

Contrastive Learning for Task-Independent SpeechLLM-Pretraining

Maike Züfle Jan Niehues

Karlsruhe Institute of Technology, Germany
{maike.zuefle, jan.niehues}@kit.edu

Abstract

Large language models (LLMs) excel in natural language processing but adapting these LLMs to speech processing tasks efficiently is not straightforward. Direct task-specific fine-tuning is limited by overfitting risks, data requirements, and computational costs. To address these challenges, we propose a scalable, two-stage training approach: (1) A task-independent speech pretraining stage using contrastive learning to align text and speech representations over all layers, followed by (2) a task-specific fine-tuning stage requiring minimal data. This approach outperforms traditional ASR pretraining and enables the model to surpass models specialized on speech translation and question answering while being trained on only 10% of the task-specific data.

1 Introduction

Large language models (LLMs) have demonstrated impressive capabilities in language understanding and translation tasks (OpenAI et al., 2024; Dubey et al., 2024). However, their reliance on written text limits their application in real-world scenarios, where communication often occurs in the form of speech. Extending LLMs to spoken language could enable them to reason about and process spoken content effectively.

A straightforward approach is to use a cascaded pipeline (Shen et al., 2023) where speech is first converted to text using an automatic speech recognition (ASR) model, and the text is subsequently processed by an LLM. While simple, this approach has limitations: it cannot be trained end-to-end, ASR errors can propagate through the pipeline, and critical paralinguistic information, such as speaker features or noise, is lost during transcription.

To address these issues, recent works have proposed SpeechLLMs (Wang et al., 2024a; Tang et al.,

<https://github.com/MaikeZuefle/contr-pretraining>

2024; Nguyen et al., 2024; Chu et al., 2024), which connect speech encoders to LLMs via a trainable projector in an end-to-end manner. However, training them effectively remains a key challenge due to the scarcity of labelled data for downstream tasks (Xu et al., 2023; Ünlü Menevşe et al., 2024).

Pretraining SpeechLLMs before finetuning on task-specific data can help overcome this challenge, reducing the need for large downstream datasets. A straightforward strategy is to use ASR data, training on the ASR task itself (Tang et al., 2024; Wang et al., 2024b). However, this might lead to overfitting to this pretraining task (Tang et al., 2024).

Our work introduces contrastive learning as a task-independent pretraining strategy for SpeechLLMs using ASR data. Contrastive learning aligns the representations of paired speech and text inputs while encouraging representations of unrelated inputs to remain distinct. We demonstrate that applying a contrastive loss to the representations after every layer and pretraining on just 400 hours of parallel speech-text data serves as an effective, task-agnostic foundation for SpeechLLMs.

To evaluate our approach, we finetune the pre-trained model on three downstream tasks: ASR, speech translation (ST), and speech question answering (SQA). We compare its performance against that of traditional ASR pretraining and next-word prediction (NWP) pretraining using mixed speech-text input. The contrastive strategy not only surpasses the other methods but also matches the performance of task-specific state-of-the-art (SOTA) models on nearly all metrics. We also analyse combinations of the pretraining methods to explore potential complementary interactions.

Scaling up contrastive pretraining to 1,400 hours yields further performance gains, with our models surpassing specialized models and current SpeechLLMs. Even in simulated low-resource scenarios, using only 10% of the finetuning data, our method consistently outperforms existing SpeechLLMs.

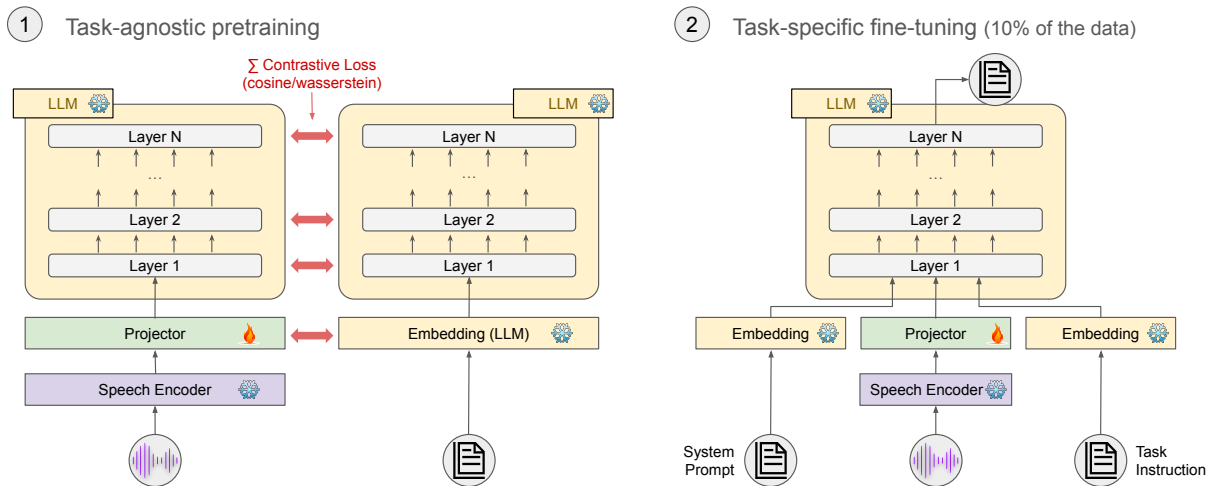


Figure 1: An overview of our two-stage training process. First we perform task-agnostic pretraining, by leveraging ASR data to calculate a contrastive loss over a selected set of layers. Then we finetune on task-specific data. The speech encoder and the LLM remain frozen (indicated by the snowflake), we only train the projector.

Our SpeechLLM architecture is simple yet flexible, consisting of a speech encoder, a projector, and an LLM. Only the projector is trained, making the model parameter-efficient and enabling maximum flexibility in adapting to the ever-evolving landscape of LLMs and speech encoders. Moreover, this approach preserves the LLM’s text-to-text capabilities, ensuring its functionality across both speech and text-based tasks. An overview of our approach is provided in Fig. 1.

The main contributions of our paper are:

1. We propose contrastive pretraining for SpeechLLMs, achieving state-of-the-art results in ST and SQA.
2. We demonstrate the effectiveness of our pretraining method in low-resource settings.
3. We provide a comprehensive analysis of contrastive learning for SpeechLLMs and its interaction with other pretraining methods.
4. We show that contrastive pretraining preserves a model’s ability to recognise paralinguistic speech features.

2 Related Works

SpeechLLMs. SpeechLLM architectures typically consist of three main components: a speech encoder, a pretrained language model, and a projector linking the two. Training strategies differ based on which components are finetuned.

Chu et al. (2024) adopt a comprehensive strategy by finetuning all components. Similarly, other works explore finetuning both the projector and the encoder, alongside applying Low-Rank Adaptation

(LoRA; Hu et al. 2022) to the LLM (Wu et al., 2023; Gong et al., 2024; Fathullah et al., 2023). In contrast, some methods freeze the LLM entirely, focusing on training only the encoder and projector (Wang et al., 2024b; Held et al., 2024; Fathullah et al., 2024). Several other approaches focus on finetuning the projector while introducing LoRA layers to the LLM, leaving the speech encoder and LLM frozen (Peng et al., 2024; Tang et al., 2024). An even more lightweight approach freezes both the encoder and the LLM, finetuning only the projection layers (Wang et al., 2024a).

In this work, we use the latter approach, using a frozen speech encoder and LLM while finetuning only the projection layers. This preserves the text-to-text abilities of the LLM, while also being parameter efficient.

Text and Speech Alignment. While ‘alignment’ has various meanings (Hämmerl et al., 2024), in this work, we use it to refer to the process of creating meaningfully similar representations for text and speech. One approach to achieve this alignment is to directly finetune on the desired tasks without explicitly modelling alignment (Fang et al., 2024; Fathullah et al., 2023). An alternative is to use a downstream task to pretrain the SpeechLLM before finetuning on desired tasks, for example pretraining on ASR (Tang et al., 2024; Wang et al., 2024b) or ASR variants (Wang et al., 2024a).

Other works achieve alignment by encouraging paired speech and text embeddings closer together using an appropriate loss term. For example, Held et al. (2024) use an L2 loss on speech and text em-

beddings in addition to a KL-divergence loss on the output distributions, and [Chuang et al. \(2020\)](#) use a cosine distance loss between paired embeddings. Optimal transport has also been used to align paired embeddings in encoder-decoder speech translation systems ([Le et al., 2023](#); [Tsiamas et al., 2024](#)).

Contrastive training is another widely used approach, bringing parallel text and speech representations closer while pushing non-parallel representations apart. It has been used in knowledge distillation for spoken language understanding ([Zhu et al., 2022](#); [Cappellazzo et al., 2024](#)), as an auxiliary loss for encoder-decoder ST systems ([Ye et al., 2022](#); [Ouyang et al., 2023](#)), and for music ([Manco et al., 2022](#)). It has also been applied to classification tasks, such as emotion recognition ([Sachidananda et al., 2022](#); [Zhou et al., 2024](#)).

Other approaches mix text and speech in the input to help models learn both modalities. For example, [Fang et al. \(2022\)](#) align the output distributions of speech and mixed text-speech inputs, while [Nguyen et al. \(2024\)](#) and [Peng et al. \(2024\)](#) mix speech and text sequences for next-word prediction and question-answering tasks.

In this work, we leverage the previous success of contrastive learning on speech models and apply it to SpeechLLMs. We show how contrastive learning can be applied successfully in this new scenario and combined with other losses. In contrast to other works, we experiment with using optimal transport for the contrastive loss function and apply the loss to several hidden layers of the SpeechLLM.

3 Methods

We propose contrastive learning as a task-independent pretraining method for aligning speech and text representations, leveraging abundantly available ASR data. We compare with two ASR-based pretraining baselines: traditional ASR pretraining and next-word prediction with mixed text-speech spans, which we also introduce as a pretraining approach in this paper.

Notation. Let $\mathbf{t} \in \mathbb{R}^{H \times N}$ represent a text sequence with N tokens, where each token is embedded into a H -dimensional space. The speech sequence is embedded into M vectors, with $\mathbf{s} \in \mathbb{R}^{H \times M}$. Note that usually $M > N$.

3.1 Contrastive Training

Contrastive training brings representations of matched speech-text pairs closer whilst pushing

representations of mismatched pairs apart. Harnessing both matched and mismatched pairs is often a very efficient way of using a dataset and allows one to leverage abundantly available ASR data.

In this paper, we employ the InfoNCE loss ([van den Oord et al., 2019](#)) similar to other works ([Manco et al., 2022](#)). InfoNCE loss is defined as:

$$\mathcal{L}_{\text{InfoNCE}} = -\frac{1}{|B|} \sum_{i \in B} \log \frac{\exp\left(\frac{\text{sim}(\mathbf{s}_i, \mathbf{t}_i)}{\tau}\right)}{\sum_{j \in B} \exp\left(\frac{\text{sim}(\mathbf{s}_i, \mathbf{t}_j)}{\tau}\right)}, \quad (1)$$

where $\text{sim}(\cdot, \cdot)$ denotes a similarity measure, \mathbf{s}_i and \mathbf{t}_i denote speech and text embeddings respectively for the i -th example in the batch, and τ is the temperature parameter controlling the sharpness of the similarity scores. Summing over B in the denominator includes the aligned speech-text pair $(\mathbf{s}_i, \mathbf{t}_i)$, as well as negative (contrastive) pairs. We choose to use all contrastive pairs within a mini-batch to allow for an efficient implementation.

The contrastive loss can also be calculated for multiple layers throughout the network and then summed. [Fig. 1](#) illustrates this setup. We find that this gives the best results (see [Section 5](#)).

We adopt two different similarity measures:

Cosine Similarity. One way to compute $\text{sim}(\mathbf{s}_i, \mathbf{t}_i)$ is to use cosine similarity. Since the length of the speech sequences, M , is typically much larger than the length of the text sequences, N , we cannot directly compute cosine similarity on the entire sequence. Instead, we average the embeddings over the sequence length to obtain a single embedding for each ([Ye et al., 2022](#)). We denote this pretraining strategy by *contr-cos*.

Wasserstein Distance. To avoid aggregating the sequences, we can leverage the Wasserstein (or Optimal Transport) distance between two sequences ([Peyré and Cuturi, 2019](#); [Le et al., 2023](#)). In our case, the Wasserstein distance involves finding the optimal way to "transport" mass from one normalized set of embeddings to another, enabling us to compute the alignment between text and speech embeddings of differing lengths.

More precisely, the optimal transport distance is calculated as follows: a mass of $1/N$ is placed at each text embedding location (points in a H -dimensional space) and we want to move this total mass of 1 to the M speech embedding locations, $1/M$ to each. There is a cost C_{ij} associated with moving a unit of mass from location t_i to location s_j for all $i \in \{1, \dots, N\}$ and $j \in \{1, \dots, M\}$.

The cost is given by the L^p distance in the embedding space. This can be written as:

$$d_{\text{Wasser}}(\mathbf{s}, \mathbf{t}) = \min_{\mathbf{Z}} \{\mathbf{C}(\mathbf{s}, \mathbf{t}) \odot \mathbf{Z}\}, \text{ s.t.} \\ \mathbf{Z}\mathbf{1}^N = \mathbf{1}^N/N, \mathbf{Z}^T\mathbf{1}^M = \mathbf{1}^M/M, \mathbf{Z} \geq \mathbf{0}, \quad (2)$$

where \mathbf{Z} is the transportation matrix with row and column sums $\frac{1}{N}$ and $\frac{1}{M}$ respectively, $\mathbf{C}(\mathbf{s}, \mathbf{t})$ is the cost matrix for sequences \mathbf{s} and \mathbf{t} , $\mathbf{1}^N$ denotes a vector of 1's, and \odot denotes the entrywise dot product. For a more detailed description, we refer the reader to [Le et al. \(2023\)](#).

Using the Wasserstein distance in the loss is expensive, so we resort to an upper-bound approximation that can be computed efficiently using the Sinkhorn algorithm ([Sinkhorn and Knopp, 1967](#)). Note, however, that even with this approximation, using Wasserstein distance is more computationally expensive ($O(n^2/\epsilon^2)$ ([Pham et al., 2020](#))) than using cosine similarity ($O(n)$).

We use the negative distance as the similarity measure for $\mathcal{L}_{\text{InfoNCE}}$ and denote this pretraining strategy by *contr-wasser*.

3.2 Other Pretraining Approaches

To evaluate our contrastive pretraining approach we compare it with traditional ASR pretraining and a mixed text-speech NWP pretraining approach.

ASR Pretraining. We use task-specific finetuning on the ASR task, i.e. following the ASR-instruction prompt, for pretraining ([Tang et al., 2024](#); [Wang et al., 2024b](#)) and use it as a baseline.

Mixed Next Word Prediction. Inspired by [Nguyen et al. \(2024\)](#), who add mixed NWP to their finetuning data mix for downstream tasks, we propose mixed text and speech NWP as a pretraining strategy. We mix text and speech embeddings across randomly selected spans in the sequence, maintaining coherence with word boundaries. A sample sequence might look as follows: $t_0^0, t_1^0, s_2^1, s_3^1, s_4^2, s_5^2, s_6^2, t_7^3$, where superscripts represent word boundaries, and subscripts represent subwords within the text or speech embedding. Since speech embeddings typically differ in length from text embeddings, we compute the NWP loss exclusively on text tokens. Details on how we select the speech and text spans are provided in [Appendix E](#).

We also experiment with using these mixed sequences for contrastive learning: instead of aligning speech and text, we align mixed text-speech representations with text representations.

4 Experiments

Our experiments follow a two-stage training strategy. First, we perform text-speech alignment pretraining using ASR data to adapt the model to handle speech. In the second stage, we finetune the model jointly on a mix of task-specific data. This approach leverages abundant ASR data to reduce reliance on large task-specific datasets, enabling effective adaptation even in low-resource settings.

4.1 Model Architecture

Our SpeechLLM architecture consists of three core components: a speech encoder, a text-based LLM, and a projector bridging the two. We select HuBERT ([Hsu et al., 2021](#)) as our speech encoder due to its ability to process audio inputs longer than 30 seconds, which is crucial for SQA tasks. Recent research ([Hassid et al., 2024](#); [Nguyen et al., 2024](#)) has demonstrated strong performance with HuBERT variants. We specifically utilize the facebook/hubert-large-1s960-ft variant. We use Llama-3.1-8B-Instruct ([Dubey et al., 2024](#)) as our LLM. As a projector, we choose Q-Former ([Li et al., 2023](#)), a window-level query Transformer, which already proved successful for other SpeechLLMs ([Tang et al., 2024](#); [Held et al., 2024](#)).

For contrastive learning, we compare speech and text embeddings without the use of system prompts. This creates a positional mismatch between contrastive pretraining and finetuning. To address this potential overfitting problem, we experiment with varying the absolute starting position in the positional embeddings. However, this does not consistently improve model performance ([Appendix D](#)).

4.2 Training Parameters

In both pretraining and task-specific finetuning, we keep the LLM and speech encoder frozen, adjusting only the speech projector, which has 42.5 million parameters. Training is conducted on four NVIDIA A100-SXM4-40GB GPUs. ASR pretraining takes approximately 9h, contrastive pretraining with cosine similarity over all layers takes around 5h. The task-specific finetuning takes around 30h with 100% of the data. Detailed training parameters and model specifics are listed in [Appendix A](#).

4.3 Data

Pretraining. For pretraining, we use the English portion of the widely-used MustC-v1 ([Di Gangi et al., 2019](#)) data (approximately 400 hours) to

demonstrate effective performance with limited data. MustC includes TED talks with diverse speakers, accents, and topics, along with manually curated transcriptions. Additionally, we test pretraining on a larger dataset, adding the M-version of Gigaspeech (Chen et al., 2021) to our pretraining data, which includes 1,000 hours of speech from audiobooks, podcasts, and YouTube videos.

Finetuning. We jointly finetune and evaluate our contrastive pretraining method on three diverse tasks covering a range of speech-text model abilities: Automatic Speech Recognition (ASR), Speech Translation (ST), and Speech Question Answering (SQA). For ASR and ST, we use the MustC-v1 dataset for training and testing. As before, English transcriptions serve as the ASR data, and we use language pairs en-de, en-fr, en-it, and en-es for ST. For question answering, we use Spoken-SQuAD (Lee et al., 2018), an English dataset based on Wikipedia, derived from the Stanford Question Answering Dataset (Rajpurkar et al., 2016, SQuAD).

While there is overlap between pretraining and finetuning data for the ASR and ST tasks, we consider this a realistic scenario: In practice, when task-specific data is limited, leveraging all available data—even if it appears in both stages—is often necessary. This is particularly relevant for speech translation, where transcripts are abundant but parallel translations are much scarcer.

Paralinguistic Feature Analysis. We also analyse whether contrastive pretraining harms a model’s ability to capture paralinguistic features, such as speaking pace, gender or noise. We use two datasets: the *mls-eng-speaker-descriptions* dataset (Pratap et al., 2020; Lacombe et al., 2024; Lyth and King, 2024), which involves reasoning about speaking rate, gender, background noise, and other features, and the MuST-SHE (Bentivogli et al., 2020) dataset, where gender information is needed to correctly translate the segment.

More information about the train and test sets can be found in Appendix B and details on prompts in Appendix C.

4.4 Evaluation

We report the following evaluation metrics: WER (Word Error Rate, using *jiwer*) for ASR, sacreBLEU (Post, 2018) and COMET₂₂^{DA} (Rei et al., 2020) for ST, and exact match accuracy (EM) and F1 score for SQA. Scores for ST are averaged over the language pairs. For the best performing models,

we report results for individual language pairs in Appendix H. Moreover, we calculate the average across the three tasks using WER, COMET, and F1 scores. Since these scores have different ranges, care has to be taken when averaging. We therefore normalize the scores using lower (lb) and upper (ub) bounds as detailed in Section 4.5. We then calculate the normalized average as follows:

$$\text{Norm. Avg.} = \frac{1}{|T|} \sum_{t \in T} \frac{\text{score}(t) - \text{lb}(t)}{\text{ub}(t) - \text{lb}(t)}, \quad (3)$$

where T is the set of tasks, and $\text{score}(t)$ is the model’s score on task t .

We also report the contrastive loss on the test set for the embedding layer in order to evaluate the effectiveness of aligning speech and text.

4.5 Baselines

We compare our approach against four baselines: (1) Firstly, we take three SOTA models, each specialized on one of the three tasks to set a goal for our models. Specifically, we use Whisper (Radford et al., 2022) for ASR, Seamless (Communication et al., 2023) for ST, and You et al. (2022) for SQA¹. These SOTA results also serve as an upper bound for normalizing the respective scores when calculating the overall (normalized average) performance scores. (2) As the second baseline, we use the cascaded HuBERT-Llama-3.1-8B-Instruct model, which also serves as the lower bound for calculating overall performance scores.

Our primary objective is to assess the impact of contrastive pretraining. To do so, we need a controlled environment, where we keep the model architecture, data, and training setup as consistent as possible. To this end, we include two additional baselines using our training setup: (3) No pretraining and (4) Standard ASR pretraining.

5 Results

The goal of our experiments is to find the best setup for using contrastive learning for SpeechLLMs. We investigate various contrastive pretraining losses and where to apply them, combinations with other pretraining losses, and the impact of additional data. We then compare our best-performing model to recent SpeechLLMs. Finally, we investigate how text-speech alignment affects the detection of paralinguistic features.

¹More recently, Manakul et al. (2024) report SQA results for Gemini-1.5-Pro, which underperforms You et al. (2022).

FT Data	Model	ASR	ST		SQA		Contr loss on test		Norm.
		WER \downarrow *	BLEU \uparrow	COMET \uparrow *	EM \uparrow	F1 \uparrow *	cos.	wasser.	avg. \uparrow
N/A	Specialized	6.54	30.99	80.02	64.19	77.10	N/A	N/A	100
N/A	HuBERT + Llama	18.38	19.85	73.92	36.15	54.76	N/A	N/A	0
100%	no pretrain	12.09	28.84	79.94	64.13	76.53	1.37	1.26	83.08
10%	no pretrain	23.78	19.87	69.72	34.72	46.07	1.38	1.36	-51.13
	ASR pretrain	12.21	24.82	75.70	49.48	63.36	1.37	1.30	39.93
	contr-cos-emb	13.98	25.47	76.32	54.86	68.49	0.97	0.72	46.00
	contr-wasser-emb	15.99	26.10	77.09	54.64	67.83	1.08	0.62	43.52
	contr-cos-all	13.06	27.19	78.29	60.48	73.90	1.08	0.91	67.39
	contr-wasser-all	12.92	29.07	80.52	64.06	77.22	1.07	0.64	84.96

Table 1: Results of models with no pretraining, ASR pretraining, and contrastive pretraining using cosine similarity (*contr-cos*) or Wasserstein distance (*contr-wasser*) on embedding (*-emb*) or all (*-all*) layers, followed by fine-tuning on a 10% subset of task-specific data. The specialized baselines are Whisper (Radford et al., 2022) for ASR, Seamless (Communication et al., 2023) for ST, and You et al. (2022) for SQA. Metrics with * contribute to the normalized average as in Eq. (3).

5.1 Contrastive Pretraining

Contrastive Loss on the Embedding Layer.

We start by calculating the contrastive loss on the embedding layer and comparing the two loss functions: Wasserstein (*contr-wasser-emb*), and cosine similarity (*contr-cos-emb*). After pretraining, the models are finetuned on three downstream tasks (jointly) using a random 10% subset of the data, simulating a low-resource setting (Table 1).

Our contrastive pretraining models consistently outperform the cascaded Hubert+Llama and no-pretraining baselines across all metrics. In ST and SQA, they also surpass standard ASR pretraining, by over 5% in SQA. However, ASR-pretrained models remain better for ASR, which is unsurprising, given they are pretrained with more data on this task. Analysing the contrastive loss on the test set, we find that it is significantly lower after contrastive pretraining than after ASR pretraining or direct finetuning on the task, showing that indeed there is more alignment. Interestingly, while contrastive models exhibit lower test-set contrastive loss on the embedding layer, the losses for no-pretraining and ASR-pretraining are similar despite a large performance gap.

Contrastive Loss on Different Layers. Keeping with the low-resource setting, we investigate calculating the contrastive loss at different LLM layers. We find that applying contrastive loss to deeper layers significantly increases downstream performance, however, the best layer to use is not consistent. Therefore, we also explore using a

summed loss over all (multiples of 5²) layers in Table 1. The methods *contr-cos-all* and *contr-wasser-all* achieve normalized task averages of 67.39 and 84.96, respectively, significantly outperforming the ASR-pretraining baseline and the *contr-emb* models. *Contr-wasser-all*, despite using only 10% of the finetuning data, even matches the performance of the specialized models on two out of three tasks and surpasses the no pretraining model finetuned on 100% of the task-specific data. This highlights the suitability of contrastive pretraining for scenarios with limited task-specific data. We analyse the impact of training on different layers in more detail in Appendix D, including the impact on individual task performance.

5.2 Mixed-Text-Speech Input

The previous section shows that contrastive learning is an effective pretraining strategy, outperforming ASR pretraining. We now analyse, whether aligning mixed speech-text and text sequences can further enhance performance.

Mixed Contrastive Pretraining. We adapt our approach to use mixed speech-text sequences (*contr-*-+mixed*). Using the *contr-cos* loss on the embedding layer, our model matches the performance of standard contrastive pretraining, outperforming both no pretraining and ASR pretraining (Table 2). However, with the *contr-wasser* loss results degrade significantly. Consequently, we continue with *contr-cos-all+mixed* to test the loss

²Initial experiments found that using every 5th layer efficiently approximates using all layers (Table 11).

Model	ASR	ST	SQA	Norm.
	WER _↓ *	COMET _↑ *	F1 _↑ *	avg. _↑
no pretrain	23.78	69.72	46.07	-51.13
ASR pretrain	12.21	75.70	63.36	39.93
contr-cos-emb	13.98	76.32	68.49	46.00
+ mixed	14.26	76.61	68.40	46.61
contr-wasser-emb	15.99	77.09	67.83	43.52
+ mixed	14.70	75.05	61.48	26.53
contr-cos-all	13.06	78.29	73.90	67.39
+ mixed	17.63	73.83	55.12	2.14
mixed-nwp	15.27	76.05	67.68	39.67

Table 2: Comparison of models pretrained on mixed speech-text input and finetuned on 10% of the task-specific data.

Model	ASR	ST	SQA	Norm.
	WER _↓ *	COMET _↑ *	F1 _↑ *	avg. _↑
ASR pretrain	12.21	75.70	63.36	39.93
contr-cos-emb	13.98	76.32	68.49	46.00
+ asr loss	13.18	76.89	68.67	51.64
contr-wasser-emb	15.99	77.09	67.83	43.52
+ asr loss	13.19	77.40	70.44	57.03
contr-cos-all	13.06	78.29	73.90	67.39
+ asr loss	11.68	78.94	76.12	78.17
contr-wasser-all	12.92	80.52	77.22	84.96
+ asr loss	11.23	80.26	77.19	88.24

Table 3: Comparison of models pretrained on a combination of contrastive and ASR losses and then finetuned on a 10% subset of task-specific data.

across all layers. This model also significantly underperforms its non-mixed counterpart.

Mixed NWP. We also propose using next-word prediction (NWP) with mixed speech-text inputs (*mixed-nwp*) for pretraining, inspired by Nguyen et al. (2024). While this approach surpasses the no pretraining and ASR pretraining baselines (Table 2, last row), it underperforms contrastive pretraining.

Further ablations of contrastive pretraining with mixed speech-text inputs and of *mixed-nwp*, including two different forced-alignment algorithms, do not yield additional improvements (detailed results, including all metrics, are provided in Appendix F).

5.3 Combining losses

Building on the previous sections, which analyse contrastive pretraining and mixed speech-text pretraining in comparison to traditional ASR pretraining, we now explore the potential of combining these approaches. We sum the individual losses, assigning equal weight to each.

Contr. + ASR Loss. We find that combining the ASR pretraining loss with the contrastive losses on the embedding layer leads to improved performance for ASR and SQA, as seen in Table 3. Similarly, this enhances the performance of the *contr-*-all* models, closing the gap between our contrastive pretrained models and ASR pretraining on ASR.

Contr. + mixed-NWP Loss. In contrast, combining the mixed-NWP loss with the contrastive losses does not yield further improvements. Detailed results are provided in Appendix G.

5.4 Additional pretraining data

In the previous sections, we use Must-C ASR data for pretraining. Given the abundance of ASR data, we expand our pretraining dataset by incorporating GigaSpeech (Chen et al., 2021), effectively tripling the amount of data. The results demonstrate the significant impact of larger-scale pretraining.

Low-resource setting. As before, we use a 10% finetuning subset to simulate a low-resource scenario. In preliminary experiments with lightweight contrastive pretraining on the embedding layer (Appendix H), we find that increasing the pretraining data leads to much larger improvements for *contr-cos* loss (+32.45 Norm. avg.) than for *contr-wasser* loss (+6.60). Therefore, we report *contr-cos* results for subsequent experiments; results for *contr-wasser* can be found in Appendix H. We perform contrastive learning on all layers with and without ASR loss, as these settings yielded the best results in the previous sections. As shown in Table 4, both models *contr-cos-all+giga* and *contr-cos-all+asr+giga* achieve an overall improvement of more than 19 points (Norm. avg.) compared to the same models trained with less pretraining data, reaching 97.12 and 102.15 points respectively. Remarkably, despite being finetuned on only 10% of the task-specific data, these models surpass both specialized models, the Seamless model (Communication et al., 2023) for ST and the You et al. (2022) model for SQA. Moreover, they outperform the ASR-pretraining baseline even on the ASR task.

Benchmarking our best variants. We evaluate our best-performing models in a higher resource setting. We finetune *contr-cos-all+giga* and *contr-cos-all+asr+giga* on 100% of the task-specific data. As expected, this leads to even stronger results. But while combining the contrastive and ASR losses provides significant improvements in the low-resource settings, when finetuning with 100% of the available data, the contrastive loss

FT Data	Model	ASR		ST		SQA		Contr loss on test		Norm.
		WER↓*	BLEU↑	COMET↑*	EM↑	F1↑*	cos.	wasser.	avg.↑	
N/A	Specialized	6.54	30.99	80.02	64.19	77.10	N/A	N/A	100	
N/A	BLSP-lslm-7b	44.51	28.70	78.68	5.60	21.82	N/A	N/A	-96.72	
N/A	Qwen2-Audio-7b	12.03	21.57	74.81	27.79	46.75	N/A	N/A	10.79	
10%	ASR pretrain	12.21	24.82	75.70	49.48	63.36	1.37	1.30	39.93	
	+ giga	12.16	27.44	78.38	62.23	73.17	1.37	1.31	69.37	
	contr-cos-all	13.06	27.19	78.29	60.48	73.90	1.08	0.91	67.39	
	+ giga	10.94	29.95	81.21	65.45	79.12	1.06	0.71	97.12	
	contr-cos-all + asr	11.68	27.61	78.94	63.81	76.12	1.08	0.91	78.17	
	+ giga	11.12	29.90	81.29	72.16	82.52	0.98	0.65	102.15	
100%	ASR pretrain	10.28	30.31	80.98	69.87	80.92	1.37	1.22	100.24	
	+ giga	11.48	30.46	81.03	72.25	82.56	1.37	1.23	99.77	
	contr-cos-all	12.56	30.69	81.48	71.71	82.50	1.25	1.02	99.06	
	+ giga	9.31	31.36	81.98	75.12	84.67	1.10	0.83	114.18	
	contr-cos-all + asr	10.04	30.82	81.56	72.40	82.47	1.19	1.03	106.57	
	+ giga	10.03	31.54	81.99	76.11	84.94	1.10	0.79	112.66	

Table 4: Comparison of pretraining only on Must-C data or the combination with the bigger Giga dataset (+giga). We compare our models against BLSP (Wang et al., 2024a) and the Qwen2-Audio-7B-Instruct model (Chu et al., 2024). Metrics with * contribute to the normalized average as in Eq. (3). Results for *contr-wasser* can be found in Appendix H and do not outperform *contr-cos*.

alone performs on par with the combined loss. We hypothesize that the ASR loss is more sensitive to the out-of-domain Giga pretraining data than the contrastive loss, leading to no improvement on the ASR task.

We compare these models against `blsp_lsm_7b` (Wang et al., 2024a), which has a similar setup to ours in that it only trains the projector while keeping other components frozen, and Qwen2-Audio-7b-Instruct (Chu et al., 2024), a strong baseline SpeechLLM. Notably, BLSP is finetuned on Must-C translation data, whereas for all other tasks, BLSP and Qwen2-Audio results reflect zero-shot performance. Across all tasks, our model outperforms both of these SpeechLLMs.

Results for *contr-wasser* experiments can be found in Appendix H, these yield similar results but do not outperform *contr-cos*.

5.5 Impact of text-speech alignment on capturing paralinguistic features

Speech data is much richer than a sequence of words, and the pretrained speech encoder (HuBERT) might capture additional paralinguistic features. By aligning the richer speech embeddings more closely with text embeddings, the pretraining methods we explore could potentially train the projector to discard these additional features.

To assess whether this occurs, we examine our

contrastively pretrained models on a) a paralinguistic classification task and b) a generation task where paralinguistic information is needed.

For a) we finetune our models on paralinguistic classification tasks before and after pretraining. In addition, we use the speech encoder (HuBERT) with a linear classification head as a baseline. The results demonstrate that models with contrastive and ASR-based pretraining perform comparably to the baseline HuBERT model in predicting paralinguistic annotations, indicating that alignment-focused pretraining does not affect the retention of paralinguistic information significantly. Detailed accuracy results are presented in Table 5. In addition, Table 16 shows that incorporating these features as auxiliary finetuning tasks does not notably alter model performance.

For b) we test our models on the translation dataset MuST-SHE (Bentivogli et al., 2020), where the speakers’s gender determines correct translations. We include a cascaded model (HuBERT + Llama) as a baseline where all paralinguistic features are lost due to the cascaded setup. We find that the gender accuracy of the translations is higher for contrastively pretrained models than for ASR pretrained models or the cascaded model as can be seen in Table 6, further supporting that contrastive pretraining does not hinder a model’s ability to capture paralinguistic features.

Model	Speaking Rate	Gender	Noise	Reverberation	Speech Monotony	SDR Noise	PESQ Speech Quality	Avg.
random	14.29	50.00	14.29	20.00	20.00	16.67	16.67	22.25
HuBERT + Class. Head	57.38	64.69	27.89	28.12	39.35	42.85	27.22	41.07
no pretrain	57.07	64.26	<u>27.62</u>	27.38	37.60	33.86	<u>23.72</u>	38.79
ASR pretrain	57.23	<u>57.95</u>	28.87	29.05	36.35	<u>30.19</u>	<u>21.86</u>	<u>37.36</u>
contr-cos-all	56.46	62.17	28.60	29.27	<u>35.39</u>	34.17	24.12	38.60
contr-wasser-all	56.12	62.30	28.26	27.57	37.70	34.25	24.30	38.64
contr-cos-all + asr	<u>55.00</u>	61.29	28.73	<u>26.24</u>	<u>35.29</u>	<u>30.94</u>	24.86	<u>37.48</u>
contr-wasser-all + asr	<u>56.86</u>	<u>59.62</u>	28.31	29.21	36.64	40.25	26.37	39.61

Table 5: Accuracy in % results for different meta-speech categories on the mls-eng-speaker-descriptions dataset. Bold results are the best or not significantly worse than the best, underlined results are the weakest or not significantly worse than the weakest result using pairwise t-test ($p < 0.01$).

Model	Avg. Gender Acc in %	COMET	
		Correct	Diff
HuBERT + LLama	72.21	77.24	0.69
no pretrain	70.61	70.07	0.61
ASR pretrain	74.11	77.46	0.84
contr-cos-all	78.42	82.20	1.09
contr-wasser-all	76.87	81.46	1.01
contr-cos-all + asr	76.10	80.01	0.95
contr-wasser-all + asr	76.21	81.40	0.99

Table 6: Results for the MuST-SHE dataset (Bentivogli et al., 2020) for different models. Following Bentivogli et al. (2020), we evaluate the overall gender accuracy and the COMET score difference between the scores of correctly and incorrectly gendered references. Results are average over three language pairs, individual results can be found in Table 17.

6 Conclusion

We introduce a contrastive learning-based pretraining strategy for SpeechLLMs, aligning text and speech embeddings in a parameter-efficient manner. This approach establishes an effective task-agnostic foundation for SpeechLLMs and demonstrates state-of-the-art performance across multiple downstream tasks, even when little finetuning data is available.

Importantly, this pretraining method preserves the model’s ability to understand paralinguistic information, offering a promising path forward for advancing SpeechLLMs.

7 Limitations

While our work provides valuable insights into pretraining of LLMs, it has some limitations that offer opportunities for future research:

More Expensive Training Scenarios: Our study focuses on a lightweight and flexible ap-

proach, keeping the LLM and speech encoder frozen while fine-tuning only the projector. Although this minimizes computational costs, future research could explore scenarios where the LLM and encoder or adapter layers are also fine-tuned, potentially unlocking additional capabilities.

Limited Incorporation of Meta-Speaker data:

We analyze whether contrastively pretrained models, designed to align speech and text, retain their ability to capture meta-speech information such as speaker characteristics or noise. While we find that these models do not lose this capacity in both classification and generation tasks, extending this analysis to additional generation tasks would provide a deeper understanding of their impact.

Low-resource Scenarios: Our study simulates low-resource settings by using subsets of the fine-tuning data, finding that models perform well even with only 10% of the task-specific data. However, extending this analysis to truly low-resource languages or tasks would offer richer insights into the model’s adaptability and performance under resource constraints.

Analysis of Contrastive Examples: In this work, we use all contrastive pairs within a mini-batch to maintain an efficient implementation. However, an alternative approach could involve selecting contrastive examples that are more challenging by ensuring they are closer to the positive examples, as demonstrated by Ye et al. (2022).

Potential Risks: Enabling Language Models to understand speech makes them more easily and widely accessible. However, this also lowers the barrier for the potential misuse these models. In addition, SpeechLLMs can reproduce biases seen in the training data. We try to minimize this risk by not training the LLM, so any safeguard measures

trained into the LLM also apply to the SpeechLLM. However, the model could still be affected by adversarial attacks.

Acknowledgments

This work has received funding from the European Union’s Horizon research and innovation programme under grant agreement No 101135798, project Meetween (My Personal AI Mediator for Virtual MEETtings BETWEEN People). We gratefully acknowledge Poland’s high-performance Infrastructure PLGrid ACC Cyfronet AGH for providing computer facilities.

References

- Alexei Baevski, Yuhao Zhou, Abdelrahman Mohamed, and Michael Auli. 2020. [wav2vec 2.0: A framework for self-supervised learning of speech representations](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 12449–12460. Curran Associates, Inc.
- Luisa Bentivogli, Beatrice Savoldi, Matteo Negri, Mattia A. Di Gangi, Roldano Cattoni, and Marco Turchi. 2020. [Gender in danger? evaluating speech translation technology on the MuST-SHE corpus](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6923–6933, Online. Association for Computational Linguistics.
- Umberto Cappellazzo, Enrico Fini, Muqiao Yang, Daniele Falavigna, Alessio Brutti, and Bhiksha Raj. 2024. [Continual contrastive spoken language understanding](#). In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 3727–3741, Bangkok, Thailand. Association for Computational Linguistics.
- Guoguo Chen, Shuzhou Chai, Guanbo Wang, Jiayu Du, Wei-Qiang Zhang, Chao Weng, Dan Su, Daniel Povey, Jan Trmal, Junbo Zhang, Mingjie Jin, Sanjeev Khudanpur, Shinji Watanabe, Shuaijiang Zhao, Wei Zou, Xiangang Li, Xuchen Yao, Yongqing Wang, Yujun Wang, and 2 others. 2021. [Gigaspeech: An evolving, multi-domain asr corpus with 10,000 hours of transcribed audio](#). In *Proc. Interspeech 2021*.
- Yunfei Chu, Jin Xu, Qian Yang, Haojie Wei, Xipin Wei, Zhifang Guo, Yichong Leng, Yuanjun Lv, Jinzheng He, Junyang Lin, Chang Zhou, and Jingren Zhou. 2024. [Qwen2-audio technical report](#). *Preprint*, arXiv:2407.10759.
- Shun-Po Chuang, Tzu-Wei Sung, Alexander H. Liu, and Hung-yi Lee. 2020. [Worse WER, but better BLEU? leveraging word embedding as intermediate in multi-task end-to-end speech translation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5998–6003, Online. Association for Computational Linguistics.
- Seamless Communication, Loïc Barrault, Yu-An Chung, Mariano Cora Meglioli, David Dale, Ning Dong, Paul-Ambroise Duquenne, Hady Elsahar, Hongyu Gong, Kevin Heffernan, John Hoffman, Christopher Klaiber, Pengwei Li, Daniel Licht, Jean Maillard, Alice Rakotoarison, Kaushik Ram Sadagopan, Guillaume Wenzek, Ethan Ye, and 49 others. 2023. [Seamlessm4t: Massively multilingual & multimodal machine translation](#). *Preprint*, arXiv:2308.11596.
- Mattia A. Di Gangi, Roldano Cattoni, Luisa Bentivogli, Matteo Negri, and Marco Turchi. 2019. [MuST-C: a Multilingual Speech Translation Corpus](#). 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 2012–2017, Minneapolis, Minnesota. Association for Computational Linguistics.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, Arun Rao, Aston Zhang, and 516 others. 2024. [The llama 3 herd of models](#). *Preprint*, arXiv:2407.21783.
- Qingkai Fang, Shoutao Guo, Yan Zhou, Zhengrui Ma, Shaolei Zhang, and Yang Feng. 2024. [Llama-omni: Seamless speech interaction with large language models](#). *arXiv preprint arXiv:2409.06666*.
- Qingkai Fang, Rong Ye, Lei Li, Yang Feng, and Mingxuan Wang. 2022. [STEMM: Self-learning with speech-text manifold mixup for speech translation](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 7050–7062, Dublin, Ireland. Association for Computational Linguistics.
- Yassir Fathullah, Chunyang Wu, Egor Lakomkin, Junteng Jia, Yuan Shangguan, Ke Li, Jinxi Guo, Wenhan Xiong, Jay Mahadeokar, Ozlem Kalinli, Christian Fuegen, and Mike Seltzer. 2023. [Prompting large language models with speech recognition abilities](#). *Preprint*, arXiv:2307.11795.
- Yassir Fathullah, Chunyang Wu, Egor Lakomkin, Ke Li, Junteng Jia, Yuan Shangguan, Jay Mahadeokar, Ozlem Kalinli, Christian Fuegen, and Mike Seltzer. 2024. [AudioChatLlama: Towards general-purpose speech abilities for LLMs](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 5522–5532, Mexico City, Mexico. Association for Computational Linguistics.
- Yuan Gong, Hongyin Luo, Alexander H. Liu, Leonid Karlinsky, and James R. Glass. 2024. [Listen, think,](#)

- and understand. In *The Twelfth International Conference on Learning Representations*.
- Katharina Hämmerl, Jindřich Libovický, and Alexander Fraser. 2024. [Understanding cross-lingual Alignment—A survey](#). In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 10922–10943, Bangkok, Thailand. Association for Computational Linguistics.
- Michael Hassid, Tal Remez, Tu Anh Nguyen, Itai Gat, Alexis Conneau, Felix Kreuk, Jade Copet, Alexandre Defossez, Gabriel Synnaeve, Emmanuel Dupoux, Roy Schwartz, and Yossi Adi. 2024. [Textually pretrained speech language models](#). *Preprint*, arXiv:2305.13009.
- William Held, Ella Li, Michael Ryan, Weiyang Shi, Yanzhe Zhang, and Diyi Yang. 2024. [Distilling an end-to-end voice assistant without instruction training data](#). *Preprint*, arXiv:2410.02678.
- Wei-Ning Hsu, Benjamin Bolte, Yao-Hung Tsai, Kushal Lakhota, 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*, PP:1–1.
- Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. [LoRA: Low-rank adaptation of large language models](#). In *International Conference on Learning Representations*.
- Ludwig Kurzinger, Dominik Winkelbauer, Lujun Li, Tobias Watzel, and Gerhard Rigoll. 2020. [Ctc-segmentation of large corpora for german end-to-end speech recognition](#). In *International Conference on Speech and Computer*.
- Yoach Lacombe, Vaibhav Srivastav, and Sanchit Gandhi. 2024. Data-speech. <https://github.com/ylacombe/dataspeech>.
- Phuong-Hang Le, Hongyu Gong, Changan Wang, Juan Pino, Benjamin Lecouteux, and Didier Schwab. 2023. [Pre-training for speech translation: Ctc meets optimal transport](#). *Preprint*, arXiv:2301.11716.
- Chia-Hsuan Lee, Szu-Lin Wu, Chi-Liang Liu, and Hung-yi Lee. 2018. Spoken squad: A study of mitigating the impact of speech recognition errors on listening comprehension. *Proc. Interspeech 2018*, pages 3459–3463.
- Junnan Li, Dongxu Li, Silvio Savarese, and Steven C. H. Hoi. 2023. [Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models](#). In *International Conference on Machine Learning*.
- Dan Lyth and Simon King. 2024. [Natural language guidance of high-fidelity text-to-speech with synthetic annotations](#). *Preprint*, arXiv:2402.01912.
- Potsawee Manakul, Guangzhi Sun, Warit Sirichotedumrong, Kasima Tharnpipitchai, and Kunat Pipatanakul. 2024. [Enhancing low-resource language and instruction following capabilities of audio language models](#). *Preprint*, arXiv:2409.10999.
- Iliaria Manco, Emmanouil Benetos, Elio Quinton, and György Fazekas. 2022. Contrastive audio-language learning for music. In *Proceedings of the 23rd International Society for Music Information Retrieval Conference (ISMIR)*.
- Tu Anh Nguyen, Benjamin Muller, Bokai Yu, Marta R. Costa-jussa, Maha Elbayad, Sravya Popuri, Christophe Ropers, Paul-Ambroise Duquenne, Robin Algayres, Ruslan Mavlyutov, Itai Gat, Mary Williamson, Gabriel Synnaeve, Juan Pino, Benoit Sagot, and Emmanuel Dupoux. 2024. [Spirit lm: Interleaved spoken and written language model](#). *Preprint*, arXiv:2402.05755.
- OpenAI, Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoní Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, Red Avila, Igor Babuschkin, Suchir Balaji, Valerie Balcom, Paul Baltescu, Haiming Bao, Mohammad Bavarian, Jeff Belgum, and 262 others. 2024. [Gpt-4 technical report](#). *Preprint*, arXiv:2303.08774.
- Siqi Ouyang, Rong Ye, and Lei Li. 2023. [WACO: Word-aligned contrastive learning for speech translation](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3891–3907, Toronto, Canada. Association for Computational Linguistics.
- Yifan Peng, Krishna C. Puvvada, Zhehuai Chen, Piotr Zelasko, He Huang, Kunal Dhawan, Ke Hu, Shinji Watanabe, Jagadeesh Balam, and Boris Ginsburg. 2024. [Voicetextblender: Augmenting large language models with speech capabilities via single-stage joint speech-text supervised fine-tuning](#). *Preprint*, arXiv:2410.17485.
- Gabriel Peyré and Marco Cuturi. 2019. [Computational optimal transport: With applications to data science](#). *Foundations and Trends® in Machine Learning*, 11:355–206.
- Khiem Pham, Khang Le, Nhat Ho, Tung Pham, and Hung Bui. 2020. [On unbalanced optimal transport: An analysis of sinkhorn algorithm](#). *Preprint*, arXiv:2002.03293.
- Matt Post. 2018. [A call for clarity in reporting BLEU scores](#). In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 186–191, Brussels, Belgium. Association for Computational Linguistics.
- Vineel Pratap, Andros Tjandra, Bowen Shi, Paden Tomasello, Arun Babu, Sayani Kundu, Ali Elkahky, Zhaoheng Ni, Apoorv Vyas, Maryam Fazel-Zarandi, Alexei Baevski, Yossi Adi, Xiaohui Zhang, Wei-Ning

- Hsu, Alexis Conneau, and Michael Auli. 2023. [Scaling speech technology to 1,000+ languages](#). *Preprint*, arXiv:2305.13516.
- Vineel Pratap, Qiantong Xu, Anuroop Sriram, Gabriel Synnaeve, and Ronan Collobert. 2020. Mls: A large-scale multilingual dataset for speech research. *ArXiv*, abs/2012.03411.
- Alec Radford, Jong Wook Kim, Tao Xu, Greg Brockman, Christine McLeavey, and Ilya Sutskever. 2022. [Robust speech recognition via large-scale weak supervision](#). *Preprint*, arXiv:2212.04356.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. [Squad: 100,000+ questions for machine comprehension of text](#). *Preprint*, arXiv:1606.05250.
- Ricardo Rei, Craig Stewart, Ana C Farinha, and Alon Lavie. 2020. [COMET: A neural framework for MT evaluation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2685–2702, Online. Association for Computational Linguistics.
- Vin Sachidananda, Shao-Yen Tseng, Erik Marchi, Sachin Kajarekar, and Panayiotis Georgiou. 2022. [Calm: Contrastive aligned audio-language multirate and multimodal representations](#). *Preprint*, arXiv:2202.03587.
- Yongliang Shen, Kaitao Song, Xu Tan, Dongsheng Li, Weiming Lu, and Yueting Zhuang. 2023. [Hugging-GPT: Solving AI tasks with chatGPT and its friends in hugging face](#). In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Richard Sinkhorn and Paul Knopp. 1967. [Concerning nonnegative matrices and doubly stochastic matrices](#). *Pacific Journal of Mathematics*, 21:343–348.
- Jianlin Su, Yu Lu, Shengfeng Pan, Ahmed Murtadha, Bo Wen, and Yunfeng Liu. 2023. [Roformer: Enhanced transformer with rotary position embedding](#). *Preprint*, arXiv:2104.09864.
- Changli Tang, Wenyi Yu, Guangzhi Sun, Xianzhao Chen, Tian Tan, Wei Li, Lu Lu, Zejun MA, and Chao Zhang. 2024. [SALMONN: Towards generic hearing abilities for large language models](#). In *The Twelfth International Conference on Learning Representations*.
- Ioannis Tsiamas, Gerard Gállego, José Fonollosa, and Marta Costa-jussà. 2024. [Pushing the limits of zero-shot end-to-end speech translation](#). In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 14245–14267, Bangkok, Thailand. Association for Computational Linguistics.
- Merve Ünlü Menevşe, Yusufcan Manav, Ebru Arisoy, and Arzucan Özgür. 2024. [Dealing with data scarcity in spoken question answering](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 4449–4455, Torino, Italia. ELRA and ICCL.
- Aaron van den Oord, Yazhe Li, and Oriol Vinyals. 2019. [Representation learning with contrastive predictive coding](#). *Preprint*, arXiv:1807.03748.
- Francesco Verdini, Pierfrancesco Melucci, Stefano Perna, Francesco Ciriaggi, Marco Gaido, Sara Papi, Szymon Mazurek, Marek Kasztelnik, Luisa Bentivogli, Sébastien Bratières, Paolo Meriardo, and Simone Scardapane. 2024. [How to connect speech foundation models and large language models? what matters and what does not](#). *Preprint*, arXiv:2409.17044.
- Chen Wang, Minpeng Liao, Zhongqiang Huang, Jinliang Lu, Junhong Wu, Yuchen Liu, Chengqing Zong, and Jiajun Zhang. 2024a. [Blsp: Bootstrapping language-speech pre-training via behavior alignment of continuation writing](#). *Preprint*, arXiv:2309.00916.
- Xiong Wang, Yangze Li, Chaoyou Fu, Yunhang Shen, Lei Xie, Ke Li, Xing Sun, and Long Ma. 2024b. [Freeze-omni: A smart and low latency speech-to-speech dialogue model with frozen llm](#). *Preprint*, arXiv:2411.00774.
- Jian Wu, Yashesh Gaur, Zhuo Chen, Long Zhou, Yimeng Zhu, Tianrui Wang, Jinyu Li, Shujie Liu, Bo Ren, Linqun Liu, and Yu Wu. 2023. [On decoder-only architecture for speech-to-text and large language model integration](#). *Preprint*, arXiv:2307.03917.
- Chen Xu, Rong Ye, Qianqian Dong, Chengqi Zhao, Tom Ko, Mingxuan Wang, Tong Xiao, and Jingbo Zhu. 2023. [Recent advances in direct speech-to-text translation](#). In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence, IJCAI '23*.
- Rong Ye, Mingxuan Wang, and Lei Li. 2022. [Cross-modal contrastive learning for speech translation](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5099–5113, Seattle, United States. Association for Computational Linguistics.
- Chenyu You, Nuo Chen, Fenglin Liu, Shen Ge, Xian Wu, and Yuexian Zou. 2022. [End-to-end spoken conversational question answering: Task, dataset and model](#). In *Findings of the Association for Computational Linguistics: NAACL 2022*, pages 1219–1232, Seattle, United States. Association for Computational Linguistics.
- Jianing Zhou, Ziheng Zeng, Hongyu Gong, and Suma Bhat. 2024. [CLASP: Cross-modal alignment using pre-trained unimodal models](#). In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 11518–11531, Bangkok, Thailand. Association for Computational Linguistics.

Yi Zhu, Zexun Wang, Hang Liu, Peiying Wang, Mingchao Feng, Meng Chen, and Xiaodong He. 2022. [Cross-modal transfer learning via multi-grained alignment for end-to-end spoken language understanding](#). In *Interspeech 2022*, pages 1131–1135.

Appendix Overview

The appendix includes the following information:

- Hyperparameter and model details (§A)
- Data details (§B)
- Prompt details (§C)
- Contrastive training over multiple layers (§D)
- Details on the mixed-nwp pretraining (§E)
- Details on combining pretraining losses (§G)
- Details on pretraining with more data (including translation results for all language pairs) (§H)
- Results for capturing paralinguistic features (§I) on classification as well as generation tasks

A Hyperparameters and Model Details

All hyperparameters, model and training parameters are listed in Table 7.

Our SpeechLLM architecture consists of three core components: a speech encoder, a text-based LLM, and a projector bridging the two. We select facebook/hubert-large-ls960-ft (Hsu et al., 2021) as our speech encoder, Llama-3.1-8B-Instruct (Dubey et al., 2024) as our LLM and Q-Former (Li et al., 2023) as the projector.

HuBERT has an apache-2.0 license, the llama model has a llama3.1 license. We adhere to both license requirements, when using the models for research.

For contrastive learning, we compare speech and text embeddings without the use of system prompts. This creates a positional mismatch between contrastive pretraining and fine-tuning. When using ROPE embeddings (Su et al., 2023), this impacts contrastive pretraining at layers beyond the embedding layer. To address this potential overfitting problem, we experiment with varying the absolute starting position in RoPE by randomly sampling it during pretraining for each example.

B Data

We fine-tune and evaluate our contrastive pretraining method on three task: Automatic Speech Recognition (ASR), Speech Translation (ST), and Speech Question Answering (SQA).

training parameters	Q-Former Num Query Token	4
	Q-Former Num Hidden Layers	4
	Q-Former Num Attention Heads	12
	Q-Former Seconds per Window	1/3
	per device batch size	10
	gradient accumulation steps	2
	num GPUs	4
	learning rate	1e-4
	warmup ratio	0.03
	optimizer	adamw_torch
	learning rate scheduler type	cosine
model max length	2048	
gradient clipping	1	
pretraining specific	num epochs	5
	contrastive τ cos + wasser	0.1
	contrastive τ nwp	0.5
	sinkhorn loss p	2
	sinkhorn loss blur	0.5
finetuning	num epochs	2

Table 7: Hyperparameters for the trainings, which are conducted on four NVIDIA A100-SXM4-40GB GPUs, mostly following Verdini et al. (2024)

For ASR and ST, we use the MustC-v1 dataset (Di Gangi et al., 2019). We use the en-de portion as English ASR data, and we use language pairs en-de, en-fr, en-it, and en-es for ST. We exclude examples longer than 45 seconds for training and use the *tst-common* and *tst-he* test sets for evaluation.

For question answering, we use Spoken-SQuAD (Lee et al., 2018), which is based on Stanford Question Answering Dataset (Rajpurkar et al., 2016, SQuAD). To allow for higher batch sizes during training, we exclude examples exceeding 150 seconds.

For pretraining, we use the MustC ASR data (approximately 400 hours), but also experiment with adding the M-version of Gigaspeech (Chen et al., 2021), which includes 1,000 hours of speech from audiobooks, podcasts, and YouTube.

Contrastive pretraining aligns text and speech, which could potentially lead to the loss of paralinguistic features in the embeddings. Therefore, we also train and test our models on *mlls-eng-speaker-descriptions* (Pratap et al., 2020; Lacombe et al., 2024; Lyth and King, 2024). We use the *mlls-eng-speaker-descriptions* train and test set. The available training data is quite large, so we subsample it to make training feasible. Since the test set is pre-defined, we sample the training data such that the test data size is 5% of the data. We let the model predict the following features: Speaking rate, gender, noise, reverberation, speech monotony, SDR noise, and PESQ speech quality. Moreover, we test

our models on MuST-SHE (Bentivogli et al., 2020), a translation test set, where the speaker gender is needed to translate the segment correctly. We use the en-es (), en-it and en-fr portion of MuST-SHE.

Details on prompts can be found in Appendix C.

C Prompts

In this sections we report the prompt we used for the different tasks during inference. During training, for each example, the prompt is sampled from a list of multiple prompts. Here, we report the prompts used for testing, for the training prompts, we refer the reader to our codebase.

Task	Prompt
ASR	Can you transcribe this audio?
SQA	Listen to the audio and answer this question:

Table 8: ASR and SQA Prompts used for inference. We finetuned on a range of different prompts, please refer to our codebase for these.

Prompt
en Can you translate this audio from <sl> into <tl>?
fr Pouvez-vous traduire cet audio de <sl> en <tl>?
it Puoi tradurre questo audio da <sl> in <tl>?
es ¿Puedes traducir este audio de <sl> al <tl>?
de Können Sie dieses Audio von <sl> zu <tl> übersetzen?

Table 9: ST Prompts used for inference. We finetuned on a range of different prompts, please refer to our codebase for these. For <sl> the source language is added, for <tl> the target language is added.

Prompt
Evaluate the audio file and provide details in the following order: speaking rate, gender, noise level, reverberation, speech monotony, SDR noise, and PESQ speech quality. Format the response as: <speaking_rate>, <gender>, <noise>, <reverberation>, <speech_monotony>, <sdr_noise>, <pesq_speech_quality>.

Table 10: Prompts for the MLS speaker description dataset used for inference. We finetuned on a range of different prompts, please refer to our codebase for these.



(a) *contr-cos* pretrained models(b) *contr-wasser* pretrained models

Figure 2: Analysis of contrastive pretraining performed on different layers. We report the normalized average over tasks as in Eq. (3) for contrastive pretraining and finetuning on 10% of task-specific data. The last point on the x-axis refers to contrastive pretraining using the loss summed over all layers.

FT Data	Model	ASR		ST		SQA		Contr loss on test		Norm. avg.↑
		WER↓*	BLEU↑	COMET↑*	EM↑	F1↑*	cos.	wasser.		
N/A	Specialized	6.54	30.99	80.02	64.19	77.10	N/A	N/A	100	
N/A	HuBERT + Llama	18.38	19.85	73.92	36.15	54.76	N/A	N/A	0	
100%	no pretrain	12.09	28.84	79.94	64.13	76.53	1.37	1.26	83.08	
10%	contr-cos-all (multiple of 5)	13.06	27.19	78.29	60.48	73.90	1.08	0.91	67.39	
	contr-cos-all	13.08	27.25	78.39	61.45	74.91	1.14	1.05	69.41	

Table 11: Initial ablations show that calculating the loss over multiples of 5 layers approximates calculating the loss over all layers efficiently. Results of models with contrastive pretraining using cosine similarity (*contr-cos*) on all (*-all*) layers, followed by fine-tuning on a 10% subset of task-specific data. The specialized baselines are Whisper (Radford et al., 2022) for ASR, Seamless (Communication et al., 2023) for ST, and You et al. (2022) for SQA. Metrics with * contribute to the normalized average as in Eq. (3).

D Contrastive Loss over Multiple Layers

Contrastive Loss over Different Layers For the contrastively pretrained model, we experiment with calculating the loss over different layers (Fig. 2), using the final representation of each layer in the LLM. We find that applying contrastive loss to deeper layers significantly increases downstream performance, however, the best layer to use is not consistent. Thus, using a general approach, such as calculating the loss over multiple layers, is the most likely to be robust to different settings. However, calculating the loss over all layers is computationally expensive. Specifically, when using the more computationally expensive Wasserstein distance, the total pretraining time increases by a factor of 4 when applied to all layers. In contrast, for the more efficient cosine similarity, the runtime only increases by 10%. Initial experiments showed that using every 5th layer approximates the results well (Table 11), so we adopted this more efficient method for all experiments.


Contrastive Loss over Different Layers on Different Tasks

Further analysis reveals that contrastive learning on deeper layers benefit the ST and SQA task the most, while improvements for ASR are less consistent. Fig. 3 shows the finetuning performance on 10% on task-specific data after pretraining on different layers for the tasks ASR, ST and SQA.


Contrastive Loss Analysis during Pretraining

Fig. 4 shows the contrastive loss in the pretraining phase. The later the layer, the lower the contrastive loss, with an exception of the embedding layer. For cosine contrastive training, these differences are significantly bigger than for wasserstein contrastive training. Here the losses for pretraining seem to correlated mostly with the performance of the model after finetuning: Models perform better after contrastive pretraining at later layers. However, for ASR, this trend is not so noticeable.


Position Shifts during Pretraining To address positional mismatches between pretraining and finetuning (detailed in Appendix A), we experiment




(a) ASR: contr-cos pretrained




(b) ASR: contr-wasser pretrained




(c) ST: contr-cos pretrained



(d) ST: contr-wasser pretrained



(e) SQA: contr-cos pretrained



(f) SQA: contr-wasser pretrained

Figure 3: Analysis of a contrastive pretraining performed on multiple layers for different tasks. The performance is reported after finetuning on 10% of the task specific data. The last point on the x-axis refers to contrastive pretraining using the loss sum over all layers.

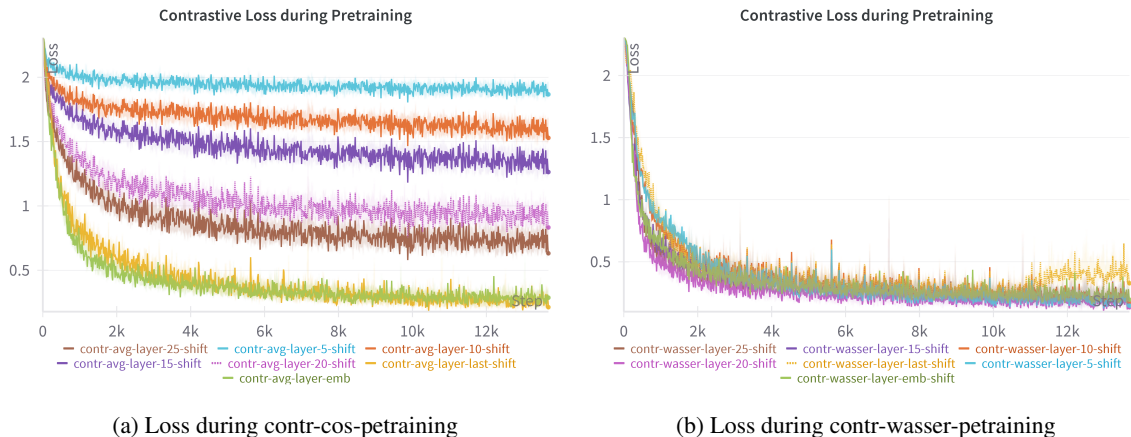


Figure 4: Contrastive loss during pretraining performed on multiple layers for models trained with contrastive average loss (left) and contrastive wasserstein loss (right).

FT Data	Model	ASR	ST		SQA		Norm.
		WER*	BLEU	COMET*	EM	F1*	avg.
10%	no pretrain	23.78	19.87	69.72	34.72	46.07	-51.13
	ASR pretrain	12.21	24.82	75.70	49.48	63.36	39.93
10%	mixed-nwp w/o punctuation (p)	14.78	23.85	75.78	47.35	66.39	37.69
	mixed-nwp w/o punctuation (k)	15.27	25.22	76.05	53.57	67.68	39.67
	mixed-nwp with punctuation (k)	15.76	24.23	75.09	47.65	61.36	23.59
	mixed nwp w/o punctuation + giga (k)	14.91	26.07	77.17	56.99	70.94	51.67
10%	contr-cos-emb + mixed (k)	14.26	25.96	76.61	51.97	68.40	46.61
	contr-wasser-emb + mixed (k)	14.70	24.15	75.05	48.01	61.48	26.53
	contr-cos-last + mixed (k)	19.02	23.02	73.92	46.29	58.56	3.86
	contr-wasser-last + mixed (k)	16.03	23.09	74.14	43.84	58.20	12.95
	contr-cos-all + mixed (k)	17.63	22.90	73.83	42.15	55.12	2.14

Table 12: Ablation studies for models with mixed speech and text input. Metrics with * contribute to the normalized average as in Eq. (3). We experiment with two different forced alignment algorithms, one by [Pratap et al. \(2023, \(p\)\)](#) and one by [Kurzinger et al. \(2020, \(k\)\)](#). Since the latter performs better in our scenario, we use this aligner for all subsequent experiments.

with position shifts during pretraining. Specifically, we experiment with varying the absolute starting position in RoPE by randomly sampling it during pretraining for each example. However, these do not consistently improve performance (Fig. 2).

E Mixed Next Word Prediction

For alignment between text and speech, we utilize the forced alignment algorithm described by [Kurzinger et al. \(2020\)](#), which leverages wav2vec 2.0 ([Baevski et al., 2020](#)). We also experimented with the alignment algorithm by [Pratap et al. \(2023\)](#), which performed slightly worse in our scenario (Table 12).

We then randomly select spans from a sequence

and assign them either text or speech embeddings, ensuring that sequences begin with text or speech equally often. Because speech embeddings are generally longer than text embeddings for the same content, we select spans of 2 to 5 words for speech and 4 to 10 words for text, inspired by [Nguyen et al. \(2024\)](#) who use a similar distribution. However, unlike [Nguyen et al. \(2024\)](#), we do not introduce modality-specific tokens, as their framework uses these tokens to predict speech output — a capability we do not require.

In the alignment process for speech and text, the aligner generates timestamps for each word. If there is a pause between words, this pause is typically excluded from the alignment. However,

FT Data	Model	ASR		ST	SQA		Norm.
		WER↓*	BLEU↑	COMET↑*	EM↑	F1↑*	avg.↑
10%	ASR pretrain	12.21	24.82	75.70	49.48	63.36	39.93
10%	contr-cos-emb	13.98	25.47	76.32	54.86	68.49	46.00
	+ asr loss	13.18	25.77	76.89	56.12	68.67	51.64
	+ mixed + nwp	14.68	25.28	76.18	48.86	63.92	36.41
	contr-wasser-emb	15.99	26.10	77.09	54.64	67.83	43.52
	+ asr loss	13.19	26.55	77.40	59.02	70.44	57.03
	+ mixed + nwp	16.00	23.98	74.59	45.00	59.11	16.87
10%	contr-cos-all	13.06	27.19	78.29	60.48	73.90	67.39
	+ asr loss	11.68	27.61	78.94	64.16	76.12	78.17
	+ mixed + nwp	22.08	23.16	74.20	48.47	62.13	2.07
	contr-wasser-all	12.92	29.07	80.52	64.06	77.22	84.96
	+ asr loss	11.23	29.06	80.26	66.73	77.19	88.24
	+ mixed + nwp	15.72	24.13	75.00	51.21	64.25	27.52

Table 13: Comparison of models with combinations of alignment losses. Metrics with * contribute to the normalized average as in Eq. (3).

to prevent the model from encountering unexpected silence or noise before a word during fine-tuning, we include any preceding silence or noise in the audio segment before the word.

F Mixed Speech-Text Inputs

We experiment with mixed speech-text input as described in Section 3.2, and report results in Table 12.

For the mixed speech-text input, we experiment with two different forced alignment algorithms, one by Pratap et al. (2023, (p)) and one by Kurzinger et al. (2020, (k)). Since the latter performs better in our scenario, we use this aligner for all subsequent experiments.

Pretraining with mixed speech-text next word prediction (nwp) achieves similar results as ASR pretraining. Here, removing punctuation seems to help the model understand the NWP task better.

We also test contrastive pretraining with mixed speech-text input (contr-*+mixed) and find that aligning the representation on the embedding layer works best (however, contrastive learning without mixed input still works better).

G Combined Losses

This section explores the combination of different losses. We combine 1) ASR and contrastive loss

and 2) NWP and contrastive loss. The result can be found in Table 13. While 1) leads to significant improvements, 2) does not seem to work together.

H Adding more pretraining data

Results on 10% of the data We analyze the effect of contrastive pretraining with more data, adding the Giga dataset to our pretraining data. Results for finetuning on 10% of the data can be found in Table 14.

Results on 100% of the data The results on 100% of the data can be found in Table 4

Results on different translation directions For the models we discuss in Section 5.4, we also report detailed results for the translation tasks. Table 15 contains the results for the different translation directions.

FT Data	Model	ASR	ST		SQA		Contr loss on test		Norm.
		WER↓*	BLEU↑	COMET↑*	EM↑	F1↑*	cos.	wasser.	avg.↑
N/A	Specialized	6.54	30.99	80.02	64.19	77.10	N/A	N/A	100
N/A	BLSP-lslm-7b	44.51	28.70	78.68	5.60	21.82	N/A	N/A	-96.72
N/A	Qwen2-Audio-7b	12.03	21.57	74.81	27.79	46.75	N/A	N/A	10.79
10%	ASR pretrain	12.21	24.82	75.70	49.48	63.36	1.37	1.30	39.93
	+ giga	12.16	27.44	78.38	62.23	73.17	1.37	1.31	69.37
	contr-cos-emb	13.98	25.47	76.32	54.86	68.49	0.97	0.72	46.00
	+ giga	14.24	28.76	80.18	63.94	76.60	0.97	0.47	78.45
	contr-wasser-emb	15.99	26.10	77.09	54.64	67.83	1.08	0.62	43.52
	+ giga	16.36	26.47	77.59	58.05	71.09	1.06	0.58	50.12
	contr-cos-all	13.06	27.19	78.29	60.48	73.90	1.08	0.91	67.39
	+ giga	10.94	29.95	81.21	65.45	79.12	1.06	0.71	97.12
	contr-wasser-all	12.92	29.07	80.52	64.06	77.22	1.07	0.64	84.96
	+ giga	12.29	29.30	80.79	68.13	80.29	1.09	0.60	92.75
	contr-cos-all + asr	11.68	27.61	78.94	63.81	76.12	1.08	0.91	78.17
+ giga	11.12	29.90	81.29	72.16	82.52	0.98	0.65	102.15	
contr-wasser-all + asr	11.23	29.06	80.26	66.29	77.19	1.09	0.66	88.24	
+ giga	15.13	30.10	81.45	70.47	82.21	1.03	0.30	91.26	
100%	ASR pretrain	10.28	30.31	80.98	69.87	80.92	1.37	1.22	100.24
	+ giga	11.48	30.46	81.03	72.25	82.56	1.37	1.23	99.77
	contr-cos-all	12.56	30.69	81.48	71.71	82.50	1.25	1.02	99.06
	+ giga	9.31	31.36	81.98	75.12	84.67	1.10	0.83	114.18
	contr-wasser-all	10.35	30.62	81.59	68.65	80.21	1.24	0.91	102.46
	+ giga	12.10	30.90	81.65	75.49	84.52	1.25	0.87	104.32
	contr-cos-all + asr	10.04	30.82	81.56	72.40	82.47	1.19	1.03	106.57
	+ giga	10.03	31.54	81.99	76.11	84.94	1.10	0.79	112.66
	contr-wasser-all + asr	12.08	30.75	81.62	73.13	83.24	1.33	0.93	102.31
+ giga	9.50	31.20	82.00	74.27	84.39	1.11	0.47	113.37	

Table 14: Comparison of pretraining only on Must-C data or the combination with the bigger Giga dataset (+giga). We compare our models against BLSP (Wang et al., 2024a) and the Qwen2-Audio-7B-Instruct model (Chu et al., 2024). Metrics with * contribute to the normalized average as in Eq. (3).

FT Data	Model	BLEU \uparrow				COMET \uparrow^*			
		en-de	en-es	en-fr	en-it	en-de	en-es	en-fr	en-it
N/A	Seamless	26.25	33.10	36.49	28.15	79.03	80.38	79.92	80.75
N/A	BLSP-lslm-7b	23.3	27.4	31.9	23.2	76.6	78.7	77.7	78.2
N/A	Qwen2-Audio-7b	19.60	21.02	27.83	17.83	74.44	74.47	75.69	74.65
10%	ASR pretrain	20.31	26.94	28.34	21.71	74.32	76.30	74.56	76.03
	+ giga	23.25	29.81	32.42	24.28	77.56	79.12	77.68	79.17
	contr-cos-all	23.23	29.20	32.05	24.27	77.42	78.83	77.67	79.22
	+ giga	25.29	32.68	35.27	26.55	80.44	81.64	80.97	81.79
	contr-cos-all + asr	23.51	29.85	32.62	24.46	78.25	79.42	78.43	79.66
	+ giga	25.56	32.34	35.01	26.67	80.63	81.70	80.97	81.87
100%	ASR pretrain	25.92	32.61	35.64	27.07	80.20	81.36	80.64	81.70
	+ giga	26.05	32.55	35.90	27.33	80.40	81.25	80.61	81.87
	contr-cos-all	26.06	33.20	36.09	27.40	80.91	81.88	81.07	82.04
	+ giga	26.63	33.81	36.90	28.20	81.23	82.30	81.74	82.63
	contr-cos-all + asr	26.12	33.23	35.98	27.93	80.86	82.02	81.24	82.11
	+ giga	26.88	33.64	37.15	28.49	81.37	82.34	81.70	82.56

Table 15: Comparison of translation performance of pretraining only on Must-C data or the combination with the bigger Giga dataset (+giga). We compare our models against BLSP (Wang et al., 2024a) and the Qwen2-Audio-7B-Instruct model (Chu et al., 2024).

I Impact of text-speech alignment of capturing paralinguistic features.

Speech data is much richer than a sequence of words, and the pretrained speech encoder (HuBERT) might capture additional paralinguistic features. However, contrastive and ASR pretraining could potentially lead to discarding these paralinguistic features, when pulling speech embeddings closer to text embeddings. To assess this, we test our models on a) on paralinguistic classification tasks and b) a generation task where paralinguistic information is needed.

Classification Task on Paralinguistic Features

For a) we finetune our models on paralinguistic classification tasks before and after pretraining. We also compare our models against a HuBERT + linear head baseline, to assess which features the encoder captures in the first place.

To this end, we use the *mls-eng-speaker-descriptions* train and test set (Pratap et al., 2020; Lacombe et al., 2024; Lyth and King, 2024). This dataset contains classification tasks for rate, gender, background noise, and other features (more details in Appendix B). The results demonstrate that models with contrastive and ASR-based pretraining perform comparably to the plain HuBERT model in predicting paralinguistic annotations, indicating that alignment-focused pretraining does not significantly affect the retention of paralinguistic information. Table 5 shows the accuracies for the different classification tasks for the different models.

In a second step, we assess whether adding paralinguistic features during training (again using the *mls-dataset*), influences the models performance on downstream tasks. We finetune on a 10% subset of finetuning data to simulate a low-resource scenario. We find that including the paralinguistic data does

not lead to significantly different results, which is not surprising given that the tasks that we test do not rely on paralinguistic information. Results can be found in Table 16.

Generation Task that Requires Paralinguistic Features

For b) we test our models on MuST-SHE (Bentivogli et al., 2020), a translation dataset where the speaker’s gender is necessary to correctly translate. We test the same models as for a) but also include a cascaded model (HuBERT + Llama31Instruct) as a baseline where all paralinguistic features are lost due to the nature of the cascaded model.

The results can be found in Table 17. Following Bentivogli et al. (2020), we report the average gender accuracy using the evaluation script that comes with the data, and sensitivity, the score difference between the scores of correctly and incorrectly generated references. We use the COMET difference (COMET Diff) here.

Across all language pairs, the contrastive models consistently outperform the cascaded model and the ASR pretrained model in terms of gender accuracy. On average across the language pairs, ASR pretraining performs 1.9% points above the baseline, whereas the best performing contrastive model (cos-all) performs 6.2% points above the cascaded baseline. The COMET Diffs are also larger for the contrastively pretrained models than for ASR pretraining and the baseline.

This underlines our findings that contrastive learning does not mitigate a model’s ability to reason about paralinguistic features. In contrast, ASR pretraining is not as effective at preserving this ability.

FT Data	Model	ASR	ST		SQA		Norm.
		WER↓*	BLEU↑	COMET↑*	EM↑	F1↑*	avg.↑
10%	ASR pretrain	12.21	24.82	75.70	49.48	63.36	39.93
	+ meta	12.11	24.91	75.70	49.47	63.31	40.15
10%	contr-cos-all	13.06	27.19	78.29	60.48	73.90	67.39
	+ meta	13.53	27.10	78.21	61.02	74.75	66.93
	contr-wasser-all	12.92	29.07	80.52	64.06	77.22	84.96
	+ meta	11.97	29.28	80.48	66.12	78.25	88.92
10%	contr-cos-all + asr	11.68	27.61	78.94	63.81	76.12	78.17
	+ meta	12.01	27.82	78.92	63.11	75.98	76.90
	contr-wasser-all + asr	11.23	29.06	80.26	66.29	77.19	88.24
	+ meta	11.16	29.34	80.55	65.40	78.26	91.60

Table 16: Comparison of models with and without pretraining, and with and without meta-speaker data in the finetuning process (on 10%). Metrics with * contribute to the normalized average as in Eq. (3).

Lang. Pair	Model	Avg. Gender Acc in %	COMET correct	COMET diff
en-es	HuBERT + LLama	73.52	78.69	0.7882
	no pretrain	69.67	70.89	0.6510
	ASR pretrain	73.19	78.27	0.8523
	contr-cos-all	77.64	82.97	1.1382
	contr-wasser-all	76.73	82.13	1.0784
	contr-cos-all + asr	76.14	80.67	1.0183
	contr-wasser-all + asr	75.50	82.30	1.0809
en-fr	HuBERT + LLama	70.15	75.13	0.5278
	no pretrain	70.34	68.66	0.4591
	ASR pretrain	72.64	75.94	0.6852
	contr-cos-all	77.83	82.92	1.1118
	contr-wasser-all	75.48	79.97	0.8856
	contr-cos-all + asr	75.73	78.57	0.8594
	contr-wasser-all + asr	75.25	79.96	0.8909
en-it	HuBERT + LLama	72.96	77.90	0.7401
	no pretrain	71.83	70.66	0.7281
	ASR pretrain	76.50	78.16	0.9684
	contr-cos-all	79.79	80.71	1.0103
	contr-wasser-all	78.40	82.29	1.0551
	contr-cos-all + asr	76.44	80.78	0.9682
	contr-wasser-all + asr	77.90	81.93	0.9858

Table 17: Results for the MuST-SHE dataset (Bentivogli et al., 2020) for different models. Following Bentivogli et al. (2020), we evaluate the overall gender accuracy (using the evaluation script available with the dataset), and the score difference between the scores of correctly and incorrectly gendered references. We use the COMET difference (COMET Diff) here.