

# Assessing Post-editing Effort in the English-Hindi Direction

Arafat Ahsan, Vandan Mujadia and Dipti Misra Sharma

IIIT, Hyderabad, India

{arafat.a, vandan.mu}@research.iiit.ac.in

dipti@iiit.ac.in

## Abstract

We present findings from a first in-depth post-editing effort estimation study in the English-Hindi direction along multiple effort indicators. We conduct a controlled experiment involving professional translators, who complete assigned tasks alternately, in a translation from scratch and a post-edit condition. We find that post-editing reduces translation time (by 63%), utilizes fewer keystrokes (by 59%), and decreases the number of pauses (by 63%) when compared to translating from scratch. We further verify the quality of translations thus produced via a human evaluation task in which we do not detect any discernible quality differences.

## 1 Introduction

Translation workflows that are based on post-editing of Machine Translation output are being increasingly adopted in the industry (Gaspari et al., 2015). Gains that accrue from a post-editing based workflow, measured over multiple post-editing effort indicators, have been reported to be considerably significant by a number of previous studies over multiple language combinations (Plitt and Masselot, 2010; C. M. de Sousa et al., 2011; Green et al., 2013). But to extend post-editing beyond its current silos it is imperative to put new and less-studied language pairs under the lens to make a case for wider adoption via empirically backed evidence.<sup>1</sup>

Post-editing effort is often quantified across three different dimensions, each focusing in turn on a different aspect of post-editing behaviour (Krings, 2001). The dimensions studied are the following: *Temporal*– understood as the time taken to complete a translation task, often reported per segment or word; *Technical*– estimate of the physical labour

of the translation activity, measured in terms of keystrokes logged or edit operations performed; and *Cognitive*– an indirect estimate of the extent of cognitive processes underlying the translation task, inferred from keylogging pause or eye-tracking data as it is not possible to observe these directly (Moorkens et al., 2015).

If it can be shown that post-editing machine translation output is temporally efficient, technically less laborious, and cognitively less demanding, then it can be recommended as the default workflow for large translation jobs. But this first calls for a comparison between machine translation based post-editing behaviour (henceforth PE) and unaided human translation from scratch (henceforth HT). Thus, the research questions that we pose are the following:

- Is post-editing effort as measured on temporal, technical and cognitive dimensions *lesser* in the PE condition than the HT condition for the English-Hindi direction?
- Is the quality of post-edited segments *equal to* translated segments as ascertained by human raters?
- Do automatic MT evaluation metrics *correlate* with PE effort indicators, when both measured at the segment level?

Most of this paper will focus on answering the first question in some detail. We are equally interested in the other two as well, but will only be presenting some initial results from a first attempt at tackling them.

The rest of this paper is organized as follows: Section 2 discusses some past studies including those that have previously studied the English-Hindi PE direction. In Section 3 we detail our experimental setup. Section 4 presents our results

<sup>1</sup>Gaspari et al. (2015)'s survey reveals a heavy skew towards English and other European language combinations.

and analysis and in Section 5 we draw our conclusions and sketch the outlines of our future work.

## 2 Related Work

We now take a more detailed look at some of the past efforts towards contrasting the two settings. [Plitt and Masselot \(2010\)](#) compared HT and PE when translating from English into 4 European languages (French, Italian, German, and Spanish) and reported an overall productivity gain of 74% which converted into time savings of 43%. They also observed a 70% reduction in keyboard time and 31% in pause time for the PE setting. [C. M. de Sousa et al. \(2011\)](#) also report PE to be 40% faster than HT in the English-Portuguese direction when translating movie subtitles.

Other studies however ([Läubli et al., 2013](#)), have reported more modest gains, with estimated time savings of 15–20% when translating between a European language pair (German-French) within a *realistic* translation environment.<sup>2</sup> [Garcia \(2011\)](#) also finds only marginal productivity gains when studying the English-Chinese pair and additionally reports an impact of directionality when source and target languages are switched.

All of these earlier experiments, however, were based on the output of Phrase based Statistical Machine Translation (PBSMT) systems. With Neural Machine Translation (NMT) and its subsequent iterations being the current state of the art and outperforming PBSMT ([Bahdanau et al., 2014](#); [Vaswani et al., 2017](#); [Castilho et al., 2018](#)), this shift in technology paradigm from PBSMT to NMT must then be addressed in post-editing studies as well.

[Läubli et al. \(2019\)](#) conduct such a study, this time utilizing the output of an NMT system to compare PE with HT in the German-French and German-Italian translation directions.<sup>3</sup> They report significant overall productivity gains, but with marked differences between the pairs: 59.74% for the former and only 9.26% for the latter. Another interesting comparison of HT, PBSMT, and NMT post-editing settings performed on a literary text (chapter from a novel) reports an increase in productivity by 36% for the NMT based setting over HT ([Toral et al., 2018](#)).

<sup>2</sup>Experimental settings for these studies may deploy specialized interfaces for accurate measurements or make use of environments already familiar to professional translators.

<sup>3</sup>The HT condition is aided by a domain specific translation memory (TM).

We have seen in previous studies that throughputs vary considerably depending on the language pair under the lens ([Green et al., 2013](#); [Läubli et al., 2019](#)). We now discuss some earlier efforts that have included an Indian language in their experiment.

[Shah et al. \(2015\)](#) conducted an experiment where students post-edited parts of a specialized English language textbook on *bioelectromagnetism* into 7 languages, 3 of them being Indian languages including Hindi. They reported an increase in post-edit time by a factor of 3–5 when the target language was an Indian language. They put this down to greater terminological distance between English and Indian languages compared to other languages in their experiment. They do not study and compare against the HT condition, or report on technical or cognitive indicators.

[Carl et al. \(2016\)](#) also report results on Hindi (amongst 6 languages) comparing the HT and PE conditions. Their English-Hindi results are based on an existing multilingual translation database that contains experimental data around translators' activities in both conditions. They find in favour of the PE condition across all languages when measured on temporal indicators, but report translating into Hindi to be the slowest amongst the 6. They do not quantify average throughput gain or time savings.

[Meetei et al. \(2020\)](#) compare PE behaviour when translating from English into 3 Indian languages (Manipuri, Mizo and Hindi). They conduct light post-editing and report Hindi to be the fastest to post-edit amongst the three languages.<sup>4</sup> They ascribe it to the availability of relatively mature MT systems in the English-Hindi direction compared to Mizo and Manipuri, which are low-resource languages. They use student volunteers and do not investigate cognitive indicators.

[Ahmad et al. \(2018\)](#) present an industry perspective and claim a 2–3 fold increase in productivity for users using their tools in combination with MT. However, they base this on longitudinal tracking of their users.

While all these past studies have certainly helped in providing insights into post-editing behaviour in Indian languages in general and in Hindi in particular, we sense a need for a more in-depth look along all PE effort indicators within at least one

<sup>4</sup>It is also often referred to as *good enough* translations and is lower than publishable quality translations.



Figure 1: Shows a snapshot of the translator workbench used in the study. Depicts an example in the PE condition pre-filled with an MT proposal; for the HT condition the text area is left blank.

Indian-language setting. With our current work, we seek to address this gap.

We see our contribution differentiated from previous related work as follows: *(i)* we present in-depth results from all three primary indicators of PE effort (temporal, technical and cognitive) for the English-Hindi direction; *(ii)* we account for per translator and per item variation with the use of mixed-effects models; *(iii)* we utilize professional translators in order to accurately gauge the impact of contrasting conditions; *(iv)* we conduct a human quality-rating exercise comparing target text produced in both conditions; *(v)* we present correlations of automatic MT metrics with PE effort indicators.

### 3 Experimental Setup

We conducted this experiment under a 2 (translation conditions)  $\times$  200 (source segments) mixed design. All subjects saw both factor levels (HT and PE), but only one combination for each level as having been exposed to a source segment in one condition would have affected their translation in the other. The study was conducted online over 5 consecutive days with 2 sessions per day. The sessions were not time bound. Subjects translated 2 files of 10 segments each in alternating conditions in each session.

All subjects participated in a warm-up translation exercise a week prior to the start of the actual task. This was done to establish familiarity with the interface used in the study. We chose to adapt an existing beta version of a web-based translation workbench by adding extensive keystroke logging features along with some other minor tweaks.<sup>5</sup> The UI itself was kept clean and uncluttered, serving one segment at a time to the translators. This

<sup>5</sup><https://indictranslate.in/>

meant that while translators had previous context of the text under translation, they could not navigate ahead for context. A timer was displayed once a translator navigated to each new segment. This was meant to prompt the translator to focus on the activity at hand. Figure 1 shows the workbench interface as seen by translators under the PE condition.

We instructed the participants to aim for publishable translation quality. They were free to conduct web searches and consult online or offline dictionaries, but were discouraged from spending too much time doing so.<sup>6</sup> It was deemed acceptable to transliterate any technical terms or terminology into Hindi if they could not find its translation even with the aid of resources available to them. However, they were strictly prohibited from consulting any online MT engines during the task. Subjects were encouraged to complete each task (consisting of 10 segments) in a single sitting without a break.

Previous studies, such as those discussed earlier, have noted the impact of a number of different variables (language pairs, MT paradigms, text domains, translation environments, translator competencies) on translation throughputs. This calls for not only careful experiment design, but also utilization of techniques that can help with the testing and inference of results. Green et al. (2013) utilized one of the first such designs for post-editing productivity studies and deployed mixed-effects models (Baayen et al., 2008) to account for inter-language, inter-subject, and inter-item variability.

Mixed-effects models are able to model this variability in two ways: *(i)* through random intercepts, that can account for the differences between translators seen in their differing throughputs (or differences between linguistic items due to the features inherent to them); *(ii)* and through a random slope that accounts for how different subjects may experience the change of condition differently. Accounting for these variabilities allows us to isolate the effect of condition, generalize our findings beyond our sample, and avoid the "language-as-fixed-effect fallacy" (Clark, 1973).

In fitting our mixed-effects models we follow a methodology similar to the one described by Baayen et al. (2008) and followed by Green et al. (2013) and later Toral et al. (2018). Maximal models were fit when possible (Barr et al., 2013); in

<sup>6</sup>This was done as technical terminology related difficulties have previously been noted for this language direction (Shah et al., 2015).

Day	S1-T1	S1-T2	S2-T1	S2-T2
Day 1	20.20 (A1)	26.30 (A1)	26.80 (A1)	21.30 (A2)
Day 2	20.90 (A2)	22.30 (A2)	24.90 (A2)	23.20 (A3)
Day 3	26.40 (A4)	24.60 (A4)	26.30 (A4)	22.40 (A4)
Day 4	22.40 (A4)	20.30 (A4)	18.00 (A5)	18.70 (A5)
Day 5	22.50 (A5)	16.00 (A5)	20.30 (A5)	23.00 (A5)

Table 1: Average sentence lengths (in words) per session-task block as presented to translators. Also shown in parenthesis are the source articles used for each block.

case of convergence failure, a less complex model was fit by successively removing the random slopes of the by-subject and by-segment random effects component. Models thus obtained were compared via likelihood ratio tests. We also refit our final models after filtering data points with residuals deviating more than 2.5 standard deviations. This helps check for the influence of any atypical outliers (Baayen et al., 2008). We verify the residual plots for normality and homoscedasticity. We utilized the `lme4` package in R (Bates et al., 2015) for all mixed-effects models related analyses.

### 3.1 Data

We assembled a corpus of recent English language news articles from two distinct online sources. The choice of news as a domain was motivated by observations of terminology-related difficulties in more specialized domains, as reported by earlier studies (Shah et al., 2015). Each news article was segmented into sentences using the NLTK library and divided into *blocks* of 10 segments.<sup>7</sup> Only those blocks were used that fell within a  $MEAN \pm SD$  of the corpus mean (Table 1). We prioritized the continuity of a news article across blocks when making *block* selections.<sup>8</sup> This methodology yielded a total of 200 unique source segments divided into 20 blocks of 10 segments each, spanning 5 different news articles: A1–A5. Conditions were counter-balanced to handle order effects.

### 3.2 Participants

The participants of our study are self-declared professional translators. We contacted a professional translation service provider to help assemble the pool.<sup>9</sup> A short questionnaire accompanied the registration form for the task. Of our participant pool of 10 subjects, 70% reported 2–5 years of experience translating in the English-Hindi direction,

<sup>7</sup><https://www.nltk.org/api/nltk.tokenize.html>

<sup>8</sup>In 2 cases out of 20 we had to skip the subsequent block, owing to short average sentence lengths of the blocks.

<sup>9</sup><http://www.ebhashasetu.com/>

while 30% reported 0–2 years of experience. The same percentage breakdown was observed for a question related to previous post-editing experience. All subjects were paid the going market rates for the task regardless of the condition (PE, HT).

### 3.3 MT System

The English→Hindi MT engine used for the task is a transformer based neural machine translation system. This subword-based NMT system is trained on cleaned WAT 2021<sup>10</sup> English-Hindi training corpus using the Opennmt-py toolkit (Klein et al., 2020). The system also utilizes forward and backward translations on the IndicCorp monolingual corpus to obtain synthetic data for training.<sup>11</sup> It uses subwords as the basic translation unit with 20,000 merge operations on both source and target languages. The system obtained a BLEU score of 35.46 on cleaned WAT 2021 English-Hindi test data.

## 4 Results and Discussion

### 4.1 Pre-processing

Once we processed the activity logs for all 10 subjects across all 200 segments, they yielded 2000 unique observations. We found that 7 of these items (all from the HT condition) did not contain a final translation, so we discarded those. We think that in these cases the subjects may have accidentally navigated to the next segment without having completed a translation. In the PE condition we found that one subject *POI* had not touched 68% of the MT outputs and had accepted them without modifications. This was almost 3 times the next highest proportion we detected across all other subjects. We decided to remove all data points generated by this subject. We were thus left with 1793 observations on which we base these results.

We calculated time per segment, source segment lengths (in words and characters), number of keystrokes (total, as well as those belonging to different categories: content, navigation and deletion), average pause duration, initial pause duration, and number of pauses. We also computed H-BLEU (Papineni et al., 2002), H-TER (Snover et al., 2006) and H-chrF (Popović, 2015) metrics on the post-edited segments.<sup>12</sup>

<sup>10</sup><http://lotus.kuee.kyoto-u.ac.jp/WAT/indic-multilingual/>

<sup>11</sup><https://indicnlp.ai4bharat.org/corpora/>

<sup>12</sup>H signifies that scores were computed using the reference generated in the PE condition by the same subject.






Figure 2: Individual translation throughputs in words per hour and average throughput in each contrasting condition.

Having assembled this data, we set out to answer the research questions posed earlier in Section 2.

## 4.2 Temporal Effort

We first present a view of temporal effort in terms of productivity measured as words per hour. We see productivity improvements in the PE condition across the board except for subject *P07*. Overall, this translates into a throughput increase from 359 words/hour to 979 words/hour. We thus observe an overall productivity gain of 172%, which amounts to 63% in time savings.

This is more than twice the 74% gain reported by (Plitt and Masselot, 2010) when studying European-language pairs and the 59.74% reported by (Läubli et al., 2019) recently. But we note that in the first case, the experiments were conducted on PBSMT outputs, and in the second, while NMT was used, the control condition was aided by a TM, thus pushing up the baseline throughputs. With this context in mind, the average productivity gain seen in our study does not appear to be unrealistic.

Figure 2 shows a comparison of individual throughputs in contrasting task conditions along with means aggregated for the two conditions. We also note a great variation in productivity gain amongst subjects ranging from -7% to 410%.

Table 2 helps interpret this further. We contrast the number of unedited and edited MT proposals per subject and their individual productivity gain percentages. It follows that higher the acceptance of MT proposals without modifications by a subject, greater the gain in individual productivity. While this may point to high quality MT output, it also demands a closer scrutiny of the quality of translations generated in each condition. We address this in Section 4.5.

Participant	Unedited (%)	Edited (%)	Productivity Gain (%)
P02	20	80	411
P03	3	97	73
P04	3	97	92
P05	18	82	275
P06	25	75	404
P07	0	100	-7
P08	2	98	105
P09	25	75	234
P10	16	84	169

Table 2: MT segments accepted without modifications and with modifications per subject along with individual productivity gain percentages.

We now report the mixed-effects regression results. Plotting temporal data showed a right-skewed distribution, so we log transform all time data before proceeding further. As our goal is to predict translation time and establish the significance of conditions, we fit a linear mixed-effects regression model with two fixed-effect predictors (condition and segment length) and two random-effect predictors (subjects and segments), where on the subject predictor we also include a random-slope for task condition.

In the final model, we observe a significant main effect for both segment length as well as translation condition.<sup>13</sup> Temporal effort significantly increases with segment length, but decreases for the PE condition. Table 3 shows the significance levels and direction for each predictor in our final models across all PE effort dimensions that we study.

## 4.3 Technical Effort

We measure technical effort as the number of keystrokes used to generate the target text. We normalize it per source segment character. Figure 3 shows 1.33 keystrokes used per source character in the HT condition and 0.54 keystrokes in the PE condition, amounting to an effort reduction of 59%. Contrast this with the 23% reduction reported by Toral et al. (2018) when post-editing a literary text. Again, except for subject *P07* all participants show reduced effort in the PE condition.

We also classified each keystroke based on the type of the keystroke logged. We classify these into content, navigation, and deletion categories and report the percentage breakdown of the total into these categories in Table 4. We observe higher navigation and deletion operations in the PE condition (28% and 26%) than the HT condition (8% and 14%), while content keystrokes register a higher percentage in HT (77%) compared to PE (46%).

<sup>13</sup>We utilize the *lmerTest* package that extends results with *p*-values for models built with *lme4*.

Predictor	Temporal	Technical	Cognitive		
			number	average duration	initial duration
<i>Segment length</i>	↑***	↑***	↑***	↑***	↑***
<i>Condition (PE vs. HT)</i>	↓***	↓***	↓***	—	—

Significance levels: — ( $p > 0.1$ ), ( $p < 0.1$ ), \* ( $p < 0.05$ ), \*\* ( $p < 0.01$ ), \*\*\* ( $p < 0.001$ ).

Direction: (↑ ↓) arrows depict whether the predictor has a negative or positive correlation with the dependent variable.

Table 3: Significance levels of predictors in our final models across all modeled PE effort dimensions.




Figure 3: Technical effort estimated as number of keystrokes needed to generate target text per source character.

Subject *P08* is an interesting case as they register the highest number of delete operations (they have high navigation numbers too) in either condition amongst all participants. This could point to frequent revisions made on the text.

As number of keystrokes is expressed as counts, we fit a Poisson generalized linear mixed-effects model to predict technical effort. We follow the same methodology as described in the previous section. We again find a significant main effect both for segment length as well as translation condition (Table 3), similar to what we saw for the temporal dimension earlier. Technical effort increases with increase in segment length, and decreases for the change in condition to PE.

#### 4.4 Cognitive Effort

Post-editing effort estimation studies based on eye-tracking data use *fixations* as a proxy to estimate cognitive load; the idea being that greater the number and duration of fixations, greater the cognitive load (O’Brien, 2011). In the absence of eye-tracking data, the use of *pauses* as a proxy for cognitive load is also well established (O’Brien, 2006). We report on three such cognitive indicators: number of pauses, pause duration, and initial pause duration. Findings related to the first two




Figure 4: Cognitive effort estimated as average number of pauses per source segment.

have been reported in previous post-editing literature (Green et al., 2013; Toral et al., 2018).<sup>14</sup> The third (initial pause duration), we introduce in order to gauge differences in reaction times from when a subject first navigates to a new segment displayed in either condition to their first action on it.

We calculate the time difference between two subsequent key events and consider all observations above *1000ms* to be pauses following (O’Brien, 2006; Koehn, 2009).

Figure 4 shows the differences in the frequency of pauses for each subject in the two conditions. We notice a reduction of 63% in the PE condition from 31 pauses per segment in HT to 12 pauses per segment in PE. This points to a much reduced cognitive load when post-editing.

However, a similar exercise on pause duration data reveals an increase of approximately 12% in the PE condition compared to the HT condition (Figure 5). Although, it is not significant, this is in line with findings reported previously comparing these two specific cognitive indicators (Green et al., 2013).

We finally compare initial pause duration between PE and HT. We expect this initial load to be higher for the PE condition given that there are two

<sup>14</sup>There is also an indicator reported as pause ratio which we eschew in favour of initial pause time.

Participant	HT (%)			PE (%)		
	Content	Navigation	Deletion	Content	Navigation	Deletion
P02	81	9	11	41	32	27
P03	94	1	5	58	8	34
P04	86	9	5	48	41	11
P05	85	5	10	56	30	14
P06	90	1	9	54	28	18
P07	66	11	23	56	12	32
P08	24	25	51	10	27	63
P09	78	12	10	35	52	13
P10	94	1	5	59	19	22
Mean±SD	77.55 ± 21.80	8.15 ± 7.89	14.30 ± 14.63	46.52 ± 15.91	27.69 ± 13.65	25.8 ± 16.1

Table 4: Types of keystrokes generated by subjects in each condition.




Figure 5: Cognitive effort estimated as average pause duration per source segment.




Figure 6: Cognitive effort estimated as average initial pause duration per source segment.

segments displayed to the subject in this condition: the source segment and the MT proposal, which have to be read and comprehended before starting the post-editing activity. This seems to hold, but not significantly, as we see only a small increase of about 5% for the PE condition (Figure 6). One explanation could be that in the PE condition, the MT proposal in spite of registering a higher cognitive load initially also later acts as a helpful prompt for the subject. An eye-tracking based experiment might prove useful in teasing apart these two opposite effects.

When comparing the means<sup>15</sup> for pause duration and initial pause duration we find pause duration (6.77s for HT and 7.71s for PE) to be considerably lower than initial pause duration (35.02s for HT and 36.67s for PE). The translator therefore, takes a longer initial pause to start formulating a response, but once they start the activity, they take considerably shorter pauses.

We go on to fit three mixed-effects models to validate these findings. A Poisson generalized mixed-effects model to estimate pause counts finds signifi-

<sup>15</sup>After transforming back from log scale. Also, note that pause duration does not include initial pause duration as a component, and is the duration of pauses after post-editing starts.

cant main effects for segment length and condition. Cognitive effort (measured as count of pauses) increases with segment length and decreases for the change in condition to PE.

The other two linear mixed-effects models fit to predict average pause duration, and initial pause duration find a significant effect only for segment length and not for condition (Table 3). This shows that while cognitive effort certainly increases with segment length, the change in condition to PE, does not have a discernible effect on cognitive effort, when measured by the average and initial time duration indicators.

#### 4.5 Quality Judgements

To evaluate whether the quality of texts created in the PE condition matched those created in the HT condition, we conducted a human judgement based pairwise ranking task (Callison-Burch et al., 2011) on a small sample. We randomly sampled 3 target segments per condition from each subject. For each target text thus obtained, we paired it with another random sample after constraining on condition. We thus obtained 60 HT-PE pairs for evaluation. As we discovered issues (discussed earlier in Section 4.1) with subject P01's data after the quality eval-

Evaluator	PE vs. HT		
	Win	Loss	Tie
<i>E1</i>	18	14	15
<i>E2</i>	27	15	5
<i>E3</i>	7	7	33
<i>E4</i>	14	11	22
<i>E5</i>	13	11	23

Table 5: Pairwise quality judgements on sampled target texts reported as win, loss, and ties for PE against HT.

uation exercise had already been completed, we removed any pairs that had a segment translated by the subject. We report our results on this filtered set that consists of 47 pairs. Five evaluators were asked to judge each pair. Ties were allowed.

Table 5 shows the judgements from evaluators represented as win-loss statistics on the PE condition. We notice a high number of ties and a slight preference for the PE condition. However, the preference does not test to be significant on a sign test ignoring ties ( $p$ -value = 0.08). We conclude that translating in either condition produces similar quality target segments. However, we realise that the sample size was quite small compared to the number of possible combinations across all participants. We hope to conduct a more thorough review of quality in future.

#### 4.6 Automatic Quality Metrics

Finally, we investigate the correlations of some popular automatic MT evaluation metrics with the post-editing effort indicators reported so far in this study. We generated metric scores by comparing the MT proposal against its post-edited reference. We calculate scores for H-(BLEU, TER, and chrF).

Table 6 shows moderate correlations for all 3 MT metrics on the temporal indicator, similar to what [Tatsumi \(2009\)](#) also reported for this indicator. Correlations then get stronger on the technical indicator and then fade for the cognitive indicator.

We believe this may be because cognitive effort is the only one out of the three PE dimensions we studied that is not directly observed (instead, inferred from pause frequency and pause duration data), whereas the technical and temporal indicators can be measured more directly. This is similar to findings previously reported by [Moorkens et al. \(2015\)](#) who note a similar correlation trend across the three PE effort dimensions. The technical effort indicator appears to be the one most strongly correlated with automatic metrics.

The other two cognitive indicators (average and

MT Metric	PE Indicator		
	temporal	technical	cognitive (# pauses)
$(H)BLEU$	$r = -.56, ***$	$r = -.71, ***$	$r = -.48, ***$
$(H)TER$	$r = +.54, ***$	$r = +.70, ***$	$r = +.49, ***$
$(H)chrF$	$r = -.56, ***$	$r = -.73, ***$	$r = -.49, ***$

Note: All coefficients for r(898). For TER lower is better hence the positive correlation. The other two cognitive indicators (average and initial pause duration) did not show any correlation with any of the metrics – coefficients were close to 0.

Table 6: Correlations of PE Effort indicators with automatic MT metrics.

initial pause duration), which did not test significant as per our mixed-effects models, also do not show any correlation with any of the MT metrics – coefficients obtained were close to 0. We omit reporting them in Table 6 due to space constraints.

## 5 Conclusion

We conducted a post-editing effort assessment study and presented detailed analysis of effort indicators along the temporal, technical and cognitive dimensions. We observed that in the temporal dimension, post-editing reduced translation time by 63%; in the technical dimension it reduced number of key strokes by 59%; and in the cognitive dimension, it reduced the frequency of pauses by 63%. However, it increased average pause duration by 12% and average initial pause duration by 5%.

We then compared the quality of translations generated in each condition and found them to be similar.

And finally, we detected moderate to strong correlations for 3 automatic MT evaluation metrics across all PE effort indicators, with technical effort most strongly correlating with automatic MT metrics.

The last two observations regarding human quality judgement, and MT metrics and their correlations demand a closer look, which was not possible owing to time and space constraints. We expect to undertake this as part of our future work. We also propose to extend this study by including a third condition in future, either as an additional MT engine to check if MT quality differences show up in PE effort indicators ([Toral et al., 2018](#)), or by the use of translation aids (TM) to gauge their impact in a similar manner ([Läubli et al., 2013, 2019](#)).

We also intend to study other language pairs, especially those within the multilingual Indian context.



## References

- Rashid Ahmad, Priyank Gupta, Nagaraju Vuppala, Sanket Kumar Pathak, Ashutosh Kumar, Gagan Soni, Sravan Kumar, Manish Shrivastava, Avinash K Singh, Arbind K Gangwar, et al. 2018. Transzaar: Empowers human translators. In *2018 18th International Conference on Computational Science and Applications (ICCSA)*, pages 1–8. IEEE.
- R Harald Baayen, Douglas J Davidson, and Douglas M Bates. 2008. Mixed-effects modeling with crossed random effects for subjects and items. *Journal of memory and language*, 59(4):390–412.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*.
- Dale J Barr, Roger Levy, Christoph Scheepers, and Harry J Tily. 2013. Random effects structure for confirmatory hypothesis testing: Keep it maximal. *Journal of memory and language*, 68(3):255–278.
- Douglas Bates, Martin Machler, Ben Bolker, and Steve Walker. 2015. Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*, 67(1):1–48.
- Chris Callison-Burch, Philipp Koehn, Christof Monz, and Omar Zaidan. 2011. Findings of the 2011 workshop on statistical machine translation. In *Proceedings of the Sixth Workshop on Statistical Machine Translation*, pages 22–64.
- Michael Carl, Akiko Aizawa, and Masaru Yamada. 2016. English-to-japanese translation vs. dictation vs. post-editing: Comparing translation modes in a multilingual setting. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC’16)*, pages 4024–4031.
- Sheila Castilho, Joss Moorkens, Federico Gaspari, Rico Sennrich, Andy Way, and Panayota Georgakopoulou. 2018. Evaluating mt for massive open online courses. *Machine translation*, 32(3):255–278.
- Herbert H Clark. 1973. The language-as-fixed-effect fallacy: A critique of language statistics in psychological research. *Journal of verbal learning and verbal behavior*, 12(4):335–359.
- Ignacio Garcia. 2011. Translating by post-editing: is it the way forward? *Machine Translation*, 25(3):217–237.
- Federico Gaspari, Hala Almaghout, and Stephen Doherty. 2015. A survey of machine translation competences: Insights for translation technology educators and practitioners. *Perspectives*, 23(3):333–358.
- Spence Green, Jeffrey Heer, and Christopher D Manning. 2013. The efficacy of human post-editing for language translation. In *Proceedings of the SIGCHI conference on human factors in computing systems*, pages 439–448.
- Guillaume Klein, François Hernandez, Vincent Nguyen, and Jean Senellart. 2020. The opennmt neural machine translation toolkit: 2020 edition. In *Proceedings of the 14th Conference of the Association for Machine Translation in the Americas (AMTA 2020)*, pages 102–109.
- Philipp Koehn. 2009. A process study of computer-aided translation. *Machine Translation*, 23(4):241–263.
- Hans P Krings. 2001. *Repairing texts: Empirical investigations of machine translation post-editing processes*, volume 5. Kent State University Press.
- Samuel Lüubli, Chantal Amrhein, Patrick Düggelin, Beatriz Gonzalez, Alena Zwahlen, and Martin Volk. 2019. Post-editing productivity with neural machine translation: an empirical assessment of speed and quality in the banking and finance domain. *arXiv preprint arXiv:1906.01685*.
- Samuel Lüubli, Mark Fishel, Gary Massey, Maureen Ehrensberger-Dow, Martin Volk, Sharon O’Brien, Michel Simard, and Lucia Specia. 2013. Assessing post-editing efficiency in a realistic translation environment.
- Loitongbam Sanayai Meetei, Thoudam Doren Singh, Sivaji Bandyopadhyay, Mihaela Vela, and Josef van Genabith. 2020. English to manipuri and mizo post-editing effort and its impact on low resource machine translation. In *Proceedings of the 17th International Conference on Natural Language Processing (ICON)*, pages 50–59.
- Joss Moorkens, Sharon O’Brien, Igor AL Da Silva, Norma B de Lima Fonseca, and Fabio Alves. 2015. Correlations of perceived post-editing effort with measurements of actual effort. *Machine Translation*, 29(3-4):267–284.
- Sharon O’Brien. 2006. Pauses as indicators of cognitive effort in post-editing machine translation output. *Across Languages and Cultures*, 7(1):1–21.
- Sharon O’Brien. 2011. Towards predicting post-editing productivity. *Machine translation*, 25(3):197–215.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, pages 311–318.
- Mirko Plitt and François Masselot. 2010. A productivity test of statistical machine translation post-editing in a typical localisation context. *The Prague bulletin of mathematical linguistics*, 93(1):7–16.

- Maja Popović. 2015. chrf: character n-gram f-score for automatic mt evaluation. In *Proceedings of the Tenth Workshop on Statistical Machine Translation*, pages 392–395.
- Ritesh Shah, Christian Boitet, Pushpak Bhattacharyya, Mithun Padmakumar, Leonardo Zilio, Ruslan Kalitvianski, Mohammad Nasiruddin, Mutsuko Tomokiyo, and Sandra Milena Castellanos Pérez. 2015. Post-editing a chapter of a specialized textbook into 7 languages: importance of terminological proximity with english for productivity. In *Proceedings of the 12th International Conference on Natural Language Processing*, pages 325–332.
- Matthew Snover, Bonnie Dorr, Richard Schwartz, Linnea Micciulla, and John Makhoul. 2006. A study of translation edit rate with targeted human annotation. In *Proceedings of the 7th Conference of the Association for Machine Translation in the Americas: Technical Papers*, pages 223–231.
- Sheila C. M. de Sousa, Wilker Aziz, and Lucia Specia. 2011. [Assessing the post-editing effort for automatic and semi-automatic translations of DVD subtitles](#). In *Proceedings of the International Conference Recent Advances in Natural Language Processing 2011*, pages 97–103, Hissar, Bulgaria. Association for Computational Linguistics.
- Midori Tatsumi. 2009. Correlation between automatic evaluation metric scores, post-editing speed, and some other factors. *The Twelfth Machine Translation Summit (MT-Summit XII)*, pages 332–339.
- Antonio Toral, Martijn Wieling, and Andy Way. 2018. Post-editing effort of a novel with statistical and neural machine translation. *Frontiers in Digital Humanities*, 5:9.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008.