

# Preparing an Endangered Language for the Digital Age: The Case of Judeo-Spanish

Alp Öktem<sup>1</sup>, Rodolfo Zevallos<sup>1,2</sup>, Yasmin Moslem<sup>3</sup>, Güneş Öztürk<sup>1</sup>, Karen Şarhon<sup>4</sup>

<sup>1</sup>CollectivaT, Barcelona, Spain <sup>2</sup>Universitat Pompeu Fabra, Barcelona, Spain

<sup>3</sup>Dublin City University, Dublin, Ireland <sup>4</sup>Sephardic Center of Istanbul, Istanbul, Turkey

alp@collectivat.cat, rodolfojoel.zevallos@upf.edu, yasmin.moslem@adaptcentre.ie

ozgurgunes@collectivat.cat, karensarhon@gmail.com

## Abstract

We develop machine translation and speech synthesis systems to complement the efforts of revitalizing Judeo-Spanish, the exiled language of Sephardic Jews, which survived for centuries, but now faces the threat of extinction in the digital age. Building on resources created by the Sephardic community of Turkey and elsewhere, we create corpora and tools that would help preserve this language for future generations. For machine translation, we first develop a Spanish to Judeo-Spanish rule-based machine translation system, in order to generate large volumes of synthetic parallel data in the relevant language pairs: Turkish, English and Spanish. Then, we train baseline neural machine translation engines using this synthetic data and authentic parallel data created from translations by the Sephardic community. For text-to-speech synthesis, we present a 3.5 hour single speaker speech corpus for building a neural speech synthesis engine. Resources, model weights and online inference engines are shared publicly.

**Keywords:** Extremely low-resource language, Machine Translation, Data-augmentation, Text-to-Speech, Judeo-Spanish

## 1. Introduction

In this paper, we present our ongoing language technology-related efforts for preparing Judeo-Spanish to the digital age. We embark upon creating open language corpora and tools that would serve for language documentation, assisting language learners and development of advanced applications. We focus on two main tools, machine translation (MT) and text-to-speech synthesis (TTS). In our extremely low-resource setup, we get use of Judeo-Spanish’s proximity to Spanish by using transfer learning methodologies. For MT, we build a rule-based machine translation engine that allows us to convert Spanish text to Judeo-Spanish. Using this system, we create large synthetic pre-training data from publicly available English, Turkish and Spanish parallel corpora and train neural machine translation systems. For TTS, we do transfer learning from pre-trained Spanish and English engines using a small single-speaker speech corpus. During the development of these tools, we have packaged various types of raw resources into training-ready language data and models and shared them in our project’s data portal *Ladino Data Hub*<sup>1</sup>. The complete list of output of this work can be presented as follows:

1. A monolingual news corpus,
2. Authentic and synthetic parallel corpora in English, Spanish and Turkish paired with Judeo-Spanish,
3. A Spanish to Judeo-Spanish rule-based machine translation system,

4. Neural machine translation models between Judeo-Spanish and English, Spanish and Turkish,
5. A 3.5 hour single speaker speech corpus,
6. Neural network-based speech synthesis model,
7. Web application for MT and TTS<sup>2</sup>.

## 2. Background

Judeo-Spanish, also referred to as Ladino or Judezmo (ISO 639-3 *lad*), is a descendant of old Castilian Spanish from the 15th century (Sefardiweb del CSIC, 2022). It is the historical and predominant language of the Sephardic Jews, who were expelled from their homes by the Spanish Inquisition (1492) and welcomed into the Ottoman Empire, where they retained the language, as well as France, Italy, the Netherlands, Morocco and England, where they shifted to the dominant language. It has traces of numerous Iberian languages of the 15th century like Old Aragonese, Astur-Leonese, Old Catalan, Galician-Portuguese and Mozarabic with Castilian Spanish forming its basis vocabulary (Minervini, 2006). After 530 years, Judeo-Spanish still survives as a language of Ottoman Sephardic Jews in more than 30 countries, with most speakers residing in Israel. Although it has survived and evolved over the centuries, it is currently classified as a severely endangered language by UNESCO (Moseley, 2010).

The digital age has a direct effect on endangered languages like Judeo-Spanish. There is currently a growing digital divide between languages with sufficient resources and languages with fewer resources, further

<sup>1</sup><http://data.sefarad.com.tr>

<sup>2</sup><http://translate.sefarad.com.tr>

exacerbating the danger of digital extinction for them (Kornai, 2013). For the dominant languages the process of generating artificial intelligence tools is much easier due to their large web-presence. However, many marginalized languages do not have sufficient material and human resources to power the creation of such tools. Lack of state support, public visibility, as well as societal and institutional oppression are direct causes of these languages being deprioritized in the digital spaces of today (V et al., 2020).

The Sephardic community of Turkey has been active in promoting their language heritage in many ways. These include: publishing the only newspaper in the world entirely in Judeo-Spanish *El Amaneser*, giving language lessons, writing and performing plays in Judeo-Spanish, creating language learning content, collecting speech corpora and publishing dictionaries, music albums and books.

The aim of this work is to build data-centric technology for Judeo-Spanish for it to gain digital ground. Besides building new and compiling already existing corpora for this purpose, we create first machine translation and text-to-speech synthesis systems for the language. Machine translation makes it possible for the language to be interpretable by non-speakers and is also proposed as a way of language documentation (Bird and Chiang, 2012). It is now also considered as an attractive tool for many language learners in addition to dictionaries and thesauri (Clifford et al., 2013). Even though it is difficult to obtain high performance in low resource settings, it has been used to strike interest in language and collect translations and corrections from the community. The second language tool we focus on, text-to-speech synthesis (TTS), makes it possible building of tools like virtual assistants and screen readers. In the context of language learning, one can learn how a certain word or sentence is pronounced in a language without the help of an instructor or a speaker.

### 3. Judeo-Spanish Resources

We explain our various data compilation efforts in this section. All data presented are published with *CC BY-SA 4.0 license*<sup>3</sup> on Ladino Data Hub. We also provide the scripts we have used in developing these resources with GPL-licenses for facilitating expansion and reproducibility<sup>4</sup>.

#### 3.1. Monolingual text corpus

Text corpora have been used both in language technology and in linguistic research. They are an essential part of creating statistical language models that are used in applications such as optical character recognition, handwriting recognition, machine translation and spelling correction.

<sup>3</sup><http://creativecommons.org/licenses/by-sa/4.0/>

<sup>4</sup><http://github.com/CollectivaT-dev/judeo-espanyol-resources>

For this task, we automatically scraped the articles published in the weekly online newspaper *Şalom*<sup>5</sup>. As of now, we have collected 397 articles totaling to 176, 843 words.

#### 3.2. Parallel corpus

The type of data that is needed to build a MT system is parallel data, which consists of a collection of sentences in a language together with their translations. We have only detected two publicly available corpora of Judeo-Spanish in the commonly used OPUS portal<sup>6</sup>: Wikimedia corpus consisting of 18 sentences and Tatoeba corpus of 872 sentences.

In order to expand on this set, we gathered translations made by the Sephardic Center of Istanbul. These covered topics like news articles, online shop strings, recipes and cultural event announcements. We automatically segmented the text into sentences getting use of punctuation and then manually verified alignments. We also digitized the language learning material *Fraza del dia*<sup>7</sup>, where daily a Judeo-Spanish phrase is presented with their translations in another language. The sizes of parallel corpora created for each language pair is listed in Table 1.

Language pair (Judeo-Spanish and)	#Sentences	Total #tokens
English	3333	41,508
Spanish	977	12,712
Turkish	845	15,781

Table 1: Parallel data compiled from Tatoeba and translations by Sephardic community.

#### 3.3. Spanish Judeo-Spanish Dictionary

We developed a digital Spanish–Judeo-Spanish dictionary from the sources listed in Table 2. To process the dictionaries shown in Table 2 which were in PDF format, we used the Python programming language, where we aligned the Spanish word with Judeo-Spanish word and eliminated irrelevant information like example sentences. Once the dictionaries were processed, the data were stored in a plain text file under the following structure:  $\langle word\text{-spanish}, word\text{-judeospanish} \rangle$ .

#### 3.4. Single-speaker speech corpus

We built a single-speaker speech corpus of 3 hours and 24 minutes to be used in the creation of Judeo-Spanish TTS system. We had our native Judeo-Spanish speaking author read 30 articles from the weekly newspaper *El Amaneser*. The articles are about different topics, ranging from historical issues, current affairs, cultural events and politics. The recordings had an

<sup>5</sup><http://www.salom.com.tr>

<sup>6</sup><http://opus.nlpl.eu/>

<sup>7</sup><http://sefarad.com.tr/judeo-espanyolladino/frazadeldia/>

Dictionary	# Entries
Diksionario de Ladino a Espanyol (Güler and Tinoco, 2003)	2523
Diksionario de Djudeo-Espanyol a Castellano (Orgun and Tinoco, 2009)	4215

Table 2: Dictionaries used for the construction of the digital dictionary.

average length of 6 minutes. To obtain TTS training data material, we had to divide the audios into smaller segments. For this task, we developed an automatic aligner<sup>8</sup> based on Coqui Speech-to-text (Coqui, 2022). The pre-trained Spanish model performed well enough to optimize the process. Nevertheless, to ensure completely matching audio and transcription pairs, we manually verified each pair and performed corrections where needed. The resulting corpus consists of 1987 16-bit, single-channel WAV audio files sampled at 16kHz with their transcriptions.

## 4. Machine Translation

In this section, we present our experiments for Judeo-Spanish machine translation. To account for the lack of data, we first build a rule-based Spanish to Judeo-Spanish translator and then use that to obtain the data needed to train neural baseline models.

### 4.1. Rule-Based machine translation

In the following, we describe the procedure of our rule-based machine translation system from Spanish to Judeo-Spanish based on the dictionaries available in Table 2. The Python-based scripts and documentation are provided with GNU General Public License in our Github repository<sup>9</sup>.

The first step in the translation process is to tokenize the input Spanish phrase, for which we use the Python library Stanza<sup>10</sup> (Qi et al., 2020). This library, in addition to tokenizing the phrase, obtains the part-of-speech (POS) and lemmas of each token. As a second step, each token is looked up in the Spanish–Judeo-Spanish dictionary. If the token is found in the dictionary, its corresponding Judeo-Spanish token is obtained, otherwise, the dictionary is searched for its lemmatized form of the token. If the lemmatized token is found

<sup>8</sup>[https://github.com/CollectivaT-dev/Judeo-Spanish\\_STT](https://github.com/CollectivaT-dev/Judeo-Spanish_STT)

<sup>9</sup><https://github.com/CollectivaT-dev/Espanyol-Ladino-Translation>

<sup>10</sup>Stanza is a collection of tools for the linguistic analysis (Tokenization, Part-of-Speech, Lemmatization, etc.) of many human languages, including Spanish. <https://stanfordnlp.github.io/stanza/>

in the dictionary, the corresponding Judeo-Spanish token is obtained and is conjugated according to its POS. Our method transforms a Spanish token to a conjugated Judeo-Spanish form using an algorithm based on conjugation rules specified in (Perahya, 2012). We also convert the verb form Present Perfect (e.g. spa.“he cocinado”) to past indefinite (e.g. spa.“cociné” lad. “gizi”) as the former form is not common in Judeo-Spanish. In case the lemmatized token is not found in the dictionary, it is processed by a Judeo-Spanish correction method, which follows established orthographic rules of the language<sup>11</sup>. Finally, in step three, phrase forms that do not exist in Judeo-Spanish are corrected into their right form using a phrase correction dictionary. For example, “*tengo ke*” (from spa.“*tengo que*” eng.“*I have to*”) is corrected to “*debo de*”, or “*ay ke*” (from spa.“*hay que*” eng.“*one must*”) is corrected to “*Kale*”. Some example translations are listed in Table 3. Automatic evaluation results are presented in Table 5.

### 4.2. Data augmentation

We introduce a data augmentation method based on creating synthetic parallel data using the rule-based MT system presented in Section 4.1. We first collect publicly available parallel data in pairs English-Spanish and Turkish-Spanish from the OPUS collection. Then, we translate the Spanish portions into Judeo-Spanish using the rule-based MT. This yields Turkish–Judeo-Spanish and English–Judeo-Spanish synthetic parallel data. Finally, the Spanish portions of two sets are then merged to create Spanish–Judeo-Spanish synthetic parallel data. The statistics and sources for synthetic data augmentation are listed in Table 4.

### 4.3. Neural machine translation

We used the OpenNMT-py toolkit (Klein et al., 2018) to train the models. The model consists of an eight-head Transformer “big” (Vaswani et al., 2017a) with six-layer hidden units of 512 unit size. It uses Relative Position Representations (Shaw et al., 2018) with a clipping distance  $k=16$ . A token-batch size of 1,024 was selected. Adam optimizer (Kingma and Ba, 2015) was selected with 4,000 warm-up steps. Trainings were performed until no further improvement was recorded in development set perplexity in the last five validations.

We used the synthetic parallel data we created as training data. As for development and test sets, we used the authentic data mixes presented in Section 3.2. As English portion was about three times larger than Spanish and Turkish, we used a commercial machine translation engine to translate the extra data available for

<sup>11</sup>Orthographic structure of Judeo-Spanish compared to Spanish available in [https://github.com/CollectivaT-dev/judeo-espanyol-resources/blob/main/resources/Gramatica\\_Ladino.doc](https://github.com/CollectivaT-dev/judeo-espanyol-resources/blob/main/resources/Gramatica_Ladino.doc)

Spanish input	Judeo-Spanish translation
Me gusta leer.	Me plaze meldar.
¿No has leído el libro?	No meldates el livro?
Bebo café turco después del almuerzo.	Bevo kafe turko despues del komida de midi.
Tengo dos niños; una hija y un hijo.	Tengo dos kriaturas; una ija i un ijo.
Tengo que cocinar para mañana.	Devo de gizar para amanyana.

Table 3: Example translations obtained with the rule-based machine translation system.

ENG-SPA	#sentences
Books	93,470
Europarl	615,626
News-commentary	49,089
OpenSubtitles	4,652,910
SciELO	164,500
TED2013	157,895
WMT-News	14,522
<b>TOTAL ENG-SPA</b>	<b>5,748,012</b>
SPA-TUR	
EUBookshop	19,914
GlobalVoices	7,461
OpenSubtitles	4,000,000
TED2020	370,465
Tatoeba	28,829
WikiMatrix	147,352
<b>TOTAL SPA-TUR</b>	<b>4,574,021</b>
<b>TOTAL SPA</b>	<b>10,322,033</b>

Table 4: Publicly available parallel data used for synthetic data creation.

English to Spanish and Turkish and added them to the mixes. Finally, we reserved 500 sentences from Spanish and Turkish and 750 sentences from English mix as test data and used the rest as validation data during training. Development, test sets, training configuration files, subword models and training logs are provided for reproducibility<sup>12</sup>. Model weights are made available in Ladino Data Hub.

	ENG	SPA	TUR
<b>LAD</b> → <i>lang</i>	34.96	47.13	20.14
<i>lang</i> → <b>LAD</b>	26.03	44.85	21.03
<b>Rule-based SPA</b> → <b>LAD</b>	-	45.80	-

Table 5: Automatic evaluation results in 6 language directions and also on rule-based system. BLEU scores were calculated on lowercased output and reference with SACREBLEU toolkit with Moses tokenizer (Post, 2018).

We report our test set BLEU-scores (Papineni et al., 2002) for each translation direction in Table 5. As future work, we will also perform human evaluations on

<sup>12</sup>[https://github.com/CollectivaT-dev/judeo-espanyol-resources/tree/main/MT\\_devtest\\_configs](https://github.com/CollectivaT-dev/judeo-espanyol-resources/tree/main/MT_devtest_configs)

additional data to have a fairer judgment of the translation qualities.

## 5. Text-to-Speech

In this section, we present our experiments for the development of a Judeo-Spanish speech synthesizer. We use Glow-TTS model (Kim et al., 2020) for our experiments and the Griffin-Lim algorithm (Griffin and Lim, 1984) to avoid using vocoder, which we intend to develop for future research. We trained three Glow-TTS models with our 3.5 hour dataset. The first model was trained from scratch; the second and third, by fine-tuning English and Spanish TTS models. For all experiments, we do not enable the use of phonemes or the phonemizer as in Paniv (2021). We follow the settings for the mel-spectrogram of Prenger et al. (2019).

**Training Judeo-Spanish from scratch** First, we evaluate the performance of the model trained only using our single-speaker dataset. During training, like Kim et al. (2020) we set the standard deviation to 1. Our model was trained for 5,000 iterations with a batch size of 32, using the Adam optimizer (Kingma and Ba, 2014) with the Noam learning rate program (Vaswani et al., 2017b). This required only 4 days with an 8GB NVIDIA GPU.

**Fine-Tuning from English and Spanish** To train the Glow-TTS model from the pre-trained English (ljspeech/glow-tts) and Spanish (mai/tacotron2-DDC) models (Coqui, 2022), we used the same setup as the Glow-TTS model trained only with Judeo-Spanish data, but added the use of the phonemes and phonemizer corresponding to each language from the pre-trained models. Each model required only 4 days with an 8GB NVIDIA GPU.

### 5.1. Evaluation

Our best model was the one where fine-tuning was applied to the pre-trained English model, achieving a much better intelligibility and naturalness than the other models, reducing even more the “metallic” sound that appears in some consonants. On the other hand, the Spanish fine-tuned model did achieve a better naturalness for Judeo-Spanish phonemes but did not achieve a good intelligibility, perhaps due to the amount of training data. Likewise, the from-scratch model did achieve an excellent naturalness in the phonemes but a very poor intelligibility.

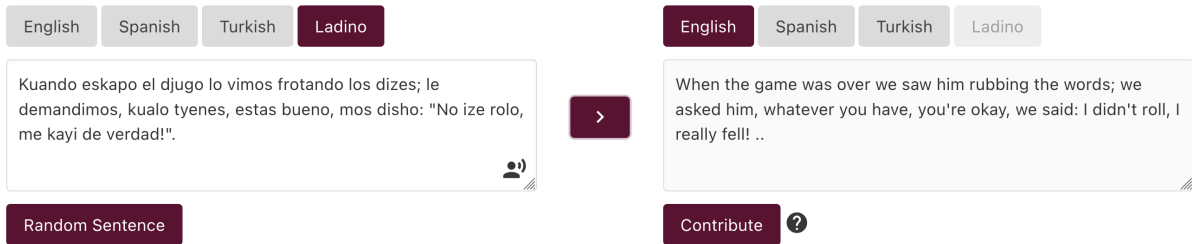


Figure 1: Web application for MT and TTS available in <http://translate.sefarad.com.tr>

We selected the best performing model (fine-tuned from English) for human evaluation. We used Mean opinion score (MOS) of intelligibility and naturalness with a 5-point scale: 5 for excellent, 4 for good, 3 for fair, 2 for poor and 1 for bad. An evaluation survey consisting of ten out-of-corpus samples were published in a closed Ladino speaker community of Istanbul and 12 native speakers participated. The average scores among ten samples are listed in Table 6. The configurations of all pre-trained models as well as audio samples are made available online for reproducibility<sup>13</sup>. Model weights are shared in Ladino Data Hub.

Model	Intelligibility	Naturalness
Glow-TTS (f.t. on English)	4.04	3.61

Table 6: Judeo-Spanish text-to-speech system evaluation results for intelligibility and naturalness (MOS)

## 6. Web Application

Our web application for serving the machine translation and speech synthesis systems can be seen in Figure 5. It allows translation between English, Spanish, Turkish and Ladino and makes it possible to listen to synthesized Ladino text. For Spanish, we integrated the rule-based system translating to Ladino and our model translating to Spanish. For the rest of the translation directions, we chained open source OPUS-MT translation models (Tiedemann and Thottingal, 2020) to these two systems to get translation to and from English and Turkish.

We also added a participation feature to make Judeo-Spanish speakers be part of future developments. By clicking the "Contribute" button, users can correct the translations and then submit to our database to be stored as parallel data for future trainings.

## 7. Conclusion

In this work, we introduced baseline systems of machine translation and speech synthesis for Judeo-Spanish. First, we developed a rule-based machine

<sup>13</sup>[https://github.com/CollectivaT-dev/Ladino\\_TTS](https://github.com/CollectivaT-dev/Ladino_TTS)

translator from Spanish to Judeo-Spanish. This base translator was used to apply a data augmentation technique. Second, we developed three bidirectional machine translation models between Judeo-Spanish and Spanish, Turkish and English, being the first neural-based systems for this language. Although some of our models do not perform optimally, we believe that this work is the basis for future research regarding this language, as well as motivating research for extremely low-resourced languages using data augmentation strategies. Third, we developed speech synthesis models for Judeo-Spanish, achieving an acceptable result by fine-tuning on an English model. Data, model checkpoints, development and test sets and configuration files are shared openly on project's data portal Ladino Data Hub and our Github repository. Finally, we created a web-application for machine translation with voice to help language learners, researchers and linguists who want to study Judeo-Spanish.

## 8. Acknowledgements

This paper was written as part of the project "Judeo-Spanish: Connecting the two ends of the Mediterranean" carried out by CollectivaT and Sephardic Center of Istanbul within the framework of the "Grant Scheme for Common Cultural Heritage: Preservation and Dialogue between Turkey and the EU (CCH-II)" implemented by the Ministry of Culture and Tourism of the Republic of Türkiye with the financial support of the European Union. The content of this paper is the sole responsibility of the authors and does not necessarily reflect the views of the European Union.

Yasmin Moslem contributed to this work while she was pursuing her PhD degree, supported by the Science Foundation Ireland Centre for Research Training in Digitally-Enhanced Reality (d-real) under Grant No. 18/CRT/6224, and the ADAPT Centre for Digital Content Technology which is funded under the Science Foundation Ireland (SFI) Research Centres Programme (Grant No. 13/RC/2106) and is co-funded under the European Regional Development Fund.

We would like to thank the excellent volunteer work by Brian Russell, who helped us digitize *Fraza del dia* content.

## 9. Bibliographical References

- Bird, S. and Chiang, D. (2012). Machine translation for language preservation. In *Proceedings of COLING 2012: Posters*, pages 125–134, Mumbai, India, December. The COLING 2012 Organizing Committee.
- Clifford, J., Merschel, L., and Munné, J. (2013). Surveying the landscape: What is the role of machine translation in language learning? @*tic. revista d'innovació educativa*, (10).
- Coqui. (2022). Coqui TTS. <https://github.com/coqui-ai/TTS>.
- ∇, Nekoto, W., Marivate, V., Matsila, T., Fasubaa, T., Kolawole, T., Fagbohunge, T., Akinola, S. O., Muhammad, S. H., Kabongo, S., Osei, S., et al. (2020). Participatory research for low-resourced machine translation: A case study in African languages. *Findings of EMNLP*.
- Griffin, D. and Lim, J. (1984). Signal estimation from modified short-time fourier transform. *IEEE Transactions on acoustics, speech, and signal processing*, 32(2):236–243.
- Güler, P. and Tinoco. (2003). Dikcionario de ladino a espanyol. In *Dikcionario de Ladinokomunita*.
- Kim, J., Kim, S., Kong, J., and Yoon, S. (2020). Glowtts: A generative flow for text-to-speech via monotonic alignment search. *Advances in Neural Information Processing Systems*, 33:8067–8077.
- Kingma, D. P. and Ba, J. (2014). Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Kingma, D. P. and Ba, J. (2015). Adam: A method for stochastic optimization. In Yoshua Bengio et al., editors, *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Klein, G., Kim, Y., Deng, Y., Nguyen, V., Senellart, J., and Rush, A. (2018). OpenNMT: Neural machine translation toolkit. In *Proceedings of the 13th Conference of the Association for Machine Translation in the Americas (Volume 1: Research Track)*, pages 177–184, Boston, MA, March. Association for Machine Translation in the Americas.
- Kornai, A. (2013). Digital language death. *PLOS ONE*, 8(10):1–11, 10.
- Minervini, L. (2006). El desarrollo histórico del judeoespañol. *Revista Internacional de Lingüística Iberoamericana*, 4(2 (8)):13–34.
- Moseley, C. (2010). *Atlas of the World's Languages in Danger*. UNESCO Publishing, 3 edition.
- Orgun, P. and Tinoco. (2009). Dikcionario de djudeoespanyol a castellano. In *Dikcionario de Ladinokomunita*.
- Paniv, Y. (2021). Ukrainian TTS (text-to-speech) using coqui TTS. <https://github.com/robinhad/ukrainian-tts>.
- Papineni, K., Roukos, S., Ward, T., and Zhu, W.-J. (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, Philadelphia, Pennsylvania, USA, July. Association for Computational Linguistics.
- Perahya, K. (2012). *DIKSYONARYO JUDEO ESPANYOL - TURKO LADINO - TÜRKÇE SÖZLÜK*. Sentro de Investigaciones sobre la Cultura Sefardi Otomana - Turka, 2 edition.
- Post, M. (2018). A call for clarity in reporting BLEU scores. In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 186–191, Brussels, Belgium, October. Association for Computational Linguistics.
- Prenger, R., Valle, R., and Catanzaro, B. (2019). Waveglow: A flow-based generative network for speech synthesis. In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 3617–3621. IEEE.
- Qi, P., Zhang, Y., Zhang, Y., Bolton, J., and Manning, C. D. (2020). Stanza: A Python natural language processing toolkit for many human languages. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*.
- Sefardiweb del CSIC. (2022). El Judeoespañol o Ladino. <http://www.proyectos.cchs.csic.es/sefardiweb/node/10>. Accessed: 2022-05-24.
- Shaw, P., Uszkoreit, J., and Vaswani, A. (2018). Self-attention with relative position representations. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 464–468, New Orleans, Louisiana, June. Association for Computational Linguistics.
- Tiedemann, J. and Thottingal, S. (2020). OPUS-MT — Building open translation services for the World. In *Proceedings of the 22nd Annual Conference of the European Association for Machine Translation (EAMT)*, Lisbon, Portugal.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., and Polosukhin, I. (2017a). Attention is all you need. In Isabelle Guyon, et al., editors, *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, 4-9 December 2017, Long Beach, CA, USA*, pages 5998–6008.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., and Polosukhin, I. (2017b). Attention is all you need. *Advances in neural information processing systems*, 30.