

Resolving Transcription Ambiguity in Spanish: A Hybrid Acoustic-Lexical System for Punctuation Restoration

Xiliang Zhu*, Chia-Tien Chang*, Shayna Gardiner, David Rossouw, Jonas Robertson

Dialpad Canada Inc.

{xzhu, karol.chang, sgardiner, davidr, jonas}@dialpad.com

Abstract

Punctuation restoration is a crucial step after Automatic Speech Recognition (ASR) systems to enhance transcript readability and facilitate subsequent NLP tasks. Nevertheless, conventional lexical-based approaches are inadequate for solving the punctuation restoration task in Spanish, where ambiguity can be often found between unpunctuated declaratives and questions. In this study, we propose a novel hybrid acoustic-lexical punctuation restoration system for Spanish transcription, which consolidates acoustic and lexical signals through a modular process. Our experiment results show that the proposed system can effectively improve F1 score of question marks and overall punctuation restoration on both public and internal Spanish conversational datasets. Additionally, benchmark comparison against LLMs (Large Language Model) indicates the superiority of our approach in accuracy, reliability and latency. Furthermore, we demonstrate that the Word Error Rate (WER) of the ASR module also benefits from our proposed system.

1 Introduction

Automatic Speech Recognition (ASR) systems are applied in a variety of industry applications such as voice assistance and conversation analysis. However, typical ASR systems avoid producing punctuation marks in the transcripts, which leads to poor readability and causes ambiguity in the context (Jones et al., 2003). Therefore, a post-processing step to restore punctuation marks in transcripts is critical for speech-based commercial products.

Lexical-based approaches have been extensively studied in punctuation restoration tasks (Păiș and Tufiş, 2021). One major advantage of using lexical features is the availability of a massive amount of text data that is often well punctuated, such as

Wikipedia. Most of the existing work on punctuation restoration focuses on English. Spanish is little studied, although it is the world’s second largest mother tongue and even has more native speakers than English. Although a handful of work has addressed Spanish punctuation restoration using BERT-based approaches in recent years (González-Docasal et al., 2021; Zhu et al., 2022a), one major challenge in restoring punctuation marks for languages like Spanish has not been fully tackled: the rich morphology of Spanish allows speakers to omit subject pronouns and order words in sentences more freely than in English, which forces Spanish speakers to rely more on prosodic features when distinguishing questions from declarative sentences. These characteristics present a unique challenge from an NLP perspective when written transcripts are the main source of information for models.

In order to address the challenges in predicting Spanish question marks and improve the overall punctuation restoration accuracy, we introduce a hybrid punctuation restoration system leveraging both acoustic and lexical signals for Spanish conversations. While previous work on multimodal methodologies often requires large-scale, parallel audio-text data (Klejch et al., 2017), or additional audio encoding and fusion steps (Sunkara et al., 2020), our approach employs the conventional modular ASR-NLP setup in industry applications with no additional computational cost. Moreover, our system allows independent training of ASR and NLP modules, eliminating the need for massive parallel training resources. The main contributions of this paper are as follows:

1. Evaluate the impact of including punctuation in Spanish ASR training data on Word Error Rate (WER).
2. Propose a hybrid system for Spanish punctuation restoration leveraging ASR and NLP sequentially.

*These authors contributed equally to this work

3. Demonstrate the effectiveness of our system by achieving up to a 2.1% relative reduction in Word Error Rate (WER) for the Spanish ASR decoder, improving question mark prediction F1 score by over 4% absolute, and consequently enhancing overall punctuation restoration accuracy on internal and public datasets from the Linguistic Data Consortium (LDC) (Graff et al., 2010a,b), also outperforming top LLMs (Large Language Model) in terms of accuracy, reliability and latency.

2 Background

2.1 Related Work

Punctuation restoration is often formulated as a sequence labeling task, where punctuation marks are predicted at appropriate positions in a sequence of words. Early studies used lexical-based methods such as n-gram language models (Gravano et al., 2009) and Conditional Random Fields (CRF) (Lu and Ng, 2010). More recently, long short-term memory (LSTM) (Hochreiter and Schmidhuber, 1997) and pre-trained large language models have been used (Xu et al., 2016; Devlin et al., 2019; Fu et al., 2021). Some works (Guan, 2020; O’Neill et al., 2021) proposed a speech recognition system with direct punctuation output, but it is unclear whether this approach is more effective than a traditional lexicon-based approach. For Spanish, a multilingual LSTM-based approach was studied in (Li and Lin, 2020). (Zhu et al., 2022a) proposed a transformer-based architecture with transfer learning to overcome the Spanish resource limitation and (González-Docasal et al., 2021) integrated silence embedding into BERT; however, as far as we are aware, no study has yet investigated the use of a hybrid approach incorporating acoustic input for the task of Spanish punctuation restoration.

Recent advances in LLMs such as ChatGPT¹, GPT4² and PaLM2³ have reshaped the approaches for many NLP tasks. (Qin et al., 2023) found that ChatGPT performs well on tasks favouring reasoning capabilities while still faces challenges in sequence tagging tasks. (Lai et al., 2023) studied the multilingual capability of ChatGPT and found that it shows less optimal performance compared to

task-specific models in different languages. However, the application of LLM in punctuation restoration task has not been studied yet to the best of our knowledge.

2.2 Ambiguity in Unpunctuated Spanish Text

Identifying Spanish questions from unpunctuated text is a challenging task. There are three relevant sociolinguistic features to consider for question identification in Spanish.

First, declarative sentences can occasionally become questions on intonation and context alone; e.g. *ustedes no pueden mandar un cheque con la orden* can be either a declarative or a question. This is true even for its English counterpart – both *Can’t you send a cheque with the order?* and *You can’t send a cheque with the order?* are well-formed – but the phenomenon is extremely common in Spanish (Brown and Rivas, 2011; Raymond, 2015; Cuza, 2016). In fact, in Caribbean Spanish, it is becoming increasingly more common to see questions like *¿ustedes no pueden mandar un cheque con la orden?* (*You can’t send a cheque with the order?*), which has typical declarative syntax, rather than *¿no pueden ustedes mandar un cheque con la orden?* (*Can’t you send a cheque with the order?*), which uses subject-verb inverted order (Brown and Rivas, 2011).

Second, Spanish morphosyntax also allows the reverse to occur: that is, declarative sentences can have subject-verb inversion too (Mackenzie, 2021), meaning that the question *no pueden ustedes mandar un cheque con la orden* above is also a perfectly well-formed declarative sentence.

Third, Spanish is a pro-drop language: due to richly-inflected morphology, it is possible to drop a subject pronoun entirely, using a verb’s suffix alone to identify its subject – and removing the possibility of subject-verb inversion. For instance, the above example could easily become *no pueden mandar un cheque con la orden* or *¿no pueden mandar un cheque con la orden?*. In a small survey of our own data, we reviewed 200 utterances, in which there were 180 questions, of which 125 (69%) were pro-dropped – leaving only 55 questions with a fully realized subject noun or pronoun.

These facts make subject-verb inversion a much less helpful tool for definitively identifying questions in Spanish than it is for English, which consequently limits the performance of lexical-based NLP models in the Spanish punctuation restoration task. We also know that Spanish speakers them-

¹<https://openai.com/blog/chatgpt>

²<https://openai.com/gpt-4>

³<https://blog.google/technology/ai/google-palm-2-ai-large-language-model/>

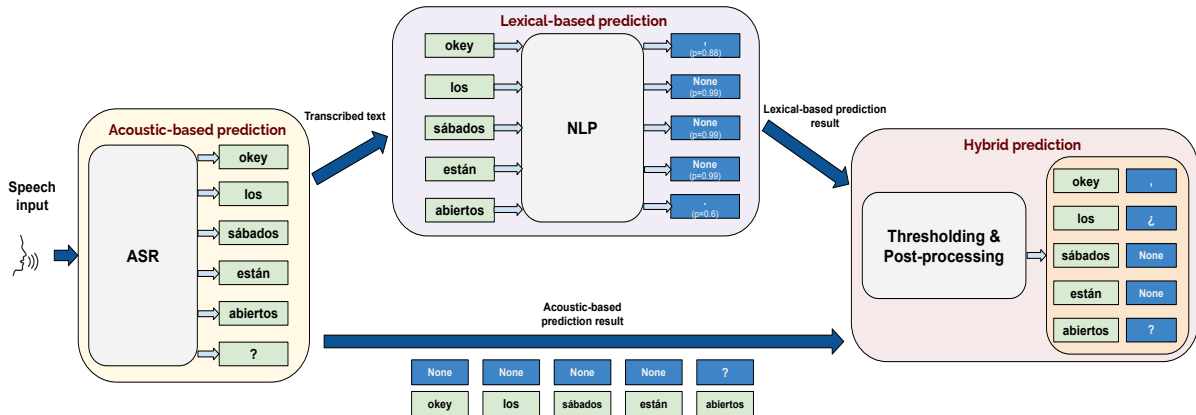


Figure 1: Overview of our hybrid punctuation restoration system, showing the example of an ambiguous unpunctuated utterance "okey los sábados están abiertos" (can be interpreted as "OK, they are open on Saturdays." or "OK, are they open on Saturdays?") processed as "Okey, ¿los sábados están abiertos?"

selves do not rely on lexical information alone to distinguish questions from declaratives: evidence suggests that acoustic features are measurably different when the speaker intends an utterance as a question rather than a declarative, and that this is true across many varieties of Spanish (Face, 2005; Willis, 2007; Lee and M.A., 2010; Armstrong, 2017).

3 Method

3.1 System Overview

Our hybrid punctuation restoration system is used in a Spanish call center product. In the real-time system pipeline, the audio from customer support phone calls is first transcribed by the ASR module, then text output is fed into the downstream NLP module. Instead of adding an extra acoustic encoder and combining it with a lexical encoder as proposed in (Sunkara et al., 2020; Zhu et al., 2022b), we directly train the ASR decoder to predict target punctuation marks. The ASR punctuation predictions (acoustic-based) are combined with the NLP module predictions (lexical-based) via a probability thresholding process. A heuristic-based post-processing step is then applied to make corrections in the prediction as the final step. Our system is illustrated in Figure 1. The set of Spanish punctuation marks predicted by the system are: OPEN_QUESTION (¿)⁴, CLOSE_QUESTION (?), COMMA (,), PERIOD (.), and NONE (for tokens that have no associated punctuation marks).

⁴An open question mark (¿) is used at the start position of a question in Spanish

3.2 Acoustic-based Prediction

Acoustic features, including intonation and prosody, play an important role in distinguishing declarative and interrogative sentences in Spanish, as described in section 2.2. In order to leverage our ASR module to directly predict punctuation marks from the speech signal, we keep each target punctuation mark in our ASR training data and treat it as an individual token by separating it from surrounding words; an example of predicting CLOSE_QUESTION is shown in Figure 1. Note that we omit OPEN_QUESTION from the training data since it can mostly be restored by heuristics in the following post-processing step.

We use an End-to-End based ASR system provided by the Nemo toolkit (Kuchaiev et al., 2019). The applied Conformer-CTC architecture is slightly different from the original Conformer architecture (Gulati et al., 2020), where the LSTM decoder is replaced with a linear decoder. The encoder uses CTC (Connectionist Temporal Classification) loss (Graves et al., 2006) instead of RNNT (RNN-Transducer) (Graves, 2012) which makes it a non-autoregressive model. For word prediction, we use an in-house streaming decoder with language model shallow fusion.

3.3 Lexical-based Prediction

The lexical-based approach is capable of predicting all supported punctuation marks outlined in section 3.1, which consumes unpunctuated transcribed text emitted from the ASR module as shown in Figure 1. For the NLP module utilized in this lexical-based prediction, we follow the similar structure as uti-

lized in (Zhu et al., 2022a) and (Fu et al., 2021), which is a fine-tuned mBERT (multilingual BERT) (Devlin et al., 2019) with an additional token classification head. The output of the prediction indicates the appropriate punctuation marks to be attached to the corresponding input token. Additionally, a probability score (illustrated as p in Figure 1) is also computed for each token using a softmax layer on top of the prediction logits, which reflects the confidence of each predicted punctuation mark in the lexical-based prediction.

Algorithm 1 Thresholding algorithm

Input:
 $Pred_a$: acoustic-based prediction
 $Pred_l$: lexical-based prediction
 P_l : probability score of $Pred_l$
 $T_{question}$ & $T_{declarative}$: hyperparameters
Output:
 $Pred_c$: consolidated prediction
Where:
 C_Q : CLOSE_QUESTION

```

if  $Pred_a == C\_Q$  and  $Pred_l$  in  $\{PERIOD, COMMA\}$ 
then
  if  $P_l \leq T_{declarative}$  then
     $Pred_c \leftarrow C\_Q$ 
  else
     $Pred_c \leftarrow Pred_l$ 
  end if
else if  $Pred_a \neq C\_Q$  and  $Pred_l == C\_Q$  then
  if  $P_l \leq T_{question}$  then
     $Pred_c \leftarrow PERIOD$ 
  else
     $Pred_c \leftarrow C\_Q$ 
  end if
else
   $Pred_c \leftarrow Pred_l$ 
end if

```

3.4 Hybrid Prediction

To consolidate the results from both acoustic-based and lexical-based predictions, we introduce a probability thresholding step based on the probability score generated by the lexical-based prediction. Our approach focuses on improving Spanish question prediction, which employs a set of threshold values $T_{question}$ and $T_{declarative}$ as hyperparameters. These thresholds represent the minimal probability score the lexical-based prediction needs to have when conflicting with acoustic-based prediction. The detailed thresholding algorithm is illustrated in Algorithm 1. The optimal values of $T_{question}$ and $T_{declarative}$ are identified through grid search towards the development dataset in our experiment⁵.

⁵We found $[0.7, 0.8]$ is usually a reasonable range to start with for both $T_{question}$ and $T_{declarative}$ in our experiments.

A heuristic-based post-processing step (details in Appendix A.1) is also applied after probability thresholding to mitigate the error caused by unmatched OPEN_QUESTION and CLOSE_QUESTION in the prediction. For example, as illustrated in the hybrid prediction result in Figure 1, an OPEN_QUESTION is added on the first token of the word chunk *los sábados están abiertos?* after an unmatched CLOSE_QUESTION is created after the thresholding process.

4 Experiment

4.1 Datasets

We conduct our experiment using a variety of data resources. Since the proposed system is used in our call center product, the in-domain data resource is our internal audio recording and human-annotated transcripts from real customer support calls in Spanish. This internal data resource consists of 50 hours of audio and around 10,000 rows of corresponding transcribed utterances (more statistical detail is available in Appendix A.2). Apart from our internal dataset, Linguistic Data Consortium (LDC) Spanish Fisher corpora (Graff et al., 2010a,b) is also added as a supplementary resource for real-life human conversations, which has approximately 160 hours of audio with 130,000 rows of transcribed utterances from Spanish telephone conversations. Out of both LDC and our internal data, we leave out 10% and 5% as test and development sets respectively in our experiments. Note that in order to evaluate the performance of our system on reference transcripts, we leverage Levenshtein distance to align punctuation marks from each ASR hypothesis to reference transcript in acoustic-based prediction during our evaluation process on the test set.

Additionally, the open-sourced Spanish datasets from Openslr (Guevara-Rukoz et al., 2020; Kolobov et al., 2021) and Common-voice (Ardila et al., 2019) are used in ASR training as well, which collectively provide 1200 hours of audio. A subset of 80,000 utterances were also randomly sampled from the Spanish OpenSubtitle corpus (Lison and Tiedemann, 2016) and added as an extra text-only dataset into the NLP module training process to improve the accuracy of lexical-based prediction. All text-based resources are also used in the language model for the in-house streaming ASR decoder.

	Baseline	ASR w/ C_Q	ASR w/ all
Public	15.77	15.44 (-2.1%)	16.53 (+4.8%)
Internal	26.81	26.36 (-1.7%)	27.95 (+4.3%)

Table 1: WER and relative changes compared to the baseline on public (LDC) and internal datasets, where the *Baseline* performance is evaluated by the ASR module trained without punctuation. *ASR w/ C_Q*: ASR module trained with only CLOSE_QUESTION; *ASR w/ all*: ASR module trained with CLOSE_QUESTION, PERIOD and COMMA.

	Reliability	Latency (s)
ChatGPT-few	92.4%	1.13
ChatGPT-zero	87.9%	1.10
PaLM2-few	28.7%	0.56
PaLM2-zero	28.6%	0.49
Our system (excl. ASR)	-	0.04

Table 2: Reliability and Latency comparison between LLM APIs (with both zero- and few- shot prompting) and our system (excluding ASR latency), averaged over all internal and public test samples. Latency shown as "seconds per input utterance".

4.2 Experiment Setup

For the ASR module, we use the Nemo (Kuchaiev et al., 2019) Spanish model `STT_Es_Conformer CTC_Large` as the pre-trained model. The presented model is fine-tuned for 20 epochs, with the Adam optimizer (Kingma and Ba, 2014) and no weight decay. The Noam learning scheduler (Vaswani et al., 2017) is used with a warmup of 100 steps and a learning rate of 0.01.

In consideration of the real-time inference speed, we take only the bottom 6 layers of `bert-base-multilingual-cased` from Hugging Face (Wolf et al., 2020) library as the backbone of our NLP module. The 6-layer mBERT is then fine-tuned through a token classification task using all lexical training data described in section 4.1. The NLP module is trained using the Adam optimizer with 4 epochs and a learning rate of $3e-5$.

In the subsequent sections, all assessments are performed utilizing a single Intel Xeon 2.20GHz CPU, 1.5G memory and under identical network connection condition.

5 Results

5.1 Evaluation on Speech Recognition

We first evaluate the performance impact by introducing CLOSE_QUESTION prediction in our ASR module. Word-Error-Rate (WER) is a standard metric for the ASR system. A lower word error rate shows superior accuracy in speech recognition, compared with a higher word error rate. To accurately determine the word error rate of

the ASR module, free from punctuation interference, we exclude all punctuation marks in both the ASR hypothesis and reference transcripts while evaluating. Table 1 shows WER on both test sets. Compared to our baseline, the ASR module trained only with CLOSE_QUESTION shows 2.1% and 1.7% WER improvement in public (LDC) and the internal test set respectively, which indicates that our ASR module can learn better acoustic features of Spanish interrogative sentences by keeping CLOSE_QUESTION in training data. In addition to predicting CLOSE_QUESTION, we also conduct a second experiment to keep all CLOSE_QUESTION, COMMA and PERIOD in the ASR module, but this unexpectedly increases the WER by up to 4.8%, which is not a tolerable performance deterioration for our production use. Therefore, we only focus on CLOSE_QUESTION prediction from the ASR module in our design.

5.2 Evaluation on Punctuation Restoration

In order to assess the comprehensive proficiency of our system in restoring Spanish punctuation, we conduct a benchmark test against some leading LLMs available on the market. First, we evaluate and compare the runtime performance of producing Spanish punctuation marks from unpunctuated transcripts between (a) utilizing commercial LLM APIs (ChatGPT and PaLM2) and (b) executing our proposed system. Our evaluation criteria for this analysis include two metrics: (1) *Reliability*: the percentage of the results where the original input words can be extracted without any modification or reordering (except casing changes), to measure the impact of the LLM hallucination or other undesired outcomes. (2) *Latency*: the elapsed time to receive responses from API calls for ChatGPT and PaLM2, as well as the total execution time of our lexical and hybrid prediction (excluding ASR latency), under the same environment setting as stated in section 4.2. Table 2 presents the *Reliability* and *Latency* comparison, it is clear that except our proposed system, all LLM APIs exhibit various levels of reliability concerns. Additionally, our system shows much lower latency compared to API calls. It is also noteworthy that *Reliability* of PaLM2 stands at a mere 28% in both zero- and few- shot prompting, suggesting that it is not suitable for the Spanish punctuation restoration task. Details on the API call setup and prompts are listed in Appendix A.3 and A.4.

Table 3 presents the comprehensive F1 score

	Public (LDC) data					Internal data				
	Lexical	Acoustic	ChatGPT-zero	ChatGPT-few	Hybrid	Lexical	Acoustic	ChatGPT-zero	ChatGPT-few	Hybrid
C_Q	54.0	51.5	13.5	13.9	58.2	47.3	28.2	24.4	28.6	51.7
O_Q	50.6	-	11.1	11.7	52.7	44.5	-	22.9	25.4	47.0
COMMA	60.8	-	43.5	51.0	60.7	68.9	-	47.2	60.4	69.0
PERIOD	87.7	-	58.6	58.8	88.0	83.6	-	59.5	72.4	83.8
Overall ¹	74.49	-	44.92	49.57	74.83	72.29	-	48.04	61.20	72.61

¹Micro average of all punctuation

Table 3: F1 score comparison over all punctuation marks with different approaches. *C_Q*: CLOSE_QUESTION; *O_Q*: OPEN_QUESTION; *Lexical*: lexical-based prediction; *Acoustic*: acoustic-based prediction; *Hybrid*: our proposed hybrid system with consolidated prediction; *ChatGPT-few/zero*: ChatGPT with few/zero-shot prompting, details in Appendix A.

	Public (LDC) data				Internal data			
	Lexical	Acoustic	Union	Threshold	Lexical	Acoustic	Union	Threshold
Precision	48.0	65.3 ²	47.0	53.1	63.7	98.4 ³	66.0	66.1
Recall	61.7	44.1	71.1	64.3	37.7	16.4	42.1	42.4
F1	54.0	51.5	56.5	58.2	47.3	28.2	51.4	51.7

²19.3% of the True Positive prediction is ambiguous in unpunctuated text, and not identified as questions by Lexical.

³32.3% of the True Positive prediction is ambiguous in unpunctuated text, and not identified as questions by Lexical.

Table 4: F1, precision and recall comparison on **CLOSE_QUESTION** using different approaches. *Lexical*: lexical-based prediction; *Acoustic*: acoustic-based prediction; *Union*: the union of **CLOSE_QUESTION** predictions from both lexical and acoustic prediction; *Threshold*: our proposed thresholding process to consolidate lexical and acoustic predictions.

performance of our punctuation restoration system and LLM API⁶ on both public and internal datasets. Note that we also show the performance of the standalone lexical module which represents the conventional BERT-based lexical-only structure used in recent punctuation restoration studies (Zhu et al., 2022a; Fu et al., 2021). It is clear that both lexical and hybrid predictions demonstrate a substantial accuracy advantage over ChatGPT. Moreover, the hybrid approach, enhanced by the improvements in **CLOSE_QUESTION** of up to 4.4%, exhibits varied degrees of F1 score improvement for all other punctuation marks after our thresholding and post-processing step outlined in section 3.4. Consequently, our proposed hybrid system outperforms the lexical-only approach by 0.34% and 0.32% absolute in overall F1 score respectively on public and internal datasets.

To better illustrate the enhancement on **CLOSE_QUESTION** from our hybrid system, we additionally provide precision, recall and F1 score details on **CLOSE_QUESTION** in Table 4. Although with a lower F1 score, acoustic-based prediction exhibits a higher precision in predicting **CLOSE_QUESTION** compared to lexical prediction in both testing datasets. In addition, up to 32.3% of True Positives from the acoustic predic-

tion is ambiguous in unpunctuated text and does not overlap with that in lexical prediction. In order to demonstrate the effectiveness of our proposed thresholding process to consolidate acoustic and lexical predictions as described in section 3.4, we compare it with a naive union of the two on **CLOSE_QUESTION**. Table 4 shows that *Thresholding* consistently outperforms *Union* in both datasets. As a result, with our thresholding approach, the F1 score for **CLOSE_QUESTION** is noticeably improved by 4.2% and 4.4% compared to lexical-only prediction across public and internal datasets respectively.

6 Future Work

From the evaluation result in section 5.1, contrary to the WER improvement when predicting only **CLOSE_QUESTION** by the ASR module, we discovered a WER deterioration when adding **COMMA** and **PERIOD** to the prediction. Future work may focus on establishing a possible cause for this change. In addition, lexical ambiguity between questions and declarations exists beyond Spanish; thus, a natural next step would be evaluating our system in other human languages.

7 Conclusion

In this study, we propose a hybrid acoustic-lexical punctuation restoration system for Spanish conversational transcripts, with a focus to address the ambiguity in unpunctuated Spanish questions. The

⁶Only reliable outcomes from ChatGPT are evaluated. PaLM2 is left out in this evaluation as it cannot produce reliable results of a large enough size to establish a meaningful comparison, due to its low *Reliability*.

proposed system leverages an ASR decoder to make direct predictions of Spanish question marks, which are later consolidated with lexical predictions from an NLP module. We evaluate the system on both internal and public datasets and show that it can effectively enhance Spanish question marks prediction, and consequently improve the overall punctuation restoration accuracy. Additional benchmark indicates that our proposed system outperforms some top LLMs in accuracy, latency and reliability. Furthermore, we demonstrate that keeping question marks in the ASR decoder vocabulary results in an improved WER of the ASR module alone.

8 Ethical Considerations

During our internal data collection process, we implement a data retention policy for all our users, such that user consent is obtained prior to any data collection. In addition, we have ensured that all the annotators involved in the transcription process of our internal dataset are paid with adequate compensation. Moreover, to protect the privacy and confidentiality of individuals, the dataset underwent further processing to remove any sensitive, personal, or identifiable information.

References

- Rosana Ardila, Megan Branson, Kelly Davis, Michael Henretty, Michael Kohler, Josh Meyer, Reuben Morais, Lindsay Saunders, Francis M. Tyers, and Gregor Weber. 2019. [Common Voice: A Massively-Multilingual Speech Corpus](#).
- Meghan E. Armstrong. 2017. [Accounting for intonational form and function in puerto rican spanish polar questions](#). *Probus*, 29(1):1–40.
- Esther Brown and Javier Rivas. 2011. [Subject-verb word order in spanish interrogatives: A quantitative analysis of puerto rican spanish](#). *Spanish in Context*, 8.
- Alejandro Cuza. 2016. [The status of interrogative subject-verb inversion in spanish-english bilingual children](#). *Lingua*, 180.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Timothy L. Face. 2005. [F0 peak height and the perception of sentence type in castilian spanish](#). *Revista Internacional de Lingüística Iberoamericana*, 3:49–65.
- Xue-Yong Fu, Cheng Chen, Md Tahmid Rahman Laskar, Shashi Bhushan, and Simon Corston-Oliver. 2021. [Improving punctuation restoration for speech transcripts via external data](#). In *Proceedings of the Seventh Workshop on Noisy User-generated Text (W-NUT 2021)*, pages 168–174, Online. Association for Computational Linguistics.
- Ander González-Docasal, Aitor García-Pablos, Haritz Arzelus, and Aitor Álvarez. 2021. [Autopunct: A BERT-based Automatic Punctuation and Capitalisation System for Spanish and Basque](#). *Procesamiento del Lenguaje Natural*, 67(0):59–68.
- David Graff, Shudong Huang, Ingrid Cartagena, Kevin Walker, and Christopher Cieri. 2010a. [Fisher Spanish - Transcripts LDC2010T04](#). Web Download. Philadelphia: Linguistic Data Consortium.
- David Graff, Shudong Huang, Ingrid Cartagena, Kevin Walker, and Christopher Cieri. 2010b. [Fisher Spanish Speech LDC2010S01](#). Web Download. Philadelphia: Linguistic Data Consortium.
- Agustin Gravano, Martin Jansche, and Michiel Bacchiani. 2009. [Restoring punctuation and capitalization in transcribed speech](#). In *2009 IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 4741–4744.
- A. Graves, S. Fernandez, F. Gomez, and J. Schmidhuber. 2006. [Connectionist temporal classification: Labelling unsegmented sequence data with recurrent neural nets](#). In *ICML '06: Proceedings of the International Conference on Machine Learning*.
- Alex Graves. 2012. [Sequence transduction with recurrent neural networks](#). *CoRR*, abs/1211.3711.
- Yushi Guan. 2020. [End to End ASR System with Automatic Punctuation Insertion](#).
- Adriana Guevara-Rukoz, Isin Demirsahin, Fei He, Shan-Hui Cathy Chu, Supheakmungkol Sarin, Knot Pitsrisawat, Alexander Gutkin, Alena Butryna, and Oddur Kjartansson. 2020. [Crowdsourcing Latin American Spanish for Low-Resource Text-to-Speech](#). In *Proceedings of The 12th Language Resources and Evaluation Conference (LREC)*, pages 6504–6513, Marseille, France. European Language Resources Association (ELRA).
- Anmol Gulati, James Qin, Chung-Cheng Chiu, Niki Parmar, Yu Zhang, Jiahui Yu, Wei Han, Shibo Wang, Zhengdong Zhang, Yonghui Wu, and Ruoming Pang. 2020. [Conformer: Convolution-augmented Transformer for Speech Recognition](#).
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. [Long Short-term Memory](#). *Neural computation*, 9:1735–80.

- Douglas Jones, Florian Wolf, Edward Gibson, Elliott Williams, Evelina Fedorenko, Douglas Reynolds, and Marc Zissman. 2003. [Measuring the readability of automatic speech-to-text transcripts](#).
- Diederik P. Kingma and Jimmy Ba. 2014. [Adam: A Method for Stochastic Optimization](#).
- Ondřej Klejch, Peter Bell, and Steve Renals. 2017. [Sequence-to-sequence models for punctuated transcription combining lexical and acoustic features](#). In *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5700–5704.
- Rostislav Kolobov, Olga Okhapkina, Andrey Platonov, Olga Omelchishina, Roman Bedyakin, Vyacheslav Moshkin, Dmitry Menshikov, and Nikolay Mikhaylovskiy. 2021. [MediaSpeech: Multilanguage ASR Benchmark and Dataset](#).
- Oleksii Kuchaiev, Jason Li, Huyen Nguyen, Oleksii Hrinchuk, Ryan Leary, Boris Ginsburg, Samuel Kriman, Stanislav Beliaev, Vitaly Lavrukhin, Jack Cook, et al. 2019. [Nemo: a toolkit for building ai applications using neural modules](#). *arXiv preprint arXiv:1909.09577*.
- Viet Dac Lai, Nghia Trung Ngo, Amir Poursan Ben Veyseh, Hieu Man, Franck Dernoncourt, Trung Bui, and Thien Huu Nguyen. 2023. [Chatgpt beyond english: Towards a comprehensive evaluation of large language models in multilingual learning](#).
- Su Ar Lee and B.A. and M.A. 2010. [Absolute interrogative intonation patterns in Buenos Aires Spanish](#).
- Xinxing Li and Edward Lin. 2020. [A 43 Language Multilingual Punctuation Prediction Neural Network Model](#). In *Proc. Interspeech 2020*, pages 1067–1071.
- Pierre Lison and Jörg Tiedemann. 2016. [OpenSubtitles2016: Extracting large parallel corpora from movie and TV subtitles](#). In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC’16)*, pages 923–929, Portorož, Slovenia. European Language Resources Association (ELRA).
- Wei Lu and Hwee Tou Ng. 2010. [Better Punctuation Prediction with Dynamic Conditional Random Fields](#). In *Proceedings of the 2010 Conference on Empirical Methods in Natural Language Processing*, pages 177–186, Cambridge, MA. Association for Computational Linguistics.
- Ian Mackenzie. 2021. [The linguistics of spanish](#).
- Patrick K. O’Neill, Vitaly Lavrukhin, Somshubra Majumdar, Vahid Noroozi, Yuekai Zhang, Oleksii Kuchaiev, Jagadeesh Balam, Yuliya Dovzhenko, Keenan Freyberg, Michael D. Shulman, Boris Ginsburg, Shinji Watanabe, and Georg Kucsko. 2021. [SPGISpeech: 5,000 hours of transcribed financial audio for fully formatted end-to-end speech recognition](#).
- Vasile Păiș and Dan Tufiș. 2021. [Capitalization and punctuation restoration: a survey](#). *Artificial Intelligence Review*, 55:1681 – 1722.
- Chengwei Qin, Aston Zhang, Zhuosheng Zhang, Jiaao Chen, Michihiro Yasunaga, and Diyi Yang. 2023. [Is chatgpt a general-purpose natural language processing task solver?](#)
- Chase Wesley Raymond. 2015. [Questions and responses in spanish monolingual and spanish–english bilingual conversation](#). *Language and Communication*, 42:50–68.
- Monica Sunkara, Srikanth Ronanki, Dhanush Bekal, Sravan Bodapati, and Katrin Kirchhoff. 2020. [Multimodal Semi-supervised Learning Framework for Punctuation Prediction in Conversational Speech](#).
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems*, pages 5998–6008.
- Erik Willis. 2007. [Utterance signaling and tonal levels in dominican spanish declaratives and interrogatives](#). *Journal of Portuguese Linguistics*, 6:179.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Kaituo Xu, Lei Xie, and Kaisheng Yao. 2016. [Investigating LSTM for punctuation prediction](#). In *2016 10th International Symposium on Chinese Spoken Language Processing (ISCSLP)*, pages 1–5.
- Xiliang Zhu, Shayna Gardiner, David Rossouw, Tere Roldán, and Simon Corston-Oliver. 2022a. [Punctuation restoration in Spanish customer support transcripts using transfer learning](#). In *Proceedings of the Third Workshop on Deep Learning for Low-Resource Natural Language Processing*, pages 80–89, Hybrid. Association for Computational Linguistics.
- Yaoming Zhu, Liwei Wu, Shanbo Cheng, and Mingxuan Wang. 2022b. [Unified Multimodal Punctuation Restoration Framework for Mixed-Modality Corpus](#).

A Appendix

A.1 Heuristic-based post-processing

As mentioned in 3.4, we apply the following heuristic-based steps to post-process the prediction result:

1. Convert all unmatched OPEN_QUESTION to NONE in the prediction.
2. For all unmatched CLOSE_QUESTION, change the prediction of the first token in the continuous word chunk (the longest continuous word sequence where no punctuation is predicted in-between) to OPEN_QUESTION.

A.2 Description of our internal dataset

Our internal data is collected from audio recordings of Spanish customer support conversations, covering a large range of domains such as retail, technology, automotive and professional services. Our primary focus lies within the North American region, including both Mexican and American accents. The audio duration of the dataset totals around 50 hours. For ASR training purposes, each individual audio clip is broken down into segments based on audio silence, with a maximum of 2 minutes and averaging approximately 18 seconds. Audio clips are also transcribed by the annotators to create text data for NLP training. We provide the statistical summary on the length of the transcribed utterances in Table 5.

	mean	medium	min	max	std
num of words	43.4	38.0	1.0	231.0	25.4

Table 5: Statistical summary on length of the utterances in our internal dataset.

A.3 API call setup

We use `gpt-3.5-turbo` for ChatGPT and `text-bison@001` for PaLM2 in API calls. For both models, *temperature* is set as 0.2 while the maximum token length of output (named as *max_tokens* in ChatGPT and *maxOutputTokens* in PaLM2) is configured as 1024.

A.4 Prompt

The following prompts are used in our experiments when calling LLM APIs:

Few-shot prompting:

Without any explanation or modification, add punctuation to the following Spanish transcript from human conversations, use only punctuation marks from this list: `comma(,)`, `period(.)`, `open_question(¿)` and `close_question(?)`. Return the punctuated utterance only. Here are some examples:

```
### Input: {Unpunctuated Spanish Utterance 1}
```

```
### Output: {Punctuated Spanish Utterance 1}
```

```
### Input: {Unpunctuated Spanish Utterance 2}
```

```
### Output: {Punctuated Spanish Utterance 2}
```

```
### Input: {Unpunctuated Spanish Utterance 3}
```

```
### Output: {Punctuated Spanish Utterance 3}
```

Now, add punctuation marks to:

```
### Input: {text}
```

```
### Output:
```

Zero-shot prompting:

Without any explanation or modification, add punctuation to the following Spanish transcript from human conversations, use only punctuation marks from this list: `comma(,)`, `period(.)`, `open_question(¿)` and `close_question(?)`. Return the punctuated utterance only.

Add punctuation marks to:

```
### Input: {text}
```

```
### Output:
```

where we put the unpunctuated test utterance in the `text` field. Note that in all of our experiments, we use three in-context examples for few-shot prompting. In addition, we make sure to sample utterances with presence of all targeted punctuation marks in these three in-context examples. Note that both zero-shot and few-shot prompting are used in the evaluation results as presented in Table 3 and Table 4.