Measuring the Effects of Human and Machine Translation on Website Engagement

Geza Kovacs geza@alumni.stanford.edu
John DeNero john@lilt.com
Lilt Inc, San Francisco, CA, USA

Abstract

With the internet growing increasingly multilingual, it is important to consider translating websites. However, professional translators are much more expensive than machines, and machine translation quality is continually increasing, so we must justify the cost of professional translation by measuring the effects of translation on website engagement, and how users interact with translations. This paper presents an in-the-wild study run on 2 websites fully translated into 15 and 11 languages respectively, where visitors with non-English preferred languages were randomized into being shown text translated by a professional translator, machine translated text, or untranslated English text. We find that both human and machine translations improve engagement, users rarely switch the page language manually, and that in-browser machine translation is often used when English is shown, particularly by users from countries with low English proficiency. We also release a dataset of interaction data collected during our studies, including 3,332,669 sessions from 190 countries across 2 websites.

1 Introduction

The userbase of the internet is becoming increasingly linguistically diverse (Group, 2020). As a result, publishers increasingly need to translate websites to make content available to global audiences. Professional translators can be expensive, so localization decisions involve many tradeoffs. Should we show translations generated by professional translators, or is machine translation quality sufficiently high that showing machine translations will not negatively impact user engagement? If a human translation is not available in the user's preferred language, is it preferable to show English or a machine translation? Does the website need to be translated at all, given that browsers have built-in machine translation functionality? To make informed decisions, we need to know the effects of each form of translation on website engagement.

This paper quantifies the effects of translation on engagement metrics and how users interact with translations, by running large-scale (over 3 million sessions), in-the-wild A/B tests on the homepages of two different open-source software projects. Our contributions in this paper are a set of studies quantifying how users interact with website-provided translations and in-browser machine translations, interface recommendations on how to incorporate translations into websites based on these findings, and a public dataset of interactions and code that can be used to reproduce our results and conduct follow-up research.

2 Related Work

Previously published work has not evaluated the effects of machine translation, human translation, and non-translation on in-the-wild website user engagement in a randomized, A/B test fashion.

While machine translation quality has historically been negatively perceived (Läubli and Orrego-Carmona, 2017), some machine translation systems claim to have reached parity in translation quality with professional translators in certain settings (Hassan et al., 2018; Barrault et al., 2019; Popel et al., 2020). Other work has questioned these claims of human parity (Läubli et al., 2018; Toral et al., 2018; Toral, 2020). Website translation has been found to benefit search engine optimization by attracting users who search in their native language (Cappelli, 2007). Machine translation of product listings has been found to help increase purchases on eBay (Brynjolfsson et al., 2019). This work seeks to measure the effects of human and machine translation on user engagement in the context of software-centric websites.

3 Research Questions

The studies in this paper aim to answer the following research questions:

- Does showing a human translation result in better website engagement than showing a machine translation or an untranslated page?
- Should we automatically show machine translations based on browser language preferences, or show English by default and let users view a translation by clicking a button?
- When users are shown untranslated English pages, do they end up using their browser's built-in machine translation system?

4 Methodology

We ran A/B tests on two sites, both of which are open-source software projects with a single-page design, shown in Figure 1. Both sites had been translated from English to several languages[1] by professional translators two years before we began running this experiment. We will call these *supported languages* for each site. We refer to the preferred language in the user's browser settings as the *preferred language*. Site 1 had 3,298,635 sessions total, of which 1,233,841 (37.4%) had a supported non-English preferred language. Site 2 had 34,034 sessions total, of which 9,316 (27.4%) had a supported non-English preferred language. Data was gathered from Nov 25, 2020 to Jan 14, 2022 (415 days). We obtained machine translated text from Google Translate in Nov 2020. See the Appendix for the demographics of visitors to the websites.

4.1 Experiment Conditions

Each session with a non-English preferred language for which translations are available are randomized into one of five conditions:

- *UE (Untranslated English)*: Shows the page in English only; user cannot switch languages.
- *HTT (Human Translation, show Translation by default)*: Shows a human translation to the user's preferred language by default. Users can switch to English or to a human translation in any supported language via the language switcher. See right side of Figure 1.
- *HTE (Human Translation, show English by default)*: Shows English by default. Users can switch to a human translation in their preferred language or any supported language via the language switcher. See left side of Figure 1.
- *MTT (Machine Translation, show Translation by default)*: Shows a machine translation to the user's preferred language by default. Users can switch to English or any supported language via the language switcher.
- *MTE (Machine Translation, show English by default)*: Shows English by default. Users can

[1] Site 1 is translated into 15 languages: Chinese (Simplified), Chinese (Traditional), Danish, Dutch, English, French, German, Greek, Hebrew, Hungarian, Italian, Malaysian, Portuguese, Spanish, and Turkish. Site 2 is translated into 11 languages: Chinese (Simplified), Chinese (Traditional), Czech, Dutch, English, French, German, Greek, Portuguese, Spanish, and Turkish.

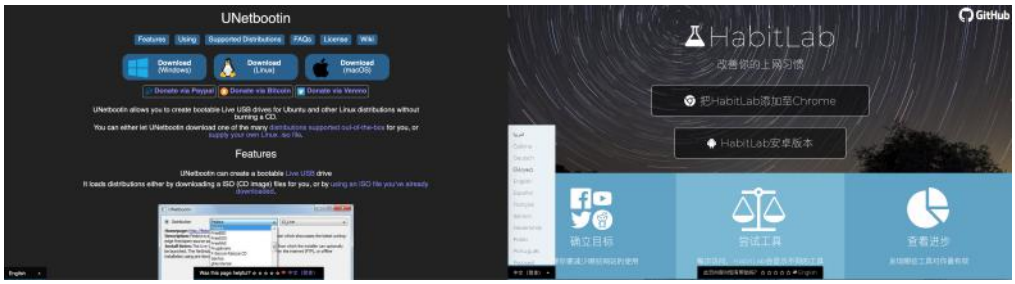

Figure 1: Sites we ran our study on; left is Site 1 (`https://unetbootin.github.io`), right is Site 2 (`https://habitlab.github.io`). On the bottom-left of each screenshot is the language-selection widget from the Transifex Live localization toolkit; it shows the current language by default (left); if clicked it allows the user to select a language from a list (right). The switcher in the middle includes a button to switch between English and the user's preferred language with a single click; the preferred language is Simplified Chinese in these screenshots.

Site 1 Engagement Metrics	HTT vs UE	HTT vs MTT	MTT vs UE
Download link clicked	$\chi^2 = 14.25, p = 0.0002$	$\chi^2 = 4.190, p = 0.041$	$\chi^2 = 2.968, p = 0.085$
Non-download link clicked	$\chi^2 = 191.9, p < 0.0001$	$\chi^2 = 13.60, p = 0.0002$	$\chi^2 = 103.2, p < 0.0001$
User scrolled	$\chi^2 = 190.6, p < 0.0001$	$\chi^2 = 10.43, p = 0.001$	$\chi^2 = 111.7, p < 0.0001$
Visit duration $\geq$ 17 seconds	$\chi^2 = 321.0, p < 0.0001$	$\chi^2 = 1.595, p = 0.207$	$\chi^2 = 277.0, p < 0.0001$
Site 2 Engagement Metrics	HTT vs UE	HTT vs MTT	MTT vs UE
Download link clicked	$\chi^2 = 2.477, p = 0.116$	$\chi^2 = 0.010, p = 0.921$	$\chi^2 = 2.963, p = 0.085$
Non-download link clicked	$\chi^2 = 0.025, p = 0.874$	$\chi^2 = 0.042, p = 0.838$	$\chi^2 = 0.216, p = 0.642$
User scrolled	$\chi^2 = 0.341, p = 0.560$	$\chi^2 = 1.196, p = 0.274$	$\chi^2 = 2.933, p = 0.087$
Visit duration $\geq$ 17 seconds	$\chi^2 = 7.373, p = 0.007$	$\chi^2 = 0.025, p = 0.874$	$\chi^2 = 8.479, p = 0.004$

Table 1: χ^2 tests for engagement metrics on Sites 1+2, comparing the HTT (Human Translation, show Translation by default), MTT (Machine Translation, show Translation by default), and UE (Untranslated English) conditions. See Figure 2 and Figure 3 for mean values.

switch to a machine translation in their preferred language or any supported language via the language switcher.

5 Study 1: Effects of human and machine translation availability on engagement

We study the effect of showing a human translation, machine translation, or no translation on the following engagement metrics:

- Percent of sessions where the user clicks on a download link.
- Percent of sessions where the user clicks on a non-download link.
- Percent of sessions where the user scrolls.
- Percent of sessions where the visit duration is in the top quartile of visit durations. (This is 17 or more seconds, for both Site A and B).

We chose these metrics as they are conversion events and proxies for reading. Clicking the download button is the "conversion event" for these sites, or what the site's most salient call to action is attempting to get the user to do. Clicking a non-download link suggests that the user is reading the text, as the non-download links are text-only, so users would presumably only click them if they read the associated link text and understood what the link points to. Scrolling likewise suggests that the user is reading, as both websites are long, single-page documents that

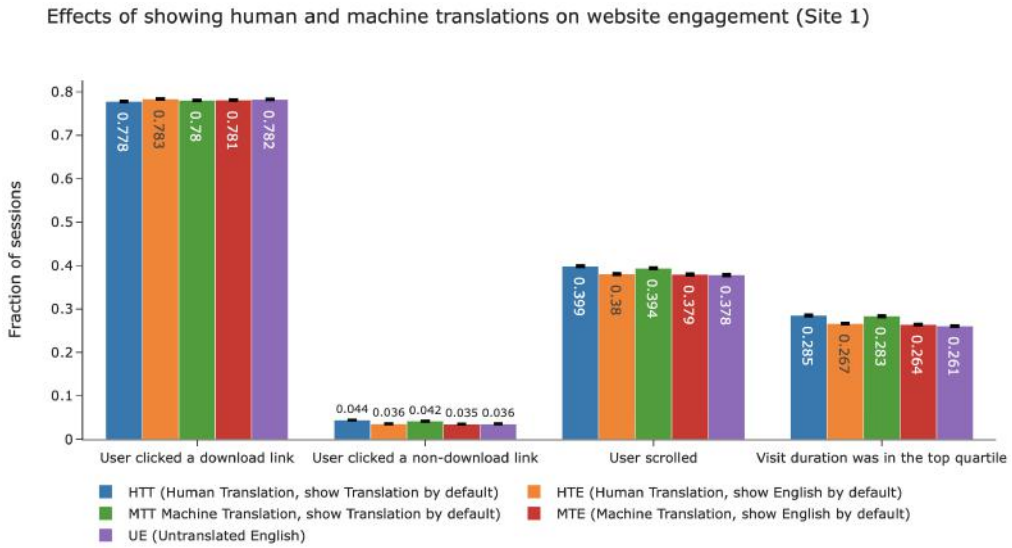

Figure 2: Engagement metrics for Site 1, for users whose preferred language is supported and not English. Error bars indicate 95% confidence intervals. Study 1 analyzes only the HTT, MTT, and UE conditions; the HTE and MTE conditions are analyzed in Study 2.

have the majority of textual content describing the software below the fold when viewed on a standard monitor. A visit duration in the top quartile (17 seconds or more) suggests the user is likely reading the page, as there are only images and no videos on the sites, so the non-textual portions of the websites can likely be skimmed in less than 17 seconds.

Results showing engagement metrics for Site 1 are in Figure 2. For Site 1, which had 1.23 million sessions from users with a supported non-English preferred language, χ^2 tests indicate there are significant differences for all 4 measured engagement metrics between the HTT, MTT, and UE conditions. These conditions have a significant effect on the proportion of sessions where the user clicks a download link (χ^2 omnibus test: $\chi^2 = 24.95, p < 0.0001$), the proportion of sessions where the user clicks a non-download link ($\chi^2 = 376.0, p < 0.0001$), the proportion of sessions where the user scrolls ($\chi^2 = 326.4, p < 0.0001$), and the proportion of sessions where the visit duration is in the top quartile—17 seconds or more ($\chi^2 = 562.3, p < 0.0001$). Post-hoc χ^2 analysis results are shown in Table 1. Users are significantly more likely to click a non-download link, scroll, or have a visit duration in the top quartile when shown a translation (HTT, MTT) than when shown untranslated English (UE). Scrolling and clicking on non-download links increases with human translations (HTT) over machine translations (MTT), but other engagement metrics do not significantly differ. Surprisingly, users are significantly more likely to click the download link if shown an untranslated English page (UE), than if shown a translation (HTT, MTT); perhaps this is due to the graphically salient nature of the download button.

Results for Site 2 are shown in Figure 3. For Site 2, which had only 9,316 sessions from users with a non-English preferred language, χ^2 tests indicate that between the HTT, MTT, and UE conditions, there are no significant differences in clicking on download links (χ^2 omnibus test: $\chi^2 = 9.285, p = 0.054$), clicking on non-download links ($\chi^2 = 0.392, p = 0.983$), or scrolling ($\chi^2 = 3.704, p = 0.447$). The only engagement metric with significant differences between conditions is whether the visit duration was in the top quartile – 17 seconds or more (χ^2 omnibus test: $\chi^2 = 13, 37, p = 0.0096$). Post-hoc χ^2 analysis results are shown in Table 1. Users are significantly more likely to stay on-page for at least 17 seconds if shown translations

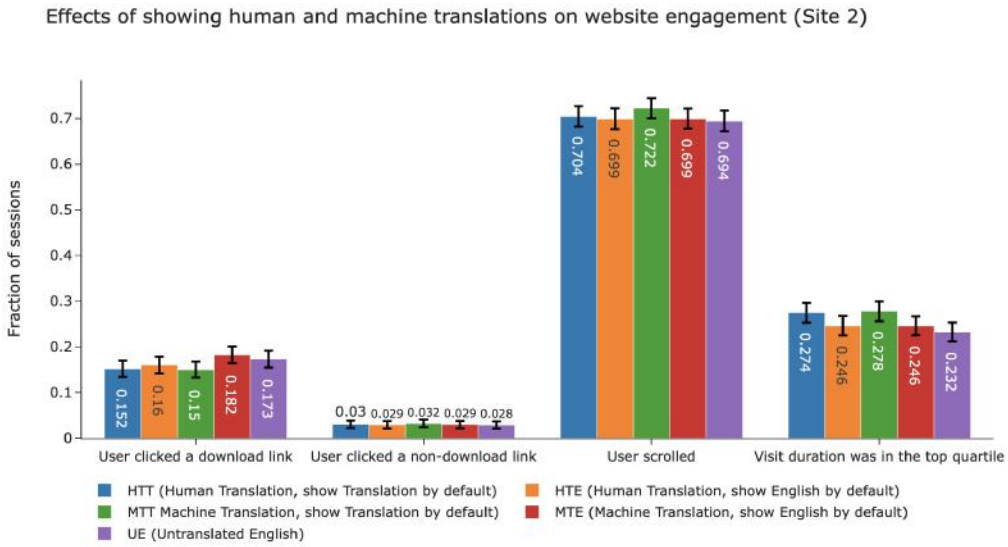

Figure 3: Engagement metrics for Site 2, for users whose preferred language is supported and not English. Error bars indicate 95% confidence intervals. Study 1 analyzes only the HTT, MTT, and UE conditions; the HTE and MTE conditions are analyzed in Study 2.

Site 1 Engagement Metrics	HTT vs HTE	MTT vs MTE	HTE vs UE	MTE vs UE
Download link clicked	$\chi^2 = 21.13, p < 0.0001$	$\chi^2 = 0.285, p = 0.594$	$\chi^2 = 0.679, p = 0.410$	$\chi^2 = 1.409, p = 0.235$
Non-download link clicked	$\chi^2 = 180.3, p < 0.0001$	$\chi^2 = 118.8, p < 0.0001$	$\chi^2 = 0.155, p = 0.694$	$\chi^2 = 0.539, p = 0.463$
User scrolled	$\chi^2 = 152.8, p < 0.0001$	$\chi^2 = 94.31, p < 0.0001$	$\chi^2 = 2.028, p = 0.154$	$\chi^2 = 0.737, p = 0.391$
Visit duration $\geq$ 17 seconds	$\chi^2 = 185.2, p < 0.0001$	$\chi^2 = 198.6, p < 0.0001$	$\chi^2 = 18.33, p < 0.0001$	$\chi^2 = 6.548, p = 0.010$
Site 2 Engagement Metrics	HTT vs HTE	MTT vs MTE	HTE vs UE	MTE vs UE
Download link clicked	$\chi^2 = 0.347, p = 0.556$	$\chi^2 = 5.970, p = 0.015$	$\chi^2 = 0.858, p = 0.354$	$\chi^2 = 0.428, p = 0.513$
Non-download link clicked	$\chi^2 = 0.005, p = 0.943$	$\chi^2 = 0.094, p = 0.759$	$\chi^2 = 0.000, p = 1.000$	$\chi^2 = 0.004, p = 0.950$
User scrolled	$\chi^2 = 0.083, p = 0.773$	$\chi^2 = 2.024, p = 0.155$	$\chi^2 = 0.064, p = 0.801$	$\chi^2 = 0.075, p = 0.785$
Visit duration $\geq$ 17 seconds	$\chi^2 = 3.222, p = 0.073$	$\chi^2 = 4.160, p = 0.041$	$\chi^2 = 0.749, p = 0.387$	$\chi^2 = 0.757, p = 0.384$

Table 2: χ^2 tests for engagement metrics on Sites 1+2, comparing the HTT (Human Translation, show Translation by default), MTT (Machine Translation, show Translation by default), HTE (Human Translation, show English by default), MTE (Machine Translation, show English by default), and UE (Untranslated English) conditions. See Figure 2 and Figure 3 for mean values.

(HTT, MTT) than if shown untranslated pages (UE).

Thus, we observe that while the availability of a translation increases the proportion of users who are retained for at least 17 seconds on both sites, the improvements in non-download link click rates and scrolling were observed only on Site 1 and not Site 2. Contrary to our expectations, we did not observe an increase in download rates due to translations being shown on either site. One explanation is that users may have read about the software in their native language elsewhere (referral logs indicated many visitors were coming from non-English sites), so they have no need to read the site's contents before downloading. Another explanation is that the download button is visually prominent, so perhaps some users just download and try the software to learn how it works, instead of reading about it first.

6 Study 2: Effects of default language choice and language switcher use

Websites commonly display translations in two ways: one is to determine users' preferred language via browser settings or their geographic region, and automatically show pages translated

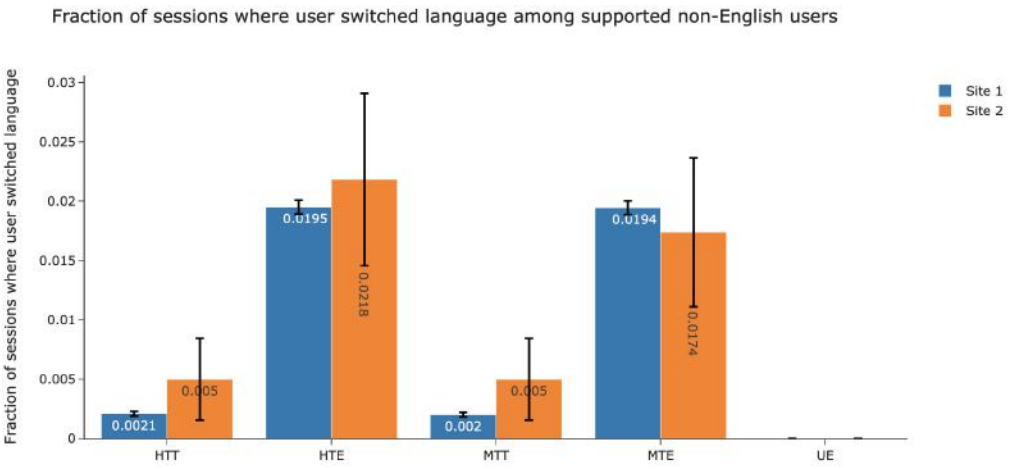

Figure 4: Use of the language switcher on Sites 1+2 in each condition, for users whose preferred language is supported and not English. Error bars indicate 95% confidence intervals.

to that language. Another approach is to show untranslated pages by default, and ask the user to use a language selector or click a link in order to see the translation. We wished to see how often users would use a language selector, and the effects of requiring language selection on user engagement. The language selector interface used in this experiment is shown in Figure 1.

Figure 4 shows the fraction of users who used the language selector in each condition on Sites 1 and 2. On both sites, there is a significant increase in usage of the language selector when English is shown by default (HTE, MTE) vs when a translation is shown by default (HTT, MTT). On Site 1, the language switcher is used in 1.95% of sessions in HTE vs 0.21% of sessions in HTT ($\chi^2 = 3095, p < 0.0001$). On Site 1, the language switcher is used in 1.94% of sessions in MTE vs 0.20% of sessions in MTT ($\chi^2 = 3123, p < 0.0001$). On Site 2, the language switcher is used in 2.18% of sessions in HTE vs 0.50% of sessions in HTT ($\chi^2 = 15.81, p < 0.0001$). On Site 2, the language switcher is used in 1.74% of sessions in MTE vs 0.50% of sessions in MTT ($\chi^2 = 10.14, p < 0.005$). Users hardly engage with language switchers on either site—even though it takes only a single button click to see the page in the user's preferred language, only 2% of users click it when English is shown by default. Switching away from the preferred language is even less frequent. Thus, automatically detecting the user's preferred language and showing translations accordingly is important—otherwise users will not see translations.

Showing English by default and requiring the user to use a language selector to view the page in their preferred language is detrimental to engagement. As we can see in Figure 2, the improvement in engagement metrics from translation we had previously observed for Site 1 are lost in the HTE and MTE conditions. χ^2 tests are shown in Table 2; there is a significant decrease in engagement (clicking non-download links, scrolling, and visit durations in the top quartile) when users need to click to switch to their preferred language. For most metrics, the HTE and MTE conditions do not display significantly higher engagement than the UE condition (Table 2). The same χ^2 tests were not significant on Site 2 (Table 2).

There are no significant differences in language switcher use between HTE vs MTE (Site 1: $\chi^2 = 0.021, p = 0.886$; Site 2: $\chi^2 = 0.616, p = 0.432$) or between HTT vs MTT (Site 1: $\chi^2 = 0.342, p = 0.559$; Site 2: $\chi^2 = 0, p = 1$). The fact that on both sites less than 0.6% in the MTT condition switch to English, and this is not significantly higher than the HTT condition, suggests that the machine translation was sufficiently usable that users did not switch to English.

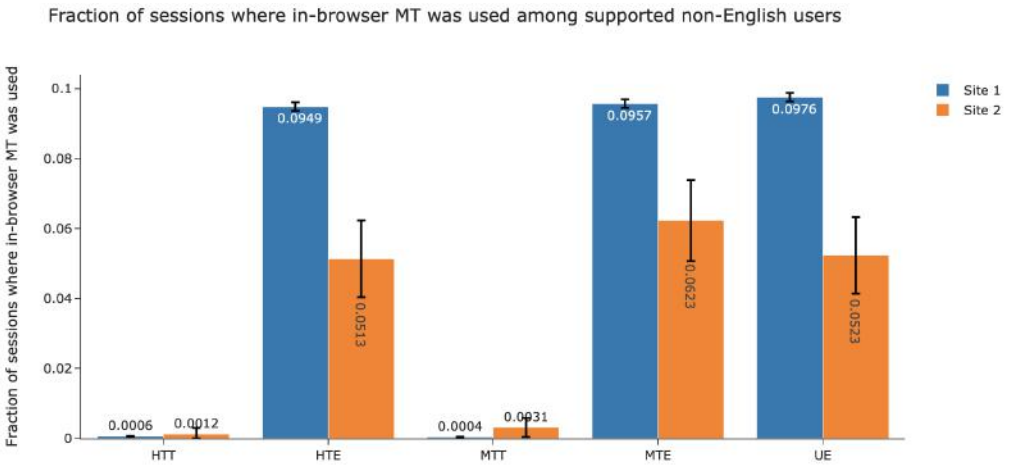

Figure 5: Use of in-browser machine translation on Sites 1+2 in each condition, for users whose preferred language is supported and not English. Error bars indicate 95% confidence intervals.

7 Study 3: In-browser machine translation use

We were surprised by our results from Study 1 that there was high engagement with untranslated English pages, even though the user's browser language preference was a non-English language that the site had been translated to, and our result from Study 2 that users rarely switch the page to their preferred language. We suspected that users may be viewing pages through the machine translation functionality that is built into browsers, so we measured the use of in-browser machine translation. Results are shown in Figure 5 for Sites 1 and 2. When a page is only available in English, 9.76% of users use in-browser machine translation on Site 1, and this number is 5.23% on Site 2. When a human translation is available in the user's preferred language, but the page is shown in English by default, 9.49% of users will use in-browser machine translation on Site 1, and this number is 5.13% on Site 2. This is higher than the 2% of users who used the language switcher from Study 2, meaning that if users are required to click a button to see a translation, users will be more likely to end up seeing their browser's machine translation than clicking the button to see a human translation.

We were curious whether the user's language and country has an influence on their use of in-browser machine translation. Is in-browser machine translation used at roughly equal levels everywhere, or is it used much more in some countries than others? Perhaps users whose preferred language is a language where the quality of machine translation from English is high will be more likely to use machine translation? Perhaps users whose preferred language has high lexical overlap with English, or uses the same Latin alphabet as English, will need to rely on machine translation less? Perhaps users from countries with a high level of English proficiency will be less likely to use machine translation? We display the in-browser machine translation usage broken down by the target language from Site 1 in Figure 6. Site 2 results are in Figure 8 of the Appendix.

Interestingly, it does not appear to be the case that machine translation use is highest for languages where the machine translation quality is better. For example, observe that Vietnamese (vi) has the highest proportion of users using the browser's built-in machine translation on Site 1 (14.8%) and is the fourth highest on Site 2 (11.3%), despite it being a low-resource language for which parallel data for training machine translation systems is scarce (Ngo et al., 2020). We can also observe the contrast in machine translation use between Spanish (es, 11.80% on Site 1

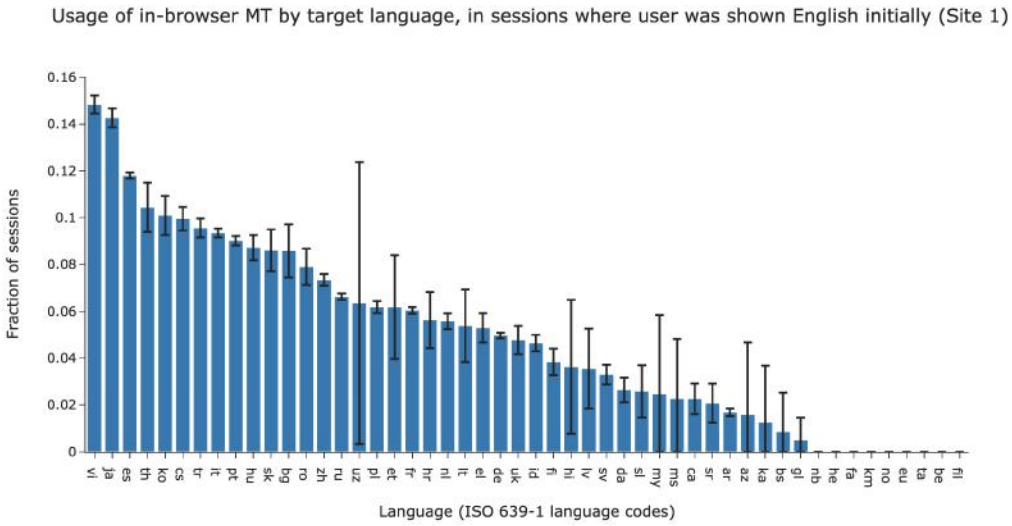

Figure 6: Use of in-browser machine translation on Site 1 by language, for users whose preferred language is not English. Error bars indicate 95% confidence intervals.

and 6.41% on Site 2) vs French (fr, 6.05% on Site 1 and 2.90% on Site 2)—despite both being high-resource Romance languages with high machine translation quality from English, and high lexical overlap with English for technical vocabulary. The type of script the language is written in does not seem to be the main influence on machine translation use either—the top 4 languages with the most machine translation use in Site 1 and Site 2 include both those written in non-Latin (Japanese, Thai, Bulgarian) and Latin scripts (Vietnamese, Spanish, Romanian).

Perhaps we should consider countries rather than the properties of the languages themselves? We show the in-browser machine translation usage broken down by country from Site 1 in Figure 7. Site 2 results are in Figure 8 of the Appendix. Since we suspected that English proficiency may influence machine translation usage, we group countries by their English proficiency ranking on the EF English Proficiency Index 2020 (Index, 2020). We find that among users with non-English preferred languages, those from countries with lower English proficiency tend to use machine translation more. On Site 1, 9.84% of sessions with non-English preferred languages from countries with "moderate" or lower English proficiency use in-browser machine translation, whereas only 6.10% of such sessions from counties with "high" or "very high" English proficiency used in-browser machine translation; this difference was statistically significant ($\chi^2 = 3797, p < 0.0001$). On Site 2, 6.33% of sessions with non-English preferred languages from countries with "moderate" or lower English proficiency use in-browser machine translation, whereas only 4.05% of such sessions from counties with "high" or "very high" English proficiency used in-browser machine translation; this difference was statistically significant ($\chi^2 = 18.67, p < 0.0001$).

8 Dataset and Code

To help replicate this study and enable researchers to run further analyses, we publicly release the datasets for both websites and the notebooks to reproduce our results at `https://transabtest.github.io/` under the Creative Commons Attribution license. The dataset represents the full 3,298,635 sessions from Site 1, and the full 34,034 sessions from Site 2, and the site content and translations.

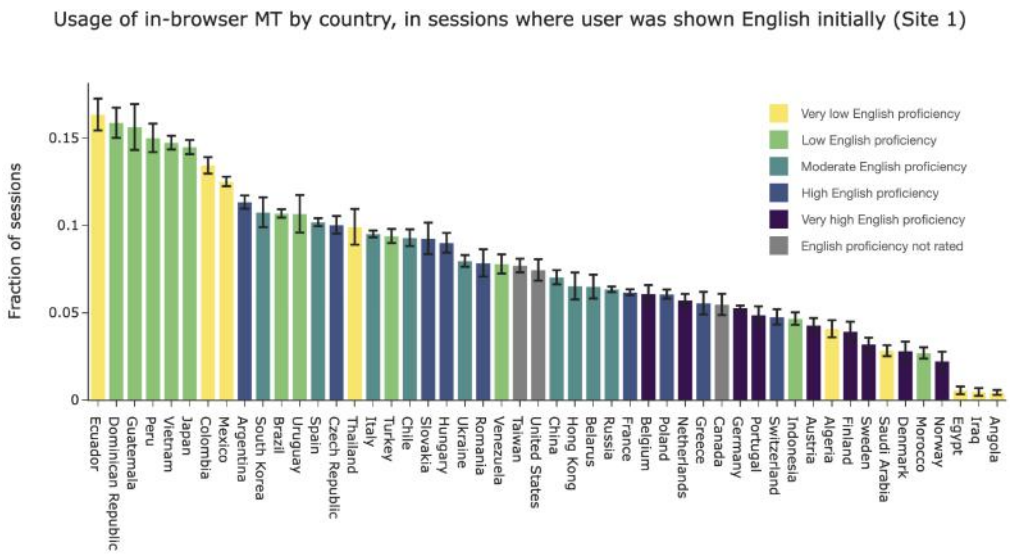

Figure 7: Use of in-browser machine translation on Site 1 by country, for users whose preferred language is not English. Colors indicate English proficiency ratings by the EF English Proficiency Index 2020. Error bars indicate 95% confidence intervals.

9 Conclusion

As the population of internet users grows increasingly linguistically diverse, it is increasingly important for websites to consider whether they should translate their websites, and how to present translations. Machine translation quality has considerably increased over the years, and in-browser machine translation has likewise become integrated into major browsers, which leads us to ask how much benefit websites can expect from translating their website with human translators, as opposed to showing machine translations or leaving pages untranslated.

Through an A/B test we run on the homepages of two open-source software projects, we find that both human and machine translations improve engagement metrics that are indicative of reading the page, though download rates remain high regardless of whether a translation is shown or not. Compared to machine translations, human translations increase two engagement metrics (scrolling and clicking on non-download links), but had no effect on others. If we require users to click a button to switch to their preferred language, they rarely do so, and the gains in engagement metrics we observed from translation are negated. A significant minority of users with non-English preferred languages use machine translation systems integrated into browsers, especially from countries with low English proficiency. This, along with our finding that only a small fraction of users switch to English if shown a machine translation to their preferred language, suggests that machine translations are often of acceptable quality to some users.

While showing machine translations performed well on the software-centric sites we studied, our findings may not generalize to other types of content; perhaps human translations will have more visible benefits on more text-centric content such as news. The effects of translation on engagement will also change over time, as machine translation quality improves, the web's userbase becomes increasingly multilingual, and as foreign language reading abilities change. Localization decision makers choosing between human and machine translations should also consider website demographics such as their visitors' countries and preferred languages, and the quality of machine translation for the languages they are considering supporting.

References

Barrault, L., Bojar, O., Costa-Jussa, M. R., Federmann, C., Fishel, M., Graham, Y., Haddow, B., Huck, M., Koehn, P., Malmasi, S., et al. (2019). Findings of the 2019 conference on machine translation (wmt19). In *Proceedings of the Fourth Conference on Machine Translation (Volume 2: Shared Task Papers, Day 1)*, pages 1–61.

Brynjolfsson, E., Hui, X., and Liu, M. (2019). Does machine translation affect international trade? evidence from a large digital platform. *Management Science*, 65(12):5449–5460.

Cappelli, G. (2007). The translation of tourism-related websites and localization: problems and perspectives. *Rassegna italiana di linguistica applicata*, 39(1/2):97.

Group, M. M. (2020). Interet world stats usage and population statistics. *Retrieved April*, 27:2021.

Hassan, H., Aue, A., Chen, C., Chowdhary, V., Clark, J., Federmann, C., Huang, X., Junczys-Dowmunt, M., Lewis, W., Li, M., et al. (2018). Achieving human parity on automatic chinese to english news translation. *arXiv preprint arXiv:1803.05567*.

Index, E. E. P. (2020). Ef english proficiency index a ranking of 100 countries and regions by english skills. *Retrieved April*, 27:2021.

Läubli, S. and Orrego-Carmona, D. (2017). When google translate is better than some human colleagues, those people are no longer colleagues.

Läubli, S., Sennrich, R., and Volk, M. (2018). Has machine translation achieved human parity? a case for document-level evaluation. *arXiv preprint arXiv:1808.07048*.

Ngo, T.-V., Nguyen, P.-T., Ha, T.-L., Dinh, K.-Q., and Nguyen, L.-M. (2020). Improving multilingual neural machine translation for low-resource languages: French-, english-vietnamese. *arXiv preprint arXiv:2012.08743*.

Popel, M., Tomkova, M., Tomek, J., Kaiser, Ł., Uszkoreit, J., Bojar, O., and Žabokrtskỳ, Z. (2020). Transforming machine translation: a deep learning system reaches news translation quality comparable to human professionals. *Nature communications*, 11(1):1–15.

Toral, A. (2020). Reassessing claims of human parity and super-human performance in machine translation at wmt 2019. *arXiv preprint arXiv:2005.05738*.

Toral, A., Castilho, S., Hu, K., and Way, A. (2018). Attaining the unattainable? reassessing claims of human parity in neural machine translation. *arXiv preprint arXiv:1808.10432*.

A Appendix: Demographics and Supplementary Figures

Country from IP geolocation: *Site 1*: Top 20 countries are USA (12.411%), Germany (6.397%), Russia (5.501%), India (5.48%), Italy (5.044%), France (4.579%), Brazil (4.041%), Spain (3.736%), Mexico (3.293%), United Kingdom (2.829%), Vietnam (2.314%), Indonesia (2.219%), Poland (2.105%), Canada (2.018%), Ukraine (1.675%), Japan (1.618%), Argentina (1.557%), Turkey (1.462%), Netherlands (1.412%), Colombia (1.165%). *Site 2*: Top 20 countries are USA (19.047%), Germany (8.65%), France (5.109%), Italy (4.91%), India (4.719%), Russia (4.633%), Spain (3.514%), United Kingdom (3.179%), Brazil (2.917%), Canada (2.574%), Poland (2.248%), Mexico (2.212%), Netherlands (1.669%), Ukraine (1.654%), Australia (1.595%), Vietnam (1.325%), Argentina (1.257%), Japan (1.257%), China (1.143%), Turkey (1.134%).
Preferred Language: *Site 1*: Top 20 are English (46.115%), Spanish (11.831%), Russian (6.567%), French (6.021%), German (5.932%), Italian (4.624%), Portuguese (4.088%), Chinese (2.032%), Polish (1.732%), Vietnamese (1.646%), Japanese (1.407%), Arabic (1.305%), Turkish (1.03%), Dutch (0.873%), Indonesian (0.688%), Czech (0.688%), Hungarian (0.542%), Swedish (0.356%), Korean (0.253%), Ukrainian (0.25%). *Site 2*: Top 20 are English (53.679%), Spanish (8.259%), German (6.529%), Russian (5.809%), French (5.248%), Italian (4.51%), Portuguese (2.738%), Chinese (2.271%), Polish (1.828%), Arabic (1.275%), Japanese (0.97%), Vietnamese (0.92%), Dutch (0.855%), Turkish (0.77%), Czech (0.526%), Hungarian (0.461%), Korean (0.42%), Swedish (0.373%), Indonesian (0.291%), Ukrainian (0.253%).
Gender from Google Analytics: *Site 1*: 24.6% female, 75.4% male. *Site 2*: 29.2% female, 70.8% male.
Age from Google Analytics: *Site 1*: 30.14% of users are 18-24 years old, 30.44% are 25-34, 15.68% are 35-44, 11.78% are 45-54, 6.57% are 55-64, 5.40% are over 65. *Site 2*: 20.9% of users are 18-24 years old, 28.7% are 25-34 years old, 19.1% are 33-44, 13.9% are 45-54, 8.70% are 55-65, 8.70% are over 65.

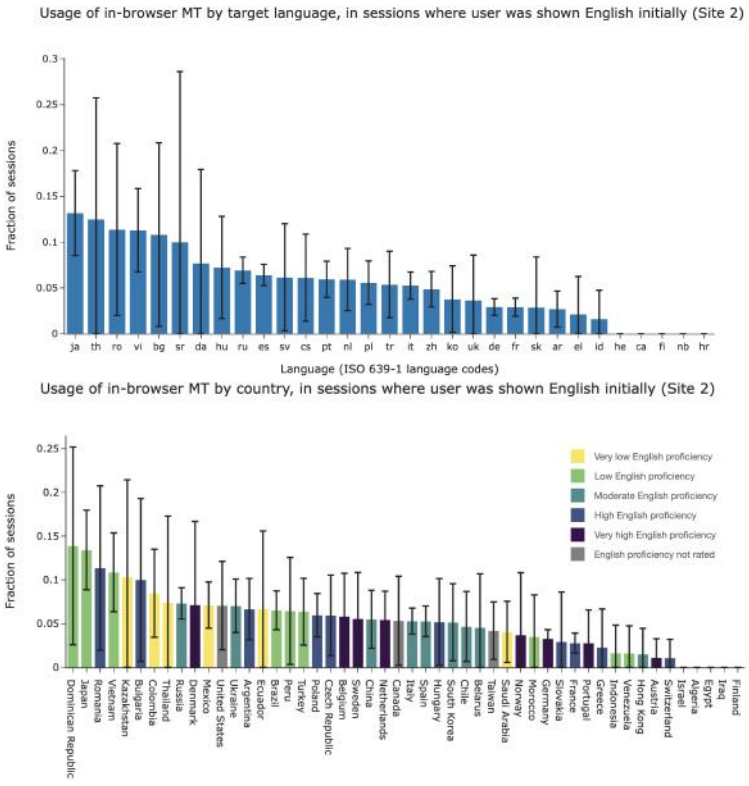

Figure 8: Use of in-browser machine translation on Site 2 for users whose preferred language is not English. Error bars show 95% confidence intervals. Top: by language. Bottom: by country.