

# NAVER LABS Europe Submission to the Instruction-following Track

Beomseok Lee<sup>1,2,3\*</sup> Marceley Zanon Boito<sup>2,\*</sup> Laurent Besacier<sup>2</sup> Ioan Calapodescu<sup>2</sup>

<sup>1</sup> University of Trento

<sup>2</sup> NAVER LABS Europe

<sup>3</sup> Fondazione Bruno Kessler

contact: marceley.zanon-boito@naverlabs.com

## Abstract

In this paper we describe NAVER LABS Europe submission to the instruction-following speech processing short track at IWSLT 2025. We participate in the constrained settings, developing systems that can simultaneously perform ASR, ST, and SQA tasks from English speech input into the following target languages: Chinese, Italian, and German. Our solution leverages two pretrained modules: (1) a speech-to-LLM embedding projector trained using representations from the SeamlessM4T-v2-large speech encoder; and (2) LoRA adapters trained on text data on top of Llama-3.1-8B-Instruct. These modules are jointly loaded and further instruction-tuned for 1K steps on multilingual and multimodal data to form our final system submitted for evaluation.

## 1 Introduction

Large language models (LLMs) have demonstrated remarkable success across various text-based natural language processing tasks (Achiam et al., 2023; Touvron et al., 2023; Jiang et al., 2024; Yang et al., 2024; Alves et al., 2024; Martins et al., 2024a), motivating research into extending them to other modalities. This has led to the development of multimodal LLMs capable of processing speech, audio, images and video (Team et al., 2023; Driess et al., 2023; Rubenstein et al., 2023; Liu et al., 2023; Tang et al., 2023; Défossez et al., 2024; Hu et al., 2024; Laurençon et al., 2024; Huang et al., 2024; Nguyen et al., 2025; Ambilduke et al., 2025).

This year *IWSLT Instruction-following Speech Processing Track* focuses on the leveraging of LLMs and speech foundation models (SFM) to build solutions capable to perform multilingual tasks from English speech input and textual multilingual instructions (Abdulmumin et al., 2025). NAVER LABS Europe (NLE)

participates in the constrained setting of the *short track*, where the tasks proposed are automatic speech recognition (ASR), speech translation (ST) and multilingual spoken question answering (SQA). The target languages for ST and multilingual SQA are Chinese, Italian and German. The participants are allowed to use the speech backbone SeamlessM4T-v2-large (Barraut et al., 2023) and the text LLM Llama-3.1-8B-Instruct (Grattafiori et al., 2024) for both training and data generation.

Our submitted systems leverage all the available data from the constrained settings, together with data automatically obtained using both backbones. We train two types of systems in parallel: (1) speech-to-text ASR/ST/SQA projectors that project the averaged speech representation from the SeamlessM4T-v2-large encoder to the embedding space of a frozen Llama-3.1-8B-Instruct; (2) text-only LoRA adapters (Hu et al., 2022), plugged on top of the same frozen LLM. Once both systems are separately trained, we show that we can merge them, increasing overall speech performance, by fine-tuning for only 1K steps on multimodal multilingual data.

This system paper is organized as follows. Section 2 describes the preprocessing applied to the data used in this challenge. Sections 3 and 4 describe our training pipeline and experimental settings, respectively. Section 5 presents our experiments and discussion. Section 6 presents the submitted system. Section 7 concludes the paper.

## 2 Data

For training our models, we leverage the data from the constrained setting: CoVoST2 (Wang et al., 2020), EuroParlST (Iranzo-Sánchez et al., 2020) and SpokenSQuAD (Lee et al., 2018). With the agreement of the organizers, we also take advantage of the SeamlessM4T-v2-large to

\* Equal contribution.

produce extra synthetic speech data (Seamless TTS) and multilingual text data (Seamless MT). Llama-3.1-8B-Instruct is used to rephrase SQA answers. ACL 60-60 (Salesky et al., 2023) is used for validation and evaluation only. We now present our data preprocessing (Section 2.1) and prompt format (Section 2.2).

## 2.1 Data Preprocessing

We produce both speech-to-text and text-to-text instructions to train our systems. For the aforementioned datasets, we produce the following splits, where \* denotes synthetic splits obtained via SeamlessM4T-v2-large MT; † indicates splits generated with SeamlessM4T-v2-large TTS; and ‡ marks those derived through Llama-3.1-8B-Instruct-based rephrasing.

- **CoVoST2:**
  - ASR and ST/MT (en-de; en-zh)
- **EuroParlST:**
  - ASR and ST/MT (en-de; en-it)
- **SpokenSQuAD:**
  - ASR<sup>†</sup> and MT (en<sup>†</sup>-de\*; en<sup>†</sup>-it\*; en<sup>†</sup>-zh\*)
  - SQA/QA (en<sup>†</sup>-en; en<sup>†</sup>-de\*; en<sup>†</sup>-it\*; en<sup>†</sup>-zh\*)
  - *fluent* SQA/QA (en<sup>†</sup>-en<sup>‡</sup>; en<sup>†</sup>-de<sup>\*,‡</sup>; en<sup>†</sup>-it<sup>\*,‡</sup>; en<sup>†</sup>-zh<sup>\*,‡</sup>)
- **ACL 60-60:**
  - ASR and ST/MT (en-de; en-it\*; en-zh)

Below we detail dataset-specific preprocessing. Statistics are presented in Table 1.

**CoVoST2 and EuroParlST** CoVoST2 covers English to German and Simplified Chinese language directions. EuroParlST covers English to German and Italian. ASR splits for these datasets were built by merging the existing language splits and deduplicating the audio files. For both, language-specific ST and MT splits are created by aligning translations to English speech and reference transcriptions, respectively.

**SpokenSQuAD** The SpokenSQuAD dataset is organized into two jsons, train and test. Each split is organized in themes, each with several paragraphs. For each paragraph, TTS audio files are available, aligned at the sentence level.<sup>1</sup> For each

sentence, questions (each with several answers) are available. We performed the following modifications to this dataset:

- **Duplicated answers:** we removed duplicated answers to the same question using exact string matching, as well as any questions that required more than one audio file to answer, since we are participating in the SHORT track.
- **Validation set:** we created a validation set by selecting the first 20 themes of the training set and removing them from training (3,102 entries).
- **New TTS audio:** we generate new TTS audio files for the training set using SeamlessM4T-v2-large. Resynthesizing the source audio for this dataset was necessary due to existing dataset misalignment, which we detail in Appendix Section A.1.
- **Multilingual SQA/QA:** we created multilingual SQA/QA sets by first translating questions and answers to target languages using SeamlessM4T-v2-large. We then use reference-free COMET<sup>2</sup> (Rei et al., 2022) to filter out all pairs of questions and answers that do not both score at least 0.85.
- **Invalid splits:** we created the *invalid* SQA sets by deliberately mismatching context and question themes, thereby creating unanswerable examples. The corresponding answers were labeled as “Not answerable” in four languages (English, Italian, German and Chinese), following the guideline answer provided by the task organizers. While we acknowledge that a small, unknown subset of these reassigned questions may still be answerable, we hypothesize that ensuring a thematic mismatch between the reference context and the question is the most effective strategy for minimizing this issue.
- **Fluent SQA/QA version:** we created an alternative SQA/QA training set using Llama-3.1-8B-Instruct to regenerate the dataset original answers as fluent sentences. The motivation behind this was the observation that, since the original answers are made of an exact extract of the reference audio/text,

<sup>1</sup>The format is themeID\_paragraphID\_sentenceID.

<sup>2</sup>Unbabel/wmt22-cometkiwi-da

Dataset	Task	Language	# Samples	
CoVoST2	ASR	en	289,413	
	ST/MT	en-de	289,413	
		en-zh	289,413	
EuroParlST	ASR	en	35,372	
	ST/MT	en-de	32,628	
		en-it	29,552	
SpokenSQuAD	ASR	en	34,003	
		en-de	39,362	
		MT	en-it	55,030
			en-zh	25,078
	SQA/QA	en-en	34,003 <sup>†</sup>	
		en-de	6,574 <sup>†</sup>	
		en-it	16,767 <sup>†</sup>	
		en-zh	7,093 <sup>†</sup>	
	fluent SQA/QA	en-en	32,320	
		en-de	4,169	
en-it		13,712		
en-zh		3,424		

Table 1: Training sets statistics by task. For ST/MT sets, target side is duplicated. For SpokenSQuAD, † highlights that the source speech is used twice (valid and invalid questions, as described in Section 2.1).

the model had a difficult time answering some out-of-domain questions fluently.<sup>3</sup> More details are presented in Appendix Section A.2.

**ACL 60-60** We use SeamlessM4T-v2-large to generate Italian translations, since the data shared only contained en-de and en-zh splits. We leverage the *dev* set for checkpoint selection during training. The *eval* set is used for testing.

## 2.2 Prompt Format

The goal of the short track of this challenge is to produce a model that is capable to 1) transcribe English speech; 2) translate English speech into Italian, German and Chinese; 3) Answer multilingual questions using English speech as input. In this setting, the language of the question must match the language of the answer.

To develop a model capable of smoothly switching between different tasks, we designed task prompts with a consistent structure: regardless of the task (ASR, ST, or SQA), the user turn begins by encapsulating the speech embeddings within textual tags. This is followed on a new line by a task-specific instruction formulated as a question, and finally, another line containing a common

<sup>3</sup>In the official IWSLT 2025 test set we observed examples as the following. Question: “What are the names of the speakers?” Our model’s answer: “yin and my colleague jiang”. While the model answer is an exact and correct extraction from the audio, we were unsure about how this would be considered during evaluation.

suffix. The list of templates used is available in Appendix Table 5.

## 3 Training Pipeline

Our training pipeline is illustrated in Figure 1. We first train a speech projector on speech tasks (A), and text LoRA weights on textual tasks (B). These modules are then reloaded and adapted together on both speech and textual tasks (C). In this section we describe the key components used and the data sampling strategy.

**Foundation Models** For speech, we leverage SeamlessM4T-v2-large model, extracting speech representations for all our audio data from its 24th speech encoder layer (i.e. the last layer). Prior to training, we average every 3 consecutive frame vectors, reducing significantly the sequence length. This simple trick allows us to train our models with larger batches, while maintaining good performance in speech tasks. All our models are built on top of a frozen Llama-3.1-8B-Instruct.

**Speech Projector Architecture** The speech projector consists of 4 Transformer encoder layers, each with 8 attention heads. The input dimension is set to 1,024, the feed-forward network dimension to 2,048, and the output dimension to 4,096 to align with the embedding size of Llama-3.1-8B-Instruct. A dropout rate of 0.1 is applied throughout, and the model employs pre-layer normalization.

**LoRA Adapters** LoRA adaptation (Hu et al., 2022) is applied to both the self-attention (Q/K values, output projection) and feed-forward modules, and across all LLM layers. We use  $rank = 8$ ,  $\alpha = 16$ . We do not use dropout.

**Data Sampling Strategy** For training all our models, we define an epoch as  $X$  steps across the dataset, with  $X = \frac{|\text{speech\_examples}|}{\text{batch\_size}}$ . To construct the training data for each epoch, we sample batches by first applying the predefined task-level sampling ratios, followed by sampling based on the internal domain-level splits within each task. In the case of multimodal training (speech and text tasks mixed), we consider speech as our *main* modality, using it for defining epoch size and task ratio. Each time we sample a task and language split that has a textual equivalent (e.g., ST corresponds to MT; SQA to QA), we also sample a batch from the corresponding textual task. In practice, this means that every

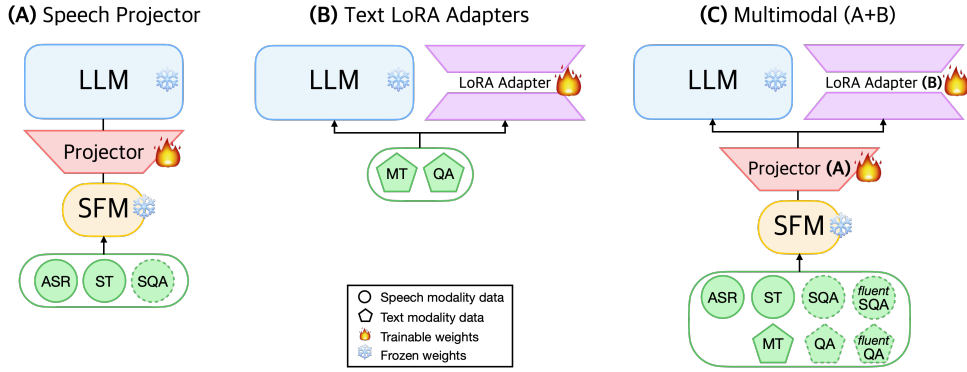


Figure 1: Our training pipeline. A speech projector (A) and text LoRA adapters (B) are trained in parallel using speech-to-text and text-to-text data, respectively. These modules are then integrated during a brief multimodal adaptation step (C).

time a batch from ST en-de split is sampled, a batch from MT en-de follows it. We hypothesize that interleaving similar speech and textual batches during training provides regularization benefits, with the text data serving as a stabilizing signal for learning (Pikabea et al., 2025).

## 4 Experimental Settings

**Codebase** We train our models using an internal fork of torchtune (torchtune maintainers and contributors, 2024), which allows us to process interleaved representations of text and high-dimensional vectors within the user turn during instruction tuning. The high-dimensional vectors pass through our speech projector, while the text prefix and suffix user prompts are processed by the LLM embedding layer. The obtained speech and text embeddings are both concatenated and fed into the first layer of the LLM which is trained on the masked input with standard cross-entropy loss. Different learning rate schedulers and optimizers are employed for the speech projector and the LoRA weights, allowing for more controlled and effective training of these distinct model components.

**Inference Settings** We perform inference using torchtune, with a batch size of 1 and greedy decoding. The maximum number of new tokens was limited to 300. Unless explicitly stated otherwise, this decoding strategy was consistently applied across all experimental settings. Additional discussion regarding multimodal inference is present in Appendix Section B.1.

**Evaluation Metrics** We evaluate our models on speech (ASR, ST, SQA) and text (MT, QA) tasks when relevant. For ASR, we score word error rate

(WER) using HuggingFace evaluate library with default settings and MMS normalization (Pratap et al., 2024). For ST/MT we present two evaluation metrics: BLEU4 computed with sacrebleu library (Post, 2018)<sup>4</sup>, and COMET (Rei et al., 2022).<sup>5</sup> For SQA/QA, we use LLM-as-a-judge evaluation scripts from the bergen library<sup>6</sup> (Rau et al., 2024). We use their “yes/no” quality assessment evaluation format including the English reference text, the multilingual questions and the generated answers. We report average accuracy across four LLMs: EuroLLM-9B-Instruct (Martins et al., 2024b), Gemma3-12B-Instruct, Gemma3-27B-Instruct (Team et al., 2025), and Llama3.1-70B-Instruct.

**Baselines** We compare our results with both backbones we use for training. We evaluate MT and QA using the reference transcripts and Llama-3.1-8B-Instruct in zero-shot settings, and we evaluate SeamlessM4T-v2-large in both ASR and ST.

## 5 Experiments

We now present our results for ASR, ST, and SQA. Section 5.1 introduces the models used in our experiments, followed by results and discussion in Section 5.2.

### 5.1 Our Models

In this section, we describe the models used in our experiments. Additional hyperparameter details are provided in Appendix C.

<sup>4</sup>The signature is “nrefs:llcase:mixedleff:noltok:TOKENsmooth:explversion:2.3.1”, with  $TOK = zh$  for Chinese, and  $13a$  for the other languages.

<sup>5</sup>Unbabel/wmt22-comet-da

<sup>6</sup><https://github.com/naver/bergen>



**A.1 Speech Projector (ASR/ST)** This version of our speech projector focuses on ASR and ST tasks, which do not require any complex reasoning capability of the LLM. For ASR, we sample CoVoST2 and EuroParlST with probabilities 0.8 and 0.2 respectively. For ST, we sample proportionally to the datasets size, and set language sampling probabilities for the pairs en-de, en-zh, and en-it to 0.3, 0.4, and 0.3 respectively. We trained for 4 epochs using AdamW with learning rate of  $1e - 4$ , a constant learning rate scheduler, and gradient accumulation of 16. This model trains for 4.6 days in a single A100-80GB, and the best checkpoint is obtained after 18.2 hours.

**A.2 Speech Projector (ASR/ST/SQA)** This version of our projector extends A.1, including the SQA task. We use task sampling probabilities of 0.4, 0.35, and 0.25 for ASR, ST and SQA respectively. For the ASR and ST tasks, we use the same data ratios as defined above. For the SQA task, we leveraged both valid and invalid splits, with language-specific sampling probabilities for English, German, and Italian set to 0.4, 0.3 and 0.3, respectively. We do not train with Chinese SQA. This model trains for 4.75 days in a single A100-80GB, and the best checkpoint is obtained after 27.36 hours. The best checkpoints for both versions of the speech projector are selected using its average ST performance on the ACL 60-60 dev split across all language directions. Additional hyper-parameters for A.1 and A.2 are presented in Appendix Section C.1.

**B. Text-only LoRA (MT/QA)** We train LoRA weights on top of Llama-3.1-8B-Instruct using all text-to-text data from Table 1, and by using probability sampling of 0.6 and 0.4 for MT and QA respectively. We train for one epoch using AdamW with learning rate of  $3e - 4$ , weight decay of 0.1, and 100 warm-up steps. Batch size of 10, and gradient accumulation of 8 is used. This model trains for approximately 4 days in a single A100-80GB. We select the last checkpoint.

**C. Multimodal (A.x + B)** We restart training by using both one of the speech projectors detailed above, and the text-only LoRA weights. We adapt our models using all speech (ASR/ST/SQA) and textual (MT/QA) tasks. We experiment with two versions of the SQA/QA training sets: the original short lowercase answers, and the *fluent* SQA/QA version we created. In preliminary experiments, we

observed that as little as 100 steps were enough to successfully integrate the projector representation to the LoRA weights, but the best performance gains were obtained with 1K steps, which is the value we adopt for the experiments presented in the next section. We use learning rate of  $1e - 5$  for the speech projector, and of  $3e - 4$  for the LoRA weights. We use a batch size of 16, and gradient accumulation of 16. This model trains for approximately 6 hours in a single A100-80GB. We select the last checkpoint.

## 5.2 Results and Discussion

Table 2 presents our results for ASR, ST/MT and SQA/QA. ACL 60-60 *eval set* is used for ASR and ST/MT. SpokenSQuAD official test set is used for English SQA/QA. A smaller automatically obtained version is used for multilingual SQA/QA.<sup>7</sup> In the top portion, we present results for Llama-3.1-8B-Instruct before and after LoRA fine-tuning on text-only data. The middle portion of the table presents the speech backbone (SeamlessM4T-v2-large) and our projector-only models: for these rows, the only adaptation is the training of a speech projector that is plugged to a frozen Llama-3.1-8B-Instruct. Finally, the bottom portion of the table presents results from the merging of our text backbone (text-only LoRA) and the projectors of the middle portion via multimodal training. Additionally, ACL 60-60 ASR/ST dev results are presented in Appendix Table 7.

**Performance of Text-only Models (topline)** We observe that zero-shot Llama-3.1-8B-Instruct presents strong performance in both MT and QA tasks. By adding LoRA adapters on top of it, we increase translation performance in detriment of QA performance. We partially attribute this drop in QA performance to the SpokenSQuAD answer format, that is very short and might be judged as incomplete by the LLM evaluation. However, we also scored ROUGE1 (Lin, 2004) recall, which measures the intersection between the reference answer tokens and the produced one, finding that those scores were similar to the LLM-as-a-judge metric.<sup>8</sup> This result confirms that the QA performance is worse after text adaptation.

<sup>7</sup>Statistics for the multilingual test set are presented in Appendix Table 3.

<sup>8</sup>For en, de, it and zh splits ROUGE1 recall scores were respectively: 81.4%, 63.1%, 69.5%, 79.0%.

Model (fine-tuning tasks)	ASR (WER)			ST/MT (BLEU)			ST/MT (COMET)			SQA/QA (LLM-AS-A-JUDGE)			
	en	en-de	en-it	en-zh	en-de	en-it	en-zh	en-en	en-de	en-it	en-zh		
<b>Text-only Models (MT/QA)</b>													
Llama-3.1-8B-Instruct (zero-shot)	-	23.88	35.51	45.89	0.779	0.806	0.809	<b>91.8%</b>	<b>92.0%</b>	<b>88.6%</b>	<b>84.6%</b>		
<b>B.</b> Text-only LoRA (MT/QA)	-	<b>41.69</b>	<b>48.31</b>	<b>53.65</b>	<b>0.838</b>	<b>0.863</b>	<b>0.867</b>	83.4%	75.7%	71.4%	69.5%		
<b>Speech-only Models (ASR/ST/SQA)</b>													
SeamlessM4T-v2-large	<b>17.6</b>	<b>27.95</b>	<b>43.54</b>	33.58	0.737	0.788	0.753	-	-	-	-		
<b>A.1</b> Speech Projector (ASR/ST)	19.8	27.58	36.30	40.62	<b>0.760</b>	0.796	<b>0.793</b>	-	-	-	-		
<b>A.2</b> Speech Projector (ASR/ST/SQA)	19.9	27.20	36.60	<b>40.72</b>	<b>0.760</b>	<b>0.797</b>	0.792	0.7%	0.5%	0.3%	0.6%		
<b>Multimodal Models (ASR/ST/SQA)</b>													
<b>A.1 + B</b> (ASR/ST/MT/SQA/QA)	<b>17.7</b>	30.37	<b>41.22</b>	42.76	0.758	<b>0.791</b>	0.795	79.8%	71.9%	69.4%	65.5%		
<b>A.1 + B</b> (ASR/ST/MT/ <i>fluent</i> SQA/ <i>fluent</i> QA)	18.6	<b>30.75</b>	40.48	42.51	0.755	0.788	0.789	90.3%	85.2%	82.9%	76.4%		
<b>A.2 + B</b> (ASR/ST/MT/SQA/QA)	18.2	29.91	38.13	43.12	0.759	0.786	<b>0.799</b>	80.5%	74.9%	68.0%	66.7%		
<b>A.2 + B</b> (ASR/ST/MT/ <i>fluent</i> SQA/ <i>fluent</i> QA)	18.7	29.68	32.28	<b>43.38</b>	<b>0.763</b>	0.782	0.798	<b>91.1%</b>	<b>87.3%</b>	<b>84.8%</b>	<b>78.0%</b>		

Table 2: Results for the different models and backbones used in this work. ASR and ST scores are obtained using ACL 60-60 eval set, while SQA/QA scores are obtained using SpokenSQuAD test set.

**Speech Projectors** We observe that both A.1 and A.2 are equally capable of performing ASR and ST, which is consistent with the fact that both are trained on the same data. However, we observe that A.2, the model that trains with SQA data, is unable to produce SQA output. We believe this is a limitation of the projector approach: we train a model capable of biasing the output of the LLM, which works well for content tasks such as ASR and ST. For a reasoning task, further adaptation might be required in order to force the model to comply to the instruction. Additional results for CoVoST2 and EuroParlST are presented in Appendix Table 6, and they confirm that both models are very similar in ST performance.

**Multimodal Training** We observe that our efficient multimodal training is beneficial, consistently increasing scores for ASR and ST. Our multimodal models outperform the speech projector models by 1-2 WER points in the ASR task. For ST task, in BLEU scores, the multimodal models always outperform speech projector models while in COMET there are some mixed results, with some models presenting slight deterioration for some language pairs. Finally, in the SQA task, while the speech projector model A.2 failed to learn the task effectively, our multimodal models achieved strong results, outperforming the text topline (B) performance across all language settings, and reaching scores that are close to the Llama-3.1-8B-Instruct topline, despite working from the speech signal.

**ASR Performance** We observe that our models are competitive with SeamlessM4T-v2-large, scoring slightly worse than the baseline for some

configurations. Overall, for ACL 60-60 we observe quite elevated WER scores, compared to the ones we obtained for the training ASR datasets (see Appendix Table 6). We believe this is partially due to the nature of the dataset. We manually inspected some of the data, observing that some of the audio files contain *leftover fragments* from previous sentences. Moreover, looking at the transcriptions, we observed that these faithfully reproduced the audio, without applying any normalization or removing disfluencies (e.g. repeated and filler words were kept). We find that disfluent transcriptions are hardly produced by LLM-based models or SeamlessM4T-v2-large, that both tend to translate the content into a clean format. Therefore, we believe the WER scores presented in our results table are not representative of the models real ASR capabilities.

**German MT/ST Performance** Across all experiments, we observed that our models consistently performed poorly on German, the language for which we had the largest amount of training data. This could be attributed to the LLM’s inherent limitations in handling German, as reflected by its relatively lower performance in zero-shot settings (Llama-3.1-8B-Instruct) for the en-de pair for both BLEU and COMET, compared to en-it and en-zh. Nonetheless, training the model on multilingual speech-only data for the ST task led to an improvement of 3–4 BLEU points. Incorporating multimodal data (i.e., both ST and MT) yielded an additional gain of 2–3 BLEU points, further enhancing performance. For COMET, speech-only and multimodal training does not improve COMET scores over the text topline.

**Overall ST Performance** We observe that the speech projectors (A.1 and A.2) outperform the backbone SeamlessM4T-v2-large for en-zh using BLEU, and in all languages using COMET. Then, multimodal training further increases the BLEU scores, but in some cases, this training slightly hurts the COMET scores. Since this difference in COMET is very small (en-de 0.005; en-it 0.008; en-zh 0.004), and since the BLEU scores increase, we attribute this to some formatting bias that could be happening during adaptation. The Appendix Section C.2 discusses the matter further.

**SQA performance** Overall, we observe that by replacing SQA by *fluent SQA*, we drastically increased our SQA/QA scores. However, results for ASR slightly deteriorate. We hypothesize that this is due to the old SQA task being closer to the ASR task. For the original SQA, the answer is always a direct transcript of a portion of the input text, which as a task has a better synergy with ASR. In intermediate experiments, we observed that adding the original SQA data to the multimodal training was always beneficial for the ASR performance of the model.

**Final Discussion** Overall our results show that it is possible to train text (B) and speech adaptation (A) in parallel, and then to align both via joint instruction tuning (C). By separating the pre-training of both components, we are able to focus hyper-parameter search at the merging stage, using two components that are already competent in their respective modalities. Despite improvements in scores over speech-only models, our best models do not beat the topline working from text for ST, but they do outperform SeamlessM4T-v2-large across all language pairs and metrics. Moreover, for SQA, we highlight that the obtained scores are in some cases very close (en-en, en-it) to the text topline, despite using speech as input context.

## 6 Submitted Model

Table 2 presented the results for our multimodal models. Due to the reasons highlighted in Section 2.1, we only consider two models for submission: the ones which were trained with *fluent SQA*. This is because we believe that these models will suffer the least from domain shift, since they are capable of producing full fluent sentences for SQA.

We observe that both models (A.1 and A.2-based) seem to be equally capable of ASR (18.6

and 18.7). We evaluated language confusion for these two splits, finding that the output produced was English in 98% and 98.5% of the cases respectively. These models differ more in terms of ST performance: they obtained averages of 37.9 and 35.1 BLEU score points respectively. Finally, these models present the following average accuracy for SQA: 83.7% and 85.3%. In summary, while the A.1-based model seem much more capable in ST, the A.2-based model has a very slight edge in ASR and SQA.

Therefore, we decided to select the A.1-based model as our primary submission model. For producing the decoding of the test set, we transform the instructions into our own prompt format, and submit the output of the same model, with the same decoding settings for all splits: greedy decoding using 150 as maximum number of tokens. A brief post-submission discussion is provided in Appendix Section C.2.1.

## 7 Conclusion

In this paper, we presented NLE’s submission to the instruction-following speech processing short track at IWSLT 2025, constrained setting. We developed multimodal models that simultaneously performed ASR, ST, and SQA tasks from English speech input into Chinese, German and Italian. Our approach is simple yet effective: we decoupled training, training a speech projector on speech-to-text tasks, and LoRA adapters on text-to-text tasks. Then, both modules are loaded and the resulting multimodal model was instruction-tuned for a few steps on multilingual and multimodal data to produce the final system submitted for evaluation.

## Acknowledgments

We thank our NLE colleagues Vassilina Nikoulina, for the implementation of the LLM-as-judge scripts using bergen, and Christian Wolf, for kindly verifying some of our German ST/SQA outputs. This work was partially funded by the European Horizon 2022 project UTTER (Unified Transcription and Translation for Extended Reality), under grant agreement No 101070631.

## References

Idris Abdulmumin, Victor Agostinelli, Tanel Alumäe, Antonios Anastasopoulos, Ashwin, Luisa Bentivogli, Ondřej Bojar, Claudia Borg, Fethi Bougares, Roldano Cattoni, Mauro Cettolo, Lizhong Chen, William

- Chen, Raj Dabre, Yannick Estève, Marcello Federico, Marco Gaido, Dávid Javorský, Marek Kasztnik, and 30 others. 2025. Findings of the iwslt 2025 evaluation campaign. In *Proceedings of the 22nd International Conference on Spoken Language Translation (IWSLT 2025)*, Vienna, Austria (in-person and online). Association for Computational Linguistics. To appear.
- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altmenschmidt, Sam Altman, Shyamal Anadkat, and 1 others. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Duarte Miguel Alves, José Pombal, Nuno M Guerreiro, Pedro Henrique Martins, João Alves, Amin Farajian, Ben Peters, Ricardo Rei, Patrick Fernandes, Sweta Agrawal, Pierre Colombo, José G. C. de Souza, and Andre Martins. 2024. [Tower: An open multilingual large language model for translation-related tasks](#). In *First Conference on Language Modeling*.
- Kshitij Ambilduke, Ben Peters, Sonal Sannigrahi, Anil Keshwani, Tsz Kin Lam, Bruno Martins, Marcelly Zanon Boito, and André FT Martins. 2025. From tower to spire: Adding the speech modality to a text-only llm. *arXiv preprint arXiv:2503.10620*.
- Loïc Barrault, Yu-An Chung, Mariano Cora Meglioli, David Dale, Ning Dong, Paul-Ambroise Duquenne, Hady Elsahar, Hongyu Gong, Kevin Heffernan, John Hoffman, and 1 others. 2023. Seamless4t: Massively multilingual & multimodal machine translation. *arXiv preprint arXiv:2308.11596*.
- Alexandre Défossez, Laurent Mazaré, Manu Orsini, Amélie Royer, Patrick Pérez, Hervé Jégou, Edouard Grave, and Neil Zeghidour. 2024. Moshi: a speech-text foundation model for real-time dialogue. *arXiv preprint arXiv:2410.00037*.
- Danny Driess, Fei Xia, Mehdi SM Sajjadi, Corey Lynch, Aakanksha Chowdhery, Brian Ichter, Ayzaan Wahid, Jonathan Tompson, Quan Vuong, Tianhe Yu, and 1 others. 2023. Palm-e: An embodied multimodal language model. *arXiv preprint arXiv:2303.03378*.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, and 1 others. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- 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*.
- Shujie Hu, Long Zhou, Shujie Liu, Sanyuan Chen, Lingwei Meng, Hongkun Hao, Jing Pan, Xunying Liu, Jinyu Li, Sunit Sivasankaran, and 1 others. 2024. Wavllm: Towards robust and adaptive speech large language model. *arXiv preprint arXiv:2404.00656*.
- Rongjie Huang, Mingze Li, Dongchao Yang, Jiaotong Shi, Xuankai Chang, Zhenhui Ye, Yuning Wu, Zhiqing Hong, Jiawei Huang, Jinglin Liu, and 1 others. 2024. Audiogpt: Understanding and generating speech, music, sound, and talking head. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 23802–23804.
- J. Iranzo-Sánchez, J. A. Silvestre-Cerdà, J. Jorge, N. Roselló, A. Giménez, A. Sanchis, J. Civera, and A. Juan. 2020. Europarl-st: A multilingual corpus for speech translation of parliamentary debates. In *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 8229–8233.
- Albert Q Jiang, Alexandre Sablayrolles, Antoine Roux, Arthur Mensch, Blanche Savary, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Emma Bou Hanna, Florian Bressand, and 1 others. 2024. Mixtral of experts. *arXiv preprint arXiv:2401.04088*.
- Hugo Laurençon, Léo Tronchon, Matthieu Cord, and Victor Sanh. 2024. [What matters when building vision-language models?](#) *arXiv preprint. ArXiv:2405.02246*.
- 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.
- Chin-Yew Lin. 2004. [ROUGE: A package for automatic evaluation of summaries](#). In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2023. [Visual Instruction Tuning \(LLaVA\)](#). *arXiv preprint. ArXiv:2304.08485 [cs]*.
- Pedro Henrique Martins, Patrick Fernandes, João Alves, Nuno M Guerreiro, Ricardo Rei, Duarte M Alves, José Pombal, Amin Farajian, Manuel Faysse, Mateusz Klimaszewski, and 1 others. 2024a. Eurollm: Multilingual language models for europe. *arXiv preprint arXiv:2409.16235*.
- Pedro Henrique Martins, Patrick Fernandes, João Alves, Nuno M. Guerreiro, Ricardo Rei, Duarte M. Alves, José Pombal, Amin Farajian, Manuel Faysse, Mateusz Klimaszewski, Pierre Colombo, Barry Haddow, José G. C. de Souza, Alexandra Birch, and André F. T. Martins. 2024b. [Eurollm: Multilingual language models for europe](#). *Preprint, arXiv:2409.16235*.
- Tu Anh Nguyen, Benjamin Muller, Bokai Yu, Marta R Costa-Jussa, Maha Elbayad, Sravya Popuri, Christophe Ropers, Paul-Ambroise Duquenne, Robin Algayres, Ruslan Mavlyutov, and 1 others. 2025. Spirit-lm: Interleaved spoken and written language model. *Transactions of the Association for Computational Linguistics*, 13:30–52.



- Iñigo Pikabea, Iñaki Lacunza, Oriol Pareras, Carlos Escolano, Aitor Gonzalez-Agirre, Javier Hernando, and Marta Villegas. 2025. Breaking language barriers in visual language models via multilingual textual regularization. *arXiv preprint arXiv:2503.22577*.
- 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, Belgium, Brussels. 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, and 1 others. 2024. Scaling speech technology to 1,000+ languages. *Journal of Machine Learning Research*, 25(97):1–52.
- David Rau, Hervé Déjean, Nadezhda Chirkova, Thibault Formal, Shuai Wang, Stéphane Clinchant, and Vasilina Nikoulina. 2024. [BERGEN: A benchmarking library for retrieval-augmented generation](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 7640–7663, Miami, Florida, USA. Association for Computational Linguistics.
- Ricardo Rei, José G. C. de Souza, Duarte Alves, Chrysoula Zerva, Ana C Farinha, Taisiya Glushkova, Alon Lavie, Luisa Coheur, and André F. T. Martins. 2022. [COMET-22: Unbabel-IST 2022 submission for the metrics shared task](#). In *Proceedings of the Seventh Conference on Machine Translation (WMT)*, pages 578–585, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- 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, and 1 others. 2023. Audiopalm: A large language model that can speak and listen. *arXiv preprint arXiv:2306.12925*.
- Elizabeth Salesky, Kareem Darwish, Mohamed Al-Badrashiny, Mona Diab, and Jan Niehues. 2023. [Evaluating multilingual speech translation under realistic conditions with resegmentation and terminology](#). In *Proceedings of the 20th International Conference on Spoken Language Translation (IWSLT 2023)*, pages 62–78, Toronto, Canada (in-person and online). Association for Computational Linguistics.
- Changli Tang, Wenyi Yu, Guangzhi Sun, Xianzhao Chen, Tian Tan, Wei Li, Lu Lu, Zejun Ma, and Chao Zhang. 2023. Salmonn: Towards generic hearing abilities for large language models. *arXiv preprint arXiv:2310.13289*.
- Gemini Team, Rohan Anil, Sebastian Borgeaud, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M Dai, Anja Hauth, Katie Millican, and 1 others. 2023. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*.
- Gemma Team, Aishwarya Kamath, Johan Ferret, Shreya Pathak, Nino Vieillard, Ramona Merhej, Sarah Perrin, Tatiana Matejovicova, Alexandre Ramé, Morgane Rivière, Louis Rouillard, Thomas Mesnard, Geoffrey Cideron, Jean bastien Grill, Sabela Ramos, Edouard Yvinec, Michelle Casbon, Etienne Pot, Ivo Penchev, and 197 others. 2025. [Gemma 3 technical report](#). Preprint, arXiv:2503.19786.
- torch tune maintainers and contributors. 2024. [torch tune: Pytorch’s finetuning library](#).
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shrutu Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, and 49 others. 2023. [Llama 2: Open foundation and fine-tuned chat models](#). Preprint, arXiv:2307.09288.
- Changhan Wang, Anne Wu, and Juan Pino. 2020. [Covost 2: A massively multilingual speech-to-text translation corpus](#). Preprint, arXiv:2007.10310.
- An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, and 1 others. 2024. Qwen2.5 technical report. *arXiv preprint arXiv:2412.15115*.

	en	de	it	zh
# lines	8,026	3,272	783	627

Table 3: Statistics for the SpokenSQuAD SQA/QA test sets. The multilingual version is obtained via SeamlessM4T-v2-large translation of questions and answers, with posterior quality filtering based on COMET scores.

## A Data Preprocessing and Prompts

### A.1 SpokenSQuAD TTS data

In this section we explain the misalignment we found in the SpokenSQuAD train split, as well as our procedure for creating the new TTS split.

**Train Split Misalignment** During the preprocessing of the train split of SpokenSQuAD, we witnessed cases of misalignment: for some given paragraphs, the corresponding audios were shifted by a factor varying between 1 and 3. For instance, the first audio in a given paragraph was incorrectly named as “1”, instead of “0”, shifting all the paragraph’s alignment. We listened to all cases we were able to flag, manually correcting them. However, we believe this hinted to a deeper alignment issue, as the obtained training set seemed to be difficult to learn. We observed that models trained with this training set included were unable to generalize to SpokenSQuAD’s validation and test sets, always producing random Wikipedia sentences when receiving the SpokenSQuAD’s TTS voice as input.

**New TTS Source Audio Generation** We use SeamlessM4T-v2-large to produce new source audio for the SpokenSQuAD training set (34,003 sentences). For each entry in this set, we re-synthesize its SQuAD reference text by randomly sampling one of the 200 speakers present in SeamlessM4T. This results in a more diverse training set, since the original TTS used a single female voice for all sentences. We also generated extra speech data using all different questions present in this training set, producing a second ASR set containing speech for 28,000 questions.

### A.2 SpokenSQuAD Answers Regeneration

SpokenSQuAD answers are direct extracts from the reference text, formatted as lowercase text without punctuation. We discovered this presented a limitation when training our models on SQA. The trained models over-fitted to that format, solving SQA as

a transcription task of the relevant portion of the source audio. While conceptually correct, this approach can result in generalization issues if more than one extract is required to answer the question, as the model never observed such a setting during training, and it will thus have the tendency to transcribe everything between the two points of interest.

We use Llama-3.1-8B-Instruct to regenerate all answers in our multilingual training set. The prompt used for regeneration is presented in Table 4. After regeneration, we remove answers generated in the wrong language using an automatic language identification tool. Statistics are presented in Table 1.

### A.3 Data Statistics

Table 3 presents the statistics for the multilingual test set of SpokenSQuAD.

### A.4 Our Prompts

Table 5 presents our prompt format. We designed our prompts to be very similar, independently of the target task. The language of the question defines the answer’s target language.

## B Additional Results

**CoVoST2 and EuroParlST Results** Table 6 presents results for relevant models on the in-domain test sets from CoVoST2 and EuroParlST. We observe that the multimodal adaptation improves the speech projectors’ ASR and BLEU scores, while slightly decreasing COMET scores.

**ACL 60-60 Dev Results** Table 7 presents ACL 60-60 dev split scores for some of the models presented in the main results table (Table 2).

**SQA/QA BERT Scores** Table 8 presents BERT scores for the multimodal models presented in the main results table (Table 2), computed after the evaluation period and using the same settings from [Abdulmumin et al. \(2025\)](#). We observe that scores for languages other than English over the valid test set are considerably lower than the LLM-as-judge scores in Table 2. Prior to the evaluation period we had evaluated our models using BERT score and xlm-roberta-large, which yielded much higher scores, similar to those obtained in the LLM-as-judge evaluation. Those scores are presented in Table 9.

Context: [REFERENCE TEXT]  
 Question: [QUESTION]  
 Answer: [ANSWER]  
 Instruction: Reformulate the answer to be a natural sounding sentence that answers the question in the correct language. Produce text in the same language of the question and answer. Do not make it too long, or add too much information. Don't add anything else to your answer.

Table 4: Regeneration prompt we gave to Llama-3.1-8B-Instruct to regenerate the answers in the training set of SpokenSQuAD.

User Prompt	
<b>Speech Prefix</b>	Content: <speech>[SPEECH EMBEDDINGS]</speech>\n
<b>Text Prefix</b>	Content: <text>[SPEECH TRANSCRIPTION]</text>\n
<b>ASR</b>	Question: Can you transcribe the Speech content into English text?\n
<b>ST/MT (de)</b>	Question: Können Sie den Inhalt der Rede in den deutschen Text übersetzen?\n
<b>ST/MT (it)</b>	Question: Puoi tradurre il contenuto del discorso in testo italiano?\n
<b>ST/MT (zh)</b>	Question: 你能把演讲内容翻译成中文吗?\n
<b>SQA/QA</b>	Question: [QUESTION]\n
<b>Suffix</b>	Your answer:

Table 5: The user turn prompt template used for training our models. For speech tasks, the user prompt is given by **Speech Prefix+Task+Suffix**, for text tasks, the user prompt is given by **Text Prefix+Task+Suffix**. ST/MT instructions were obtained by translating the instruction “Can you translate the Speech content into [German/Italian/Chinese] text?” to corresponding target languages using SeamlessM4T-v2-large.

	ASR				ST							
	CoVoST2		EuroParlST		CoVoST2				EuroParlST			
	WER	CER	WER	CER	en-de		en-zh		en-de		en-it	
				BLEU	COMET	BLEU	COMET	BLEU	COMET	BLEU	COMET	
A.1 (ASR/ST)	7.84	3.70	11.21	7.30	<b>30.24</b>	<b>0.789</b>	<b>40.14</b>	<b>0.806</b>	<b>25.06</b>	<b>0.840</b>	<b>27.55</b>	0.860
A.2 (ASR/ST/SQA)	<b>7.21</b>	<b>3.44</b>	<b>10.98</b>	<b>7.18</b>	30.68	<b>0.789</b>	40.63	0.807	25.21	0.841	28.36	<b>0.859</b>
A.1 + B (ASR/ST/MT/fluentlySQA/fluentlyQA)	7.59	3.55	11.20	7.20	31.42	0.770	42.19	0.802	27.16	0.809	27.86	0.842
A.2 + B (ASR/ST/MT/fluentlySQA/fluentlyQA)	7.01	3.25	10.82	7.11	31.83	0.772	42.36	0.804	26.70	0.812	27.77	0.848

Table 6: ASR and ST results for the test sets of CoVoST2 and EuroParlST.

Model (fine-tuning data)	ASR (WER)	ST/MT (BLEU)			ST/MT (COMET)		
	en	en-de	en-it	en-zh	en-de	en-it	en-zh
<b>Text-only models (MT/QA)</b>							
Llama-3.1-8B-Instruct (zero-shot)	-	21.27	33.81	44.01	0.732	0.757	0.755
<b>B</b> Text-only LoRA (MT/QA)	-	<b>36.94</b>	<b>49.81</b>	<b>52.33</b>	<b>0.782</b>	<b>0.815</b>	<b>0.822</b>
<b>Speech-only models (ASR/ST/SQA)</b>							
SeamlessM4T-v2-large	25.3	23.77	<b>37.84</b>	28.17	0.669	0.713	0.663
<b>A.1</b> Speech Projector (ASR/ST)	<b>15.9</b>	26.77	36.34	38.65	0.718	<b>0.750</b>	<b>0.753</b>
<b>A.2</b> Speech Projector (ASR/ST/SQA)	<b>15.9</b>	<b>26.85</b>	36.09	<b>38.85</b>	<b>0.720</b>	0.749	<b>0.753</b>
<b>Multimodal models (ASR/ST/SQA)</b>							
<b>A.1 + B</b> (ASR/ST/MT/SQA/QA)	<b>13.9</b>	28.74	41.73	40.72	0.716	<b>0.759</b>	0.756
<b>A.1 + B</b> (ASR/ST/MT/fluentlySQA/fluentlyQA)	14.4	27.20	<b>42.01</b>	41.23	0.712	0.752	0.753
<b>A.2 + B</b> (ASR/ST/MT/SQA/QA)	17.2	<b>29.18</b>	39.14	40.99	0.719	0.755	0.762
<b>A.2 + B</b> (ASR/ST/MT/fluentlySQA/fluentlyQA)	15.5	27.62	33.22	<b>41.74</b>	<b>0.726</b>	0.749	<b>0.764</b>

Table 7: ACL 60-60 dev set results for the different models and backbones used in this work.

	Valid Questions				Invalid Questions			
	en-en	en-de	en-it	en-zh	en-en	en-de	en-it	en-zh
<b>A.1 + B</b> (ASR/ST/MT/ <i>fluent</i> SQA/ <i>fluent</i> QA)	0.975	0.532	0.541	0.665	0.999	0.984	0.990	0.988
<b>A.2 + B</b> (ASR/ST/MT/ <i>fluent</i> SQA/ <i>fluent</i> QA)	0.975	0.536	0.546	0.666	0.999	0.990	0.989	0.991

Table 8: BERT scores for SpokenSQuAD test sets computed using the same settings from the organizers (default model, bert\_score version 0.3.13). Invalid questions corresponds to a version of the test set in which the questions are impossible to answer given the speech context.

	Valid Questions				Invalid Questions			
	en-en	en-de	en-it	en-zh	en-en	en-de	en-it	en-zh
<b>A.1 + B</b> (ASR/ST/MT/ <i>fluent</i> SQA/ <i>fluent</i> QA)	0.857	0.869	0.863	0.899	0.996	0.997	0.996	0.997
<b>A.2 + B</b> (ASR/ST/MT/ <i>fluent</i> SQA/ <i>fluent</i> QA)	0.857	0.869	0.863	0.896	0.995	0.998	0.998	0.998

Table 9: BERT scores for SpokenSQuAD test sets computed using xlm-roberta-large. All other settings are equal to Table 8. Invalid questions corresponds to a version of the test set in which the questions are impossible to answer given the speech context.

**Invalid Questions** In Table 8 we also present BERT scores for an invalid multilingual test set, made of the same reference audio files from the English test split (8,026 examples), but with incorrect questions. We observe that the BERT scores for this invalid split is very high, showcasing that our models are fully capable of respecting the instruction format for incorrect answers.

*eralisierung testen, könnte es mir so vorkommen, als ob wir in diesem Fall die Mädchen schlafen, und ich sehe, dass sie schlafen, und ich sehe, dass sie schlafen, (..) und ich sehe, dass sie schlafen, und ich sehe, dass*”

## B.1 Inference Issues

Through manual inspection of our model outputs, we observed that in a small number of cases, inference degenerates, resulting in repeated words or sentences until the maximum token limit is reached. We experimented with various inference strategies, greedy decoding, top-p, and top-k sampling, as well as different temperature settings, but were unable to identify a configuration that fully eliminated the issue. We hypothesize that a lightweight post-processing model could offer a simple and effective solution to mitigate this problem. Below we give some examples of inference degeneration for Chinese and German.

An example of inference degeneration in Chinese:

“国家语言理解模型从各种知识来源中提取出来，例如，通常通过预训练获得的参数中包含的知识，通常通过预训练获得的知识，通常通过预训练获得的知识，通常通过预训练获得的知识，(..)通常通过预训练”

An example of inference degeneration in German:

“In Zusammenhang mit der semantischen Parsierung, wenn wir nach der kompositionalen Gen-



## C Models Hyperparameters

Table 12 lists the data splits used for each model presented in results Table 2. Table 13 presents the probability sampling employed during training.

### C.1 Speech Projector (A) Hyperparameters

**Architecture** We explored multiple architectures to map speech embeddings from SeamlessM4T-v2-large to Llama-3.1-8B-Instruct. To train the speech projector, speech features extracted from SeamlessM4T-v2-large are input to the projector, and the resulting outputs are passed through a frozen Llama-3.1-8B-Instruct model. The speech projector, initialized with random weights, is trained using cross-entropy loss on an ASR task. Preliminary experiments demonstrated that the Transformer encoder architecture consistently outperformed both the Conformer and Multi-Layer Perceptron architectures of similar parameter sizes. Consequently, we adopt the Transformer encoder architecture for all experiments presented in this work.

**Averaged Features** As mentioned, in preliminary experiments, we experimented using the original output of SeamlessM4T-v2-large, as well as performing average every 2 or 3 frames. We observe that averaging every 3 frames results in models that are considerably faster to train, while maintaining similar performance to the original output.

**Data Ratio** For the ASR task, preliminary experiments revealed that training solely on EuroparlST ASR data resulted in poor generalization, whereas incorporating CoVoST ASR data significantly improved model robustness. For the ST task, we defined the data sampling ratios according to the target language distribution across the CoVoST and EuroparlST datasets.

**Batch Size** ASR and ST tasks use a batch size of 16, while SQA is batched with size 8, due to the longer user prompts.

**Checkpoint Selection** Checkpoints were selected based on development set performance across three or four configurations: ASR-best, ST-best, SQA-best (A.2-only), and an All-best checkpoint combining all tasks. We only present results for ST-best checkpoints, which we found to produce the best scores in ST compared to the other versions, while only marginally decreasing scores

in ASR compared to the ASR-best checkpoint. We do not consider SQA-best checkpoints, as the overall SQA performance of projector-only models is very low regardless of the checkpoint selection method.

**Exclusion of Chinese SQA Data** During the training of the speech projector (A.2), we excluded Chinese SQA data. This was due to parallel observation in B models (text-only), in which we observed that the LLM failed to generate coherent Chinese answers. While later we were able to confirm that the issue did not come from the Chinese split itself, this model was obtained simultaneously to that finding, explaining why the data was not included in this setting.

### C.2 Multimodal Models (C) Hyperparameters

In this section we present some ablation experiments for our multimodal adaptation setup. The experiments are performed by producing variants of the  $A.I+B$  model, which is the model we submitted to the challenge. Table 10 present ACL 60-60 dev split ASR and ST results that are discussed in the next paragraphs.

**Impact of Parameters Count** During our multimodal merging step, we combine text-only LoRA weights with our speech projector, yielding better scores. Since this increase in scores could be simply due the additional parameters, we trained a variant of our model in which the merging step is performed using a randomly initialized LoRA. We observe that our training setup indeed benefits from *any* additional weights during adaptation: the models trained with a randomly initialized LoRA outperform the speech projector backbone (A.1). Adding textual tasks in this setting does not help the system, which we hypothesize is due to the LoRA weights not being pretrained on the textual task. Finally, adapting using a pretrained LoRA model further improves ASR and ST scores for two out of three language directions (en-it and en-zh).

**Impact of Textual Tasks** For the multimodal models presented in the main paper, we adapt pretrained modules leveraging speech and textual tasks. We thus investigated the impact of having aligned speech and textual tasks during this adaptation. We observe that incorporating textual tasks has little impact on ASR performance, while substantially improving ST performance for Italian. The results are less favorable for German and Chi-

	ASR		ST (BLEU)			ST (COMET)		
	WER	CER	en-de	en-it	en-zh	en-de	en-it	en-zh
<b>A.1</b> (ASR/ST)	15.9	7.7	26.77	36.34	38.65	0.718	0.750	0.753
<b>A.1 + random LoRA</b> (ASR/ST/SQA)	17.4	8.2	<b>29.07</b>	37.04	<b>41.47</b>	<b>0.732</b>	<b>0.760</b>	<b>0.761</b>
<b>A.1 + random LoRA</b> (ASR/ST/SQA) + (MT/QA)	<b>16.4</b>	<b>7.7</b>	28.07	<b>38.88</b>	41.24	0.725	0.755	0.760
<b>A.1 + B</b> (ASR/ST)	<b>13.8</b>	<b>6.4</b>	<b>27.91</b>	28.88	<b>41.70</b>	<b>0.728</b>	<b>0.743</b>	<b>0.760</b>
<b>A.1 + B</b> (ASR/ST) + (MT)	14.9	7.2	26.09	<b>32.97</b>	41.10	0.715	0.741	0.755
<b>A.1 + B</b> (ASR/ST/SQA)	<b>14.1</b>	<b>6.5</b>	<b>29.02</b>	30.08	<b>41.85</b>	<b>0.722</b>	0.743	<b>0.763</b>
<b>A.1 + B</b> (ASR/ST/SQA) + (MT/QA)	14.4	<b>6.5</b>	27.20	<b>42.01</b>	41.23	0.712	<b>0.752</b>	0.753
<b>A.1 + B</b> (ASR/ST/SQA) + (MT/QA) <b>No synthetic data</b>	18.8	9.2	28.47	32.97	40.03	0.717	0.741	0.753
<b>A.1 + B</b> (ASR/ST/SQA) + (MT/QA) <b>Only synthetic data</b>	<b>13.9</b>	<b>6.6</b>	<b>29.71</b>	<b>39.80</b>	<b>42.04</b>	<b>0.721</b>	<b>0.751</b>	<b>0.754</b>
<b>A.1 + B</b> (ASR/ST/SQA) + (MT/QA) <b>2K steps</b>	14.3	6.4	27.09	40.53	41.53	0.713	0.753	0.755

Table 10: ACL 60-60 dev set ASR and ST scores for variants of our best model (A.1+B).

ASR/ST/SQA	average WER	average BLEU
0.2 / 0.4 / 0.4	<b>16.03</b>	<b>37.58</b>
0.2 / 0.5 / 0.3	17.72	36.36
0.2 / 0.6 / 0.2	16.24	36.62
0.25 / 0.5 / 0.25	16.58	37.08
0.3 / 0.5 / 0.2	16.23	37.30

Table 11: ACL 60-60 dev and test set average WER and BLEU scores for our best model (A.1+B) by varying the ASR/ST/SQA ratios.

nese: in German, the addition of textual tasks leads to a performance drop, whereas in Chinese, the decline is minimal. Overall, these findings suggest that the textual modality may be particularly beneficial in low-resource settings. Italian, which has the fewest training examples in our dataset, appears to benefit the most from this adaptation.

**Impact of SQA** Examining the results in Table 10, we observe that incorporating the SQA and QA tasks leads to improvements in both ASR and ST performance. We hypothesize that the SQA task enhances the model’s adherence to the prompt by encouraging it to *attend to* the information provided, thereby reducing both task and language confusion.

**Inclusion of Synthetic textual Data** Table 11 presents the results of our investigation into the inclusion of potentially noisy synthetic textual data during training. We observe that excluding this synthetic data (*No synthetic data*) negatively affects both BLEU and WER scores. Conversely, training exclusively with synthetic data (*Only synthetic data*) yields improved performance across all metrics. We attribute this to the fact that the

target text in the non-synthetic data is duplicated from the speech task (i.e. the MT set is built from the ST set), leading to reduced data diversity. Removing this duplicated data introduces greater variability during the adaptation phase, which appears beneficial. Finally, using both types of data leads to improved performance for all languages except German. As discussed in the main paper, we hypothesize that this discrepancy is due to issues in the German training data sourced from EuroParlST and CoVoST2.

**Number of Adaptation Steps** Table 10 presents results for a version of our model trained for twice as long (2K steps). At this point, we observe signs of training saturation: differences in ASR and COMET scores across all metrics are minimal, and BLEU scores drop for both German and Italian. These results suggest that our adaptation step does not require a considerable number of training steps.

**Task Ratios** On our preliminary experiments, we tested different task ratios, selecting the one with best average WER and BLEU scores over both ACL 60-60 dev and test set. Table 11 presents those results.

### C.2.1 Post-submission Discussion

Due to time constraints, many of our ablation studies were conducted after the initial submission. Upon analyzing these results (Table 10), we hypothesize that there may be an issue with the German training data: the more the model is exposed to it during training, the worse the COMET scores become. This hypothesis is supported by our *Only synthetic data* results, which show improved BLEU scores for German when we exclude textual data from EuroParlST and CoVoST2. Additionally, our

	CoVoST			EuroParlST			SpokenSQuAD			
	ASR	ST	MT	ASR	ST	MT	ASR	MT	SQA/QA	fluent SQA/QA
<b>B</b> Text-only LoRA	✗	✗	✓	✗	✗	✓	✗	✓	✓	✗
<b>A.1</b> Speech Projector (ASR/ST)	✓	✓	✗	✓	✓	✗	✓	✗	✗	✗
<b>A.2</b> Speech Projector (ASR/ST/SQA)	✓	✓	✗	✓	✓	✗	✓	✗	✓ (no zh)	✗
<b>A.1 + B</b> Multimodal model (ASR/ST/MT/SQA/QA)	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗
<b>A.1 + B</b> Multimodal model (ASR/ST/MT/fluentSQA/fluentQA)	✓	✓	✓	✓	✓	✓	✓	✓	✗	✓
<b>A.2 + B</b> Multimodal model (ASR/ST/MT/SQA/QA)	✓	✓	✓	✓	✓	✓	✓	✓	✓	✗
<b>A.2 + B</b> Multimodal model (ASR/ST/MT/fluentSQA/fluentQA)	✓	✓	✓	✓	✓	✓	✓	✓	✗	✓

Table 12: List of datasets and splits used for each model presented in Table 2. Statistics for number of examples can be seen in Table 1.

	ASR task ratio	ST/MT			SQA/QA (valid/invalid)					
		task ratio	en-de	en-it	en-zh	task ratio	en-en	en-de	en-it	en-zh
Text-only LoRA	x	0.6	0.4	0.3	0.3	0.4	0.2 / 0.05	0.2 / 0.05	0.2 / 0.05	0.2 / 0.05
A.1	0.4	0.6	0.3	0.4	0.3	x	x	x	x	x
A.2	0.4	0.35	0.3	0.4	0.3	0.25	0.2 / 0.2	0.15 / 0.15	x	0.15 / 0.15
A.1 + B	0.2	0.4	0.4	0.3	0.3	0.4	0.2 / 0.05	0.2 / 0.05	0.2 / 0.05	0.2 / 0.05
A.2 + B	0.2	0.4	0.4	0.3	0.3	0.4	0.2 / 0.05	0.2 / 0.05	0.2 / 0.05	0.2 / 0.05

Table 13: Two-level sampling ratio for each model.

ablations suggest that using textual data selectively, rather than uniformly, may be more effective. In particular, textual supervision appears to be most beneficial for Italian, with more limited gains observed for the other two language directions.

Regarding SQA, we were surprised to find that the evaluation setup provided by the organizers yields scores that differ significantly from those obtained with our own evaluation protocol (see Tables 8 and 9). These discrepancies also extend to the scores we obtain using LLM-as-judge (Table 2). We plan to further investigate the limitations of our current evaluation setup to better understand these inconsistencies.