

# ProsodyFM: Unsupervised Phrasing and Intonation Control for Intelligible Speech Synthesis

Xiangheng He<sup>1</sup>, Junjie Chen<sup>3</sup>, Zixing Zhang<sup>4\*</sup>, Björn Schuller<sup>1,2</sup>

<sup>1</sup>GLAM – Group on Language, Audio, & Music, Imperial College London, UK

<sup>2</sup>CHI – Chair of Health Informatics, MRI, Technical University of Munich, Germany

<sup>3</sup>Department of Computer Science, The University of Tokyo, Japan

<sup>4</sup>College of Computer Science and Electronic Engineering, Hunan University, China

{x.he20, bjoern.schuller}@imperial.ac.uk, junjiechen@g.ecc.u-tokyo.ac.jp, zixingzhang@hnu.edu.cn

## Abstract

Prosody contains rich information beyond the literal meaning of words, which is crucial for the intelligibility of speech. Current models still fall short in phrasing and intonation; they not only miss or misplace breaks when synthesizing long sentences with complex structures but also produce unnatural intonation. We propose ProsodyFM, a prosody-aware text-to-speech synthesis (TTS) model with a flow-matching (FM) backbone that aims to enhance the phrasing and intonation aspects of prosody. ProsodyFM introduces two key components: a Phrase Break Encoder to capture initial phrase break locations, followed by a Duration Predictor for the flexible adjustment of break durations; and a Terminal Intonation Encoder which learns a bank of intonation shape tokens combined with a novel Pitch Processor for more robust modeling of human-perceived intonation change. ProsodyFM is trained with no explicit prosodic labels and yet can uncover a broad spectrum of break durations and intonation patterns. Experimental results demonstrate that ProsodyFM can effectively improve the phrasing and intonation aspects of prosody, thereby enhancing the overall intelligibility compared to four state-of-the-art (SOTA) models. Out-of-distribution experiments show that this prosody improvement can further bring ProsodyFM superior generalizability for unseen complex sentences and speakers. Our case study intuitively illustrates the powerful and fine-grained controllability of ProsodyFM over phrasing and intonation.

## Code and demo —

<https://github.com/XianghengHee/ProsodyFM>

## Extended version with Appendix —

<https://arxiv.org/abs/2412.11795>

## Introduction

Prosody, which encompasses various properties of speech such as phrasing, intonation, prominence, and rhythm, can convey rich information beyond the literal meaning of words (Xu 2019). It plays a crucial role in the intelligibility of speech. Although recent TTS models have achieved great progress in synthesizing intelligible speech, they still lack

in many prosody aspects. In this study, we focus on two prosody aspects in English: phrasing and intonation.

Phrasing refers to grouping words into chunks. An intonational phrase contains a chunk of words with their own intonation pattern. **Phrase break** in this paper refers to the perceivable acoustic pause at the end of intonational phrases. Phrase break plays an important role in enhancing speech intelligibility (Futamura et al. 2021). It implies the phrasal organization in the sentence, allowing listeners to accurately discern the syntactic structure of the sentence and deduce its correct meaning. For example, the sentence “I saw the man with the telescope.” can be interpreted differently depending on whether there is a break after the word “man”. When the sentence is spoken without the break, “with the telescope” modifies “the man”, suggesting that the man observed by the speaker had a telescope. When the break is introduced after “man”, it implies the speaker used a telescope to see the man. This demonstrates that incorrect phrasing can lead to incorrect interpretation of the sentence, thereby impairing speech intelligibility. However, due to the difficulty in obtaining break labels and the variability of break duration, current TTS systems usually miss or misplace breaks when synthesizing complex sentences.

Intonation, especially terminal intonation, is essential for synthesizing intelligible speech. We refer to the intonation pattern of the last word in an intonational phrase as the **terminal intonation**. Terminal intonation carries many linguistic and paralinguistic information in English. A rising terminal intonation at the end of a sentence usually signals uncertainty or a request for clarification, while a falling intonation typically indicates certainty or is used to make statements and assertions (Lieberman 1975). A rising terminal intonation in the middle of a sentence usually indicates the speaker is not finished yet, while a falling tone indicates the end of the thought (Bolinger 1998). This information of intonation change can be represented through the change in the pitch contour (Cole, Steffman, and Tilsen 2022). However, instead of modeling the relative change in the pitch contour, previous TTS systems directly model the absolute pitch value. This design choice hinders their ability to accurately capture natural intonation, as pitch tracking and predicting absolute pitch values is inherently challenging.

To address these issues, we propose ProsodyFM, a novel

\*Corresponding author  
Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Prosody-aware TTS model based on a Flow-Matching (FM) backbone that enhances both phrasing and intonation aspects of prosody in an unsupervised manner, resulting in more intelligible synthesized speech. For the break labeling issue, we introduce a Phrase Break Encoder to capture initial break locations, followed by a Duration Predictor to adjust break durations, enabling flexible and accurate modeling of phrase breaks. For the intonation modeling issue, we employ a novel Pitch Processor and learn a bank of intonation shape tokens, which effectively mitigates pitch tracking errors, enables more robust modeling of pitch shapes, and aligns more closely with human perception of intonation changes. ProsodyFM is trained without any prosodic labels and yet can uncover a wide range of break durations and intonation patterns. The main contributions of this paper are as follows:

- We propose ProsodyFM, a prosody-aware TTS model with strong generalizability and fine-grained prosody control, capable of synthesizing speech with natural phrasing and intonation, leading to greater intelligibility than existing systems.
- We provide novel and effective solutions for the break labeling issue and the intonation modeling issue.
- We release our demo, code, and model checkpoints to facilitate further research.

## Related Works

### The Break Labeling Issue

Breaks in speech can be roughly divided into punctuation-based and respiratory breaks (Hwang, Lee, and Lee 2023). Unlike punctuation-based breaks which are marked by punctuations, respiratory breaks have no explicit label on the text side. Most of the current TTS systems (Mehta et al. 2024; Li et al. 2024) have only considered punctuation-based breaks, resulting in many non-final phrase breaks being overlooked or misplaced (Taylor 2009). Some TTS systems model phrase breaks explicitly. These models (Hwang, Lee, and Lee 2023; Abbas et al. 2022; Yang et al. 2023) use manually designed thresholds combined with the Montreal Forced Aligner (MFA) (McAuliffe et al. 2017) to obtain break labels in an unsupervised manner. The frequency and duration of the phrase break are shaped by both the linguistic phrase structure and a speaker’s speaking style (Hwang, Lee, and Lee 2023). However, due to the variability of break duration and its dependence on speaker information, the handcraft threshold-based methods can hardly account for speaker-specific variations in break durations.

ProsodyFM tackles this issue by designing a Fusion Encoder to integrate initial break locations obtained from the Phrase Break Encoder with speaker information, and then adjusting the break durations with a Duration Predictor, enabling flexible modeling of phrase breaks.

### The Intonation Modeling Issue

Annotating intonation pattern labels is a high-cost task and often yields unreliable results (Lee and Kim 2019) due to the complexity of current annotation systems (Silverman et al.




Figure 1: Pitch contours extracted from 5 pitch tracking methods (blue) and our pitch smoothing method (orange).

1992). Almost all the existing intonation-aware TTS systems (Ren et al. 2021; Min et al. 2021; Huang et al. 2022; Li et al. 2024) directly model the absolute pitch values obtained from some pitch tracking methods. However, pitch tracking is inherently challenging, and existing methods frequently yield errors like pitch doubling/halving and incorrect unvoiced/voiced flags (Hirst and de Looze 2021), leading to unreliable results. Figure 1 illustrates the pitch tracking results across 5 different methods. From top to bottom are Harvest (Morise 2017), DIO (Morise, Kawahara, and Katayose 2009), SWIPE (Camacho and Harris 2008), pYIN (Mauch and Dixon 2014), Praat (Boersma 2001) and Praat with smoothing. We can clearly observe frequent prediction errors in pitch values and unvoiced/voiced flags, as well as inconsistencies across these five methods. Some recent findings from human perceptual studies offer a potential basis for this issue; the authors in (Chodroff and Cole 2019; Cole, Steffman, and Tilsen 2022) have shown that **compared to the detailed pitch values, the shape of the pitch contour is more important for human perception of intonation change.**

ProsodyFM introduces a novel Pitch Processor that interpolates, smooths, and perturbs raw pitch values to highlight their shape, and subsequently learns a set of intonation shape tokens to model perceptually aligned intonation change instead of directly modeling absolute pitch values. The orange line in Figure 1 shows an example after our smoothing process. Our method alleviates pitch tracking errors, enables more robust modeling of pitch shapes, and aligns more closely with human perception of intonation change.

## Method

ProsodyFM is designed to extract phrasing and terminal intonation patterns from reference speech and adjust these patterns to match the target text. Following the MatchaTTS (Mehta et al. 2024) backbone, ProsodyFM is trained using the Optimal-Transport Conditional Flow Matching (OT-CFM) (Lipman et al. 2023). The formulation and training algorithm of ProsodyFM can be found in Appendix A of the extended version. ProsodyFM predicts Mel-spectrograms

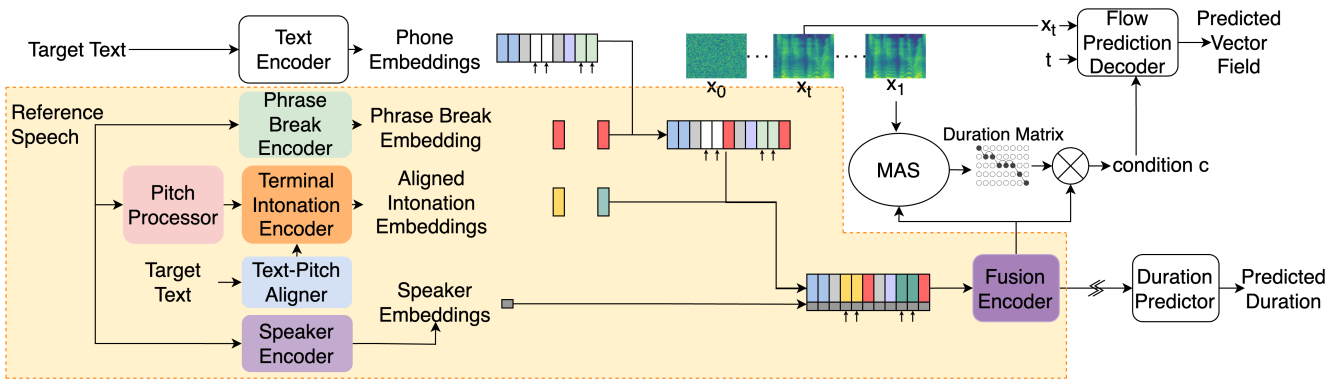


Figure 2: The model architecture of the proposed ProsodyFM during training. The components outlined by the yellow shaded area are unique to ProsodyFM and differ from those in MatchaTTS.

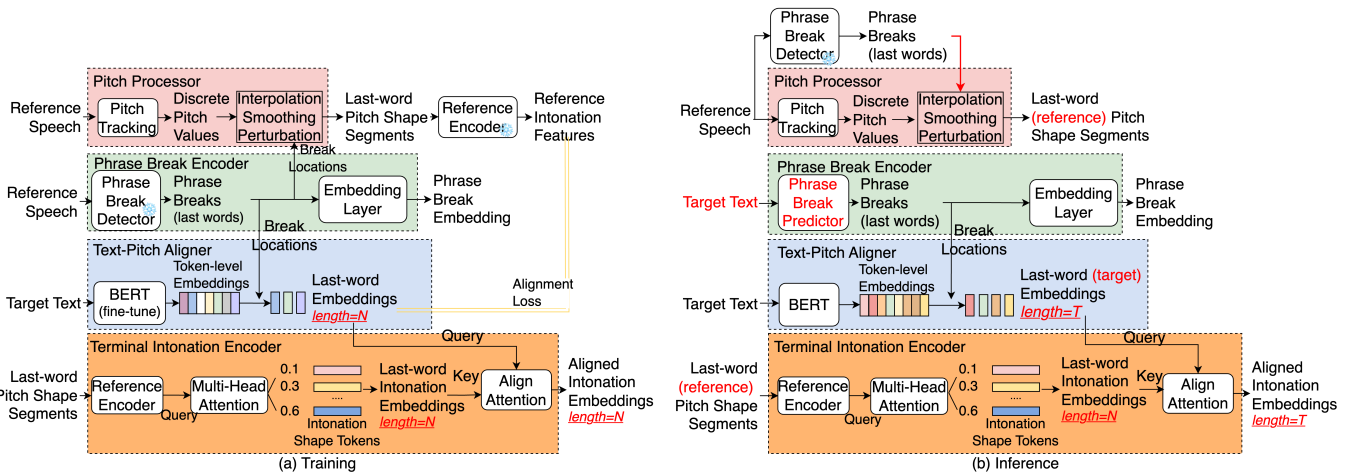


Figure 3: The key components of the proposed ProsodyFM in the training (a) and inference (b) phrases. The red markings highlight the differences. The snowflake mark means the module is frozen during training.

from raw text, which are then converted to waveforms using the HifiGAN vocoder (Kong, Kim, and Bae 2020).

The given target text aligns with the reference speech during training but may differ during inference. During training, the reference speech serves as the ground truth and the target text matches its transcript, while during inference, the target text may not match the transcript of the reference speech.

Figure 2 illustrates the overall structure of ProsodyFM, highlighting the proposed components within the yellow-shaded area. Details of four key components are presented in Figure 3: (1) the Pitch Processor extracts robust pitch shape segments of the last words; (2) the Phrase Break Encoder predicts initial phrase break locations, which are then combined with speaker information and refined for duration by the Duration Predictor; (3) the Text-Pitch Aligner estimates intonation patterns from the target text to guide the selection of reference intonation patterns; and (4) the Terminal Intonation Encoder models terminal intonation patterns that are properly aligned with the target text.

## Pitch Processor

The Pitch Processor (pink box in Figure 3) extracts robust pitch shape segments of the last words through three operations: interpolation, smoothing, and perturbation. First, it interpolates and smooths the discrete, unreliable raw pitch values from pitch tracking into continuous contours. Then, to emphasize pitch shape over absolute values, it subtracts a random offset (uniformly sampled from  $[f_{min}, f_{max}]$ ) from each contour point, preserving the shape patterns while perturbing its specific value information.

## Phrase Break Encoder

The Phrase Break Encoder (green box in Figure 3) predicts where phrase breaks occur, thus allowing it to locate the last word of each intonational phrase. These last-word locations guide the Pitch Processor and the Text-Pitch Aligner in selecting the corresponding pitch shape segments and word embeddings.

During training, the Phrase Break Encoder uses a pre-trained, frozen Phrase Break Detector to identify phrase breaks from reference speech. During inference, when no

aligned reference speech is available, the Phrase Break Encoder relies on a Phrase Break Predictor fine-tuned from T5 (Ni et al. 2022) to infer breaks directly from plain target text. The performance of this Phrase Break Predictor is reported in Appendix B of the extended version.

### Text-Pitch Aligner

The Text-Pitch Aligner (blue box in Figure 3) predicts intonation patterns from the target text, even without matched speech during inference. We fine-tune BERT by minimizing the L2 loss between BERT-derived word embeddings and the reference intonation features extracted by the Reference Encoder. The Reference Encoder is identical to the one in the Terminal Intonation Encoder, but detached to prevent gradient flow. The predicted BERT embeddings then guide the selection of suitable reference intonation patterns in the Terminal Intonation Encoder.

### Terminal Intonation Encoder

The Terminal Intonation Encoder (orange box in Figure 3) extracts the terminal intonation patterns that are aligned with the target text. The Reference Encoder compresses the pitch shape segments of the last word in the reference speech into a fixed-length intonation feature, used as the query for the Multi-head Attention module. This attention module learns a similarity measure between the reference intonation features and a bank of **intonation shape tokens**. These tokens serve as a learnable codebook designed to capture and represent various intonation patterns. Trained with OT-CFM loss alone, these tokens require no annotated intonation labels. The Multi-head Attention module generates weights for these tokens, and their weighted sum forms the last-word intonation embedding of the reference speech.

However, during inference, the reference speech may not be aligned with the target text, resulting in a different number of last words in the reference speech compared to the target text. We use scaled dot-product attention (Align Attention module) to select the terminal intonation patterns from the reference speech that best suit the target text. Specifically, we treat the last-word intonation embeddings (of the reference speech) as the key (and value) and the last-word embeddings (of the target text) as the query. This alignment enables ProsodyFM to autonomously choose the terminal intonation pattern based on both the reference speech and the target text during the inference phase.

### Mel-spectrogram Generation

During inference, the Fusion Encoder combines the phrase break and aligned intonation embeddings with speaker and phone embeddings to produce phone-level prior statistics. The Duration Predictor (instead of the MAS during training) then determines the optimal durations of each phone and phrase break to obtain the frame-level condition  $c$ . Given  $c$ , a sampled time  $t$ , and  $x_t$ , the Flow Prediction Decoder predicts the target vector field. Finally, the ODE solver uses this predicted vector field to generate the Mel-spectrogram.

## Experimental Details

**Model Configurations** For a fair comparison, we utilize the same model architecture and hyperparameters as MatchaTTS (Mehta et al. 2024) except for the following modules. For our Terminal Intonation Encoder, we employ the attention module in (Wang et al. 2018) with 4 attention heads and 6 64-D tokens. We replace the complex reference encoder in (Wang et al. 2018) with a single-layer LSTM with 128-D hidden size to speed up training. For the Phrase Break Detector, we use the released checkpoint of PSST (Roll, Graham, and Todd 2023) without fine-tuning. For the Phrase Break Predictor, we fine-tune T5 (Ni et al. 2022) independent from ProsodyFM using LoRA (Hu et al. 2022) with 16 ranks and consider the phrase breaks obtained from the PSST as the ground truth labels when fine-tuning. For the Text-Pitch Aligner, we initialize BERT<sup>1</sup> with pre-trained weights, using its original tokenizer to process input text and obtain 768-D token-level embeddings. We then select the tokens corresponding to each last word, average their embeddings, and pass it through a fully connected layer to produce a final 192-D embedding for each last word. For the Pitch Processor, we use Praat (Boersma 2001) to extract discrete pitch values and use a customized Praat script modified from (Cangemi 2015) to interpolate and smooth them into a continuous pitch contour. For the Speaker Encoder, we extract the same external speaker embedding as in (Casanova et al. 2022) for each speech sample and add two fully connected layers to transform the 512-D d-vector to the final 64-D speaker embeddings.

ProsodyFM and its ablated variants in Table 4 are trained for 350 epochs on an NVIDIA A100 GPU with batch size 64 and learning rate  $1e-4$ .

**Datasets** We perform the experiments in Table 1, Table 2, Table 4, and Figure 4 on the LibriTTS corpus (Zen et al. 2019). We randomly split (speakers-independent) the audio samples in the train-clean-100, dev-clean, and test-clean sections of LibriTTS into 40421, 839, and 839 samples for our training, validation, and testing sets, respectively. The whole dataset has in total 71 hours of audio signals and 326 speakers. For the experiments in Table 3, we train the models on the VCTK corpus (Yamagishi, Veaux, and MacDonald 2019) with the same training set as in (Kim, Kong, and Son 2021) and test the models on our LibriTTS testing set. We re-sample all audio to 22050 Hz and extract Mel-spectrograms with a 1024 FFT size, 256 hop size, 1024 window length, and 80 frequency bins.

**Objective Evaluation Metrics** We conduct objective evaluations with three metrics. The implementation details of these metrics can be found in the extended version.

1) The log-scale F0 Root Mean Squared Error ( $RMSE_{f_0}$ ) measures the pitch error. Following (Birkholz and Zhang 2020), we use it to evaluate the **intonation**.

2) The F1 score of the break classification ( $F1_{break}$ ) evaluates the **phrasing** aspect of prosody.

3) The Word Error Rate ( $WER$ ) correlates well with the **intelligibility** of synthesised speech (Taylor and Richmond

<sup>1</sup><https://huggingface.co/google-bert/bert-base-uncased>

**Subjective Evaluation Metrics** We conduct a crowd-sourced Mean Opinion Score (MOS) human listening test to assess four aspects of synthesized speech, including the phrase break similarity ( $MOS_{break}$ ), the terminal intonation similarity ( $MOS_{intonation}$ ) between the synthesized and a reference speech, the intelligibility ( $MOS_{intelligibility}$ ) and the quality ( $MOS$ ) of the synthesized speech. Each MOS is assessed using a 5-point scale with 95% confidence intervals, where score 1 indicates dissimilarity, unintelligibility, or poor quality whereas score 5 signifies full similarity, intelligibility, or excellent quality. We randomly select 15 utterances (3 groups of 5) with different lengths in the testing set, and each sample is rated by 21 testers. Our testers are PhD students from four universities specializing in Computer Audition, with native languages including English, German, Chinese, and Turkish. They are all fluent in English.

To account for the possibility of non-expert testers, we provided detailed explanations of the four MOS metrics with clear definitions and examples before starting the test. The instruction page is in Appendix D of the extended version.

Considering text consistency between the reference audio and the target text to be synthesized, we perform both parallel and non-parallel MOS tests. Given the subjective nature of prosody evaluation, where individuals from different linguistic backgrounds can have varying perceptions of what constitutes ‘appropriate’ phrasing and intonation for a target text (Grover, Jamieson, and Dobrovolsky 1987), we adopt specific assumptions:

1) For the parallel subjective evaluation, we assume that the phrasing and intonation derived from the reference speech represent the ‘appropriate’ prosody. We provide labels for breaks and intonation based on the reference speech and ask testers to assess the **similarity** of the phrasing and intonation to these labels, rather than their appropriateness.

2) For the non-parallel subjective evaluation, a reference speech is still needed for similarity assessment. To maintain the relevance of prosody while accommodating different text content, we modify the words in the reference speech (the same 15 samples) transcript, ensuring that the sentence semantics and structure remain as close as possible to the original. We then transfer the break and intonation labels from the reference speech to the new target text. Testers are again asked to assess the **similarity** of the phrasing and intonation to these labels, rather than their appropriateness. This evaluation rests on the assumption that two sentences with similar semantics and structure should share same phrasing and intonation labels.

The target text and transcripts of reference speech with break and intonation labels under parallel and non-parallel settings can be found in Appendix E of the extended version.

**Comparative Models** To evaluate the performance of our model, we compare ProsodyFM with four SOTA models<sup>2</sup>: (1) StyleSpeech (Min et al. 2021): the expressive multi-speaker TTS model built on FastSpeech2 (Ren et al. 2021);

<sup>2</sup>The checkpoints used are from their official implementations.

Models	$RMSE_{f_0} \downarrow$	$WER \downarrow$	$F1_{break} \uparrow$
GT(vocoder)	0.2264	2.05%	77.02
StyleSpeech	0.4477	4.30%	60.49
GenerSpeech	0.4556	7.29%	60.30
StyleTTS2	0.3120	3.25%	58.44
MatchaTTS	0.3376	3.56%	60.08
ProsodyFM	<b>0.3068</b>	<b>3.22%</b>	<b>62.76</b>

Table 1: Objective results on the LibriTTS testing set.

(2) GenerSpeech (Huang et al. 2022): the TTS model towards high-fidelity style transfer, also extended from FastSpeech2 (Ren et al. 2021); (3) StyleTTS2 (Li et al. 2024): the expressive TTS model with human-level speech quality, improved from StyleTTS (Li, Han, and Mesgarani 2022); (4) MatchaTTS (Mehta et al. 2024): the fast and high-quality TTS model based on conditional flow matching.

To verify the effectiveness of our proposed modules, we compare ProsodyFM against three ablated variants: (5) w/o\_intonation: remove the Terminal Intonation Encoder from ProsodyFM; (6) w/o.break: remove the Phrase Break Encoder from ProsodyFM; (7) w/o\_into\_break: remove both the Terminal Intonation Encoder and the Phrase Break Encoder from ProsodyFM.

To provide a reference upper bound, we also include: (8) GT(vocoder): we extract the Mel-spectrogram from the ground truth audio and then reconstruct it using HiFiGAN.

## Results

### Model Performance

We conduct both objective and subjective evaluations to assess ProsodyFM and four SOTA models in terms of phrasing (break), intonation, and overall intelligibility. Table 1 and Table 2 show the results.

**Objective Results** As shown in Table 1, we observe that ProsodyFM outperforms the other four SOTA models across all three objective evaluation metrics. Additionally, the results for phrasing ( $F1_{break}$ ) and intonation ( $RMSE_{f_0}$ ) show a positive correlation with overall intelligibility ( $WER$ ). These results indicate that ProsodyFM exhibits superior performance in phrasing and intonation, which further contributes to its enhanced intelligibility.

**Subjective Results** As shown in Table 2, compared to the other four SOTA models, ProsodyFM obtains significantly better scores in terms of  $MOS_{intonation}$  and  $MOS_{break}$  under both parallel and non-parallel settings; it also achieves significantly better  $MOS_{intelligibility}$  under the parallel setting. In the non-parallel setting, ProsodyFM matches the  $MOS_{intelligibility}$  of StyleTTS2 and surpasses the remaining three models significantly. In the non-parallel setting, ProsodyFM shows lower speech quality ( $MOS$ ) than StyleTTS2, likely due to the smaller dataset used in ProsodyFM (71 hours) compared to StyleTTS2 (245 hours). Consistent with the objective evaluation results, we can

Models	$MOS$		$MOS_{intelligibility}$		$MOS_{intonation}$		$MOS_{break}$	
	Parallel	Non-para	Parallel	Non-para	Parallel	Non-para	Parallel	Non-para
GT(vocoder)	4.77±0.05	-	4.77±0.05	-	4.90±0.04	-	4.91±0.04	-
StyleSpeech	3.43±0.10	3.39±0.10	3.72±0.11	3.72±0.11	3.58±0.10	3.53±0.10	3.52±0.12	3.74±0.11
GenerSpeech	2.67±0.10	-	3.03±0.12	-	3.14±0.12	-	2.85±0.13	-
StyleTTS2	4.09±0.10	<b>4.26±0.10</b>	4.23±0.10	4.36±0.09	3.66±0.10	3.78±0.11	3.82±0.12	4.08±0.11
MatchaTTS	3.61±0.10	3.90±0.09	4.03±0.10	4.10±0.10	3.48±0.10	3.73±0.10	3.88±0.10	4.08±0.11
ProsodyFM	<b>4.22±0.08</b>	4.07±0.09	<b>4.47±0.07</b>	<b>4.40±0.08</b>	<b>4.36±0.08</b>	<b>4.24±0.09</b>	<b>4.42±0.08</b>	<b>4.34±0.09</b>

Table 2: MOS results with 95% confidence intervals on the LibriTTS testing set. “Parallel” and “Non-para” indicate that the transcript of the reference audio is the same with or different from the target text, respectively.

Models	$RMSE_{f_0} \downarrow$	$WER \downarrow$	$F1_{break} \uparrow$
GT(vocoder)	0.2264	2.05%	77.02%
MatchaTTS	0.4727	6.78%	55.28%
ProsodyFM	<b>0.4080</b>	<b>4.97%</b>	<b>59.95%</b>

Table 3: Objective evaluation results on the out-of-distribution (unseen long and complex sentences, unseen speakers) testing data. All the models are trained on the VCTK (short sentences) training set and tested on the LibriTTS (long and complex sentences) testing set.

still observe a positive correlation between  $MOS_{intonation}$ ,  $MOS_{break}$  and  $MOS_{intelligibility}$ , further substantiating that ProsodyFM can effectively improve the phrasing and intonation, thereby enhancing the speech intelligibility.

### Model Generalizability

To evaluate the impact of enhanced phrasing and intonation on the generalizability of models, we conduct out-of-distribution experiments on unseen complex sentences. Both MatchaTTS and ProsodyFM are trained on the same VCTK training set (short sentences) and tested on the LibriTTS testing set (long sentences). The speakers in the testing set are unseen during training. Table 3 presents the results.

We observe that although both MatchaTTS and ProsodyFM experience performance declines on the out-of-distribution testing set, the decrease in ProsodyFM’s performance is considerably smaller than that of MatchaTTS. Notably, for the  $F1_{break}$  metric, ProsodyFM in the out-of-distribution setting achieves matching performance with the four SOTA models in the in-distribution setting (as shown in Table 1). This indicates that enhanced phrasing and intonation can bring strong generalizability.

### Ablation Study

To verify the necessity and effectiveness of our proposed Phrase Break Encoder and Terminal Intonation Encoder, we compare ProsodyFM against its three ablated variants. Table 4 shows the results on our LibriTTS validation set.

We observe an improved  $F1_{break}$  score in both w/o\_intonation and w/o.break compared to w/o.into.break,

Models	$RMSE_{f_0} \downarrow$	$WER \downarrow$	$F1_{break} \uparrow$
GT(vocoder)	0.2244	2.01%	76.04%
ProsodyFM	<b>0.3047</b>	<b>2.86%</b>	<b>62.51%</b>
w/o_intonation	0.3373	3.13%	61.25%
w/o_break	0.3355	3.10%	61.06%
w/o_into_break	0.3391	3.33%	60.25%

Table 4: Ablation study on the LibriTTS validation set.

which suggests that both the break encoder and intonation encoder inject partial phrase break information. Those break information may be complementary, leading to a further boost in  $F1_{break}$  when we combine the two encoders in ProsodyFM. The  $RMSE_{f_0}$  for three ablated variants exhibit no substantial differences, which likely stems from the  $RMSE_{f_0}$  metric being calculated only for voiced segments. When both intonation and break information are present, ProsodyFM achieves considerably better  $RMSE_{f_0}$ . These observations suggest that both the Phrase Break Encoder and Terminal Intonation Encoder are essential for synthesizing highly intelligible speech.

### Case Study: Prosody Controllability

We present a case study to visually demonstrate the ability of ProsodyFM to control prosody, specifically in terms of intonation and phrasing, which are the focus of this paper. Figure 4 displays the spectrograms (truncated due to space constraints) of audio samples synthesized with controlled intonation and phrasing (can be found on our demo page). The reference speech is from our LibriTTS testing set. The corresponding transcript with the original break and intonation labels is “Quite suddenly he rolled over (falling tone) stared for a moment (rising tone)”. We also provide more cases in Appendix C of the extended version.

**Intonation Control** For the controllability of the terminal intonation, we manually modify the reference pitch shape segment of the last word in an intonational phrase and synthesize the corresponding speech using ProsodyFM (we modify the input “Last-word (reference) Pitch Shape Segments” of the Terminal Intonation Encoder in Figure 3 (b)). We apply linear adjustments to the pitch values to create ris-

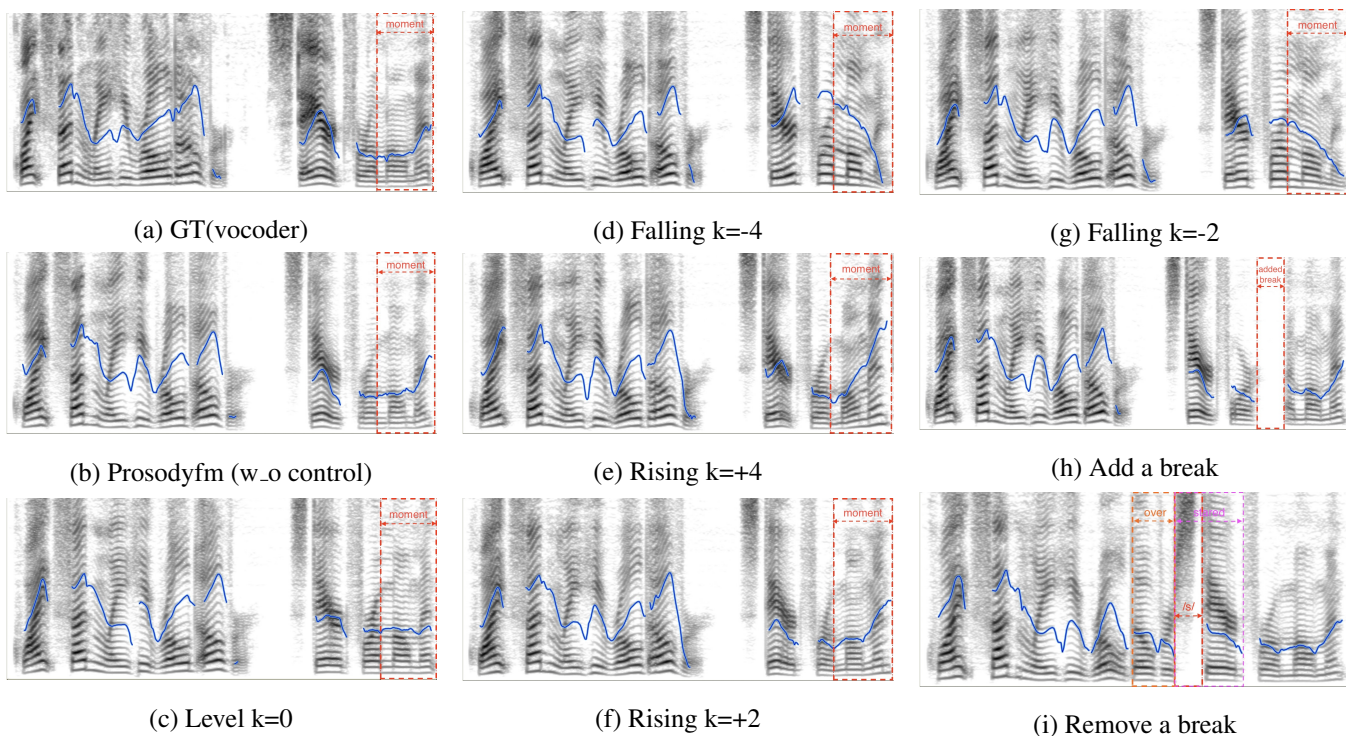


Figure 4: Part of spectrograms and pitch contours (blue line) of a reference speech (1363\_139304\_000009\_000005.wav) and speech synthesized by ProsodyFM with controlled intonation and phrasing.

ing, falling, and level tones, controlling the magnitude of these adjustments through the slope: a slope of  $k = +4$  represents a rapid rising tone,  $k = +2$  represents a gradual rising tone,  $k = -4$  indicates a rapid falling tone,  $k = -2$  indicates a gradual falling tone, and  $k = 0$  corresponds to a level tone. In Figure 4 (c-g), we modified the reference pitch shape segment of the last word “moment”.

We observe that when a level tone is provided as a reference, the pitch contour corresponding to the word “moment” in the synthesized speech remains essentially flat. Conversely, when rising or falling tones are used as references, the pitch contour for “moment” exhibits the corresponding upward or downward movement. Additionally, the pitch contour slopes in Figures 4 (d) and (e) are noticeably steeper than those in Figures 4 (g) and (f). This indicates that the pitch contour in the synthesized speech accurately reflects the same degree of reference rapid or gradual shape. These results demonstrate that our proposed Terminal Intonation Encoder effectively captures the reference intonation pattern, allowing ProsodyFM to achieve precise and fine-grained control over intonation.

**Phrasing Control** For the controllability of the phrase break, we manually add or remove a phrase break and synthesize the corresponding speech using ProsodyFM (we modify the “Phrase Breaks (last words)” in the Phrase Break Encoder in Figure 3 (b)). In Figure 4 (h), we added a phrase break after the word “for,” and in Figure 4 (i), we removed the phrase break between “over” and “stared”.

We observe that when a break is added after “for,” as shown in Figure 4 (h), the spectrogram of the synthesized speech displays a noticeable blank space between “for” and “a moment.” As shown in Figure 4 (i), when the break between “over” and “stared” is removed, a previously existing blank space disappears. This demonstrates that ProsodyFM exhibits excellent control over phrasing.

## Conclusion

We proposed ProsodyFM, a novel prosody-aware TTS model designed to enhance phrasing and intonation without requiring any prosodic labels, resulting in more intelligible synthesized speech. We addressed the intonation modeling issue by employing a novel Pitch Processor to highlight pitch shapes and training a bank of intonation shape tokens to model perceptually aligned intonation patterns instead of absolute pitch values. We tackled the break labeling issue by designing a Phrase Break Encoder to capture initial phrase break locations and then adjusting the variable break durations with a Duration Predictor. Our performance experiments demonstrated that ProsodyFM effectively improved the phrasing and intonation aspects of prosody, thereby enhancing overall intelligibility compared to four SOTA models. Our out-of-distribution experiments showed that this enhanced prosody further brought ProsodyFM strong generalizability on unseen complex sentences. Our ablation study verified the effectiveness of our proposed modules. Our case study visually demonstrated ProsodyFM’s powerful, precise and fine-grained control over phrasing and intonation.

## References

- Abbas, S. A.; Merritt, T.; Moinet, A.; Karlapati, S.; Muszynska, E.; Slangen, S.; Gatti, E.; and Drugman, T. 2022. Expressive, Variable, and Controllable Duration Modelling in TTS. In *23rd Annual Conference of the International Speech Communication Association, Interspeech 2022*, 4546–4550. ISCA.
- Birkholz, P.; and Zhang, X. 2020. Accounting for microprosody in modeling intonation. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 8099–8103. IEEE.
- Boersma, P. 2001. Praat, a system for doing phonetics by computer. *Glott. Int.*, 5(9): 341–345.
- Bolinger, D. 1998. Intonation in American English. *Intonation systems: A survey of twenty languages*, 45–55.
- Camacho, A.; and Harris, J. G. 2008. A sawtooth waveform inspired pitch estimator for speech and music. *The Journal of the Acoustical Society of America*, 124(3): 1638–1652.
- Cangemi, F. 2015. mausmooth. <https://ifl.phil-fak.uni-koeln.de/sites/linguistik/Phonetik/pdf-publications/2015/cangemi2015mausmooth.pdf>. Accessed: 2024-12-01.
- Casanova, E.; Weber, J.; Shulby, C. D.; Junior, A. C.; Gölge, E.; and Ponti, M. A. 2022. Yourtts: Towards zero-shot multi-speaker tts and zero-shot voice conversion for everyone. In *International Conference on Machine Learning*, 2709–2720. PMLR.
- Chodroff, E.; and Cole, J. S. 2019. Testing the Distinctiveness of Intonational Tunes: Evidence from Imitative Productions in American English. In *20th Annual Conference of the International Speech Communication Association, Interspeech 2019*, 1966–1970. ISCA.
- Cole, J.; Steffman, J.; and Tilsen, S. 2022. Shape matters: Machine classification and listeners’ perceptual discrimination of American English intonational tunes. In *Proceedings of the International Conference on Speech Prosody 2022*, Proceedings of the International Conference on Speech Prosody 2022, 297–301. ISCA.
- Futamata, K.; Park, B.; Yamamoto, R.; and Tachibana, K. 2021. Phrase Break Prediction with Bidirectional Encoder Representations in Japanese Text-to-Speech Synthesis. In *22nd Annual Conference of the International Speech Communication Association, Interspeech 2021*, 3126–3130. ISCA.
- Grover, C.; Jamieson, D. G.; and Dobrovolsky, M. B. 1987. Intonation in English, French and German: perception and production. *Language and Speech*, 30(3): 277–295.
- Hirst, D. J.; and de Looze, C. 2021. Measuring Speech. Fundamental frequency and pitch. *Cambridge Handbook of Phonetics*, (1): 336–361.
- Hu, E. J.; Shen, Y.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; and Chen, W. 2022. LoRA: Low-Rank Adaptation of Large Language Models. In *The Tenth International Conference on Learning Representations, ICLR 2022*. OpenReview.net.
- Huang, R.; Ren, Y.; Liu, J.; Cui, C.; and Zhao, Z. 2022. Generspeech: Towards style transfer for generalizable out-of-domain text-to-speech. *Advances in Neural Information Processing Systems*, 35: 10970–10983.
- Hwang, J.-S.; Lee, S.-H.; and Lee, S.-W. 2023. PauseSpeech: Natural Speech Synthesis via Pre-trained Language Model and Pause-Based Prosody Modeling. In *Asian Conference on Pattern Recognition*, 415–427. Springer.
- Kim, J.; Kong, J.; and Son, J. 2021. Conditional variational autoencoder with adversarial learning for end-to-end text-to-speech. In *International Conference on Machine Learning*, 5530–5540. PMLR.
- Kong, J.; Kim, J.; and Bae, J. 2020. Hifi-gan: Generative adversarial networks for efficient and high fidelity speech synthesis. *Advances in neural information processing systems*, 33: 17022–17033.
- Lee, Y.; and Kim, T. 2019. Robust and fine-grained prosody control of end-to-end speech synthesis. In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 5911–5915. IEEE.
- Li, Y. A.; Han, C.; and Mesgarani, N. 2022. Styletts: A style-based generative model for natural and diverse text-to-speech synthesis. *arXiv preprint arXiv:2205.15439*.
- Li, Y. A.; Han, C.; Raghavan, V.; Mischler, G.; and Mesgarani, N. 2024. Styletts 2: Towards human-level text-to-speech through style diffusion and adversarial training with large speech language models. *Advances in Neural Information Processing Systems*, 36.
- Liberman, M. Y. 1975. *The intonational system of English*. Ph.D. thesis, Massachusetts Institute of Technology.
- Lipman, Y.; Chen, R. T. Q.; Ben-Hamu, H.; Nickel, M.; and Le, M. 2023. Flow Matching for Generative Modeling. In *The Eleventh International Conference on Learning Representations, ICLR 2023*. OpenReview.net.
- Mauch, M.; and Dixon, S. 2014. pYIN: A fundamental frequency estimator using probabilistic threshold distributions. In *2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 659–663. IEEE.
- McAuliffe, M.; Socolof, M.; Mihuc, S.; Wagner, M.; and Sonderegger, M. 2017. Montreal forced aligner: Trainable text-speech alignment using kaldi. In *Interspeech*, volume 2017, 498–502.
- Mehta, S.; Tu, R.; Beskow, J.; Székely, É.; and Henter, G. E. 2024. Matcha-TTS: A fast TTS architecture with conditional flow matching. In *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 11341–11345. IEEE.
- Min, D.; Lee, D. B.; Yang, E.; and Hwang, S. J. 2021. Meta-styleSpeech: Multi-speaker adaptive text-to-speech generation. In *International Conference on Machine Learning*, 7748–7759. PMLR.
- Morise, M. 2017. Harvest: A High-Performance Fundamental Frequency Estimator from Speech Signals. In *18th Annual Conference of the International Speech Communication Association, Interspeech 2017*, 2321–2325. ISCA.

- Morise, M.; Kawahara, H.; and Katayose, H. 2009. Fast and reliable f0 estimation method based on the period extraction of vocal fold vibration of singing voice and speech. In *Audio Engineering Society Conference: 35th International Conference: Audio for Games*. Audio Engineering Society.
- Ni, J.; Hernandez Abrego, G.; Constant, N.; Ma, J.; Hall, K.; Cer, D.; and Yang, Y. 2022. Sentence-T5: Scalable Sentence Encoders from Pre-trained Text-to-Text Models. In *Findings of the Association for Computational Linguistics: ACL 2022*, 1864–1874. ACL.
- Ren, Y.; Hu, C.; Tan, X.; Qin, T.; Zhao, S.; Zhao, Z.; and Liu, T. 2021. FastSpeech 2: Fast and High-Quality End-to-End Text to Speech. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.
- Roll, N.; Graham, C.; and Todd, S. 2023. PSST! Prosodic Speech Segmentation with Transformers. In Jiang, J.; Reitter, D.; and Deng, S., eds., *Proceedings of the 27th Conference on Computational Natural Language Learning (CoNLL)*, 476–487. Singapore: Association for Computational Linguistics.
- Silverman, K. E.; Beckman, M. E.; Pitrelli, J. F.; Ostendorf, M.; Wightman, C. W.; Price, P.; Pierrehumbert, J. B.; and Hirschberg, J. 1992. ToBI: A standard for labeling English prosody. In *ICSLP*, volume 2, 867–870.
- Taylor, J.; and Richmond, K. 2021. Confidence Intervals for ASR-Based TTS Evaluation. In *22nd Annual Conference of the International Speech Communication Association, Interspeech 2021*, 2791–2795. ISCA.
- Taylor, P. 2009. *Text-to-speech synthesis*. Cambridge university press.
- Wang, Y.; Stanton, D.; Zhang, Y.; Ryan, R.-S.; Battenberg, E.; Shor, J.; Xiao, Y.; Jia, Y.; Ren, F.; and Saurous, R. A. 2018. Style tokens: Unsupervised style modeling, control and transfer in end-to-end speech synthesis. In *International conference on machine learning*, 5180–5189. PMLR.
- Xu, Y. 2019. Prosody, tone, and intonation. In *The Routledge handbook of phonetics*, 314–356. Routledge.
- Yamagishi, J.; Veaux, C.; and MacDonald, K. 2019. CSTR VCTK Corpus: English Multi-speaker Corpus for CSTR Voice Cloning Toolkit (version 0.92). sound.
- Yang, D.; Koriyama, T.; Saito, Y.; Saeki, T.; Xin, D.; and Saruwatari, H. 2023. Duration-aware pause insertion using pre-trained language model for multi-speaker text-to-speech. In *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 1–5. IEEE.
- Zen, H.; Dang, V.; Clark, R.; Zhang, Y.; Weiss, R. J.; Jia, Y.; Chen, Z.; and Wu, Y. 2019. LibriTTS: A Corpus Derived from LibriSpeech for Text-to-Speech. In *20th Annual Conference of the International Speech Communication Association, Interspeech 2019*, 1526–1530. ISCA.