the morpheme level. That is, CONJ_DE_SEP transformation can be applied **only if** a “-de/da” morpheme with a CONJUNCTION part-of-speech tag is found. Additionally, a small set of rules requires specialized lexicons, e.g., a list of exceptional foreign words for FOREIGN_R2_EXC. To address the former, we use a state-of-the-art morphological analyzer Dayanik et al. (2018), and the lexicons are taken from the official lists provided by TDK².

Annotation Format We use the standard GEC annotation format following Ng et al. (2013) and Bryant et al. (2019). An example annotation is given in Fig. 2. Here, S and A refer to the ungrammatical sentence and edit annotations respectively. Each A contains starting and ending indices, the error type, the corrected phrase, and the id of the annotator.

⁴We choose the probabilities intuitively after an initial analysis on web corpus and student essays.

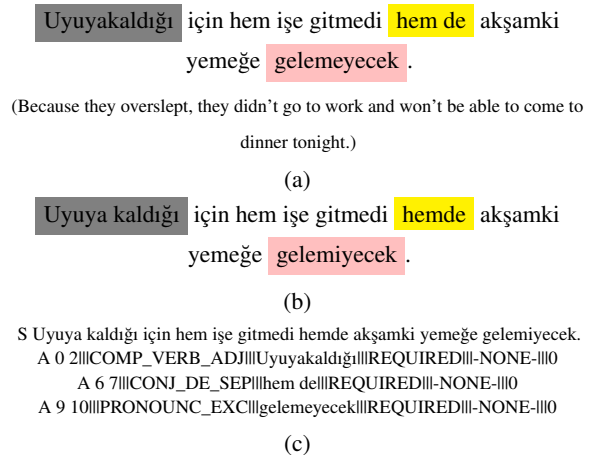


Figure 2: The grammatically correct sentence is given in (a), the transformed version is given in (b), and the annotation format is given in (c).

Postprocessing Despite the use of professionally edited source sentences, there are still some grammatically incorrect sentences that slip through. We detect these cases by taking advantage of a key property of our reverse transformations: since each f is reversible, we should see $S = f^{-1}(f(S))$ for each sentence S . Therefore, at the end of the transformation process, we perform this check on every generated, grammatically incorrect sentence. If a sentence fails this check, then we know it is problematic, and we remove it from the corpus. The final sentences are thus properly modified in the desired way, with no unintentional side effects.

3.3 Annotated Corpus

The annotated corpus includes more than 138K sentences, with 104K error annotations belonging to 25 error types given in Table 1. In this corpus, 50% of sentences are error free, in order for models to learn how to detect/correct sentences that are already grammatically correct. The generative pipeline controls the frequency of those error types, aiming to mimic the human error frequencies (see App. B). As in the dataset of CoNLL-2014 shared task (Ng et al., 2014), some error types appear more frequently than others. These frequencies are by the probability parameters p ; therefore, the difference between frequencies of error types is an intended result. Our dataset is finally split into a train/val/test set of 70%/15%/15%.

Algorithm 1 Apply f

Require: $s :=$ sentence, $M_s :=$ morphological analysis, flags, p
Ensure: tags
 tags $\leftarrow []$
 $n \leftarrow$ Number of tokens in s
 for $i = 1 \rightarrow n$ **do**
 if flags[i] **and** $is_eligible(s, M_s)$ **and** $flipCoin(p)$ **then**
 tags.insert("A i { $i + 1$ } ||| ruleID ||| sentence[i] ||| REQUIRED ||| -NONE- ||| 0")
 sentence[i] \leftarrow transformed token at index i
 flags[i] \leftarrow **False**

3.4 Curated Test Corpus

For a more realistic test setting, we used movie reviews from a popular website⁵ shared by Altinok (2023). We asked a domain expert to annotate the sentences for the proposed error types, following the standard GEC annotation format. As a result, we curated a test dataset of 300 sentences, wherein half of the sentences were grammatically correct and the other half contained errors.

4 Tasks and Models

In this paper, we consider two tasks: Grammatical Error Correction (GEC) and Grammatical Error Detection (GED).

Grammatical Error Correction (GEC) takes as input a grammatically incorrect sentence and outputs the corrected version of the sentence. Formally, given an input sentence $\mathbf{x} = (x_1, \dots, x_T)$ which may contain some grammar mistakes, the aim is to produce an output sentence $\mathbf{y} = (y_1, \dots, y_{T'})$ which contains no grammatical errors. Conditions are not imposed on how the model produces grammatically correct sentences.

Grammatical Error Detection (GED) takes a slightly different approach to this problem, with the goal of producing detailed information about the errors in the source sentence. This includes details about the type of error and the location of the error in the sentence. Formally, given an input sentence $\mathbf{x} = (x_1, \dots, x_T)$, we can represent the problem as a token-level classification task, where the output is $\mathbf{c} = (c_1, \dots, c_T)$, and c_i represents the error type of token i . Given the knowledge that an error of type c_k occurred at the location from m to n , it is then possible to apply the corresponding reverse transformation f^{-1} , and fix the error.

⁵www.beyazperde.com

4.1 Models

We introduce three models to evaluate the performance of GECTurk: An NMT baseline, a sequence tagger using BERT (Devlin et al., 2019) pretrained on Turkish, and mGPT using prefix-tuning. All models are trained using 1 Nvidia V100 GPU. We only provide the essential information about the models here. More details are available in Appendix A.

NMT Baseline: We train a vanilla transformer model (Vaswani et al., 2017) for GEC. This choice is inspired by the most recent shared task on grammatical correction (Bryant et al., 2019), where many of the winning teams used transformer-based models and modeled the problem as a Neural Machine Translation (NMT). The training dataset consists of triples $\{(x_i, y_i, a_i)\}_{i=1}^N$, where x_i is the i -th input sentence, y_i is the corresponding ground truth corrected sentence, and a_i are the annotations. During training, the model receives x_i as input and tried to predict y_i as output. Due to the nature of the formulation, NMT is only used for correction.

Sequence Tagger: Similar to recent work (Omelianchuk et al., 2020), we train a sequence tagging model using a cased BERT encoder, pretrained on Turkish text (Schweter, 2020) with default configurations and additional linear and softmax layers on the top. The BERT model uses the WordPiece tokenizer (Wu et al., 2016) that segments tokens into subwords. Therefore, each sentence in the dataset is first tokenized into subwords and passed into the BERT encoder. We only hold the first subword’s representation for words with multiple subword tokens. Then, the encoder’s representations are linearly transformed and passed to the softmax layer to classify into possible error types described in Table 1 or no error. The model is finetuned for token classification objective using cross-entropy loss.

The advantage of this model is the ability to per-

Category	Rule ID	Description	f	Frequency
-DE/-DA	1. CONJ_DE_SEP	Conjunction “-de/-da” is written separately.	Durumu [oğluna da -> oğlunada] bildirdi.	12962
	2. CONJ_DE_VH	Conjunction “-de/-da” must follow the vowel harmony	Çok [da -> de] iyi olmuş.	101
	3. CONJ_DE_AR	Conjunction “-de/-da” does not follow phonetic assimilation rules.	Sınıf [da -> ta] temizlendi.	99
	4. YADA	“-de/-da” written together with the word “ya” is always written separately.	Sen [ya da -> yada] o buradan gidecek.	472
	5. CONJ_DE_APOS	Conjunction “-de/-da” cannot be used with an apostrophe.	[Ayşe de -> Ayşe’de] geldi.	10859
	6. CASE_DE	Suffix “-de/-da” is written adjacent.	[Evde -> Ev de] hiç süt kalmamıştı.	37462
-KI	7. CONJ_KI_SEP	Conjunction “-ki” is written separately.	Bugün öyle çok [yorulmuş ki -> yorulmuşki] hemen yattı.	1817
	8. CONJ_KI_EXC	On some exceptional instances, by convention, the conjunction “-ki” is written adjacent.	[Belki -> Bel ki], [oysaki -> oysa ki], [çünkü -> çünkü]	1395
FOREIGN	9. FOREIGN_R1	Words that start with double consonants of foreign origin are written without adding an “-i” between the letters.	[gram -> gıram]	307
	15. FOREIGN_R2	Some foreign origin words undergo consonant assimilation	[sebebi -> sebepi]	327
	16. FOREIGN_R2_EXC	Exceptions to the FOREIGN_R2 rule	[evrakı -> evrağı]	422
BISYLL	13. BISYLL_HAPL_VOW	Some bisyllabic words undergo haplology when they get a suffix starting with a vowel.	[ağzı -> ağızı]	1567
	14. BISYLL_HAPL_VOW_EXC	Exception to previous rule	[içeride -> içerde]	866
LIGHT VERB	17. LIGHT_VERB_SEP	Light verbs such as “etmek, edilmek, eylemek, olmak, olunmak” are written separately in case of no phonological assimilation	[arz etmek -> arzetmek]	460
	18. LIGHT_VERB_ADJ	Light verbs are written adjacent in case of phonological assimilation e.g., liaison	[emretti -> emir etti]	547
COMPOUND	20. COMP_VERB_ADJ	Compound words formed by knowing, giving, staying, stopping, coming, and writing are written adjacent if they have a suffix starting with -a, -e, -i, -u, -ü.	[uyuyakalma -> uyuya kalmak], [gidedurmak -> gide durmak], [çıkma gelmek -> çıka gelmek]	7840
SINGLE	22. PRONOUN_EXC	Traditionally, some pronouns are written adjacent.	[hiçbir -> hiç bir], [herhangi -> her hangi]	3867
	23. SENT_CAP	The first letter of the sentence is capitalized.	[Onlar -> onlar] geldi.	2613
	24. CAPPED	Some Arabic and Persian originated words are written with capped letters.	[kâğıt -> kağıt], [karargâh -> karargah]	834
	25. ABBREV	Grammatical rules for abbreviations, such as adding suffixes to abbreviations, punctuations with abbreviations etc.	[Alm. -> Alm], [THY’de -> THY’da], [cm’yi -> cm’ye]	359
	12. PRONOUNC_EXC	Unlike its pronunciation, verbs ending with “-a/-e” do not mutate when they get a suffix other than “-yor”	[başlayacağım -> başlayacağım]	12750

Table 1: Selected list of writing rules introduced by Turkish Language Association. $f[arg1 -> arg2]$ refers to the transformation function where the correct and corrupted surface forms are given with $arg1$ and $arg2$ respectively.

Dataset	#s	#a	%e	#e
(Ours) GECTurk	138K	104K	49.7%	25
(Ours) CuratedTest	300	227	50%	25
BOUN	10K	6K	50.0%	2
BOUN complex	102	105	100.0%	2

Table 2: Datasets used for training and evaluation. #s: sentences, #a: annotations, %e: percentage of erroneous sentences, #e: number of errors types

form error detection easily, as opposed to simply error correction. Correction is simply performed with reverse transformations as described previously.

Prefix Tuning: Inspired by the recent successes of prefix tuning (Li and Liang, 2021) as an alternative to model fine-tuning, we use OpenPrompt (Ding et al., 2022) to perform prefix tuning on mGPT (Shliazhko et al., 2022). Despite being multilingual and primarily focused on other languages, mGPT achieves encouraging results on morphologically rich languages (Acikgoz et al., 2022). In prefix tuning, we append N trainable (soft) tokens to the front of each input. Therefore, given input $\mathbf{x} = (x_1, \dots, x_T)$, the new input becomes $\mathbf{x} = (s_1, \dots, s_N, x_1, \dots, x_T)$, where the s_i ’s are the added artificial tokens. We then optimize only these tokens during training, while leaving the original model frozen.

Here, we model both correction and detection tasks in the same sequence generation approach, where the corrected sentence is first generated, and then information about the violated rule, and the location of this error is generated at the end of the sentence. This allows for one trained model to output both results. In order to train this correctly, the target sentence was appended with the details of the error type and location, and used for loss calculations. An example is provided in Fig 6.

5 Experimental Setup

5.1 Datasets

The list of datasets and their statistics are given in Table 2. GECTurk and MovieReview datasets are already described in §3.3 and §3.4 accordingly. The BOUN dataset (Arikan et al., 2019) is a relatively smaller dataset of 15K training and 2K test sentences, containing only 2 error types. It also includes a complex split, a list of 100 sentences that are mentioned to be extra challenging by the authors.

5.2 Evaluation

Grammatical Error Correction Following the Omelianchuk et al. (2020) and Bryant et al. (2019), we report Precision (P), Recall (R), and $F_{0.5}$ scores using the M^2 (MaxMatch) scorer (Dahlmeier and Ng, 2012).

Grammatical Error Detection To allow for a fair comparison with the BOUN dataset from Arikan et al. (2019), we use the same metrics, namely Precision (P), Recall (R), and F_1 . Since the task is modeled as a sequence tagging problem, this aligns with the standard evaluation for sequence tagging, such as in Huang et al. (2015). For all GED results, we report macro metrics, which are computed by taking unweighted average of each classes’ result. We use macro metrics over micro ones since distribution of grammatical errors types made by humans are imbalanced. To calculate these scores, we use SeqEval (Nakayama, 2018), a common library for evaluating sequence tagging tasks and use one tag for each error type.

6 Experiments and Results

We train the baseline models using the setup explained in Section §5 with three different fixed seeds. The mean and standard deviation of their performances on GEC and GED are given in Table 3. As can be seen, both SeqTag and mGPT provide strong results over 0.94 $F_{0.5}$ score for the GEC task, compared to the NMT baseline model. On the other hand, the detection task is performed more competently by SeqTag—as expected—than mGPT, which again achieves around 0.90 F_1 score. Moreover, the experiment on the effect of dataset sizes shows that the proposed dataset is more challenging than existing ones, with a steeper learning curve due to the larger number of error types. More information on the dataset size experiments can be found in Appendix C.

Successful detection does not always translate into successful correction because erroneous edits undermine the grammatical accuracy of the sentence, even when the grammatical error is successfully detected. Here, the M^2 scorer does identify this anomaly and explains the decrease in the correction scores. It is also worth noting that GED and GEC are two separate tasks, and models handle them differently. For example, generative models such as mGPT generate both the corrections and detections by predicting the next tokens, so there

	GECTurk					
	Detection			Correction		
	P	R	F_1	P	R	$F_{0.5}$
NMT	-	-	-	0.50 ± 0.01	0.84 ± 0.01	0.55 ± 0.01
SeqTag	0.90 ± 0.003	0.90 ± 0.013	0.90 ± 0.006	0.98 ± 0.001	0.98 ± 0.001	0.98 ± 0.001
mGPT	0.52 ± 0.02	0.38 ± 0.004	0.41 ± 0.01	0.95 ± 0.01	0.92 ± 0.03	0.94 ± 0.01
Curated Test Data						
NMT	-	-	-	0.31	0.62	0.35
SeqTag	0.94	0.87	0.89	0.85	0.80	0.84
mGPT	0.73	0.52	0.59	0.75	0.61	0.72

Table 3: Detection and Correction results of the baselines on GECTurk (in-domain) and curated test dataset (out-of-domain).

isn’t necessarily a strong correlation between what is generated for each of them. On the other hand, SeqTag uses a pipeline approach, where it detects the errors first, then applies reverse transformations to fix them. Hence there is a stronger correlation between the detection and correction performances for SeqTag, as expected.

In order to test whether the performance of our models would transfer to different domains, we perform zero-shot experiments on the curated test set from the movie domain. We use our highest performing checkpoints on our synthetic test set, and evaluate without any additional training on the hand-annotated corpus as given in Table 3, second row. For detection, SeqTag performs similarly to synthetic setting, while mGPT’s performance *increases dramatically*, proving the importance of being exposed to real-life data during pretraining. However, SeqTag still outperforms mGPT by a large margin due to its classification objective. On the other hand, both models perform significantly worse on the correction task for the curated test data compared to the synthetic setting, suggesting larger room for improvement on this more challenging test set.

6.1 Knowledge Transfer

Next, we investigate the transfer capacity of our models on unseen datasets using a different set of errors (i.e., mostly a subset) originally introduced to our models. We first evaluate our pretrained models on the BOUN (Arikan et al., 2019) standard and complex test splits to see their zero-shot ability, seen in Table 4, first row. Surprisingly, our best model, SeqTag, achieves 0.80 F_1 , on-par with state-of-the-art for the standard split. It also surpasses state-of-the-art accuracy on the complex split by a large margin together with the mGPT model. This suggests that the error type knowledge

	BOUN Zero-Shot						
	Detection				Correction		
	P	R	F_1	Acc	P	R	$F_{0.5}$
NMT	-	-	-	-	0.19	0.56	0.22
SeqTag	0.91	0.72	0.80	0.99	0.81	0.63	0.77
mGPT	0.60	0.56	0.58	0.97	0.75	0.67	0.73
BOUN Complex Split Zero-Shot							
NMT	-	-	-	-	0.71	0.73	0.72
SeqTag	0.99	0.93	0.96	0.99	0.99	0.93	0.98
mGPT	0.61	0.96	0.90	0.93	0.79	0.75	0.78
BOUN Full Training							
NMT	-	-	-	-	0	0	0
SeqTag	0.97	0.86	0.91	0.99	0.97	0.86	0.94
mGPT	0.61	0.48	0.53	0.97	0.85	0.73	0.82
BOUN Results	0.92	0.82	0.87	0.71*	-	-	-

* This result is for BOUN Complex Split Zero-Shot task.

Table 4: Performance metrics of various models on the BOUN dataset. The table is divided into three sections: models trained on the GECTurk dataset and evaluated zero-shot on two different BOUN splits, and models exclusively trained and evaluated on BOUN.

is mostly transferrable to other domains. Similar to our results on GECTurk, mGPT scores much lower on detection compared to SeqTag. However, the performance of mGPT is higher on BOUN dataset due to the small number of error types that are relatively more balanced. We note that despite the claims made by the authors of the BOUN dataset, our results suggest no additional complexity in the “complex” split as shown in Table 4, second row.

Finally, we investigate the effectiveness of our general approach by training our proposed models from scratch on the BOUN (Arikan et al., 2019) training split, given in Table 4, BOUN Full Training. For this setup, the NMT model was not able to converge, and just produced noise, hence, shown as 0.

Following our previous results, SeqTag achieves F_1 score of 0.91, surpassing the state-of-the-art

by 0.04 pp, and its zero-shot performance by a large margin (0.11 pp). This suggests that pipeline approach is able to transfer a lot of knowledge. However, there is still a large gap that can be compensated by directly training on the actual domain and error types. On the other hand, the high scores provide cues for the strength of the proposed model. Surprisingly, prefix tuning of mGPT directly on the BOUN training dataset does not increase the performance compared to the zero-shot setting. This suggests two things: i) synthetic dataset such as ours, GECTurk, provides quality prior knowledge on the Turkish grammatical error types and ii) pre-trained models have a considerably larger transfer capacity compared to training from scratch, as expected. Furthermore, for languages where the error types are mostly identified and can be fixed by a set of rules, a pipeline approach such as SeqTag proves more effective, efficient and robust.

7 Conclusion and Future Work

In this work, we have presented an annotated dataset for Grammatical Error Correction (GEC) and Detection (GED), GECTurk, containing more than 20 Turkish writing error types proposed by Turkish language experts. We have also introduced a flexible and extensible data generation pipeline that can be used to create a synthetic dataset from grammatically correct sentences. We used this pipeline to create a large-scale dataset using multiple opinion columns from Turkish newspapers. In addition, we have manually constructed a more challenging test set by annotating the movie-reviews with the proposed error types.

Finally, we implemented a diverse set of strong baseline models, by training from scratch, fine-tuning, or using prefix tuning. Our results show that smaller models focusing on the simpler problem of detecting the error types outperform large pre-trained models on both the synthetic and real-life datasets, especially for the detection task. On the other hand, we observe that pretraining helps the models to handle more realistic cases, even though they still lag behind the simpler models. Our out-of-domain results suggest that training on the synthetic data gives a strong prior to both smaller and larger models.

Acknowledgements

This work has been supported by the Scientific and Technological Research Council of

Türkiye (TÜBİTAK) as part of the project “Automatic Learning of Procedural Language from Natural Language Instructions for Intelligent Assistance” with the number 121C132. We also gratefully acknowledge KUIS AI Lab for providing computational support. We thank our anonymous reviewers and the members of GGLab who helped us improve this paper.

Limitations

One key issue is that mGPT is very computationally intensive to work with, even when only doing prefix tuning. This meant we could only train for 1 epoch. Another limitation is that the data generation pipeline is very time-consuming due to the use of a morphological analyzer. This means many resources are needed for very large-scale datasets. Additionally, needing hand-crafted rules and reverse transformations makes adding new rules slow.

Another key limitation is the need of dictionaries for exceptions to grammatical rules. Words in Turkish originating from other languages (notably Persian, French, and Arabic) don’t always follow normal grammar rules, and thus require special lists of exceptions. While we included as many as possible, our list is not exhaustive, which can lead to rare edge cases where our pipeline fails. While dictionaries allow for adding learned knowledge directly, and is an invaluable part of our pipeline, these edge cases can cause problems during dataset generation.

Ethical Considerations

One ethical issue is the misuse of grammatical correction models for cheating. Having such models and datasets mean that students can more easily use these to score better than normal on assignments and exams. This is bad for the student’s learning, and also affects others who can be negatively impacted by this artificial success.

Despite this, we believe that grammatical error correction models are more beneficial than harmful. Many people, from authors to language learners, can benefit from having grammar corrections. By introducing a dataset and demonstrating models on Turkish, an under-served language in the NLP community, more people will be able to take advantage of this, similar to the many existing tools for English.

References

- Emre Can Acikgoz, Tilek Chubakov, Muge Kural, Gözde Şahin, and Deniz Yuret. 2022. **Transformers on multilingual clause-level morphology**. In *Proceedings of the The 2nd Workshop on Multi-lingual Representation Learning (MRL)*, pages 100–105, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- Duygu Altinok. 2023. **A diverse set of freely available linguistic resources for Turkish**. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 13739–13750, Toronto, Canada. Association for Computational Linguistics.
- M. Fatih Amasyalı and Banu Diri. 2006. Automatic Turkish text categorization in terms of author, genre and gender. In *Natural Language Processing and Information Systems*, pages 221–226, Berlin, Heidelberg. Springer Berlin Heidelberg.
- Ugurcan Arikan, Onur Güngör, and Suzan Uskudarli. 2019. **Detecting clitics related orthographic errors in Turkish**. In *Proceedings of the International Conference on Recent Advances in Natural Language Processing, RANLP 2019, Varna, Bulgaria, September 2-4, 2019*, pages 71–76. INCOMA Ltd.
- Christopher Bryant, Mariano Felice, Øistein E. Andersen, and Ted Briscoe. 2019. **The BEA-2019 shared task on grammatical error correction**. In *Proceedings of the Fourteenth Workshop on Innovative Use of NLP for Building Educational Applications*, pages 52–75, Florence, Italy. Association for Computational Linguistics.
- Ender Can and Mehmet Fatih Amasyalı. 2016. **Text2arff: A text representation library**. In *2016 24th Signal Processing and Communication Application Conference (SIU)*, pages 197–200.
- Çagri Çoltekin, A. Seza Dođruöz, and Özlem Çetinoglu. 2023. **Correction to: Resources for Turkish natural language processing: A critical survey**. *Lang. Resour. Evaluation*, 57(1):489.
- Luís Morgado da Costa. 2021. **Using rich models of language in grammatical error detection**. Ph.D. thesis, Nanyang Technological University.
- Daniel Dahlmeier and Hwee Tou Ng. 2012. **Better evaluation for grammatical error correction**. In *Proceedings of the 2012 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 568–572, Montréal, Canada. Association for Computational Linguistics.
- Erenay Dayanik, Ekin Akyürek, and Deniz Yuret. 2018. **Morphnet: A sequence-to-sequence model that combines morphological analysis and disambiguation**. *CoRR*, abs/1805.07946.
- 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, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 4171–4186. Association for Computational Linguistics.
- Ning Ding, Shengding Hu, Weilin Zhao, Yulin Chen, Zhiyuan Liu, Haitao Zheng, and Maosong Sun. 2022. **Openprompt: An open-source framework for prompt-learning**. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics, ACL 2022 - System Demonstrations, Dublin, Ireland, May 22-27, 2022*, pages 105–113. Association for Computational Linguistics.
- Banu Diri and Mehmet Fatih Amasyalı. 2003. Automatic author detection for Turkish texts.
- Dan Flickinger and Woodley Packard. Robust parsing in hpsg: Bridging the coverage chasm. Poster presented at the 22nd International Conference on HPSG. 2015.
- Dan Flickinger and Jiye Yu. 2013. **Toward more precision in correction of grammatical errors**. In *CoNLL Shared Task*.
- Zhiheng Huang, Wei Xu, and Kai Yu. 2015. **Bidirectional LSTM-CRF models for sequence tagging**. *CoRR*, abs/1508.01991.
- Kemik NLP Group. 2022. Our datasets. <http://www.kemik.yildiz.edu.tr/data/File/2500koseyazisi.rar>. Online; accessed 26-November-2022].
- Diederik P. Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980.
- G. Lewis. 1985. *Turkish Grammar*. Oxford University Press.
- Xiang Lisa Li and Percy Liang. 2021. Prefix-tuning: Optimizing continuous prompts for generation. *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, abs/2101.00190.
- Ilya Loshchilov and Frank Hutter. 2019. **Decoupled weight decay regularization**. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.
- Hiroki Nakayama. 2018. **seqeval: A python framework for sequence labeling evaluation**. Software available from <https://github.com/chakki-works/seqeval>.

- Hwee Tou Ng, Siew Mei Wu, Ted Briscoe, Christian Hadiwinoto, Raymond Hendy Susanto, and Christopher Bryant. 2014. [The CoNLL-2014 shared task on grammatical error correction](#). In *Proceedings of the Eighteenth Conference on Computational Natural Language Learning: Shared Task*, pages 1–14, Baltimore, Maryland. Association for Computational Linguistics.
- Hwee Tou Ng, Siew Mei Wu, Yuanbin Wu, Christian Hadiwinoto, and Joel Tetreault. 2013. [The CoNLL-2013 shared task on grammatical error correction](#). In *Proceedings of the Seventeenth Conference on Computational Natural Language Learning: Shared Task*, pages 1–12, Sofia, Bulgaria. Association for Computational Linguistics.
- Kemal Oflazer. 2014. [Turkish and its challenges for language processing](#). *Lang. Resour. Evaluation*, 48(4):639–653.
- Kostiantyn Omelianchuk, Vitaliy Atrasevych, Artem Chernodub, and Oleksandr Skurzhanyskiy. 2020. [GECToR – grammatical error correction: Tag, not rewrite](#). In *Proceedings of the Fifteenth Workshop on Innovative Use of NLP for Building Educational Applications*, pages 163–170, Seattle, WA, USA → Online. Association for Computational Linguistics.
- Sascha Rothe, Jonathan Mallinson, Eric Malmi, Sebastian Krause, and Aliaksei Severyn. 2021. [A simple recipe for multilingual grammatical error correction](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 702–707, Online. Association for Computational Linguistics.
- Stefan Schweter. 2020. [Berturk - bert models for Turkish](#).
- Oleh Shliakhko, Alena Fenogenova, Maria Tikhonova, Vladislav Mikhailov, Anastasia Kozlova, and Tatiana Shavrina. 2022. [mGPT: Few-shot learners go multilingual](#). *CoRR*, abs/2204.07580.
- 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*, volume 30. Curran Associates, Inc.
- Yonghui Wu, Mike Schuster, Zhifeng Chen, Quoc V. Le, Mohammad Norouzi, Wolfgang Macherey, Maxim Krikun, Yuan Cao, Qin Gao, Klaus Macherey, Jeff Klingner, Apurva Shah, Melvin Johnson, Xiaobing Liu, Lukasz Kaiser, Stephan Gouws, Yoshikiyo Kato, Taku Kudo, Hideto Kazawa, Keith Stevens, George Kurian, Nishant Patil, Wei Wang, Cliff Young, Jason Smith, Jason Riesa, Alex Rudnick, Oriol Vinyals, Greg Corrado, Macduff Hughes, and Jeffrey Dean. 2016. [Google’s neural machine translation system: Bridging the gap between human and machine translation](#). *CoRR*, abs/1609.08144.
- Linting Xue, Noah Constant, Adam Roberts, Mihir Kale, Rami Al-Rfou, Aditya Siddhant, Aditya Barua, and Colin Raffel. 2021. [mT5: A massively multilingual pre-trained text-to-text transformer](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 483–498, Online. Association for Computational Linguistics.
- Hasan Öztürk, Alperen Değirmenci, Onur Güngör, and Suzan Uskudarli. 2020. [The role of contextual word embeddings in correcting the ‘de/da’ clitic errors in Turkish](#). In *2020 28th Signal Processing and Communications Applications Conference (SIU)*, pages 1–4.

A Model Details

A.1 NMT Baseline

For tokenization, we used BerTurk-based (Schweter, 2020) tokenizer, passed to the NMT model. The transformer model has 6 encoders with embedding size 512, 6 decoder layers, and 8 heads. A dropout of 0.1 is used directly after the positional embeddings. For training, an Adam (Kingma and Ba, 2014) optimizer with $\beta_1 = 0.9$, $\beta_2 = 0.98$, and $\varepsilon = 1e - 9$, and a learning rate of $1e - 4$ is used. We used batch size of 32, and trained the model for 100 epochs on a single V100. We use a standard cross-entropy loss during training, as follows:

$$\mathcal{L}_{GEC}(\hat{y}, y) = - \sum_{n=1}^N \sum_{c=1}^V \log \frac{\exp(\hat{y}_{n,c}) y_{n,c}}{\sum_{i=1}^V \exp(\hat{y}_{n,i})} \quad (1)$$

Here, N is the batch size, V is the number of error classes, x is the model output, and y is the target. For the data size experiments, we used the same architecture but with slightly different hyperparameters. For both the 75% and 100% experiments, the model was trained for 100 epochs. For the 50% experiment, we only trained for 50 epochs. When training 10% and 25%, the Adam optimizer is used with the same β values, a learning rate of $5e - 4$, and a weight decay of $1e - 4$, for 100 epochs. The zero-shot testing on the BOUN (Arikan et al., 2019) dataset is tokenized with the same tokenizer, and the best pre-trained model from GECTurk is used for evaluation.

A.2 Sequence Tagger

For training, we used the AdamW (Loshchilov and Hutter, 2019) optimizer for 3 epochs, using batch size 16, learning rate $2e - 5$, weight decay 0.01, $\beta_1 = 0.9$, and $\beta_2 = 0.999$.

A.3 Prefix Tuning

We used the standard mGPT tokenizer and the OpenPrompt prefix tuning template. All experiments use 5 soft tokens at the beginning. Teacher forcing is used during training, and both the correction and detection tasks are formulated as a sequence generation problem. Following the settings from Acikgoz et al. (2022), we don't use weight decay for the bias and LayerNorm weights. The AdamW optimizer is used, with an initial learning rate of $5e - 5$, linearly decaying to 0 over the entire training. We clip the norm of the gradient at 1.0.




Figure 3: Number of sentences with each writing rule type.

Due to the computational requirements of mGPT, we only train on GECTurk for a single epoch on all experiments. However, on the smaller BOUN dataset, we train for 5 epochs. For inference, we also follow the hyperparameters from Acikgoz et al. (2022), using a temperature of 1.0, top p of 0.9, no repetition penalty, and a beam search of 5 beams. For all experiments, a batch size of 3 was used. The max sequence length, including soft tokens, is set to 512.

B GECTurk Error Frequencies

Fig. 3 shows the frequencies of each error type in GECTurk dataset.

C Effect of Dataset Size

In order to obtain a better understanding of how important the dataset size is for this task, we conducted training on 1, 10, 25, 50, 75, and 100 percent of GECTurk and evaluated each model using the same evaluation measures. Fig. 4 shows how the performance of the models vary with more training data. NMT reaches its top point with around 75% of the training data, while SeqTag and mGPT achieve high $F_{0.5}$ scores with 25% of the training split. However, as discussed before, correction scores can be misleadingly high, since high frequency and easy to correct errors will push the results much higher. Hence we also plot the F1 scores for the detection task both on TurkishGEC and BOUN datasets in Fig. 5. The plot shows that the GECTurk dataset is richer than the BOUN, since SeqTag and mGPT F1 performances are much steeper on the former.




Figure 4: Correction performance of each model trained on varying sizes of GECTurk




Figure 5: F_1 score performance of SeqTag and mGPT on GECTurk and Boun training splits.

S Ne yani Atatürk kadına ilgi duymuyormuydu?
A 5 6|||rule_11|||duymuyor muydu?|||REQUIRED|||-NONE-||
|0

→ M-GPT →

Ne yani Atatürk kadına ilgi duymuyor muydu?
;11;5;6

Figure 6: Example output of the mGPT model. The first line performs the grammatical error correction, and subsequent lines allow for detection. The annotations are included for comparison with model outputs, but are not actually provided to the model.