

EmphAssess : a Prosodic Benchmark on Assessing Emphasis Transfer in Speech-to-Speech Models

Maureen de Seyssel^{*1,2} Antony D’Avirro¹ Adina Williams¹ Emmanuel Dupoux^{1,2}

¹Meta AI Research

²ENS, EHESS, CNRS, PSL University, France

maureen.deseysse@gmail.com {adavirro, adinawilliams, dpoux}@meta.com

Abstract

We introduce EmphAssess, a prosodic benchmark designed to evaluate the capability of speech-to-speech models to encode and reproduce prosodic emphasis. We apply this to two tasks: speech resynthesis and speech-to-speech translation. In both cases, the benchmark evaluates the ability of the model to encode emphasis in the speech input and accurately reproduce it in the output, potentially across a change of speaker and language. As part of the evaluation pipeline, we introduce EmphaClass, a new model that classifies emphasis at the frame or word level.

1 Introduction

In recent years, significant advancements have been made in the development of Self-Supervised Learning (SSL) models for speech, extending beyond the traditional text-only methods prevalent in the field (Mohamed et al., 2022). Such speech-based models find successful application across various domains from generative language modelling (Lakhotia et al., 2021; Borsos et al., 2023; Nguyen et al., 2023b) to speech-to-speech translation (S2ST) (Jia et al., 2019, 2022; Lee et al., 2021; Rubenstein et al., 2023; Barrault et al., 2023). Unlike text-only models, they exploit additional cues present in the speech signal which are absent in textual input.

One crucial speech-only cue is prosody. Also termed the “music of speech” (Wennerstrom, 2001), prosody is marked by the perceived loudness, rhythm, and pitch of speech. Prosody not only adds naturalness to an utterance but also has the capacity to modify the meaning of the conveyed message, both at a global level, such as in the expression of different emotions, and at a local level, by influencing the interpretation of individual phrases or words (Cutler et al., 1997; Dahan, 2015). For

instance, slower speech may suggest hesitation, while altering something like pause placement can actually change the segmentation into words or syntactic constituents, with downstream consequences for the meaning. Hence, accurately capturing these prosodic elements is essential in SSL speech models for any application (Avila and Ward, 2023).

To address this, Kharitonov et al. (2021) proposed explicitly adding prosodically-relevant information such as fundamental frequency and duration to the speech representations models learn, while others aimed at explicitly modelling emotions in such representations (Gan et al., 2022; Duret et al., 2023). Although some progress has been made, robust evaluation metrics for prosody remain scarce, and human evaluation, while insightful, is subjective - which can limit reproducibility; as well as being expensive and time intensive - which can hinder its utility in large-scale applications.

Objective evaluations of prosody fall into two main categories: one focuses on utterance-level features like emotion and speech rate to assess *global prosody*, and the other examines *local prosody*, which is concerned with prosodic effects at the level of a word or a phrase, such as breaks, turn ends and emphasis. In addition, one may address prosody for two classes of models: generative decoder-only models (the speech equivalent of GPT (Radford et al., 2018) (e.g. GSLM, Lakhotia et al., 2021; AudioLM, Borsos et al., 2023; dGSLM, (Nguyen et al., 2023b)), and speech-to-speech (encoder-decoder) approaches, which take speech as input and produce output in a different voice (speech resynthesis) or a different language (S2ST). In this paper, we address the second class of models.

In the context of speech-to-speech (S2S) models, evaluating global prosody can be relatively straightforward, as the features are not directly related to the lexical content. The assessment of local prosody, however, presents more of a challenge, as

^{*}Currently at Apple

it necessitates mapping at the lexical level. This can be relatively feasible in the context of speech resynthesis, where the model directly reconstructs the input signal and, therefore, preserves lexical content (e.g., by correlating prosodic attributes such as duration and fundamental frequency (F0) between input and output utterances; Suni et al., 2020). However, this becomes more complicated when evaluating S2ST models, as one needs to ensure the correct prosodic feature is applied to the correct word(s) (Duret et al., 2023) (alignment problem).

Although scarce, there have been recent efforts made to establish benchmarks in the prosodic evaluation of speech models allowing models comparison, including evaluation corpora and pipelines, both at the global prosodic level (pragmatic information : Lin et al. (2023)) and at the local prosodic level (prosodic pauses: de Seyssel et al. (2023)). Yet, there is a need for more benchmarks to cover other aspects of prosody, and all types of speech models.

In this work, we introduce the EmphAssess benchmark, which is focused on local prosody for speech-to-speech models and includes: (i) a new, automatic pipeline for emphasis evaluation that is modular, handles multiple languages and kinds of outputs (including paraphrases and translations, (ii) a novel dataset, the EmphAssess test set, for evaluating model emphasis preservation in English and Spanish according to our pipeline, and (iii) EmphaClass, an emphasis classifier that we finetuned with English data over an existing multilingual SSL model to support our pipeline.

2 Background

Emphasis as a prosodic feature. Emphasis, the phonetically-realized importance given to particular words or phrases, is critical for interpreting language. Some of the most important correlates of emphasis are fundamental frequency (f0), duration, and amplitude (Terken and Hermes, 2000; Mo, 2008), although the weight and behaviour of each can vary across languages (Ladd and Arvaniti, 2023). These acoustic attributes collectively shape the prosodic contours that signal emphasis in speech. Altering the emphasis in a sentence such as “I never said he stole my bag” from “he” to “stole” can drastically change its meaning. Such nuances are essential for models to process, if they are to have an accurate representation of speech, be they generative language models or S2ST systems.

In fact, the issue of accurate emphasis transfer in S2ST models has attracted some research attention over the years. Studies by Tsiartas et al. (2013); Do et al. (2016, 2018) approach this topic using cascaded models (with separate Automatic Speech Recognition, Machine Translation, and Text-to-Speech models). A more recent approach by Huang et al. (2023) integrates the two first components into a single encoder module capable of multilingual embeddings. Similar to other prosodic features, emphasis in S2S models is primarily evaluated through human evaluation (Tsiartas et al., 2013; Huang et al., 2023), although Do et al. (2016, 2018) proposed leveraging an emphasis classification algorithm to calculate F1 scores by matching emphasised words in the input and output utterances. Yet, this method is limited to a single language pair and cannot handle variations in translation outputs, only recognising one “gold” translation per dataset utterance. Consequently, this metric is ill-suited for comprehensive automatic benchmarking across various models.

Word-level emphasis classification. As suggested by Do et al. (2016, 2018), a robust word-level emphasis classification system is critical in automatic evaluation of emphasis transfer in S2ST models. Existing algorithms, predominantly designed for text-to-speech applications, often rely on traditionally engineered features (e.g. MFCCs or Fbanks), sometimes augmented with other prosodic-related information (e.g. F0, duration) (Do et al., 2016; Heba et al., 2017; Ning et al., 2017; Zhang et al., 2018). Some also incorporate lexical information from textual transcripts (Brenier et al., 2005; Zhou et al., 2020). However, these models frequently suffer from limited generalisability across different datasets, voice types, and languages. There is a compelling argument for using the speech waveform directly as input to enhance generalisability. To our knowledge, the only study to have adopted this approach is that of Vaidya et al. (2022), which employed a CRNN framework for classifying emphasis in children’s speech; their work, however, was limited to a single language (and is not open-sourced). We propose that leveraging pretrained models trained on multilingual datasets could result in significant advancements in this field.

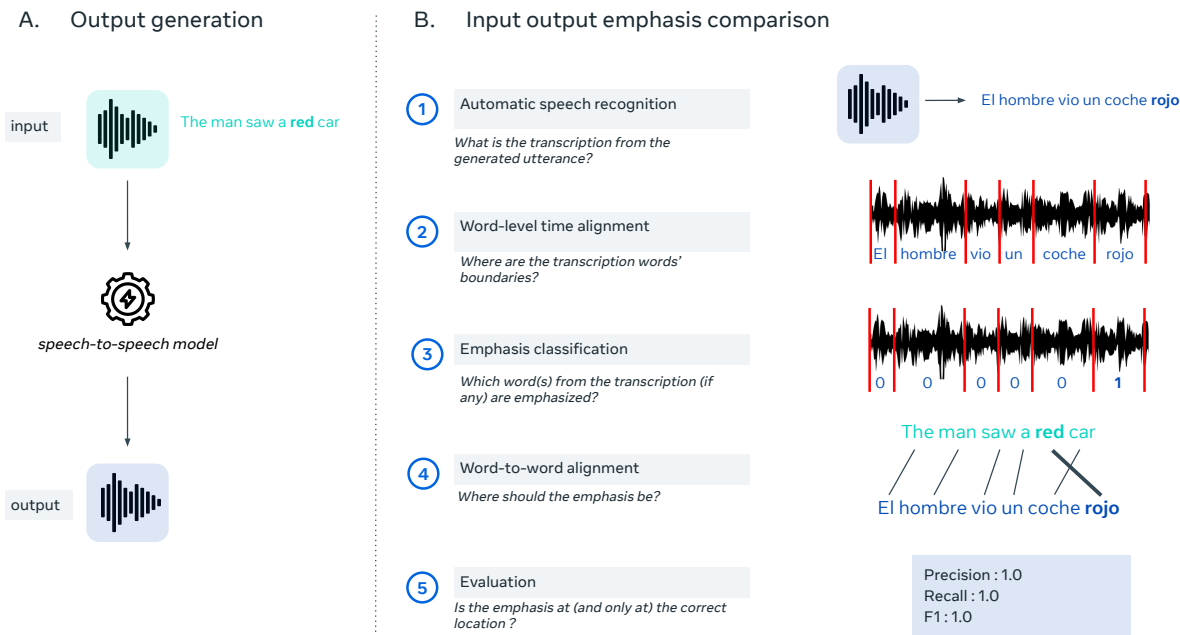


Figure 1: **Overview of the EmphAssess evaluation pipeline.** Left panel : Output generation. Right panel : Input-output emphasis comparison.

3 Introducing EmphAssess

In this study, we introduce EmphAssess, a versatile automatic benchmark for evaluating emphasis preservation in S2S models, including S2ST ones. Essentially, this benchmark comprises a carefully curated dataset of English utterances with emphasised words, accompanied by an automatic evaluation pipeline, and results on some of the most recent S2S SSL models. Our evaluation framework, inspired by the methodology of Do et al. (2016, 2018), assesses emphasis alignment between the source and the model’s output utterances. Our benchmark’s novelty lies in its capacity to handle various output types, including paraphrases and translations.

Guided by the data we have for setting optimal baselines, the EmphAssess benchmark is specifically designed for English-to-English and English-to-Spanish S2S models. However, our work goes further, laying the groundwork for extending this benchmark to other language pairs. Moreover, the evaluation pipeline itself is already capable of being applied to a broad spectrum of language pairs. Also, while we focus here on unsupervised speech language models, EmphAssess is versatile enough to be applied to any S2S framework.

The EmphAssess evaluation pipeline’s modular structure is a key feature, with each module

designed to function independently and allow for straightforward modifications. We leverage a suite of distinct open-source models, each finetuned for particular tasks. The pipeline can therefore be upgraded to incorporate improvements in each module seamlessly. Although such enhancements may necessitate a re-evaluation of the models within our benchmark, this inherent adaptability is a considerable benefit, ensuring EmphAsses can remain current with the latest research for years to come.

Finally, we introduce and open-source, as part of this automatic evaluation pipeline, a novel emphasis classifier at the word level: EmphaClass. This classifier is finetuned over an existing multilingual SSL model with the hope of enhancing its robustness across multiple languages and variability.

The evaluation code, emphasis classifier and dataset introduced in this paper are available in our related repository ¹.

4 The EmphAssess Dataset

The EmphAssess dataset comprises synthetically generated speech utterances, each containing at least one emphasised word. Accompanying these utterances are metadata detailing the transcription, the positional index of the emphasised word(s), and information about the synthetic voice employed for

¹<https://github.com/facebookresearch/emphassess>

synthesis. In total, the dataset boasts 3652 speech samples derived from 913 unique transcripts (with each transcript being rendered in 4 distinct voices).

The dataset generation started with a selection of transcripts from a list of handwritten transcripts with emphasis annotations² previously created for company-internal Text-to-Speech purposes. Transcripts containing characters beyond letters or specific punctuation marks³ or those featuring proper nouns (identified using the NLTK toolkit; Bird 2006) were excluded, to ensure the translations are as straightforward as possible. Moreover, we ensured a minimum of two distinct versions with different emphases for string identical sentences (those with matching word tokens but possibly differing emphasis position indices). This approach was adopted to mitigate any bias should a model exhibit a preference for emphasising a particular word over others. Finally, we filtered out transcripts that could face alignment challenges with emphasised words during translation. We set up an algorithm to assess the difficulty of aligning emphasised words in an English sentence with their counterparts in multiple target languages, using the SimAlign word-alignment tool (Sabet et al., 2020). Simply put, if an emphasised word in the source matched consistently to a corresponding word across a list of other languages (German, French, Spanish, and Chinese), the sentence was labelled “easy”; otherwise, it was deemed “difficult.” Only “easy” transcripts were retained for our dataset. We were left with 913 distinct transcriptions (with varying emphases) derived from a pool of 299 unique transcriptions. We ensured that the distribution of transcripts was well balanced, in terms of where the emphasis was located.

Next, we employed an internal Text-to-Speech (TTS) tool with a 16 kHz sample rate to synthesise all 913 transcripts, each in the four distinct open-source English Espresso voices (Nguyen et al., 2023a), namely ex01, ex02, ex03 and ex04, resulting in a comprehensive set of 3,652 speech samples.

Finally, we compiled a dataset that is available as part of the benchmark. This dataset comprised four columns: an id column that denotes the unique identifier for each speech segment, a src_sentence column that contains the corresponding tokenised text transcript presented in list

format, a gold_emphasis column that highlights the index of the emphasised word(s) also in list format, and a voice column that specifies the particular Espresso voice employed for the synthesis.

5 The EmphaAssess Evaluation Pipeline

The evaluation pipeline, as illustrated in Figure 1, is divided into two main stages. The first one (left panel) corresponds to the generation of utterances from the evaluated S2S model. That is, for each utterance from the EmphaAssess dataset, we need to generate the corresponding utterance output from the evaluated model. Hence, this inference stage is dependent on the model tested, and we will not expand on it here.

In the second stage (right panel), we perform the automatic evaluation by comparing the input and output utterances. The objective is twofold: firstly, to ascertain whether the emphasis is *retained* in the generated utterance, and secondly, to determine whether the emphasis is *correctly positioned* on the corresponding word. At this stage, available resources include the input (original) utterance, the corresponding output utterance, and the tokenised transcript of the input with the location of the emphasised word(s) identified. A schematic overview of the evaluation pipeline is shown in the right panel of Figure 1. Initially, we obtain a transcription of the generated utterance (1) and the time-aligned word boundaries (2). This information can be used in addition to the raw waveform to detect emphasis at the word level in the output utterance using a classifier (3). At this stage, we must determine which word(s) in the generated utterance should be emphasised to obtain evaluation scores (4). We use word-to-word alignment at the text level to address this, a technique borrowed from the machine translation field. Finally, we can use this information to compute precision, recall and F1 score (5). We will now detail our methodology for each of these steps.

5.1 Automatic speech Recognition and word-level forced time-alignment

To achieve accurate transcription of the generated utterance and its associated word-level time-alignments, we utilise the WhisperX system (Bain et al., 2023). This system, which relies on the weakly supervised speech recognition model Whisper (Radford et al., 2023) for speech transcription, allows retrieval of accurate word-level timestamps,

²The emphasis could be applied to any sentence constituents, but it followed a contrastive pattern.

³Retained punctuation characters include: [,:;!()]

in a variety of languages.

5.2 Word Emphasis Classification

As the next step requires detecting emphasis at the word level from the waveform and its corresponding transcription, we propose EmphaClass, a new model for emphasis classification. Our approach was centred around finetuning a pretrained SSL speech model through a frame-classification task to classify a frame as either emphasised or not. We can then aggregate frame-level scores to derive word-level emphasis classifications.

Data. We utilised speech sourced from the English Expressive Espresso dataset (Nguyen et al., 2023a). Indeed, this dataset comprises utterances that contain emphasised words, accompanied by their annotations, presented in a diverse range of speaking styles. We retained only those utterances that had at least one word emphasised. We divided the four speakers into two for validation (ex03 and ex04) and two for the test set (ex01 and ex02). Additionally, we had utterances from six other speakers recorded under identical conditions and with similar emphasis annotations. These were utilised to create an internal training set, amounting to 2.06 hours of speech. We then used the Montreal Forced Aligner to align the transcription with the audio and obtain reliable word boundaries (McAuliffe et al., 2017). We subsequently processed the data to provide annotations at the frame level regarding emphasis. We deem a frame as ‘emphasised’ if it falls within a word annotated as such, with each frame corresponding to 20ms of speech.

Emphasis classifier architecture. We finetuned the multilingual SSL speech model, XLS-R (Babu et al., 2021), grounded in the Wav2Vec 2.0 architecture (Baevski et al., 2020). This finetuning encompassed a binary frame classification task using cross-entropy loss, and was carried out using the Wav2Vec2ForAudioFrameClassification method from HuggingFace (Wolf et al., 2019). Our choice of the XLS-R model for extended training and evaluations stemmed from its exceptional performance metrics and promising potential for cross-language generalisation.

Evaluation. We use F1 score as the primary metric for evaluating our emphasis classifier, both at the frame and word level. For word-level classification, we compute the average accuracy of the frames within the boundaries of each word. A word was

deemed emphasised if more than 50% of its frames were classified as such. A representative example of this classification is illustrated in Figure 2.

We evaluate the classifier on our test set split of the Espresso dataset, but also on the utterances used in our EmphAssess dataset. Results are presented in Table 1. The scores suggest that the model performs well at classifying emphasis in both the Espresso dataset 78.4% and the Emphasess dataset 93.48%. The lower scores from the Espresso dataset, compared to the EmphAssess dataset, can be attributed to two factors. Firstly, the Espresso dataset incorporates utterances with speaking styles where the emphasis is notably challenging to discern, such as whispering and laughing. Secondly, using synthetic voices in EmphAssess might offer more consistent and clearer patterns of emphasis than the natural utterances from Espresso, making it easier for the classifier to discern, and thus leading to higher accuracy scores.

Test data	Frame-level (%)			Word-level(%)		
	F1	Prec.	Rec.	F1	Prec.	Rec.
EmphAssess	89.77	89.71	91.72	93.48	93.81	94.04
Espresso EN	75.52	60.82	76.90	78.40	56.93	76.90

Table 1: Results of EmphaClass on The EmphAssess dataset and a subset of the Espresso dataset. F1 score, precision and recall

We also ran cross-languages analyses, testing the model on other languages, which results showed that the model can, to some extent, classify other languages. This suggests our research may have utility beyond just the English and Spanish languages we explicitly support. More information is presented in Appendix A.

5.3 Word-to-word alignment

Returning to the automatic emphasis evaluation pipeline, we can detect which word(s) is emphasised in an output utterance with the classifier described above, given a waveform, its transcriptions and word boundaries. At this point, we need to identify which word(s) should be emphasised in the output utterance to compute a score for the quality of emphasis transfer. This step is vital because it lets us evaluate any output utterance, including paraphrases and translations, without being limited to a “gold” output. To do this, we use a word-to-word alignment algorithm, often seen in machine translation, especially the SimAlign one (Sabet et al., 2020). This tool can align words

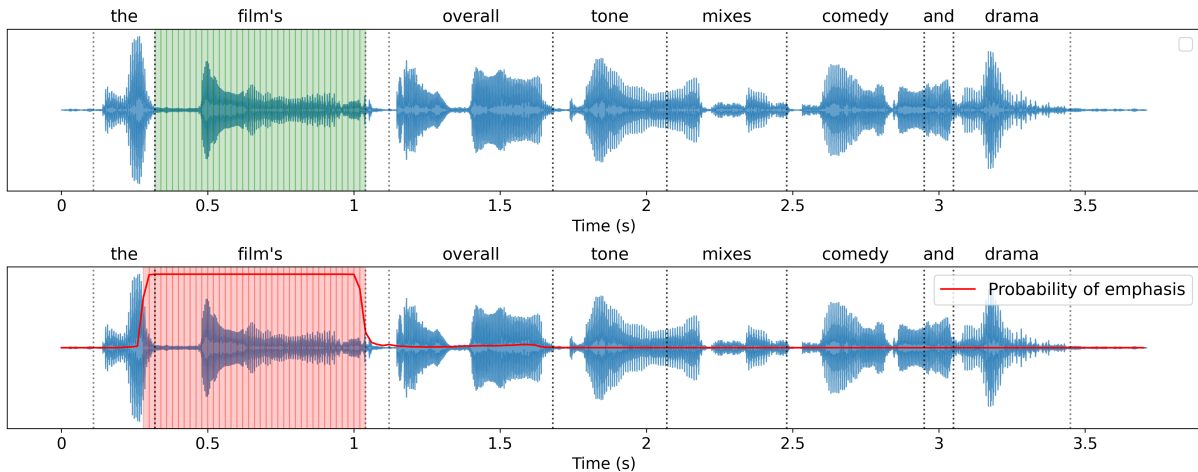


Figure 2: **Illustrative example of emphasis classification with the trained classifier.** Top: gold annotations. Bottom: Emphasis classifier predictions.

between two text sentences. Although typically used in machine translation, it’s also effective for paraphrases in the same language. A key benefit of SimAlign is that it works across many languages without requiring finetuning. For our needs, we compare the original text input with the output utterance transcription from the ASR to see which word(s) match the emphasised word in the original sentence.

5.4 Metrics

In the final step, we compare the words that were meant to be emphasised (from the previous step) with the words that were actually emphasised (from the emphasis classification phase). By doing this comparison, we can determine precision, recall, and F1 scores for the whole dataset.

6 Results

We benchmarked a series of models on the EmphAssess evaluation, both within language (English to English) and using translations (English to Spanish).

6.1 English S2S models

We first present results on models that generate speech with the target and source language being identical, here English (left panel of Figure 3). This encompasses models that undergo an encoding-decoding method, simply resynthesising the learnt units and those which can learn paraphrases.

For a topline evaluation, we matched the input utterances from EmphAssess with themselves (that is, we pretended the output utterances were the same

as the input ones). This gave us an insight into the best achievable scores, with any potential loss in performance due to problems in the dataset or the various comparison stages. This topline produced an F1 score of 89%, indicating that our cascaded pipeline performs well. It should also be noted that we consider chance-level to yield scores of 0, corresponding to a model which does not encode emphasis and thus should not produce any emphasis.

We first assessed the generative GSLM model (Nguyen et al., 2023b), specifically the HuBERT, 100 units version. This model initially encodes speech into continuous forms using HuBERT (Hsu et al., 2021), which are then quantised into units for language modelling. Subsequently, a synthesiser converts these units back to speech. In our study, we extracted the quantised representations from our EmphAssess dataset’s speech samples and directly resynthesised them, bypassing the generative language modelling phase. Despite scoring notably lower than the topline with an F1 of 42%, the model successfully transferred some emphasis to the output utterances. This indicates the presence of prosodic information within these units learned from SSL speech model, a finding supported by de Seyssel et al. (2022, 2023).

We also assessed the pGSLM variant, which incorporates extra prosodic features during training to enhance prosody modelling (Kharitonov et al., 2021)⁴. Notably, the pGSLM models achieved scores close to the topline, with an F1 of 88%,

⁴We opted for the variant with continuous input and shift, as it was the top performer in de Seyssel et al. (2023).




Figure 3: **Precision, recall and F1 scores on the EmphAssess benchmark.** Left : English-to-English models and English Emphasis classifier. Right : English-to-Spanish models and Spanish Emphasis classifier.

highlighting their excellent proficiency in encoding emphasis accurately.

Finally, we assessed the Seamless M4T model (Barrault et al., 2023), forcing it to generate outputs in English. Contrary to the previous models, which generate output constrained in their lexical input, this one is primarily a S2ST model and can output paraphrases. We did not expect these models to encode any prosodic information given to their architecture, an expectation which was actually supported by a very low score on EmphAssess (18%).

6.2 Generalising the pipeline to S2S translation

We now want to discuss how we can adapt our pipeline to S2ST capabilities. While most target languages can be evaluated directly using the existing pipeline, there are several considerations to remember. Firstly, it is essential to establish a validated topline. In other words, when introducing a new target language, we require validated translated utterances of the input English dataset in the desired language to have a topline in this target language. This process necessitates human validation, not only for the text translation, but also to either synthesise or record this translation with the correct emphasis, depending on the available resources. This new set of utterances can additionally serve as an input test set when we want to modify the source language to one other than English.

Furthermore, we might want to modify or adapt some of the stages of the automatic evaluation pipeline in order to be better suited to the new language. For example, we have gathered evidence indicating that the emphasis classifier performs better when trained in the specific language it will be evaluated in. Thus, retraining it with emphasis

data in the target language can prove advantageous, albeit demanding the corresponding larger dataset.

We undertook a two-step process to modify our evaluation for English-to-Spanish translation. Firstly, external annotators translated the input sentences into Spanish, ensuring the inclusion of emphasis annotations. Subsequently, these translated sentences were synthesised into Spanish using our in-house TTS (Text-to-Speech) voices designed for Spanish, with a focus on retaining emphasis. Additionally, we adjusted the emphasis classifier to one specifically trained for Spanish as it yielded better results on Spanish data (see Appendix A).

As depicted in the right panel of Figure 3, the ‘topline,’ which aligns the English input with the synthesised Spanish voices as the output, achieved a score of 58%. While this result is reasonable, it notably lags behind the English topline. This decline may be attributed to various factors, including challenges in the synthesised voices, as we observed that our Spanish TTS voices do not emphasise as effectively as desired. Furthermore, issues in different stages of our automatic evaluation pipeline might contribute (for instance, the Spanish emphasis classifier’s performance on Spanish is not as optimal as its English counterpart on English data). Additionally, linguistic differences could play a role, with Spanish emphasis potentially being less prominent than in English or conveyed through alternative means, possibly paraphrastically in the text itself. Nonetheless, having this topline facilitates the comparison of other models and the assessment of their relative performance. Subsequently, we evaluated the Seamless M4T model (Barrault et al., 2023) in its English-to-Spanish translation capability, which yielded an F1 score of 14%. This result, akin to its English-to-

English counterpart, suggests that the M4T model does not effectively capture emphasis.

6.3 Human Evaluation

To gauge human performance on the task, we conducted an evaluation with expert annotators. These annotators were presented with an utterance and its word-tokenised transcription, and were tasked with marking words they considered to be emphasised. Importantly, they were not obliged to mark any word as emphasised if they didn't perceive any. This evaluation was carried out on a subset of the data, incorporating both English and Spanish utterances, with native annotators for each language. Figure 3 shows precision, recall, and F1 scores for English-to-English and English-to-Spanish, respectively⁵. These metrics were calculated by comparing the annotators' identification of emphasis against the 'gold standard' annotation with which we synthesised the utterances.

Focusing first on the English dataset, the annotators achieved a commendable precision score of 86%, although this was offset by a lower recall score (50%). The lower recall could be attributed to annotators not perceiving emphasis in numerous sentences (Note: it is often harder to perceive emphasis in utterances taken out of their general, wider context); nonetheless, the high precision score is encouraging. Turning our attention to the Spanish dataset, both recall and precision scores were lower. This aligns with our hypothesis that the quality of voice synthesis in Spanish was not up to par - with the larger drop of recall compared to the topline could be explained by the Spanish emphasis classifier model picking up very subtle cues that are not obvious to the human ear. It may also suggest that the nuances of emphasis might be linguistically specific, thereby differing between English and Spanish.

7 Conclusion

We have introduced an evaluation framework for emphasis in speech-to-speech (S2S) models. This framework comprises an English dataset, an automated evaluation pipeline, and a results benchmark focusing on English-to-English and English-to-Spanish models. Crucially, our framework offers a generalisable approach applicable to other language pairs, the only major requirement being

⁵For English-to-Spanish, the human topline is set using a subset of the Spanish utterances synthesized the Spanish topline

the acquisition of a relevant dataset to establish a reliable gold standard.

Additionally, we have open-sourced an emphasis-classification model that has been finetuned on English data. The model builds on a multilingual SSL architecture and has shown impressive accuracy in classifying emphasised speech in English on our dataset, along with reasonable performance in other languages (for further details, refer to the Appendix). The model's robustness in English makes it a plausible starting point for finetuning classifiers in other languages, potentially minimising the volume of data needed for training. Interestingly, the fact that the successful results were achieved without retraining the encoder, suggests that the inherent features in the original XLS-R model were adequate for emphasis classification.

There is an existing agenda for future research centring around the evaluation of prosody within SSL models. Firstly, on the subject of emphasis, we aim to scrutinise its functional role more closely—specifically, its ability to convey importance. We intend to investigate whether such a function is intrinsically represented within these models. Beyond emphasis, other aspects of prosody, such as turn-taking and speech grouping, merit attention. We are interested in determining whether these elements, too, are encoded within SSL models. Improved benchmarks and evaluations for these prosodic features could pave the way for the development of more expressive and nuanced models.

To conclude, the EmphAssess benchmark sets a new standard for the evaluation of prosodic features in S2S models, offering both methodological contributions and actionable insights that could pave the way for more natural and effective machine-generated speech across various applications.

8 Limitations

While pioneering in its approach to evaluating emphasis in S2S models, our study encounters certain limitations. First, the emphasis classifier presented in this paper was made to be used with this exact dataset, and we recommend constraining its use to this particular use case (that is, with the presented benchmark and evaluation pipeline). Indeed, further testing is required to enhance its robustness and ensure its efficacy in detecting more nuanced forms of emphasis across other datasets.

Furthermore, the robustness of our evaluation

process relies on the quality of multiple pipeline components, including Automatic Speech Recognition, forced alignment, and word-to-word alignment. Therefore, it is crucial to be mindful that errors could arise at various stages. Yet, the modular nature of the pipeline allows for continual improvements and assures that inter-model comparisons remain valid.

Another limitation of our work lies in the use of synthesised speech to create our dataset. While this approach provides a more controlled and consistent dataset—for instance, by enabling the synthesis of identical textual content with varying word emphases and voices—it may fail to capture the full range of characteristics found in natural speech. Consequently, this limitation could affect how well the benchmark results can be applied to practical use cases.

Lastly, our study is currently limited to binary categorisation of emphasis. Future endeavours could explore varying degrees of emphasis, although this would require more advanced models. For instance, capturing subtle differences in emphasis between the input and output of an S2S system could be a valuable addition to this line of research.

Acknowledgements

ED in his EHESS capacity has been funded by the Agence Nationale pour la Recherche (ANR-17-EURE-0017 Frontcog, ANR-10-IDEX-0001-02 PSL*, ANR-19-P3IA-0001 PRAIRIE 3IA Institute) and a grant from CIFAR (Learning in Machines and Brains).

References

- Jonathan E Avila and Nigel G Ward. 2023. Towards cross-language prosody transfer for dialog. *arXiv preprint arXiv:2307.04123*.
- Arun Babu, Changhan Wang, Andros Tjandra, Kushal Lakhota, Qiantong Xu, Naman Goyal, Kritika Singh, Patrick von Platen, Yatharth Saraf, Juan Pino, et al. 2021. Xls-r: Self-supervised cross-lingual speech representation learning at scale. *arXiv preprint arXiv:2111.09296*.
- Alexei Baevski, Yuhao Zhou, Abdelrahman Mohamed, and Michael Auli. 2020. wav2vec 2.0: A framework for self-supervised learning of speech representations. *Advances in neural information processing systems*, 33:12449–12460.
- Max Bain, Jaesung Huh, Tengda Han, and Andrew Zisserman. 2023. Whisperx: Time-accurate speech transcription of long-form audio. *arXiv preprint arXiv:2303.00747*.
- Loïc Barrault, Yu-An Chung, Mariano Cora Meglioli, David Dale, Ning Dong, Paul-Ambroise Duquenne, Hady Elsahar, Hongyu Gong, Kevin Heffernan, John Hoffman, et al. 2023. Seamless4t-massively multilingual & multimodal machine translation. *arXiv preprint arXiv:2308.11596*.
- Steven Bird. 2006. Nltk: the natural language toolkit. In *Proceedings of the COLING/ACL 2006 Interactive Presentation Sessions*, pages 69–72.
- Zalán Borsos, Raphaël Marinier, Damien Vincent, Eugene Kharitonov, Olivier Pietquin, Matt Sharifi, Dominik Roblek, Olivier Teboul, David Grangier, Marco Tagliasacchi, et al. 2023. Audioldm: a language modeling approach to audio generation. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*.
- Jason M Brenier, Daniel M Cer, and Daniel Jurafsky. 2005. The detection of emphatic words using acoustic and lexical features. In *Ninth European Conference on Speech Communication and Technology*.
- Anne Cutler, Delphine Dahan, and Wilma Van Donseelaar. 1997. Prosody in the comprehension of spoken language: A literature review. *Language and speech*, 40(2):141–201.
- Delphine Dahan. 2015. Prosody and language comprehension. *Wiley Interdisciplinary Reviews: Cognitive Science*, 6(5):441–452.
- Maureen de Seyssel, Marvin Lavechin, Yossi Adi, Emmanuel Dupoux, and Guillaume Wisniewski. 2022. Probing phoneme, language and speaker information in unsupervised speech representations. In *InterSpeech 2022*.
- Maureen de Seyssel, Marvin Lavechin, Hadrien Titeux, Arthur Thomas, Gwendal Virlet, Andrea Santos Revilla, Guillaume Wisniewski, Bogdan Ludusan, and Emmanuel Dupoux. 2023. Prosaudit, a prosodic benchmark for self-supervised speech models. In *Interspeech 2023*.
- Quoc Truong Do, Sakriani Sakti, and Satoshi Nakamura. 2018. Sequence-to-sequence models for emphasis speech translation. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 26(10):1873–1883.
- Quoc Truong Do, Tomoki Toda, Graham Neubig, Sakriani Sakti, and Satoshi Nakamura. 2016. Preserving word-level emphasis in speech-to-speech translation. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 25(3):544–556.
- Jarod Duret, Benjamin O’Brien, Yannick Estève, and Titouan Parcollet. 2023. Enhancing expressivity transfer in textless speech-to-speech translation. *arXiv preprint arXiv:2310.07279*.

- Wendong Gan, Bolong Wen, Ying Yan, Haitao Chen, Zhichao Wang, Hongqiang Du, Lei Xie, Kaixuan Guo, and Hai Li. 2022. Iqddubbing: Prosody modeling based on discrete self-supervised speech representation for expressive voice conversion. *arXiv preprint arXiv:2201.00269*.
- Abdelwahab Heba, Thomas Pellegrini, Tom Jorquera, Régine André-Obrecht, and Jean-Pierre Lorré. 2017. Lexical emphasis detection in spoken french using f-banks and neural networks. In *International Conference on Statistical Language and Speech Processing*, pages 241–249. Springer.
- Wei-Ning Hsu, Benjamin Bolte, Yao-Hung Hubert Tsai, Kushal Lakhotia, Ruslan Salakhutdinov, and Abdelrahman Mohamed. 2021. Hubert: Self-supervised speech representation learning by masked prediction of hidden units. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 29:3451–3460.
- Wen-Chin Huang, Benjamin Peloquin, Justine Kao, Changhan Wang, Hongyu Gong, Elizabeth Salesky, Yossi Adi, Ann Lee, and Peng-Jen Chen. 2023. A holistic cascade system, benchmark, and human evaluation protocol for expressive speech-to-speech translation. In *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1–5. IEEE.
- Ye Jia, Michelle Tadmor Ramanovich, Tal Remez, and Roi Pomerantz. 2022. Translatotron 2: High-quality direct speech-to-speech translation with voice preservation. In *International Conference on Machine Learning*, pages 10120–10134. PMLR.
- Ye Jia, Ron J Weiss, Fadi Biadsy, Wolfgang Macherey, Melvin Johnson, Zhifeng Chen, and Yonghui Wu. 2019. Direct speech-to-speech translation with a sequence-to-sequence model. *arXiv preprint arXiv:1904.06037*.
- Eugene Kharitonov, Ann Lee, Adam Polyak, Yossi Adi, Jade Copet, Kushal Lakhotia, Tu-Anh Nguyen, Morgane Rivière, Abdelrahman Mohamed, Emmanuel Dupoux, et al. 2021. Text-free prosody-aware generative spoken language modeling. *arXiv preprint arXiv:2109.03264*.
- D Robert Ladd and Amalia Arvaniti. 2023. Prosodic prominence across languages. *Annual Review of Linguistics*, 9:171–193.
- Kushal Lakhotia, Eugene Kharitonov, Wei-Ning Hsu, Yossi Adi, Adam Polyak, Benjamin Bolte, Tu-Anh Nguyen, Jade Copet, Alexei Baevski, Abdelrahman Mohamed, et al. 2021. On generative spoken language modeling from raw audio. *Transactions of the Association for Computational Linguistics*, 9:1336–1354.
- Ann Lee, Hongyu Gong, Paul-Ambroise Duquenne, Holger Schwenk, Peng-Jen Chen, Changhan Wang, Sravya Popuri, Yossi Adi, Juan Pino, Jiatao Gu, et al. 2021. Textless speech-to-speech translation on real data. *arXiv preprint arXiv:2112.08352*.
- Guan-Ting Lin, Chi-Luen Feng, Wei-Ping Huang, Yuan Tseng, Tzu-Han Lin, Chen-An Li, Hung-yi Lee, and Nigel G Ward. 2023. On the utility of self-supervised models for prosody-related tasks. In *2022 IEEE Spoken Language Technology Workshop (SLT)*, pages 1104–1111. IEEE.
- Michael McAuliffe, Michaela Socolof, Sarah Mihuc, Michael Wagner, and Morgan Sonderegger. 2017. Montreal forced aligner: Trainable text-speech alignment using kald. In *Interspeech*, volume 2017, pages 498–502.
- Yoonsook Mo. 2008. Acoustic correlates of prosodic prominence for naïve listeners of american english. In *Annual Meeting of the Berkeley Linguistics Society*, volume 34, pages 257–267.
- Abdelrahman Mohamed, Hung-yi Lee, Lasse Borgholt, Jakob D Havtorn, Joakim Edin, Christian Igel, Katrin Kirchhoff, Shang-Wen Li, Karen Livescu, Lars Maaløe, et al. 2022. Self-supervised speech representation learning: A review. *IEEE Journal of Selected Topics in Signal Processing*.
- Tu Anh Nguyen, Wei-Ning Hsu, Antony d’Avirro, Bowen Shi, Itai Gat, Maryam Fazel-Zarani, Tal Remez, Jade Copet, Gabriel Synnaeve, Michael Hassid, et al. 2023a. Espresso: A benchmark and analysis of discrete expressive speech resynthesis. *arXiv preprint arXiv:2308.05725*.
- Tu Anh Nguyen, Eugene Kharitonov, Jade Copet, Yossi Adi, Wei-Ning Hsu, Ali Elkahky, Paden Tomasello, Robin Algayres, Benoit Sagot, Abdelrahman Mohamed, et al. 2023b. Generative spoken dialogue language modeling. *Transactions of the Association for Computational Linguistics*, 11:250–266.
- Yishuang Ning, Zhiyong Wu, Runnan Li, Jia Jia, Mingxing Xu, Helen Meng, and Lianhong Cai. 2017. Learning cross-lingual knowledge with multilingual blstm for emphasis detection with limited training data. In *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5615–5619. IEEE.
- Alec Radford, Jong Wook Kim, Tao Xu, Greg Brockman, Christine McLeavey, and Ilya Sutskever. 2023. Robust speech recognition via large-scale weak supervision. In *International Conference on Machine Learning*, pages 28492–28518. PMLR.
- Alec Radford, Karthik Narasimhan, Tim Salimans, Ilya Sutskever, et al. 2018. Improving language understanding by generative pre-training.
- Paul K Rubenstein, Chulayuth Asawaroengchai, Duc Dung Nguyen, Ankur Bapna, Zalán Borsos, Félix de Chaumont Quitry, Peter Chen, Dalia El Badawy, Wei Han, Eugene Kharitonov, et al. 2023. Audiopalm: A large language model that can speak and listen. *arXiv preprint arXiv:2306.12925*.
- Masoud Jalili Sabet, Philipp Dufter, François Yvon, and Hinrich Schütze. 2020. Simalign: High quality

word alignments without parallel training data using static and contextualized embeddings. *arXiv preprint arXiv:2004.08728*.

Antti Suni, Sofoklis Kakouros, Martti Vainio, and Juraj Šimko. 2020. Prosodic prominence and boundaries in sequence-to-sequence speech synthesis. *arXiv preprint arXiv:2006.15967*.

Jacques Terken and Dik Hermes. 2000. The perception of prosodic prominence. In *Prosody: Theory and experiment: Studies presented to Gösta Bruce*, pages 89–127. Springer.

Andreas Tsiartas, Panayiotis G Georgiou, and Shrikanth S Narayanan. 2013. A study on the effect of prosodic emphasis transfer on overall speech translation quality. In *2013 IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 8396–8400. IEEE.

Mithilesh Vaidya, Kamini Sabu, and Preeti Rao. 2022. Deep learning for prominence detection in children’s read speech. In *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 8157–8161. IEEE.

Ann Wennerstrom. 2001. *The music of everyday speech: Prosody and discourse analysis*. Oxford University Press.

Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, et al. 2019. Huggingface’s transformers: State-of-the-art natural language processing. *arXiv preprint arXiv:1910.03771*.

Long Zhang, Jia Jia, Fanbo Meng, Suping Zhou, Wei Chen, Cunjun Zhang, and Runnan Li. 2018. Emphasis detection for voice dialogue applications using multi-channel convolutional bidirectional long short-term memory network. In *2018 11th International Symposium on Chinese Spoken Language Processing (ISCSLP)*, pages 210–214. IEEE.

Suping Zhou, Jia Jia, Long Zhang, Yanfeng Wang, Wei Chen, Fanbo Meng, Fei Yu, and Jialie Shen. 2020. Inferring emphasis for real voice data: an attentive multimodal neural network approach. In *MultiMedia Modeling: 26th International Conference, MMM 2020, Daejeon, South Korea, January 5–8, 2020, Proceedings, Part II 26*, pages 52–62. Springer.

A Cross-language generalisation in the classifier

Using a Spanish company-internal variant of the Espresso dataset, we trained and tested the classifier on Spanish data in an identical manner to our approach with English. We should however note that the version of the data we had was of lesser recording quality than the English one.

The classifier’s outcomes when evaluated on both the English and Spanish train sets are presented in Table 2. The most important observation from the results is the classifier’s superior performance when trained and tested on the same language. Cross-language assessments, especially from English-trained models tested on Spanish data, manifested a decline in performance. Nevertheless, despite the noted challenges, the results demonstrate that the classifier is able to detect emphasis, even across languages. It is also worth that the Spanish dataset was of considerably lower quality than the English one and is just used here for demonstration purposes. It is plausible that this quality might have affected the model’s performance. Therefore, a more definitive assessment of its cross-language generalisation potential would necessitate testing on datasets of other languages, ideally of comparable quality to the English version.

We also extended the evaluation of the English and Spanish emphasis classifiers to additional languages, using internal datasets to compile test sets mirroring the structure of the English ones, each featuring 2 to 3 speakers. These are summarised in Table 2. Intriguingly, the Spanish classifier outperformed across all tested languages, a finding readily attributable to linguistic similarities in the case of Italian, French, and Portuguese, but less so for Vietnamese. Furthermore, in some instances, performance on non-native test sets was on par with, or even surpassed, native datasets; for example, a word-level F1 score of 84.4% was achieved on the Portuguese test set. These observations imply the feasibility of applying classifiers to languages they were not specifically trained on, particularly when sufficient training data is lacking, and suggest the merit in experimenting with classifiers based on different languages. Additional results could potentially advocate for the benefits of multi-language training approaches. An additional point of interest arises from the performance of the Vietnamese test sets. Vietnam’s tonal nature,

which distinctly shapes its emphasis patterns, ostensibly diverges from the prosodic systems used in Romance and Germanic languages. Despite these fundamental differences, the fact that the Spanish-trained classifier achieved commendable results with Vietnamese indicates that it may be recognising universal features of emphasis that transcend language-specific prosodic systems.

Test data	Train data	Frame-level metrics (%)			Word-level metrics (%)		
		F1 score	Precision	Recall	F1 score	Precision	Recall
English	English	75.52	77.48	76.9	78.4	78.96	79.46
English	Spanish	67.36	68.74	71.95	68.66	66.73	75.21
Spanish	English	55.75	60.82	55.16	56.14	56.93	57.92
Spanish	Spanish	72.52	73.26	75.12	73.92	74.21	76.32
Vietnamese	English	61.65	68.98	61.51	64.59	70.63	63.7
Vietnamese	Spanish	71.21	71.82	76.32	75.48	77.69	78.2
Italian	English	56.79	70.61	52.86	56.12	57.18	57.61
Italian	Spanish	64.72	72.64	63.46	67.81	68.42	70.41
French	English	60.18	62.81	63.31	65.08	65.85	67.07
French	Spanish	62.50	63.09	68.05	68.17	67.64	72.41
Portuguese	English	71.84	83.56	68.41	72.86	73.17	74.69
Portuguese	Spanish	79.84	82.93	80.08	84.4	84.15	87.1

Table 2: Performance metrics of the emphasis classifier across multiple languages, benchmarked using F1 score, precision, and recall. The classifier is trained either on English or Spanish data sets. Rows highlighted in grey represent instances where the training and test data languages are identical.