



Late Fusion and Multi-Level Fission Amplify Cross-Modal Transfer in Text-Speech LMs

Santiago Cuervo, Adel Moumen, Yanis Labrak, Sameer Khurana, Antoine Laurent, Mickael Rouvier, Phil Woodland, Ricard Marxer

► To cite this version:

Santiago Cuervo, Adel Moumen, Yanis Labrak, Sameer Khurana, Antoine Laurent, et al.. Late Fusion and Multi-Level Fission Amplify Cross-Modal Transfer in Text-Speech LMs. 2025. hal-05378288

HAL Id: hal-05378288

<https://hal.science/hal-05378288v1>

Preprint submitted on 23 Nov 2025

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Late Fusion and Multi-Level Fission Amplify Cross-Modal Transfer in Text-Speech LMs

Santiago Cuervo*¹

Adel Moumen*²

Yanis Labrak^{4,6}

Sameer Khurana³

Antoine Laurent⁵

Mickael Rouvier⁴

Phil Woodland²

Ricard Marxer*¹

¹Université de Toulon, Aix-Marseille Université, CNRS, LIS, France

²Department of Engineering, University of Cambridge, UK

³Mitsubishi Electric Research Laboratories (MERL), Cambridge, MA, USA

⁴LIA, Avignon Université, France

⁵LIUM, Le Mans Université, France

⁶Zenidoc, Marseille, France

Abstract

Text-Speech Language Models (TSLMs)—language models trained to jointly process and generate text and speech—are commonly trained through an early modality fusion/fission approach, in which both modalities are fed and predicted from a shared backbone via linear layers. We hypothesize that this approach limits cross-modal transfer by neglecting feature compositionality—specifically, the finer-grained nature of speech representations compared to text—preventing the emergence of a shared feature hierarchy within model layers. In this paper, we argue that this limitation can be addressed through late fusion and fission, with a fission process that accesses both high- and low-level features for speech generation. Our models implementing these principles, SMOLTOLK, rival or surpass state-of-the-art TSLMs trained with orders of magnitude more compute, and achieve significantly improved cross-modal performance relative to early fusion/fission baselines. Representation analyses further suggest that our method enhances the model’s ability to abstract higher-level, more semantic features from speech, and leads to increasingly shared representation spaces across layers.

1 Introduction

Generative auto-regressive language modeling has proven to be a powerful method well beyond the textual domain. Since the paradigm-shifting emergence of Large Language Models (LLMs), a highly prolific line of research has focused on extending this modeling approach to discrete representations (tokens) of other modalities, such as images [Chen et al., 2020], video [Yu et al., 2024a], and audio [Lakhotia et al., 2021, Borsos et al., 2023]. This unified framework has enabled the development of multimodal generative LMs capable of jointly modeling multiple modalities [Aghajanyan et al., 2022, Chameleon Team, 2024, Rubenstein et al., 2023], paving the way for AI systems that learn patterns across heterogeneous data and transfer knowledge between modalities [Aghajanyan et al., 2023].

Given the remarkable capabilities of text LLMs and the relative abundance of high-quality textual datasets compared to other modalities, enabling cross-modal knowledge transfer from text LLMs to

*Main contributor. Correspondence to santiago.cuervo@lis-lab.fr.

lower-resource modalities has become a central goal. A common approach is to fine-tune pretrained text LLMs on multimodal data. However, this raises a fundamental question: *How should a model originally designed to process and generate text adapt to handle non-textual modalities so as to promote cross-modal knowledge transfer?* On the input side, this challenge is known as **modality fusion**—the problem of integrating different modalities into a shared representation space. On the output side, a related but less studied challenge arises: **multimodal fission**², which refers to the process of transforming a shared multimodal representation into modality-specific outputs.

Methods for modality fusion are typically categorized as either late or early fusion. In **late fusion**, non-textual modalities are first encoded by specialized modality-specific encoders—such as Whisper for speech [Yu et al., 2024a] or CLIP for images [Liu et al., 2023]—and optionally aligned using modules such as the Q-Former [Li et al., 2023], before their representations are passed to the text LLM. This approach was common in early multimodal LLMs designed for perception tasks with text-only generation. However, these architectures often rely on continuous representations from bidirectional encoders, which are not suitable for causal language modeling. More recent models for multimodal generation—sometimes referred to as *omnimodal*—have largely adopted **early fusion** instead, where non-textual modalities are fed directly in their native discrete representations [Chen et al., 2020], or as vector-quantized embeddings from self-supervised models not explicitly aligned to the text input space [Chameleon Team, 2024, Nguyen et al., 2024]. These models typically adopt **early fission**, predicting modality-specific outputs linearly from the shared transformer backbone.

Early fusion/fission models benefit from a simpler, unified, and more end-to-end architecture. However, this comes at the cost of burdening the language model backbone with the task of internal modality alignment. In this paper, **we hypothesize that early fusion and fission architectures are fundamentally ill-suited for effective cross-modal knowledge transfer**. Our argument rests on three key observations:

- (i) Cross-modal transfer from text LLMs aims to leverage functions learned during text pre-training to process shared, “text-like” structure in non-textual modalities—such as linguistic content in speech or semantic scene descriptions in images.
- (ii) Textual features (e.g., words or subwords) are inherently higher-level abstractions than features in other modalities. A single text token can denote a complex image object composed of many pixels or patches, or a spoken word be composed of many discrete speech tokens, which are usually phonetic or sub-phonetic in nature.
- (iii) Language models build increasingly compositional features across layers, with deeper layers refining these abstractions into lower-level features optimized for next-token prediction [Valeriani et al., 2023, Cheng et al., 2025].

From these premises, we argue:

- (a) To fulfill (i), non-textual inputs must reach the level of abstraction expected by pretrained text functions—which, per (iii), increases with depth. Given (ii), this requires composing low-level modality-specific features into higher-level representations. However, in early fusion, all modalities pass through the same number of layers despite starting at different abstraction levels, leaving non-textual features under-composed. For instance, early layers that combine words into part-of-speech units in text would first need to build syllabic or word-level structures from speech tokens. This mismatch suggests the need for an input adapter that performs additional composition prior to the text backbone—i.e., late fusion.
- (b) The text LM output layer is tuned to predict next word-like tokens. These predictions contain valuable information for prediction in other modalities—e.g., the next speech tokens depend on the next word, or the next image patches may correlate with it. However, as per (ii), non-textual targets are typically of lower abstraction. To predict them from backbone outputs, a form of decomposition is needed—mirroring how word-level predictions are refined from more abstract representations (iii). This motivates an output adapter that transforms LM outputs into low-level non-textual targets—we call such process **late fission**.
- (c) Because of (ii), text-like features underlying non-textual modalities vary less frequently across time or space than non-textual features. As a result, next-word predictive features at the output of the backbone—optimized to model text dynamics—are likely most informative at sparse intervals (e.g., near word boundaries in speech). For much of the input (e.g.,

²Perhaps a non-standard term, but with precedents in the literature [Landragin, 2007, Marinov et al., 2023].

within a spoken word), non-textual next-token prediction relies instead on lower-level features abstracted earlier in the network. For example, given the speech utterance “we submitted the paper to the”, predicting upcoming speech tokens initially benefits from high-level features (e.g., anticipating “conference”). However, while generating “conference”, accurate prediction depends on fine-grained context—e.g., knowing the current phone ‘f’ is necessary to predict the next phone ‘e’. This suggests the need for a fission mechanism that accesses both low- and high-level features, possibly in an input-dependent manner. We refer to such a mechanism as *multi-level fission*.

In this work, we test these hypotheses in the context of text-speech language models. We focus on speech due to its high mutual information with text—making it especially suitable for cross-modal transfer from text-based LMs—and the prevalence of early fusion/fission approaches in the field [Hassid et al., 2023, Nguyen et al., 2024, Défossez et al., 2024, Zeng et al., 2024]. At the same time, there is strong practical motivation: unimodal speech LMs scale poorly, requiring orders of magnitude more compute to reach comparable linguistic competence [Cuervo and Marxer, 2024]. This inefficiency suggests that simply scaling models and data may not be viable, making cross-modal transfer not only promising but perhaps necessary for achieving high-quality speech generation.

Our main contributions are:

- We propose an architecture featuring late fusion and late fission via simple causal adapters, and multi-level fission through an attention-like mechanism across layers, combined with a residual connection from early to late layers.
- We apply this architecture to the SmolLM family of text LMs [Allal et al., 2025], ranging from 135M to 1.7B parameters. The resulting models, SMOLTOLK, consistently outperform early fusion/fission baselines. Our largest variant, SMOLTOLK-2B, rivals or surpasses state-of-the-art models trained with orders of magnitude more compute.
- We conduct ablations and representational analyses to isolate the impact of each architectural component. Results confirm that all proposed elements contribute to improved cross-modal transfer and performance, and reveal representation dynamics that support our hypotheses.

2 Text-Speech Language Models

TSLMs model the joint probability of text and speech token sequences as

$$P(\mathbf{w} = w_1, \dots, w_n) = \prod_{i=1}^n P(w_i | w_1, \dots, w_{i-1}), \quad (1)$$

where $w_i \in \mathcal{V}_t \cup \mathcal{V}_s$, with \mathcal{V}_t and \mathcal{V}_s denoting text and speech vocabularies. TSLMs are typically decoder-only transformers [Vaswani et al., 2017] optimized to minimize the negative log-likelihood:

$$\mathcal{L}_{LM} = - \sum_{i=1}^n P(w_i | w_1, \dots, w_{i-1}). \quad (2)$$

Tokens are mapped to embeddings via a linear function $E \in \mathbb{R}^{(|\mathcal{V}_t|+|\mathcal{V}_s|) \times d}$, where d is the embedding dimension. The sequence $E(w_1), \dots, E(w_n)$ is processed by a stack of decoder-only transformer layers, producing contextual representations $(\mathbf{c}_1, \dots, \mathbf{c}_n)$, where $\mathbf{c}_i \in \mathbb{R}^d$ and each \mathbf{c}_i depends on $\mathbf{c}_{\leq i}$. A linear projection $U \in \mathbb{R}^{d \times (|\mathcal{V}_t|+|\mathcal{V}_s|)}$ maps these to logits defining $P(w_{i+1} | \mathbf{c}_i)$.

Training. TSLMs are typically trained via *vocabulary expansion and speech fine-tuning* of text LMs. Vocabulary expansion extends the embedding function and output projections over \mathcal{V}_t to include \mathcal{V}_s , while the rest of the LM remains unchanged. Fine-tuning methods vary in data mixture: Rubenstein et al. [2023] train on mixed speech-text tasks (TTS, ASR, speech-to-speech translation), while Chou et al. [2023] use word-level alignments to switch modalities within a sequence. This interleaved text-speech strategy was shown to be crucial for cross-modal transfer, later validated and scaled up by Nguyen et al. [2024] and Zeng et al. [2024], achieving state-of-the-art speech LM performance.

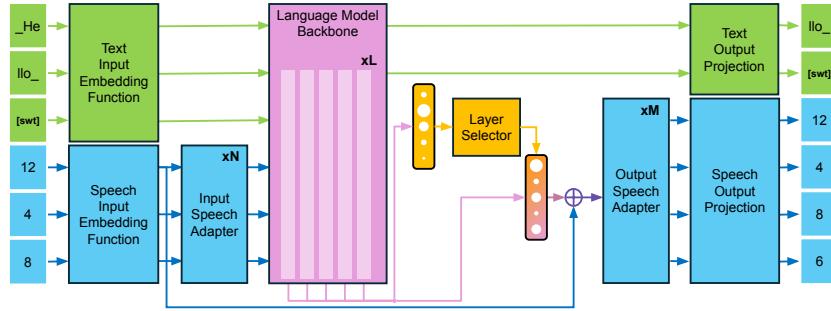


Figure 1: Proposed architecture. The model processes interleaved text-speech sequences. The [swt] token denotes a modality switch. **Late fusion:** Speech inputs (blue) are processed by speech-specific layers before merging with text embeddings (green) in the text LM backbone. **Multi-level fission:** an input speech residual and an average across layers’ representations with input-dependent weights produce multi-layer representations. **Late fission:** These are passed through output speech-specific layers to predict speech tokens. Text tokens are predicted from the final backbone layer.

3 Proposed method

Our proposed method is illustrated in Figure 1, and described in detail below.

3.1 Model

The embedding function E is applied to the input sequence composed of text and speech tokens \mathbf{w} , yielding a sequence of text and speech embeddings $(E(w_1), \dots, E(w_n)) = (\mathbf{z}_1, \dots, \mathbf{z}_n)$.

Late fusion. Contiguous chunks of speech embeddings are processed through an input adapter $A_{\text{in}} : \mathbb{R}^d \rightarrow \mathbb{R}^d$, designed to compose speech token embeddings into higher-level representations suitable as input to the text backbone. Unlike most late fusion architectures, which typically rely on Whisper-like pre-trained bidirectional encoders followed by downsampling and non-linear projections [Yu et al., 2024b, Chu et al., 2023], we employ a simple stack of decoder-only transformer layers applied directly to the speech tokens, maintaining a unified input/output representation. For a contiguous chunk of speech embeddings $(\mathbf{z}_i, \dots, \mathbf{z}_{i+k})$, the input adapter outputs a sequence $(\mathbf{z}'_i, \dots, \mathbf{z}'_{i+k})$.

The output of the input adapter and the text embeddings are fed into the text LM transformer at their respective positions in the input. For instance, given the input sequence $(w_1, w_2, w_3, w_4, w_5)$, where only the third element is a text token, the sequence passed to the transformer layers after the adapter would be $(\mathbf{z}'_1, \mathbf{z}'_2, E(w_3), \mathbf{z}'_4, \mathbf{z}'_5)$. For each transformer layer l , a sequence of contextual representations $(\mathbf{c}_1^{(l)}, \dots, \mathbf{c}_n^{(l)})$ is obtained. As in standard text LMs and TSLMs, the text output logits are computed by applying U to the contextual representations $\mathbf{c}_i^{(L)}$ at the last layer L . The speech output logits are computed as described next.

Multi-level fission. As stated in Section 1, hypothesis (c), we posit that speech language modeling requires the model to alternate between two modes of operation: one for generating speech within a word, and another for initiating a new word. Within a word, the model should rely on low-level representations encoding the current word and its internal speech token, as these determine the next speech token. When starting a new word, the model should instead draw on higher-level features predictive of upcoming words, such as those found in the later layers of the text LM. To support this behavior, we introduce a learnable mechanism that attends to representations from different layers in an input-dependent manner. A linear layer selector $S : \mathbb{R}^d \rightarrow \mathbb{R}^L$ maps a contextual representation \mathbf{c}'_i to a vector of weights $\omega_i = (\omega_i^{(1)}, \dots, \omega_i^{(L)})$, which are softmax-normalized and used to compute a weighted average over layer representations, yielding the contextual multi-level representation $\bar{\mathbf{c}}_i$:

$$\omega_i = \text{Softmax}(S(\mathbf{c}'_i)) \quad , \quad \bar{\mathbf{c}}_i = \sum_{l=1}^L \omega_i^{(l)} \mathbf{c}_i^{(l)} \quad (3)$$

A key question is which representation to use as \mathbf{c}'_i . We found that the last layer representations $\mathbf{c}_i^{(L)}$ are not well suited as they often resulted in the selector collapsing to select a single layer. Rather than searching for the best layer—which would likely vary across different architectures—we use a weighted average of the contextual representations with learned input-independent weights:

$$\mathbf{c}'_i = \sum_{l=1}^L \phi^{(l)} \mathbf{c}_i^{(l)} \quad (4)$$

where $\phi = (\phi^{(1)}, \dots, \phi^{(L)})$, $\phi^{(l)} \in \mathbb{R}$ are learned weights.

To provide information about the current speech token, a residual connection from the speech input embeddings is added to the multi-level contextual representation:

$$\bar{\mathbf{c}}'_i = \bar{\mathbf{c}}_i + \mathbf{z}_i \quad (5)$$

Here, $\bar{\mathbf{c}}'_i$ contains both the layer-pooled features and the current speech token.

Late fission. The output adapter $A_{out} : \mathbb{R}^d \rightarrow \mathbb{R}^d$, a stack of decoder transformer layers, takes $\bar{\mathbf{c}}'_i$ as input and refines it into a representation predictive of upcoming speech tokens, upon which the speech output logits are computed by applying U .

3.2 Training

We train our model to optimize Eq. 2 on a data mixture similar to Chou et al. [2023], including unimodal speech, text, and interleaved text-speech samples. Unlike prior work, we exclude ASR and TTS samples (except in the experiment in Appendix D.4).

Two stage training. Training follows a two-stage process. In the first stage, the text LM backbone is frozen, and only the newly added modules are trained on interleaved text-speech modeling for approximately 3% of the total training iterations. This stage is introduced to mitigate text capability forgetting, as suggested by preliminary experiments. In the second stage, the full model is trained on the complete data mixture for the remaining iterations.

Preventing layer selector collapse. In larger models, we observed that the layer selector S sometimes collapses early in training, attending to a single layer. To mitigate this, an entropy maximization term is added to the loss to encourage diversity in the model layer selection:

$$\mathcal{L} = \mathcal{L}_{LM} + \beta \frac{1}{n} \sum_{i=1}^n \sum_{l=1}^L \omega_i^{(l)} \ln(\omega_i^{(l)}) \quad (6)$$

where β is a hyperparameter that controls the entropy term.

4 Experimental setup

4.1 Models and training

We use SmolLM models [Allal et al., 2025] as text LM backbones, available in three sizes: 135 million, 360 million, and 1.7 billion parameters. We refer to the models resulting from applying our method to the SmolLM backbones as SMOLTOLK-150M, SMOLTOLK-400M and SMOLTOLK-2B.

The input and output adapters consist of transformer layers matching the text backbone architecture. In initial experiments, we observed that using more than two adapter layers led to only marginal or no improvements; thus, we use two layers for all our models. This aligns with findings from Turetzky and Adi [2024] on adapting speech representations to frozen text LMs. Across models, adding the adapters introduces a parameter overhead between 11% and 15%.

For speech tokenization, we follow Hassid et al. [2023], using the same tokenizer, which quantizes HuBERT representations extracted at 25 Hz into a 500-token vocabulary. As is common in speech LM training, adjacent token repetitions are collapsed to a single token.

All LMs are optimized with AdamW [Loshchilov and Hutter, 2019] and use a weight decay of 0.1. Models use a constant learning rate of 3e-4 for text and backbone parameters in the 135M, 360M, and 1.7B baselines, and 1e-4 for SMOLTOLK-2B. SMOLTOLK models apply a trapezoidal schedule for speech-specific parameters in stage 1 (1B tokens, 1000 steps), with a 100-step linear warmup and a maximum learning rate 10 \times that of text parameters, decaying over the final 20% of stage 1. For SMOLTOLK-400M and SMOLTOLK-2B, we set $\beta = 0.01$ in Equation 6. We use a batch size of 1 million tokens with the full 2048-token context. Each batch contains equal proportions of speech, text, and interleaved text-speech. All models are trained for 16 billion tokens, except the largest models, which are trained for up to 32 billion.

For more details on models and training see Appendix A.

4.2 Evaluation

Metrics. For downstream evaluation, we use standard zero-shot speech language modeling metrics. Syntactic knowledge is assessed using the sBLIMP benchmark [Nguyen et al., 2020], which measures the model’s accuracy in selecting a syntactically correct utterance over an incorrect one based on estimated likelihood. Semantics and commonsense reasoning are evaluated using the sStoryCloze and tStoryCloze benchmarks [Hassid et al., 2023], which measure accuracy in selecting the correct continuation of a given context based on predicted likelihood. To measure cross-modal transfer, following Nguyen et al. [2024] and Zeng et al. [2024], we evaluate sStoryCloze and tStoryCloze in four settings: speech context to speech continuation (S), text context to speech continuation (T \rightarrow S), speech context to text continuation (S \rightarrow T), and text context to text continuation (T).

We also report text performance on MMLU [Hendrycks et al., 2021] before (pre) and after (post) speech training to assess whether fine-tuning causes forgetting of text capabilities, as seen in other TSLMs [Nguyen et al., 2024, Défossez et al., 2024]. MMLU is evaluated following the guidelines for the SmoILM models³.

Baselines. We compare SMOLTOLK to models trained with the same text LM backbones and data but using early fusion/fission, referring to these simply as *baselines*. We also compare against state-of-the-art early fusion/fission TSLMs: SpiritLM [Nguyen et al., 2024], Moshi [Défossez et al., 2024], and the 1.5-billion and 9-billion GLM-4-Voice models [Zeng et al., 2024].

4.3 Data

Speech datasets. We train on a collection of public English speech datasets: LibriSpeech [Panayotov et al., 2015], LibriLight [Kahn et al., 2020], SWC [Baumann et al., 2019], Tedlium [Hernandez et al., 2018], People [Galvez et al., 2021], Vox Populi [Wang et al., 2021], and sTinyStories [Cuervo and Marxer, 2024], totaling 160k hours of audio (10.89B speech tokens).

Text datasets. We use a 12-billion-token subset of the SmoILM corpus [Allal et al., 2025]. Unlike Nguyen et al. [2024], we include math and code data, aiming to better preserve text capabilities. Our data distribution matches that used for pre-training SmoILM models⁴.

Text-Speech datasets. We use the forced aligner from Pratap et al. [2024] to obtain word-level time alignments for the LibriHeavy [Kang et al., 2024], sTinyStories, and SWC datasets. Interleaved samples are generated on the fly during batch sampling by randomly switching modalities within the input sequence. Following Nguyen et al. [2024], we randomly select word spans so that each text sequence contains 10–30 words and each speech sequence 5–15 words, balancing the proportion of speech and text tokens in each sample.

For more detail on the datasets see Appendix B.

5 Experiments and results

Table 1 presents the benchmark results. For brevity, we report only the results for the 1.7-billion baseline. Smaller baseline models underperformed relative to the larger one. An expanded results table is provided in Appendix D.5. See Appendix D.3 for cross-modal generation samples.

³<https://huggingface.co/HuggingFaceFW/ablation-model-fineweb-edu#evaluation>

⁴As reported in <https://github.com/huggingface/smollm/blob/main/pre-training/>

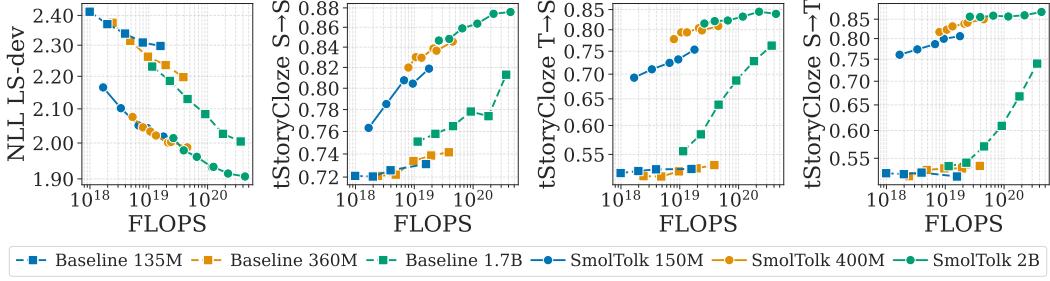


Figure 2: Scaling of the LibriSpeech dev set negative log-likelihood (NLL) and tStoryCloze accuracy across modalities with respect to training compute (in FLOPs).

Our method significantly outperforms the early fusion/fission baselines across all tasks. One might attribute this difference to the parameter increase induced by the added modules. To better account for this factor, Figure 2 shows the scaling behavior of the negative log-likelihood (NLL) on the LibriSpeech dev set and the tStoryCloze benchmark as a function of compute, which helps adjust for model size differences and enables a fairer comparison. The figure demonstrates that, across the entire compute range, our models consistently outperform their respective baselines. Notably, SMOLTOLK-150M surpasses the 360-million baseline and rivals or surpasses the 1.7-billion baseline, proving that factors beyond model size drive the performance difference. Interestingly, in the baselines, cross-modal alignment scales only in the 1.7-billion parameter model. In contrast, our architecture demonstrates this capability even at smaller model sizes, indicating a more efficient approach to cross-modal learning. Scaling behavior for other metrics is presented in Appendix D.6.

Model	Params.	Tokens	BLIMP		tStoryCloze				sStoryCloze				MMLU	
			T	S	T	S	T → S	S → T	T	S	T → S	S → T	T (post/pre)	
<i>Previous TSLMs</i>														
SPIRIT LM [Nguyen et al., 2024]	7B	~175B	73.3	59.7	95.8	90.5	78.6	94.3	74.0	66.3	64.7	71.7	37.7 / 39.0	
Moshi [Défossez et al., 2024]	7.7B	2.1T	—	58.8	—	83.0	—	—	—	60.8	—	—	49.8 / 54.3	
GLM-4-Voice [Zeng et al., 2024] 1.5B	1.5B	1T	—	—	—	77.5	81.4	90.1	—	55.4	58.6	64.0	—	
GLM-4-Voice [Zeng et al., 2024] 9B	9B	1T	—	—	—	83.0	85.0	93.6	—	62.4	63.2	76.3	—	
<i>Ours</i>														
Baseline 1.7B	1.7B	16B	79.9	56.3	92.8	77.5	72.6	67.3	72.5	53.0	57.0	57.6	40.0 / 40.0	
Baseline 1.7B	1.7B	32B	79.8	58.1	92.9	81.3	76.3	74.0	73.5	55.1	59.0	59.2	39.2 / 40.0	
SMOLTOLK-150M	150M	16B	79.4	58.0	88.4	82.0	75.2	81.0	64.1	55.0	58.8	58.4	30.0 / 30.2	
SMOLTOLK-400M	400M	16B	79.8	59.4	91.3	84.6	80.9	85.0	68.4	57.5	62.3	62.1	34.0 / 34.2	
SMOLTOLK-2B	2B	16B	80.2	61.4	92.6	87.5	83.9	86.0	73.2	60.0	64.0	63.4	40.0 / 40.0	
SMOLTOLK-2B	2B	32B	80.2	61.9	92.6	87.6	84.3	87.1	73.6	61.4	64.2	64.2	40.1 / 40.0	

Table 1: Downstream evaluations. The **best model** in each task is shown in bold and underlined. The **second best** is shown in bold. For SPIRIT LM we report the results for the open-weights version. For other models we present the results reported by the authors. For MMLU, the best model is defined as the one exhibiting the least degradation.

Compared to state-of-the-art TSLMs, SMOLTOLK-2B outperforms GLM-4-Voice-1.5B—the closest model in size—while using over 20x less training compute. Notably, even SMOLTOLK-400M surpasses GLM-4-Voice-1.5B on tStoryCloze S and sStoryCloze S and T → S. SMOLTOLK-2B performs comparably to larger models on most tasks, except for S → T, where the gap is wider. It also achieves the best performance on the sBLIMP syntactic task.

Our setup, including the baselines, exhibits less deterioration in text MMLU. We attribute this to our decision to use a text fine-tuning distribution that matches the one used during pre-training.

5.1 Ablation study

Overall, the results indicate that our design choices enhance multimodal performance. To better understand each component contribution, we conduct an ablation study in Table 2 on the medium-size model by systematically removing elements and evaluating their impact. We observe that removing any component degrades performance across most metrics, confirming the importance of our design

Model	tStoryCloze				sStoryCloze			
	T	S	T \rightarrow S	S \rightarrow T	T	S	T \rightarrow S	S \rightarrow T
SMoTALK-400M	91.3	84.6	80.9	85.0	68.4	57.5	62.3	62.1
-Dyn. pooling	90.8	84.0	80.1	83.9	68.2	57.5	60.9	61.6
-Layer pooling	91.4	82.6	77.8	82.0	68.8	57.3	60.1	60.1
-In Adapter	90.5	82.3	70.1	75.0	68.2	55.4	56.2	57.1
-Out Adapter	90.7	80.7	76.1	84.1	68.1	54.9	60.0	60.0
-Adapters	89.9	77.6	58.8	63.7	68.0	52.6	51.5	54.8
-Residual	91.0	83.1	80.7	82.3	68.8	56.6	61.4	60.8
Baseline 360M	90.4	74.1	53.1	53.8	68.4	54.0	52.1	53.1

Table 2: Architecture ablation study. “-” denotes removal. “-Dyn. pooling” uses fixed learned weights instead of dynamic ones from the layer selector, while “-Layer pooling” entirely disables multi-layer pooling, relying only on the last text LM layer.

choices. Eliminating all adapters results in the steepest drop—especially in cross-modal transfer, highlighting the importance of accounting for the compositionality hierarchy. The input adapter seems to be of greater importance for cross-modal transfer than the output adapter, underscoring the importance of late fusion. Multi-level representations and the residual connection also provide consistent gains, demonstrating the benefits of multi-level fission.

5.2 Representation analyses

The results above show that our proposed architecture improves multimodal performance but offer limited insight into the representation-level hypotheses (a), (b), and (c) from Section 1. To probe the mechanisms underlying these gains, we analyze learned representations across ablated model variants, focusing on two aspects: (1) the model’s ability to abstract high-level features, and (2) the alignment of text and speech representation spaces. To assess (1), we follow Valeriani et al. [2023] and use *intrinsic dimensionality* (ID) as a proxy for compositionality and semantic abstraction. To evaluate (2), we compute the top-k principal components of speech and text representations on paired data and measure how much variance is explained when projecting one onto the other. The intuition is that cross-modal transfer is reflected in the degree to which shared subspaces are used to represent semantically equivalent content. We apply these analyses to the medium-sized models and their ablations from Table 2. Experimental details are provided in Appendix C.

Figure 3 (top) shows that different components affect the model’s ability to compose higher-level features. The early fusion/fission baseline yields the lowest intrinsic dimensionality (ID). All models—except the one without multi-level representations—achieve higher ID, indicating that layer pooling is crucial for compositionality. This could be predicted from hypothesis (c): both low- and high-level features are needed for speech generation. Without multi-level representations, the model relies solely on the final layer, which must retain low-level information (see Appendix D.2 for evidence)—relevant for most steps—potentially hindering higher-level abstraction. Late fusion/fission also plays a role: removing input adapters lowers early-layer ID, while removing output adapters reduces ID overall.

Figure 3 (bottom) shows that multi-level fission and late fusion/fission also affect the subspace overlap between modalities. Notably, as with ID, the absence of multi-level representations yields the lowest degree of alignment (excluding the baseline). While we cannot claim a causal link between ID and alignment, the evidence is consistent with hypothesis (a): insufficient compositionality may prevent alignment. The absence of input adapters results in low overlap in earlier layers, supporting the idea that late fusion performs the required composition. Similarly, removing output adapters reduces overlap in later layers, which aligns with hypothesis (b): the model would repurpose later layers—originally optimized for next-word prediction—to generate representations for upcoming speech tokens. In this study, the speech input residual has minimal effect.

What is dynamic layer pooling learning? Hypothesis (c) yields a testable prediction: the model should alternate attention between low-level and next-word predictive representations, with the latter gaining weight near word boundaries. Figure 4 shows the weights assigned by the layer selector

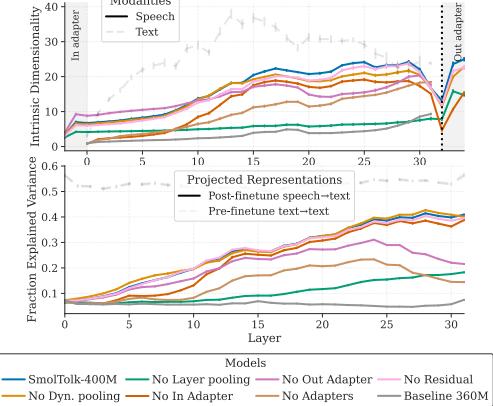


Figure 3: Ablations. *Top*: Intrinsic dimensionality of representations. *Bottom*: Fraction of variance explained by cross-modal projections.

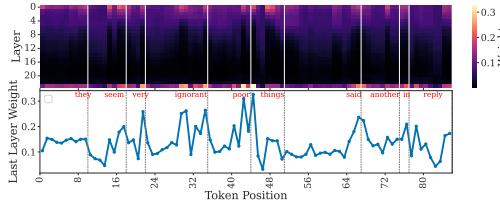


Figure 4: Selector S layer weights across a speech input sequence for SMOLTOLK-2B. Vertical bars indicate word endings.

Model	Precision	Recall	F-1	R-val
SCPC [Bhati et al., 2022]	30.3	20.3	24.5	40.5
SMOLTOLK-150M	50.5	46.7	48.4	56.8
SMOLTOLK-400M	37.6	37.2	37.4	46.8
SMOLTOLK-2B	33.3	29.4	31.2	43.4

Table 3: Word-segmentation scores on the TIMIT test set using dynamic-pooling last-layer weights as word-boundary predictors.

S across layers (top) and for the final layer—specialized in next-token prediction—(bottom), over a speech input with word boundaries. Visual inspection suggests a pattern consistent with our hypothesis: the model shifts attention between low- and high-level features, with spikes in final-layer weights often aligning with word boundaries. To quantify this, we apply a peak detector to the final-layer weights and treat peaks as predicted word boundaries (see Appendix C for details). Table 3 compares this segmentation method to SCPC [Bhati et al., 2022], a state-of-the-art unsupervised speech segmentation model, on the TIMIT test split [Garofolo et al., 1993]. Our approach outperforms SCPC, supporting hypothesis (c). Note that we do not claim our segmentation method is unsupervised, as training on interleaved text-speech data implicitly provides word boundary cues. Interestingly, segmentation quality decreases with model size—we hypothesize that larger models rely more on output adapters and less on backbone features. Additional examples are given in Appendix D.7.

6 Related work

Several prior works adopt late fusion architectures to improve modality alignment. In the speech domain, relevant examples include Yu et al. [2024b], Chu et al. [2023], Tang et al. [2024]. A more recent wave of TSLMs has focused on decoupling speech generation from the text LM backbone via additional speech-specific output modules [Xu et al., 2025] or modality-specific forward passes [Xie and Wu, 2024]—an idea aligned with our late fission approach. Unlike these models, we use simpler causal adapters that support a unified input-output representation (discrete speech tokens) naturally suited to auto-regressive language modeling, and we do not rely on paired text conditioning for speech generation. In addition, we introduce multi-level fission—a novel contribution to the best of our knowledge—and provide new insights from representation analyses into the role and impact of modality fusion and fission mechanisms.

In terms of training and evaluation methodology, our work is most closely aligned with Nguyen et al. [2024] and Zeng et al. [2024], with whom we compare. Compared to them, we use smaller, less diverse open speech datasets and significantly less computational resources, while achieving comparable—or in some cases, superior—performance.

7 Limitations

Scale. Our models are small by current standards, and the observed phenomena may differ at larger scales. Figure 2 shows that cross-modal transfer improves in larger models with early fusion/fission—suggesting that scale might help close the gap with our architecture. However, as shown in Appendix D.1 (Figure 5), even when compared to models at the scale of our largest baseline and SPIRIT LM (7B), our model exhibits qualitatively different representation-space behavior, consistent with small-scale observations—namely, higher and increasing cross-modal alignment. Based on this and the arguments in Section 1, we expect our contributions to remain beneficial at larger scales.

Practical usage of SMOLTOLK. Since our model is not instruction fine-tuned, its direct applicability to speech LM tasks of interest, such as spoken conversational agents, is limited. Future work could explore instruction tuning and broader evaluations. Additionally, using HuBERT-based tokens limits our model expressivity [Nguyen et al., 2023]—one of speech LMs’ key advantages over cascaded approaches combining text LMs with ASR and TTS. Future work could apply our architecture to model more expressive speech representations.

8 Conclusions

We argued that early fusion/fission architectures pose limitations for cross-modal transfer in multi-modal LLMs. Focusing on text-speech language models, we introduced an alternative architecture based on late fusion/fission and a fission mechanism sensitive to both low- and high-level features. Our SMOLTOLK models, which embody these ideas, achieved state-of-the-art-level multimodal performance using significantly less data and computational resources than prior work. Representation analyses showed that these components enhance abstraction and improve alignment between speech and text representations. Overall, our results highlight the importance of architectures that respect the compositional hierarchy of multimodal features—an insight that may be key to efficiently adapting text LMs to other perceptual modalities.

Author contributions

Santiago proposed and led the project, defining the hypotheses, model, and experimental design, and developed most of the training, modeling, and analysis codebase. Adel played a key role in data preprocessing and synthesis, text and cross-modal evaluations, and model inference, while also contributing insights on model design. He maintained the codebase with Santiago’s assistance. Ricard provided supervision throughout and contributed significantly to the data pipeline, experimental codebase, and analysis. Yanis assisted with experiment implementation, execution, and discussions. Phil contributed to the current manuscript, discussions, and to Adel’s supervision. Other authors contributed to early-stage discussions and supervision during the JSALT 2024 workshop.

Acknowledgements

Santiago and Ricard are grateful to the French National Research Agency for their support through the ANR-20-CE23-0012-01 (MIM) grant, as well as to the Institute of Convergence ILCB, which is supported by grants from France 2030 (ANR-16-CONV-0002) and the Excellence Initiative of Aix-Marseille University (A*MIDEX). Adel is funded by the Cambridge Trust and is grateful to his former supervisor in Avignon, professor Yannick Estève, for involving him in this project and providing sponsorship. Other team members have also received funding from the European Union’s Horizon 2020 research and innovation program under the Marie Skłodowska-Curie grant agreement No. 101007666.

This work was supported by HPC GENCI-IDRIS under grants AD011014044R1, A0161014876, AD011015344, A0161015099, AD011013061R3, and A0161014871.

Finally, we extend our gratitude to the JSALT 2024 workshop and the Center for Language and Speech Processing at Johns Hopkins University for generously hosting the event, fostering the emergence of fruitful collaborations and friendships in a vibrant and stimulating environment.

References

- Armen Aghajanyan, Bernie Huang, Candace Ross, Vladimir Karpukhin, Hu Xu, Naman Goyal, Dmytro Okhonko, Mandar Joshi, Gargi Ghosh, Mike Lewis, and Luke Zettlemoyer. CM3: A Causal Masked Multimodal Model of the Internet. *CoRR*, abs/2201.07520, 2022. URL <https://arxiv.org/abs/2201.07520>.
- Armen Aghajanyan, Lili Yu, Alexis Conneau, Wei-Ning Hsu, Karen Hambardzumyan, Susan Zhang, Stephen Roller, Naman Goyal, Omer Levy, and Luke Zettlemoyer. Scaling laws for generative mixed-modal language models. In *Proceedings of the 40th International Conference on Machine Learning*, ICML’23. JMLR.org, 2023.
- Loubna Ben Allal, Anton Lozhkov, Elie Bakouch, Gabriel Martín Blázquez, Guilherme Penedo, Lewis Tunstall, Andrés Marafioti, Hynek Kydliček, Agustín Piqueres Lajarín, Vaibhav Srivastav, Joshua Lochner, Caleb Fahlgren, Xuan-Son Nguyen, Clémentine Fourrier, Ben Burtenshaw, Hugo Larcher, Haojun Zhao, Cyril Zakka, Mathieu Morlon, Colin Raffel, Leandro von Werra, and Thomas Wolf. SmolLM2: When Smol Goes Big – Data-Centric Training of a Small Language Model, 2025. URL <https://arxiv.org/abs/2502.02737>.

Alan Baade, Puyuan Peng, and David Harwath. SyllableLM: Learning coarse semantic units for speech language models. In *The Thirteenth International Conference on Learning Representations*, 2025. URL <https://openreview.net/forum?id=dGSOn7sdWg>.

Timo Baumann, Arne Köhn, and Felix Hennig. The spoken wikipedia corpus collection: Harvesting, alignment and an application to hyperlistening. *Lang. Resour. Eval.*, 53(2):303–329, jun 2019. ISSN 1574-020X. doi: 10.1007/s10579-017-9410-y. URL <https://doi.org/10.1007/s10579-017-9410-y>.

Saurabhchand Bhati, Jesús Villalba, Piotr Źelasko, Laureano Moro-Velazquez, and Najim Dehak. Unsupervised speech segmentation and variable rate representation learning using segmental contrastive predictive coding. *IEEE/ACM Trans. Audio, Speech and Lang. Proc.*, 30:2002–2014, June 2022. ISSN 2329-9290. doi: 10.1109/TASLP.2022.3180684. URL <https://doi.org/10.1109/TASLP.2022.3180684>.

Zalán Borsos, Raphaël Marinier, Damien Vincent, Eugene Kharitonov, Olivier Pietquin, Matt Sharifi, Dominik Roblek, Olivier Teboul, David Grangier, Marco Tagliasacchi, and Neil Zeghidour. Audiolm: A language modeling approach to audio generation. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 31:2523–2533, 2023. doi: 10.1109/TASLP.2023.3288409.

Chameleon Team. Chameleon: Mixed-Modal Early-Fusion Foundation Models, 2024. URL <https://arxiv.org/abs/2405.09818>.

Mark Chen, Alec Radford, Rewon Child, Jeffrey Wu, Heewoo Jun, David Luan, and Ilya Sutskever. Generative pretraining from pixels. In Hal Daumé III and Aarti Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 1691–1703. PMLR, 13–18 Jul 2020. URL <https://proceedings.mlr.press/v119/chen20s.html>.

Emily Cheng, Diego Doimo, Corentin Kervadec, Iuri Macocco, Jade Yu, Alessandro Laio, and Marco Baroni. Emergence of a High-Dimensional Abstraction Phase in Language Transformers, 2025. URL <https://arxiv.org/abs/2405.15471>.

Ju-Chieh Chou, Chung-Ming Chien, Wei-Ning Hsu, Karen Livescu, Arun Babu, Alexis Conneau, Alexei Baevski, and Michael Auli. "toward joint language modeling for speech units and text". In Houda Bouamor, Juan Pino, and Kalika Bali, editors, *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 6582–6593, Singapore, December 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.findings-emnlp.438. URL <https://aclanthology.org/2023.findings-emnlp.438/>.

Yunfei Chu, Jin Xu, Xiaohuan Zhou, Qian Yang, Shiliang Zhang, Zhijie Yan, Chang Zhou, and Jingren Zhou. Qwen-audio: Advancing universal audio understanding via unified large-scale audio-language models, 2023. URL <https://arxiv.org/abs/2311.07919>.

Santiago Cuervo and Ricard Marxer. "scaling properties of speech language models". In Yaser Al-Onaizan, Mohit Bansal, and Yun-Nung Chen, editors, *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 351–361, Miami, Florida, USA, November 2024. Association for Computational Linguistics. doi: 10.18653/v1/2024.emnlp-main.21. URL <https://aclanthology.org/2024.emnlp-main.21/>.

Tri Dao. Flashattention-2: Faster attention with better parallelism and work partitioning. In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=mZn2Xyh9Ec>.

Francesco Denti, Diego Doimo, Alessandro Laio, and Antonietta Mira. The generalized ratios intrinsic dimension estimator. *Scientific Reports*, 12(1):20005, November 2022. ISSN 2045-2322. doi: 10.1038/s41598-022-20991-1. URL <https://doi.org/10.1038/s41598-022-20991-1>.

Alexandre Défossez, Laurent Mazaré, Manu Orsini, Amélie Royer, Patrick Pérez, Hervé Jégou, Edouard Grave, and Neil Zeghidour. Moshi: a speech-text foundation model for real-time dialogue, 2024. URL <https://arxiv.org/abs/2410.00037>.

Daniel Galvez, Greg Diamos, Juan Manuel Ciro Torres, Juan Felipe Cerón, Keith Achorn, Anjali Gopi, David Kanter, Max Lam, Mark Mazumder, and Vijay Janapa Reddi. The people’s speech: A large-scale diverse english speech recognition dataset for commercial usage. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 1)*, 2021. URL <https://openreview.net/forum?id=R8CwidgJ0yT>.

John S. Garofolo, Lori F. Lamel, William M. Fisher, Jonathan G. Fiscus, David S. Pallett, Nancy L. Dahlgren, and Victor Zue. TIMIT acoustic-phonetic continuous speech corpus, 1993. URL <https://catalog.ldc.upenn.edu/LDC93S1>. Catalog Number: LDC93S1, Philadelphia, PA, USA.

Aldo Glielmo, Iuri Macocco, Diego Doimo, Matteo Carli, Claudio Zeni, Romina Wild, Maria d’Errico, Alex Rodriguez, and Alessandro Laio. DADAPy: Distance-based analysis of data-manifolds in python. *Patterns*, 3(10):100589, 2022. doi: 10.1016/j.patter.2022.100589. URL <https://doi.org/10.1016/j.patter.2022.100589>.

Michael Hassid, Tal Remez, Tu Anh Nguyen, Itai Gat, Alexis Conneau, Felix Kreuk, Jade Copet, Alexandre Défossez, Gabriel Synnaeve, Emmanuel Dupoux, Roy Schwartz, and Yossi Adi. Textually pretrained speech language models. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. URL <https://openreview.net/forum?id=UlHueVjAKr>.

Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. Measuring massive multitask language understanding. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=d7KBjmI3GmQ>.

François Hernandez, Vincent Nguyen, Sahar Ghannay, Natalia Tomashenko, and Yannick Estève. TED-LJUM 3: Twice as much data and corpus repartition for experiments on speaker adaptation. In Alexey Karpov, Oliver Jokisch, and Rodmonga Potapova, editors, *Speech and Computer*, pages 198–208, Cham, 2018. Springer International Publishing. ISBN 978-3-319-99579-3.

J. Kahn, M. Rivière, W. Zheng, E. Kharitonov, Q. Xu, P.E. Mazaré, J. Karadayi, V. Liptchinsky, R. Collobert, C. Fuegen, T. Likhomanenko, G. Synnaeve, A. Joulin, A. Mohamed, and E. Dupoux. Libri-light: A benchmark for asr with limited or no supervision. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 7669–7673, 2020. doi: 10.1109/ICASSP40776.2020.9052942.

Wei Kang, Xiaoyu Yang, Zengwei Yao, Fangjun Kuang, Yifan Yang, Liyong Guo, Long Lin, and Daniel Povey. Libriheavy: A 50,000 hours asr corpus with punctuation casing and context. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 10991–10995, 2024. doi: 10.1109/ICASSP48485.2024.10447120.

Felix Kreuk, Joseph Keshet, and Yossi Adi. Self-supervised contrastive learning for unsupervised phoneme segmentation. In *Interspeech*, pages 3705–3709, 2020. URL https://www.isca-archive.org/interspeech_2020/kreuk20_interspeech.pdf.

Kushal Lakhotia, Eugene Kharitonov, Wei-Ning Hsu, Yossi Adi, Adam Polyak, Benjamin Bolte, Tu-Anh Nguyen, Jade Copet, Alexei Baevski, Abdelrahman Mohamed, and Emmanuel Dupoux. On generative spoken language modeling from raw audio. *Transactions of the Association for Computational Linguistics*, 9:1336–1354, 2021. doi: 10.1162/tacl_a_00430. URL <https://aclanthology.org/2021.tacl-1.79>.

Frédéric Landragin. Physical, semantic and pragmatic levels for multimodal fusion and fission. In *Proceedings of the Seventh International Workshop on Computational Semantics (IWCS-7)*, pages 346–350. Universitätsverlag Tilburg, January 2007. URL <https://shs.hal.science/halshs-00138503>.

Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. BLIP-2: bootstrapping language-image pre-training with frozen image encoders and large language models. In *Proceedings of the 40th International Conference on Machine Learning*, ICML’23. JMLR.org, 2023.

Guan-Ting Lin, Prashanth Gurunath Shivakumar, Aditya Gourav, Yile Gu, Ankur Gandhe, Hung yi Lee, and Ivan Bulyko. Align-SLM: Textless spoken language models with reinforcement learning from ai feedback, 2024. URL <https://arxiv.org/abs/2411.01834>.

Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. In A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine, editors, *Advances in Neural Information Processing Systems*, volume 36, pages 34892–34916. Curran Associates, Inc., 2023. URL https://proceedings.neurips.cc/paper_files/paper/2023/file/6dcf277ea32ce3288914faf369fe6de0-Paper-Conference.pdf.

Yan Liu, Renren Jin, Ling Shi, Zheng Yao, and Deyi Xiong. Finemath: A fine-grained mathematical evaluation benchmark for chinese large language models, 2024. URL <https://arxiv.org/abs/2403.07747>.

Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=Bkg6RiCqY7>.

Gallil Maimon, Avishai Elmakies, and Yossi Adi. Slamming: Training a speech language model on one gpu in a day, 2025. URL <https://arxiv.org/abs/2502.15814>.

Zdravko Marinov, Simon Reiß, David Kersting, Jens Kleesiek, and Rainer Stiefelhagen. Mirror u-net: Marrying multimodal fission with multi-task learning for semantic segmentation in medical imaging. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) Workshops*, pages 2283–2293, October 2023.

Tu Anh Nguyen, Maureen de Seyssel, Patricia Rozé, Morgane Rivière, Evgeny Kharitonov, Alexei Baevski, Ewan Dunbar, and Emmanuel Dupoux. The zero resource speech benchmark 2021: Metrics and baselines for unsupervised spoken language modeling. *CoRR*, abs/2011.11588, 2020. URL <https://arxiv.org/abs/2011.11588>.

Tu Anh Nguyen, Wei-Ning Hsu, Antony D’Avirro, Bowen Shi, Itai Gat, Maryam Fazel-Zarani, Tal Remez, Jade Copet, Gabriel Synnaeve, Michael Hassid, Felix Kreuk, Yossi Adi, and Emmanuel Dupoux. Expresso: A benchmark and analysis of discrete expressive speech resynthesis. In *INTERSPEECH 2023*, pages 4823–4827, 2023. doi: 10.21437/Interspeech.2023-1905.

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. Spirit LM: Interleaved spoken and written language model, 2024. URL <https://arxiv.org/abs/2402.05755>.

Vassil Panayotov, Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur. Librispeech: An ASR corpus based on public domain audio books. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5206–5210, 2015.

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. Scaling speech technology to 1,000+ languages. *Journal of Machine Learning Research*, 25(97):1–52, 2024. URL <http://jmlr.org/papers/v25/23-1318.html>.

Alec Radford, Jong Wook Kim, Tao Xu, Greg Brockman, Christine McLeavey, and Ilya Sutskever. Robust speech recognition via large-scale weak supervision, 2022. URL <https://arxiv.org/abs/2212.04356>.

Okko Johannes Räsänen, Unto Kalervo Laine, and Toomas Altosaar. An improved speech segmentation quality measure: The r-value. In *Tenth Annual Conference of the International Speech Communication Association (Interspeech 2009)*, pages 1851–1854, 2009. URL http://www.isca-speech.org/archive/interspeech_2009/i09_1851.html. Funded by EU FP6 FET project ACORNS (FP6-034362).

Paul K. Rubenstein, Chulayuth Asawaroengchai, Duc Dung Nguyen, Ankur Bapna, Zalán Borsos, Félix de Chaumont Quirky, Peter Chen, Dalia El Badawy, Wei Han, Eugene Kharitonov, Hannah Muckenheim, Dirk Padfield, James Qin, Danny Rozenberg, Tara Sainath, Johan Schalkwyk, Matt Sharifi, Michelle Tadmor Ramanovich, Marco Tagliasacchi, Alexandru Tudor, Mihajlo Velimirović, Damien Vincent, Jiahui Yu, Yongqiang Wang, Vicky Zayats, Neil Zeghidour, Yu Zhang, Zhishuai

Zhang, Lukas Zilka, and Christian Frank. AudioPaLM: A Large Language Model That Can Speak and Listen, 2023. URL <https://arxiv.org/abs/2306.12925>.

Changli Tang, Wenyi Yu, Guangzhi Sun, Xianzhao Chen, Tian Tan, Wei Li, Lu Lu, Zejun MA, and Chao Zhang. SALMONN: Towards generic hearing abilities for large language models. In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=14rn7HpKVk>.

Arnon Turetzky and Yossi Adi. Last: Language model aware speech tokenization, 2024. URL <https://arxiv.org/abs/2409.03701>.

Lucrezia Valeriani, Diego Doimo, Francesca Cuturello, Alessandro Laio, Alessio ansuini, and Alberto Cazzaniga. The geometry of hidden representations of large transformer models. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. URL <https://openreview.net/forum?id=cCYvakU5Ek>.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is All you Need. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30, 2017. URL https://proceedings.neurips.cc/paper_files/paper/2017/file/3f5ee243547dee91fb053c1c4a845aa-Paper.pdf.

Changhan Wang, Morgane Riviere, Ann Lee, Anne Wu, Chaitanya Talnikar, Daniel Haziza, Mary Williamson, Juan Pino, and Emmanuel Dupoux. VoxPopuli: A large-scale multilingual speech corpus for representation learning, semi-supervised learning and interpretation. In Chengqing Zong, Fei Xia, Wenjie Li, and Roberto Navigli, editors, *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 993–1003, Online, August 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.acl-long.80. URL <https://aclanthology.org/2021.acl-long.80>.

Zhifei Xie and Changqiao Wu. Mini-omni: Language models can hear, talk while thinking in streaming, 2024. URL <https://arxiv.org/abs/2408.16725>.

Jin Xu, Zhifang Guo, Jinzheng He, Hangrui Hu, Ting He, Shuai Bai, Keqin Chen, Jialin Wang, Yang Fan, Kai Dang, Bin Zhang, Xiong Wang, Yunfei Chu, and Junyang Lin. Qwen2.5-omni technical report, 2025. URL <https://arxiv.org/abs/2503.20215>.

Lijun Yu, Jose Lezama, Nitesh Bharadwaj Gundavarapu, Luca Versari, Kihyuk Sohn, David Minnen, Yong Cheng, Agrim Gupta, Xiuye Gu, Alexander G Hauptmann, Boqing Gong, Ming-Hsuan Yang, Irfan Essa, David A Ross, and Lu Jiang. Language model beats diffusion - tokenizer is key to visual generation. In *The Twelfth International Conference on Learning Representations*, 2024a. URL <https://openreview.net/forum?id=gzqrANCF4g>.

Wenyi Yu, Changli Tang, Guangzhi Sun, Xianzhao Chen, Tian Tan, Wei Li, Lu Lu, Zejun Ma, and Chao Zhang. Connecting speech encoder and large language model for asr. In *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 12637–12641, 2024b. doi: 10.1109/ICASSP48485.2024.10445874.

Aohan Zeng, Zhengxiao Du, Mingdao Liu, Lei Zhang, Shengmin Jiang, Yuxiao Dong, and Jie Tang. Scaling Speech-Text Pre-training with Synthetic Interleaved Data, 2024. URL <https://arxiv.org/abs/2411.17607>.

A Models details

Table 4 describes the SmolLM [Allal et al., 2025] backbones architectural hyperparameters. All SMOLTOlk models use two layer input and output adapters with the same architecture as the backbone layers. Table 5 describes the resulting models after vocabulary expansion.

As mentioned in Section 3.1, we use a linear layer with bias as the dynamic layer selector S . We experimented with simple non-linear MLP selectors; however, these were prone to collapse and

Model	Num. Layers	Num. Heads	Num. KV Heads	Emb. Dim.	Hidden Dim.
SmolLM-135M	30	9	3	576	1536
SmolLM-360M	32	15	5	960	2560
SmolLM-1.7B	24	32	32	2048	8192

Table 4: Backbones architectural hyperparameters.

Model	Backbone	Num. Layers	Text params	Speech Params
Baseline-135M	SmolLM-135M	30	135M	0.29M
Baseline-360M	SmolLM-360M	32	360M	0.48M
Baseline-1.7B	SmolLM-1.7B	24	1.7B	1M
SMOLTOLK-150M	SmolLM-135M	34	135M	15M
SMOLTOLK-400M	SmolLM-360M	36	360M	40M
SMOLTOLK-2B	SmolLM-1.7B	28	1.7B	270M

Table 5: Models description.

resulted in worse overall performance. That said, a more carefully designed non-linear selector could potentially perform better. We also explored alternative ways to define the contextual representation c'_i (Equation 3) used as input for the layer selector. Instead of a learned weighted average, we tried concatenating low-dimensional linear projections from each layer’s representations, but this performed more poorly.

Regarding training, we tuned the learning rate for each model, including baselines, to be as high as possible without causing instabilities or increasing text data validation loss, which we considered a sign of text capability forgetting. We also experimented with learning rate schedules for text/backbone parameters, but a constant rate performed better.

For our experiments, we used NVIDIA H100 GPU nodes, each featuring four NVIDIA H100 80GB SXM5 GPUs, dual Intel Sapphire Rapids 48-core processors, 512GB of RAM, and four NVIDIA ConnectX-7 400Gb/s InfiniBand network adapters. SMOLTOLK-2B used a per-GPU batch size of 8 million tokens and was trained on 64 GPUs across 16 nodes, taking approximately 10 hours to process 32 billion tokens. All models were trained using bf16 mixed precision with FlashAttention-2 [Dao, 2024], packed sequences, and PyTorch compile.

B Datasets details

Table 6 presents the size of the datasets used for each modality and their respective sampling ratios. In each batch, each modality was sampled with equal probability. We experimented with reducing the proportion of text samples per batch to as low as 10%, as their primary purpose is merely to preserve text capabilities. However, we observed that this significantly impacted text downstream performance. The sampling ratios of text datasets were set according to the pre-training setup of SmolLM models, as specified in <https://github.com/huggingface/smollm/blob/main/pre-training/>. Speech datasets were sampled based on their size—that is, uniformly across tokens.

For interleaved text-speech datasets, we upsampled the Spoken Wikipedia Corpora (SWC), as it is a knowledge-dense dataset, and we hypothesized that this would enhance downstream performance on knowledge-based tasks. Additionally, we expected it to better align with the distribution of article-style samples in the text pre-training data, potentially facilitating cross-modal transfer. However, we did not explicitly evaluate these capabilities in this work. The sampling ratios for LibriHeavy and sTinyStories were defined in proportion to their size.

Modality	Dataset	Tokens		Sampling ratio
		Text	Speech	
Text	FineWeb-Edu [Allal et al., 2025]	4B	—	0.7
	Cosmopedia-v2 [Allal et al., 2025]	4B	—	0.15
	Python-Edu [Allal et al., 2025]	2B	—	0.08
	FineMath [Liu et al., 2024]	2B	—	0.06
Speech	LibriSpeech [Panayotov et al., 2015]	—	67M (960 hours)	—
	LibriLight [Kahn et al., 2020]	—	3.7B (53k hours)	—
	SWC [Baumann et al., 2019]	—	32M (1k hours)	—
	Tedlium [Hernandez et al., 2018]	—	0.1B (1.6k hours)	—
	People [Galvez et al., 2021]	—	0.5B (7k hours)	—
	Vox Populi [Wang et al., 2021]	—	1.6B (24k hours)	—
Interleaved text-speech	sTinyStories [Cuervo and Marxer, 2024]	—	4.8B (72k hours)	—
	LibriHeavy [Kang et al., 2024]	313M	3.1B (50k hours)	0.37
	SWC [Baumann et al., 2019]	800M	4.8B (72k hours)	0.53
		3.6M	26M (800 hours)	0.1

Table 6: Datasets statistics. Speech datasets were sampled according to their size.

C Representation analyses setup

C.1 Intrinsic dimensionality and subspace overlap

The intrinsic dimensionality and subspace overlap are estimated using five batches of 10k representations each, totaling 50k samples. We compute each metric per batch and report average and standard deviations across batches. To obtain each representation, we extract random subsequences of 20 words and use the final representation in the sequence as the sequence representation. For the intrinsic dimensionality text samples are randomly drawn from FineWeb-Edu, while speech samples are taken from the full set of speech datasets. For the subspace overlap we use paired samples from the sTinyStories dataset.

Since transformer architectures exhibit large activation outliers, we truncate feature elements (i.e., individual activations) that exceed the 95th percentile across the entire 50k sample set.

Intrinsic Dimensionality. To estimate the intrinsic dimensionality we use the *Generalized Ratios Intrinsic Dimension Estimator* (GRIDE) [Denti et al., 2022] implementation in dadapy [Glielmo et al., 2022] and follow the procedure described by Cheng et al. [2025].

Cross-modal subspace overlap. Let $\mathbf{X} \in \mathbb{R}^{n \times d_X}$ be the *target*-modality embeddings (e.g. text) and $\mathbf{Y} \in \mathbb{R}^{n \times d_Y}$ the *source*-modality embeddings (e.g. speech), both mean-centred. We measure how much of \mathbf{Y} can be expressed inside the k -dimensional principal subspace of \mathbf{X} :

- (i) **PCA on \mathbf{X} .** Compute the top- k eigenvectors $\mathbf{U}_X^{(k)} \in \mathbb{R}^{d_X \times k}$ of the covariance of \mathbf{X} and the associated variance fraction $v_X(k) = \sum_{i=1}^k \lambda_i^{(X)}$.
- (ii) **Project & reconstruct \mathbf{Y} .**

$$\hat{\mathbf{Y}} = \mathbf{Y} \mathbf{U}_X^{(k)} \mathbf{U}_X^{(k)\top}.$$

- (iii) **Fraction of variance in \mathbf{Y} explained.**

$$f(k) = 1 - \frac{\|\mathbf{Y} - \hat{\mathbf{Y}}\|_F^2}{\|\mathbf{Y}\|_F^2} \in [0, 1].$$

A score of 1 indicates that the \mathbf{Y} variance lies entirely within the dominant \mathbf{X} span; lower values signal weaker linear alignment between the two modalities. In all experiments we use $k = 50$.

C.2 Word segmentation

We apply a peak detector to the sequences of last-layer dynamic weights, $\omega^{(L)}$, using SciPy’s `find_peaks` tool. Performance is evaluated as a binary prediction task, where a peak indicates the

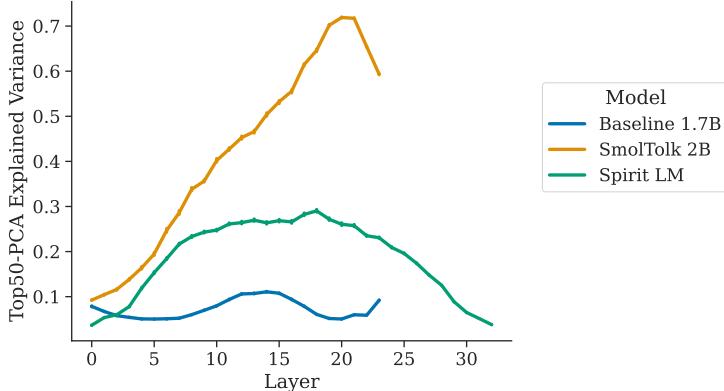


Figure 5: Cross-modal alignment score in our largest models and SPIRIT LM

prediction of a boundary at a given position with a tolerance of one token to account for noise in the boundary annotations. As in standard binary prediction tasks, recall, precision, and the F1-score are used as performance metrics. Additionally, we use the R-value [Räsänen et al., 2009], which penalizes trivial over-segmentation solutions. To optimize performance, we tune the prominence parameter of the peak detector over a grid $(0, 0.15]$ with steps of 0.01 so as to maximize the R-value, following Kreuk et al. [2020].

D Additional results

D.1 Cross-modal subspace overlap for larger models

We compute the cross-modal representation space alignment score described in Section 5.2 for our largest models—namely, the 1.7B baseline, SMOLTOLK-2B, and SPIRIT LM, a much larger model (7B parameters trained on over 170B tokens). Unlike Figure 3, where we compared models with the same representation dimensionality, here the models have different dimensions, and thus may exhibit different spectral tails. To account for this, we normalize the overlap score by the fraction of variance captured by the \mathbf{Y} subspace:

$$\frac{f(k)}{v_Y(k)}.$$

In Figure 5, a normalized score of 1 indicates that the dominant variance of \mathbf{Y} lies entirely within the dominant subspace of \mathbf{X} .

Several observations are worth noting: larger models with early fusion/fission—such as the baseline and SPIRIT LM—exhibit increased alignment in the middle layers. We hypothesize that, by this point, these models have developed sufficiently high-level representations for some text-learned functions at those layers to become useful. This likely contributes to the improved cross-modal performance observed in the 1.7B baseline (Figure 2). As suggested in Valeriani et al. [2023], Cheng et al. [2025], models are expected to reach a peak of compositionality, after which some layers may be dedicated to processing high-level features. We further expect that deeper models have more capacity for such processing (rather than immediate refinement toward next-token prediction), which may help explain SPIRIT LM’s strong performance. Greater depth gives speech representations more opportunities to undergo composition and approximate the abstraction level needed to exploit text-learned functions in later layers.

Nonetheless, alignment remains substantially lower than in SMOLTOLK-2B. As observed at smaller scales, early fusion models struggle to leverage text-learned functionality due to insufficient composition in early layers. Similarly, the absence of late multi-level fission limits abstraction and may force these models to use their final layers for extracting speech-predictive features, thereby preventing or disrupting the use of next-word predictive knowledge encoded in those layers. Given these observations, we expect that at a similar scale to SPIRIT LM, a model with our architecture would outperform it.

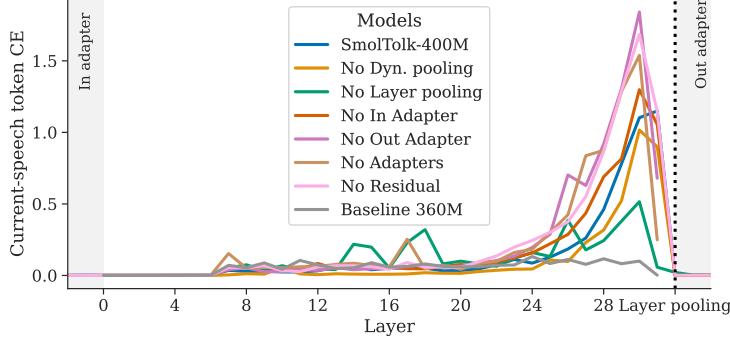


Figure 6: Forgetting of the current speech token across layers.

Setup	Prompt	Generation
T→S	The world is [swt]	[Hu24] [Hu12] [Hu341] ... [Hu500] ➡ too small for all the ideas of mine...
T→S	Here's a story about a cat that can speak: [swt]	[Hu43] [Hu412] [Hu433] ... [Hu500] ➡ This cat was very special because it could...
S→T	[Hu324] [Hu72] [Hu405] ... [Hu341] [swt] ➡ Illustration: the lemon [swt]	is a fruit that grows on trees it is round and yellow and tastes sour...
S→T	[Hu14] [Hu201] [Hu81] ... [Hu14] [swt] ➡ The Earth has undoubtedly entered upon a new orbit, but she is now encountering a [swt]	new phase of its history to be followed by a new and marvelous system of life upon The Earth...

Table 7: SMOLTOLK-2B cross-modal generations. The [swt] token denotes a modality switch in the generation process. HuBERT speech tokens are denoted as [Hu]. SMOLTOLK-2B displays coherent and consistent semantic cross-modal abilities.

D.2 Low-level speech information across layers

Figure 7 shows the cross-entropy loss for affine probes trained to predict the speech token w_i from the representation $c_i^{(l)}$ at each layer l of the models in Table 2. This measures how much low-level speech information is linearly preserved across the model’s layers. Notably, models without multi-level fission tend to preserve more of such information throughout layers.

D.3 SMOLTOLK cross-modal generations

As illustrated in Table 7, SMOLTOLK can generate semantically consistent content when prompted with either speech or text tokens. Audio prompts are taken from unseen LibriSpeech test-clean samples, while text prompts were manually crafted to resemble those in Table 1 of Nguyen et al. [2024], for comparison purposes.

D.4 SMOLTOLK + ASR and TTS

A common application of TSLMs is traditional speech processing tasks such as ASR and TTS. While our focus is on text LLM-like linguistic capabilities rather than these applications, we perform an additional experiment where, in addition to the tasks described in Section 3.1—speech modeling, text modeling, and interleaved text-speech modeling—we also include ASR and TTS training. This follows similar multi-task setups in prior work [Chou et al., 2023, Nguyen et al., 2024, Zeng et al., 2024, Défossé et al., 2024], and allows us to assess how expanding the task mix affects model performance.

In this setup, each training batch allocated one-twelfth of samples to ASR, one-twelfth to TTS, one-sixth to interleaved samples, and the remaining two-thirds to speech and text modeling. Table 8 reports zero-shot and few-shot results on LibriSpeech. For TTS evaluation, we follow Nguyen

Model	Params.	Tokens	# Shots	LS clean		LS other	
				ASR↓	TTS↓	ASR↓	TTS↓
SPIRIT LM [Nguyen et al., 2024]	7B	~175B	10	6.0	6.7	11.0	7.9
SMOLTOLK-2B+ASR+TTS	2B	16B	0	8.0	12.0	19.2	14.5
SMOLTOLK-2B+ASR+TTS	2B	16B	10	7.8	10.0	17.4	12.9

Table 8: Automatic Speech Recognition (ASR) and Text-to-Speech (TTS) performance on the LibriSpeech (LS) dataset. ASR scores are reported as Word Error Rate (WER, %), and TTS scores as Character Error Rate (CER, %). SMOLTOLK-2B is evaluated with a maximum context length of 2048 tokens. Few-shot examples are sampled from the corresponding LibriSpeech development sets.

Model	Params.	Tokens	tStoryCloze				sStoryCloze			
			T	S	T→S	S→T	T	S	T→S	S→T
SMOLTOLK-2B	2B	16B	92.6	87.5	83.9	86.0	73.2	60.0	64.0	63.4
SMOLTOLK-2B+ASR+TTS	2B	16B	92.3	87.0	83.1	86.0	72.6	60.6	62.2	62.6

Table 9: StoryCloze evaluation results comparing SMOLTOLK-2B and the more multitask variant trained with a mixture of interleaved, pure speech, pure text, ASR, and TTS samples.

et al. [2024], using Whisper Medium [Radford et al., 2022] to transcribe speech synthesized by the HiFiGAN vocoder from Hassid et al. [2023]. Despite being 3.5× smaller and trained on 11× less data than SPIRIT LM, SMOLTOLK-2B+ASR+TTS achieves competitive ASR and TTS performance. Results improve with more in-context examples, highlighting the model’s ability to adapt in context.

Finally, to assess whether adding ASR and TTS compromises cross-modal reasoning, we evaluate on the StoryCloze benchmark. As shown in Table 9, performance across other tasks is only minimally affected.

D.5 Scaling of other metrics

Figure 7 illustrates the scaling of several metrics for the baseline and SMOLTOLK models.

D.6 Extended benchmark table

Extended benchmark results, including the smaller baselines and other speech LMs from the literature, are presented in Table 10.

D.7 More multi-layer pooling samples and weights details

Figures 8, 9, and 10 provide additional results on the analysis of multi-layer pooling for SMOLTOLK-2B, SMOLTOLK-400M, and SMOLTOLK-150M, respectively.

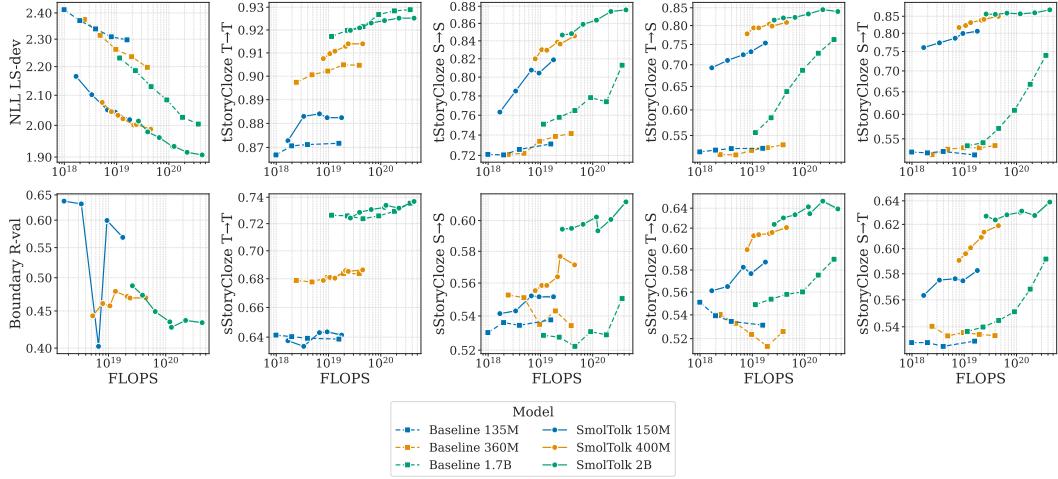


Figure 7: Scaling of the LibriSpeech dev set negative log-likelihood (NLL), R-val word segmentation score, and tStoryCloze and sStoryCloze accuracy across modalities with respect to training compute (in FLOPs).

Model	Params.	Tokens	BLIMP		tStoryCloze				sStoryCloze				MMLU	
			T	S	T	S	T→S	S→T	T	S	T→S	S→T	T (post/pre)	
<i>Textless Speech LMs</i>														
GSLM [Lakhota et al., 2021]	100M	—	∅	54.2	∅	66.6	∅	∅	∅	53.3	∅	∅	∅	∅
AudiolM [Borsos et al., 2023]	150M	—	∅	64.7	∅	—	∅	∅	∅	—	∅	∅	∅	∅
Hassid et al. [2023] cold-init 1.3B	1.3B	10.8B	∅	56.5	∅	—	∅	∅	∅	—	∅	∅	∅	∅
TWIST [Hassid et al., 2023] 7B	1.3B	10.8B	∅	57.0	∅	70.6	∅	∅	∅	52.4	∅	∅	∅	∅
TWIST [Hassid et al., 2023] 13B	7B	36B	∅	59.0	∅	74.1	∅	∅	∅	55.3	∅	∅	∅	∅
TWIST [Hassid et al., 2023] 13B	13B	36B	∅	59.2	∅	76.4	∅	∅	∅	55.4	∅	∅	∅	∅
Cuervo and Marxer [2024] best	823M	82B	∅	61.3	∅	78.0	∅	∅	∅	56.7	∅	∅	∅	∅
SyllableLM [Baade et al., 2025]	300M	1.2B	∅	63.7	∅	75.4	∅	∅	∅	—	∅	∅	∅	∅
AlignSLM [Lin et al., 2024] 7B	7B	—	∅	62.3	∅	86.8	∅	∅	∅	61.1	∅	∅	∅	∅
Slam (scaled) [Maimon et al., 2025]	358M	16.7B	∅	61.1	∅	84.2	∅	∅	∅	61.3	∅	∅	∅	∅
<i>Previous Text-Speech LMs</i>														
SPIRIT LM [Nguyen et al., 2024]	7B	~175B	73.3	59.7	95.8	90.5	78.6	94.3	74.0	66.3	64.7	71.7	37.7 / 39.0	
LAST [Turetzky and Adi, 2024]	~390M	—	—	56.8	—	—	—	—	—	—	—	—	—	
Moshi [Défossez et al., 2024]	7.7B	2.1T	—	58.8	—	83.0	—	—	—	60.8	—	—	49.8 / 54.3	
GLM-4-Voice [Zeng et al., 2024] 1.5B	1.5B	1T	—	—	—	77.5	81.4	90.1	—	55.4	58.6	64.0	—	
GLM-4-Voice [Zeng et al., 2024] 9B	9B	1T	—	—	—	83.0	85.0	93.6	—	62.4	63.2	76.3	—	
<i>Ours</i>														
Baseline 135M	135M	16B	79.0	52.0	87.0	73.2	53.3	52.7	63.9	54.0	53.8	53.7	30.3 / 30.2	
Baseline 360M	360M	16B	79.8	52.4	90.4	74.1	53.1	53.8	68.4	54.0	52.1	53.1	34.5 / 34.0	
Baseline 1.7B	1.7B	16B	79.9	56.3	92.8	77.5	72.6	67.3	72.5	53.0	57.0	57.6	40.0 / 40.0	
Baseline 1.7B	1.7B	32B	79.8	58.1	92.9	81.3	76.3	74.0	73.5	55.1	59.0	59.2	39.2 / 40.0	
SMOLTOLK-150M	150M	16B	79.4	58.0	88.4	82.0	75.2	81.0	64.1	55.0	58.8	58.4	30.0 / 30.2	
SMOLTOLK-400M	400M	16B	79.8	59.4	91.3	84.6	80.9	85.0	68.4	57.5	62.3	62.1	34.0 / 34.2	
SMOLTOLK-2B	2B	16B	80.2	61.4	92.6	87.5	83.9	86.0	73.2	60.0	64.0	63.4	40.0 / 40.0	
SMOLTOLK-2B	2B	32B	80.2	61.9	92.6	87.6	84.3	87.1	73.6	61.4	64.2	64.2	40.1 / 40.0	

Table 10: Downstream evaluations. The **best model** in each task is shown in bold and underlined. The **second best** is shown in bold. For SPIRIT LM we report the results for the open-weights version. For other models we present the results reported by the authors.

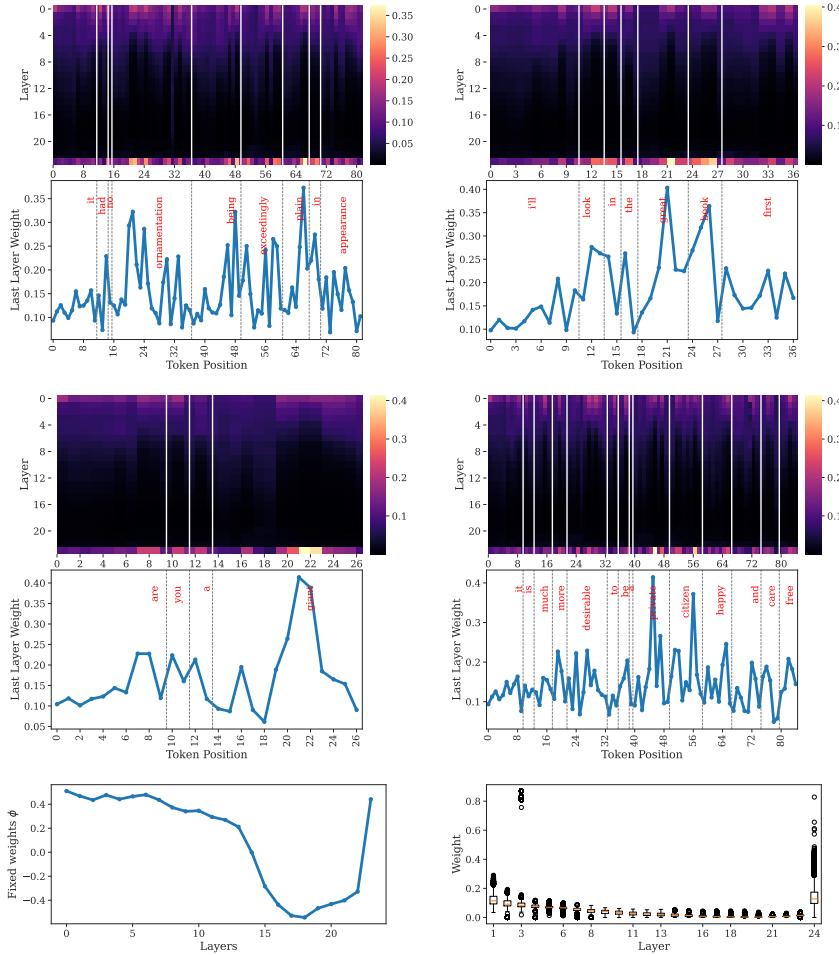


Figure 8: Selector S dynamic layer weights ω across several speech input sequences for SMOLTOLK-2B. Vertical bars indicate word boundaries (Top and middle). Non-input dependent learned weights ϕ (bottom-left). Distribution of dynamic weights ω across layers (bottom-right).

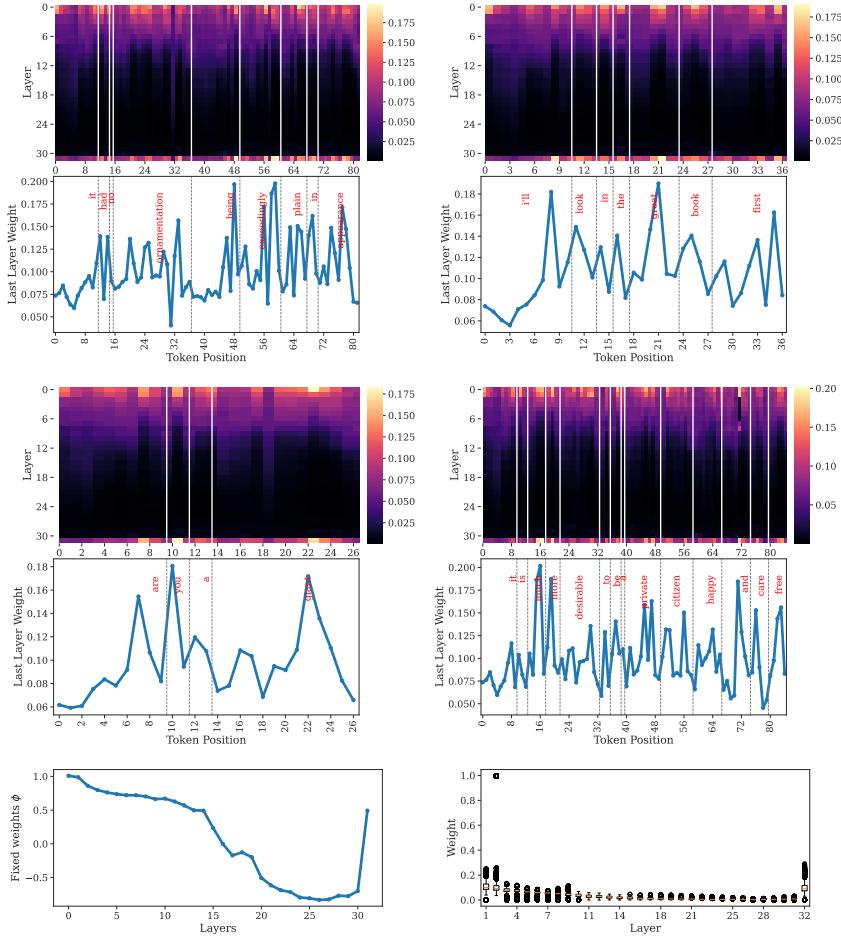


Figure 9: Selector S dynamic layer weights ω across several speech input sequences for SMOLTOLK-400M. Vertical bars indicate word boundaries (Top and middle). Non-input dependent learned weights ϕ (bottom-left). Distribution of dynamic weights ω across layers (bottom-right).

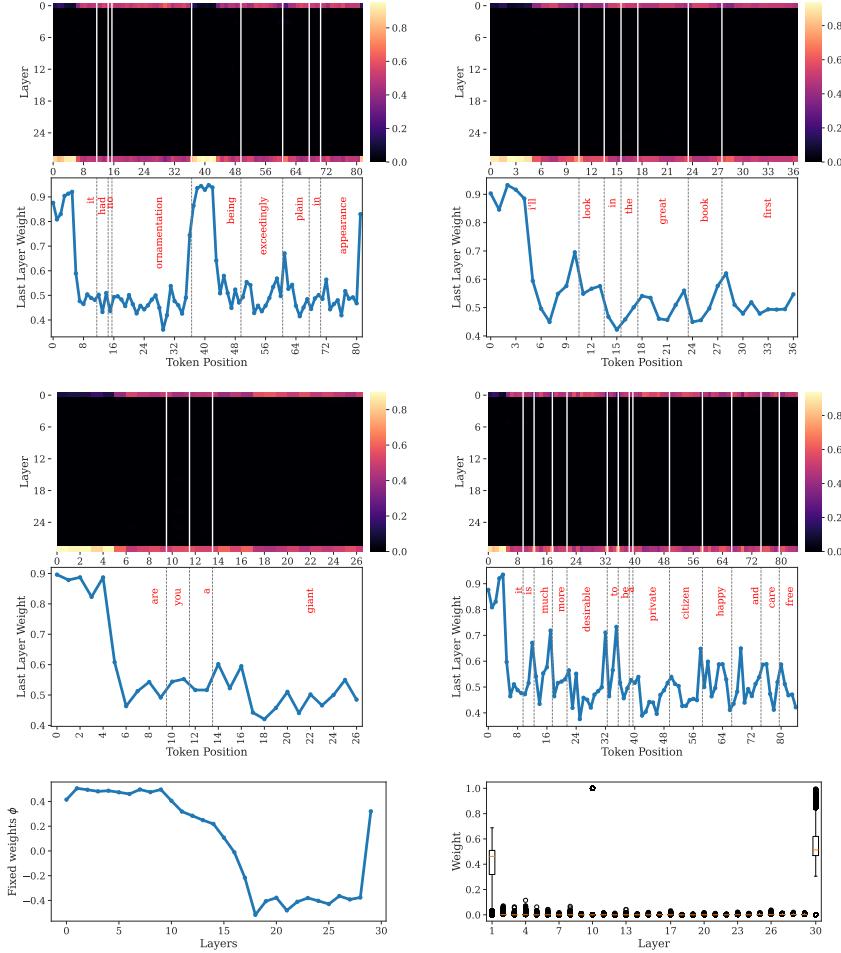


Figure 10: Selector S dynamic layer weights ω across several speech input sequences for SMOLTOLK-150M. Vertical bars indicate word boundaries (Top and middle). Non-input dependent learned weights ϕ (bottom-left). Distribution of dynamic weights ω across layers (bottom-right).