

# Automatic Correction of Human Translations

Jessy Lin<sup>◇♣</sup> Geza Kovacs<sup>♡</sup> Aditya Shastry<sup>♣</sup>  
Joern Wuebker<sup>♣</sup> John DeNero<sup>◇♣</sup>

<sup>◇</sup> University of California, Berkeley    <sup>♡</sup> Google    <sup>♣</sup> Lilt  
jessy\_lin@berkeley.edu, joern@lilt.com, john@lilt.com

## Abstract

We introduce translation error correction (TEC), the task of automatically correcting human-generated translations. Imperfections in machine translations (MT) have long motivated systems for improving translations post-hoc with automatic post-editing. In contrast, little attention has been devoted to the problem of automatically correcting human translations, despite the intuition that humans make distinct errors that machines would be well-suited to assist with, from typos to inconsistencies in translation conventions. To investigate this, we build and release the ACED corpus with three TEC datasets<sup>1</sup>. We show that human errors in TEC exhibit a more diverse range of errors and far fewer translation fluency errors than the MT errors in automatic post-editing datasets, suggesting the need for dedicated TEC models that are specialized to correct human errors. We show that pre-training instead on synthetic errors based on human errors improves TEC F-score by as much as 5.1 points. We conducted a human-in-the-loop user study with nine professional translation editors and found that the assistance of our TEC system led them to produce significantly higher quality revised translations.

## 1 Introduction

Despite recent progress in machine translation (MT), a tremendous amount of translated content in the world is still written by humans (DePalma, 2021). Humans are often assumed to produce trusted, high-quality translations. In reality, they do make errors, including spelling, grammar, and translation errors (Hansen, 2009). This paper introduces the task of **translation error correction** (TEC). Given a source sentence  $s$  and a human-generated translation  $t$ , the goal of TEC is to produce an improved translation  $t'$  by correcting all errors in  $t$ .

<sup>1</sup>Dataset available at: [github.com/lilt/tec](https://github.com/lilt/tec)

“Translation correction” has long been studied in the MT community through the task of automatic post-editing (APE), which aims to correct errors in machine-generated translations (Simard et al., 2007). TEC is structurally identical to APE. However, it requires modeling a different data distribution: errors made by humans, which differ from those made by MT systems (Freitag et al., 2021). We characterize the error distribution in TEC by building, analyzing, and releasing the ACED corpus, a collection of three TEC datasets from varying domains, with a total of 35,261 English–German translations produced and corrected by professional translators in the natural course of their work. While APE is dominated by the fluency errors that are characteristic of MT systems (74% of sentences), our TEC corpus exhibits a broader distribution of errors that human translators are prone to make.

Using this error analysis, we propose an approach for TEC that pre-trains on synthetic corruptions more similar to errors made by humans, outperforming models that were developed for the related tasks of MT, grammatical error correction, and APE on all ACED datasets.

The task of TEC is often currently performed by humans, e.g. translators hired to review and edit translations (“reviewers”). Can a TEC system help reviewers edit faster, or produce higher quality final translations than they would have without assistance? We ran a human-in-the-loop user study with nine professional translators using our best-performing TEC model. We found that the reviews produced when assisted with a TEC system were rated as higher quality than those produced without, and produced with less manual effort. Qualitatively, users commented that trust and consistency of the suggestions were critical. They speculated future automated assistance could be helpful for onboarding to new content, spotting technical errors, and improving their own awareness of errors to catch.

Error Type	Example Text
<b>Monolingual: typos</b>	<i>s</i> : Do your feet roll inwards when running?
	<i>t</i> : Kl <b>l</b> ppen deine Füße beim Laufen nach innen?
	<i>t'</i> : Kl <b>l</b> ppen deine Füße beim Laufen nach innen?
<b>Monolingual: grammar</b>	<i>s</i> : Own tough winter runs in the ...
	<i>t</i> : Bei harten Winterläufe sorgt der ...
	<i>t'</i> : Bei harten Winterläufen <b>n</b> sorgt der ...
<b>Monolingual: fluency</b>	<i>s</i> : The traffic emerges from the VPN server and ...
	<i>t</i> : Der <b>Verkehr</b> wird vom VPN-Server ausgegeben und ...
	<i>t'</i> : Der <b>Datenverkehr</b> wird vom VPN-Server ausgegeben und ...
<b>Bilingual</b>	<i>s</i> : Quad Core XEON E3-1501M, 2.9GHz
	<i>t</i> : Quad Core XEON 2,9 GHz
	<i>t'</i> : Quad Core XEON <b>E3-1501M</b> , 2,9 GHz
<b>Preferential</b>	<i>s</i> : VersaMax I / O auxiliary spring clamp style
	<i>t</i> : VersaMax Zusatz-E / A <b>Federklemmenart</b>
	<i>t'</i> : VersaMax Zusatz-E / A <b>Federklemmenbauform</b>

Table 1: Error taxonomy for the ACED corpus, with examples from the dataset.

Looking forward, a natural question arises of whether the research community should focus on learning to revise *model* outputs (APE) or *human* outputs (TEC). With recent improvements in MT, it has been increasingly difficult for APE models to improve model output that is already high quality (Chollampatt et al., 2020). On the other hand, we should expect that humans will continue to make errors. TEC models will continue to provide benefit as a way of assisting humans, whether for professional translation work or everyday language learners. TEC is synergistic with continuing advancements in MT: improved MT will lead to improved error correction for human-generated translations. While APE pits models against models, TEC is an opportunity to combine the best of humans and models *because* humans and models make different errors.

In sum, this paper revisits the notion of “translation correction” conceived narrowly as the MT-centered task of APE, with an empirical investigation of translation error correction (TEC), the task of learning to correct human translations. Our contributions are:

1. We release ACED, the first corpus for TEC containing three datasets of human translations and revisions generated naturally from a commercial translation workflow.
2. We analyze the kinds of errors humans make in ACED, finding that while APE is dominated by correcting translation *fluency*, TEC focuses on correcting a broader range of errors that

appear in translation.

3. We propose a pre-training approach for TEC that outperforms approaches developed for similar tasks such as APE. Together, our results suggest the need for distinct approaches to correct human translation errors.
4. We perform a human-in-the-loop user study, finding that professional translators produce *higher quality* translations when assisted by a TEC model.

## 2 The ♠ ACED Corpus for TEC

Given a source language sentence  $s$  and a human-generated translation  $t$ , the goal of TEC is to produce a corrected target language sentence  $t'$ .

We introduce the ACED corpus, a set of three TEC datasets: ASICS, EMERSON, and DIGITALOCEAN (DO), each consisting of English–German sentence triples  $(s, t, t')$  from varying domains.

ACED is a *real-world benchmark*, containing naturalistic data from a task humans perform, rather than manually annotated data. All translations were created by professional translators working with Lilt, a localization services provider. All translators have at least 5 years of professional translation experience and experience working with the customer and domain. Each document was translated from scratch (i.e. not post-edited) by a human translator using an interactive neural MT system. Each translated document was then reviewed by a reviewer, who Lilt selects as one of the more senior translators. As a result, the examples in our corpus exhibit

Dataset	Domain	# sentences			% edited	# edits (mean)	Edit distance (mean)
		train	dev	test			
ASICS	Marketing	1395	525	616	29 %	1.6	7.5
EMERSON	Technical	4287	1255	1662	20 %	1.5	5.8
DIGITALOCEAN (DO)	Technical	11773	6104	7644	8 %	1.5	7.1

Table 2: Corpus statistics for each dataset in the ACED corpus, including edit statistics on % of sentences that edited, and for edited sentences, the average number of edits and average edit distance.

real errors that translators make, and the corrected translations are publication quality.

Secondly, ACED is *diverse*, with the three datasets from varying domains exhibiting different error distributions and difficulty for initial work on the TEC task. Information for each dataset is shown in Table 2. The ASICS dataset consists of marketing content with product names and descriptions for an activewear company. EMERSON consists of industrial product names for a manufacturing company. DO consists of software engineering tutorials. The various content types pose different challenges for translators and thus for TEC systems, which we discuss in Section 2.1 and Section 2.2.

Duplicate sentences with the same source  $s$  were removed. A portion of sentences were rewritten by the reviewer rather than being edited (a relative edit distance of more than 25% and a minimum of two edited words). We replace  $t$  in these sentences with  $t'$  in the corpus, so that training and evaluation focuses on local edits rather than re-translations. The train, dev, and test splits were constructed by splitting along document boundaries.

## 2.1 How do TEC and APE errors differ?

To understand how the human errors in the TEC task differ from model errors in APE, we compare the types of errors in ACED with 100 randomly sampled errors in the WMT 2021 APE shared task dev set<sup>2</sup>, which we then annotate with error types. We define an error taxonomy that classifies each edit as one of three types: (1) **Monolingual edits** are identifiable from only the target-side text. We divide these further into subcategories that highlight different capabilities needed to correct edits: *typos* (including spelling, punctuation, spacing, orthographic issues), *grammar*, and *fluency* (awkward phrasing, word choice, or non-native-sounding disfluencies); (2) **Bilingual edits** concern mismatches between the source and target text, e.g. over- or under-translation, mis-translations; (3) **Preferen-**

**tial edits** correct text that is inconsistent with the preferences of the customer, as described in extralinguistic project requirements (e.g. terminology or stylistic preferences). Examples of each error type are shown in Table 1. Our error taxonomy closely mirrors those of previous analyses of human translation errors (Specia and Shah, 2014; Yuan and Sharoff, 2020; Gupta et al., 2021), and we confirm their findings that human translation errors differ from MT errors. However, while previous work focuses on error *detection* and quality estimation, TEC is concerned with error *correction*. Our error types are intended to isolate the capabilities that models need to learn to correct edits (e.g., target-side language models can learn to correct monolingual errors, but cannot do well on bilingual edits).

We annotate and release error labels for all test sentences in ASICS to enable its use as a diagnostic set for per-type evaluation of models. On the larger EMERSON and DO datasets, we randomly sample 50 errors to annotate for this analysis. Error types were annotated by a professional German translator. Each segment can have multiple error types. In Table 3, we report the percentage of sentences with at least one error of each type. 74% of sentences in APE exhibit a fluency error, in contrast to up to 22% of sentences in ACED, while other types like monolingual grammar, bilingual, and preferential errors are notably underrepresented in APE. We also note that all sentences in the APE shared task are edited, while a key feature of TEC is identifying when a sentence does *not* need to be edited. The error distributions suggest that different modeling techniques may shine in each: while APE challenges models to correct disfluent translations characteristic of MT systems, our task is designed to focus on identifying and correcting the typos, mismatches, and grammatical errors more commonly exhibited by humans. Guided by this observation, we describe an approach in Section 3 that pre-trains on synthetic edits that are more representative of this

<sup>2</sup><https://www.statmt.org/wmt21/ape-task.html>

	TEC		APE	
	ASICS	EMERSON	DO	WMT '21
Monolingual				
typos	13	16	22	16
grammar	41	4	2	6
fluency	22	0	20	74
Bilingual	22	70	32	5
Preferential	7	24	40	6

Table 3: Percentages of erroneous sentences that contain at least one error of each type for ASICS, EMERSON, DO, and the dev set of the WMT 2021 APE shared task. As a task, APE exhibits many more fluency errors than TEC.

error distribution.

## 2.2 How difficult is it to learn to edit?

To quantify how difficult it may be to learn the correct edits in ACED, we report statistics on *edit overlap*: what proportion of edits that we expect models to perform (e.g. adding a comma) appear exactly in the training set? We use the *errant* toolkit to identify discrete edits (Bryant et al., 2017). Each edit is represented as a tuple (*original span*, *replacement span*), e.g. (“auf”, “an”) to replace “auf” with “an.” Edit statistics are reported in Table 4: in ASICS and DO, ~20% of the total number of edits in dev and test appear in the training set, while EMERSON has ~60% of dev and test edits appear in the training set.

While the edit overlap rate provides a relative sense of scale for precision and recall numbers, it does not provide an upper bound on recall. It is possible to learn edits that do not exactly appear in the training set. For example, capitalizing product names (“winterized” → “WINTERIZED”) is a learnable pattern that would appear as many distinct edits. Additionally, some errors can be corrected without fine-tuning because they are generic typo, grammatical, fluency, or bilingual errors. Conversely, it is also possible that it is wrong to make an edit that appears in the training set, depending on the surrounding sentential context.

## 3 Approaches to TEC

We propose a TEC model and compare it to several models designed for related tasks to determine whether they are also effective for TEC. An overview of the differences between the models is shown in Figure 1. All models use the Transformer neural architecture (Vaswani et al., 2017) that generates the target sequence  $t'$  from left to right. All are pre-trained on 36M sentences from the WMT18

		ASICS	EMERSON	DO
Train	Total Edits	606	1436	1212
	Unique Edits	418	486	940
Dev	Total Edits	246	381	1004
	% in train	14	63	21
Test	Total Edits	287	364	766
	% in train	23	60	21

Table 4: Edit and edit overlap statistics for each ACED dataset: total number of edits, unique edits in each train split, and percentage of total edits in each dev and test split that appear in the train split.

translation task<sup>3</sup> and fine-tuned on ACED, unless indicated otherwise. All pre-training and fine-tuning data is pre-processed by normalizing punctuation with the Moses toolkit (Koehn et al., 2007).

All models have 6 encoder and decoder layers, model dimension of 256, feed-forward dimension of 512, and 8 attention heads. We use a joint English–German vocabulary with 33k byte pair encoding subwords (Sennrich et al., 2016). During pre-training, we set the dropout to 0.1 and use the Adam optimizer (Kingma and Ba, 2015) with a learning rate of 0.0002. During fine-tuning, we decrease the learning rate to 0.0001 and reset the Adam momentum parameters. We select the best fine-tuning checkpoint with edit-level  $F_{0.5}$  score on our dev set. We use greedy inference.

## 3.1 Dual-Source Encoder-Decoder Model

We first describe the dual-source encoder-decoder we use for the APE and TEC models. Formally, the original Transformer architecture (Vaswani et al., 2017) takes a sequence of  $J$  source tokens  $s_{1...J}$  and predicts a sequence of  $I'$  target tokens  $t'_{1...I'}$ . We adapt the architecture to additionally encode the original translation  $t$ , a sequence of  $I$  tokens,  $t_{1...I}$ . We independently project  $t$  into the embedding space, add an offset vector  $\mathbf{o}$ , and then concatenate the embedding with the embedding of the source  $s$  to form the encoder input. To allow the dual-source model to copy tokens from the original translation  $t$ , we implement the copy-mechanism proposed by Zhao et al. (2019), which augments the model with an additional encoder-decoder attention layer. An expanded description of the model can be found in Appendix Section A.

<sup>3</sup><http://www.statmt.org/wmt18>



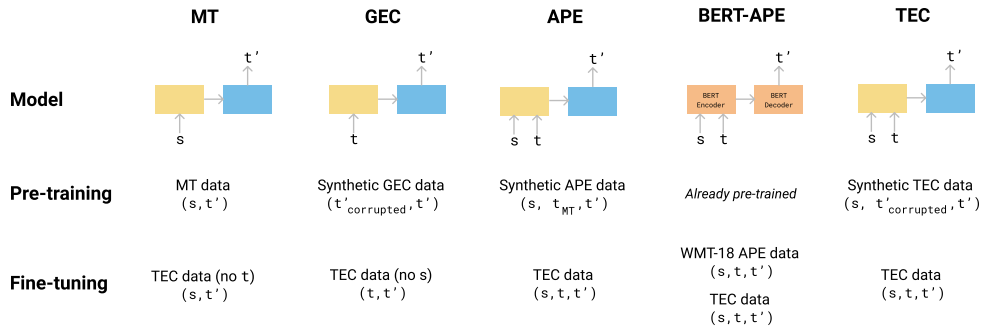


Figure 1: Overview of model architectures, pre-training data, and fine-tuning data for each approach to TEC. Transformer encoders and decoders are depicted as yellow and blue rectangles, respectively. The GEC and TEC models are pre-trained with synthetic corruptions of  $t'$  ( $t'_{\text{corrupted}}$ ), as detailed in the description of the TEC and GEC models. APE uses MT-generated translations of  $s$  ( $t_{\text{MT}}$ ) as synthetic data. BERT-APE is a state-of-the-art pre-trained APE model made available by [Correia and Martins \(2019\)](#).

### 3.2 Synthetic Data Generation

For the TEC and GEC model, we generate synthetic triples  $(s, t, t')$  for pre-training. We generate a synthetic  $t$  by corrupting the German side of the translation data into  $t'_{\text{corrupted}}$ . For each sentence, we sample the probability of corruption  $p_c \sim \mathcal{N}(\mu = 0.01, \sigma = 0.04)$  clipped at 0. On each character and word in that sentence, with probability  $p_c$ , we randomly select one of the following perturbations to apply at that position: insertion, deletion, transposition, repetition.

### 3.3 TEC Models

The five approaches we compare are:

**TEC (this work)** We implement the dual-source encoder-decoder model that encodes two inputs  $(s, t)$  and outputs  $t'$ , as described in Section 3.1, and then pre-train on synthetic data generated with the procedure in Section 3.2. We then fine-tune on ACED.

**MT** We train an English-German neural machine translation model (with the standard architecture described previously) and fine-tune it on  $(s, t')$  ACED pairs, ignoring the original translation  $t$ .

**GEC** We evaluate a encoder-decoder (monolingual) GEC model that takes an incorrect German sentence  $t$  as input and outputs a corrected  $t'$ . We use the same copy mechanism to attend to  $t$  as our TEC model. To pre-train, we perturb  $t'$  using the procedure described in Section 3.2, throwing away the source side to obtain  $(t, t') = (t'_{\text{corrupted}}, t')$  pairs. We then fine-tune on the ACED corpus, ignoring  $s$ .

**APE** We implement a dual-source encoder-decoder model that is identical to our TEC model. Following common practice in APE ([Junczys-Dowmunt and Grundkiewicz, 2016](#); [Negri et al., 2018](#)), we pre-train on synthetic “post-editing” triples  $(s, t, t')$  where  $t = t_{\text{MT}}$  is generated by translating  $s$  with an MT system. We split the training dataset into two parts, train an MT model on each half, and use each model to translate the other half of the dataset not seen during training. We then fine-tune on ACED.

**BERT-APE** We also evaluate whether a state-of-the-art APE model can be directly applied to our task. We evaluate the BERT-based encoder-decoder of [Correia and Martins \(2019\)](#), on which the WMT 2019 shared task winner was based ([Lopes et al., 2019](#)). Following [Correia and Martins \(2019\)](#), we fine-tuned on 23K English–German SMT triplets from the WMT18 shared task<sup>4</sup>. We reproduce their results on the APE shared task test sets, and continue fine-tuning this model on ACED. Following their paper, the inputs are pre-processed by tokenizing and joining the two inputs with a separator to form  $(s [\text{SEP}] t, t')$  pairs.

## 4 Results & Discussion

The primary metric for TEC is MaxMatch scores ( $M^2$ ) ([Dahlmeier and Ng, 2012](#)) computed with the *errant* toolkit ([Bryant et al., 2017](#)).  $M^2$  is a standard metric for GEC that aligns  $t$  and  $t'$  to extract discrete “edits.” We choose to follow the GEC evaluation practice of up-weighting precision by comparing  $F_{0.5}$ , since the original translation is

<sup>4</sup><http://www.statmt.org/wmt18/ape-task.html>, with 12K train triplets from 2016 and 11k from 2017

Model	ASICS			EMERSON			DO		
	Prec.	Rec.	F <sub>0.5</sub>	Prec.	Rec.	F <sub>0.5</sub>	Prec.	Rec.	F <sub>0.5</sub>
MT	3.3	<b>31.4</b>	4.0	16.2	<b>78.6</b>	19.2	1.2	<b>41.9</b>	1.5
GEC	52.8	6.6	22.0	78.0	53.3	71.4	20.5	2.0	7.1
APE	51.2	7.3	23.3	78.1	54.4	71.8	14.4	1.7	5.8
BERT-APE	6.8	10.8	7.3	32.0	57.8	35.1	2.3	3.8	2.5
TEC	<b>57.4</b>	9.4	<b>28.4</b>	<b>82.1</b>	57.2	<b>75.5</b>	<b>21.7</b>	2.0	<b>7.2</b>

Table 5: Main results on ACED. Our fine-tuned TEC model outperforms on F<sub>0.5</sub>. The fine-tuned MT model scores highest on recall because it makes many edits, but at the cost of unacceptably low precision.

mostly correct: it is better to suggest few correct edits than potentially introduce new errors.

Table 5 shows that TEC achieves the best overall F<sub>0.5</sub> score on all datasets, from +0.1 (on DO) up to +5.1 (on ASICS) above the next-best model. Fine-tuning on actual human corrections provides substantial gains; results without fine-tuning can be found in Appendix Section B.

Both the MT model (which ignores  $t$ ) and the GEC model (which ignores  $s$ ) underperform TEC. The MT model’s high edit recall can be attributed to the fact that it proposes many edits, greatly trading off precision. Without conditioning on  $t$ , direct MT translations of the source diverge from the reference. The GEC model obtains high precision but underperforms on recall. Conditioning on  $s$  not only makes it possible to propose bilingual edits, but also provides additional information to correct monolingual edits, as we show in Section 4.1.

### Can APE models be directly adapted for TEC?

Since our task is structurally identical to APE, a natural question is whether models that are trained on the APE objective can be directly adapted for TEC. The APE and TEC models differ only in pre-training, but the performance difference between them is substantial, indicating that the more GEC-like data synthesis procedure is a better fit for TEC than APE-style data synthesis via MT. Even more, the BERT-APE model, which is first fine-tuned to achieve state-of-the-art on APE before fine-tuning on ACED, achieves a particularly low F<sub>0.5</sub> score because it makes too many edits (low precision). Although future work may find insights in APE, our results emphasize that models that excel at correcting machine errors cannot be assumed to work well on human translations.

#### 4.1 Fluency & Per-category Error Analysis

We perform a more in-depth comparison using alternative metrics on ASICS, which includes an-

notated error labels as a diagnostic tool. First, to understand how much models are editing, we look at  $n$ -gram overlap with the GLEU metric (Napoles et al., 2015), a variant of BLEU used in GEC evaluation to measure the fluency of holistic rewrites (Sakaguchi et al., 2016). Next, we compare *sentence-level accuracy*, which measures exact match with  $t'$ . We compute overall sentence-level accuracy, which includes unedited sentences (which some models may incorrectly edit). We also report accuracy per error type over (edited) sentences annotated with that error type. These metrics need to be interpreted carefully: a no-edit baseline achieves a GLEU score of 87.85 (since original translations are mostly close in edit distance to the final) and sentence-level accuracy of 70.62% (the % of unedited sentences), outperforming models like MT and BERT-APE that make too many incorrect edits.

Our TEC model achieves the best score overall on both alternative metrics over all sentences, but various models outperform on specific error types. The full results are shown in Table 6. Examples of system outputs for different error types can be found in Appendix Section C.

Notably, APE models lose the most accuracy relative to our model on monolingual typo edits. This may be because neural MT decoders much less frequently introduce target-side errors that would be similar to typos (compared to the frequency of fluency errors). Still, the observation that different models do well at different errors suggests that future work can improve on TEC by leveraging the strengths of different models, e.g. using MT models to propose alternative translations.

## 5 User Study: Assisting Professional Translators with TEC

Our automatic evaluation shows how our TEC model can outperform other baseline systems, but

Model	GLEU	Sentence-level Accuracy (%)						
		Overall	Unedited	Mono. Typos	Mono. Grammar	Mono. Fluency	Bilingual	Preferential
		/616	/435	/16	/78	/41	/41	/14
No-edit	87.85	70.62	(100)	-	-	-	-	-
MT	44.79	31.66	40.00	0.00	11.54	<b>2.44</b>	<b>24.39</b>	<b>14.29</b>
GEC	88.46	71.10	97.01	<b>31.25</b>	14.10	0.00	0.00	0.00
APE	88.39	71.43	<b>97.47</b>	12.50	16.67	<b>2.44</b>	0.00	0.00
BERT-APE	82.35	39.77	52.18	6.25	<b>19.23</b>	0.00	4.88	0.00
TEC (this work)	<b>88.81</b>	<b>71.92</b>	97.01	25.00	17.95	0.00	7.32	0.00

Table 6: Additional analysis of n-gram overlap (via GLEU) and exact-match sentence accuracy over all test sentences (Overall) and per error category for the ASICS dataset. Number of sentences with an error of each type are indicated in gray, with some sentences containing errors from multiple categories. For each error category, we report the percentage of those sentences with at least one error of that type that a system predicted exactly.

we are ultimately interested in whether any TEC system is indeed useful in practice. Presently, TEC is done manually by humans. To investigate whether TEC systems can already be useful to humans—improving the quality, speed, or ease of human review—we performed a human-in-the-loop user study with our TEC model.

## 5.1 Methodology

We recruited 9 professional translators to serve as reviewers. None of them had prior experience with ASICS content. They were allowed to read and reference the sentences in the ASICS training set to familiarize themselves with the content and preferred terminology. Then, they were each assigned to review 74 sentences from the test set of ASICS. Of the 74 sentences, our TEC system predicted a suggested edit for 57 sentences, and for the remaining 17 sentences our TEC system did not predict any edits.

We opt for a within-subjects design to control for speed and experience differences between reviewers. For each reviewer, the 74 sentences were randomized such that half were in the “assisted condition” showing the TEC suggestion if available for the sentence, and the other half were in the “unassisted condition” where no TEC suggestion was shown. The reviewing interface is shown in Figure 2. If the sentence has suggestions available, the reviewer is asked to first accept or reject each of the suggestions. Then, they are asked to make edits to the text until they are satisfied with the translation. They then click a button to confirm their translation and move to the next sentence.

During the review process, we track:

1. Whether the TEC suggestion, if shown, was accepted or declined

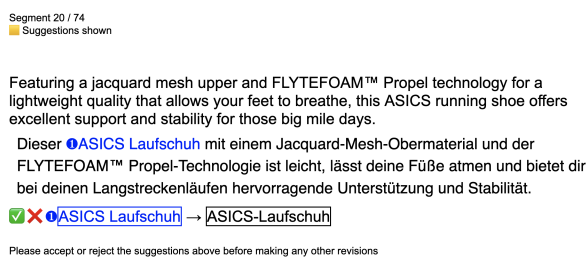


Figure 2: The interface used to show suggestions to reviewers in our user study.

2. Total time spent reviewing each sentence
3. Number of edit operations (insertions and deletions) the user made
4. Levenshtein edit distance from the original text to the final text

Finally, to evaluate whether TEC has an effect on quality, we asked a 10th translator to compare the quality of the reviewed sentences by ranking the 9 reviewed translations, with ties allowed. This translator was the translator who had reviewed the reference translations in the corpus, as the explicit goal is to ensure consistency with conventions in the training documents.

## 5.2 Results

In the user study, **79%** of TEC suggestions were accepted. For the purpose of analyzing effects of the TEC suggestions on time spent and translation quality, we will focus on only the 255 sentences (across 9 reviewers) where a TEC suggestion exists. Results are shown in Table 7. For all statistical significance tests, we use the Mann-Whitney (MW) U-test for testing statistical significance as all quantities are neither normally distributed nor log-normal.

Suggestions	Review Time (ms/char)	Inserts + Deletes	Levenshtein Distance
Hidden	361	0.0625	0.0347
Shown ( <i>Overall</i> )	367	0	0.0185
<i>Accepted</i>	328	0	0.0176
<i>Declined</i>	841	0.0625	0.0426

Table 7: Length-normalized medians from our user study for review times, number of characters the user inserted and deleted, and final Levenshtein edit distances from the original to the final translations. Data is split based on whether the suggestion was hidden or shown, and “Suggestion Shown” is further broken down according to whether the user accepted or declined the suggestion.

### 5.2.1 Effects of Suggestions on Time Spent During the Review Process

We first analyze how TEC suggestions influence the time spent reviewing a sentence. We compare the time durations normalized by the length of the sentence that needed to be reviewed (the *length-normalized review time*), as longer sentences require more time to be read and reviewed.

There is no significant difference in length-normalized review time when the suggestion is hidden vs. shown (MW  $U = 31654, p = 0.460$ ). When suggestions are shown, the length-normalized review time is significantly less on sentences where reviewers accepted the suggestion, compared to sentences where they declined (MW  $U = 3555, p < 0.0005$ ).

A potential explanation for these results is that when reviewers are shown incorrect TEC suggestions, they are distracted and slowed down, providing some evidence that precision should indeed be emphasized in automatic evaluations of TEC.

### 5.2.2 Effects of Suggestions on Edits Made During the Review Process

We also analyze the effects of suggestions on the editing effort, as measured by the number of characters the reviewer had to insert and delete, as well as how different the final reviewed sentences were from the original.

When suggestions are shown vs. hidden, there is a significant reduction in the number of insertions+deletions (MW  $U = 41348.5, p < 0.0001$ ). There is also a significant reduction when a shown suggestion is accepted vs. declined (MW  $U = 4007.5, p < 0.005$ ). There is no significant difference in the Levenshtein distance from the original translation to the final translation, between when a suggestion is shown vs. hidden (MW

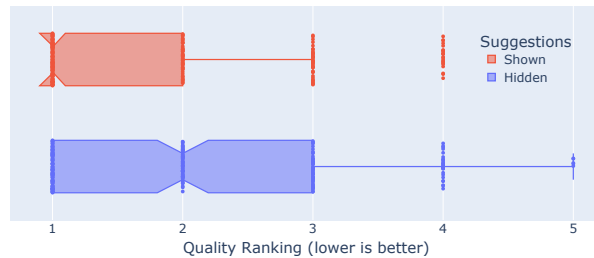


Figure 3: Box plot showing quality rankings of segments reviewed with suggestions hidden vs shown. Lower is better. Notches indicate median quality rankings. Bars indicate the upper fence ( $3^{rd}$  quartile +  $IQR \cdot 1.5$ ).

$U = 33750.0, p = 0.611$ ), or between when a shown suggestion is accepted vs. declined (MW  $U = 5485.0, p = 0.783$ ).

Thus, the TEC system suggestions help to significantly reduce the amount of manual typing that the user must perform.

### 5.2.3 Effects of suggestions on reviewed sentence quality

To assess the effects of TEC assistance on quality, we used the quality rankings produced by the independent reviewer. Quality rankings were not normally distributed, so we use the Mann-Whitney U-test for testing statistical significance. A box plot of the quality rankings is shown in Figure 3. The median quality ranking when the suggestion is shown is 1, vs. 2 when the suggestion is hidden. The quality ranking is significantly lower (meaning quality is higher) when the suggestion is shown, vs. hidden (MW  $U = 28738.0, p < 0.01$ ).

This suggests that showing TEC suggestions may be helping reviewers correct errors they may not have otherwise noticed, or help nudge them towards desired corrections.

## 5.3 Qualitative Findings

We also conducted a post-study survey for reviewers to report qualitative feedback. To understand common themes in the responses, we present all themes that at least two reviewers mention in their commentary.

### 5.3.1 The Role of Reliability and Trust

Five reviewers commented that *reliability is critical*: it was difficult to trust the system when they noticed some suggestions were incorrect, or the system did not reliably make an edit when applicable (e.g. always hyphenating when appropriate):



“Because I wasn’t sure I can trust the suggestions (because I saw several incorrect ones) so it took me longer to think/check whether the suggestion is right. And I have to read the entire sentence again anyway to check for other errors the suggestion didn’t catch...would only work if I knew 100% that the suggestions are always right”

These comments are in concordance with our quantitative findings. Perhaps unlike other assistive applications, it is not enough to only have high precision: if reviewers cannot trust that the system has caught most or all errors, they will not save time as they still have to read the entire sentence carefully. Conversely, a high-recall, low-precision system is not only distracting, but also leads reviewers to be suspicious of whether suggestions are correct in general. In general, future TEC systems must manage this balance of precision and recall for user trust.

### 5.3.2 Use Cases for TEC

Many reviewers highlighted scenarios where trustworthy TEC systems could be particularly useful.

Two reviewers said TEC is helpful for **corrections and typos**, similar to the use cases for GEC in the wild (Omelianchuk et al., 2021):

“If the tool would manage to reliably show missing punctuation marks, or numbers, or that the translation contains different numbers than the source, that would be helpful and save time.”; “recurring mistakes”

On the other hand, three reviewers mentioned that they hoped such a system would make more substantial corrections in order to save a non-negligible amount of time, although of course these edits may come at the expense of precision:

“There were not many suggestions, and they only offered small improvements...Not clear whether I would save time or not.”

Three reviewers commented that a TEC system could be a **memory aid or substitute for researching client-specific requirements**, which is often an intensive part of the production translation process. One reviewer pointed out it could be particularly useful as an *instructive* tool for translators who are new to a client:

“if I am new to an account and don’t yet know whether this client wants hyphens or not (always an issue with German). So usually I have to research... (or guess), but if the QA suggestions knew this client’s preference and would tell me, that would save me time.”

Finally, three reviewers commented that it could be useful as an **attention-directing tool** by making them aware of what errors they might look out for, especially in repetitive content where it may be easy to miss details:

“makes you more sensitive for spotting similar errors”; “makes you aware of what kind of errors to look for in upcoming segments”; “maybe it helps with [repetitive sentences] that you would otherwise just quickly glance at.”

## 6 Conclusion & Future Work

We introduced the task of translation error correction (TEC) and released the ACED corpus to study automatic correction of human translations, consisting of three TEC datasets across varying domains. In our analysis of TEC data, we showed how the errors that humans make differ from those made by MT systems, suggesting that this task warrants different approaches from those previously studied in the task of automatic post-editing. We confirm this empirically by proposing a synthetic data generation procedure that more closely matches the distribution of human translation errors and showing that our TEC model, pre-trained on this data, consistently outperforms models developed for APE, as well as those for MT and GEC. Finally, we showed how our TEC system is helpful to real humans, assisting professional reviewers and leading them to produce higher quality reviewed translations.

Future work may improve on our TEC system by investigating how to leverage the strengths of recent MT systems (e.g. for initializing systems or proposing edits) or developing more sophisticated synthetic data generation techniques (e.g. using the source sentence or linguistic knowledge). Beyond our benchmark, it would be interesting to apply TEC systems to other settings in which human translation errors appear, e.g., to correct translations written by language learners, denoise MT training sets, or clean up MT evaluation sets.

From the perspective of human-AI interaction, TEC presents a real-world use case and testbed to study how to assist experts with modern NLP systems, hinting at the opportunity to combine the best of humans and machines.

## Acknowledgments

We thank Sai Gouravajhala, Yunsu Kim, Eric Wallace, and the other members of the Lilt research team and Berkeley NLP group for helpful discussion and feedback. We thank Morgan Raymond and Spence Green for their support in releasing the dataset. Finally, we are grateful to the professional translators who annotated the dataset and participated in the user study.

## References

- Christopher Bryant, Mariano Felice, and Ted Briscoe. 2017. [Automatic annotation and evaluation of error types for grammatical error correction](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 793–805, Vancouver, Canada. Association for Computational Linguistics.
- Shamil Chollampatt, Raymond Hendy Susanto, Liling Tan, and Ewa Szymanska. 2020. [Can automatic post-editing improve NMT?](#) In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2736–2746, Online. Association for Computational Linguistics.
- Gonçalo M. Correia and André F. T. Martins. 2019. [A simple and effective approach to automatic post-editing with transfer learning](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, Florence, Italy. Association for Computational Linguistics.
- 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.
- Donald A. DePalma. 2021. [The language sector in eight charts](#). Accessed: 2022-05-01.
- Markus Freitag, George Foster, David Grangier, Viresh Ratnakar, Qijun Tan, and Wolfgang Macherey. 2021. [Experts, errors, and context: A large-scale study of human evaluation for machine translation](#).
- Prabhakar Gupta, Ridha Juneja, Anil Kumar Nelakanti, and Tamojit Chatterjee. 2021. Detecting over/under-translation errors for determining adequacy in human translations. *ArXiv*, abs/2104.00267.
- Gyde Hansen. 2009. [A classification of errors in translation and revision](#). In *CIUTI-Forum: Enhancing Translation Quality: Ways, Means, Methods*.
- Marcin Junczys-Dowmunt and Roman Grundkiewicz. 2016. [Log-linear combinations of monolingual and bilingual neural machine translation models for automatic post-editing](#). In *Proceedings of the First Conference on Machine Translation: Volume 2, Shared Task Papers*, pages 751–758, Berlin, Germany. Association for Computational Linguistics.
- Marcin Junczys-Dowmunt, Roman Grundkiewicz, Shubha Guha, and Kenneth Heafield. 2018. [Approaching neural grammatical error correction as a low-resource machine translation task](#). 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 595–606, New Orleans, Louisiana. Association for Computational Linguistics.
- Diederik P. Kingma and Jimmy Ba. 2015. [Adam: A method for stochastic optimization](#). In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Philipp Koehn, Hieu Hoang, Alexandra Birch, Chris Callison-Burch, Marcello Federico, Nicola Bertoldi, Brooke Cowan, Wade Shen, Christine Moran, Richard Zens, Chris Dyer, Ondřej Bojar, Alexandra Constantin, and Evan Herbst. 2007. [Moses: Open source toolkit for statistical machine translation](#). In *Proceedings of the 45th Annual Meeting of the Association for Computational Linguistics Companion Volume Proceedings of the Demo and Poster Sessions*, pages 177–180, Prague, Czech Republic. Association for Computational Linguistics.
- António V. Lopes, M. Amin Farajian, Gonçalo M. Correia, Jonay Trénous, and André F. T. Martins. 2019. [Unbabel’s submission to the WMT2019 APE shared task: BERT-based encoder-decoder for automatic post-editing](#). In *Proceedings of the Fourth Conference on Machine Translation (Volume 3: Shared Task Papers, Day 2)*, pages 118–123, Florence, Italy. Association for Computational Linguistics.
- Courtney Napoles, Keisuke Sakaguchi, Matt Post, and Joel Tetreault. 2015. [Ground truth for grammatical error correction metrics](#). In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 588–593, Beijing, China. Association for Computational Linguistics.
- Matteo Negri, Marco Turchi, Rajen Chatterjee, and Nicola Bertoldi. 2018. [ESCAPE: a large-scale synthetic corpus for automatic post-editing](#). In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).
- Kostiantyn Omelianchuk, Vipul Raheja, and Oleksandr Skurzhashnyi. 2021. [Text Simplification by Tagging](#). In *Proceedings of the 16th Workshop on Innovative Use of NLP for Building Educational Applications*, pages 11–25, Online. Association for Computational Linguistics.
- Keisuke Sakaguchi, Courtney Napoles, Matt Post, and Joel Tetreault. 2016. [Reassessing the goals of grammatical error correction: Fluency instead of grammaticality](#). *Transactions of the Association for Computational Linguistics*, 4:169–182.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. [Neural machine translation of rare words with subword units](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1715–1725, Berlin, Germany. Association for Computational Linguistics.

- Michel Simard, Nicola Ueffing, Pierre Isabelle, and Roland Kuhn. 2007. [Rule-based translation with statistical phrase-based post-editing](#). In *Proceedings of the Second Workshop on Statistical Machine Translation*, pages 203–206, Prague, Czech Republic. Association for Computational Linguistics.
- Lucia Specia and Kashif Shah. 2014. [Predicting human translation quality](#). In *Proceedings of the 11th Conference of the Association for Machine Translation in the Americas: MT Researchers Track*, pages 288–300, Vancouver, Canada. Association for Machine Translation in the Americas.
- 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.
- Yu Yuan and Serge Sharoff. 2020. [Sentence level human translation quality estimation with attention-based neural networks](#). In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 1858–1865, Marseille, France. European Language Resources Association.
- Thomas Zenkel, Joern Wuebker, and John DeNero. 2019. [Adding interpretable attention to neural translation models improves word alignment](#). *arXiv preprint arXiv:1901.11359*.
- Wei Zhao, Liang Wang, Kewei Shen, Ruoyu Jia, and Jingming Liu. 2019. [Improving grammatical error correction via pre-training a copy-augmented architecture with unlabeled data](#). 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 156–165, Minneapolis, Minnesota. Association for Computational Linguistics.

## A Transformer Architecture Background and Model Description

### A.1 Transformer Architecture

The neural models implemented in this work are based on the self-attentional Transformer architecture (Vaswani et al., 2017). Formally, given a sequence of source tokens (encoded as one-hot vectors)  $\mathbf{s}_{1\dots J} = (s_1, \dots, s_J)$ ,  $s_j \in \mathbf{V}$ , the goal is to predict a sequence of target tokens  $\mathbf{t}'_{1\dots I'} = (t'_1, \dots, t'_{I'})$ ,  $t'_i \in \mathbf{V}$ , that is a translation of the source sequence, where  $\mathbf{V}$  is the vocabulary. The model has two main components, the *encoder* and the *decoder*. The encoder transforms the source sequence  $\mathbf{s}_{1\dots J}$  into a sequence of hidden states by first mapping each individual token into a continuous embedding space, adding a positional embedding and then processing it through a sequence of self-attention and feed-forward layers:

$$\mathbf{x}_{1\dots J} = \mathbf{E}\mathbf{s}_{1\dots J} + \mathbf{p}_{1\dots J} \quad (1)$$

$$\mathbf{h}_{1\dots J}^{\text{enc}} = \text{encoder}(\mathbf{x}_{1\dots J}), \quad (2)$$

where  $x_j \in \mathbf{V}$ ,  $j \in (1, \dots, J)$ ,  $\mathbf{E}$  is the embedding matrix for vocabulary  $\mathbf{E}$  and  $\mathbf{p}_{1\dots J}$  is the sequence of positional embeddings described in Sec. 3.5 of (Vaswani et al., 2017). At a given time step  $i$ , the decoder defines a probability distribution  $P_i$  over all vocabulary items in  $\mathbf{V}$ :

$$\mathbf{y}'_{1\dots i-1} = \mathbf{E}\mathbf{t}'_{1\dots i-1} + \mathbf{p}_{1\dots i-1} \quad (3)$$

$$\mathbf{h}_i^{\text{dec}} = \text{decoder}(\mathbf{y}'_{1\dots i-1}, \mathbf{h}_{1\dots J}^{\text{enc}}) \quad (4)$$

$$P_i(t'_i) = \text{softmax}(\mathbf{h}_i^{\text{dec}}\mathbf{E}^\top) \quad (5)$$

where we assume a single shared vocabulary  $\mathbf{V}$  and embedding matrix  $\mathbf{E}$ . At training time we optimize the cross-entropy loss

$$\mathcal{L}_{\text{CE}}(P) = - \sum_i \log(P_i(t'_i)). \quad (6)$$

### A.2 Dual-Source Encoder-Decoder Model

Given an additional input sequence  $\mathbf{t}_{1\dots I} = (t_1, \dots, t_I)$ . the dual-source model used for the APE and TEC models is implemented by independently projecting  $\mathbf{t}_{1\dots I}$  into the embedding space, adding an offset vector  $\mathbf{o}$  and concatenating the embedding sequences. Equations 2 and 4 are rewritten as

$$\mathbf{y}_{1\dots I} = \mathbf{E}\mathbf{t}_{1\dots I} + \mathbf{p}_{1\dots I} + \mathbf{o} \quad (7)$$

$$\mathbf{h}_{1\dots(J+I)}^{\text{enc}} = \text{encoder}([\mathbf{x}_{1\dots J}; \mathbf{y}_{1\dots I}]) \quad (8)$$

$$\mathbf{h}_i^{\text{dec}} = \text{decoder}(\mathbf{y}'_{1\dots i-1}, \mathbf{h}_{1\dots(J+I)}^{\text{enc}}), \quad (9)$$

where  $\mathbf{o}$  is a single learned vector that is broadcast to all positions  $i \in (1, \dots, I)$  and  $[\cdot; \cdot]$  denotes the concatenate operation.

### A.3 Copy-Attention Mechanism

The new output probability distribution for the next target token  $P_i(t'_i)$  is a weighted sum of the probability of generating and the probability of copying token  $t'_i$ :

$$\hat{P}_i(t'_i) = (1 - \alpha_i^{\text{copy}})P_i(t'_i) + \alpha_i^{\text{copy}}P_i^{\text{copy}}(t'_i), \quad (10)$$

where the copy probabilities are calculated from the attention matrix of an additional encoder-decoder attention layer that is added on top of the final decoder layer,  $\mathbf{A}_i$ :

$$P_i^{\text{copy}}(t'_i) = \text{softmax}(\mathbf{A}_i) \quad (11)$$

The copy probability weight  $\alpha_i^{\text{copy}}$  is determined with the attention context vector  $\mathbf{c}_i$ , computed as a weighted sum of the attention values (i.e. linearly transformed encoder states) where the weights are defined by  $\mathbf{A}_i$ :

$$\alpha_i^{\text{copy}} = \text{sigmoid}(\mathbf{W}^\top \mathbf{c}_i). \quad (12)$$

This copy-attention layer applies a source-side mask so that it only attends to the positions  $(J + 1, \dots, J + I)$  that correspond to the second input sequence  $\mathbf{t}_{1\dots I}$ , and its implementation follows Zenkel et al. (2019). In particular, it uses a single attention head, no skip connection, and contains a separate output layer that predicts the target word based on its context vector with probability distribution  $P_i^{\text{align}}(\cdot)$ . At training time both output layers are optimized jointly by defining the overall loss  $\mathcal{L}$  as the weighted sum of both cross-entropy losses:

$$\mathcal{L} = \mathcal{L}_{\text{CE}}(\hat{P}) + \lambda \mathcal{L}_{\text{CE}}(P^{\text{align}}) \quad (13)$$

$\lambda$  is set to 0.05 in all experiments. We further apply source-word dropout (Junczys-Dowmunt et al., 2018), setting the full embedding vector for words in  $t$  to  $1/p_{\text{src}}$  with probability  $p_{\text{src}} = 0.05$ .

## B Results without finetuning

See Table 8.

## C Examples of System Output

See Table 9, Table 10, Table 11.

## D Full User Study Results

See Table 12.



Model	ASICS			EMERSON			DO		
	Prec.	Rec.	F <sub>0.5</sub>	Prec.	Rec.	F <sub>0.5</sub>	Prec.	Rec.	F <sub>0.5</sub>
MT	0.8	<b>12.2</b>	1.0	1.6	<b>16.9</b>	2.0	0.5	<b>23.4</b>	0.6
GEC	5.9	1.4	3.6	0.7	0.6	0.7	0.3	1.2	0.3
APE	2.7	7.3	3.1	3.6	3.3	3.5	0.4	5.0	0.5
BERT-APE	0.5	5.9	0.6	0.1	2.2	0.1	0.3	9.2	0.4
TEC	<b>41.7</b>	1.7	<b>7.5</b>	<b>12.2</b>	4.2	<b>8.8</b>	<b>5.3</b>	2.5	<b>4.3</b>

Table 8: Results without finetuning.

Type:	Monolingual: technical
<i>s</i> :	Do your feet roll inwards when running?
<i>t</i> :	Klppen deine FüÙe beim Laufen nach innen ?
Reference <i>t'</i> :	Kiippen deine FüÙe beim Laufen nach innen ?
Model	Predicted <i>t'</i>
MT	Klppen Rollen deine FüÙe beim Laufen nach innen ?
APE	Correctly predicts <i>t'</i>
BERT-APE	No change to <i>t</i>
GEC	Correctly predicts <i>t'</i>
TEC (ours)	Correctly predicts <i>t'</i>

Table 9: A monolingual technical error the APE, GEC and TEC models edit correctly.

Type:	Monolingual: technical
<i>s</i> :	Run further than you thought possible in the GEL-NIMBUS™ 21 LS running shoe for men by ASICS - packed full of plush cushioning and special technologies to respond to your natural stride.
<i>t</i> :	Mit mit dem ASICS GEL-NIMBUS™ 21 LS Laufschuh für Herren mit weicher Dämpfung und speziellen Technologien , durch die sich der Schuh an deine natürlichen Schritte anpasst , läufst du weiter , als du es selbst je für möglich gehalten hast .
Reference <i>t'</i> :	Mit dem ASICS GEL-NIMBUS™ 21 LS Laufschuh für Herren mit weicher Dämpfung und speziellen Technologien , durch die sich der Schuh an deine natürlichen Schritte anpasst , läufst du weiter , als du es selbst je für möglich gehalten hast .
Model	Predicted <i>t'</i>
MT	Mit mit dem Laufe weiter , als du es im GEL-NIMBUS™ 21 LS Laufschuh für Herren von ASICS für möglich gehalten hättest - mit weicher Dämpfung und speziellen Technologien , durch die sich der Schuh an deine natürlichen Schritte anpasst , läufst du weiter , als du es selbst je für möglich gehalten hast . um auf deinen natürlichen Schritt zu reagieren .
APE	No change to <i>t</i>
BERT-APE	Correctly predicts <i>t'</i>
GEC	Correctly predicts <i>t'</i>
TEC (ours)	Correctly predicts <i>t'</i>

Table 10: A monolingual technical error the BERT-APE, GEC and TEC models edit correctly.

Type:	Bilingual
<i>s</i> :	The DUOMAX™ midsole offers smooth overpronation control by combining two different density materials to reduce the risk of flat feet and bunions.
<i>t</i> :	Die DUOMAX™ -Mittelssole bietet mühelos Halt bei Überpronation , indem zwei unterschiedliche Dichtematerialien kombiniert werden , um das Risiko von flachen Füßen und Fußballen zu verringern .
Reference <i>t'</i> :	Die DUOMAX™ -Mittelssole bietet mühelos Halt bei Überpronation , indem zwei unterschiedliche Dichtematerialien kombiniert werden , um das Risiko von Plattfüßen und Ballenzehen zu verringern .
<b>Model</b>	
<b>MT</b>	Die DUOMAX™ -Mittelssole bietet mühelos Halt bei eine reibungslose Überpronation , indem sie zwei unterschiedliche Dichtematerialien kombiniert , um das Risiko von flachen Füßen und Fußballen Bündchen zu verringern .
<b>APE</b>	No change to <i>t</i>
<b>BERT-APE</b>	Die DUOMAX™ -Mittelssole bietet mühelos Halt bei Überpronation , indem zwei unterschiedliche Dichtematerialien kombiniert werden , um das Risiko von flachen Füßen und Fußballen Baseballen zu verringern .
<b>GEC</b>	No change to <i>t</i>
<b>TEC (ours)</b>	No change to <i>t</i>

Table 11: A bilingual error all models fail to edit correctly.

	Suggestion Hidden	Suggestion Shown	Suggestion Shown and Accepted	Suggestion Shown and Declined
Review Time (median)	34.0475 sec	33.103 sec	26.606 sec	50.524 sec
Review Time (length-norm, median)	361 ms/char	367 ms/char	328 ms/char	841 ms/char
Inserts (median)	1.5 chars	0 chars	0 chars	4 chars
Inserts (length-norm, median)	0.0248	0	0	0.0345
Deletes (median)	2 chars	0 chars	0 chars	4 chars
Deletes (length-norm, median)	0.0294	0	0	0.0375
Inserts+Deletes (median)	4 chars	0 chars	0 chars	9 chars
Inserts+Deletes (length-norm, median)	0.0625	0	0	0.0625
Levenshtein Dist (median)	2 chars	1 char	1 char	6 chars
Levenshtein Dist (length-norm, median)	0.0347	0.0185	0.0176	0.0426

Table 12: All data from our user study about review times, number of characters the user inserted and deleted, and final levenshtein distances from the original. Data shown are medians (raw and length-normalized) across the segments, based on whether the suggestion was hidden or shown. “Suggestion Shown” is further broken down according to whether the user accepted or declined the suggestion.