

Document downloaded from:

<http://hdl.handle.net/10251/176103>

This paper must be cited as:

Domingo-Ballester, M.; Casacuberta Nolla, F. (2020). Modernizing historical documents: A user Study. *Pattern Recognition Letters*. 133:151-157.
<https://doi.org/10.1016/j.patrec.2020.02.027>



The final publication is available at

<https://doi.org/10.1016/j.patrec.2020.02.027>

Copyright Elsevier

Additional Information

Modernizing Historical Documents: a User Study

Miguel Domingo and Francisco Casacuberta

Pattern Recognition and Human Language Technology Research Center
Universitat Politècnica de València - Camino de Vera s/n, 46022 Valencia, Spain

midobal@prhlt.upv.es, fcn@prhlt.upv.es

Abstract

Accessibility to historical documents is mostly limited to scholars. This is due to the language barrier inherent in human language and the linguistic properties of these documents. Given a historical document, modernization aims to generate a new version of it, written in the modern version of the document’s language. Its goal is to tackle the language barrier, decreasing the comprehension difficulty and making historical documents accessible to a broader audience. In this work, we proposed a new neural machine translation approach that profits from modern documents to enrich its systems. We tested this approach with both automatic and human evaluation, and conducted a user study. Results showed that modernization is successfully reaching its goal, although it still has room for improvement.

1 Introduction

Historical documents are an important part of our cultural heritage. However, the nature of human language, which evolves with the passage of time, and the linguistic properties of these documents—due to the lack of a spelling convention, orthography changes depending on the time period and author—increase the difficulty of comprehending them. For this reason, historical documents are mostly accessible to scholars. Thus, in order to preserve them and make them reachable to a broader audience, a

scholar is typically in charge of producing a comprehensive contents document which allows non-experts to locate and gain a basic understanding of a given document (e.g., Monk, 2018).

Modernization aims to tackle this language barrier by generating a new version of a historical document, written in the modern version of the document’s original language. Fig. 1 shows an example of modernizing a document. In this case, part of the language structures and rhymes have been lost. However, the modern version is easier to read and comprehend by a broader audience. This problem is also present in poetry translation since the entwinement between sound and word and sense cannot be truly replicated in a different language (Ilonka, 2018). However, translating a poem from one language into another is a way of sharing cultural practices and ideologies across languages (Rajvanshi, 2015).

Modernization can be a controversial topic since it implies an alteration of the original document (e.g., the manual modernization of *El Quijote* rose a controversy in Spain (Flood, 2015)). However, it is manually applied to classic literature in order to make works that had been relegated to scholars due to the hardness of their comprehension, understandable to contemporary readers (Rodríguez Marcos, 2015).

Finally, while the language richness present in historical documents is also part of our cultural heritage, the goal of modernization is to make historical documents accessible to a general audience. Other research topics on historical document are focused

O Romeo, Romeo! Wherefore art thou Romeo?
Deny thy father and refuse thy name.
Or, if thou wilt not, be but sworn my love,
And I'll no longer be a Capulet.

With love's light wings did I o'erperch these walls,
For stony limits cannot hold love out,
And what love can do, that dares love attempt.
Therefore thy kinsmen are no stop to me.

Oh, Romeo, Romeo, why do you have to be Romeo?
Forget about your father and change your name.
Or else, if you won't change your name, just swear you love me
and I'll stop being a Capulet.

I flew over these walls with the light wings of love.
Stone walls can't keep love out.
Whatever a man in love can possibly do, his love will make him try to do it.
Therefore your relatives are no obstacle.

Figure 1: Example of modernizing a historical document. The original text is composed of fragments from *Romeo and Juliet* by *William Shakespeare*. The modernized version was obtained from Crowther (2003).

on different aspects of their richness. For example, historical manuscripts are automatically transcribed and digitized (Toselli et al., 2010, 2017). Orthography is normalized to account for the lack of a spelling convention (Laing, 1993; Porta et al., 2013). Search queries can find all occurrences of one or more words (Rogers and Willett, 1991; Ernst-Gerlach and Fuhr, 2006). Word frequency lists are generated (e.g., Baron et al., 2009). And natural language processing tools provide automatic annotations to identify and extract linguistic structures such as relative clauses (Hundt et al., 2011) or verb phrases (Fiebranz et al., 2011; Pettersson et al., 2013).

In this work, we followed a machine translation (MT) approach to tackle the modernization problem. Similarly to Domingo and Casacuberta (2018), we profited from modern documents to enrich the modernization systems. However, we applied a data selection technique to take better profit of these documents, selecting only the most relevant sentences for each task. We evaluated our approach both automatically and with the help of 4 scholars specialized in classic Spanish literature. Additionally, we conducted a user study with 42 people to assess whether or not modernization is able to decrease the difficulty of comprehending historical documents. Our main contributions are as follows:

- We proposed a new neural machine translation (MT) approach that successfully profits from modern documents to enrich its modernization systems.
- We tested our proposal using 3 datasets from different languages and time periods.
- We assessed the quality of our proposal us-

ing both automatic and human evaluation, conducted by 4 scholars specialized in classic Spanish literature.

- First time, to the best of our knowledge, in which an NMT modernization approach behaves similarly or better than a statistical machine translation (SMT) modernization approach.
- We conducted a study with 42 users to assess whether modernization successfully decreases the difficulty of comprehending historical documents.

The rest of this document is structured as follows: Section 2 introduces the related work. Then, Section 3 presents the modernization approach. After that, in Section 4, we describe the experimental framework of our work. Next, in Section 5, we present and discuss the evaluation conducted in order to assess our approach. Section 6 describes and presents the user study. Finally, in Section 7, conclusions are drawn.

2 Related work

Modernization has been manually applied to literature for centuries. One of the most well-known examples is *The Bible*, which has been adapted and translated for generations in order to preserve and transmit its contents (Given, 2015). Classic literature is also frequently modernized in order to bring it closer to a contemporary audience (e.g., *No Fear*

*Shakespeare*¹; *Odres Nuevos*²; *El Quijote* (Trapiello, 2015)). However, on the literature we find that, while normalizing orthography to account for the lack of a spelling convention has been extensively researched for years (Laing, 1993; Baron and Rayson, 2008; Porta et al., 2013; Hämäläinen et al., 2018), automatic modernization of historical documents is a young research field.

One of the first related works was a shared task for translating historical text to contemporary language (Tjong Kim Sang et al., 2017). The task was focused on normalizing the document’s spelling. However, they also approached document modernization using a set of rules. Domingo et al. (2017) proposed a modernization approach based on SMT. An NMT approach was proposed by Domingo and Casacuberta (2018). Finally, Sen et al. (2019) augmented the training data by extracting pairs of phrases and adding them as new training sentences.

3 Modernization approaches

In this section, we present the state-of-the-art SMT modernization approach and our NMT-based proposal. Both approaches rely on MT which, given a source sentence \mathbf{x} , aims at finding the most likely translation $\hat{\mathbf{y}}$ (Brown et al., 1993):

$$\hat{\mathbf{y}} = \arg \max_{\mathbf{y}} Pr(\mathbf{y} | \mathbf{x}) \quad (1)$$

3.1 SMT approach

For years, SMT has been the prevailing approach to compute Eq. (1), using models that rely on a log-linear combination of different models (Och and Ney, 2002): namely, phrase-based alignment models, re-ordering models and language models; among others (Zens et al., 2002; Koehn et al., 2003).

In this approach, modernization is tackled as a conventional translation task: training an SMT system from a parallel corpora in which, for each sentence of the original document, its corresponding modernized

version is available. For training this system, the language of the original document is considered as the source language, and its modernized version as the target language.

3.2 NMT approach

NMT models Eq. (1) with a neural network which usually follows an encoder-decoder architecture, in which the source sentence is projected into a distributed representation at the encoding step. Then, at the decoding step, the decoder generates its most likely translation—word by word—using a beam search method (Sutskever et al., 2014).

The system’s input is a word sequence in the source language. An embedding matrix linearly projects each word to a fixed-size real-valued vector. These words embeddings are, then, fed into a bidirectional (Schuster and Paliwal, 1997) long short-term memory (LSTM) (Hochreiter and Schmidhuber, 1997) network. As a result, a sequence of annotations is produced by concatenating the hidden states from the forward and backward layers. An attention mechanism (Bahdanau et al., 2015) allows the decoder to focus on parts of the input sequence, computing a weighted mean of annotated sequences. A soft alignment model computes these weights, weighting each annotation with the previous decoding state. Another LSTM network—conditioned by the representation computed by the attention model and the last word generated—is used for the decoder. Finally, a distribution over the target language vocabulary is computed by the deep output layer (Pascanu et al., 2013). The model is trained by applying stochastic gradient descent jointly to maximize the log-likelihood over a bilingual parallel corpus.

As the SMT approach (see Section 3.1), our proposal tackles modernization as a conventional translation task but using NMT instead of SMT. Additionally, since NMT systems need larger quantities of training data, and a frequent problem when working with historical documents is the scarce availability of parallel training data (Bollmann and Søgaard, 2016), we created synthetic data in order to profit from modern documents to enrich the NMT models. First, we applied feature decay algorithm (Biçici and Yuret,

¹<https://www.sparknotes.com/shakespeare/>.

²https://www.castalia.es/libros?tipo=coleccion&letra=0&numero=49&other_page=1.

2015) to select those documents which are closer to the ones we have to modernize. After that, we followed a backtranslation approach (Sennrich et al., 2015) to create a parallel synthetic corpus. Backtranslation has become the norm when building state-of-the-art NMT systems—especially in resource-poor scenarios (Poncelas et al., 2018). Given a monolingual corpus in the target language and an MT system trained to translate from the target language to the source language, the synthetic data is generated by translating the monolingual corpus with the MT system—the resulting data is used as the source part of the corpus, and the monolingual data as the target part.

4 Experimental framework

In this section, we describe the MT systems, corpora and evaluation metrics from our experimental framework.

4.1 MT systems

SMT systems were trained with **Moses** (Koehn et al., 2007), following the standard procedure: we estimated a 5-gram language model—smoothed with the improved KneserNey method—using **SRILM** (Stolcke, 2002), and optimized the weights of the log-linear model with **MERT** (Och, 2003). SMT systems were used both for the SMT modernization approach and for generating synthetic data (see Section 3).

We built NMT systems using **OpenNMT-py** (Klein et al., 2017). We used long short-term memory units (Gers et al., 2000), with all model dimensions set to 512. We trained the system using **Adam** (Kingma and Ba, 2014) with a fixed learning rate of 0.0002 and a batch size of 60. We applied label smoothing of 0.1 (Szegedy et al., 2015). At inference time, we used beam search with a beam size of 6. In order to reduce vocabulary, we applied joint byte pair encoding (BPE) (Sennrich et al., 2016) to all corpora, using 32,000 merge operations. NMT systems were trained using synthetic data and, then, were fine-tuned with the training data.

4.2 Corpora

Dutch Bible (Tjong Kim Sang et al., 2017): A collection of different versions of the Dutch Bible. Among others, it contains a version from 1637—which we consider as the original version—and another from 1888—which we consider as the modern version (using 19th century Dutch as if it were *modern Dutch*).

El Quijote (Domingo and Casacuberta, 2018): the well-known 17th century Spanish novel by Miguel de Cervantes, and its correspondent 21st century version.

OE-ME (Sen et al., 2019): contains the original 11th century English text *The Homilies of the Anglo-Saxon Church* and a 19th century version—which we consider as *modern English*.

As reflected in Table 1, the corpora sizes are small. Thus, the use of synthetic data to profit from modern documents and increase the training data (see Section 3.2). As *modern documents*, we made use of the collection of Dutch books available at the *Digitale Bibliotheek voor de Nederlandse letteren*³, for Dutch; and **OpenSubtitles** (Lison and Tiedemann, 2016)—a collection of movie subtitles in different languages—for Spanish and English.

4.3 Metrics

Modernization adopted evaluation metrics from MT. In order to assess our proposal, we made use of:

Translation Error Rate (TER) (Snover et al., 2006): number of word edit operations (insertion, substitution, deletion and swapping), normalized by the number of words in the final translation.

BiLingual Evaluation Understudy (BLEU) (Papineni et al., 2002): geometric average of the modified n-gram precision, multiplied by a brevity factor.

³<http://dbnl.nl/>.

Table 1: Corpora statistics. $|S|$ stands for number of sentences, $|T|$ for number of tokens and $|V|$ for size of the vocabulary. *Modern documents* refers to the monolingual data used to create the synthetic data. M denotes millions and K thousands.

| | | Dutch Bible | | El Quijote | | OE-ME | |
|------------------|-------|-------------|------------|------------|------------|----------|------------|
| | | Original | Modernized | Original | Modernized | Original | Modernized |
| Train | $ S $ | 35.2K | | 10K | | 2716 | |
| | $ T $ | 870.4K | 862.4K | 283.3K | 283.2K | 64.3K | 69.6K |
| | $ V $ | 53.8K | 42.8K | 31.7K | 31.3K | 13.3K | 8.6K |
| Validation | $ S $ | 2000 | | 2000 | | 500 | |
| | $ T $ | 56.4K | 54.8K | 53.2K | 53.2K | 12.2K | 13.3K |
| | $ V $ | 9.1K | 7.8K | 10.7K | 10.6K | 4.2K | 3.2K |
| Test | $ S $ | 5000 | | 2000 | | 500 | |
| | $ T $ | 145.8K | 140.8K | 41.8K | 42.0K | 11.9K | 12.9K |
| | $ V $ | 10.5K | 9.0K | 8.9K | 9.0K | 4.1K | 3.2K |
| Modern documents | $ S $ | 3.0M | | 2.0M | | 6.0M | |
| | $ T $ | 76.1M | 74.1M | 22.3M | 22.2M | 67.5M | 71.6M |
| | $ V $ | 1.7M | 1.7M | 210.1K | 211.7K | 290.2K | 287.4K |

We used **sacreBLEU** (Post, 2018) in order to ensure consistent BLEU scores. Additionally, we applied approximate randomization tests (Riezler and Maxwell, 2005)—with 10,000 repetitions and using a p -value of 0.05—to determine whether two systems presented statistically significance.

5 Evaluation

In order to assess the quality of our modernization approaches, we started by performing an automatic evaluation. Then, with the help of 4 scholars, we conducted a human evaluation.

5.1 Automatic evaluation

Table 2 presents the results of the experimental session. All approaches significantly improved the modernization quality. Differences between the SMT and NMT approaches were only statistically significant for *Dutch Bible*. In that case, the NMT approach yielded the best results: an overall improvement of 46.8 points according to TER and 67.7 points according to BLEU; and an improvement of 0.4 and 2.9 points according to TER and BLEU respectively, with respect to the SMT approach.

To the best of our knowledge, this is the first

time that an NMT modernization approach is able to achieve these kinds of results. Domingo and Casacuberta (2018) already tried to profit from modern documents to enrich the neural models. However, their approach only improved the modernization quality in some cases—and never enough to reach the quality of the SMT approach—while in others it lowered it significantly. Our approach was based on theirs, but we used a data selection technique to help us filtered the monolingual data in order to generate synthetic data more suitable for each task.

5.2 Human evaluation

The human evaluation was performed by 4 scholars specialized in classic Spanish literature. For this reason, it was conducted using *El Quijote*. We randomly selected 100 sentences, checking that modernizations were different to the original sentences. We showed each sentence together with its modernization—50 sentences modernized with the SMT approach and another 50 with the NMT approach—and asked the scholars to give a rating according to the quality of the following aspects: fluency, lexical meaning, syntax, semantic and modernization. To avoid any bias, we shuffled the sentences and did not give any detail to the evaluators about how modernizations had been produced. Table 3 shows the results of the eval-

Table 2: Experimental results. *Baseline* system corresponds to considering the original document as the modernized version. † indicates statistically significance between the SMT/NMT approaches and the baseline. ‡ indicates statistically significance between the NMT and SMT approaches. [↓] indicates that the lowest the value the highest the quality. [↑] indicates that the highest the value the highest the quality.

| Approach | Dutch Bible | | El Quijote | | OE-ME | |
|----------|-------------|----------|------------|----------|---------|----------|
| | TER [↓] | BLEU [↑] | TER [↓] | BLEU [↑] | TER [↓] | BLEU [↑] |
| Baseline | 57.9 | 12.9 | 44.2 | 36.3 | 91.0 | 2.8 |
| SMT | 11.5† | 77.5† | 30.7† | 58.3† | 39.6† | 39.6† |
| NMT | 11.1†‡ | 80.6†‡ | 31.9† | 57.3† | 44.3† | 35.9† |

uation.

While the automatic evaluation (see Section 5.1) did not show any significant differences between the SMT and NMT approaches, the human evaluators slightly preferred SMT over NMT. Scores vary considerably depending on the evaluator—scholar₁ and scholar₄ gave higher scores than scholar₂ and scholar₃. However, all evaluators agreed that fluency is the strongest point of both approaches. In general, scores are above the average, which seems to correlate with the automatic evaluation.

When we asked evaluators about their opinion, they commented that the main problems were related with punctuation and diacritical marks. They also mentioned that, sometimes, part of the sentence was lost in the modernization—a known issue related with NMT (Wu et al., 2016). Additionally, scholar₁ commented that, overall, the quality of the modernization was acceptable. However, scholar₂ commented that if they had to correct the mistakes, they would have preferred to do the modernization from scratch.

6 User study

In order to assess whether modernization is able to decrease the difficulty of comprehending historical documents and, thus, making them accessible to a broader audience, we conducted a user study using *El Quijote*. 42 participants took part in this study. Considering that *El Quijote* is well-known in Spain, we asked participants about their familiarity with it. Fig. 2 shows some information about the user’s age and their familiarity with *El Quijote*.

The majority of the participants were between 20 and 50 years old, but there were also older and younger people. With one exception, all participants were familiar with *El Quijote* to some extent. In fact, 35.7% of them had read the original version of the novel.

The study consisted in several questions in which we showed two sentences to the user—the original sentence and its modernized version (either by the SMT or the NMT approach)—and asked them to select which sentence was easier for them to read and comprehend, if both of them had the same difficulty, or if they thought that both sentence did not have the same meaning. The selected sentences were the same used in the human evaluation (see Section 5.2). In order to avoid any bias, the order in which sentences appeared (i.e., the original sentence and its modernized version) was randomized, as well as the use of the different approaches. Fig. 3 shows an example of a question.

Table 4 presents the results of the study. Despite the users’ familiarity with *El Quijote*, modernization succeed in making it easier to comprehend. No matter the modernization approach, users selected the modernized version in the majority of the cases. In most of the remaining cases, users did not find any significant difference with respect to the original sentence.

When comparing both approaches, we observe that the SMT approach yielded better results: Users selected 61.4% of their modernized versions, while they only selected a 50.9% of the sentences modernized by the NMT approach. Additionally, the SMT approach only introduced errors in 7.8% of the cases—the NMT

Table 3: Results of the human evaluation. Values correspond to the average score for all sentences of each approach. 1 is the lowest score and 5 is the highest.

| Scholar | SMT approach | | | | | NMT approach | | | | |
|----------------------|--------------|-----------------|--------|----------|---------------|--------------|-----------------|--------|----------|---------------|
| | Fluency | Lexical meaning | Syntax | Semantic | Modernization | Fluency | Lexical meaning | Syntax | Semantic | Modernization |
| Scholar ₁ | 5.0 | 4.3 | 4.3 | 4.6 | 3.9 | 4.8 | 4.0 | 4.0 | 4.1 | 4.0 |
| Scholar ₂ | 2.1 | 1.9 | 2.0 | 2.1 | 2.0 | 2.0 | 1.9 | 1.9 | 1.9 | 1.9 |
| Scholar ₃ | 3.2 | 3.1 | 2.9 | 2.9 | 3.1 | 3.3 | 3.2 | 2.9 | 3.0 | 3.1 |
| Scholar ₄ | 4.5 | 3.9 | 4.6 | 4.3 | 4.0 | 3.8 | 3.5 | 3.7 | 3.7 | 3.5 |
| Average | 3.7 | 3.3 | 3.4 | 3.5 | 3.2 | 3.4 | 3.1 | 3.1 | 3.2 | 3.1 |




Figure 2: Information about study participants.

Select the sentence which is easier for you to read and comprehend:

- Y, leuantandose, dexó de comer, y fue a quitar la cubierta de la primera imagen, que mostro ser la de San Iorge puesto a cauhallo, con vna serpiente enroscada a los pies, y la lança atrauessada por la boca, con la fiereça que suele pintarse.
- Y levantándose, dejó de comer, y fue a quitar la cubierta de la primera imagen, que mostró ser la de San Jorge puesto a caballo, con una serpiente enroscada a los pies, y la lanza atravesada por la boca, con la fiereça que suele pintarse.
- Indifferent.
- Both sentences do not have the same meaning.

Figure 3: Example of a question.

introduced them in 20.3% of the cases—and its modernized versions were harder to comprehend only in 3.2% of the cases—versus a 6.4% of the cases for the NMT approach. Therefore, despite neither the automatic nor the human evaluation were able to find significant differences between both approaches, the user study showed that the SMT approach produced versions easier to read and comprehend more successfully than the NMT approach.

6.1 Qualitative analysis

In this section, we show some behavioral examples of the modernization approach. The example from Fig. 3 shows a successfully modernized sentence. Except for one small mistake (*fiereça*, which should be *fiereza*), orthography has been successfully modernized, making the sentence easier to read. (Note that, in this case, orthography is the only thing that needs

Table 4: Results [%] of the user study. *Original* means that users understood better the original version. *Modernized* means that users understood better the modernized version. *Indifferent* means that users did not found any significant differences between the original and modernized versions. *Not equal* means that users felt that the meaning between both version differ.

| SMT | | | | NMT | | | |
|----------|------------|-------------|-----------|----------|------------|-------------|-----------|
| Original | Modernized | Indifferent | Not equal | Original | Modernized | Indifferent | Not equal |
| 3.2 | 61.4 | 27.6 | 7.8 | 6.4 | 50.9 | 22.3 | 20.3 |

to be modified in order to achieve a modern Spanish version.)

Fig. 4 shows an example in which there is not any significant difference between the modernized and the original version. Only three words have been modified—and one of them (*huéolo*) is not even a real word but a mistake introduced by the use of BPE. Despite this, there are people who found the modernized version easier to read; a great majority that found no difference between them; and a few people that either preferred the original version or considered that they did not have the same meaning.

In Fig. 5, we can see an example in which the original sentence is easier to understand than its modernized version. While users considered both versions to have the same meaning, the modernized one is harder to comprehend since the first half of the sentence does not make much sense. In fact, looking at the human evaluation, scholars considered the modernized version to be more or less fluent, but with a poor lexical meaning, syntax and semantic.

Finally, Fig. 6 shows an example in which the modernization went very bad. On the one hand, the modernized version is way shorter than the original version. On the other hand, its meaning has no relation with the original one.

7 Conclusions and future work

In this work, we proposed a new NMT modernization approach in order to tackle the language barrier inherent in historical documents. We tested this approach on three different historical datasets from three different languages and time periods, comparing it with the state-of-the-art SMT approach.

An automatic evaluation showed that our approach improved the results achieved by the SMT approach on one dataset. Results were not statistically different than the SMT ones for the other two datasets. Additionally, we conducted a human evaluation for the Spanish dataset. This evaluation involved 4 scholars specialized in classical Spanish literature. Its results correlated with the automatic evaluation.

Finally, we conducted a user study to evaluate whether modernization—both SMT and NMT approaches—was able to decrease the difficulty of comprehending historical documents and, thus, increase their accessibility to a broader audience. 42 volunteers, of different age and background, participated in this study. The study was conducted using the same Spanish subset than for the human evaluation. Results showed that modernization successfully decreased the comprehension difficulty. In most of the cases, users chose the modernized version as the easiest to read and comprehend. However, there is still room for improvement. Sometimes, the modernization introduced errors that made users feel that the meaning had been changed. Other times, users did not find any significant difference between the original version and its modernization. When comparing the SMT and NMT approaches, the NMT approach made a bigger number of errors and the user chose its modernized versions as the best option fewer times than with the SMT approach.

While results showed that modernization had successfully improved the understanding of historical documents, we have to take into consideration that language-related losses may appear during the process (e.g., Fig. 1 shows an example in which part of the language structures and rhymes disappear). However, the goal of modernization is limited to

Original version: Huuolo de conceder don Quixote, y assi lo hizo.

Modernized version: Huéolo de conceder don Quijote, y así lo hizo.

Figure 4: Example of modernization in which the modernized version is similar to the original version.

Original version: Ofreciosele el gallardo pastor, pidiole que se viniessse con el a sus tiendas;

Modernized version: Se le rosó el gallardo pastor, pile dio que se viniese con él a sus tiendas;

Figure 5: Example of modernization in which users preferred the original version over the modernized one.

Original version: Que me maten si los encantadores que me persiguen no quieren enredarme en ellas, y detener mi camino, como en vengança de la riguridad que con Altissidora he tenido.

Modernized version: – Con mucho gusto?

Figure 6: Example of modernization in which the modern version differs with respect to the original.

bringing understanding of historical documents to a general audience.

As a future work, we would like to tackle the main problems pointed out during the human evaluation and the user study. Mainly, punctuation, diacritical marks, the introduction of non-existent words and loosing part of the sentence. We would also like to conduct a new human evaluation involving more scholars and more languages and datasets, and a new user study for different languages and datasets. Finally, we would like to apply the field of interactive machine translation to modernization, in order to assist scholars to achieve an error-free modernization.

Acknowledgments

The authors wish to thank the anonymous reviewers for their careful reading and in-depth criticisms and suggestions. The research leading to these results has received funding from the European Union through *Programa Operativo del Fondo Europeo de Desarrollo Regional (FEDER)* from Comunitat Valenciana (2014–2020) under project *Sistemas de fabricación inteligentes para la industria 4.0* (grant agreement IDIFEDER/2018/025); from Ministerio de Economía y Competitividad (MINECO) under project *MISMIS-FAKEHATE* (grant agreement PGC2018-096212-B-C31); from Fundación BBVA under project *Carabela* (grant agreement CARABELA); and from Generalitat Valenciana (GVA) un-

der project *DeepPattern* (grant agreement PROMETEO/2019/121). We gratefully acknowledge the support of NVIDIA Corporation with the donation of a GPU used for part of this research, and Andrés Trapiello and Ediciones Destino for granting us permission to use their book in our research. Additionally, we would like to thank all the volunteers that took part in the user study, and the scholars from *Prolope* that took part in the human evaluation.

References

- Bahdanau, D., Cho, K., Bengio, Y., 2015. Neural machine translation by jointly learning to align and translate. *arXiv:1409.0473*.
- Baron, A., Rayson, P., 2008. VARD2: A tool for dealing with spelling variation in historical corpora. *Postgraduate conference in corpus linguistics* .
- Baron, A., Rayson, P., Archer, D., 2009. Word frequency and key word statistics in corpus linguistics. *Anglistik* 20, 41–67.
- Biçici, E., Yuret, D., 2015. Optimizing instance selection for statistical machine translation with feature decay algorithms. *IEEE/ACM Transactions on Audio, Speech and Language Processing* 23, 339–350.
- Bollmann, M., Søgaard, A., 2016. Improving historical spelling normalization with bi-directional lstms

- and multi-task learning, in: Proceedings of the International Conference on the Computational Linguistics, pp. 131–139.
- Brown, P.F., Pietra, V.J.D., Pietra, S.A.D., Mercer, R.L., 1993. The mathematics of statistical machine translation: Parameter estimation. *Computational Linguistics* 19, 263–311.
- Crowther, J., 2003. No Fear Shakespeare: Romeo and Juliet. SparkNotes.
- Domingo, M., Casacuberta, F., 2018. A machine translation approach for modernizing historical documents using back translation, in: Proceedings of the International Workshop on Spoken Language Translation, pp. 39–47.
- Domingo, M., Chinea-Rios, M., Casacuberta, F., 2017. Historical documents modernization. *The Prague Bulletin of Mathematical Linguistics* 108, 295–306.
- Ernst-Gerlach, A., Fuhr, N., 2006. Generating search term variants for text collections with historic spellings, in: European Conference on Information Retrieval, pp. 49–60.
- Fiebranz, R., Lindberg, E., Lindström, J., Ågren, M., 2011. Making verbs count: the research project ‘gender and work’ and its methodology. *Scandinavian Economic History Review* 59, 273–293.
- Flood, A., 2015. Modern version of Don Quixote declared ‘crime against literature’. <https://www.theguardian.com/books/2015/aug/19/modern-version-of-don-quixote-declared-against-literature>.
- Gers, F.A., Schmidhuber, J., Cummins, F., 2000. Learning to forget: Continual prediction with LSTM. *Neural computation* 12, 2451–2471.
- Given, M.D., 2015. A discussion of bible translations and biblical scholarship. <http://courses.missouristate.edu/markgiven/re1102/bt.htm>.
- Hämäläinen, M., Säily, T., Rueter, J., Tiedemann, J., Mäkelä, E., 2018. Normalizing early english letters to present-day english spelling, in: Proceedings of the Workshop on Computational Linguistics for Cultural Heritage, Social Sciences, Humanities and Literature, pp. 87–96.
- Hochreiter, S., Schmidhuber, J., 1997. Long short-term memory. *Neural computation* 9, 1735–1780.
- Hundt, M., Denison, D., Schneider, G., 2011. Retrieving relatives from historical data. *Literary and Linguistic Computing* 27, 3–16.
- Ilonka, K., 2018. Is it possible for machines to translate poetry, when humans can barely do it? <https://electricliterature.com/is-it-possible-for-machines-to-translate-poetry-when-humans-can-barely-do-it/>.
- Kingma, D.P., Ba, J., 2014. Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 .
- Klein, G., Kim, Y., Deng, Y., Senellart, J., Rush, A.M., 2017. OpenNMT: Open-Source Toolkit for Neural Machine Translation, in: Proceedings of the Association for Computational Linguistics: System Demonstration, pp. 67–72.
- Koehn, P., Hoang, H., Birch, A., Callison-Burch, C., Federico, M., Bertoldi, N., Cowan, B., Shen, W., Moran, C., Zens, R., Dyer, C., Bojar, O., Constantin, A., Herbst, E., 2007. Moses: Open source toolkit for statistical machine translation, in: Proceedings of the Annual Meeting of the Association for Computational Linguistics, pp. 177–180.
- Koehn, P., Och, F.J., Marcu, D., 2003. Statistical phrase-based translation, in: Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics on Human Language Technology, pp. 48–54.
- Laing, M., 1993. The linguistic analysis of medieval vernacular texts: Two projects at Edinburgh’, in: *Corpora across the Centuries: Proceedings of the First International Colloquium on English Diachronic Corpora*, edited by M. Rissanen, M. Kytd, and S. Wright. St Catharine’s College Cambridge, pp. 121–141.

- Lison, P., Tiedemann, J., 2016. Opensubtitles2016: Extracting large parallel corpora from movie and tv subtitles, in: Proceedings of the International Conference on Language Resources Association, pp. 923–929.
- Monk, C., 2018. Custumale Roffense: An overview of the thirteenth-century custumal of Rochester Cathedral priory. <https://www.themedievalmonk.com/the-rochester-customs-book.html>.
- Och, F.J., 2003. Minimum error rate training in statistical machine translation, in: Proceedings of the Annual Meeting of the Association for Computational Linguistics, pp. 160–167.
- Och, F.J., Ney, H., 2002. Discriminative training and maximum entropy models for statistical machine translation, in: Proceedings of the Annual Meeting of the Association for Computational Linguistics, pp. 295–302.
- Papineni, K., Roukos, S., Ward, T., Zhu, W.J., 2002. BLEU: a method for automatic evaluation of machine translation, in: Proceedings of the Annual Meeting of the Association for Computational Linguistics, pp. 311–318.
- Pascanu, R., Gulcehre, C., Cho, K., Bengio, Y., 2013. How to construct deep recurrent neural networks. arXiv preprint arXiv:1312.6026 .
- Pettersson, E., Megyesi, B., Tiedemann, J., 2013. An SMT approach to automatic annotation of historical text, in: Proceedings of the workshop on computational historical linguistics at NODALIDA 2013; May 22-24; 2013; Oslo; Norway. NEALT Proceedings Series 18, pp. 54–69.
- Poncelas, A., Shterionov, D., Way, A., Maillette de Buy Wenniger, G., Passban, P., 2018. Investigation backtranslation in neural machine translation, in: Proceedings of the Annual Conference of the European Association for Machine Translation, pp. 249–258.
- Porta, J., Sancho, J.L., Gómez, J., 2013. Edit transducers for spelling variation in old spanish, in: Proceedings of the workshop on computational historical linguistics, pp. 70–79.
- Post, M., 2018. A call for clarity in reporting bleu scores, in: Proceedings of the Third Conference on Machine Translation, pp. 186–191.
- Rajvanshi, T., 2015. The poetics of poetry translation. <https://unravellingmag.com/articles/poetics-poetry-translation/>.
- Riezler, S., Maxwell, J.T., 2005. On some pitfalls in automatic evaluation and significance testing for mt, in: Proceedings of the workshop on intrinsic and extrinsic evaluation measures for machine translation and/or summarization, pp. 57–64.
- Rodríguez Marcos, J., 2015. Un ‘Quijote’ moderno. https://elpais.com/cultura/2015/05/27/babelia/1432726379_211033.html.
- Rogers, H.J., Willett, P., 1991. Searching for historical word forms in text databases using spelling-correction methods: Reverse error and phonetic coding methods. *Journal of Documentation* 47, 333–353.
- Schuster, M., Paliwal, K.K., 1997. Bidirectional recurrent neural networks. *IEEE Transactions on Signal Processing* 45, 2673–2681.
- Sen, S., Hasanuzzaman, M., Ekbal, A., Bhattacharyya, P., Way, A., 2019. Take help from elder brother: Old to modern english nmt with phrase pair feedback, in: Proceedings of the International Conference on Computational Linguistics and Intelligent Text Processing. In press.
- Sennrich, R., Haddow, B., Birch, A., 2015. Improving neural machine translation models with monolingual data. arXiv preprint arXiv:1511.06709 .
- Sennrich, R., Haddow, B., Birch, A., 2016. Neural machine translation of rare words with subword units, in: Proceedings of the Annual Meeting of the Association for Computational Linguistics, pp. 1715–1725.

- Snover, M., Dorr, B., Schwartz, R., Micciulla, L., Makhoul, J., 2006. A study of translation edit rate with targeted human annotation, in: Proceedings of the Association for Machine Translation in the Americas, pp. 223–231.
- Stolcke, A., 2002. SRILM - an extensible language modeling toolkit, in: Proceedings of the International Conference on Spoken Language Processing, pp. 257–286.
- Sutskever, I., Vinyals, O., Le, Q.V., 2014. Sequence to sequence learning with neural networks, in: Proceedings of the Advances in Neural Information Processing Systems, pp. 3104–3112.
- Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., Rabinovich, A., 2015. Going deeper with convolutions, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 1–9.
- Tjong Kim Sang, E., Bollmann, M., Boschker, R., Casacuberta, F., Dietz, F., Dipper, S., Domingo, M., van der Goot, R., van Koppen, M., Ljubešić, N., Östling, R., Petran, F., Pettersson, E., Scherrer, Y., Schraagen, M., Sevens, L., Tiedemann, J., Vanallemeersch, T., Zervanou, K., 2017. The CLIN27 shared task: Translating historical text to contemporary language for improving automatic linguistic annotation. *Computational Linguistics in the Netherlands Journal* 7, 53–64.
- Toselli, A.H., Leiva, L.A., Bordes-Cabrera, I., Hernández-Tornero, C., Bosch, V., Vidal, E., 2017. Transcribing a 17th-century botanical manuscript: Longitudinal evaluation of document layout detection and interactive transcription. *Digital Scholarship in the Humanities* 33, 173–202.
- Toselli, A.H., Romero, V., Pastor, M., Vidal, E., 2010. Multimodal interactive transcription of text images. *Pattern Recognition* 43, 1814–1825.
- Trapiello, A., 2015. Don Quijote de la Mancha Puesto en castellano actual íntegra y fielmente por Andrés Trapiello. Ediciones Destino.
- Wu, Y., Schuster, M., Chen, Z., Le, Q.V., Norouzi, M., Macherey, W., Krikun, M., Cao, Y., Gao, Q., Macherey, K., Klingner, J., Shah, A., Johnson, M., Liu, X., Kaiser, L., Gouws, S., Kato, Y., Kudo, T., Kazawa, H., Stevens, K., Kurian, G., Patil, N., Wang, W., Young, C., Smith, J., Riesa, J., Rudnick, A., Vinyals, O., Corrado, G., Hughes, M., Dean, J., 2016. Google’s neural machine translation system: Bridging the gap between human and machine translation. *arXiv:1609.08144*.
- Zens, R., Och, F.J., Ney, H., 2002. Phrase-based statistical machine translation, in: Proceedings of the Annual German Conference on Advances in Artificial Intelligence, pp. 18–32.