

A Thai Spoken Dialogue Generation System Incorporating Generative Adversarial Networks

Ya Mansha^{1,*}

¹Department of Training College, Guangxi Police College, Nanning, 530028, Guangxi, China

Corresponding author: Ya Mansha (e-mail: jenny1213ya@sina.com).

ABSTRACT Traditional Thai dialogue generation systems, under maximum likelihood training, often favor high-frequency "safe" sentence structures, resulting in monotonous style and poor multi-turn contextual coherence. Similar problems exist in automated quality control within the textile industry: traditional models, trained on large amounts of defect-free data, struggle to distinguish subtle or rare defects, and are difficult to capture long-range dependencies in multi-sensor data streams. This paper proposes a Thai spoken dialogue generation system based on RelGAN and Transformer. It captures multi-turn context through a segmented relative position self-attention encoder and segmented recurrent memory, while the decoder embeds a relational memory module to track dialogue states and long-term entity slots, simultaneously generating text and prosodic markers. The generated results are prosodic mapped using FastSpeech2 and synthesized into high-fidelity speech using HiFi-GAN. Experiments show that the system achieves a Distinct-2 score of 0.11, an intent retention rate of 97.5%, an average response time of 98.44 ms, and a MOS of 4.05. This method has significant potential for cross-industry applications, enabling the construction of highly consistent real-time diagnostic systems in the textile industry, processing continuous multi-cycle sensor data, and achieving rapid fault identification and feedback.

INDEX TERMS Thai Spoken Language; Dialogue Generation System; RelGAN-Transformer Algorithm; Generating Response Diversity; Textile Industry

I. INTRODUCTION

WITH the widespread use of online language teaching and intelligent tutoring systems in educational settings, spoken dialogue generation systems are playing an increasingly important role in classroom practice, oral assessment, and autonomous learning. Analogously, in the textile industry, sophisticated AI-driven dialogue or guidance systems are becoming essential for intelligent quality control and machine maintenance, providing real-time, step-by-step guidance for technicians, enabling precise fault diagnosis, and supporting continuous autonomous optimization of production lines [1-3]. To meet the demand for natural, coherent, and spoken-style interactions in teaching scenarios, the system must not only correctly respond to semantics but also closely resemble authentic native Thai expressions in terms of style, prosody, and cross-turn coherence. However, existing Thai spoken dialogue generation systems still have significant deficiencies in terms of word segmentation problems caused by writing without spaces, processing colloquial expressions and slang, and maintaining multi-round context, which directly affects the quality of teaching interaction and learner experience. In spoken dialogue generation systems, existing research has explored

naturalness, response diversity, and emotional understanding. Zhai C's dialogue system based on WGAN (Wasserstein Generative Adversarial Network) performs well in modeling humor, empathy, and cultural sensitivity, but is weak in modeling discrete text and long-term context [4]. Le A C combines deep reinforcement learning with adversarial learning to improve anthropomorphic expression, but the training complexity and cross-round consistency still need to be improved [5]. EP-GAN (Empathy and Persona Aware Generative Adversarial Network) proposed by Firdaus M encodes character information through memory vectors to generate emotionally sensitive and personalized responses, but it is not well adapted to low-resource languages and spoken rhythm. Although its multimodal slot framework enhances coherence and information integrity, it increases data and engineering costs [6-7]. Fuad A's AraConv uses mT5 (Multilingual T5) and cross-language joint training to alleviate the low resource problem, but it relies on cross-language data and joint training strategies, and is not well adapted to spoken language details such as intonation and word segmentation[8]; Li X integrates Transformer and BERT (Bidirectional Encoder Representations from Transformers) to improve contextual understanding, but the computational

cost is high and it is not robust to spoken language noise[9]; Zhao Y explicitly integrated multi-hop concept entity flow with post-perception knowledge encoding to construct a dialogue generation model, improving fluency and response diversity, but it was difficult to model dynamic semantic relationships between multi-hop entities[10]. To address the bottlenecks of existing generators, which suffer from a single spoken style and loss of multi-turn context, this paper constructs a generation system that combines RelGAN (Relational Generative Adversarial Networks) with the Transformer. On the data side, a hybrid character-subword representation, ASR (Automatic Speech Recognition)-style noise injection, and prosody token annotation are employed. On the generator side, Transformer self-attention is used as the backbone, and relational memory slots are embedded in the decoder to explicitly maintain cross-turn information. A pre-training/fine-tuning process using MLE (Maximum Likelihood Estimation) + coverage is combined with adversarial fine-tuning using a WGAN-GP (Wasserstein Generative Adversarial Networks with Gradient Penalty) multi-view Transformer discriminator (for content, style, prosody, and fluency) under continuous relaxation of Gumbel-Softmax, along with feature matching and auxiliary prosody/style supervision. Ultimately, the prosodic output of the generated text was combined with FastSpeech2+HiFi-GAN for teaching demonstration speech. The results demonstrate that through joint optimization using data-level augmentation, explicit memory, and multi-view adversarial methods, the diversity and cross-turn coherence of generated spoken Thai responses can be significantly improved while maintaining semantics and fluency, meeting the practical requirements of classroom demonstrations and student imitation. This optimization strategy is also highly applicable to the textile industry, where it can enhance the performance of automated quality inspection by improving the detection diversity of subtle or rare fabric defects and ensuring diagnostic coherence across continuous production batches, thereby maintaining the high standards required for real-time quality assurance and process control. The core innovation of this study is not the originality of a single component, but rather the first proposal and validation of a specific architectural combination and its cross-domain transfer paradigm to address the unique challenges of spoken Thai dialogue. The innovations are reflected in three aspects: (1) Collaborative architectural innovation: The relational memory module of RelGAN is embedded between Transformer decoder layers (instead of simple concatenation), forming a dual-path long-term dependency model with segmented recurrent memory, so that the entity coverage rate remains at 65.45% after 10 rounds of dialogue (42.95 percentage points higher than SeqGAN); (2) Multimodal joint optimization: For the first time, text and prosodic markers are modeled simultaneously in Thai dialogue generation, and adversarial fine-tuning is performed by combining a four-view WGAN-GP discriminator (content/style/prosodic/fluency), which solves the problems of Thai word segmentation ambiguity and

spoken rhythm modeling, so that the intent retention rate reaches 97.5%, which is significantly better than cGAN-Transformer-CVAE (+10.1%); (3) Cross-industry architectural migration: The relational memory mechanism for dialogue state tracking is redefined as a temporal consistency maintenance module for multi-sensor streams, providing the first interpretable long-range diagnostic framework for industrial scenarios such as textile quality inspection. The combination achieved a statistically significant improvement ($p < 0.001$) on the key Distinct-2 (0.11) metric, validating the performance breakthrough brought about by architectural collaboration rather than individual components.

II. RELGAN-TRANSFORMER THAI SPOKEN DIALOGUE GENERATION MODEL

A. THAI SPOKEN CORPUS PREPROCESSING AND DATA AUGMENTATION

This paper integrates public dialogue corpus, customer service conversations, classroom speech transcriptions, and manually collected teaching conversations, and unifies them into UTF-8 encoding. Unicode normalization, full-width/half-width unification, punctuation normalization, and number regularization are performed. A mapping table is created for common spelling variants and replaced to eliminate noise-type bias. In the character-subword hybrid representation, the paper uses BPE (Byte Pair Encoding) sentence fragments to train the subword vocabulary while retaining the character-level sequence [11-12]. For each input sentence, it generate both the character sequence and the subword sequence, and use a learnable weighted fusion in the embedding layer, as shown in formula (1).

$$e_t = \alpha E_{\text{char}}(c_t) + (1 - \alpha) E_{\text{sub}}(s_{\pi(t)}) \quad (1)$$

In formula (1), E_{char} represents the character embedding matrix, E_{sub} represents the subword embedding matrix, $\pi(t)$ represents the character-to-subword alignment mapping, and α represents the trainable weights. c represents the character sequence, s represents the subword sequence, and e_t represents the fused vector at time t . Based on the statistical editing model, the process of generating noise samples from the original sentence is simulated by controlling the perturbation, that is, random editing at independent positions. The probabilities of retention, replacement and deletion at each position are $p_{\text{keep}}(i)$, $p_{\text{sub}}(i)$ and $p_{\text{del}}(i)$, respectively, satisfying the following constraint (2):

$$p_{\text{keep}}(i) + p_{\text{sub}}(i) + p_{\text{del}}(i) = 1 \quad (2)$$

These probabilities are estimated by statistically comparing aligned ASR outputs with artificial transcriptions. Integrating noisy samples proportionally into the training set allows the model to directly confront typical ASR error distributions during training, thereby improving the fault tolerance of downstream generation. In Equation (2), sub represents the replacement function based on the perplexity matrix, and del represents deletion. For the subset with

speech alignment, Librosa is used to extract the fundamental frequency, duration, and energy. The fundamental frequency is then logarithmically transformed and quantized. The quantized prosodic tokens are then inserted into the text sequence (for pauses, stresses, etc.). These tokens can be used as part of the training objective, enabling the generator to learn the joint text-prosodic distribution and improve the natural presentation of the generated text in TTS (Text to Speech).

The study used a pre-trained Transformer fine-tuned CRF (Conditional Random Field) model for named entity recognition [13-15]. Key entities (names, places, times, and course terms) were replaced with placeholders. Both the original and placeholder versions were retained in the training data, enabling the model to learn common response patterns while also filling in real entities during inference via retrieval or external KGs (Knowledge Graphs), reducing factual errors and enhancing cross-round information tracking. The paper also expanded the scarce teaching corpus and used a bidirectional Thai-English/Thai-Chinese sequence-to-sequence translation model to back-translate sentences and generate diverse variants. It combined a synonym dictionary with syntactic template replacement (maintaining slot occupancy), and manually sampled the generated samples.

B. DESIGN OF A CONTEXT-AWARE DIALOGUE GENERATOR

This paper uses the Transformer as the backbone [16-17], divides the entire session into segments, and constructs segment-level recursive memory. Let the input matrix of the current segment be $\tilde{X}$, the hidden state of the previous segment be H^{mem} , and the key-value pair matrix $K = [H^{\text{mem}}, \tilde{X}]$ be constructed. The query is $Q = \tilde{X}$. To support relative position encoding, a relative position representation is adopted, and a position term is added to the attention calculation. The single-head attention calculation is shown in formula (3) [18-20].

$$A = \text{softmax}\left(\frac{QK^T + QR^T}{\sqrt{d_k}}\right) \quad (3)$$

In formula (3), R represents the relative position matrix encoding, and d_k represents the key dimension. Multi-head attention is implemented by parallelizing multiple groups (Q, K, V) and concatenating them after the head. Segment-level recursion allows key information (entities, slots, etc.) from the previous segment to be retained and directly accessed by the current segment, alleviating the problem of cross-round long-range dependency decay. The segment size was determined using a strategy that balanced dialogue structure analysis with validation set performance. Statistical analysis of the dialogue turn length distribution in the training corpus revealed that 85% of teaching dialogue scenarios completed a full topic exchange within 8 turns, with topic transitions typically occurring in even-numbered turns. Based on this, different segment sizes of 2 to 12 turns (step size 2) were tested on the validation set, with




FIGURE 1. Impact of Segment Size on Model Performance

entity F1 score, perplexity, and average response time as optimization objectives. Through multi-objective weighted evaluation (entity F1 weight 0.5, perplexity weight 0.3, response time weight 0.2), 6 turns was determined to be the optimal segment size. This configuration balances computational efficiency and context coverage while preserving the consistency of key teaching entities, avoiding cross-segment information fragmentation caused by excessively small segments and attention dilution caused by excessively large segments. The impact of segment size on model performance is shown in Figure 1.

In Figure 1, segment size has a non-linear impact on the model's long-range dependency modeling and dialogue coherence. When the segment size is less than 4 rounds, frequent cross-segment information transfer leads to a significant decrease in entity F1 (only 78.23% at 2 rounds), making it difficult for the system to maintain topic coherence in multi-round dialogues. When the segment size reaches 6 rounds, entity F1 reaches a peak of 94.15%, while maintaining low perplexity (18.95) and an acceptable response time (98.44ms), indicating that this configuration can effectively capture the typical topic span in instructional dialogues. When the segment size exceeds 8 rounds, although dialogue coherence is slightly maintained, computational complexity increases sharply, response time increases by nearly 60ms, and the attention mechanism begins to dilute key entity signals, causing entity F1 to decrease by about 5 percentage points. This trade-off reveals the essence of segmented design: it needs to be aligned with the dialogue structure features of the target domain. In the teaching scenario, the 6-round segmentation just covers the typical teaching round of "question-explanation-example-practice-feed back-summary", enabling the relational memory module to efficiently capture cross-round dependencies while maintaining entity consistency. A relational memory module M_O is introduced after each layer of the standard

Transformer decoder. The internal operations of each decoding step t include reading and writing the memory based on the current decoding state. The read operation is defined as a gated attention, and the calculation is shown in formula (4).

$$\begin{cases} q_t = W_q[h_t; m_o^{\text{agg}}], & K_M = W_k M_O, & V_M = W_v M_O, \\ a_t = \text{softmax}\left(\frac{q_t K_M^T}{\sqrt{d_k}}\right), & r_t = a_t V_M \end{cases} \quad (4)$$

In formula (4), m_o^{agg} represents the aggregate vector stored in the previous memory, W_q , W_k , and W_v represent the corresponding projection matrices, r_t represents the memory readout vector, and h_t represents the current decoding state. The readout vector is concatenated with the decoder output to form the final logits, enhancing the decoder's ability to directly access long-term stored information. Memory write operations utilize gated updates, as shown in formula (5). This writing mechanism only injects information related to the current decoding and gated into the memory, avoiding noise from covering long-term slots and ensuring that important entities or course terms in teaching scenarios are permanently retained.

$$\begin{cases} u_t = \sigma(W_u[h_t; r_t] + \beta_u), & \hat{c}_t = \tanh(W_{\hat{c}}[h_t; r_t] + \beta_{\hat{c}}), \\ M_O \leftarrow M_O + u_t \odot \hat{c}_t \end{cases} \quad (5)$$

In formula (5), σ represents element-wise sigmoid, $\odot$ represents element-wise product, W_u and $W_{\hat{c}}$ represent the trainable weights of the corresponding gates, and β_u and $\beta_{\hat{c}}$ represent biases. Compared to standard sequence generative adversarial networks (GANs) or policy gradient reinforcement learning methods, RelGAN's relational memory module offers three structural advantages in dialogue state tracking and text generation. In dialogue state tracking, relational memory replaces the implicit tracking of hidden states in traditional RNNs with an explicit memory slot mechanism, enabling the system to directly access and update discrete entity-relation structures, rather than relying on fuzzy representations in continuous vector spaces. This design effectively addresses the problem of forgetting key information (such as course terminology, names, and locations) in multi-turn dialogues, while avoiding the instability of state tracking caused by sparse rewards in reinforcement learning methods. In text generation, the collaborative mechanism between the relational memory module and the multi-view discriminator provides fine-grained adversarial feedback. Unlike sequence GANs that only focus on overall sequence authenticity or policy gradient-dependent delayed rewards, this system can simultaneously optimize four dimensions: content semantics, spoken style, prosody, and fluency. This allows the generated response to naturally incorporate Thai's unique intonation patterns while maintaining teaching accuracy. In terms of training stability, the gating update mechanism of relational memory combined with the

feature matching loss of adversarial training significantly alleviates the pattern collapse problem in traditional GANs in discrete sequence generation, and also avoids the high variance and reward sparsity challenges common to policy gradient methods. This architecture is particularly well-suited to the stringent requirements of educational scenarios for both content accuracy and natural expression, ensuring that the system can maintain factual consistency in long dialogue sequences while generating diverse expressions that conform to Thai spoken language habits. The relational memory module, through its unique memory slot structure and dynamic read/write mechanism, intrinsically achieves structured tracking of dialogue entities and slots.

Unlike traditional symbolic state trackers that rely on pre-defined rules and explicit state tables, this system's relational memory module automatically learns to identify key entities (such as people, locations, and times) and dialogue slots (such as course topics and teaching objectives) during training, and establishes the relationship topology between them in the memory slots through a multi-head attention mechanism. When a new dialogue round is input, the memory read operation activates the relevant memory slots according to the current context, while the gated write mechanism selectively updates or maintains long-term stored entity relationships, thereby maintaining entity consistency and slot integrity in continuous dialogue. This mechanism allows the system to adaptively track key information changes in multi-turn dialogues without relying on an external knowledge base, and dynamically adjust dialogue strategies to maintain the necessary contextual coherence in teaching scenarios. Through adversarial training reinforcement, the relational memory module further learns to distinguish between core teaching entities and temporary information, ensuring that important concepts are not forgotten or distorted in long dialogue sequences. The decoder simultaneously predicts the joint distribution of text tokens and prosodic tokens at the output layer. The joint logits are calculated for time step t , as shown in formula (6).

$$z_t = W_o[h_t; r_t] + \beta_o \quad (6)$$

The joint probability of text and rhythm is obtained by factorization, as shown in formula (7).

$$P(y_t^{\text{text}}, y_t^{\text{pros}} | y_{<t}, X, M) = P(y_t^{\text{text}} | y_{<t}, X, M) \cdot P(y_t^{\text{pros}} | y_{<t}^{\text{text}}, y_{<t}^{\text{pros}}, X, M) \quad (7)$$

In formula (7), y_t^{text} represents a text token, y_t^{pros} represents a prosodic token, and X represents the encoder output. Prosodic markers employ discretized acoustic feature representations, comprising four types of markers: pause boundary markers (indicating pauses at sentence/phrase boundaries), stress intensity markers (indicating syllable stress levels), fundamental frequency contour markers (indicating intonation trends), and duration stretching markers

(indicating syllable extension levels). These markers are generated through a rigorous acoustic processing procedure. First, the original fundamental frequency, frame-level energy, and phoneme duration are extracted from the speech waveform using the Librosa tool. Then, the fundamental frequency undergoes a logarithmic transformation to conform to human auditory perception characteristics, and continuous features are discretized into 16-level quantized markers based on the quantile distribution of the training corpus. Pause positions are determined by aligning the pauses with the prosodic boundary markers using a silence detection algorithm. Within the GAN training framework, prosodic quality is evaluated through a triple mechanism: the prosodic consistency perspective in the discriminator specifically supervises the matching degree between the generated markers and the distribution of the real corpus; feature matching loss constrains the distribution consistency of generated samples in the acoustic feature space; and the human evaluation stage introduces professional speech teachers to blindly score the prosodic naturalness of the synthesized speech, focusing on intonation appropriateness, stress accuracy, and emotional expression suitability in teaching scenarios. In addition, during the joint fine-tuning stage, the system verifies the mapping error between the generated prosodic markers and the target acoustic features through a forced alignment tool, ensuring that the prosodic representation can faithfully reflect the unique tonal characteristics of Thai language and serve the clear pronunciation guidance required for teaching demonstrations.

C. ADVERSARIAL TRAINING MECHANISM BASED ON A MULTI-VIEW DISCRIMINATOR

In the discriminator, the input sample is a concatenated sequence of text tokens and prosodic tokens. The discriminator uses a single shared Transformer encoder to perform computations on, obtaining the encoded representation $\hat{H}$ at each position. The discriminator's input structure includes a complete dialogue history and a combination of the currently generated response, rather than just receiving isolated generated content. The input sequence consists of three parts: the history of the previous N rounds of dialogue (including alternating sequences of teacher questions and student answers), the generated response of the current round, and a special separator between the two. This structural design allows the discriminator to simultaneously observe the interaction between the context and the generated content. For evaluating long-term consistency, the discriminator employs a dual-path attention mechanism: a content scoring perspective uses segmented attention pooling to explicitly calculate the semantic overlap between key entities in the history and those mentioned in the generated response; a fluency scoring perspective analyzes the n -gram transition probability distribution across rounds to capture logical coherence in the dialogue flow. To strengthen long-term dependency modeling, the encoder's underlying attention head is specifically allocated to the

history-response interaction region, enabling the model to identify deep coherence issues such as cross-round topic drift, entity conflict, and emotional inconsistency, rather than being limited to surface-level language fluency evaluation. This input structure and evaluation mechanism design ensures that adversarial training signals effectively guide the generator to maintain the multi-round semantic consistency required for instructional dialogue. This paper sets four discriminative perspectives: content, spoken style, prosodic consistency, and fluency. Each perspective consists of a specific pooling operation and a single-layer feedforward classification/scoring head. For content scoring (semantics and coherence), the segment-level attention pooling (in the form of CLS vector or weighted sum) is performed on the last layer of the encoder to obtain the semantic vector, and the real-valued score is obtained by linear mapping, as shown in formula (8).

$$f_{\text{cont}}(x) = w_{\text{cont}}^T g_{\text{cont}} + \beta_{\text{cont}} \quad (8)$$

In formula (8), w_{cont}^T represents the weight, β_{cont} represents the corresponding bias, and g_{cont} represents the semantic vector. In the oral style scoring (multi-class auxiliary classification), $\hat{H}$ is mean-pooled to obtain g_{style} , and the oral style probability (formal/colloquial/slang/politeness level) is output through the softmax classification head, and the classification logits are also used as the scoring signal, as shown in formula (9).

$$p_{\text{style}}(x) = \text{softmax}(W_{\text{style}} g_{\text{style}} + \beta_{\text{style}}) \quad (9)$$

In formula (9), the style directly takes the maximum class probability. In rhythm consistency scoring: the position vector set $\hat{H}_{\text{pros}}$ of the rhythm token in the encoder is used as input, and the rhythm consistency score is output through a small feedforward network, as shown in formula (10).

$$f_{\text{pros}}(x) = \text{MLP}_{\text{pros}}(\text{Pool}(\hat{H}_{\text{pros}})) \quad (10)$$

In formula (10), Pool represents position aggregation, and the output measures the degree of consistency between the distribution of the text and the prosodic token in the corpus. In the fluency score (based on the discriminator's perplexity estimate), the local n -gram matching statistics of the intermediate features $\hat{H}^{(l)}$ at each layer of the encoder are used, and the fluency score is obtained through linear mapping, as shown in formula (11).

$$f_{\text{flu}}(x) = w_{\text{flu}}^T \text{Pool}_{\text{local}}(\hat{H}^{(1:L)}) + \beta_{\text{flu}} \quad (11)$$

In formula (11), $\text{Pool}_{\text{local}}$ represents the aggregation of local features to approximate the language model perplexity. This paper linearly combines the raw scores of the four views into a total discriminant score for Wasserstein-style adversarial evaluation, as shown in formula (12).

$$F(x) = \sum_{i \in \{\text{cont}, \text{style}, \text{pros}, \text{flu}\}} \omega_i f_i(x) \quad (12)$$

In formula (12), ω_i represents the adjustment weight. To maintain training stability and obtain meaningful continuous scores, the Wasserstein discriminator objective is adopted and a gradient penalty is added, as shown in formula (13).

$$\text{Loss}_D(\phi) = \mathbb{E}_{\tilde{x} \sim p_g}[F_\phi(\tilde{x})] - \mathbb{E}_{x \sim p_{\text{data}}}[F_\phi(x)] + \lambda_{\text{gp}} \mathbb{E}_{\hat{x} \sim p_{\hat{x}}}[(\|\nabla_{\hat{x}} F_\phi(\hat{x})\|_2 - 1)^2] \quad (13)$$

In formula (13), $\hat{x} = \delta x + (1 - \delta)\tilde{x}$ and λ_{gp} represent the gradient penalty coefficients. x represents the interpolated samples along the line connecting the real samples, $\tilde{x}$ represents the interpolated samples along the line connecting the fake samples, and $\nabla_{\hat{x}} F_\phi(\hat{x})$ represents the gradient of the input. The goal of the generator in the adversarial dimension is to maximize the critic's score on the generated samples, as shown in formula (14).

$$\text{Loss}_{\text{adv}}(\theta) = -\mathbb{E}_{\tilde{x} \sim p_g}[F_\phi(\tilde{x})] \quad (14)$$

The RelGAN-Transformer algorithm framework is shown in Figure 2.

As shown in Figure 2, the RelGAN-Transformer algorithm framework consists of two main components: a context-aware generator that leverages segment-level encoding, relational memory, and autoregressive decoding to achieve multi-round joint modeling of context and prosodic information. The generated results are fed into a multi-view discriminator for adversarial evaluation across four dimensions: content consistency, spoken style, prosodic coordination, and fluency. The generator is then optimized through comprehensive feedback. The training process consists of three stages: language pre-training, dialogue fine-tuning, and adversarial fine-tuning. Stabilization scheduling and multi-loss constraints are employed to enhance robustness. The generated responses of the RelGAN-Transformer algorithm framework can be combined with the TTS module to perform prosodic mapping, acoustic modeling, and high-fidelity synthesis. The resulting output is classroom demonstration speech that can be directly used in Thai language teaching scenarios. A fallback and monitoring mechanism is also incorporated to ensure the naturalness, coherence, and controllability of the dialogue system in real-world teaching.

D. MULTI-OBJECTIVE LOSS FUNCTION COMBINATION AND TRAINING PROCESS

The overall training of the experiment is divided into three stages: language pre-training stage (minimizing cross entropy to reduce perplexity and initialize weights); dialogue fine-tuning stage (fine-tuning on dialogue data with MLE + coverage loss to improve repetition/omission); and adversarial fine-tuning stage (using WGAN-GP-style adversarial loss to combine feature matching and auxiliary loss for refinement training) [21]. Parameters are shared between stages, and the dialogue fine-tuning stage provides a stable starting point for the adversarial fine-tuning stage, alleviating the instability of the early stages of adversarial training.

In the language pre-training stage and the dialogue fine-tuning stage, MLE and coverage loss are shown in formula (15) [22-23]. Coverage loss is used to penalize repeated attention/omission.

$$\begin{cases} \text{Loss}_{\text{MLE}}(\theta) = -\sum_{t=1}^T \log P_\theta(y_t | y_{<t}, x), \\ \text{Loss}_{\text{cov}}(\theta) = \sum_{t=1}^T \sum_{i=1}^L \min(A_{t,i}, \hat{a}_i) \end{cases} \quad (15)$$

In formula (15), $\hat{a}_i$ represents the cumulative attention vector, and P_θ represents the conditional probability distribution of the generator at time step t . In the adversarial phase, a differentiable approximation is performed on the discrete tokens, and the Gumbel-Softmax approximation sample is used for the decoder's logits vector at time step t , as shown in formula (16) [24-25].

$$\tilde{y}_t = \text{softmax}\left(\frac{\log \pi_t + \rho_t}{\varrho}\right) \quad (16)$$

In formula (16), each component of ρ_t belongs to Gumbel(0,1), and ϱ represents the temperature parameter. Gradually reducing ϱ to approach discrete sampling in the later stages of training allows the discriminator to continuously input generated samples, thereby propagating the adversarial gradient back to the generator. To prevent the discriminator from being too strong and causing the adversarial signal to be unstable, a feature matching loss is introduced, where the output of an intermediate layer of the shared encoder is set to the mapping $\phi(\cdot)$. The feature matching loss is defined as shown in formula (17).

$$\text{Loss}_{\text{fm}}(\theta) = \|\mathbb{E}_{x \sim p_{\text{data}}}[\phi(x)] - \mathbb{E}_{\tilde{x} \sim p_g}[\phi(\tilde{x})]\|_2^2 \quad (17)$$

The auxiliary loss includes rhythm regression loss and style classification cross entropy. For the labeled sample set, the auxiliary loss formula is shown in (18).

$$\text{Loss}_{\text{aux}}(\theta) = \mathbb{E}_{(x,y) \in D_{\text{lab}}} [\|\text{pros}(G_\theta(x)) - y_{\text{pros}}\|^2 - \zeta \log p_{\text{style}}(y_{\text{style}} | G_\theta(x))] \quad (18)$$

In formula (18), pros represents the generator's regression output for the prosodic token, p_{style} represents the style classification probability, and ζ represents the corresponding weight. This paper uses a regularization term in the form of maximum mutual information to encourage mutual information between the output and input, addressing the issues of "safe sentences" and simplification, as shown in formula (19).

$$\text{Loss}_{\text{div}}(\theta) = \log P_\theta(y | x) - \iota \log P_\theta(y) \quad (19)$$

In formula (19), ι controls the intensity of the prior penalty. In practical implementation, this paper employs a batch statistical estimation method to approximate the




FIGURE 2. RelGAN-Transformer Algorithm Framework

mutual information regularization term. For each training batch, the marginal probability $P_\theta(y)$ is estimated using the statistical distribution of samples within the batch, and the frequency distribution of each output label in the current batch is used as an approximation of $P_\theta(y)$. This batch-based estimation method avoids the need to train independent marginal language models while ensuring computational feasibility. During optimization, the batch statistics are smoothed using an exponential moving average technique to reduce the variance of small batch estimates. In this way, this paper can effectively achieve the goal of maximizing mutual information without introducing additional model components, thereby alleviating the "safe statement" problem while maintaining the stability and reproducibility of the training process. The final goal of the generator is the weighted sum of the various losses, as shown in formula (20).

$$\begin{aligned} \text{Loss}_G(\theta) = & \text{Loss}_{\text{MLE}} + \lambda_{\text{cov}} \text{Loss}_{\text{cov}} + \lambda_{\text{adv}} \text{Loss}_{\text{adv}} \\ & + \lambda_{\text{fm}} \text{Loss}_{\text{fm}} + \lambda_{\text{aux}} \text{Loss}_{\text{aux}} + \lambda_{\text{div}} \text{Loss}_{\text{div}} \end{aligned} \quad (20)$$




FIGURE 3. Training Curves (a) Discriminator and Generator Loss Curves over Iterations; (b) Comparison of Perplexity and Distinct-1 Curves over Iterations

In formula (20), λ_{cov} , λ_{adv} , λ_{fm} , λ_{aux} , and λ_{div} represent the corresponding loss weights, and $Loss_G(\theta)$ represents the total loss. The following training schedule is used to ensure convergence: during the adversarial phase, five discriminator updates are performed before each generator update, and the gradient penalty term is enabled in the discriminator update. The AdamW optimizer (learning rate warmup + cosine decay) [26–28] is used, and the norm of the gradient is clipped to prevent explosion. During adversarial training, comprehensive metrics on the validation set are used as early stopping criteria to avoid overfitting and style shift. If the discriminator loss rapidly decreases and the quality of generated samples degrades during training, the weight of λ_{fm} is increased or the discriminator learning rate is temporarily reduced to restore balance. Early in training, the weight of λ_{fm} is increased to prioritize representation alignment, and later, λ_{aux} is gradually increased to strengthen style and prosodic constraints. The training curve is shown in Figure 3. In Figure 3, this study divides the training process into three stages. Stage 1 (Phase 1, iterations 1–3000) is the pre-training stage for the generator and discriminator, mainly used to allow the model to learn basic language patterns and prosodic features. Stage 2 (Phase 2, iterations 3001–6000) is the dialogue fine-tuning stage, which enhances multi-turn contextual coherence and entity slot tracking capabilities through maximum likelihood and coverage constraints. Stage 3 (Phase 3, iterations 6001–10000) is the adversarial fine-tuning stage, which uses RelGAN relational memory and a multi-view discriminator for adversarial training to improve text diversity, naturalness of spoken style, and long-range dependency modeling capabilities.

As shown in Figure 3 (a), the discriminator loss decreases rapidly and stabilizes during pre-training. The generator loss gradually decreases during the dialogue fine-tuning phase. During the adversarial fine-tuning phase, it experiences slight fluctuations due to game enhancement but ultimately converges. This demonstrates that adversarial training effectively improves the robustness and expressiveness of the generator. In Figure 3(b), Perplexity decreases continuously during training and improves significantly after fine-tuning the dialogue, demonstrating improved fluency and

predictability. Furthermore, the Distinct-1 metric increases significantly during the adversarial phase, indicating that the model generates Thai spoken responses with greater diversity while maintaining naturalness. These overall results demonstrate that proposed approach achieves an optimal balance between naturalness and diversity in generated responses.

E. SPEECH SYNTHESIS INTERFACE FOR TEACHING DEMONSTRATIONS

The interface between the Generator and TTS uses the generator's "text tokens + prosody tokens" as the sole input. A forced alignment tool first creates a mapping table for text, duration, pitch (F0), and energy. This mapping informs a learnable prosodic embedding table that converts the generator's prosodic tokens into fixed-dimensional vectors. These prosodic, text, and entity slot embeddings are fused at the input of the non-autoregressive acoustic model, FastSpeech2. Prosodic input explicitly controls the duration predictor and a residual branch from the prosodic embedding is added to the F0 and energy prediction heads to constrain tone and stress output. The HiFi-GAN vocoder synthesizes the final waveform from predicted acoustic features. The training follows a two-stage process: first, the acoustic model converges independently on paired speech and text with prosodic annotations; then, joint fine-tuning is performed by freezing the backbone generator and updating only the prosodic mapping layer and the acoustic front-end. This strategy minimizes end-to-end error propagation onto the generator's semantic output while ensuring the generated text's pronunciation and intonation are reliably reproduced by the TTS.

III. EXPERIMENTAL SETUP AND EVALUATION PLAN

A. EXPERIMENTAL DATASET

This study aggregates text and speech data from multiple sources to construct training/validation/test sets, which support language pre-training, dialogue fine-tuning, ASR robustness testing, and TTS fine-tuning, respectively. For text corpus, the paper used the HSE Thai Corpus (a modern Thai written language corpus of approximately 50 million words)

for language model pre-training and subword list training. It also introduced Krathu-500 (a multi-turn post-comment dialogue corpus consisting of 574 Pantip posts and approximately 63,293 comments) as the primary text source for multi-turn dialogue fine-tuning and response style learning. For speech and spoken dialogue corpus, the paper used the Thai Dialect Corpus (approximately 840 hours of Thai reading and speaking speech corpus, covering both the central and several regional dialects) and several spontaneous spoken dialogue collections for ASR style noise modeling and TTS training. This includes a speech dataset for spontaneous smartphone conversations (322 speakers covering over 20 domains, designed to enhance spoken language diversity and speaker diversity) and a collection of approximately 102 hours of spontaneous Thai conversational speech for end-to-end acoustic training and evaluation. All data is split into training/validation/test sets at an 80%/10%/10% ratio based on conversational completeness. The test set retains manually annotated gold labels for naturalness, coherence, and prosody for subjective evaluation and final comparative validation. The Thai spoken dialogue corpus used in this study is multi-source and heterogeneous, fully covering the diverse linguistic features required in teaching scenarios. All spoken materials underwent a rigorous prosodic feature annotation process. Fundamental frequency, duration, and energy parameters were extracted using the Librosa tool, and the fundamental frequency was logarithmically transformed and quantized. Finally, prosodic markers (including pause positions, stress intensity, and intonation boundaries) were systematically embedded into the text sequence. The corpus emphasizes speaker diversity, balancing the distribution of gender, age, regional accents, and educational background, particularly reinforcing typical teacher-student interaction patterns in educational scenarios. Data partitioning strictly adheres to the principle of dialogue integrity, ensuring that continuous dialogue turns are not interrupted.

B. BASELINE MODEL AND EVALUATION METRICS

In this paper, the algorithm used is the RelGAN-Transformer algorithm, and the benchmark algorithms include SeqGAN, LeakGAN (Leak Generative Adversarial Network), MaliGAN (Maximum-Likelihood Augmented Discrete Generative Adversarial Networks), MetaEx-GAN, Transformer-XL (Transformer for Extra Long contexts), RelGAN, and cGAN-Transformer-CVAE. The evaluation indicators in this paper include BLEU, ROUGE, METEOR/ChrF, Perplexity, Distinct-1/2, Entity F1, etc. The calculation formula of METEOR is shown in formula (21).

$$\text{METEOR} = (1 - \text{Penalty}) F_{\text{mean}} \quad (21)$$

In formula (21), Penalty represents the word order penalty, and F_{mean} represents the weighted harmonic mean of precision and recall. The calculation formula for Distinct-1/2 is shown in formula (22).

$$\begin{aligned} \text{Distinct1} &= \frac{\text{Unique 1-grams}}{\text{Total 1-grams}}, \\ \text{Distinct2} &= \frac{\text{Unique 2-grams}}{\text{Total 2-grams}} \end{aligned} \quad (22)$$

In formula (22), Unique 1-grams represents the number of unique 1-grams in the generated text, and Total 1-grams represents the number of all 1-grams in the generated text.

C. EXPERIMENTAL DESIGN AND PARAMETER SETTINGS

To ensure comparability and reproducibility, all comparison models were trained using the same preprocessing and word segmentation strategies. A unified language model was first pre-trained on a large-scale Thai monolingual corpus (using the same SentencePiece subword table and character-subword mixed representation), and then trained/fine-tuned on the same multi-turn dialogue fine-tuning corpus. Each model uses the same training/validation/test split (80%/10%/10%), the same random seed set, and the same hyperparameter search budget. Adversarial methods uniformly employ temperature-annealed Gumbel-Softmax or equivalent continuous relaxation to ensure fair comparison. The training strategy records the number of training steps, early stopping points, and final checkpoint for all models. During inference, a unified beam search (beam size = 5, with adjustable length penalty) and a consistent confidence threshold backoff mechanism are used to ensure fair comparison between automated metrics and human evaluation. The evaluation is divided into two streams: automatic quantitative experiments and manual teaching adaptation evaluation. Automatic metrics include BLEU (Bilingual Evaluation Understudy), ROUGE (Recall-Oriented Understudy for Gisting Evaluation), METEOR/ChrF (Character n-gram F-score), Perplexity, Distinct-1/2, Entity F1, and robustness to noisy ASR input. For each metric, five independent training/evaluation replicates were performed on the test set, the mean value was taken, and significance tests were performed (significance level 0.05). Manual evaluation was conducted by at least 30 native Thai-speaking reviewers, each blindly reviewing at least 200 examples, with each example cross-scored by three reviewers. Deployment latency and resource consumption performance were also verified to assess engineering usability and classroom suitability. The experimental parameter settings are shown in Table 1.

TABLE 1. Experimental Parameters

Parameters	Setting values	Parameters	Setting values
Maximum sequence length	128	AdamW weight decay coefficient	1.00E-02
Embedding dimension	512	Discriminator update frequency	5
Feedforward layer dimension	2048	WGAN gradient penalty coefficient	10
Number of Transformer layers (generator)	6	Adversarial loss weight	1
Number of multi-head attention heads	8	Feature matching weight	10
Number of relational memory slots	16	Auxiliary loss weight	2
Relational memory dimension	128	Coverage loss weight	1
Number of shared encoding layers for the discriminator	4	Diversity weight	0.5
Dropout ratio	0.1	Continuous relaxation initial temperature	1
Batch size	64	Minimum temperature after annealing	0.1
Generator learning rate	1.00E-04	Beam search width	5
Discriminator learning rate	5.00E-05	Random seed	42




FIGURE 4. Overall Evaluation of Thai Spoken Dialogue Generation Performance

IV. EXPERIMENTAL RESULTS AND ANALYSIS

A. OVERALL PERFORMANCE COMPARISON

This paper compares the overall performance of several generative models on the Thai spoken dialogue task, using a unified test set and identical preprocessing. The results are shown in Figure 4.

As shown in Figure 4, the traditional generative adversarial networks, SeqGAN and MaliGAN, only achieved BLEU scores of 0.42 and 0.40, respectively. LeakGAN achieved a slightly higher score of 0.45, while MetaEx-GAN improved to 0.49. Transformer-XL further reached 0.54. RelGAN and cGAN-Transformer-CVAE achieved scores of 0.57 and 0.60, respectively. The proposed RelGAN-Transformer algorithm achieved the highest score of 0.68. The same trend is also observed in the ROUGE, METEOR, and ChrF metrics. For example, in ROUGE, the RelGAN-Transformer algorithm achieves 0.78, while SeqGAN only achieves 0.48. The ChrF metric shows a significant difference, with the RelGAN-Transformer algorithm achieving 0.83 and the MaliGAN algorithm only 0.50. This demonstrates that proposed method has a clear advantage in matching the semantics and structure of the generated text with the reference response. On the Perplexity scale, the RelGAN-Transformer algorithm achieved the lowest score of 18.95, compared to Transformer-XL's 25.1 and MetaEx-GAN's 33.55, demonstrating more fluent and natural generated text and significantly reduced model perplexity. This




FIGURE 5. Diversity Improvement Analysis

demonstrates that the generator outperforms other models in capturing multi-round context and long-range dependencies. Overall, traditional GAN models have significant shortcomings in text generation response diversity and long-range dependency modeling. However, the RelGAN-Transformer algorithm, based on the Transformer structure and enhanced relational memory, leads in all metrics.

B. VERIFICATION OF THE IMPROVED GENERATION DIVERSITY

To evaluate the performance of each model in generating response diversity, the paper used four metrics: Distinct-1, Distinct-2, Type-token ratio, and repetition rate (the proportion of repeated sentences/fragments) for quantitative comparison. The results are shown in Figure 5.

In Figure 5, the RelGAN-Transformer algorithm performs best in terms of text generation diversity, with Distinct-1 and Distinct-2 reaching 0.14 and 0.11, respectively. This represents improvements of 0.11 and 0.10, respectively, compared to SeqGAN, demonstrating a significant increase in the diversity of single-word and biword sequences. In terms of type-token ratio, the RelGAN-Transformer algorithm achieved 0.42, while SeqGAN only achieved 0.18, indicating a significant increase in the proportion of diverse vocabulary in the generated text, resulting in richer textual expression. Regarding the repetition rate metric, the RelGAN-Transformer algorithm achieved the lowest

score of 0.09, a decrease of 0.19 from LeakGAN's 0.28, indicating a significant reduction in the number of repeated sentences and fragments generated. Other models, such as Transformer-XL, RelGAN, and cGAN-Transformer-CVAE, outperform traditional adversarial models (SeqGAN and MaliGAN) in both Distinct metrics and Type-token ratios, while also reducing repetition rates. This demonstrates that the Transformer architecture and relational memory offer significant advantages in improving diversity. Traditional SeqGAN and MaliGAN rely on an RNN (Recurrent Neural Network) generator and employ only a single discriminator for sequence-level rewards. This makes it difficult to generate diverse vocabulary and sentence structures, resulting in low Distinct values and high repetition rates. LeakGAN slightly improves local diversity through hierarchical rewards, but is still limited by its overall ability to capture context. MetaEx-GAN improves generative diversity by leveraging multi-task adversarial networks and an external experience library, but lacks explicit relational memory and long-term dependency modeling, resulting in limited improvements in the type-token ratio and Distinct metrics. Transformer-XL introduces a self-attention mechanism to improve long-term dependencies. RelGAN enhances cross-round information fusion through relational memory.

cGAN-Transformer-CVAE improves the diversity of generated responses through a conditional variational mechanism. However, the RelGAN-Transformer algorithm combines multi-view adversarial training, relational memory, and Transformer self-attention. While optimizing naturalness and cross-round coherence, it also significantly increases vocabulary and sentence diversity and reduces repetition. This demonstrates that the multi-view adversarial signal and relational memory module can effectively prevent generative model collapse and promote richer and more diverse expressive styles in Thai spoken conversation. To explore the diversity advantage of proposed algorithm, it conducted a T-test. Null hypothesis H_0 : The overall mean Distinct-2 score of the RelGAN-Transformer algorithm and the comparison algorithm is equal; alternative hypothesis H_1 : The two means are not equal (two-sided test). Each algorithm was independently trained and evaluated five times under the same training/evaluation conditions (with different random seeds), and the Distinct-2 score was recorded each time. Independent-sample t-tests (assuming the variances of the two groups can be pooled) were used to compare the Distinct-2 distributions of the RelGAN-Transformer algorithm with each comparison algorithm. t-scores, two-sided p-values, and 95% confidence intervals for the differences were calculated. The significance level is 0.05. The results of the significance tests are shown in Table 2.

The significance test results in Table 2 show that the RelGAN-Transformer algorithm significantly outperforms most of the comparison algorithms on the Distinct-2 metric. Compared with SeqGAN, LeakGAN, and MaliGAN, the mean value of the RelGAN-Transformer algorithm is 0.11, while the values of the comparison algorithms are only 0.01-

0.02. Their t-values are 10.85, 9.64, and 9.50, respectively. All p-values are less than 0.001, with 95% confidence intervals of (0.08, 0.12), (0.07, 0.11), and (0.07, 0.11), respectively. Compared with MetaEx-GAN, Transformer-XL, and RelGAN, the t-values are 7.00, 5.75, and 3.58, respectively, with p-values of <0.001, <0.001, and 0.0072, respectively. All confidence intervals are positive. Compared to cGAN-Transformer-CVAE, the t-value is 1.66, the p-value is 0.1351, and the 95% confidence interval is (-0.01, 0.05). Although the difference is not significant, the mean of the RelGAN-Transformer algorithm is still slightly higher. This indicates that the proposed algorithm generally outperforms traditional generative adversarial models and most state-of-the-art Transformer-class models in generating binary text sequences, demonstrating a strong and stable advantage in terms of statistical significance.

C. CROSS-TURN DIALOGUE COHERENCE ANALYSIS

To evaluate the model's ability to maintain entity and intent consistency across multiple conversations, the paper analyzed the entity recognition performance and intent/slot metrics of each model on the entire test set. It then presented the average entity coverage over different conversation turns, reflecting the model's long-term memory and cross-turn coherence. The results are shown in Figure 6. Figure 6 shows the lateral coherence metrics, which employ a unified slot-filling + Named Entity Recognition (NER) annotation system. All dialogues are first framed using manually defined task-related slots (people/locations/times/objects, etc.). Standardized entity labels are automatically generated by a unified Thai NER/slot annotator. The model-generated text also uses the same annotator to extract entities and aligns them with the test set's golden labels to calculate precision/recall/F1. To control automatic annotation noise, the test set employs an "automatic annotation + random sampling and manual review" strategy, with sampling consistency (Cohen's kappa) > 0.85. Intent retention is based on predefined teaching intent categories (e.g., asking/explaining/confirming/feedback). The "initial intent" of each dialogue is independently determined and arbitrated by two trained native Thai annotators to form a golden label (Cohen's kappa > 0.86). In subsequent rounds, if the model's predicted intent matches the current dominant intent, it is counted as "retention"; otherwise, it is considered cross-round drift. The slot conflict rate is strictly defined as the proportion of inconsistent values (including values that do not match the previous confirmed value or multiple mutually exclusive values given in a single round) for a confirmed slot in the same session. When calculating the rate, the most recently confirmed slot-value pair is used as the comparison benchmark. No empirical threshold is set, and the count is based solely on the consistency/inconsistency rules.

In Figure 6 (a), the RelGAN-Transformer algorithm far outperforms other methods in entity recognition and retention, achieving Entity Precision of 95.23%, Entity Recall of 93.10%, and Entity F1 of 94.15%. cGAN-Transformer-

TABLE 2. Significance Test Results

Comparison algorithm	RelGAN-Transformer Mean $\pm$ std	Comparison Mean $\pm$ std	t-value	p-value	95% confidence interval (difference)
SeqGAN	0.11 $\pm$ 0.02	0.01 $\pm$ 0.01	10.85	< 0.001	(0.08, 0.12)
LeakGAN	0.11 $\pm$ 0.02	0.02 $\pm$ 0.01	9.64	< 0.001	(0.07, 0.11)
MaliGAN	0.11 $\pm$ 0.02	0.02 $\pm$ 0.01	9.50	< 0.001	(0.07, 0.11)
MetaEx-GAN	0.11 $\pm$ 0.02	0.04 $\pm$ 0.01	7.00	< 0.001	(0.05, 0.09)
Transformer-XL	0.11 $\pm$ 0.02	0.05 $\pm$ 0.01	5.75	< 0.001	(0.04, 0.08)
RelGAN	0.11 $\pm$ 0.02	0.07 $\pm$ 0.02	3.58	0.0072	(0.01, 0.07)
cGAN-Transformer-CVAE	0.11 $\pm$ 0.02	0.09 $\pm$ 0.02	1.66	0.1351	(-0.01, 0.05)




FIGURE 6. Cross-Turn Coherence Evaluation Results. (a) Entity Performance; (b) Intent Preservation Rate and Slot Conflict Rate; (c) Average Entity Coverage Variation with Dialogue Turns

CVAE achieved 89.10%/86.30%/87.68% respectively, and Transformer-XL achieved 83.20%/80.45%/81.80%. Proposed algorithm significantly outperformed traditional adversarial methods, with SeqGAN achieving only 56.34%/49.12%/52.48% Entity F1, and LeakGAN and MaliGAN achieving only 56.75% and 54.20% respectively. The simultaneous improvement in both precision and recall demonstrates that the RelGAN-Transformer algorithm not only produces fewer false positives (incorrectly labeled entities) but also misses fewer real entities. It also more accurately captures and recovers key nouns and slots that appear in multi-turn conversations. In Figure 6(b), the RelGAN-Transformer algorithm also performs best in terms of intent preservation and slot consistency, achieving an intent preservation rate of 97.5%, significantly higher than the 87.4% of cGAN-Transformer-CVAE and the 80.1% of Transformer-XL. The slot conflict rate is only 1.12%, significantly lower than the 5.10% of Transformer-XL and the 12.34% of SeqGAN.

The RelGAN-Transformer algorithm maintains extremely high consistency in cross-round information transmission. The model can basically use previously confirmed intents/slots when generating subsequent rounds. Slot inconsistencies or entity replacement errors within the same

session are rare. It is more suitable for teaching application scenarios that require stable context. Figure 6(c) shows that all models experience a decrease in entity coverage with increasing rounds, but the RelGAN-Transformer algorithm shows the smallest decrease, reaching 95.0% in round 1 and 65.45% in round 10, for a cumulative decrease of 29.55%. In contrast, Transformer-XL's coverage dropped from 85.1% to 55.0% (a decrease of 30.1%), RelGAN's from 87.5% to 56.6% (a decrease of 30.9%), and SeqGAN's from 63.0% to 22.5% (a decrease of 40.5%). Even in the 10th round of long dialogue, the RelGAN-Transformer algorithm still maintained 65.45% entity coverage, significantly outperforming most models. This demonstrates its superior long-term memory and cross-round information retrieval capabilities, enabling it to retain and correctly reproduce previously encountered entities across longer conversations.

D. SUBJECTIVE SPEECH NATURALNESS (MOS)

The experiment involved blind review by native Thai speakers, who rated the system-generated responses on a scale of 1–5 based on pedagogical suitability, spoken naturalness, pronunciation demonstration usability, semantic relevance, coherence, lexical diversity, prosodic naturalness, and comprehensibility. This was used to compare the subjective performance of the model in teaching scenarios. The review was conducted by at least 30 reviewers, with each sample cross-scored by three reviewers. The results of the manual subjective evaluation are shown in Figure 7. To verify the reliability of the subjective evaluation results, this paper conducted a statistical analysis of the consistency among reviewers. For each rating dimension, the intraclass correlation coefficient (ICC) and weighted Cohen's kappa coefficient were calculated among the three reviewers. The results show that the ICC for each subjective dimension is between 0.78 and 0.84, and the weighted kappa coefficient is above 0.75, indicating high consistency and stability among reviewers. Therefore, the cross-rating results from multiple reviewers have good reliability and can be used for effective comparison of subjective performance between models.

In Figure 7, RelGAN-Transformer performs best across all instructionally relevant subjective dimensions. Its instructional suitability scores are 3.85 ± 0.22 , spoken naturalness is 4.05 ± 0.20 , pronunciation demonstration usability is 3.95 ± 0.23 , semantic relevance is 4.10 ± 0.19 , coherence




FIGURE 7. Human Subjective Evaluation Results

is 4.00 ± 0.21 , lexical diversity is 3.90 ± 0.22 , prosodic naturalness is 3.85 ± 0.24 , and comprehensibility is 4.05 ± 0.18 . RelGAN-Transformer significantly outperformed the runner-up cGAN-Transformer-CVAE (3.35–3.65) and Transformer-XL (2.95–3.40), and also significantly surpassed the traditional adversarial models SeqGAN and MaliGAN. This demonstrates that in teaching scenarios, the Thai spoken responses generated by RelGAN-Transformer not only have significant advantages in phonetic demonstration, semantic coherence, and lexical diversity, but also have higher overall comprehensibility and naturalness, better supporting student listening and speaking training and practical teaching applications.

E. NOISE ROBUSTNESS AND ASR ADAPTABILITY ANALYSIS

This experiment tested the robustness of the system under varying signal-to-noise ratios (SNRs) and various noise types, quantifying the impact of noise on generation diversity, intent preservation, and generation quality. Environmental noise is categorized into classroom, street, and coffee shop noise; device noise is categorized into microphone, electrical appliance, and recording equipment; and speaker noise is categorized into low volume, fast speech rate, and accent. The results of the ASR robustness analysis are shown in Figure 8.

In Figure 8 (a), the Distinct-2 metric shows a significant downward trend with decreasing SNR and is more sensitive to speaker characteristics. Under high signal-to-noise ratio conditions (SNR = 30 dB), Distinct-2 values for classrooms, cafes, and recording equipment are mostly around 0.11, while street and electrical noise are slightly lower, around 0.10. When the SNR drops to 5 dB, the Distinct-2 values for classroom, cafeteria, and recording equipment drop to 0.06, 0.05, and 0.05, respectively. For fast speech and accented speech, the values are only 0.04 and 0.03, respectively, at 5 dB. This indicates that even under extreme conditions, the diversity of bigram phrases decreases somewhat, but still demonstrates a certain degree of noise immunity. In moderate noise (SNR = 20 dB), the impact of ambient noise on Distinct-2 remains small (approximately 0.09–0.10), but speaker characteristics

(low volume/fast speaking rate/accents) reduce Distinct-2 to 0.07–0.08, indicating that diversity is more easily lost when faced with changes in speaking style. In Figure 8(b), the intent preservation rate generally decreases with decreasing SNR, but the extent of the decrease varies across different types of noise. Under optimal conditions (SNR = 30 dB), RelGAN-Transformer maintains 97.5% and 96.5% intent preservation in classroom and cafeteria environments, and even in street noise, it maintains 95.5%. When the SNR drops to 10 dB, intent preservation in the classroom drops to 87%, while it drops to 76%, 77%, and 72% for low-volume speech, fast speech, and accented speech, respectively. Under the worst-case SNR of 5 dB, intent preservation in the classroom is only 78%, while it drops to 63% for speech with an accent. This shows that intent preservation is highly robust in the presence of ambient noise, but is extremely sensitive to speaker characteristics (especially accents and low volume), and intent loss increases significantly at low signal-to-noise ratios. It should be noted that the METEOR values reported in this paper were obtained under a specific teaching/service-oriented Thai oral dialogue task, which differs from the completely open-domain free chat generation. This dialogue task has a relatively clear intent space and slot structure, and each user input in the test set corresponds to multiple human reference responses (multi-reference evaluation). These responses are highly consistent semantically, differing only in surface wording and prosodic expression. In addition, the METEOR calculation employs Thai synonym expansion and stemming strategies, enabling a high matching score to be obtained while maintaining semantics but changing word form. Therefore, the scores of 0.71 (classroom) and 0.67 (street) obtained under SNR=30 dB do not mean that the model output tends to be templated, but rather reflect that the model can stably reconstruct task-related semantic content under high ASR accuracy. At the same time, diversity indicators such as Distinct-2 remain in the range of 0.10–0.11, indicating that the generated results do not significantly sacrifice expressive diversity while ensuring semantic consistency. In Figure 8(c), METEOR decreases with increasing noise, decreasing in tandem with Distinct-2 and intent, reflecting a deterioration in semantic-morphological alignment. At an SNR of 30 dB, the METEOR for classroom audio is 0.71, for street audio 0.67, and for audio with an accent 0.60. At an SNR of 5 dB, the METEOR for classroom audio drops to 0.53, while for audio with an accent it is only 0.38, indicating a significant decline in semantic reconstruction quality under the conditions of an accent. Under moderate noise conditions (SNR = 20 dB), the METEOR remains between 0.64–0.68 despite ambient noise, while the METEOR due to speaker idiosyncrasies is only 0.55–0.58. This indicates that when ASR errors represent factual substitutions or deletions, the quality of the system's semantic reconstruction (and downstream generation) is somewhat compromised.




FIGURE 8. ASR Robustness Analysis Results: (a) Distinct-2; (b) Intent Preservation Rate;(c) METEOR;

F. SYSTEM RESOURCE CONSUMPTION AND RESPONSE LATENCY EVALUATION

This experiment investigates the resource and latency evaluation of each model during inference deployment to assess its engineering usability in teaching scenarios. The resource and latency evaluation results are shown in Table 3. The energy consumption metrics (J/query) in Table 3 were measured in a unified GPU inference environment, specifically using a single NVIDIA GPU (with dynamic frequency adjustment disabled) for online inference testing. Power consumption was recorded in real-time at fixed sampling intervals using the NVIDIA System Management Interface (nvidia-smi). During the stable inference phase, the power curve was integrated over time, and finally divided by the number of queries completed to obtain the energy consumption per query (J/query). This energy consumption statistic only includes GPU computational power consumption during the model inference phase and does not include CPU-side preprocessing, data loading, or system idle power consumption. Therefore, the values mainly reflect the relative energy efficiency of the model inference itself and can be used for fair comparisons between different models.

As shown in Table 3, lightweight generative models such as SeqGAN have an average response time of only 85.12 ms, an inference latency of 70.45 ms, a throughput of 142 samples/s, a model parameter count of only 4.5M, and an energy consumption of 1.8 J/query, demonstrating extremely high real-time performance. Large Transformer models such as Transformer-XL, RelGAN, and cGAN-Transformer-CVAE have response times of 150.1, 160.88, and 170.32 ms, respectively. They have inference latencies of 130.05–150.12 ms, throughput of only 70–78.5 samples/s, 60–85M parameters, and energy consumption of 9.2–11.8 J/query, indicating a high computational burden. RelGAN-

Transformer strikes a balance between performance and resources, with an average response time of 98.44 ms and an inference latency of 88.2 ms. It achieves a throughput of 114 samples/s, 32.5M model parameters, and energy consumption of 3.6 J/query. Compared to Transformer-XL, this reduces average response time by approximately 51.66 ms while maintaining reasonable throughput and a moderate model size. This demonstrates that RelGAN-Transformer strikes a balance between generation quality and engineering usability in teaching scenarios, achieving low latency and acceptable energy consumption while maintaining high performance in multi-turn dialogue generation. The inference latency in Table 3 is not linearly related to the model parameter size. All models were tested on the same hardware platform, with the same batch size (batch size = 1) and the same beam search width (beam size = 5), without model pruning or quantization. Although RelGAN-Transformer has a larger parameter size than SeqGAN, its generation process mainly relies on highly parallel Transformer self-attention computation and does not include the additional sequential overhead caused by stepwise sampling and discriminator interaction in SeqGAN. At the same time, compared with Transformer-XL, it does not introduce cross-segment caching and long context memory updates, thus significantly reducing the computation and memory access overhead during single-step inference. Therefore, under the same inference configuration, RelGAN-Transformer exhibits latency characteristics between lightweight GAN models and large-scale Transformer models.

TABLE 3. Resource and Latency Evaluation

Algorithm	Average response time (ms)	Inference latency (ms)	Throughput (samples/s)	Model parameters (M)	Energy consumption (J/query)
SeqGAN	85.12	70.45	142	4.5	1.8
LeakGAN	95.38	80.12	125	6.2	2.1
MaliGAN	90.27	75.3	132.5	5	1.95
MetaEx-GAN	140.75	120.33	83.2	45	8.5
Transformer-XL	150.1	130.05	78.5	60	9.2
RelGAN	160.88	140.6	74	70	10.5
cGAN-Transformer-CVAE	170.32	150.12	70	85	11.8
RelGAN-Transformer	98.44	88.2	114	32.5	3.6

V. CONCLUSION

This paper constructs a Thai spoken dialogue generation system based on RelGAN-Transformer, which explicitly maintains cross-turn context using a Transformer self-attention encoder and segment-level recurrent relation memory slots on the generator side; this architectural principle is transferrable to the textile industry, where it could be utilized to maintain diagnostic coherence across continuous production sequences by tracking multi-sensor inputs (e.g., tension, vibration, image data) via a similar relational memory mechanism. It employs MLE+coverage pre-training/fine-tuning and a WGAN-GP multi-view Transformer discriminator under Gumbel-Softmax continuous relaxation, combined with feature matching and prosody/style supervision for adversarial fine-tuning. This system achieves simultaneous text and prosody generation and synthesizes high-fidelity demonstration speech using FastSpeech2+HiFi-GAN. Experimental results show that the system achieved a Distinct-2 score of 0.11 and an intent retention rate of 97.5%, a 10.1% improvement over cGAN-Transformer-CVAE. The average human MOS score for spoken naturalness was 4.05, significantly outperforming the comparison system. The system also improved the diversity of generated responses and cross-turn coherence. A limitation of this paper is that the generation system still has a certain degree of dependence on extremely long contexts and complex conversational scenarios. Future work could further improve continuous conversation performance through deeper relational memory or reinforcement learning to optimize the generation strategy. For the textile industry, this translates to leveraging deeper relational memory models to track intricate, long-term machinery wear and raw material variations, and applying reinforcement learning (RL) to continuously optimize the autonomous diagnostic and control strategy for complex processes like weaving or dyeing, maximizing product quality and minimizing resource waste over extended production cycles.

REFERENCES

- [1] T. Chen, "Design and Application of English Oral Online Dialogue System Based on Reinforcement Learning Algorithm," *Procedia Computer Science*, vol. 261, no. 1, pp. 716–723, 2025, doi: <https://doi.org/10.1016/j.procs.2025.04.325>.
- [2] S. Bibauw, W. Van den Noortgate, T. Francois, and P. Desmet, "Dialogue systems for language learning: A meta-analysis," *Language Learning & Technology*, vol. 26, no. 1, pp. 1–24, 2022, doi: <https://doi.org/10.64152/10125/73488>.
- [3] M. H. Hsu, P. S. Chen, and C. S. Yu, "Proposing a task-oriented chatbot system for EFL learners speaking practice," *Interactive Learning Environments*, vol. 31, no. 7, pp. 4297–4308, 2023, doi: <https://doi.org/10.1080/10494820.2021.1960864>.
- [4] C. Zhai and S. Wibowo, "A WGAN-based dialogue system for embedding humor, empathy, and cultural aspects in education," *IEEE Access*, vol. 11, no. 1, pp. 71940–71952, 2023, doi: <https://doi.org/10.1109/ACCESS.2023.3294966>.
- [5] A. C. Le and V. N. Huynh, "Enhancing conversational model with deep reinforcement learning and adversarial learning," *IEEE Access*, vol. 11, no. 1, pp. 75955–75970, 2023, doi: <https://doi.org/10.1109/ACCESS.2023.3297652>.
- [6] M. Firdaus, N. Thangavelu, A. Ekbal, and P. Bhattacharyya, "I enjoy writing and playing, do you?: A personalized and emotion grounded dialogue agent using generative adversarial network," *IEEE Transactions on Affective Computing*, vol. 14, no. 3, pp. 2127–2138, 2022, doi: <https://doi.org/10.1109/TAFFC.2022.3155105>.
- [7] M. Firdaus, A. Madasu, and A. Ekbal, "A unified framework for slot based response generation in a multimodal dialogue system," *Multimedia Tools and Applications*, vol. 83, no. 4, pp. 11643–11667, 2024, doi: <https://doi.org/10.1007/s11042-023-15915-8>.
- [8] A. Fuad and M. Al-Yahya, "AraConv: Developing an Arabic task-oriented dialogue system using multi-lingual transformer model mT5," *Applied Sciences*, vol. 12, no. 4, pp. 1881–1896, 2022, doi: <https://doi.org/10.3390/app12041881>.
- [9] X. Li, T. Liu, L. Zhang, F. Alqahtani, and A. Tolba, "A transformer-BERT integrated model-based automatic conversation method under English context," *IEEE Access*, vol. 12, no. 1, pp. 55757–55767, 2024, doi: <https://doi.org/10.1109/ACCESS.2024.3388100>.
- [10] Y. Zhao, B. Cheng, Y. Huang, and Z. Wan, "FluGCF: a fluent dialogue generation model with coherent concept entity flow," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 32, no. 1, pp. 853–867, 2023, doi: <https://doi.org/10.1109/TASLP.2023.3340610>.
- [11] K. Manohar and R. Rajan, "Improving speech recognition systems for the morphologically complex Malayalam language using subword tokens for language modeling," *EURASIP Journal on Audio, Speech, and Music Processing*, vol. 2023, no. 1, pp. 47–71, 2023, doi: <https://doi.org/10.1186/s13636-023-00313-7>.
- [12] M. Bhadri, "BiLSTMs and BPE for English to Telugu CLIR," *J. Electrical Systems*, vol. 20, no. 3, pp. 2022–2029, 2024, doi: <https://doi.org/10.52783/jes.1798>.
- [13] Y. Qiu, L. Dong, W. Zhang, H. Xing, and J. Huang, "A diffusion enhanced CRF and BiLSTM framework for accurate entity recognition," *Scientific Reports*, vol. 15, no. 1, pp. 1–25, 2025, doi: <https://doi.org/10.1038/s41598-025-04036-x>.
- [14] Z. Sun and X. Li, "Named entity recognition model based on feature fusion," *Information*, vol. 14, no. 2, pp. 133–145, 2023, doi: <https://doi.org/10.3390/info14020133>.
- [15] D. Mazitov, I. Alimova, and E. Tutubalina, "Named entity recognition in Russian using multi-task LSTM-CRF," *Journal of Mathematical Sciences*, vol. 273, no. 4, pp. 595–604, 2023, doi: <https://doi.org/10.1007/s10958-023-06521-y>.
- [16] H. Zhang, H. Song, S. Li, M. Zhou, and D. Song, "A survey of controllable text generation using transformer-based pre-trained language models," *ACM Computing Surveys*, vol. 56, no. 3, pp. 1–37, 2023, doi: <https://doi.org/10.1145/3617680>.

- [17] M. Thakkar and N. N. Pise, "Leveraging Transformer-based Pretrained Language model for Task-oriented dialogue system," *International Journal of Computers*, vol. 8, no. 1, pp. 1–4, 2023.
- [18] K. Nassiri and M. Akhloufi, "Transformer models used for text-based question answering systems," *Applied Intelligence*, vol. 53, no. 9, pp. 10602–10635, 2023, doi: <https://doi.org/10.1007/s10489-022-04052-8>.
- [19] A. Rahali and M. A. Akhloufi, "End-to-end transformer-based models in textual-based NLP," *Ai*, vol. 4, no. 1, pp. 54–110, 2023, doi: <https://doi.org/10.3390/ai4010004>.
- [20] S. Lafkier and N. En Nahnahi, "An end-to-end transformer-based model for Arabic question generation," *Multimedia Tools and Applications*, vol. 84, no. 20, pp. 22009–22023, 2025, doi: <https://doi.org/10.1007/s11042-024-19958-3>.
- [21] R. Bhat and R. Nanjundegowda, "A review on comparative analysis of generative adversarial networks' architectures and applications," *Journal of Robotics and Control (JRC)*, vol. 6, no. 1, pp. 53–64, 2025, doi: <https://doi.org/10.18196/jrc.v6i1.24160>.
- [22] C. Shao, Z. Ma, M. Zhang, and Y. Feng, "Beyond mle: Convex learning for text generation," *Advances in Neural Information Processing Systems*, vol. 36, no. 1, pp. 8913–8936, 2023, doi: <https://doi.org/10.52202/075280-0391>.
- [23] J. Xu, X. Liu, J. Yan, D. Cai, H. Li, and J. Li, "Learning to break the loop: Analyzing and mitigating repetitions for neural text generation," *Advances in Neural Information Processing Systems*, vol. 35, no. 11, pp. 3082–3095, 2022, doi: <https://doi.org/10.52202/068431-0223>.
- [24] I. A. M. Huijben, W. Kool, M. B. Paulus, and R. J. G. Van Sloun, "A review of the gumbel-max trick and its extensions for discrete stochasticity in machine learning," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 45, no. 2, pp. 1353–1371, 2022, doi: <https://doi.org/10.1109/TPAMI.2022.3157042>.
- [25] I. Alsmadi, N. Aljaafari, M. Nazzal, S. Alhamed, A. H. Sawalmeh, C. P. Vizcarra, et al., "Adversarial machine learning in text processing: a literature survey," *IEEE Access*, vol. 10, no. 1, pp. 17043–17077, 2022, doi: <https://doi.org/10.1109/ACCESS.2022.3146405>.
- [26] M. Reyad, A. M. Sarhan, and M. Arafa, "A modified Adam algorithm for deep neural network optimization," *Neural Computing and Applications*, vol. 35, no. 23, pp. 17095–17112, 2023, doi: <https://doi.org/10.1007/s00521-023-08568-z>.
- [27] H. Kabiri, Y. Ghanou, H. Khalifi, and G. Casalino, "AMAdam: adaptive modifier of Adam method," *Knowledge and Information Systems*, vol. 66, no. 6, pp. 3427–3458, 2024, doi: <https://doi.org/10.1007/s10115-023-02052-9>.
- [28] C. Zhang, Y. Shao, H. Sun, L. Xing, Q. Zhao, and L. Zhang, "The WuC-Adam algorithm based on joint improvement of Warmup and cosine annealing algorithms," *Math. Biosci. Eng.*, vol. 21, no. 1, pp. 1270–1285, 2024, doi: <https://doi.org/10.3934/mbe.2024054>.