



Optimizing ASR for Catalan-Spanish Code-Switching: A Comparative Analysis of Methodologies

Carlos Mena¹, Pol Serra¹, Jacobo Romero¹, Abir Messaoudi¹, Jose Giraldo¹, Carme Armentano-Oller¹, Rodolfo Zevallos¹, Ivan Meza², Javier Hernando^{1,3}

¹LangTech Lab, Barcelona Supercomputing Center, Spain; ²IIMAS, Universidad Nacional Autónoma de México, Mexico; ³Universitat Politècnica de Catalunya, Spain

carlos.hernandez@bsc.es, ivanvladimir@turing.iimas.unam.mx, javier.hernando@upc.edu

Abstract

Code-switching (CS), the alternating use of two or more languages, challenges automatic speech recognition (ASR) due to scarce training data and linguistic similarities. The lack of dedicated CS datasets limits ASR performance, as most models rely on monolingual or mixed-language corpora that fail to reflect real-world CS patterns. This issue is critical in multilingual societies where CS occurs in informal and formal settings. A key example is Catalan-Spanish CS, widely used in media and parliamentary speeches. In this work, we improve ASR for Catalan-Spanish CS by exploring three strategies: (1) generating synthetic CS data, (2) concatenating monolingual audio, and (3) leveraging real CS data with language tokens. We extract CS data from Catalan speech corpora and fine-tune OpenAI’s Whisper models, making them available on Hugging Face. Results show that combining a modest amount of synthetic CS data with the dominant language token yields the best transcription performance.

Index Terms: code-switching, speech augmentation, Catalan-Spanish, speech recognition

1. Introduction

Code-switching (CS), also referred to as *code-mixing*, *language alternation*, or various other terms in the literature [1, 2], is a phenomenon in which bilingual or multilingual speakers alternate between two or more languages within a conversation, a sentence, or even within a single word [3]. It is a common linguistic practice in communities where multiple languages co-exist, often influenced by social, cultural, and cognitive factors [1]. Code-switching can occur for various reasons, such as emphasizing a point, addressing different audiences, or filling lexical gaps [4]. According to various authors, we can distinguish between three types of alternations [5, 3]:

- **Inter-sentential:** Switching occurs between entire sentences, typically at clause boundaries, where each sentence remains grammatically independent in its respective language. (e.g., “Mañana tengo un examen muy difícil. *I hope I do well.*”)
- **Intra-sentential:** Switching happens within a single sentence, mixing words or phrases from different languages while maintaining grammatical coherence. (e.g., “Ayer fuimos al *mall* y compramos mucha ropa *on sale.*”)
- **Tag-switching:** A discourse marker, filler, or tag phrase from another language is inserted into the sentence without affecting its grammatical structure. (e.g., “Eso es increíble, *you know?*”)

In this paper, we explore different approaches to enhancing Automatic Speech Recognition (ASR) in code-switching

environments between Catalan and Spanish, which is commonly found in everyday conversations in Catalan-speaking areas. This task is particularly challenging due to the linguistic similarities between both languages.

Catalan-Spanish code-switching has been widely studied from various perspectives. The literature includes studies on the behavior and acquisition of Catalan-Spanish CS in early childhood [6, 7, 8], as well as research exploring the influence of cultures from outside Catalonia [9, 2], the impact of Catalan culture in other regions of the world [10, 11], historical perspectives on CS in different periods of history [12], and its presence in various domains, such as comedy [13, 14]. More recently, studies on this topic have been conducted in the context of Catalan television [15] and social media [16].

However, despite the extensive literature on the subject, there are surprisingly few, if any, speech datasets exclusively dedicated to Catalan-Spanish CS. For this reason, in Section 3, we include a detailed explanation of how we detect and extract CS utterances from some of the datasets presented in Section 2. This is a direct contribution of our work to the community, as these subcorpora are publicly shared through the Hugging Face platform.

2. Datasets

In the name of reproducibility in science, all the data presented in this section is publicly available on the Hugging Face platform. For a quick overview, Table 1 summarizes the datasets presented in this section, the amount of CS speech extracted from them, and their primary function (test or fine-tune).

Table 1: Summary of the datasets and their primary function (test or fine-tune).

Dataset	CS Speech	Usage
ParlamentParla	2 hours	Test
Corts Valencianes	2 hours	Test
TV3 Parla	1h45m	Test
TV3 Parla	4 hours	fine-tune
Common Voice c17.0	105 hours	fine-tune
Synthetic CS Data	17 hours	fine-tune

2.1. ParlamentParla

ParlamentParla [17] is a speech corpus for Catalan, prepared by Col·lectivaT [18], consisting of audio segments extracted from Catalan Parliament plenary sessions (2007–2018). The transcriptions were aligned with the recordings, and the data is

released under the Parliament’s terms of use¹.

Currently, no in-depth linguistic study of the corpus has been conducted. However, we assume that the predominant accent is Central Catalan, as it corresponds to the geographical area where most of Catalonia’s population resides [19].

The corpus was supported by the Catalan Department of Culture and, in v2.0, by the Barcelona Supercomputing Center through the AINA project [20]. Version 2.0 [21] includes 211 hours of clean speech and 400 hours of other quality segments, with speaker and gender tags.

2.2. Corts Valencianes

The Corts Valencianes Speech Corpus [22] is a comprehensive dataset consisting of speech recordings from the sessions of the Corts Valencianes². It features high-quality and other-quality segments, categorized into short segments (under 30 seconds) and long segments (over 30 seconds). The dataset contains a total of 2,621,096 words and covers a total of 270 of speech, with 239 hours dedicated to short segments and 31 hours to long segments.

Although no in-depth linguistic study has been conducted in this case, it can be inferred that the predominant accent is Valencian, as it is the most common in this region.

2.3. TV3 Parla

TV3 Parla [23] is a 272-hour Catalan corpus extracted from the Catalan public television channel TV3 [24]. The details of segmentation and data processing are explained in [25]. Its content is publicly available but owned by Corporació Catalana de Mitjans Audiovisuals, SA (CCMA).

As in Section 2.1, we assume for demographic reasons that the predominant accent is the Central variety, albeit to a lesser extent, as one of the objectives of Catalan public television is to promote the Catalan language in all its varieties.

2.4. Mozilla Common Voice v17.0

Mozilla CommonVoice [26] is a collaborative project aimed at creating an open-source dataset of voice recordings to improve speech recognition technologies, especially for underrepresented languages. The project collects speech data from volunteers worldwide, ensuring diversity in accents, demographics, and speaking styles.

In our experiments, we use version 17 of the Common Voice corpus [27] for both Spanish and Catalan, leveraging its high-quality transcriptions and diverse speaker contributions to enhance our ASR models.

2.5. Synthetic CS Data

The synthetic audio corpus comprises 17 hours of CS data, with sentences extracted from the TV3 Parla corpus through a process detailed in Section 3.2. In essence, this data is generated by taking a Catalan text corpus (TV3 Parla) and introducing code-switching by translating noun chunks into Spanish, ensuring a natural integration within the original linguistic context. Once the text corpus was prepared, speech was synthesized using TTS technology.

¹<https://www.parlament.cat/pcat/serveis-parlament/avis-legal/>

²<https://www.cortsvalencianes.es/>

3. Methodology

This work aims to test different methods for improving the performance of ASR systems in Catalan and Spanish CS contexts. The first method involves generating synthetic CS sentences and then synthesizing them using a Text-to-Speech (TTS) system, while the second method consists of concatenating audio recordings in both languages that contain real human voices. The third method implies the detection of language tokens in transcriptions.

3.1. Creation of the Evaluation Dataset

To obtain gold standard data to evaluate the effectiveness of our ASR models, we followed two main strategies: manual and BERT [28] detection.

3.1.1. Manual Detection of CS

For the manual detection, we extracted CS sentences from the *ParlamentParla* and *Corts Valencianes* corpora. These datasets primarily contain Catalan speech in a parliamentary discourse context, delivered by Catalan and Valencian politicians. The method used to detect CS sentences was keyword-based detection. For example, in Catalan, the conjunction “and” is “i”, while in Spanish, it is “y”. Therefore, sentences containing “y” are likely to include CS. Once sentences were detected using this and other keywords, they were manually verified, and only those containing at least three consecutive words in Spanish were accepted to be part of the evaluation set.

3.1.2. BERT Detection of CS

In the case of the BERT detection, Google’s multilingual BERT [29] was fine-tuned for token classification using a synthetic corpus of code-switched dialogues in Catalan and Spanish. During fine-tuning, each word was labelled with its corresponding language token. Once trained, the model was applied to the transcriptions of the *TV3 Parla* dataset, where it performed token-level language classification. This process resulted in a “language count” for each audio file, indicating the distribution of Catalan and Spanish within the transcription. Given that the audios were short, an audio was considered code-switched if both Catalan and Spanish were present with at least three words each. With this method, we identified a substantial portion of code-switched data, totalling approximately 5 hours and 45 minutes. This data was then split for model training, with 70% used for fine-tuning and 30% for validation and testing.

3.2. Synthetic CS Generation

The creation of synthetic corpora involved two main processes: text generation and audio generation through text-to-speech (TTS). To generate the synthetic text corpus with code-switching between Catalan and Spanish, sentences were first extracted from the TV3 Parla [23], filtering them based on their length. Then, one or two *noun chunks* per sentence were identified using *spaCy* [30] and translated into Spanish with *MarianMT* [31], replacing the original segments with `<cat>` and `<esp>` tags.

To synthesize the generated CS text, we trained a multilingual TTS based on a Catalan Matcha-TTS [32, 33] version. The datasets used for the TTS training were Librivox Spanish [34], Festcat [35] and LaFresCat [36]. We added an embedding for language and a custom front-end to handle CS text. As the syn-

Table 2: WER (%) of different fine-tuning and decoding strategies. The experiments were done on the Code-switching parts of the datasets.

Experiment	Decoding Tokens	TV3	ParlamentParla	Corts
Base model		31.96	20.96	35.43
	<ca><es>	25.12	22.39	29.81
	<es><ca>	25.49	20.13	34.16
	<ca>	26.67	22.17	30.34
Synthetic data	<es>	63.16	57.87	52.16
	<ca><es>	23.22	16.38	22.82
	<es><ca>	23.81	15.47	22.71
	<ca>	23.46	14.48	22.42
TV3	<es>	24.23	15.72	23.76
	<ca><es>	21.96	15.76	23.98
	<es><ca>	21.46	16.65	24.86
	<ca>	21.20	16.01	24.22
Tuples	<es>	51.29	50.58	47.75
	<ca><es>	28.99	17.83	25.24
	<es><ca>	29.92	21.37	29.72
	<ca>	31.42	18.31	26.06
	<es>	52.78	50.22	47.32

thetic text was generated with language tags, it allowed us to route the sentences to the proper phonemizer together with the language index of the embedding. The whole text corpus was synthesized with 10 voices rated the best in an internal listening assessment. The final step involved a data augmentation pipeline with the audiomentations³ library since the goal was to simulate a real user environment instead of the clean that outputs the TTS. This pipeline used the following transformations with randomization of the parameters:

- Additive environmental noise and speech babble noise. ($min_absolute_rms_in_db = -45.0, max_absolute_rms_in_db = -40.0$)
- Clipping ($a_min = -0.07, a_max = 0.07, p = 1.0$)
- Tanh distortion ($min_distortion = 0.3, max_distortion = 0.3, p = 0.5$)
- Gradual gain transition ($min_gain_db = -3, max_gain_db = 3, min_duration = 0.5, max_duration = 1, p = 1.0$)
- Bitcrushing ($min_bit_depth = 8, max_bit_depth = 8, p = 1.0$)

3.3. Audio Concatenation

The idea of concatenating audio from different languages as a technique to handle CS in ASR is not new at all. The NVIDIA-NeMo GitHub repository provides scripts to create synthetic code-switched data using this approach⁴. However, in this work, we implemented our version of this script to ensure that the concatenation was performed exactly as intended.

The Spanish and Catalan audio files used for concatenation were extracted from version 17 of Mozilla Common Voice. We selected 50 hours per language, ensuring that the speakers' gender was well specified (male or female) to create a gender-balanced corpus.

Audio files were concatenated in pairs (tuples), ensuring that each audio segment in one language was always paired with

another segment in the other language. If 30,000 audio clips (50 hours) from one language are concatenated with 30,000 clips (50 hours) from the other, the result is 30,000 tuples (100 hours). Of these, 15,000 will be of type ca-es (starting with Catalan and then Spanish), and 15,000 will be of type es-ca (starting with Spanish and then Catalan).

3.4. Language Tag Detection

Language tag detection is a technique that has demonstrated acceptable results when used with other languages [37]. We carried out a series of experiments to identify the most effective strategy for handling CS between Catalan and Spanish while minimizing transcription errors. Our approach builds on the original configuration of Whisper, which employs the following prompt: <|startofprev|>previous text<|startoftranscript|><|language|><|task|><|notimestamps|>, as described in [38]. We modified this setup to assess its impact on CS transcription performance.

Specifically, we investigated the effect of concatenating language tags on transcription accuracy for CS utterances. This involved explicitly enforcing two language tags (e.g., <|ca|><|es|>) within the prompt: <|startoftranscript|><|ca|><|es|><|task|><|notimestamps|>. Previous findings from [38] reported a 19%

Additionally, we examined the influence of varying the position of language tags in the decoding sequence by placing Catalan first, Spanish first, or using a single language tag (either Spanish or Catalan).

4. Experiments

We performed multiple fine-tuning and decoding experiments to evaluate different strategies for handling Catalan-Spanish CS. Our objective was to identify the most effective approach for minimizing transcription errors. To achieve this, we fine-tuned Whisper Large v3 [39] on three different datasets, as shown in Table 2.

Table 2 presents the results of our fine-tuning experiments.

³<https://github.com/iver56/audiomentations>

⁴https://github.com/NVIDIA/NeMo/tree/bc4bce71d01234f568c1327f0848001d86143b3d/scripts/speech_recognition/code_switching

The “Base model” refers to the original Whisper model without additional fine-tuning. The “Synthetic data” approach involves training on artificially generated Catalan-Spanish CS speech, as described in Section 3.2. The “TV3” dataset consists of real-world broadcast speech containing natural CS instances, making it a strong candidate for domain adaptation after fine-tuning. Finally, the “Tuples” dataset (Section 3.3) is built from read speech of multiple native speakers, and its structure encourages model generalization across different CS patterns.

In the decoding phase, we investigated multiple tokenization strategies to determine their impact on transcription accuracy. Specifically, we experimented with dual-language tokens (<ca><es> and <es><ca>), which explicitly specify both languages in the sequence. We also tested single-language tokens (<ca> or <es>), where only one language is enforced during decoding. Finally, we evaluated the model’s behavior when no explicit language token was provided.

5. Analysis of the Results

Table 2 presents the WER scores across three test sets (TV3, ParlamentParla, and Corts). Several key observations emerge from the results. First, fine-tuning significantly improves recognition accuracy compared to the base model across all datasets.

The Synthetic Data approach yields competitive results, particularly for ParlamentParla (14.48% WER), highlighting the potential of synthetic augmentation for domain adaptation.

The model fine-tuned on TV3 achieves the best performance only on the TV3 test set (21.20% WER), implying that it does not adapt well to other real-world CS scenarios. When the 4 rows corresponding to the TV3 results are removed, the best result is obtained by the Synthetic Data approach (23.22% WER).

Conversely, the Tuples-based model does not generalize as effectively, showing higher WER than both the TV3 and synthetic models. This may be because randomly concatenating audio segments introduces inconsistencies within the transcript, negatively impacting the ASR model’s performance as it relies on contextual information for transcription.

In terms of decoding strategies, the ordering of the language tags (<ca><es> vs. <es><ca>) has minimal impact, as both configurations yield comparable results. However, the best results are consistently obtained with the language token <ca>.

In contrast, the language token <es> consistently leads to substantial performance degradation, with WER increasing dramatically to 63.16% in the Base model and 51.29% in the TV3 model on the TV3 test set. This suggests that the model struggles when forced to assume Spanish transcription in a Catalan-prevalent CS scenario.

It is also worth noting that the Synthetic Data approach provides the best results using only 17 hours of data, compared to the 100 hours of read speech in the Tuples approach or the 4 hours of spontaneous speech from TV3. This is a remarkable finding, which warrants further exploration for other languages in future experiments.

However, the 4 hours of audio in the TV3 dataset may not be enough to achieve optimal results, and adding more code-switched data to the training might effectively make this fine-tuned model more generalizable to other datasets. Moreover, integrating the code-switched data in *ParlamentParla*, *Corts Valencianes* and *TV3 Parla* for fine-tuning in future research could potentially lead to even better performance.

6. Conclusions and Further Work

In this work, we have presented a broad overview of the current situation of code-switching in the Catalonia region through an extensive literature review. We have described the datasets used for our experiments, which are publicly available to the community via the Hugging Face platform. We detailed the three methodologies explored in this work: generating synthetic code-switched data, concatenating audio recordings in Catalan and Spanish, and leveraging real code-switched data by detecting language tokens in transcriptions. With this in mind, we present a series of results that suggest the best strategy for handling Catalan-Spanish code-switching is to generate a modest amount of synthetic data using the methodology described in Section 3.2, and to use the language token for the more prevalent language, in this case, Catalan (<ca>).

As future work, we aim to apply these findings to creating ASR models and datasets that involve more languages. However, this work represents a modest step in the right direction, and we hope it contributes to the proper assimilation of code-switching by ASR systems of all kinds.

7. Acknowledgements

This work is funded by the Ministerio para la Transformación Digital y de la Función Pública and Plan de Recuperación, Transformación y Resiliencia - Funded by EU – NextGenerationEU within the framework of the project ILENIA with reference 2022/TL22/00215337 and by the Government of Catalonia through the Aina project.

8. References

- [1] K. A. Woolard, “Codeswitching,” *A companion to linguistic anthropology*, pp. 73–94, 2004.
- [2] K. Hall and C. Nilep, “Code-switching, identity, and globalization,” *The handbook of discourse analysis*, pp. 597–619, 2015.
- [3] S. Poplack, “‘sometimes i’ll start a sentence in spanish y termino en español’: Toward a typology of code-switching,” *Linguistics*, vol. 51, no. s1, pp. 11–14, 2013.
- [4] C. E. Hughes, E. S. Shaunessy, A. R. Brice, M. A. Ratliff, and P. A. McHatton, “Code switching among bilingual and limited english proficient students: Possible indicators of giftedness,” *Journal for the Education of the Gifted*, vol. 30, no. 1, pp. 7–28, 2006.
- [5] S. Ruba, S. Prabhu, and K. Samayan, “Types of code-switching among young adults with bilingualism,” *Rupkatha Journal on Interdisciplinary Studies in Humanities*, vol. 15, no. 3, 2023.
- [6] E. Aguilar-Mediavilla, A. Sánchez Pedroche, L. Buil-Legaz, J. A. Pérez-Castelló, and D. Adrover-Roig, “Chapter 12. code-switching and code-mixing in bilingual spanish–catalan children with and without developmental language disorder,” in *Language Acquisition in Romance Languages*. John Benjamins Publishing Company, 2024, pp. 283–305.
- [7] E. Aguilar-Mediavilla, L. Buil-Legaz, J. Pérez-Castelló, R. López-Penadés, and D. Adrover-Roig, “Code-switching in bilingual spanish-catalan children with specific language impairment,” in *Poster presented at the meeting Bilingualism and Specific Language Impairment, Tours, France*, 2015.
- [8] N. Caballero and M. L. Celaya, “Code-switching by primary school bilingual efl learners: a study on the effect of proficiency and modality of interaction,” *International Journal of Bilingual Education and Bilingualism*, vol. 25, no. 1, pp. 301–313, 2022.
- [9] F. Ali, “Constructing identity through code choice and code-switching: Evidence from multilingual muslim women in barcelona,” *Revista Española de Lingüística Aplicada/Spanish Journal of Applied Linguistics*, vol. 36, no. 1, pp. 204–233, 2023.

- [10] L. Arnaus Gil and A. Jiménez-Gaspar, “The role of internal and external factors for code-switching: A study of early multilingualism in germany with special reference to catalan as a heritage language,” *Languages*, vol. 7, no. 4, p. 258, 2022.
- [11] C. Curcó, “Code switching and identity in the discourse of catalan immigrants in mexico,” *Aila Review*, vol. 18, no. 1, pp. 18–40, 2005.
- [12] J. A. Argenter, “Code-switching and dialogism: Verbal practices among catalan jews in the middle ages,” *Language in society*, vol. 30, no. 3, pp. 377–402, 2001.
- [13] K. A. Woolard, “Changing forms of codeswitching in catalan comedy,” *Catalan Review*, vol. 9, no. 2, pp. 223–252, 1995.
- [14] —, “Codeswitching and comedy in catalonia,” *Codeswitching: Anthropological and sociolinguistic perspectives*, vol. 48, pp. 53–76, 1988.
- [15] M. Millán, “Social meaning behind catalan-spanish code-mixing/code-switching in the catalan television series merlí,” *Treballs de Sociolingüística Catalana*, no. 30, pp. 297–309, 2020.
- [16] C. Pérez-Sabater, “Mixing catalan, english and spanish on whatsapp: A case study on language choice and code-switching,” *Spanish in Context*, vol. 19, no. 2, pp. 289–313, 2022.
- [17] B. Külebi, “ParlamentParla - Speech corpus of Catalan Parliamentary sessions,” Oct. 2021. [Online]. Available: <https://doi.org/10.5281/zenodo.5541827>
- [18] O. Günes, A. Öktem, P. Dogan Özmen, and C. Basiana, “Collectivat (official webpage),” 2024. [Online]. Available: <https://collectivat.cat/>
- [19] I. Feldhausen, *Sentential form and prosodic structure of Catalan*. John Benjamins Publishing Company, 2010.
- [20] B. Külebi, C. Armentano-Oller, C. Rodríguez-Penagos, and M. Villegas, “Parlament parla (hugging face),” 2021. [Online]. Available: https://huggingface.co/datasets/projecte-aina/parlament_parla
- [21] B. Külebi, “Parlament parla (github),” 2021. [Online]. Available: <https://github.com/CollectivaT-dev/ParlamentParla>
- [22] K. Baybars, “Corts valencianes - speech corpus for valencian asr,” 2024. [Online]. Available: https://huggingface.co/datasets/projecte-aina/corts_valencianes_asr_a
- [23] B. Külebi and A. Öktem, “Tv3 parla,” 2021. [Online]. Available: https://huggingface.co/datasets/collectivat/tv3_parla
- [24] A. Quintà and L. Prenafeta, “Tv3 (cataluña),” 1983. [Online]. Available: <https://www.3cat.cat/tv3/>
- [25] B. Külebi and A. Öktem, “Building an open source automatic speech recognition system for catalan,” in *IberSPEECH*, 2018, pp. 25–29.
- [26] R. Ardila, M. Branson, K. Davis, M. Henretty, M. Kohler, J. Meyer, R. Morais, L. Saunders, F. M. Tyers, and G. Weber, “Common voice: A massively-multilingual speech corpus,” in *Proceedings of the 12th Conference on Language Resources and Evaluation (LREC 2020)*, 2020, pp. 4211–4215.
- [27] M. Foundation, “Common voice 17 dataset,” 2024. [Online]. Available: https://huggingface.co/datasets/mozilla-foundation/common_voice_17_0
- [28] J. D. M.-W. C. Kenton and L. K. Toutanova, “Bert: Pre-training of deep bidirectional transformers for language understanding,” in *Proceedings of naacL-HLT*, vol. 1, no. 2. Minneapolis, Minnesota, 2019.
- [29] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, “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)*. Minneapolis, Minnesota: Association for Computational Linguistics, Jun. 2019, pp. 4171–4186. [Online]. Available: <https://aclanthology.org/N19-1423/>
- [30] M. Honnibal, “spacy 2: Natural language understanding with bloom embeddings, convolutional neural networks and incremental parsing,” (*No Title*), 2017.
- [31] H. Face, “Marian mt model documentation,” <https://huggingface.co/docs/transformers/model.doc/marian>, 2025, accessed: 2025-01-16.
- [32] S. Mehta, R. Tu, J. Beskow, É. Székely, and G. E. Henter, “Matcha-tts: A fast tts architecture with conditional flow matching,” in *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2024, pp. 11 341–11 345.
- [33] A. Peiró-Lilja, J. Giraldo, M. Llopart-Font, C. Armentano-Oller, B. Külebi, and M. Farrús, “Multi-speaker and multi-dialectal catalan tts models for video gaming,” in *Interspeech 2024*, 2024, pp. 999–1000.
- [34] C. D. Hernandez Mena, “Librivox spanish corpus: Audio and transcriptions taken from librivox.org,” 2020. [Online]. Available: <https://catalog.ldc.upenn.edu/LDC2020S01%7D>
- [35] A. Bonafonte, J. Adell, I. Esquerria, S. Gallego, A. Moreno, and J. Pérez, “Corpus and voices for Catalan speech synthesis,” in *Proceedings of the Sixth International Conference on Language Resources and Evaluation (LREC’08)*. Marrakech, Morocco: European Language Resources Association (ELRA), May 2008.
- [36] A. Peiró-Lilja, M. Llopart-Font, C. Armentano-Oller, J. Giraldo, I. Esquerria, M. Farrús, and B. Külebi, “Lafrescat: A catalan multi-accent speech dataset for text-to-speech,” in *IberSPEECH 2024*, 2024, pp. 206–210.
- [37] Y. Yang, Y. Peng, H. Huang, E. S. Chng, and X. Zhong, “Adapting openai’s whisper for speech recognition on code-switch mandarin-english seame and asru2019 datasets,” in *2024 Asia Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA ASC)*. IEEE, 2024, pp. 1–6.
- [38] P. Peng, B. Yan, S. Watanabe, and D. Harwath, “Prompting the hidden talent of web-scale speech models for zero-shot task generalization,” *arXiv preprint arXiv:2305.11095*, 2023.
- [39] OpenAI, “Whisper-large-v3,” 2023. [Online]. Available: <https://huggingface.co/openai/whisper-large-v3>