

ASR for Affective Speech: Investigating Impact of Emotion and Speech Generative Strategy

Ya-Tse Wu

*Department of Electrical Engineering,
National Tsing Hua University
Hsinchu, Taiwan
crowpeter@gapp.nthu.edu.tw*

Chi-Chun Lee

*National Tsing Hua University
Hsinchu, Taiwan
cclee@ee.nthu.edu.tw*

Abstract—This work investigates how emotional speech and generative strategies affect ASR performance. We analyze speech synthesized from three emotional TTS models and find that substitution errors dominate, with emotional expressiveness varying across models. Based on these insights, we introduce two generative strategies: one using transcription correctness and another using emotional salience, to construct fine-tuning subsets. Results show consistent WER improvements on real emotional datasets without noticeable degradation on clean LibriSpeech utterances. The combined strategy achieves the strongest gains, particularly for expressive speech. These findings highlight the importance of targeted augmentation for building emotion-aware ASR systems.

Index Terms—Emotional speech, Automatic speech recognition, Text-to-speech synthesis, Data augmentation

I. INTRODUCTION

Automatic speech recognition (ASR) is a foundational technology in speech-based applications and plays a key role in modern human-computer interaction [1]. For real-world deployment, robustness is a critical requirement. While much of the existing research has focused on external factors such as background noise and reverberation [2], internal factors related to the speaker, particularly emotional expression, have received comparatively less attention despite their substantial impact on speech characteristics [3]. Emotion plays a central role in human communication, not only affecting prosody, articulation, and voice quality but also characterized either by discrete categories (e.g., angry, happy, sad) or by continuous dimensions such as valence and arousal [4]. Speech produced under different emotional states can vary considerably in both acoustic and linguistic patterns, posing challenges to conventional ASR systems [5]. Such variability poses challenges to conventional ASR systems, underscoring the need for models that can accurately recognize emotional speech and remain reliable under diverse real-world conditions.

Emotion has traditionally been regarded as a source of noise in the ASR process. For example, Catania et al. showed that emotional speech negatively impacts the performance of speech recognition and user input understanding, framing emotion as a disruptive factor in human-AI interaction [6]. Other studies have also noted lexical and acoustic differences across emotional states. Gulmira et al. proposed a low-resource emotional speech recognition method that detects emotions

based on a codebook of emotionally charged words in Kazakh language [7]. These findings suggest that emotional speech is more difficult to recognize than neutral speech due to its high variability in acoustic and prosodic features. Moreover, the high cost of collecting labeled emotional speech makes emotional TTS a scalable and low-cost alternative for training emotion-aware ASR systems.

One practical approach to addressing this challenge is through data augmentation, a technique widely used to improve robustness in noise-prone ASR systems [8], [9]. Among these, controllable text-to-speech (TTS) models have emerged as a scalable tool for simulating diverse speaking conditions. Previous studies have shown that TTS-based augmentation can improve ASR performance for dysarthric speech [10], suggesting that synthetic speech can help bridge the domain gap when real-world data is limited. In this context, emotional TTS offers a promising avenue for generating emotional speech that may help improve ASR performance under affective conditions.

Recent advances in large language model-based TTS systems, such as CosyVoice2 [11] and EmoVoice [12], have made it feasible to generate highly controllable and emotionally rich synthetic speech. However, the use of synthesized emotional speech for ASR training presents several challenges. Synthesized data may introduce acoustic or prosodic artifacts that distort recognition, and the emotional salience of synthetic speech may not fully align with real-world emotional expression. Additionally, it remains unclear which types of ASR errors are most impacted by emotional content, and how different TTS models influence these error patterns has not been systematically studied.

To better understand these effects, we conduct a series of analytical experiments to examine how synthetic emotional speech affects ASR performance. Starting from synthesized data, we analyze the predominant recognition errors introduced by an ASR model. Based on these observations, we design two generative strategies that guide data selection during training. Finally, we evaluate whether these strategies generalize beyond synthetic conditions by testing on real emotional speech datasets. Through this analysis, we aim to clarify the role of emotion in ASR degradation and demonstrate how targeted data construction can enhance system robustness under emotional variability. Our main findings are:

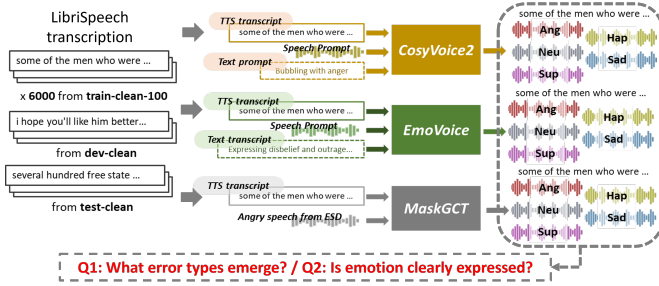


Fig. 1. Workflow for synthesized emotional speech analysis.

- 1) Synthesized emotional speech increases substitution errors, revealing emotion’s impact on phonetics.
- 2) Emotionally salient and correctly transcribed samples are more effective for ASR training.
- 3) Filtered synthetic data improves recognition on real emotional speech without harming neutral performance.

II. SYNTHESIZED DATA ANALYSIS AND EMOTION IMPACT ON ASR SYSTEM

In this section, we analyze the characteristics of emotional speech synthesized using three state-of-the-art emotion-controllable TTS models. An overview of the analysis workflow is presented in Figure 1. We conduct two main analyses: (1) examining ASR error patterns, including substitution, insertion, and deletion errors, and (2) evaluating the emotional salience of the synthesized speech based on the predicted distributions of dimensional emotion scores.

A. Emotion Controllable TTS Data Synthesis

To systematically investigate the impact of emotional variability on ASR performance, we control two key aspects in our synthetic speech generation process: transcription content and emotional expression. To minimize lexical variation and isolate the effect of emotion, we restrict the set of transcriptions used for synthesis. Specifically, we sample 6,000 utterances from the LibriSpeech train-clean-100 subset [13], yielding approximately 20 hours of speech. In addition, the full dev-clean and test-clean subsets are included as additional transcription sources.

Each transcription is synthesized in five representative emotional categories: Angry, Happy, Neutral, Sad, and Surprise. We utilize three state-of-the-art emotion-controllable TTS models to generate the speech data. For each model, the resulting dataset comprises 30,000 utterances for training, 13,470 for development, and 13,055 for testing.

1) *CosyVoice2* [11]: CosyVoice2 is a multilingual speech synthesis system that performs supervised learning over discrete speech tokens. Its architecture comprises three components: a supervised speech tokenizer, a text-to-speech language model, and a chunk-aware flow-matching module. The system is based on Qwen2.5-0.5B, a causal pre-trained language model with 24 Transformer layers and approximately 0.49 billion parameters. In our experiments, we use the instruction-tuned version of CosyVoice2 without additional fine-tuning.

Speech is synthesized by providing a speech prompt, a textual descriptor, and standard linguistic features. To control emotion, we prepend a fixed textual prompt of the form “Bubbling with {emotion}” to the input while keeping the original speech prompt unchanged to preserve speaker identity.

2) *EmoVoice* [12]: EmoVoice is an emotion-controllable TTS model that enables fine-grained emotional synthesis through natural language prompting. It is also built on Qwen2.5-0.5B but is further trained on an emotion-enriched dataset constructed using GPT-4o and GPT-4o-audio, allowing more nuanced emotional control. In our setup, we use the pretrained EmoVoice 1.5B model to synthesize speech from LibriSpeech transcriptions. To increase emotional diversity, we adopt the “freestyle” prompting strategy proposed by the authors, where a natural language emotion prompt is randomly selected from a predefined set for each utterance. The original speech prompt and transcription are preserved to maintain speaker identity and linguistic consistency. Example prompts are publicly available via the project’s Hugging Face page ¹.

3) *MaskGCT* [14]: MaskGCT is a fully non-autoregressive TTS model that uses masked generative transformers and requires no text-speech alignment or phone-level duration modeling. Unlike LLM-based models, MaskGCT is built on a LLaMA-style Transformer architecture and generates speech conditioned on a semantic speech prompt, which encodes prosodic, stylistic, and emotional information. In our experiments, we use the pretrained MaskGCT model released by the authors. For each synthesized utterance, we randomly select an emotional speech sample from the English subset of the ESD dataset [15] that shares the target emotion (e.g., Angry for Angry synthesis) to serve as the speech prompt, which is then paired with a LibriSpeech transcription to generate emotional speech.

B. Synthesized Data Analysis Methods

To comprehensively evaluate the quality of synthesized emotional speech, we adopt two complementary analysis strategies: (1) assessing transcription correctness to determine whether the synthesized speech accurately reflects the intended text, and (2) evaluating emotional salience to ensure that the generated speech conveys clear emotional expression rather than emotional neutrality.

1) *TTS Correctness Analysis*: One standard approach to assess the quality of synthesized speech is to evaluate transcription correctness using Word Error Rate (WER), which captures substitution, deletion, and insertion errors. Analyzing the distribution of these errors in synthesized emotional speech allows us to better understand how emotional content affects ASR transcription fidelity.

We use Qwen2-audio [16] as the ASR backbone for this analysis. This Speech-LLM architecture integrates Whisper-large-v3 [17] as the audio encoder and Qwen-7B [18] as the language model, jointly modeling acoustic and linguistic representations. We select Qwen2-audio due to its state-of-the-art performance on conversational English benchmarks such as

¹<https://huggingface.co/datasets/yhaha/EmoVoice-DB>

TABLE I
WER COMPARISON OF EMOTIONAL TTS AND LIBRISPEECH ON QWEN2-AUDIO.

WER on Qwen2-Audio	Train	Dev	Test
LibriSpeech	.0135	.0162	.0173
CosyVoice2	.0316	.0350	.0357
EmoVoice	.1073	.0819	.0915
MaskGCT	.0530	.0677	.0742

TABLE II
PREDICTED MOS SCORES (MEAN \pm STANDARD DEVIATION) FOR SYNTHESIZED EMOTIONAL SPEECH USING NISQA.

	LibriSpeech	CosyVoice2	EmoVoice	MaskGCT
MOS	4.06 \pm 0.55	3.79 \pm 0.63	3.85 \pm 0.56	4.40 \pm 0.42

Common Voice 15 [19], and its demonstrated ability to capture emotional nuances in speech, which is critical for evaluating ASR performance under affective conditions [20].

2) *Emotion Salience Analysis*: To assess whether the synthesized speech carries sufficient emotional salience, we employ an emotion regression model that estimates Arousal (Act), Valence (Val), and Dominance (Dom) on a 1-to-7 scale, where values near 4 indicate emotional neutrality.

The regressor is implemented as a multitask model based on the WavLM architecture, following the configuration used in the MSP-Podcast Challenge [21]. It jointly predicts Act, Val, and Dom to capture a multidimensional emotional representation. To prioritize acoustic over lexical features, the model is trained on the BIIC-Podcast dataset [22], a 157-hour Chinese podcast corpus with utterance-level annotations in all three emotion dimensions. Although the training domain differs from our synthesized English speech, we directly adopt the original hyperparameters without modification. On the BIIC-Podcast development set, the regressor achieves concordance correlation coefficients (CCC) of 0.68, 0.47, and 0.40 for Act, Val, and Dom, respectively.

C. Result and Discussion

Table I shows that all emotional TTS datasets lead to significantly higher WERs compared to the original LibriSpeech transcriptions, which have an average WER of 1.57% across splits. This supports prior findings that emotional utterances are more error-prone in ASR than neutral ones [6]. Among the synthesized datasets, CosyVoice2 yields the lowest WER at 3.74%, while MaskGCT and EmoVoice produce higher error rates at 6.50% and 9.37% in average, respectively. These results confirm that synthesized emotional speech presents additional challenges for ASR. The relatively low WER of CosyVoice2 suggests greater consistency and transcriptional clarity, whereas the high WER of EmoVoice likely stems from more expressive and variable prosody that increases recognition difficulty.

To ensure that transcription errors are not caused by degraded speech quality, we evaluate the synthesized utterances in training set using the NISQA model, a non-intrusive speech quality assessment metric [23]. As shown in Table II, all three emotional TTS systems generate high-quality audio with the

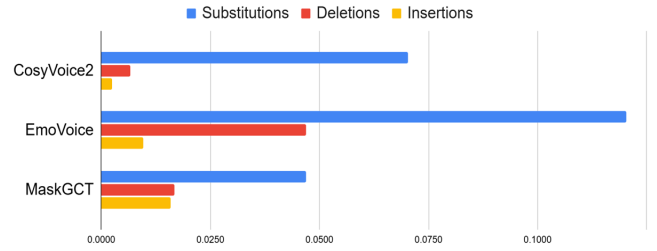


Fig. 2. Distribution of word error types in the training set of emotional TTS speech.




Fig. 3. Box plots of emotion distributions in the training set of emotional TTS speech.

average of Mean Opinion Score (MOS) estimates above 3.7, with MaskGCT achieving particularly strong results (4.40). These findings suggest that the elevated WERs are unlikely due to poor signal quality, and instead reflect challenges introduced by emotional expressiveness in the generated speech.

Figure 2 illustrates the distribution of substitution, insertion, and deletion errors on three synthesized training sets. Error rates are first normalized by the word length of each sentence and then averaged across the training set. Substitution errors are clearly dominant across all TTS systems, while insertions and deletions remain less frequent. This indicates that affective speech tends to cause semantic-level misrecognitions rather than structural disruptions. Notably, both EmoVoice and MaskGCT exhibit higher insertion and deletion rates than CosyVoice2, suggesting a higher tendency toward hallucination-related phenomena [24], such as skipped or repeated words.

Figure 3 presents the box plot distributions of Arousal (Act), Valence (Val), and Dominance (Dom) scores predicted by the emotion regressor. CosyVoice2 and EmoVoice show skewed Arousal distributions concentrated below 4, indicating weak emotional activation. MaskGCT, by contrast, produces a broader spread across the Act range. Valence distributions across all models are narrow with limited variability, although MaskGCT again demonstrates comparatively higher diversity. These results suggest that only a portion of synthesized utterances exhibit clear emotional salience.

In summary, our analysis reveals two key issues with synthesized emotional data: elevated ASR error rates, primarily due to substitution errors, and limited emotional expressiveness in a substantial portion of utterances. These findings highlight the importance of selectively leveraging high-quality, emotionally salient samples. In the next section, we present two generative strategies, informed by transcription correctness and emotion salience, to construct more effective synthetic training sets and assess their impact on ASR performance.

III. GENERATIVE STRATEGIES AND COMPARISON ON SYNTHESIZED DATA

Motivated by the findings in the previous section, we develop two generative strategies to improve the effectiveness of synthesized emotional speech for ASR training. The first strategy selects utterances based on transcription correctness, and the second identifies samples with strong emotional salience. These strategies are designed to address the specific issues observed in our earlier analysis, including high substitution error rates and limited emotional expressiveness, by constructing targeted training subsets that emphasize both linguistic accuracy and emotional clarity.

A. Generative Strategy

1) *TTS Correctness (TTS-G)*: To address emotion-induced recognition errors, we define a generative strategy that selects synthesized utterances which introduce substitution errors, identified in our earlier analysis as the dominant error type. For each synthesized utterance derived from a LibriSpeech transcription, we compare its ASR error profile to that of its original counterpart using the same text. An utterance is retained only if it results in a higher number of substitutions while maintaining equal or fewer deletion and insertion errors. Formally, the selected subset \mathcal{U}_T is defined as:

$$\mathcal{U}_T = \left\{ u \in \mathcal{U} \left| \begin{array}{l} S_s(u) > S_o(u) \wedge \\ D_s(u) \leq D_o(u) \wedge \\ I_s(u) \leq I_o(u) \end{array} \right. \right\} \quad (1)$$

where $S(u)$, $D(u)$, and $I(u)$ denote the number of substitutions, deletions, and insertions respectively, and the subscripts s and o refer to synthesized and original utterances.

2) *Emotion Salience (EMO-G)*: To ensure that the speech in the augmented data has sufficient emotional expression, we define a generative strategy based on emotion salience. For each synthesized utterance, we compute the predicted dimensional emotion scores and retain it if at least one of the three dimensions—Arousal (Act), Valence (Val), or Dominance (Dom)—deviates by more than one standard deviation from its corresponding mean. Such deviations indicate a departure from emotional neutrality and reflect stronger emotional expression. The selected subset \mathcal{U}_E is defined as:

$$\mathcal{U}_E = \left\{ u \in \mathcal{U} \left| \begin{array}{l} |\text{Act}_u - \mu_{\text{Act}}| > \sigma_{\text{Act}} \vee \\ |\text{Val}_u - \mu_{\text{Val}}| > \sigma_{\text{Val}} \vee \\ |\text{Dom}_u - \mu_{\text{Dom}}| > \sigma_{\text{Dom}} \end{array} \right. \right\} \quad (2)$$

where μ and σ denote the mean and standard deviation across the synthesized dataset.

3) *Combined Generative Strategy (TTS-EMO-G)*: To further enhance data quality, we apply both the TTS correctness and emotion salience criteria in combination. An utterance is retained only if it satisfies both conditions. This joint strategy ensures that selected samples both exhibit strong emotional expression and introduce meaningful recognition errors, increasing their potential value for ASR training. We refer to this combined approach as TTS-EMO-G in our experiments.

TABLE III
WER ACROSS SUBSETS SELECTED BY DIFFERENT GENERATIVE STRATEGIES.

WER on Qwen2-Audio	Train	Dev	Test
CosyVoice2			
Vanilla	.0316	.0350	.0357
EMO-G	.0332	.0356	.0364
TTS-G	.0333	.0367	.0369
TTS-EMO-G	.0358	.0406	.0401
EmoVoice			
Vanilla	.1073	.0819	.0915
EMO-G	.1059	.0815	.1024
TTS-G	.0434	.0447	.0468
TTS-EMO-G	.0460	.0466	.0504
MaskGCT			
Vanilla	.0530	.0677	.0742
EMO-G	.0494	.0622	.0901
TTS-G	.0381	.0464	.0447
TTS-EMO-G	.0396	.0508	.0467

B. ASR Model Fine-tuning Details

To evaluate the effectiveness of our generative strategies, we fine-tune the pretrained Qwen2-audio-7B model using the training portions of the TTS-EMO-G subsets, which are constructed following the original LibriSpeech train/dev/test splits for each TTS system (CosyVoice2, EmoVoice, and MaskGCT). Fine-tuning is performed independently for each system using only its synthesized training set, with early stopping based on the lowest development loss from the corresponding dev set. ASR transcriptions are generated using the official Qwen2-audio inference prompt: “Detect the language and recognize the speech:”

Training is conducted using ms-swift v2.6.1 [25], a flexible framework for tuning large language model-based systems. We adopt a supervised fine-tuning approach in which only a small portion of the audio encoder (AudioEnc) is updated, while the rest of the model remains frozen. Specifically, we allow gradient updates on the final 0.1% of encoder parameters, amounting to approximately 1.27 million trainable parameters.

We use the AdamW optimizer from PyTorch 2.4.1 [26], with a learning rate of 10^{-5} . All hyperparameters follow the default ms-swift configuration, including a warm-up ratio of 0.05, gradient accumulation steps of 16, weight decay of 0.1, and a maximum gradient norm of 1. The model is trained for 5 epochs using cross-entropy loss. Each fine-tuning run takes approximately 72 to 96 hours on two NVIDIA A100 or V100 GPUs. Further implementation details are available in the official ms-swift2 documentation².

C. Results and Discussion

To determine whether the proposed generative strategies provide meaningful benefits, we first evaluate their direct impact on the pretrained Qwen2-audio model without any fine-tuning. Table III presents the WER results on subsets selected by three strategies: (1) TTS-G, which filters utterances based on transcription correctness; (2) EMO-G, which selects

²<https://swift.readthedocs.io/en/stable/>

TABLE IV
WER ON SYNTHESIZED TEST SETS USING TTS-EMO-G WITH AUDIOENC FINE-TUNING.

	Qwen2-Audio	CosyVoice2	EmoVoice	MaskGCT
CosyVoice2-Test	.0401	.0196	.0291	.0285
EmoVoice-Test	.0504	.0351	.0282	.0347
MaskGCT-Test	.0467	.0414	.0991	.0284
LibriSpeech	.0173	.0169	.0207	.0184
Avg	.0386	.0283	.0443	.0275

utterances with strong emotional salience based on dimensional emotion scores; and (3) TTS-EMO-G, which combines both criteria to retain samples that are both transcriptionally informative and emotionally expressive. Compared to the original synthesized dataset (Vanilla), these strategies yield notable variations across the training, development, and testing splits. Notably, TTS-EMO-G consistently results in higher WERs than TTS-G, with average increases of 0.32%, 0.27%, and 0.26% for CosyVoice2, EmoVoice, and MaskGCT, respectively. This increase is expected, as TTS-EMO-G emphasizes samples with greater acoustic variability and emotional expressiveness, introducing additional complexity for ASR models.

Table IV further reports the results after fine-tuning under the TTS-EMO-G setting. Evaluations are conducted on the synthesized test sets from each emotional TTS model, as well as on the original LibriSpeech test-clean set. With AudioEnc fine-tuning, most configurations yield consistent WER reductions relative to the pretrained baseline, except for EmoVoice when tested on the MaskGCT set. Crucially, performance on LibriSpeech remains largely stable across all conditions, confirming that incorporating emotionally expressive synthetic data does not degrade recognition accuracy on clean, neutral speech. Specifically, WER improves by 0.34% for EmoVoice and 0.11% for MaskGCT, with only a slight increase of 0.04% observed for CosyVoice2. These results demonstrate that the TTS-EMO-G strategy can effectively mitigate substitution errors introduced by emotional variability while preserving general ASR performance.

Among the three emotional TTS models, MaskGCT achieves the best overall performance, with an average WER of 2.75%, while EmoVoice shows the least improvement, averaging 4.43%. MaskGCT’s advantage is attributed to its more balanced emotional distribution, as illustrated in Figure 3. Compared to CosyVoice2 and EmoVoice, MaskGCT exhibits fewer outliers and more centralized emotion scores across Arousal, Valence, and Dominance, closer to the neutral midpoint of 4. This balance likely contributes to more stable acoustic patterns and greater ASR robustness under affective conditions. These findings highlight the effectiveness of the TTS-EMO-G strategy in controlled synthetic environments. In the following section, we evaluate whether similar improvements can be achieved on real-world emotional speech.

IV. REAL EMOTION DATA EVALUATION

To examine whether our generative strategies generalize to real emotional speech, we evaluate the ASR models fine-

tuned on the synthesized subsets introduced in the previous section. Specifically, we assess models trained using the three generative strategies: TTS-G, EMO-G, and TTS-EMO-G, on three benchmark datasets: MSP Podcast Test1, Test2, and IEMOCAP. These datasets are used only for evaluation, and no real emotional speech is seen during training. This setup allows us to determine whether models trained solely on synthetic data can generalize to real world affective conditions and to identify which emotional regions benefit from our approach.

A. Real Speech Emotion Datasets

1) *MSP-Podcast*: We use the MSP-Podcast corpus v1.11 [27] as the real emotional speech benchmark to evaluate our fine-tuned ASR models. This dataset contains emotionally rich speech collected from online platforms, spanning diverse speakers and topics. The recordings are segmented into utterances (16 kHz, mono) and annotated with transcriptions, speaker identities, and dimensional emotion ratings for Arousal, Valence, and Dominance. Each dimension is scored on a scale from 1 to 7. Following the official evaluation protocol, we use the predefined Test 1 and Test 2 splits, which contain 30,647 and 14,815 utterances, respectively, totaling 70.94 hours of speech.

2) *IEMOCAP*: We also evaluate model performance on the IEMOCAP corpus [28], a widely used emotional speech dataset consisting of dyadic conversations recorded in a controlled lab setting. The corpus includes five sessions with ten speakers, and each utterance is annotated with a transcript and dimensional emotion scores in Arousal, Valence, and Dominance. Each dimension is rated on a scale from 1 to 5. In this study, we use the entire corpus for evaluation, comprising 10,044 utterances and approximately 12 hours of speech.

B. Result and Discussion

Table V reports the performance of all generative strategies evaluated on MSP-Podcast Test1, Test2, and IEMOCAP. The highest WER improvements observed are 0.18%, 0.99%, and 0.57% on each dataset, respectively. Among the three strategies, TTS-EMO-G consistently achieves the best performance across all benchmarks. Notably, our model was not trained on any of these datasets and had no exposure to their emotional distributions. Despite this, it demonstrates measurable improvements, indicating strong generalization. Furthermore, the training transcriptions were originally well-recognized by the baseline ASR system (WER of 1.35% as shown in Table I), suggesting that the observed gains stem from emotion-induced variability rather than lexical differences. These findings confirm that the proposed generative strategies are effective not only on synthetic emotional data but also generalize well to real-world emotional speech.

Among the three TTS methods, MaskGCT achieves the best performance on IEMOCAP, consistent with earlier findings that its well-balanced emotional distribution improves ASR robustness. On MSP Test1 and Test2, the top-performing systems are based on MaskGCT and CosyVoice2, respectively. MSP Test1, curated using emotion prediction outputs, appears

TABLE V
WER ON MSP TEST1, TEST2, AND IEMOCAP USING DIFFERENT GENERATIVE STRATEGIES WITH AUDIOENC FINE-TUNING.

	MSP Test1				MSP Test2				IEMOCAP			
Qwen2-Audio	.1578				.1541				.0927			
	Vanilla	EMO-G	TTS-G	TTS-EMO-G	Vanilla	EMO-G	TTS-G	TTS-EMO-G	Vanilla	EMO-G	TTS-G	TTS-EMO-G
CosyVoice2	.1685	.1725	.1692	.1601	.1450	.1449	.1458	.1442	.1044	.0891	.0956	.0956
EmoVoice	.2271	.2092	.1611	.1970	.2297	.1818	.1453	.1850	.1380	.1399	.1090	.1336
MaskGCT	.1620	.1743	.1577	.1560	.1528	.1519	.1598	.1517	.1116	.0990	.0892	.0870

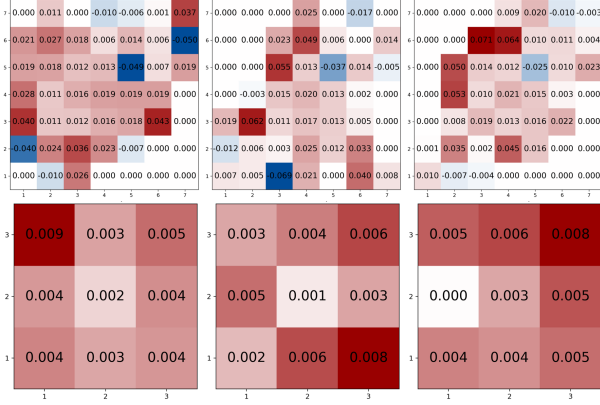


Fig. 4. Left to right: (Act, Val), (Val, Dom), and (Act, Dom) planes showing WER difference patterns for MSP-Test 2 (top) and IEMOCAP (bottom).

to benefit from MaskGCT’s expressive range. In contrast, MSP Test2 contains more emotionally neutral content, where CosyVoice2’s stable synthesis and lower hallucination rate (as shown in Table I) lead to better performance. As an additional observation, the combination of MaskGCT and TTS-EMO-G consistently improved over the pretrained baseline across all test sets, suggesting its robustness to varying emotional distributions.

To further understand where performance gains occur, we analyze WER changes across the emotional space using dimensional annotations from MSP Test2 and IEMOCAP. For MSP Test2, utterances are grouped by their rounded Arousal, Valence, and Dominance scores on a 1–7 scale, and WER is computed for each group across three 2D planes: (Act, Val), (Val, Dom), and (Act, Dom). For IEMOCAP, we adopt a coarse-grained 3×3 binning strategy. Each emotional dimension, originally scored from 1 to 5, is categorized into three levels: low (≤ 2.5), medium (> 2.5 and ≤ 3.5), and high (> 3.5). Utterances are grouped accordingly, and WER is calculated across the same three 2D planes to visualize ASR performance trends under varying emotional intensities.

Figure 4 visualizes the WER differences between the pretrained model and the best fine-tuned system. Red indicates improved WER, blue indicates degradation. Improvements are most pronounced at the emotional extremes, confirming that our strategies are especially effective for enhancing ASR performance under emotionally expressive conditions.

V. DISCUSSION AND CONCLUSION

This study investigates the impact of synthesized emotional speech on ASR performance. Analysis of three emotional TTS models shows that emotional variability mainly causes substitution errors, suggesting that affective modulation affects phonetic realization more than lexical structure. We also observe significant variation in emotional expressiveness and generative stability across TTS systems, influencing both recognition accuracy and emotion distribution.

To address these challenges, we propose two generative strategies: one based on transcription correctness and the other on emotional salience—that guide the selection of emotionally rich yet linguistically reliable training subsets. Fine-tuning with the combined strategy (TTS-EMO-G) yields consistent WER improvements on both synthetic and real emotional datasets, particularly in highly expressive regions, while maintaining performance on neutral speech. These results demonstrate the feasibility of leveraging high-quality synthetic data for emotion-aware ASR training.

This approach relies on well-trained emotional TTS systems and accurate dimensional emotion regression, limiting applicability in low-resource or cross-lingual settings. Since all synthesized data are derived from LibriSpeech transcripts, domain mismatches between synthetic and spontaneous speech may hinder generalization. The analysis also focuses on dimensional emotion representation (arousal, valence, dominance) rather than discrete categories; given that some emotions, such as sadness, may be acoustically closer to neutral, incorporating per-emotion analysis could offer complementary insights.

Beyond acoustic adaptation, incorporating lexical cues and contextual semantics may enhance robustness under emotionally complex conditions and broaden applications in multilingual and low-resource scenarios. Future work could explore language-agnostic augmentation techniques, finer-grained emotion control, and adaptive data selection strategies informed by speaker or context. Investigating lexical items particularly susceptible to recognition errors under different emotional conditions may enable targeted augmentation or model adaptation for these high-risk words.

REFERENCES

- [1] V. D. Nguyen, A. T. Do, D. K. Le, and L. V. Nguyen, “Speech recognition-based human–computer interaction: A survey,” in *Trends in Sustainable Computing and Machine Intelligence*. Springer, Singapore, 2025, pp. 289–301.

- [2] J. Li, L. Deng, Y. Gong, and R. Haeb-Umbach, "An overview of noise-robust automatic speech recognition," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 22, no. 4, pp. 745–777, 2014.
- [3] R. Munot and A. Nenkova, "Emotion impacts speech recognition performance," in *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Student Research Workshop*. Association for Computational Linguistics, 2019, pp. 16–21.
- [4] P. Laukka, P. Juslin, and R. Bresin, "A dimensional approach to vocal expression of emotion," *Cognition & Emotion*, vol. 19, no. 5, pp. 633–653, 2005.
- [5] A. Batliner, B. Schuller, D. Seppi, S. Steidl, L. Devillers, L. Vidrascu, T. Vogt, V. Aharonson, and N. Amir, "The automatic recognition of emotions in speech," in *Emotion-Oriented Systems: The Humaine Handbook*. Springer, 2010, pp. 71–99.
- [6] F. Catania, P. Crovari, M. Spitale, and F. Garzotto, "Automatic speech recognition: Do emotions matter?" in *2019 IEEE International Conference on Conversational Data & Knowledge Engineering (CDKE)*. IEEE, 2019, pp. 9–16.
- [7] G. Bekmanova, B. Yergesh, A. Sharipbay, and A. Mukanova, "Emotional speech recognition method based on word transcription," *Sensors*, vol. 22, no. 5, p. 1937, 2022.
- [8] H. Hu, T. Tan, and Y. Qian, "Generative adversarial networks based data augmentation for noise robust speech recognition," in *2018 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 2018, pp. 5044–5048.
- [9] A. Pervaiz, F. Hussain, H. Israr, M. A. Tahir, F. R. Raja, N. K. Baloch, F. Ishmanov, and Y. B. Zikria, "Incorporating noise robustness in speech command recognition by noise augmentation of training data," *Sensors*, vol. 20, no. 8, p. 2326, 2020.
- [10] W.-Z. Leung, M. Cross, A. Ragni, and S. Goetze, "Training data augmentation for dysarthric automatic speech recognition by text-to-dysarthric-speech synthesis," in *Interspeech 2024*, 2024, pp. 2494–2498.
- [11] Z. Du, Y. Wang, Q. Chen, X. Shi, X. Lv, T. Zhao, Z. Gao, Y. Yang, C. Gao, H. Wang *et al.*, "Cosyvoice 2: Scalable streaming speech synthesis with large language models," *arXiv preprint arXiv:2412.10117*, 2024.
- [12] G. Yang, C. Yang, Q. Chen, Z. Ma, W. Chen, W. Wang, T. Wang, Y. Yang, Z. Niu, W. Liu *et al.*, "Emovoice: Llm-based emotional text-to-speech model with freestyle text prompting," *arXiv preprint arXiv:2504.12867*, 2025.
- [13] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, "Librispeech: an asr corpus based on public domain audio books," in *2015 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 2015, pp. 5206–5210.
- [14] Y. Wang, H. Zhan, L. Liu, R. Zeng, H. Guo, J. Zheng, Q. Zhang, X. Zhang, S. Zhang, and Z. Wu, "MaskGCT: Zero-shot text-to-speech with masked generative codec transformer," in *The Thirteenth International Conference on Learning Representations*, 2025. [Online]. Available: <https://openreview.net/forum?id=ExuBFYtCQU>
- [15] K. Zhou, B. Sisman, R. Liu, and H. Li, "Seen and unseen emotional style transfer for voice conversion with a new emotional speech dataset," in *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 920–924.
- [16] Y. Chu, J. Xu, Q. Yang, H. Wei, X. Wei, Z. Guo, Y. Leng, Y. Lv, J. He, J. Lin *et al.*, "Qwen2-audio technical report," *arXiv preprint arXiv:2407.10759*, 2024.
- [17] A. Radford, J. W. Kim, T. Xu, G. Brockman, C. McLeavey, and I. Sutskever, "Robust speech recognition via large-scale weak supervision," in *International conference on machine learning*. PMLR, 2023, pp. 28 492–28 518.
- [18] J. Bai, S. Bai, Y. Chu, Z. Cui, K. Dang, X. Deng, Y. Fan, W. Ge, Y. Han, F. Huang *et al.*, "Qwen technical report," *arXiv preprint arXiv:2309.16609*, 2023.
- [19] R. Ardila, M. Branson, K. Davis, M. Henretty, M. Kohler, J. Meyer, R. Morais, L. Saunders, F. M. Tyers, and G. Weber, "Common voice: A massively-multilingual speech corpus," *arXiv preprint arXiv:1912.06670*, 2019.
- [20] H. Feng, S. Ueno, and T. Kawahara, "End-to-end speech emotion recognition combined with acoustic-to-word asr model," in *Interspeech 2020*, 2020, pp. 501–505.
- [21] L. Gonçalves, A. N. Salman, A. Reddy Naini, L. Moro-Velazquez, T. Thebaud, L. Paola Garcia, N. Dehak, B. Sisman, and C. Busso, "Odyssey2024 - speech emotion recognition challenge: Dataset, baseline framework, and results," in *Odyssey 2024: The Speaker and Language Recognition Workshop*, vol. To appear, Quebec, Canada, June 2024.
- [22] S. G. Upadhyay, W.-S. Chien, B.-H. Su, L. Gonçalves, Y.-T. Wu, A. N. Salman, C. Busso, and C.-C. Lee, "An intelligent infrastructure toward large scale naturalistic affective speech corpora collection," in *2023 11th International Conference on Affective Computing and Intelligent Interaction (ACII)*. IEEE, 2023, pp. 1–8.
- [23] G. Mittag, B. Naderi, A. Chehadi, and S. Möller, "Nisqa: A deep cnn-self-attention model for multidimensional speech quality prediction with crowdsourced datasets," in *Interspeech 2021*, 2021, pp. 2127–2131.
- [24] C. Du, Y. Guo, H. Wang, Y. Yang, Z. Niu, S. Wang, H. Zhang, X. Chen, and K. Yu, "Vall-t: Decoder-only generative transducer for robust and decoding-controllable text-to-speech," in *ICASSP 2025 - 2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2025, pp. 1–5.
- [25] Y. Zhao, J. Huang, J. Hu, X. Wang, Y. Mao, D. Zhang, Z. Jiang, Z. Wu, B. Ai, A. Wang *et al.*, "Swift: a scalable lightweight infrastructure for fine-tuning," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 39, no. 28, 2025, pp. 29 733–29 735.
- [26] A. Paszke, "Pytorch: An imperative style, high-performance deep learning library," *arXiv preprint arXiv:1912.01703*, 2019.
- [27] R. Lotfian and C. Busso, "Building naturalistic emotionally balanced speech corpus by retrieving emotional speech from existing podcast recordings," *IEEE Transactions on Affective Computing*, vol. 10, no. 4, pp. 471–483, 2017.
- [28] C. Busso, M. Bulut, C.-C. Lee, A. Kazemzadeh, E. Mower, S. Kim, J. N. Chang, S. Lee, and S. S. Narayanan, "Iemocap: Interactive emotional dyadic motion capture database," *Language resources and evaluation*, vol. 42, pp. 335–359, 2008.