
Scaling Open Discrete Audio Foundation Models with Interleaved Semantic, Acoustic, and Text Tokens

Potsawee Manakul^{1,2} Woody Haosheng Gan³ Martijn Bartelds¹ Guangzhi Sun⁴ William Held^{1,5} Diyi Yang¹

Abstract

Current audio language models are predominantly text-first, either extending pre-trained text LLM backbones or relying on semantic-only audio tokens, limiting general audio modeling. This paper presents a systematic empirical study of native audio foundation models that apply next-token prediction to audio at scale, jointly modeling semantic content, acoustic details, and text to support both general audio generation and cross-modal capabilities. We provide comprehensive empirical insights for building such models: (1) We systematically investigate design choices—data sources, text mixture ratios, and token composition—establishing a validated training recipe. (2) We conduct the first scaling law study for discrete audio models via IsoFLOP analysis on 64 models spanning 3×10^{18} to 3×10^{20} FLOPs, finding that optimal data grows $1.6 \times$ faster than optimal model size. (3) We apply these lessons to train SODA (Scaling Open Discrete Audio), a suite of models from 135M to 4B parameters on 500B tokens, comparing against our scaling predictions and existing models. SODA serves as a flexible backbone for diverse audio/text tasks—we demonstrate this by fine-tuning for voice-preserving speech-to-speech translation, using the same unified architecture.

1. Introduction

Building foundation models that can understand and generate audio is a key challenge in multimodal AI. Current approaches have distinct limitations. *LLM-centric* architectures, such as SALMONN (Tang et al., 2024) or Qwen3-Omni (Qwen Team, 2025b), add audio modules to a pre-

¹Stanford University ²SCB 10X ³University of Southern California ⁴University of Cambridge ⁵OpenAthena. Correspondence to: Potsawee Manakul <m.potsawee@gmail.com>, William Held <held@stanford.edu>, Diyi Yang <diyiy@stanford.edu>.

Project Page: <https://soda-audio.github.io/> for model checkpoints, discrete audio data, code, and experiment log.

trained text LLM; while effective for instruction-following, they have a “semantic bottleneck” that limits general audio-to-audio modeling. *Semantic-only* speech language models, such as TWIST (Hassid et al., 2023) or SpiritLM (Nguyen et al., 2025), are trained speech-first but discard acoustic details, limiting high-fidelity understanding and generation. *Native audio* models like Moshi (Défossez et al., 2024) or Llama-Mimi (Sugiura et al., 2025) model acoustic tokens directly but focus on specific tasks without text integration. Meanwhile, next-token prediction has enabled unified models for text and vision-language (Chameleon Team, 2024), yet analogous approaches that jointly model audio understanding and generation in a single backbone remain limited.

To bridge this gap, this paper presents a systematic empirical study of *native* audio foundation models that jointly model semantic, acoustic, and text tokens within a unified next-token prediction framework—establishing the first training recipes and scaling laws analogous to scaling study in LLMs (Kaplan et al., 2020). This design enables a range of tasks within a single model: audio continuation, semantic/acoustic understanding, cross-modal capabilities (e.g., text-to-speech and speech-to-text), and text generation. We adopt utterance-level interleaving of tokens derived from neural codecs, avoiding word-level alignment errors and enabling the use of large datasets with available transcripts. A challenge in training such audio models is the lack of established pretraining understanding: while the Chinchilla study for text LLMs (Hoffmann et al., 2022) established that model size N and training tokens D should scale equally ($N^*, D^* \propto C^{0.5}$), it is unclear if this holds for audio, where information density per token can be far lower. We address key questions for pre-training discrete audio models:

- **What training data and token design should we use?** (§4): We systematically compare speech corpora, text mixture ratios, and token compositions (semantic-only vs. semantic+acoustic vs. semantic+acoustic+text), establishing a validated training recipe.
- **How should we allocate compute, and is validation loss a reliable metric?** (§5): We show that validation loss is predictive of downstream performance, then derive scaling laws from 64 IsoFLOP models (3×10^{18} to 3×10^{20} FLOPs), finding $D^* \propto C^{0.579}$ and $N^* \propto C^{0.367}$.

- **Does scaling up work? (§6):** We train SODA (Scaling Open Discrete Audio), a suite of models from 135M to 4B parameters on 500B tokens (up to 1.3×10^{22} FLOPs), and validate them against our scaling predictions. We compare cold-start (from scratch) versus warm-start (from text LLMs) training at scale, finding that cold-start is superior and provides higher training stability. We further validate SODA as a flexible backbone by formulating voice-preserving speech-to-speech translation simply as a next-token prediction task and fine-tuning SODA.

SODA achieves competitive performance across audio and cross-modal benchmarks, with fine-tuning for S2ST demonstrating its flexibility. We release checkpoints, discrete audio data, experiment log, and code to facilitate future research.

2. Related Work

2.1. Audio & Speech Foundation Models

LLM-Centric Architectures. Models such as SALMONN (Tang et al., 2024), Llama-Omni (Fang et al., 2025), and Qwen3-Omni (Qwen Team, 2025b) warm-start from a pre-trained text LLM and add audio capability via separate encoder/decoder modules. The backbone processes text-aligned semantic representations, creating a “semantic bottleneck” where fine-grained acoustic details are compressed or lost. While effective for instruction following, these models cannot natively generate audio and often rely on separate modules like vocoders with fixed speaker embeddings, limiting their utility as end-to-end audio foundation models.

Semantic-Only Models. Approaches like TWIST (Hasid et al., 2023), SpiritLM (Nguyen et al., 2025), VoxtLM (Maiti et al., 2024), SUTLM (Chou et al., 2023), and SIMS (Maimon et al., 2025a) operate on discrete speech tokens but restrict themselves to semantic tokens (e.g., HuBERT units). VoxtLM and SUTLM combine text BPEs with HuBERT tokens and support ASR, TTS, and continuation via control tokens, but still lack acoustic detail. While SpiritLM introduces token interleaving, it focuses on semantic content, discarding the acoustic details required for acoustic understanding and high-fidelity audio generation.

Native Audio Models. The closest precursors to our work are native models that model discrete acoustic tokens directly. AudioLM (Borsos et al., 2023) pioneered the modeling of semantic and acoustic tokens but relied on a hierarchical cascaded architecture, generating semantic tokens first, followed by acoustic tokens in separate steps. Discrete audio models (such as VALLE (Wang et al., 2023), CosyVoice (Du et al., 2024), or Orpheus (Canopy Labs, 2025)) have shown success in TTS. Moshi (Défossez et al., 2024) introduced a full-duplex model for real-time dialogue, while Llama-Mimi (Sugiura et al., 2025) demonstrated that interleaving semantic and acoustic tokens within a single

Llama-3 decoder achieves the best acoustic consistency. However, these works focus on specific speech tasks without systematic investigation of training recipes or scaling behavior—gaps we address.

2.2. Scaling Laws for Foundation Models

For LLM pre-training, Kaplan et al. (2020) first established power-law relationships suggesting model size should scale faster than data ($N \propto C^{0.73}$). The Chinchilla study (Hoffmann et al., 2022) revised this, showing that for compute-optimal text LLMs, model size and training tokens should scale equally ($N^*, D^* \propto C^{0.5}$).

Recent works have attempted to extend scaling laws to the audio domain, yet remain constrained by their focus on *semantic-only* models. Both Cuervo & Marxer (2024) and Maimon et al. (2025a) limit their analysis to semantic tokens (HuBERT units), effectively discarding the acoustic details required for general speech/audio modeling. While Maimon et al. (2025a) improves upon the textless approach of Cuervo & Marxer (2024) by demonstrating that interleaving text accelerates learning, their methodology relies on a sparse number of models per compute budget, preventing reliable IsoFLOP curve fitting (N^* vs D^*). Furthermore, neither work investigates the scaling behavior of cross-modal capabilities (e.g., ASR and TTS skills), leaving a gap in understanding how audio skills emerge with scale.

3. Experimental Setup

Model Architecture: We train decoder-only Transformers using the Qwen3 architecture (Qwen Team, 2025a), which adds QK-Norm over Llama for improved training stability, with random initialization (cold-start).¹ We discretize audio using Mimi (Défossez et al., 2024) for its high reconstruction quality and semantic-acoustic separation, where the first codebook captures *semantic* content and remaining codebooks capture *acoustic* details. Mimi is a neural codec at 12.5 Hz, and we use first 8 RVQ codebooks (100 tokens/sec). Audio tokens are flattened and interleaved with text at the utterance level (shown in Figure 1(c)), with both audio-first and text-first variants for each instance. Unlike task-specific models (e.g., Orpheus for TTS (Canopy Labs, 2025), Llama-Mimi for audio continuation (Sugiura et al., 2025)), this interleaving enables the model to be *general-purpose*, learning four capabilities: (1) audio continuation, (2) text continuation, (3) audio→text, and (4) text→audio.

Training Data: For speech data, we select the largest publicly available corpora with utterance-level transcriptions: (1) *Yodas* (Li et al., 2023) (500K+ hours across 100+ lan-

¹ Section 6.3 compares cold-start vs. warm-start initialization. Training hyperparameters (optimizer, learning schedule, etc.) and hardware details are in Appendix A.3.

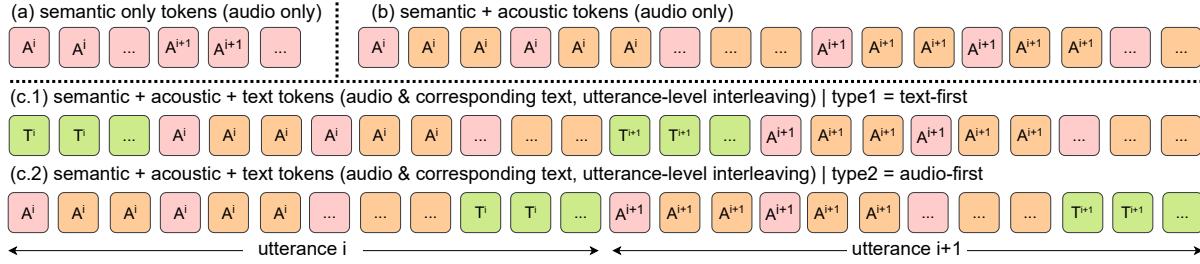


Figure 1. Three token types examined in § 4.3: (a) Semantic-only, (b) Semantic+Acoustic, and (c) Utterance-level interleaved Semantic+Acoustic+Text, where (c.1) shows text-first format, and (c.2) shows audio-first format. Superscript i denotes utterance index, where utterance(i) and utterance($i+1$) are adjacent audio segments from the same document. See Appendix A.2 for detailed data formatting.

guages; we use $\sim 165K$ hours of English), (2) *Emilia* (He et al., 2024) (101K hours of diverse, spontaneous speech; we use $\sim 140K$ hours of English), and (3) *MLS* (Pratap et al., 2020) (45K hours of audiobook speech). For text-only data, we use *Nemotron-CC* (Su et al., 2025), a large-scale web corpus widely used in LLM pre-training. Section 4 ablates these choices; full data statistics are in Appendix B.

Evaluation: We evaluate across four categories: (1) *Speech semantic knowledge* (sBLIMP, sWUGGY), (2) *Speech acoustic knowledge* (Salmon), (3) *Text knowledge* (tBLIMP, tWUGGY, HellaSwag), and (4) *Cross-modal skills* (LibriSpeech for ASR, seed-tts-eval for TTS). Full details on evaluation are provided in Appendix C.

SODA-Preliminary: As an initial exploration, we first validate whether training a vanilla transformer on our utterance-level interleaved discrete audio tokens yields meaningful capabilities. We train SODA-prelim, a 600M parameter model on 500B multilingual tokens (8 languages from Yodas), using multilingual data to maximize available training data from Yodas before processing other large corpora. The results confirm that joint semantic-acoustic-text modeling produces functional cross-modal skills (ASR, TTS) and strong acoustic understanding, but reveals limitations in semantic understanding and text knowledge (see Appendix D.1 for full details). Since standard benchmarks focus on English skills, subsequent experiments (§4) adopt English-only data to isolate multilingual effects.

4. What Data and Token Types Do We Use?

Before conducting scaling analysis, we perform empirical investigations to answer: *What speech corpora work best?* *How much text-only data should we include?* *What token composition (semantic, acoustic, text) is optimal?* These studies ensure our scaling laws reflect a well-optimized setting rather than a sub-optimal baseline. Each subsection uses a different experimental setup tailored to the question (see Appendix D.2 for rationale). To isolate data source effects, all experiments use English-only data.

4.1. What Speech Data Works the Best?

Question: Which speech data should be used for training? Building on SODA-prelim (§3), we compare available corpora to identify a better training recipe.

Setup: Comparing data sources via full training runs is too expensive. Recent work showed that *annealing experiments*, where different data mixtures are evaluated during the learning rate decay phase of pre-training, can reliably predict the relative quality of data sources at a fraction of the cost (Blakeney et al., 2024). We adopt this approach: we branch from the stable phase of SODA-prelim to compare speech datasets during annealing. We evaluate the three speech corpora introduced in §3: Yodas, Emilia, and MLS

Findings: Full results are in Table 5 (Appendix D.3); we summarize key findings here. Across semantic and acoustic tasks, all three corpora yield similar results. However, MLS shows notably poor cross-modal performance despite being a curated audiobook corpus: ASR-WER degrades to 92.6% and TTS-WER to 35.7%. We attribute this to: (i) uncased, unpunctuated transcripts creating a mismatch with standard text, and (ii) fixed-length 10–20 second chunks lacking length diversity. Between Emilia and Yodas, they show complementary strengths: Emilia achieves the best TTS results, while Yodas provides better text knowledge. Separately, we find that small-scale models trained from scratch (150M, 10B tokens), in line with Magnusson et al. (2025), confirm this result. Thus, we select **Yodas + Emilia**.

4.2. How Much Text-Only Data Should We Include?

Question: Given poor semantic understanding and general knowledge observed in SODA-prelim (see Appendix D.1), can we boost these capabilities by letting the model learn from high-quality text data in addition to noisy speech transcripts during pre-training? If so, what is the optimal ratio?

Setup: Since small-scale runs (150M, 10B tokens) yielded similar findings to expensive annealing in §4.1, we adopt this setup for the text ratio sweep. We vary the ratio from 0% to 50%, where X% text means sampling X% tokens

from Nemotron and $(100-X)\%$ from Yodas.

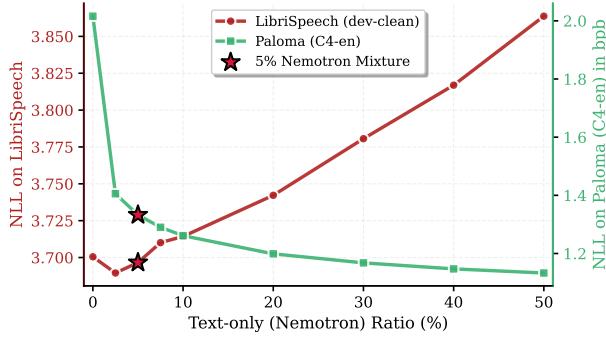


Figure 2. Impact of adding Nemotron on NLL for audio and text validation data. Full results on other metrics are shown in Figure 5.

Findings: Figure 2 reveals a trade-off between audio and text performance (we validate NLL as a reliable downstream predictor in §5.1). For NLL_{text} , adding any text yields a large improvement (drastic drop in NLL_{text} from 0% to 2.5%), with continued improvement as text ratio increases. NLL_{audio} does not degrade up to 5% text (matching the 0% baseline), but degrades beyond this point. We select **5%** as it maximizes text knowledge gains *with little to no degradation* in audio performance—practitioners prioritizing text/reasoning capabilities may increase this ratio at the cost of audio skills. Full results on downstream metrics are in Figure 5 (Appendix D.4). We fix the pre-training mixture to **5% Text (Nemotron) + 95% Speech (Yodas/Emilia)** for all subsequent experiments.

4.3. What Is a Good Token Composition?

Question: What is the impact of interleaved audio and text tokens, as well as the combination of semantic and acoustic audio tokens, on audio benchmarks?

Setup: We ablate three token types (illustrated in Figure 1) using a fixed budget of 3×10^{20} FLOPs (1.7B model, 30B data) and Yodas as speech data (without Nemotron). The three token configurations are: (1) *Semantic-only*, using only the first Mimi codebook, (2) *Semantic+Acoustic*, using up to 8 Mimi codebooks (audio only), and (3) *Semantic+Acoustic+Text*, interleave text (transcript) with audio tokens at the utterance level.

Table 1. Token ablation results. S=Semantic, A=Acoustic, T=Text. \times indicates the model lacks the capability to perform the task.

	Semantic	Acou.	Text	Cross-Modal			
	sBLI↑	sWUG↑	Salm↑	tBLI↑	tWUG↑	ASR _{WER} ↓	TTS _{WER} ↓
S	58.6	72.1	67.3	×	×	×	×
S+A	50.9	59.0	70.1	×	×	×	×
S+A+T	50.4	58.1	70.4	67.8	71.6	18.3	27.1

Findings: Table 1 reveals a trade-off: adding acous-

tic tokens improves acoustic modeling (Salmon: 67.3% \rightarrow 70.1%) but reduces semantic understanding (sBLIMP: 58.6% \rightarrow 50.9%). Interleaving text tokens has minimal further impact on these metrics while unlocking cross-modal (ASR/TTS) and text capabilities unavailable in audio-only models like our S+A model variant or Llama-Mimi.

While semantic-only models excel at semantic understanding, they lack the acoustic detail needed for high-fidelity understanding and generation. Since our goal is a *general-purpose* backbone, we adopt S+A+T, accepting the semantic trade-off for broader capabilities in a unified backbone.

5. How Should We Allocate Compute?

This section addresses: *Is validation loss a reliable metric for discrete audio models? How should we allocate compute between model size and training data?* Scaling laws for text LLMs are well-studied (Kaplan et al., 2020; Hoffmann et al., 2022; Grattafiori et al., 2024; Bi et al., 2024), yet no such analysis exists for discrete audio models. We conduct the first scaling law study for discrete audio models, investigating the extent to which the lower information density of audio tokens (100 tokens/sec vs. ~ 4 tokens/sec for text) alters the compute-optimal allocation between model size N and training data D . To determine the compute-optimal allocation, we conduct an IsoFLOP sweep training 64 models across seven compute budgets ranging from 3×10^{18} to 3×10^{20} FLOPs. For each budget C , we train models of varying sizes (from 77M to 4.2B parameters), adjusting the dataset size D to satisfy $C \approx 6ND$, with hyperparameters following existing scaling law work by Held et al. (2025).

5.1. Is Validation Loss a Reliable Metric?

Question: Before conducting IsoFLOP analysis, we assess whether validation loss (NLL) on held-out audio data is a reliable metric for evaluating discrete audio models. If this is the case, next experiments can focus on minimizing NLL.

Setup: We compute NLL on speech utterances from LibriSpeech dev-clean and analyze the correlation between NLL and downstream task performance across 64 models of varying sizes and training configurations. Results in Figure 3, with full results in Figures 6 and 7 (Appendix E).

Findings: Across all compute budgets and model sizes, validation NLL shows **strong rank correlation (Spearman ρ) with downstream performance**.

For *cross-modal skills*, NLL is highly predictive of both ASR ($\rho \approx 0.95$) and TTS quality (TTS-WER: $\rho \approx 0.96$, TTS-SIM: $\rho = 0.99$ with near-linear improvement). However, at lower loss values, improvement slows as most results from the highest compute budget (3×10^{20} FLOPs) stay worse than the regression line.

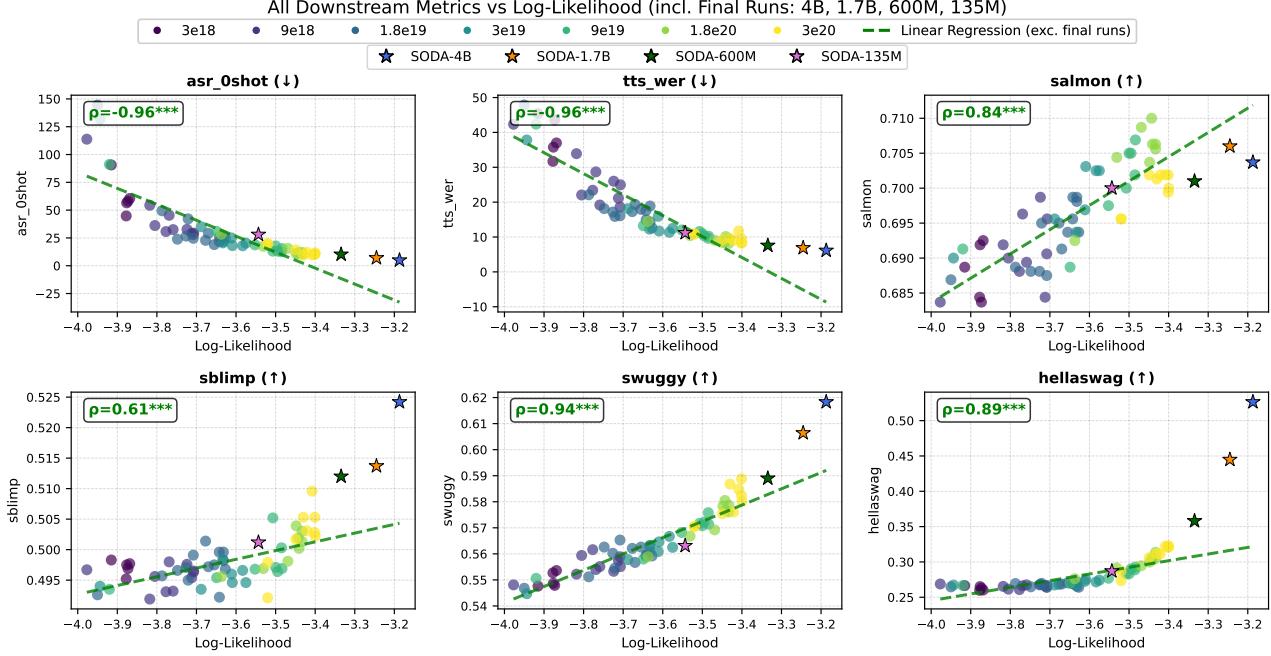


Figure 3. Validation NLL (audio+text) versus downstream task performance. Circular points: 64 IsoFLOP models (§5); star-shaped points: final SODA runs (§6). Regression lines are fitted on 64 IsoFLOP models only. Full results with other metrics are shown in Appendix E.

For *semantic and acoustic understanding* (Salmon, sBLIMP, sWUGGY), the rate of improvement is slow: as NLL decreases from 4.0 to 3.4, Salmon improves from 68.5% to 70.5%, and sWUGGY from 54% to 58%. For Salmon, we observe early signs of saturation: models at the highest compute budget (3×10^{20} FLOPs) mostly fall below the regression line, suggesting that acoustic understanding may be reaching a diminishing return point. For sBLIMP, performance remains largely pre-emergence (49.3% to 50.0%, near the 50% random baseline), suggesting this metric is not yet informative at our current scale.

For *text knowledge* tasks (tBLIMP, tWUGGY, HellaSwag; see Figure 7), NLL correlates strongly ($\rho > 0.8$) and improvement is more pronounced than speech understanding: as NLL decreases from 4.0 to 3.4, tWUGGY improves from 62% to 69% and tBLIMP from 64% to nearly 70%, with similarly no sign of saturation yet. HellaSwag ($\rho = 0.89$) shows an emergence pattern: no improvement from NLL 4.0 to 3.6, but rapid increase from 25% to 32% as NLL drops from 3.6 to 3.4.

Choice of NLL Metric: Our interleaved format admits multiple ways to compute validation NLL (e.g., all tokens from audio-first data, audio-only tokens, text-only tokens). While different variants yield similar correlations, we select **NLL on all tokens from audio+text data** as our primary metric because it is the simplest (standard NLL over all tokens) and achieves the best balance between speech and text task correlations (see Table 6 in Appendix E).

Extrapolation to Larger Scale: We validate whether these NLL–performance trends hold beyond the IsoFLOP regime by including the final SODA runs (colored points in Figure 3). These models, trained at higher compute budgets (§6), largely follow the extrapolated regression lines—confirming NLL as a reliable proxy even at scale. We discuss task-specific extrapolation patterns in detail in §6.

5.2. What Is Compute-Optimal for Discrete Audio?

Question: Having established that validation loss is a reliable metric, how should we allocate compute between model size (N) and training data (D) to minimize the loss?

Setup: Figure 4a shows validation loss versus model size N for each compute budget. Following Chinchilla’s methodology (Hoffmann et al., 2022), we fit a quadratic function $L = a(\log N)^2 + b(\log N) + c$ to the observed points and identify the compute-optimal model size N^* at the minimum of this curve. Similarly, Figure 4b shows validation loss versus training tokens D , with the compute-optimal token count D^* identified analogously.

Findings.² Given the compute-optimal (N^* , D^*) identified at each of the seven compute budgets, we fit power-laws $N^* = a_N C^{b_N}$ and $D^* = a_D C^{b_D}$ using log-linear regres-

²Another approach to derive scaling law exponents is to fit a parametric equation $L = E + A/N^\alpha + B/D^\beta$ on all data points and derive exponents as $N^* \propto C^{\beta/(\alpha+\beta)}$ and $D^* \propto C^{\alpha/(\alpha+\beta)}$. This yields similar findings; see Appendix F for details.

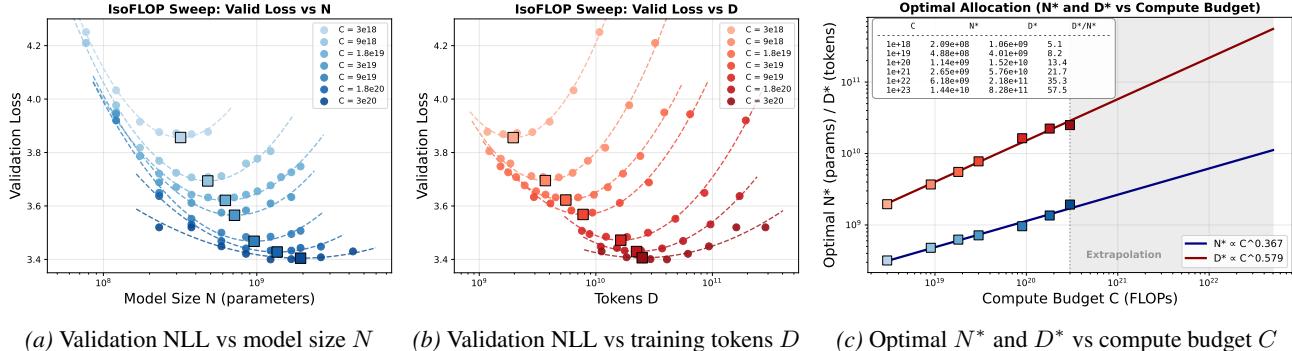


Figure 4. IsoFLOP analysis for discrete audio modeling. (a) and (b) show the loss landscape across model sizes and token counts for each compute budget. (c) shows the fitted scaling laws with extrapolation, revealing that optimal data D^* scales faster than optimal model size.

sion to derive scaling laws for discrete audio (illustrated in Figure 4c). This yields:

$$N^* \propto C^{0.367} \quad (1)$$

$$D^* \propto C^{0.579} \quad (2)$$

These exponents differ from Chinchilla’s scaling laws ($N^*, D^* \propto C^{0.5}$), with data scaling faster than model size. However, recent text LLM studies have found similar trends: Llama3 (Grattafiori et al., 2024) reports an exponent of data scaling of 0.53, and DeepSeek (Bi et al., 2024) finds exponents ranging from 0.42 to 0.55 depending on data quality. This DeepSeek work interprets higher data exponents as indicative of lower information density in training data. Our exponent ($D^* \propto C^{0.579}$) is in line with this interpretation, suggesting that discrete audio tokens at 100 tokens/sec carry less information per token than text, requiring more data to achieve the same effective learning. We also note that IsoFLOP curves tend to flatten at larger scales (Grattafiori et al., 2024), so our smaller-scale regime ($\leq 3 \times 10^{20}$ FLOPs) may contribute to the higher exponent. To our knowledge, this represents the first scaling analysis for discrete audio models, establishing a foundation for future investigations into how exponents vary with tokenizer design and token rate.³

Optimal Token-to-Parameter Ratio. This asymmetry implies that the optimal token-to-parameter ratio (D^*/N^*) is not constant but *increases* with scale. Our projections suggest an optimal ratio of around 13 tokens per parameter at 10^{20} FLOPs, rising to around 58 tokens per parameter at 10^{23} FLOPs. While the Chinchilla study suggested a constant ratio of 20 tokens per parameter, a replication attempt (Besiroglu et al., 2024) found this estimate was poorly fit and that the ratio actually *decreases* with compute for text models. Our finding of an *increasing* ratio for discrete audio represents a qualitatively different scaling behavior.

³Studying how scaling exponents change with token rate (reflecting information density) and audio data quality is an interesting direction we leave to future work.

6. Does Scaling Up Work?

This section addresses: *Does scaling up lead to competitive performance? Should we train from scratch or from text LLMs?* Guided by recipes from §4 and scaling analysis from §5, we train **SODA** (Scaling Open Discrete Audio), ranging from 135M to 4B parameters, and evaluate against our scaling predictions and existing spoken language models.

6.1. Setup and Over-Training

Training data: Following the recipe from §4, we train on 95% speech corpora (Yodas + Emilia) and 5% Nemotron text corpus (§3). This yields around 125B + 125B = 250B tokens of interleaved speech data (audio-first and text-first formats), for a total of **500B tokens** (~ 4 epochs). Prior work shows repeated data remains effective up to 4 epochs for LLM pre-training (Muennighoff et al., 2023).

Model sizes and over-training: We train models at 135M, 600M, 1.7B, and 4B parameters, all on 500B tokens. While our scaling laws define compute-optimal token counts, inference usage favors models trained beyond D^* , especially given that our 100 tokens/sec rate can make larger models slow. This results in varying over-training factors: 135M at $\sim 940 \times D^*$, 600M at $\sim 90 \times$ (similar to Llama3), 1.7B at $\sim 18 \times$ (similar to Llama2), and 4B at $\sim 4.5 \times$ (approaching compute-optimal). Training the 4B model on 500B tokens reaches 1.3×10^{22} FLOPs (~ 1 week on v5p-256 TPU). See Appendix F for detailed over-training analysis.

6.2. How Well Does SODA Perform?

6.2.1. DO OUR DESIGN CHOICES HELP?

Comparing SODA-600M-base to SODA-600M-prelim (both 600M parameters trained on 500B tokens) validates that our design choices in §4, including the shift to English-only and data mixture, yield measurable improvements. Table 2 shows that SODA-600M-base improves across all metrics: notably, ASR-WER drops from 22.0% to 10.2%,

Table 2. Results comparing SODA against existing Speech Language Models. SODA is the only model with capabilities across all skills, serving as a unified backbone. \times indicates lacks the capability; $-$ indicates results not reported due to HellaSwag contamination known in Llama2 pre-training which SpiritLM is based on (Touvron et al., 2023). † SpiritLM uses 10-shot evaluation on LibriSpeech-clean for ASR/TTS whereas ours is 0-shot with TTS on seed-tts-eval; ‡ TTS-SIM reproduced—semantic-only tokens cannot preserve voice.

Model	Speech (Semantic)		Speech (Acoustic)		Text (Knowledge)			Cross-Modal		
	sBLIMP \uparrow	sWUGGY \uparrow	Salmon \uparrow	tBLIMP \uparrow	tWUGGY \uparrow	HellaS \uparrow	ASR \downarrow	TTSWER \downarrow	TTS \downarrow	
TWIST-7B	59.0	73.9	61.6	\times	\times	\times	\times	\times	\times	
SpiritLM-base-7B	58.3	69.0	57.2	73.3	80.3	$-$	$\sim 22^\dagger$	$\sim 40^\dagger$	$\sim 0.05^\ddagger$	
SpiritLM-Expr-7B	54.2	65.0	67.1	73.6	75.8	$-$	$\sim 38^\dagger$	$\sim 50^\dagger$	$\sim 0.10^\ddagger$	
Llama-Mimi-1.3B	54.3	68.7	73.6	\times	\times	\times	\times	\times	\times	
Llama-Mimi-8B	55.1	68.8	73.2	\times	\times	\times	\times	\times	\times	
SODA-prelim-600M	50.9	57.8	69.4	69.0	71.3	26.2	22.0	9.2	0.516	
SODA-base-135M	50.1	56.3	70.0	67.4	70.7	28.7	28.1	11.2	0.500	
SODA-base-600M	51.2	58.9	70.1	70.7	73.1	35.8	10.2	7.6	0.555	
SODA-base-1.7B	51.4	60.6	70.6	70.3	74.7	44.5	7.0	6.9	0.560	
SODA-base-4B	52.4	61.8	70.4	71.3	74.8	52.6	5.0	6.1	0.560	

TTS-WER from 9.2% to 7.6%, and TTS-SIM increases from 0.516 to 0.555. Semantic understanding (sWUGGY: 57.8% \rightarrow 58.9%) and text knowledge (tWUGGY: 71.3% \rightarrow 73.1%) also improve. These gains across all skills reflect the refined recipe established in §4.

6.2.2. DOES PERFORMANCE IMPROVE WITH SCALE?

We examine how downstream task performance improves with scale, connecting the NLL correlation analysis from §5.1 to our final SODA runs (Table 2 and Figure 3). For validation loss predictions under over-training, see Appendix F.

Cross-modal tasks (ASR, TTS) showed strong NLL correlation ($\rho > 0.95$) with early saturation signs in Figure 3. The final runs confirm this: ASR-WER drops dramatically from 28.1% (135M) to 5.0% (4B), but all final-run points fall above the regression line, indicating diminishing returns at scale. TTS evaluation follows a similar pattern.

Acoustic understanding (Salmon), which showed saturation at the highest compute budgets, plateaus around 70% across all model sizes (70.0% at 135M to 70.4% at 4B). Final runs fall below the regression line in Figure 3, confirming saturation—acoustic ability appears bounded by tokenization and/or data quality rather than model capacity.

Semantic understanding shows emergence. While sWUGGY and sBLIMP showed slow improvement in IsoFLOP models in Figure 3, final SODA runs fall *above* the regression line: sWUGGY improves from 56.3% (135M) to 61.8% (4B), and sBLIMP from 50.1% to 52.4%, suggesting emergence and accelerating gains at larger scales.

Text knowledge exhibits the strongest emergence. Our analysis in Fig. 3 hinted at emergence (rapid improvement as $NLL < 3.6$), and the final runs confirm exponential gains: accuracy increases from 28.7% (135M) to 52.6% (4B), with final-run points appearing far above the regression line.

6.2.3. HOW DOES SODA COMPARE TO OTHERS?

SpiritLM uses semantic-only tokens (HuBERT), achieving strong semantic understanding (sBLIMP: 58.3%, sWUGGY: 69.0%) but weaker acoustic modeling (Salmon: 57.2–67.1%). While SODA’s semantic scores are lower with our interleaved setup, we show in Table 1 (§4.3) that a semantic-only SODA variant at 3×10^{20} FLOPs achieves sBLIMP: 58.6% and sWUGGY: 72.1%, outperforming both SpiritLM variants. This confirms that the semantic-acoustic trade-off is a design choice, not a fundamental limitation. However, semantic-only models—whether SpiritLM or our semantic-only variant—are limited in practice: they primarily model *what* is said but not *how* it is said, lacking the acoustic details necessary for general audio capabilities, reflected in extremely low TTS-SIM scores in Table 2.

Llama-Mimi achieves the highest acoustic scores (Salmon: 73.6%) by modeling audio-only sequences. Although SODA achieves lower acoustic scores, our ablation in Table 1 shows that incorporating interleaved text tokens does not degrade this capability (Salmon: 70.1% \rightarrow 70.4%), suggesting the gap is attributable to training data differences rather than model design. Nevertheless, *Llama-Mimi* lacks cross-modal capabilities entirely as it cannot perform ASR or TTS. Our utterance-level interleaving makes SODA a more general-purpose foundation, capable of both audio understanding and speech↔text tasks in a unified model.

6.3. Should We Train From Scratch or Warm-Start?

Question: Many discrete audio models (e.g., TWIST, CSM) initialize from pre-trained text LLMs. Does this *warm-start* strategy benefit over training from scratch (*cold-start*)?

Setup: We compare warm-start (from Qwen3-0.6B/1.7B-base) versus cold-start at 600M and 1.7B scales, training on 500B tokens. We provide final evaluation results as well as

throughout training trajectories, comparing warm-start vs. cold-start in Table 8 and Figure 11 in Appendix G.

Note on Training Stability: Warm-Start exhibits *instability*, with unpredictable loss spikes (including a big spike at 135K steps in 600M, ASR degrades from 21% to 34%; see Figure 10). Cold-Start shows smooth improvement throughout.

Cross-Modal Skills. Cold-Start outperforms Warm-Start on ASR from the earliest checkpoint and maintains this advantage: at 1.7B, Cold-Start achieves 19.7% vs 29.2% WER at 10K steps, widening to 7.0% vs 17.3% at completion. This suggests warm-start may interfere with learning audio→text mappings. For TTS, Cold-Start is initially better but Warm-Start catches up; both achieve comparable quality at completion ($\sim 6.5\text{--}7.5\%$ WER, ~ 0.56 SIM).

Speech Understanding. Salmon ($\sim 70\%$) and sWUGGY ($\sim 60\%$) show similar trajectories for both starts, suggesting these audio skills are learned regardless of initialization.

Text Knowledge. Warm-Start begins with a substantial advantage (at 10K steps: tWUGGY 75.4% vs 69.6%, HellaSwag 40.5% vs 29.9%). Critically, Cold-Start *never* catches up even after 500B tokens (final tWUGGY: 74.7% vs 79.2%; HellaSwag: 44.5% vs 47.1%), indicating text knowledge from LLM pre-training is not fully recoverable through audio-centric training.

Recommendation. Given the training instability and ASR degradation, we recommend **Cold-Start** as the default recipe for general audio capabilities. However, for capabilities where complex reasoning or knowledge is required, the text-knowledge advantage of Warm-Start may outweigh this. Future work could also consider hybrid approach, e.g., cold-start pre-training followed by text-enriched fine-tuning.

6.4. Could a New Audio-to-Audio Task Be Formulated as Next-Token Prediction by Fine-tuning SODA?

Question: Can SODA serve as a flexible backbone for new audio tasks? To test this, we fine-tune for *voice-preserving* speech-to-speech translation (S2ST)—an audio→audio task where translated speech must retain the source voice.

Prior work uses specialized architectures: Translatotron2 (Jia et al., 2022a) requires a speech encoder, phoneme decoder, spectrogram synthesizer, and shared speaker embeddings. SODA requires none of these—we simply format S2ST as interleaved next-token prediction (source audio → source text → target text → target audio) using the same decoder-only transformer as all other tasks.

Setup: We fine-tune on CVSS-T (Jia et al., 2022b), a multilingual voice-preserving S2ST corpus (21 languages → English) where target speech is synthesized to match the source speaker’s voice. We compare 600M models: (1) *SODA*: initialized from SODA-600M (English-only pre-training); (2)

SODA-P: initialized from SODA-prelim (multilingual pre-training, see §3); (3) *Qwen3*: initialized from Qwen3-0.6B; and (4) *Scratch*: random initialization. We evaluate on 200 held-out examples from each language using ASR-BLEU for translation quality and speaker similarity (SIM) for voice preservation. See Appendix H for details.

Table 3. Speech-to-speech (X→En) translation on CVSS-T. SODA-Prelim, despite weaker overall recipe and English performance, was pre-trained on multilingual Yodas (en, es, fr, de, th, ar, hi, zh).

Init.	Pre-training	X→En (BLEU↑)				
		es	fr	de	other	SIM↑
Scratch	None	3.5	4.5	3.5	2.7	0.349
Qwen3	Text	3.0	4.5	3.0	2.1	0.375
SODA	Text + Audio	13.5	14.7	9.4	4.3	0.466
SODA-P	Text + Audio	20.9	21.1	14.6	4.5	0.469

Findings: Table 3 shows that audio pre-training is crucial. SODA outperforms Scratch by 3–4× on BLEU and achieves substantially better voice preservation (SIM: 0.466 vs. 0.349). Text-only pre-training (Qwen3) provides no gains in BLEU over Scratch, confirming that *audio* pre-training—not just scale—drives the improvement. Notably, SODA-P outperforms SODA despite being a weaker backbone on English benchmarks (Table 2): its multilingual pre-training provides ASR capabilities in the source languages that SODA lacks. This demonstrates that task-specific skills can be enhanced by including relevant training data. For reference, prior work reports SIM scores of $\sim 0.30\text{--}0.41$ (Labiausse et al., 2025), and ASR-BLEU of $\sim 30\%$ on es, fr, de (Zheng et al., 2025) on related benchmarks, though direct comparison is difficult due to different evaluation protocols.

Takeaway: The same decoder-only transformer and NTP objective used for ASR/TTS, and audio continuation directly supports voice-preserving S2ST—validating SODA as a flexible backbone. Also, individual skills can be enhanced by adding task-specific data without architectural changes.

7. Conclusion

Given reasonable data and compute, this work answers: *how to train discrete audio models with NTP to achieve broad capabilities*, providing six findings: (1) adding text-only data improves text knowledge without degrading audio; (2) adding acoustic tokens enables high-fidelity generation but reduces semantic understanding. Future work should investigate balancing these trade-offs; (3) optimal data scales 1.6× faster than model size, consistent with lower information density as in LLM studies; (4) NLL reliably predicts downstream performance, even at larger scales; (5) cold-start outperforms warm-start for audio tasks; (6) capabilities scale differently—cross-modal and acoustic skills saturate, while semantic and text knowledge show accelerating gains.

Limitations

Limited Emergent Capabilities. Unlike large language models that exhibit emergent abilities at scale, we do not observe strong emergent skills in speech or audio capabilities from pre-training alone. Our evaluation demonstrates the model’s utility through fine-tuning on speech-to-speech translation. Future work could study how or when diverse audio capabilities emerge naturally or through audio few-shot learning.

Generalist vs. Specialist Trade-offs. As a unified foundation model, SODA prioritizes flexibility over optimizing for any single skill. While specialized models developed for specific tasks may achieve higher performance on their target benchmarks, SODA’s strength lies in its ability to formulate audio tasks as next-token prediction using the same underlying model, which could serve as a potential avenue for native audio pre-training.

Unexplored Design Choices. Building discrete audio foundation models involves numerous design decisions. While our study investigates what we believe are the primary factors, there exist other design choices that remain unexplored, for example, alternative tokenizers, different token sampling rates, RVQ token-weighted loss functions, and various architectural modifications. Future work could study these additional design choices to further refine training recipes.

Impact Statement

Positive Impacts. By releasing open models, training recipes, and scaling laws, this work democratizes access to audio foundation model research. The unified architecture supporting broad speech/audio capabilities could benefit accessibility and the use of audio interfaces in AI applications.

Potential Risks. Audio generation capabilities raise concerns about misuse, including voice cloning for impersonation or misleading content generation. High-fidelity speech synthesis could be exploited for deepfakes or fraud.

Mitigation Considerations. We encourage users to implement safeguards such as audio watermarking, voice-clone consent requirements, and synthetic speech detection.

Author Contributions

Potsawee Manakul and Will Held conceived and led the project. Potsawee Manakul led execution, implemented data processing and training, conducted experiments, and wrote the manuscript. Will Held established the infrastructure and initial training, provided direction, contributed to technical implementation, and provided feedback. Woody Gan assisted with analysis and fine-tuning explorations. Martijn Bartelds contributed to ideation, analysis, and review. Guangzhi Sun provided early direction, implemented eval-

uation, conducted sanity checks, and reviewed. Diyi Yang is the primary supervisor, providing resources, guidance, direction, and feedback throughout the project.

Acknowledgements

We thank the Marin and OpenAthena team for developing and maintaining the training infrastructure. We are particularly grateful to David Hall for his critical support with infrastructure and insights on LLM pre-training. The compute resource was supported by the Google TPU Research Cloud (TRC) and a Stanford HAI-GCP Grant as part of the Marin Project. PM is supported by SCB 10X for his joint-position at Stanford. We also thank SALT Lab members and Typhoon members for helpful discussions and feedback.

References

Besiroglu, T., Erdil, E., Barnett, M., and You, J. Chinchilla scaling: A replication attempt. *arXiv preprint arXiv:2404.10102*, 2024.

Bi, X., Chen, D., Chen, G., Chen, S., Dai, D., Deng, C., Ding, H., Dong, K., Du, Q., Fu, Z., et al. Deepseek LLM: Scaling open-source language models with longtermism. *arXiv preprint arXiv:2401.02954*, 2024.

Blakeney, C., Gururangan, S., and Jain, S. Does your data spark joy? performance gains from domain upsampling at the end of training. *arXiv preprint arXiv:2406.03476*, 2024.

Borsos, Z., Marinier, R., Vincent, D., Kharitonov, E., Pietquin, O., Sharifi, M., Roblek, D., Teboul, O., Grangier, D., Tagliasacchi, M., and Zeghidour, N. Audiolm: A language modeling approach to audio generation. *IEEE/ACM Trans. Audio, Speech and Lang. Proc.*, 31:2523–2533, June 2023. ISSN 2329-9290. doi: 10.1109/TASLP.2023.3288409. URL <https://doi.org/10.1109/TASLP.2023.3288409>.

Canopy Labs. Orpheus-tts: Human-like speech synthesis. <https://huggingface.co/canopylabs/orpheus-3b-0.1-ft>, 2025. Llama-based TTS model using SNAC tokens.

Chameleon Team. Chameleon: Mixed-modal early-fusion foundation models. *arXiv preprint arXiv:2405.09818*, 2024.

Chou, J.-C., Chien, C.-M., Hsu, W.-N., Livescu, K., Babu, A., Conneau, A., Baevski, A., and Auli, M. Toward joint language modeling for speech units and text. pp. 6582–6593, December 2023. doi: 10.18653/v1/2023.findings-emnlp.438. URL <https://aclanthology.org/2023.findings-emnlp.438/>.

Cuervo, S. and Marxer, R. Scaling properties of speech language models. In Al-Onaizan, Y., Bansal, M., and Chen, Y.-N. (eds.), *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pp. 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>.

Défossez, A., Mazaré, L., Orsini, M., Royer, A., Pérez, P., Jégou, H., Grave, E., and Zeghidour, N. Moshi: a speech-text foundation model for real-time dialogue. *arXiv preprint arXiv:2410.00037*, 2024.

Du, Z., Wang, Y., Chen, Q., Shi, X., Lv, X., Zhao, T., Gao, Z., Yang, Y., Gao, C., Wang, H., et al. Cosyvoice 2: Scalable streaming speech synthesis with large language models. *arXiv preprint arXiv:2412.10117*, 2024.

Fang, Q., Guo, S., Zhou, Y., Ma, Z., Zhang, S., and Feng, Y. LLaMA-omni: Seamless speech interaction with large language models. 2025. URL <https://openreview.net/forum?id=PYmrUQmMEw>.

Grattafiori, A., Dubey, A., Jauhri, A., Pandey, A., Kadian, A., Al-Dahle, A., Letman, A., Mathur, A., Schelten, A., Vaughan, A., et al. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024.

Hassid, M., Remez, T., Nguyen, T. A., Gat, I., Conneau, A., Kreuk, F., Copet, J., Défossez, A., Synnaeve, G., Dupoux, E., et al. Textually pretrained speech language models. *Advances in Neural Information Processing Systems*, 36: 63483–63501, 2023.

He, H., Shang, Z., Wang, C., Li, X., Gu, Y., Hua, H., Liu, L., Yang, C., Li, J., Shi, P., et al. Emilia: An extensive, multi-lingual, and diverse speech dataset for large-scale speech generation. In *2024 IEEE Spoken Language Technology Workshop (SLT)*, pp. 885–890. IEEE, 2024.

Held, W., Hall, D., Liang, P., and Yang, D. Relative scaling laws for llms. *arXiv preprint arXiv:2510.24626*, 2025.

Hoffmann, J., Borgeaud, S., Mensch, A., Buchatskaya, E., Cai, T., Rutherford, E., de Las Casas, D., Guy, L. A., et al. Training compute-optimal large language models. *arXiv preprint arXiv:2203.15556*, 2022.

Hu, S., Tu, Y., Han, X., Cui, G., He, C., Zhao, W., Long, X., Zheng, Z., Fang, Y., Huang, Y., Zhang, X., Thai, Z. L., Wang, C., Yao, Y., Zhao, C., Zhou, J., Cai, J., Zhai, Z., Ding, N., Jia, C., Zeng, G., dahai li, Liu, Z., and Sun, M. MiniCPM: Unveiling the potential of small language models with scalable training strategies. In *First Conference on Language Modeling*, 2024. URL <https://openreview.net/forum?id=3X2L2TFr0f>.

Inaguma, H., Popuri, S., Kulikov, I., Chen, P.-J., Wang, C., Chung, Y.-A., Tang, Y., Lee, A., Watanabe, S., and Pino, J. UnitY: Two-pass direct speech-to-speech translation with discrete units. In Rogers, A., Boyd-Graber, J., and Okazaki, N. (eds.), *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 15655–15680, Toronto, Canada, July 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.acl-long.872. URL <https://aclanthology.org/2023.acl-long.872>.

Jia, Y., Tadmor Ramanovich, M., Remez, T., and Pomerantz, R. Translatotron 2: High-quality direct speech-to-speech translation with voice preservation. In *International Conference on Machine Learning*, pp. 10120–10134. PMLR, 2022a.

Jia, Y., Tadmor Ramanovich, M., Wang, Q., and Zen, H. CVSS corpus and massively multilingual speech-to-speech translation. In *Proceedings of the Language Resources and Evaluation Conference (LREC)*, pp. 6691–6703, 2022b.

Kaplan, J., McCandlish, S., Henighan, T., Brown, T. B., Chess, B., Child, R., Gray, S., Radford, A., Wu, J., and Amodei, D. Scaling laws for neural language models. *arXiv preprint arXiv:2001.08361*, 2020.

Labiausse, T., Mazaré, L., Grave, E., Défossez, A., and Zeghidour, N. High-fidelity simultaneous speech-to-speech translation. 2025. URL <https://openreview.net/forum?id=fqjN8B6xVX>.

Le Godais, G., Linzen, T., and Dupoux, E. Comparing character-level neural language models using a lexical decision task. In Lapata, M., Blunsom, P., and Koller, A. (eds.), *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, Valencia, Spain, April 2017. Association for Computational Linguistics. URL <https://aclanthology.org/E17-2020/>.

Li, X., Takamichi, S., Saeki, T., Chen, W., Shiota, S., and Watanabe, S. Yodas: Youtube-oriented dataset for audio and speech. In *2023 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, pp. 1–8. IEEE, 2023.

Magnusson, I., Tai, N., Bogin, B., Heineman, D., Hwang, J. D., Soldaini, L., Bhagia, A., Liu, J., Groeneveld, D., Tafjord, O., Smith, N. A., Koh, P. W., and Dodge, J. Datadecide: How to predict best pretraining data with small experiments, 2025. URL <https://arxiv.org/abs/2504.11393>.

Maimon, G., Hassid, M., Roth, A., and Adi, Y. Scaling analysis of interleaved speech-text language models. In *Second Conference on Language Modeling*, 2025.

2025a. URL <https://openreview.net/forum?id=IXwgE8hyJs>.

Maimon, G., Roth, A., and Adi, Y. Salmon: A suite for acoustic language model evaluation. In *ICASSP 2025 - IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 1–5, 2025b. doi: 10.1109/ICASSP49660.2025.10888561.

Maiti, S., Peng, Y., Choi, S., Jung, J.-w., Chang, X., and Watanabe, S. Voxtlm: Unified decoder-only models for consolidating speech recognition, synthesis and speech, text continuation tasks. In *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 13326–13330. IEEE, 2024.

Marin Community. Marin 8b retrospective. <https://marin.readthedocs.io/en/latest/reports/marin-8b-retro/>, 2025. Accessed: 2025-10-30.

Muennighoff, N., Rush, A., Barak, B., Le Scao, T., Tazi, N., Piktus, A., Pyysalo, S., Wolf, T., and Raffel, C. A. Scaling data-constrained language models. *Advances in Neural Information Processing Systems*, 36:50358–50376, 2023.

Nguyen, T. A., de Seyssel, M., Rozé, P., Rivière, M., Kharitonov, E., Baevski, A., Dunbar, E., and Dupoux, E. The zero resource speech benchmark 2021: Metrics and baselines for unsupervised spoken language modeling. In *NeurIPS Workshop on Self-Supervised Learning for Speech and Audio Processing*, 2020.

Nguyen, T. A., Muller, B., Yu, B., Costa-jussa, M. R., El-bayad, M., Popuri, S., Ropers, C., Duquenne, P.-A., Algayres, R., Mavlyutov, R., Gat, I., Williamson, M., Synnaeve, G., Pino, J., Sagot, B., and Dupoux, E. SpiRit-LM: Interleaved spoken and written language model. *Transactions of the Association for Computational Linguistics*, 13: 30–52, 2025. doi: 10.1162/tacl_a.00728. URL <https://aclanthology.org/2025.tacl-1.2/>.

Panayotov, V., Chen, G., Povey, D., and Khudanpur, S. Librispeech: an asr corpus based on public domain audio books. In *Acoustics, Speech and Signal Processing (ICASSP), 2015 IEEE International Conference on*, pp. 5206–5210. IEEE, 2015.

Pratap, V., Xu, Q., Sriram, A., Synnaeve, G., and Collobert, R. MLS: A large-scale multilingual dataset for speech research. *ArXiv*, abs/2012.03411, 2020.

Qwen Team. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*, 2025a.

Qwen Team. Qwen3-omni technical report. *arXiv preprint arXiv:2509.17765*, 2025b.

Su, D., Kong, K., Lin, Y., Jennings, J., Norick, B., Kliegl, M., Patwary, M., Shoeybi, M., and Catanzaro, B. Nemotron-CC: Transforming Common Crawl into a refined long-horizon pretraining dataset. In Che, W., Nabende, J., Shutova, E., and Pilehvar, M. T. (eds.), *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Vienna, Austria, July 2025. Association for Computational Linguistics. URL <https://aclanthology.org/2025.acl-long.123/>.

Sugiura, I., Kurita, S., Oda, Y., and Higashinaka, R. Llama-mimi: Speech language models with interleaved semantic and acoustic tokens. *arXiv preprint arXiv:2509.14882*, 2025.

Tang, C., Yu, W., Sun, G., Chen, X., Tan, T., Li, W., Lu, L., MA, Z., and Zhang, C. 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>.

Touvron, H., Martin, L., Stone, K., Albert, P., Almahairi, A., Babaei, Y., Bashlykov, N., Batra, S., Bhargava, P., Bhosale, S., et al. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023.

Wang, C., Chen, S., Wu, Y., Zhang, Z., Zhou, L., Liu, S., Chen, Z., Liu, Y., Wang, H., Li, J., et al. Neural codec language models are zero-shot text to speech synthesizers. *arXiv preprint arXiv:2301.02111*, 2023.

Warstadt, A., Parrish, A., Liu, H., Mohananey, A., Peng, W., Wang, S.-F., and Bowman, S. R. BLiMP: The benchmark of linguistic minimal pairs for English. *Transactions of the Association for Computational Linguistics*, 8:377–392, 2020. doi: 10.1162/tacl_a.00321. URL <https://aclanthology.org/2020.tacl-1.25/>.

Zellers, R., Holtzman, A., Bisk, Y., Farhadi, A., and Choi, Y. HellaSwag: Can a machine really finish your sentence? In Korhonen, A., Traum, D., and Márquez, L. (eds.), *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 4791–4800, Florence, Italy, July 2019. Association for Computational Linguistics. doi: 10.18653/v1/P19-1472. URL <https://aclanthology.org/P19-1472/>.

Zheng, Z., Sun, X., Dinh, T., Yanamandra, A., Jain, A., Liu, Z., Hadap, S., Bhat, V., Aggarwal, M., Medioni, G., et al. Rosettaspeech: Zero-shot speech-to-speech translation from monolingual data. *arXiv preprint arXiv:2511.20974*, 2025.

A. Tokenizer and Data Formatting

This section provides details on the tokenizer and data formatting. For model architecture and audio codec details, see §3.

A.1. Tokenizer

SODA uses a custom tokenizer built on top of the Marin tokenizer⁴, which is based on Llama 3’s tokenizer, consisting of:

- **Text tokens:** 128,256 tokens from the base Marin tokenizer
- **Audio tokens:** 16,384 tokens (8 codebooks \times 2,048 codebook size) for Mimi codec
- **Special tokens:** 4 addition tokens, <|text_start|>, <|text_end|>, <|audio_start|>, <|audio_end|>
- **Total vocabulary:** 144,644 tokens

For audio tokens, we use direct codebook mapping with no BPE merges. Initially, we experimented with a merge list of up to 128K audio token merges but observed only \sim 10% token reduction. Given low reduction, we use 0 merges for simplicity.

Note: For warm-start experiments (§6.3) that initialize from Qwen3 weights, we use the Qwen3 tokenizer instead of Marin to match the pre-trained embeddings.

A.2. Data Formatting and Interleaving

A key design choice in SODA is utterance-level interleaving of text and audio tokens (illustrated in Figure 1). Each audio file (e.g., a full YouTube video) is processed through the dataset pipeline (Yodas/Emilia/MLS), which applies VAD to segment audio into chunks. For a document with N chunks (i.e., speech segments), we construct two training variants:

Text-first format:

```
<|begin_of_text|>
<|text_start|>text_1<|text_end|>
<|audio_start|>audio_1<|audio_end|>
<|text_start|>text_2<|text_end|>
<|audio_start|>audio_2<|audio_end|>
...
<|text_start|>text_N<|text_end|>
<|audio_start|>audio_N<|audio_end|>
<|end_of_text|>
```

Audio-first format:

```
<|begin_of_text|>
<|audio_start|>audio_1<|audio_end|>
<|text_start|>text_1<|text_end|>
<|audio_start|>audio_2<|audio_end|>
<|text_start|>text_2<|text_end|>
...
<|audio_start|>audio_N<|audio_end|>
<|text_start|>text_N<|text_end|>
<|end_of_text|>
```

Here, `chunk_i` and `chunk_(i+1)` are adjacent audio segments from the same document. Adjacent chunks may be temporally contiguous or separated by silence (depending on VAD); we do not treat them differently. Although this process is supposed to yield speech-only data, we observe many audio segments containing non-speech content (e.g., background noise, music, etc.), which allows our model to learn broader audio knowledge although not evaluated in this work.

⁴<https://huggingface.co/stanford-crfm/marin-tokenizer>

Sequence packing: We pack sequences to the maximum context length (4,096 tokens for most experiments) with no padding, maximizing training efficiency. Multiple documents may be packed into a single training sequence, separated by `<|end_of_text|>` tokens. Also, a single document longer than the maximum context length may be split into multiple training sequences.

This dual-format strategy ensures the model learns bidirectional cross-modal capabilities: audio→text (ASR) and text→audio (TTS), as well as audio continuation and text continuation, all within a single pre-training stage.

A.3. Training Hyperparameters

We provide training hyperparameters for reproducibility. Our small-scale (in § 4) and large-scale (in § 6) experiments use the same base configuration and optimizer settings, scaled appropriately for model size.

Optimizer. We use AdamW with $\beta_1 = 0.98$, $\beta_2 = 0.98$, $\epsilon = 10^{-16}$, weight decay 0.033, and gradient clipping with max norm 1.0. We apply a small z -loss with weight 10^{-4} to the output logits, following our previous work that improved pretraining stability text LLMs. Note that the Chameleon work (Chameleon Team, 2024) also showed that z -loss is particularly effective for multi-modal models on training stability.

Learning Rate Schedule. We use the Warmup-Stable-Decay (WSD) schedule (Hu et al., 2024) with 10% warmup, 70% stable, and 20% linear decay phases. The base learning rate is 0.003 at batch size 256 and model width 1024. For other configurations, we scale the learning rate as:

$$lr = 0.003 \times \sqrt{\frac{\text{batch_size}}{256}} \times \sqrt{\frac{1024}{\text{width}}} \quad (3)$$

where width is the model’s hidden dimension (e.g., 512 for 135M, 1024 for 600M, 2048 for 1.7B, 2560 for 4B).

Batch Size and Sequence Length. For large-scale experiments (in § 6), we use a sequence length of 4096 tokens and batch size of 512 for most experiments. The number of training steps is computed as $\lceil \text{tokens}/(\text{batch_size} \times \text{seq_len}) \rceil$.

Hardware. All experiments in this paper are conducted on Google Cloud TPU v5p. For small-scale runs, e.g., design choice ablations in § 4 and most of IsoFLOP runs in § 5, we use v5p-8 or v5p-16 depending on model size. For large-scale training (§ 6), we use v5p-32 (135M), v5p-64 (600M), v5p-128 (1.7B), and v5p-256 (4B).

For our IsoFLOP sweep runs of 64 models, please refer to exact configurations and hyperparameters outlined in our code at `experiments/audio/isoflop_audio_sweep.py`.

B. Training Data

We provide detailed statistics and descriptions for all datasets used in this work. Table 4 summarizes the key properties.

Table 4. Training data statistics. Hours and tokens are for English subsets used in our experiments. Token counts assume 100 tokens/sec for audio (8 Mimi codebooks at 12.5 Hz) and standard BPE tokenization for text. *Estimated from disk size. †Token counts include both audio-first and text-first formats (i.e., unpeated token counts would be half of this value).

Dataset	Split	Type	Hours	~Tokens	Source
Yodas (Li et al., 2023)	English	Speech	164K	~131B [†]	YouTube (100+ langs)
Emilia (He et al., 2024)	English+Yodas-English	Speech	139K	~110B [†]	YouTube (6 langs)
MLS (Pratap et al., 2020)	English	Speech	44.5K	~35B [†]	LibriVox audiobooks
Nemotron-CC (Su et al., 2025)	HQ-Actual	Text	–	~220B*	Web (CC filtered)

Speech Data. We select the largest publicly available speech corpora with utterance-level transcriptions:

Yodas (Li et al., 2023) is one of the largest open speech datasets, containing over 370K hours across 100+ languages sourced from YouTube with automatic transcripts. For our main experiments, we use the English subset (~164K hours, ~131B tokens). We note that in our preliminary experiment (SODA-prelim), we use 8 languages: English (131B tokens), Spanish (35B), French (17B), German (11B), Thai (0.5B), Hindi (0.4B), Arabic (0.3B), and Chinese (0.2B), totaling ~196B tokens. Later we wanted to isolate the impact of multilingual training, so we only used English data for SODA-prelim. The dataset

offers diverse speaking styles, recording conditions, and utterance lengths, making it well-suited for training general-purpose audio models.

Emilia (He et al., 2024) is a large-scale multilingual dataset (~ 101 K hours across 6 languages) specifically designed for speech generation. It emphasizes spontaneous, natural speech with diverse speaking styles. We use the English subset (~ 140 K hours). Emilia’s focus on spontaneity complements Yodas’s scale.

MLS (Multilingual LibriSpeech) (Pratap et al., 2020) is derived from LibriVox audiobooks, containing ~ 45 K hours of English read speech. While MLS is a high-quality curated corpus, we find in §4.1 that it performs poorly for our setting due to: (i) uncased, unpunctuated transcripts creating distribution mismatch with standard text, and (ii) fixed-length 10–20 second chunks lacking utterance length diversity. We therefore exclude MLS from our final training recipe.

Text Data. To improve semantic understanding and world knowledge, we incorporate high-quality text-only data:

Nemotron-CC (Su et al., 2025) is a large-scale filtered web corpus derived from Common Crawl, widely used in LLM pre-training. We use the high-quality actual split, containing around 220B tokens (approximated from its disk size).

Final Data Mixture. Based on our design choice investigations in §4, we use the following sampling weights for the IsoFLOP analysis (§5) and large-scale SODA training (§6): 5% text (Nemotron) and 95% speech (Yodas + Emilia). Within the speech portion, we sample proportionally to dataset size (shown in Table 4), resulting in 51.6% of Yodas English, 28.8% of Emilia-YODAS English, and 14.6% of Emilia English. Since pre-training uses standard instance sampling and packing, these ratios closely reflect the actual token distribution during training.

C. Evaluation Benchmarks

We evaluate SODA models across four categories: speech semantic knowledge, speech acoustic knowledge, text knowledge, and cross-modal skills. This section describes each benchmark and evaluation protocol.

C.1. Speech Semantic Knowledge

sBLIMP (Nguyen et al., 2020) evaluates grammatical knowledge from speech utterances. The benchmark contains minimal pairs of spoken sentences that differ in grammaticality (e.g., “The cat sleeps” vs. “The cat sleep”). Models assign likelihood to each sentence, and accuracy is computed as the percentage of pairs where the grammatical sentence receives higher likelihood. We use the full sBLIMP test set and report zero-shot accuracy. Random baseline: 50%.

sWUGGY (Nguyen et al., 2020) tests phonotactic knowledge—the ability to distinguish real words from phonologically plausible non-words in speech. Each item contains a real word and a pseudoword (e.g., “oscillation” vs. “odenacia”). Models assign likelihood to each audio token, and accuracy is the percentage where the real word receives higher likelihood. We use the full sWUGGY test set. Random baseline: 50%.

C.2. Speech Acoustic Knowledge

Salmon (Maimon et al., 2025b) evaluates acoustic modeling quality through acoustic consistency tasks. For each test sample, the benchmark compares the model’s likelihood of an original recording against a modified version where an acoustic property (e.g., speaker identity, background noise, or room acoustics) changes mid-utterance. The score is the percentage of samples where the model correctly assigns higher likelihood to the consistent recording. The metric ranges from 0–100%, with higher values indicating better acoustic modeling. Random baseline: 50%.

C.3. Text Knowledge

tBLIMP (Warstadt et al., 2020) is the original text-only version of BLIMP, evaluating grammatical knowledge from written text. The evaluation protocol is identical to sBLIMP but uses text tokens instead of audio tokens. Random baseline: 50%.

tWUGGY (Le Godais et al., 2017) is the original text-only version of WUGGY, testing lexical knowledge from written text. The evaluation protocol is identical to sWUGGY but uses text tokens. Random baseline: 50%.

HellaSwag (Zellers et al., 2019) evaluates commonsense reasoning and world knowledge. Each item presents a context and four possible continuations; models must select the most plausible one. We use the validation set and compute zero-shot accuracy by selecting the continuation with the highest likelihood. Random baseline: 25%.

C.4. Cross-Modal Skills

ASR (Speech-to-Text) measures speech-to-text transcription quality. We evaluate the ASR ability on LibriSpeech (test-clean) (Panayotov et al., 2015), reporting Word Error Rate (WER).

TTS (Text-to-Speech) evaluates zero-shot speech generation quality using the seed-tts-eval benchmark based on its official repo in <https://github.com/BytedanceSpeech/seed-tts-eval>. We measure two metrics: (1) **WER** (intelligibility): generated speech is transcribed with *Whisper-large-v3* and compared against the input text; (2) **SIM** (speaker similarity): cosine similarity between *WavLM-large* speaker embeddings of generated speech and reference audio. We use the English split of Seed-TTS-Eval.

C.5. Validation Loss Metrics

Beyond task-specific benchmarks, in Section 5.1, we compute validation negative log-likelihood (NLL) on the following held-out data. All NLL metrics are computed using the standard cross-entropy loss with the model’s vocabulary. Lower NLL indicates better predictive performance.

LibriSpeech (audio tokens + transcript tokens): We compute NLL on LibriSpeech test-clean, measuring the model’s predictive performance on this English speech corpus. We report NLL on all tokens (semantic + acoustic + text) from audio-first formatted data. Note that Section 5.1 also investigates other options for different NLL metrics.

Paloma-c4en (text tokens): We compute NLL on the C4-en subset of Paloma, a standardized text-only perplexity benchmark. Following the Marin-8B work (Marin Community, 2025), we use Paloma-c4en as our primary validation metric for tracking text modeling progress, as it has been shown to be a reliable indicator of general text understanding capabilities. We report NLL on text tokens only.

D. Design Choice Investigations: Additional Details

This section provides additional details for the design choice investigations in §4.

D.1. SODA-prelim: Preliminary Experiment

Before conducting our design choice investigations, we first validated whether training a decoder-only transformer from scratch on discrete audio tokens yields meaningful capabilities. It was unclear whether semantic and acoustic tokens from Mimi, interleaved at the utterance level with text tokens, could produce coherent speech generation or cross-modal understanding without pre-trained encoders.

Setup. We trained a 600M parameter model on 500B tokens from Yodas2, covering eight languages (English, Spanish, French, German, Thai, Hindi, Arabic, and Chinese). This model, which we call SODA-prelim, uses a Warmup-Stable-Decay learning rate schedule, enabling us to study the impact of different data sources during the annealing (fast cooling-down) phase without full retraining (Blakeney et al., 2024).

Findings. The results (shown in Table 2 in the main text) reveal both strengths and limitations:

- *Acoustic understanding*: SODA-prelim achieves 69.4% on Salmon, outperforming SpiritLM-base (57.2%) and SpiritLM-Expr (67.1%), confirming that joint semantic-acoustic modeling is effective.
- *Cross-modal skills*: SODA-prelim demonstrates functional ASR (22.0% WER zero-shot, improving to 15.2% with 2-shot prompting) and TTS (9.2% WER), validating the unified semantic, acoustic, and text token training.
- *Semantic understanding*: Performance is poor (sBLIMP: 50.9%, sWUGGY: 57.8%), suggesting the model struggles to acquire linguistic structure from audio alone.
- *Text knowledge*: Near random performance (HellaSwag: 26.2%), indicating that the initial recipe fails to transfer general knowledge effectively.

These findings motivated the subsequent design choice investigations in §4: speech data selection (§4.1), text mixture inclusion (§4.2), and token composition ablations (§4.3). SODA-prelim also serves as the base checkpoint for the annealing experiments used in speech data selection.

D.2. Experimental Setup Rationale

The three subsections in §4 use different experimental setups, each tailored to the specific question being addressed:

§4.1 (Speech Data Selection) uses *annealing experiments* (Blakeney et al., 2024), where we branch from a pre-trained checkpoint (SODA-prelim) during the learning rate decay phase to compare data sources. We expect this to be the most reliable approach for comparing data quality, as it isolates the effect of data while keeping all other factors constant. However, it requires a pre-trained checkpoint and is expensive per comparison. We also validate our findings with small-scale runs (150M models, 10B tokens); see §D.3 for full results.

§4.2 (Text Ratio) requires sweeping many ratios (0%–50%) to find the optimal balance. Given that small-scale runs (150M, 10B tokens) yielded similar findings to annealing in our previous experiment in §4.1, we adopt this cheaper setup here, allowing efficient exploration of the ratio space.

§4.3 (Token Composition) ablates token types (semantic-only vs. semantic+acoustic vs. semantic+acoustic+text). Annealing is not applicable here because different token types require different model configurations from the start. With only three configurations to compare, we can afford a larger compute budget per run. We use 1.7B models trained on 30B tokens (3×10^{20} FLOPs), a configuration informed by our concurrent IsoFLOP experiments (§5).

D.3. Speech Data Selection: Full Results

Table 5 presents the full results of our speech data selection experiments from §4.1. We evaluate three speech corpora—MLS, Emilia, and Yodas—using two experimental setups: (1) annealing experiments branching from SODA-prelim, and (2) small-scale validation runs (150M parameters, 10B tokens). MLS performs poorly on cross-modal tasks despite being a curated audiobook corpus, while Emilia and Yodas show complementary strengths. ASR = LibriSpeech-clean-test, TTS=seed-tts-eval.

Data	Semantic		Acoustic		Text		ASR	TTS
	sWUGGY↑	sBLIMP↑	Salmon↑	tWUGGY↑	tBLIMP↑	WER↓		
<i>SODA-prelim</i>								
Yodas (multilingual)	57.8	50.9	69.4	71.3	69.0	22.0	9.2	0.516
<i>Annealing from SODA-prelim (at final 20% cooldown steps)</i>								
MLS (English)	56.8	50.6	70.3	69.1	72.2	92.6	35.7	0.366
Emilia (English)	58.5	51.0	69.8	58.5	70.5	7.8	6.1	0.557
Yodas (English)	57.8	51.0	69.8	73.8	68.6	18.8	12.4	0.502
<i>Small-scale runs (150M, 10B tokens)</i>								
MLS (English)	54.9	49.7	69.5	57.2	65.6	105.6	–	–
Emilia (English)	55.8	49.4	67.9	48.6	67.1	26.8	–	–
Yodas (English)	54.9	49.6	69.4	64.9	60.9	54.1	–	–

Both experimental setups yield consistent conclusions: MLS performs poorly on cross-modal tasks (ASR WER: 92.6% in annealing, 105.6% in small-scale), while Emilia and Yodas show viable performance. The small-scale runs, despite being noisier due to limited compute, support our finding from annealing experiments that MLS should be omitted from the training mixture. We attribute MLS’s poor performance to (i) uncased, unpunctuated transcripts creating distribution mismatch with standard text, and (ii) fixed-length 10–20 second chunks lacking utterance length diversity. Between Emilia and Yodas, Emilia excels at TTS (WER: 6.1%, SIM: 0.557), while Yodas provides better text knowledge (tWUGGY: 73.8%). We therefore combine both in our final recipe.

D.4. Text-Only Data Inclusion: Full Results

Figure 5 shows the full results of the text-injection ratio sweep described in § 4.2. The 3×3 grid displays all evaluated metrics across three categories.

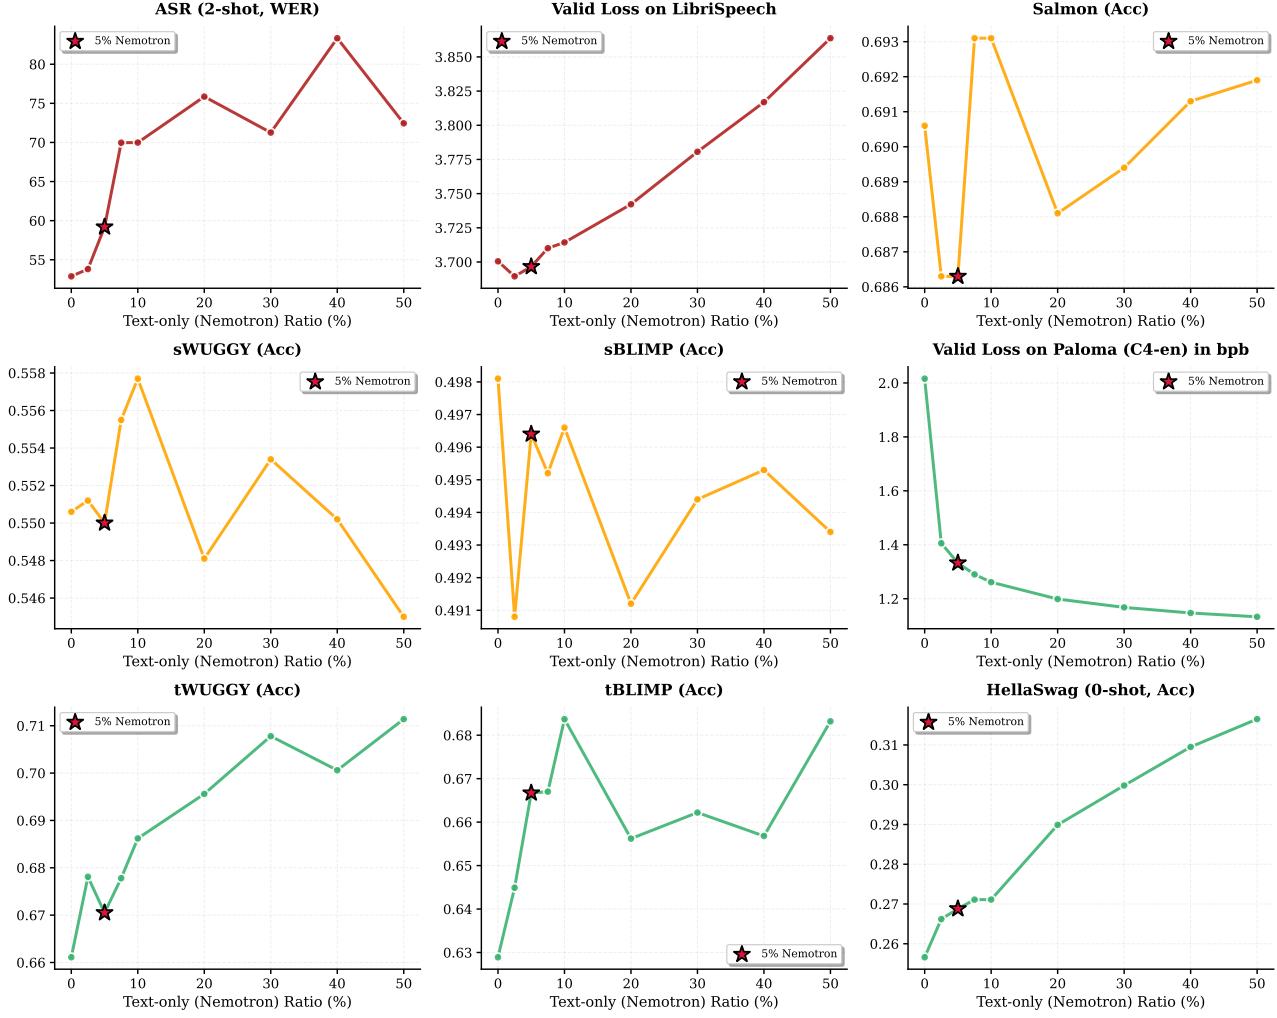


Figure 5. Full results of text-injection ratio sweep (0%–90% Nemotron text data). Colors indicate the trend as text percentage increases: **Red** (Cross-Modal skills: ASR, TTS) shows degradation in performance as text ratio increases beyond 5%. **Orange** (Semantic and Acoustic understanding: sWUGGY, sBLIMP, Salmon) indicates little variation across text ratios, suggesting these capabilities are robust to text inclusion. **Green** (Text Knowledge tasks: tWUGGY, tBLIMP, HellaSwag) shows monotonic improvement as text percentage increases, with notable gains from 0% to 2.5%.

E. NLL Correlation with Downstream Metrics: Full Results

Figures 6 and 7 show the full correlation analysis between validation NLL and all downstream metrics across 64 models from our IsoFLOP study (§ 5.1). Each subplot shows a scatter plot with Spearman rank correlation (ρ) and a fitted trend line.

Choice of NLL Metric. Our interleaved training format (semantic, acoustic, and text tokens in audio-first or text-first order) admits multiple ways to compute validation NLL. In Figures 6 and 7, different rows correspond to the following NLL metrics: (1) all tokens from audio+text data, (2) all audio tokens from text+audio data, (3) semantic tokens from text+audio data, (4) audio tokens from audio-only data, (5) text tokens from audio+text data, (6) text tokens from text-only data.

Table 6 quantifies the Spearman correlations (absolute values) between each NLL metric and downstream tasks. While all variants show similar correlations, we select **NLL on all tokens from audio+text data (Option 1)** as our primary metric because it provides the best balance: it achieves near-best correlation with speech tasks (Avg_Speech = 0.895, only 0.001 behind the best) while maintaining strong correlation with text tasks (Avg_Text = 0.864). In contrast, Option 4 (audio-only NLL) is marginally better for speech (0.896) but worse for text (0.859), while Option 6 (text-only NLL) is best for text (0.894) but substantially worse for speech (0.848). For a general-purpose model targeting both modalities, Option 1 offers the most balanced proxy.

Table 6. Spearman correlation (absolute values) between NLL metrics and downstream tasks across 64 IsoFLOP models. Avg_Speech averages 7 speech metrics; Avg_Text averages 3 text metrics. Option 1 achieves the best balance between speech and text correlations.

NLL	ASR-0	ASR-2	TTS-W	TTS-S	Sal.	sWUG	sBLI	tWUG	tBLI	Hella	Avg-Sp	Avg-Tx
NLL_1	0.959	0.967	0.955	0.989	0.841	0.941	0.614	0.879	0.824	0.887	0.895	0.864
NLL_2	0.956	0.962	0.953	0.986	0.842	0.936	0.621	0.874	0.821	0.887	0.894	0.861
NLL_3	0.954	0.959	0.952	0.985	0.842	0.929	0.620	0.866	0.819	0.884	0.892	0.856
NLL_4	0.959	0.965	0.952	0.987	0.845	0.937	0.623	0.871	0.819	0.886	0.896	0.859
NLL_5	0.962	0.961	0.960	0.986	0.820	0.944	0.602	0.899	0.827	0.904	0.891	0.877
NLL_6	0.903	0.908	0.933	0.949	0.802	0.892	0.552	0.891	0.873	0.917	0.848	0.894

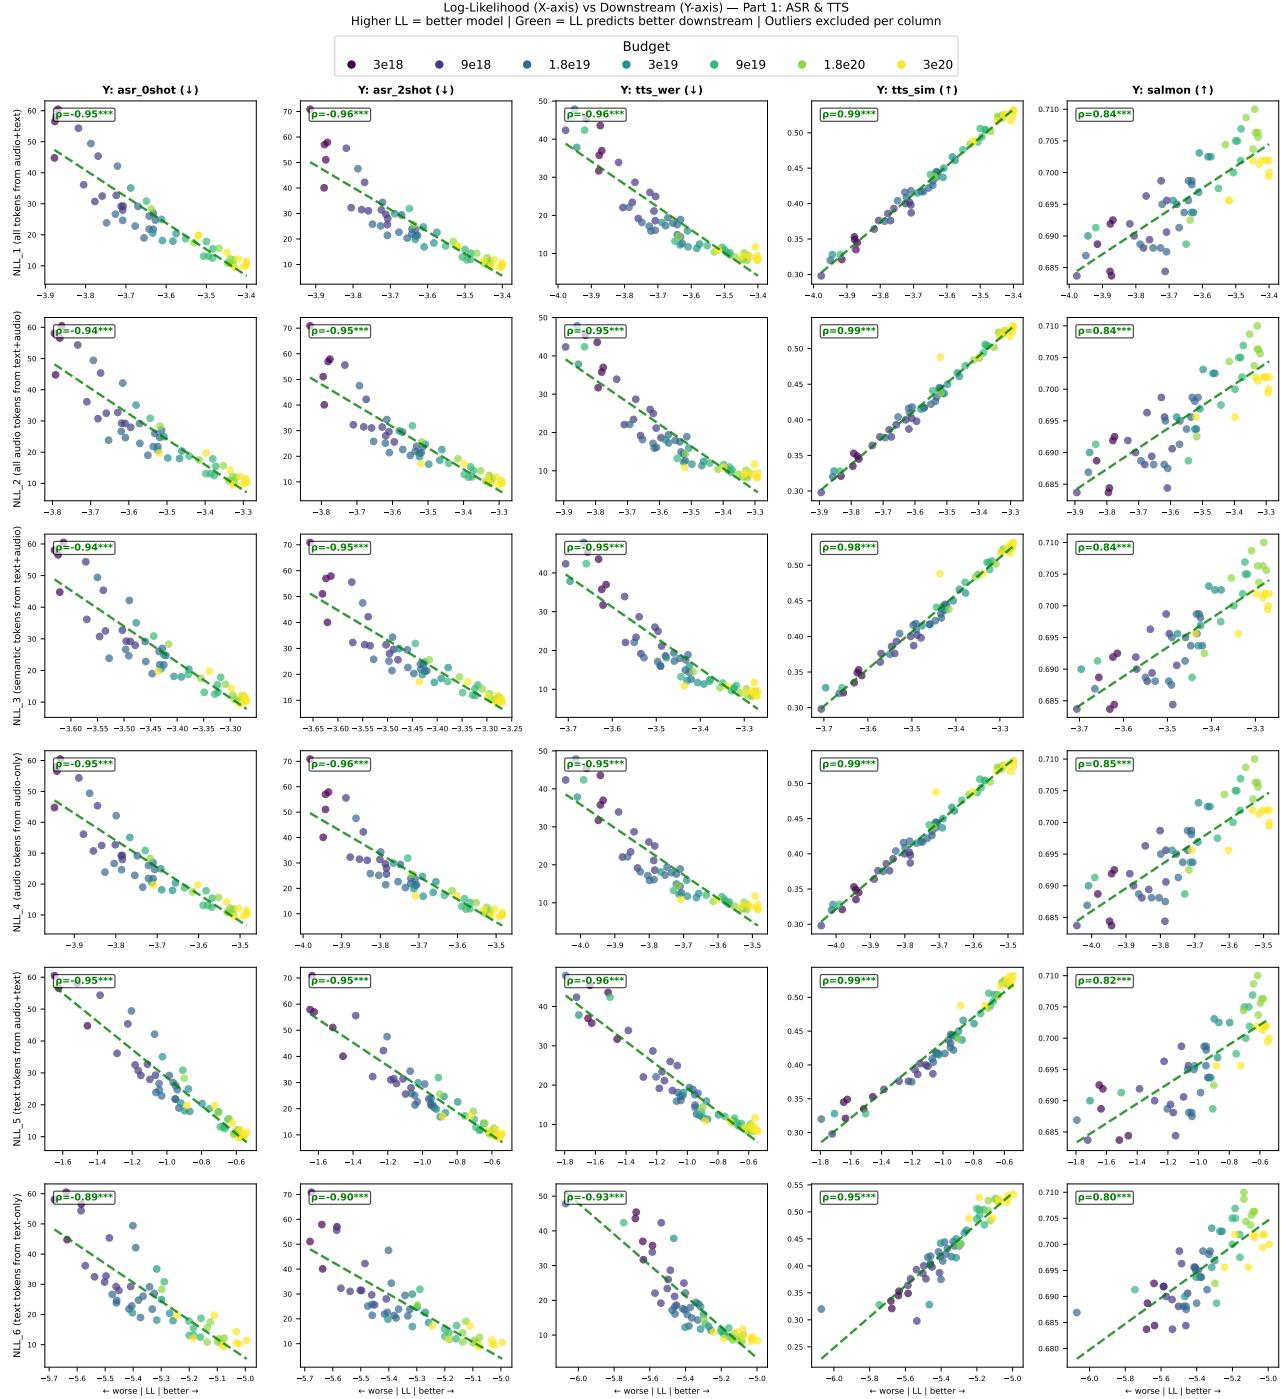


Figure 6. Validation NLL (audio-first) versus downstream task performance (Part 1). Each subplot shows one metric across 64 models from the IsoFLOP study, with Spearman ρ reported. Cross-modal metrics (ASR, TTS) show strong correlation with NLL.

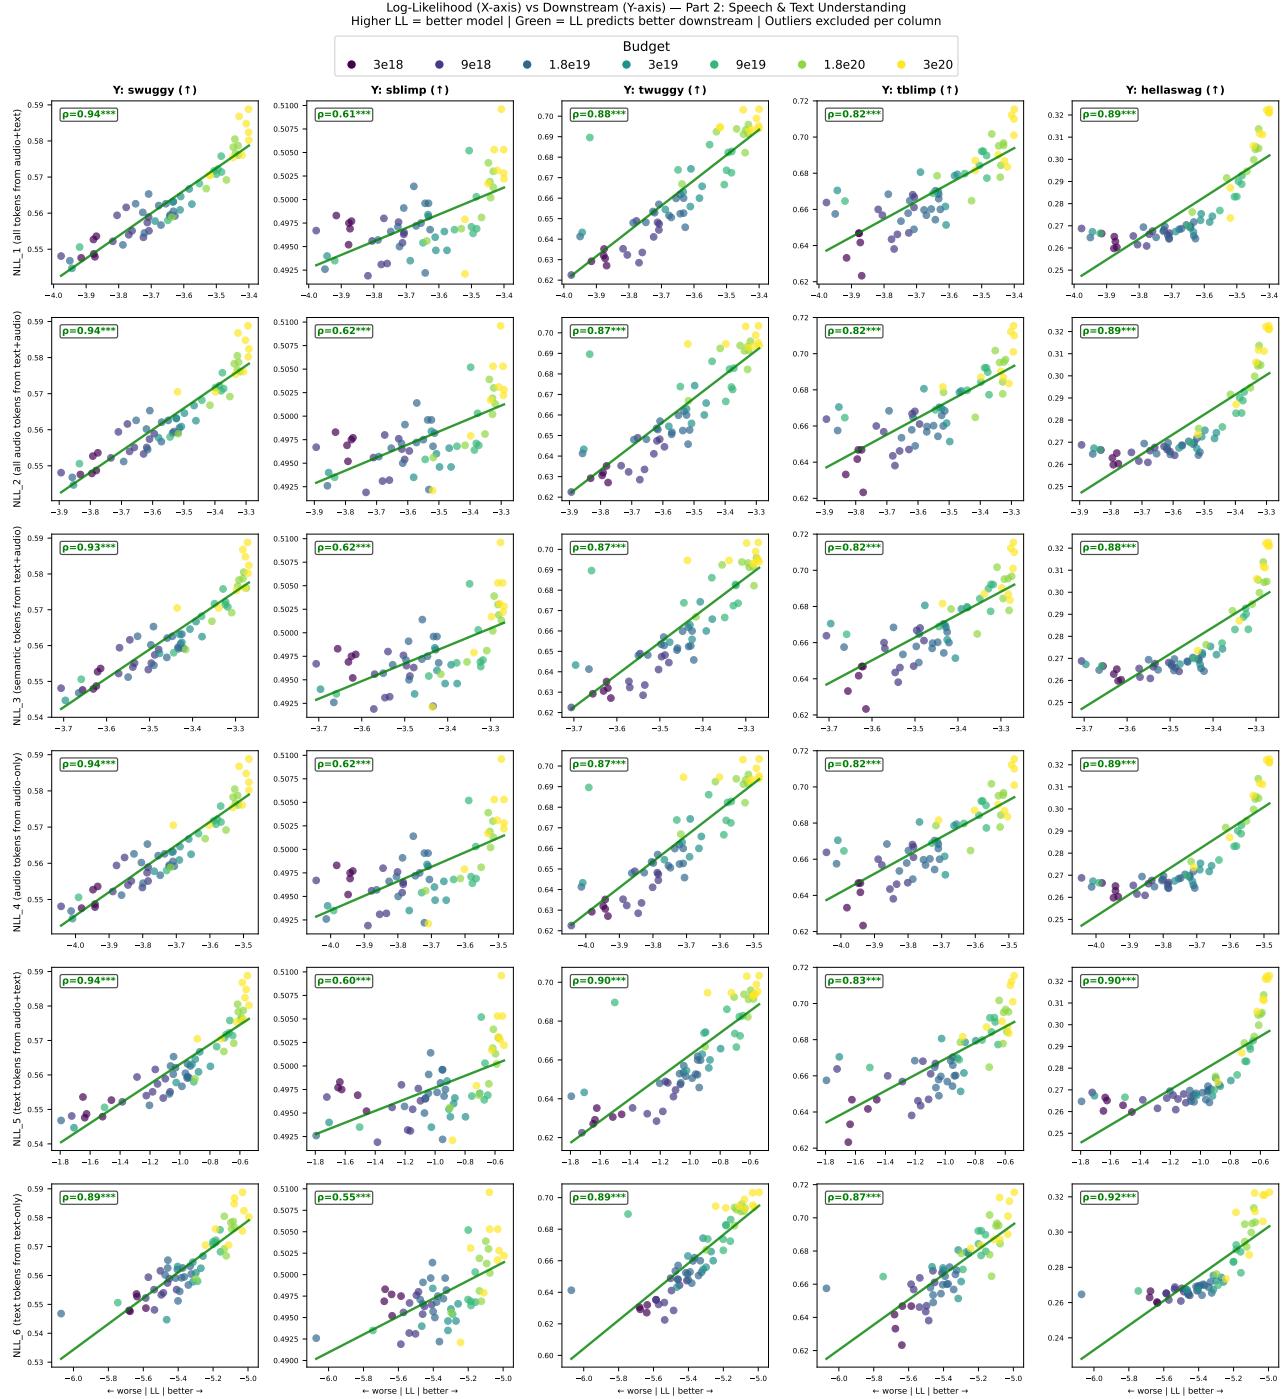


Figure 7. Validation NLL (audio-first) versus downstream task performance (Part 2). Speech semantic and acoustic metrics show continued improvement without saturation. Text knowledge tasks (tWUGGY, HellaSwag) show emergence patterns with accelerated improvement at lower NLL values.

F. Validation Loss Prediction for Over-Trained Models

This section analyzes whether our scaling laws can predict the validation loss of the final SODA models, which are heavily over-trained relative to the compute-optimal frontier.

Table 7. Training regime for SODA models. D^* is the compute-optimal token count predicted by our scaling laws (§5.2). Over-training factor indicates how many times more tokens we train on compared to optimal.

Model Size	D^*	Over-training Factor
135M	0.53B	940×
600M	5.55B	90× (~ Llama3's)
1.7B	28.6B	18× (~ Llama2's)
4B	110B	4.5×

Table 7 summarizes the over-training regime for SODA models. Smaller models are heavily over-trained (e.g., 135M at $940 \times D^*$, comparable to aggressive over-training in recent LLM practice), while the 4B model approaches the compute-optimal frontier.

F.1. Compute-Optimal Frontier Extrapolation

Figure 8 compares over-trained SODA models against the extrapolated compute-optimal frontier. We fit a power-law regression $L^* = a \cdot C^b$ on the seven compute-optimal points from our IsoFLOP study and extrapolate to higher compute budgets. All SODA models achieve higher loss than the optimal extrapolated prediction at equivalent compute, with smaller models deviating more due to heavier over-training (e.g., 135M at $940 \times D^*$). This is expected: over-trained models use smaller-than-optimal model sizes for their compute budget, trading off loss for inference efficiency.

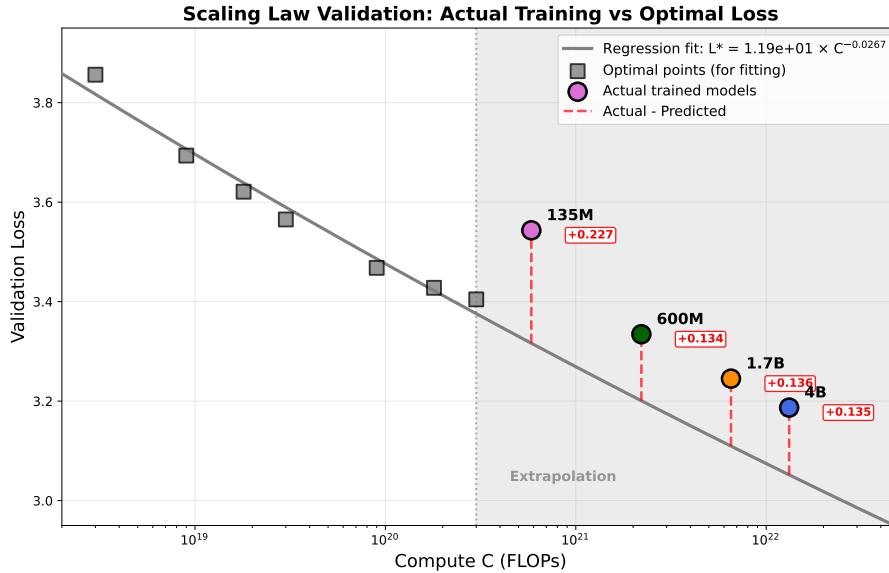


Figure 8. Validation loss vs. compute for SODA models. Grey squares show compute-optimal models from §5. We fit a power-law regression $L^* = a \cdot C^b$ on these seven points and extrapolate to higher compute (grey line). Colored circles show final SODA runs.

F.2. Parametric Loss Prediction

An alternative approach is to fit a parametric equation that captures loss as a function of both model size N and training tokens D , allowing us to directly predict loss for over-trained configurations. We fit the parametric equation from Chinchilla (Hoffmann et al., 2022):

$$L(N, D) = E + \frac{A}{N^\alpha} + \frac{B}{D^\beta} \quad (4)$$

where N is the model size (parameters), D is the number of training tokens, and E represents the irreducible loss. We fit this equation on all 64 data points from our IsoFLOP sweep (not just the compute-optimal points), obtaining:

$$L = 3.169 + \frac{215886}{N^{0.684}} + \frac{4750}{D^{0.439}} \quad (R^2 = 0.983) \quad (5)$$

Derived Scaling Exponents. The parametric fit also provides an alternative derivation of scaling exponents. From the Chinchilla derivation, the compute-optimal model size and data scale as $N^* \propto C^{\beta}/(\alpha+\beta)$ and $D^* \propto C^{\alpha}/(\alpha+\beta)$. Using our fitted $\alpha = 0.684$ and $\beta = 0.439$, we obtain:

$$N^* \propto C^{0.391} \quad (\text{vs. } C^{0.367} \text{ in IsoFLOP fit in Eqn 1 in §5.2}) \quad (6)$$

$$D^* \propto C^{0.609} \quad (\text{vs. } C^{0.579} \text{ in IsoFLOP fit in Eqn 2 in §5.2}) \quad (7)$$

These derived exponents are reasonably close to the empirically fitted values from §5.2, with both approaches showing $D^*/N^* \approx 1.57$. This asymmetry (data scaling faster than model size) is consistent with the lower information density of discrete audio tokens compared to text.

Loss Predictions. Figure 9 shows predictions at different overtraining ratios $K = D/D^*$, where $K = 1$ is compute-optimal and higher K indicates more overtraining (smaller model trained on more data). The actual SODA models (135M–4B) achieve losses *lower* than predicted, even lower than the optimal extrapolated prediction for the 1.7B and 4B models. This systematic under-prediction when extrapolating from the IsoFLOP regime (3×10^{18} – 3×10^{20} FLOPs) to higher compute (6×10^{20} – 1.3×10^{22} FLOPs) could be due to scaling law extrapolations diverging in the far tails of the compute distribution. The offset suggests that either scaling exponents evolve favorably at larger scales, or the irreducible loss E is lower than estimated from the smaller-scale regime.

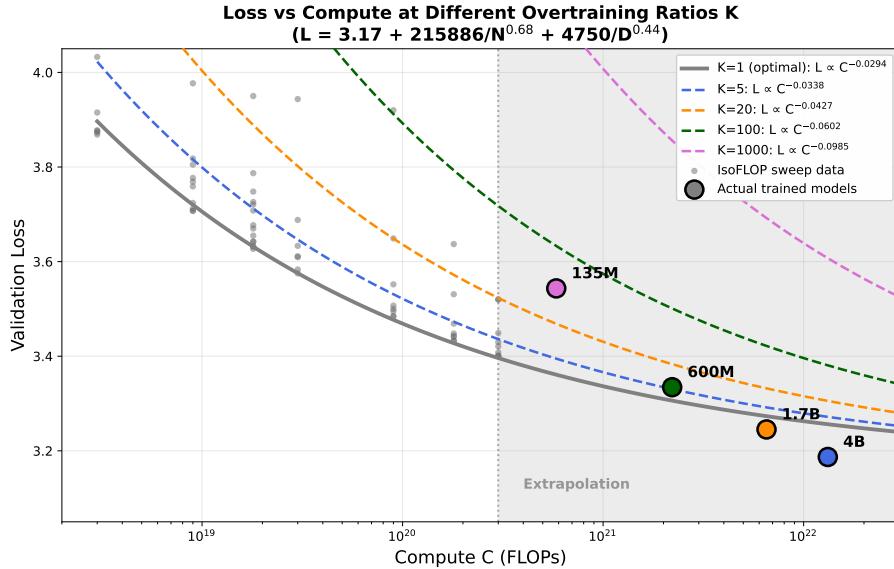


Figure 9. Loss vs. compute at different overtraining ratios K using the fitted parametric equation. Grey points show IsoFLOP sweep data. Colored curves show predicted loss at $K = 1$ (optimal), $K = 5$, $K = 20$, and $K = 100$. Colored circles show actual SODA models (135M, 600M, 1.7B, 4B), which achieve lower losses than predicted. The discrepancy shows that SODA models outperform the parametric prediction, suggesting favorable scaling dynamics or lower irreducible loss at larger scales.

F.3. Summary

Neither the compute-optimal extrapolation nor the parametric fit very accurately predicts validation loss for our final SODA runs. The compute-optimal approach predicts lower losses than observed (expected, since over-trained models are suboptimal), while the parametric approach predicts higher losses than observed. This discrepancy suggests that scaling behavior may evolve favorably beyond the IsoFLOP regime, or that our smaller-scale fits do not fully capture dynamics at larger compute budgets. Despite these prediction challenges, we note that downstream task performance scales reliably with NLL (see Figure 3 and §6), making NLL a useful proxy even when absolute loss predictions are imperfect.

G. Warm-Start vs. Cold-Start: Training Trajectories

Table 8 presents the final evaluation results, and Figure 11 shows the full training trajectories comparing Warm-Start (initialized from Qwen3) versus Cold-Start (trained from scratch) at 600M and 1.7B scales. Each subplot tracks a different metric across 500B tokens of training. Figure 10 compares training loss stability between the two initialization strategies.

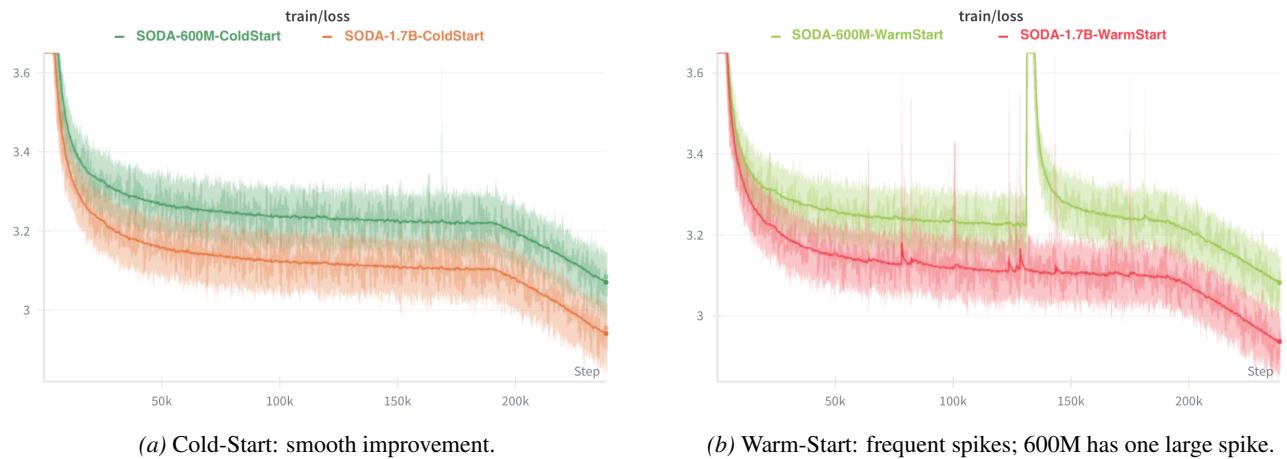


Figure 10. Training loss curve (extracted from W&B) comparison between Cold-Start and Warm-Start at 600M and 1.7B scales. Warm-Start exhibits instability with many small spikes throughout training for both model sizes. The 600M Warm-Start run shows a particularly large spike around 135K steps. We note that with our limited compute budget, we are able to train only one run for each model configuration. However, we note that better regularization or different hyperparameters could stabilize the training process of Warm-Start.

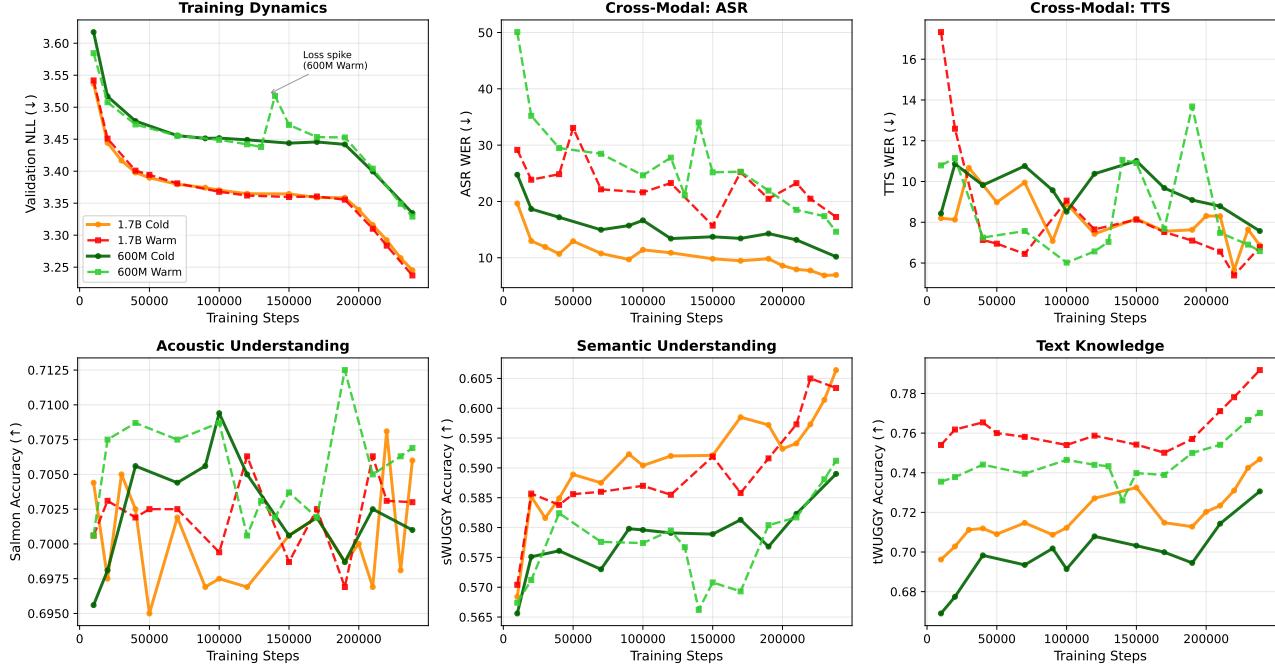


Figure 11. Warm-Start vs. Cold-Start training comparison at 600M and 1.7B scales. **Top row:** Training dynamics (NLL) and cross-modal skills (ASR, TTS). Cold-Start achieves lower ASR-WER throughout training, while TTS converges similarly. The NLL plot highlights a loss spike at 135K steps for the 600M Warm-Start run. **Bottom row:** Understanding tasks. Salmon (acoustic) and sWUGGY (semantic) show similar trajectories regardless of initialization, suggesting these are learned fresh. tWUGGY (text knowledge) shows Warm-Start maintaining a persistent advantage throughout training, with Cold-Start never catching up.

Table 8. Warm-Start vs. Cold-Start training comparison at 600M and 1.7B scales. Final evaluation results are shown.

Model	Speech (Semantic)		Speech (Acoustic)		Text (Knowledge)			Cross-Modal		
	sBLIMP↑	sWUGGY↑	Salmon↑	tBLIMP↑	tWUGGY↑	HellaS↑	ASR↓	TTSWER↓	TTS _{SIM} ↑	
SODA-600M (cold-start)	51.2	58.9	70.1	70.7	73.1	35.8	10.2	7.6	0.555	
SODA-600M (warm-start)	51.1	59.1	70.7	70.8	77.0	36.3	14.6	6.6	0.559	
SODA-1.7B (cold-start)	51.4	60.6	70.6	70.3	74.7	44.5	7.0	6.9	0.560	
SODA-1.7B (warm-start)	51.8	60.3	70.3	71.0	79.2	47.1	17.3	6.8	0.557	

H. Speech-to-Speech Translation: Dataset and Fine-tuning Details

This section provides additional details for the S2ST experiment in §6.4.

Dataset: CVSS-T. We use CVSS-T (Jia et al., 2022b), a multilingual speech-to-speech translation corpus derived from CoVoST 2. CVSS-T contains ~ 1.9 K hours of paired speech across 21 source languages translating to English. Crucially, the target English speech is synthesized to *preserve the source speaker’s voice* using a voice-cloning TTS system, making this a voice-preserving S2ST benchmark. The training set contains ~ 557 K examples.

Evaluation. Due to the computational cost of evaluating the full CVSS-T test set (~ 84 K examples), we subsample 800 examples: 200 each from Spanish (es), French (fr), German (de), and 200 from the remaining 18 languages (“other”). We report: (1) **ASR-BLEU**: BLEU score computed on ASR-transcribed model outputs against reference translations; (2) **SIM**: Speaker similarity between the generated and source speech. For both metrics, we use the same evaluation setup as seed-tts-eval (as used in our TTS evaluation), with Whisper-large-v3 for ASR and WavLM-large for speaker similarity.

Comparison Context. Direct numerical comparison with prior work is difficult due to different evaluation splits. However, for context: Unity (Inaguma et al., 2023) reports Es→En BLEU of 18.2 (S2TT+TTS cascade), Translatotron 2 (Jia et al., 2022a) achieves 25.4 (task-specific architecture). On voice preservation, Hibiki (Labiausse et al., 2025) achieves SIM ≈ 0.41 and Seamless achieves SIM ≈ 0.30 . Our SODA models achieve SIM ≈ 0.47 , competitive with or exceeding these task-specific systems.

Fine-tuning Hyperparameters. We fine-tune on the full CVSS-T training set (557K examples) for 5 epochs using:

- Learning rate: cosine schedule with peak 2×10^{-5}
- Batch size: 64
- Sequence length: 4096 tokens
- The same optimizer settings as pre-training (Appendix A.3)

Data Formatting. We format S2ST as interleaved next-token prediction: [source audio tokens] → [source text] → [target text] → [target audio tokens]. This is the same format used for ASR and TTS, with the addition of source→target text translation in between.