Joint Cultural and Grammatical Correction of Japanese Translations via a GAN-BiLSTM Framework with Policy Gradient Optimization

Shaonan Lin
Foreign Language Institute,Fujian Polytechnic Normal University, Fuqing, Fujian 350300,China
E-mail:linshaonan20250306@126.com

Keywords: Japanese translation correction, generative adversarial network, recurrent neural network, cultural consistency, semantic equivalence

Received: October 4, 2025

Japanese translations often exhibit grammatical deviations and cultural inconsistencies. This study proposes an automatic correction framework that integrates a Generative Adversarial Network with a Bi-LSTM encoder–decoder augmented by embedded cultural vectors to generate correction candidates. A convolutional discriminator jointly evaluates grammaticality and cultural appropriateness under a Wasserstein adversarial objective, while policy gradient optimization refines the generator using discriminator-based rewards. Semantic consistency is preserved through back-translation constraints. Experiments on 60,000 culturally annotated Japanese–Chinese sentence pairs from news, business, and academic domains, split into training, validation, and testing sets, compare the proposed method with RNN, Transformer, and GAN-only baselines. The model achieves 86.2 percent grammatical accuracy, a cultural alignment score of 4.52, a fluency score of 4.10, and a perplexity of 45.6, consistently outperforming baselines across sentence lengths and cultural load levels. The results demonstrate unified optimization of grammatical correction, cultural alignment, and semantic preservation in Japanese translation post-editing.

Povzetek: Študija predstavlja samodejni okvir za popravljanje Japonskih prevodov, ki z združitvijo generativne kontradiktorne mreže in modela Bi-LSTM izboljšuje slovnično pravilnost, kulturno ustreznost in semantično skladnost ter pri tem presega primerjalne pristope.

1 Introduction

In cross-language communication, Japanese translations are required to convey both cultural meaning and informational content, yet grammatical deviations and cultural misalignment remain common and directly affect acceptance in the target language environment [1-2] . When grammatical correctness is prioritized without adequate cultural adaptation, translations may satisfy formal rules while lacking pragmatic naturalness and communicative effectiveness [3-4] . Therefore, automated methods that jointly address grammatical repair and cultural expression optimization are critical for improving cross-cultural communication quality [5-6] .

Prior studies show that Japanese translation errors involve not only grammatical issues such as verb tense, particle usage, and honorific levels, but also improper handling of cultural elements including greetings, politeness strategies, and tonal conventions. Automated translation systems often struggle to preserve semantic fidelity while conforming to target-language communicative norms, a challenge observed in academic, business, and literary contexts [7-8]. Traditional statistical or rule-based correction approaches may still yield rigid expressions or pragmatic violations[9-10]. These limitations indicate that grammar checking or lexical substitution alone is insufficient for advanced cross-cultural text correction[11-12].

Various approaches have been proposed to mitigate these issues. Attention-based sequence-to-sequence models improve contextual modeling and grammatical correction, particularly in word order adjustment [13-14] . Pre-trained language models enhance fluency and expressiveness through large-scale corpus learning [15-16] . Other methods incorporate cultural knowledge bases, embedding honorific rules and greeting templates to improve cultural appropriateness [17-18] . Despite these advances, grammatical correction and cultural consistency are typically treated as separate objectives, limiting the naturalness and cultural fit of generated results [19-20] .

To address the dual requirements of grammatical accuracy and cultural compatibility, this paper proposes a GAN–RNN integrated framework. A Bi-LSTM encoder extracts semantic representations and cultural feature vectors, and a Bi-LSTM decoder with attention generates correction candidates by integrating both sources of information. A convolutional discriminator jointly evaluates grammatical validity and cultural consistency during adversarial training. Optimization combines Wasserstein loss, back-translation consistency constraints, and policy gradient stabilization, while discriminator-based re-ranking is applied at inference to achieve unified grammatical repair and cultural adaptation. This framework introduces cultural modeling into adversarial correction and enables deep integration of cultural information with target-language generation for cross-cultural automatic correction.

To clarify the methodological differences and expose the unresolved gap in joint grammatical and cultural correction, representative approaches are systematically compared in Table 1 along architectural design, optimization focus, evaluation metrics, and cultural modeling strategy.

Table 1: Structured comparison of representative grammatical correction and translation post-editing methods

Method Type	Model Architecture	Primary Optimization Target	Evaluation Metrics	Cultural Modeling	Limitation
RNN-based correction	Bi-LSTM Seq2Seq	Token-level grammatical accuracy	GEC accuracy	Absent	Cultural expressions treated as noise
Transformer-based correction	Self-attention Encoder–Decoder	Global fluency and syntax	GEC accuracy, perplexity	Absent	Pragmatic alignment not constrained
Adversarial text correction	Generator–Discriminator	Distributional naturalness	Fluency score	Implicit	Culture not explicitly represented
Culture-aware post-editing	Seq2Seq with feature injection	Cultural marker preservation	Manual cultural rating	Static features	Grammar and culture optimized separately
Proposed framework	GAN + Bi-LSTM + policy gradient	Joint grammatical–cultural optimization	GEC accuracy, cultural fit, fluency, semantic preservation	Explicit cultural vectors with adversarial reward	Addresses joint optimization gap

The comparison shows that existing state-of-the-art methods lack an integrated mechanism that simultaneously optimizes grammatical correctness and culturally appropriate expression, whereas the proposed framework directly aligns these objectives within a unified adversarial learning process.

From a research design perspective, this study defines its objectives and expectations as follows.

(i) The optimization target is the joint improvement of grammatical correction accuracy, cultural alignment quality, fluency, and semantic preservation, formulated as a unified learning objective rather than isolated evaluation dimensions.

(ii) The bidirectional recurrent encoder–decoder is adopted to ensure stable sequence-level dependency modeling, while adversarial training provides global distributional constraints on grammatical credibility and cultural acceptability. Explicit cultural feature vectors are introduced to encode pragmatic constraints beyond sequence modeling.

(iii) The expected outcome is a correction framework that consistently surpasses recurrent, Transformer-based, and adversarial baselines across grammatical, cultural, fluency, and semantic evaluations, with increased robustness under long sentences and high cultural load.

2. Model design

The overall model design explicitly defines the interactions among modules to support joint optimization of grammatical repair and cultural adaptation, as illustrated in Figure 1. Starting from Japanese translation input, the model proceeds through feature modeling, encoding–decoding, adversarial discrimination, composite optimization, and candidate reranking, forming a complete correction pipeline.

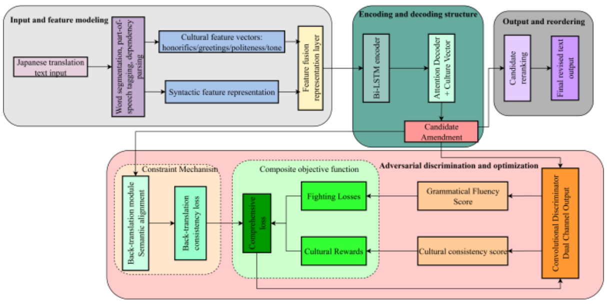

Figure 1: Overall framework of the model

As shown in Figure 1, the input corpus is first processed by word segmentation, part-of-speech tagging, and dependency parsing, after which it is fused with cultural feature vectors and fed into the encoder–decoder to generate preliminary revisions. A discriminator evaluates the generated text through both grammatical and cultural channels, while a back-translation module enforces semantic alignment. These three losses are jointly optimized to constrain both the generator and discriminator. At the output stage, candidate revisions are reranked to produce a final result that is grammatically correct, culturally consistent, and semantically stable, highlighting the central role of cultural feature modeling and composite optimization in the framework.

2.1 Data Preprocessing and Feature Modeling

During model training, the Japanese-translated parallel corpus is processed through word segmentation, part-of-speech tagging, and dependency parsing to obtain explicit grammatical and semantic structures. Word segmentation is performed using a conditional random field–based sequence tagger with large-scale lexical and statistical features, while part-of-speech tagging and dependency parsing are conducted using a neural network–based parser to generate syntactic dependency trees. These trees are incorporated as syntactic graphs in subsequent modeling, providing the encoder with long-range dependency constraints. All annotations follow a unified guideline defining honorifics, greetings, politeness markers, and tonal expressions to ensure labeling consistency. The resulting corpus is encoded using a dual sequence–graph representation, enabling joint semantic–syntactic feature learning [21–22]. Cultural labels are restricted to sentence-level linguistic realizations and exclude speaker identity or social roles, ensuring that they function purely as pragmatic markers.

Cultural feature construction maps honorifics, greetings, modal particles, and politeness levels into vector representations. A cultural dictionary is built via manual annotation of parallel corpora, with relevant segments extracted using context windows. The annotation protocol consists of three components: (1) honorific tagging with three levels (respectful, humble, polite); (2) greeting tagging for opening and closing patterns occurring in more than 1% of the corpus; and (3) politeness and tone tagging based on hierarchical combinations of sentence-ending and modal particles. Annotation is performed by three trained experts in Japanese linguistics, cross-cultural communication, and computational linguistics, achieving a double-label agreement rate above 0.92. Disagreements are resolved by a third annotator. Each cultural label is finally encoded as a multi-hot vector and projected into a 128-dimensional embedding space to ensure discriminability and learnability.To enhance discriminability, multi-hot vectors are introduced to represent different cultural dimensions, which are then mapped into a continuous space using a trainable embedding matrix. Let the cultural feature vector be :

$$C=[c_1,c_2,\ldots,c_k], c_i\in \mathrm{R}^d, \quad (1)$$

Where $k$is the number of cultural feature dimensions and d is the embedding dimension. The cultural embedding is concatenated with the word embedding via a linear projection layer, achieving the fusion of cultural and syntactic features in the same representation space. The cultural embedding vector is computed by multiplying a sentence-level multi-hot cultural label vector with a trainable embedding matrix, producing a continuous representation whose dimensions encode pragmatic formality and interactional tone learned from the annotated corpus.

In the process of jointly modeling syntactic and cultural features, a graph convolution layer is used to encode dependency relationships to capture contextual information under the constraints of syntactic structure. $X\in \mathrm{R}^{n\times d}$Graph convolution operations are constructed on word sequence embeddings and dependency matrices :$A\in \mathrm{R}^{n\times n}$

$$H=\sigma(\widetilde{D}^{\frac{1}{2}}\widetilde{A}\widetilde{D}^{\frac{1}{2}}XW), \quad (2)$$

Where $\widetilde{A}=A+I$ is $\widetilde{D}$ the degree matrix, W is the trainable weight, σ and is the activation function. This process strengthens the expressive power of syntactic dependencies while ensuring word-level contextual consistency. In Equation (2), A denotes the adjacency matrix derived from syntactic dependencies, $\widetilde{D}$ denotes the degree matrix constructed from A, X denotes the input representation of the token sequence, and W denotes the trainable parameter matrix in the graph convolution layer. These symbols follow the standard definition of graph convolution over dependency structures and remain fixed across all training steps. The graph convolutional architecture consists of two layers with 256-dimensional outputs and ReLU activations. Dependency-based graph convolution is adopted to encode syntactic relations by propagating information along dependency edges, enabling stable modeling of long-range grammatical cues in honorifics, greetings, and tonal patterns. Residual connections are introduced between layers to alleviate vanishing gradients, and symmetric normalization is applied to the adjacency matrix to ensure training stability.Tree-LSTM architectures were also explored but were found to restrict alignment between syntactic roles and cultural markers under diverse sentence structures due to their fixed branching patterns. In contrast, dependency-based graph convolution preserves syntactic depth while allowing consistent fusion with cultural embeddings. The final representation is obtained by concatenating the graph convolution output with the cultural vector, yielding a composite grammatical–cultural feature for downstream decoding.

During the cultural annotation phase, we collected statistics on the frequency of cultural features in both the training and validation sets to ensure a balanced distribution of cultural features during model training. Table 2 lists the occurrence ratios of the main cultural categories. These statistics were used to set sampling weights in subsequent training to reduce bias caused by category sparsity.

Table 2: Statistics of the distribution of cultural feature categories in the corpus

Cultural Feature Category	Training Set (%)	Validation Set (%)	Test Set (%)
Honorifics	24.5	24.7	24.6
Greetings	21.8	twenty two	21.9
Politeness Level	27.6	27.4	27.5
Tone Adjustment	26.1	25.9	26
Idioms/Metaphors	18.3	18.5	18.4

in Table 2 shows that the distribution of cultural features across datasets varies minimally, with the proportions remaining relatively stable, indicating that the corpus maintained a high degree of consistency during the partitioning phase. This balanced distribution provides reliable conditions for the model to learn cultural features during training and avoids bias caused by sparse categories.

To prevent cultural features from being diluted in high-dimensional semantic space, a gating mechanism is introduced to adjust the fusion weight of cultural and syntactic features. Let the fused features be expressed as:

$$Z=\alpha H+(1-\alpha)C'(3)$$

in H is the graph convolution output,C' is the cultural vector projection result, $\alpha=\sigma(W_g[H;C'])$ is the gating factor,W_g is a trainable parameter matrix. This mechanism adaptively adjusts the importance of different features during training, ensuring that the cultural dimension maintains a stable influence during the correction process.

2.2 Encoding and decoding structure

After the corpus is preprocessed and feature modeled, the encoder uses a bidirectional long short-term memory network to extract sequence semantics and contextual dependencies. The input sequence is composed of word embedding, syntactic embedding and cultural vector splicing. A subword tokenization pipeline based on byte-pair encoding was applied to all Japanese text prior to embedding projection, with a fixed merge table constructed from the training corpus to ensure stable subword segmentation across honorific forms, greeting expressions, and tonal markers. Let the input be $X=[x_1,x_2,\dots,x_n]$, where each $x_i\in\mathrm{R}^d$ The bidirectional network forms a global representation by superimposing the forward and backward states [23-24], and the calculation formula is:

$$\mathrm{h}_i=\vec{f}(x_i,\mathrm{h}_{i-1})+\overleftarrow{f}(x_i,\mathrm{h}_{i+1})(4)$$

in $\mathrm{h}_i\in\mathrm{R}^m$ Indicates the location i The context fusion representation of$\vec{f}$ and $\overleftarrow{f}$ are the forward and backward LSTM units. The global representation sequence obtained in the encoding phase serves as the S basis for the subsequent decoding input.

In order to ensure that the decoding end fully utilizes cultural and syntactic features during the generation process, a Bi-LSTM decoder with attention mechanism is designed [25-26]. t The hidden state of s_t, the context vector is controlled by the attention distribution:

$$\alpha_{t,i}=\frac{\exp(s_t^T W_a \mathrm{h}_i)}{\sum_j \exp(s_t^T W_a \mathrm{h}_j)},\quad c_t=\sum_i \alpha_{t,i}\,\mathrm{h}_i(5)$$

in $\alpha_{t,i}$ Indicates that at time step t Time source sequence position i The attention weight, W_a is the attention weight matrix,c_t is the context vector. Attention weights are computed by applying a trainable alignment function between the decoder hidden state and each encoder hidden state, followed by softmax normalization, and the context vector is obtained as the weighted sum of encoder states. This mechanism allows the decoder to adaptively select source language information that better aligns with semantic and cultural expressions during the generation process, avoiding reliance solely on a local window. The attention weight matrix W_a uses Xavier initialization with a dimension of 512×512 to match the dimensions of the encoder output and the decoder hidden state, ensuring stable convergence of the attention distribution. OOV items arising from domain-specific lexical forms were processed through the same subword segmentation model, and the resulting subword units were mapped to continuous embeddings through the shared embedding matrix used by both encoder and decoder.

The decoded output combines the context vector and the cultural projection vector at each time step. Let the probability distribution of the generated word be [27]:

$$P(y_t|y_{<t},X)=\mathrm{softmax}(W_o[s_t;c_t;C'])(6)$$

in W_o is the output layer parameter, C' is the projection result of the culture vector,$[s_t;c_t;C']$ Represents the concatenation of the decoded hidden state, the context vector, and cultural features. By directly incorporating cultural features into the generation phase, the candidate revisions are culturally sensitive in terms of word choice and tone. During decoding, the cultural projection vector is concatenated with the decoder hidden state and attention context at each time step and fed into the output layer to compute word probabilities, ensuring continuous influence of cultural features throughout generation rather than only at initialization. To enhance stability under long-range dependencies, residual connections and layer normalization are applied in the decoder. The encoder and decoder share a byte-pair encoding–based vocabulary, ensuring consistent representation of rare lexical variants and preservation of morphological cues relevant to cultural markers in Japanese.

2.3 Adversarial identification and training optimization

After candidate generation, an adversarial discriminator based on a convolutional architecture is introduced for evaluation and optimization. The discriminator takes as input the concatenation of word embedding sequences and cultural features, and applies multiple convolutional layers to extract local patterns. The output is divided into two channels: grammatical fluency and cultural consistency. The grammatical channel performs binary classification to distinguish candidate revisions from human-edited texts, while the cultural channel outputs a continuous regression score measuring contextual fit. Independent validation on a held-out corpus shows that the grammatical channel consistently assigns higher credibility to human-edited texts than to generator outputs. Receiver operating characteristic analysis further confirms a stable separation between the two classes, indicating a clear decision boundary when the discriminator is evaluated in isolation.

The optimization objective of adversarial training adopts an improved loss function based on Wasserstein distance. Assume that the distribution of the generator output is P_g, the manual correction text distribution is P_r, the discriminator function is D, the adversarial loss [28] is expressed as:

$$L_{adv}=E_{x\sim P_r}[D(x)]-E_{\hat{x}\sim P_g}[D(\hat{x})]+\lambda E_{\tilde{x}\sim P_{\tilde{x}}}[(\|\nabla_{\tilde{x}}D(\tilde{x})\|_2-1)^2],(7)$$

in λ is the gradient penalty coefficient,$\tilde{x}$ is a sample obtained by random interpolation between the real sample and the generated sample. This design stabilizes training and mitigates gradient vanishing and mode collapse. The adversarial objective is formulated as the Wasserstein distance between generated and reference distributions, augmented with a gradient penalty enforcing unit-norm constraints on discriminator gradients over interpolated samples. The gradient penalty coefficient λ is set to 10, following prior WRGAN studies and validation experiments. Comparative results with $\lambda \in \{5, 10, 20\}$ indicate that $\lambda = 10$ yields the best generation quality and discriminator stability, and this value is therefore adopted throughout all experiments.To preserve semantic consistency, a back-translation loss [29] is introduced. An off-the-shelf Japanese-to-source neural translation model with fixed parameters is used to retranslate generated text and align it with the original source sequence. The reverse translation employs the same tokenization and embedding pipeline as the main model, ensuring lexical compatibility during alignment.Let the original input be S, the back-translation output is $\hat{S}$, by minimizing the cross-entropy loss to ensure semantic non-drift:

$$L_{bt}=-\sum_{i=1}^{|S|}\log P(\hat{s}_i|\hat{s}_{<i},S),(8)$$

in $\hat{s}_i$ Indicates the first i . This constraint forces the generator to maintain semantic integrity during training, thereby reducing the impact of cultural adjustments on the original meaning. The semantic alignment score is computed from token-level correspondence between the reconstructed and source sequences, using the shared vocabulary of the translation system and main model to ensure stable mapping across training. The back-translation module employs a large pre-trained Transformer, translating generated text back into Japanese. It uses the same word segmentation and embedding as the main model, avoiding inconsistencies from vocabulary or encoding differences. To introduce cultural consistency into the adversarial signal, the discriminator's cultural channel outputs a score $\phi(\hat{x})\in[0,1]$, is used as a reward signal in generator training and the parameters are updated by the policy gradient method. Let the generator parameters be θ, the generated sequence is $\hat{y}$, the reward is $\phi(\hat{y})$, and its optimization goal[30] is:

$$\nabla_\theta J(\theta)=E_{\hat{y}\sim P_\theta}[\phi(\hat{y})\nabla_\theta \log P_\theta(\hat{y})].(9)$$

This formula uses the cultural consistency score to guide the generator in the search space toward outputs that better match the target context. The policy gradient reward signal $\phi(\hat{y})$is derived directly from the continuous score output by the discriminator's cultural channel, ranging from 0 to 1, indicating the cultural consistency of the candidate text. During the generator update process, this $\phi(\hat{y})$ score is treated as the expected reward and is calculated by multiplying the generation probability using the REINFORCE algorithm to calculate the gradient, thereby aligning the generator's optimization direction with the cultural score. The generator parameters are updated by maximizing the expected cultural consistency score produced by the discriminator, with gradients estimated using the REINFORCE algorithm weighted by sequence-level rewards. In the overall training, adversarial loss, back-translation consistency loss and policy gradient reward together constitute a composite objective function:

$$L=L_{adv}+\gamma L_{bt}-\eta J(\theta),(10)$$

in γ and η To adjust the weights, the effects of semantic preservation and cultural rewards are balanced respectively.

In the composite objective, the discriminator's grammatical and cultural channels provide feedback to the generator, while the back-translation constraint introduces a cross-entropy loss for semantic alignment. Combined with the adversarial loss, these form the overall optimization framework, illustrated in Figure 2.

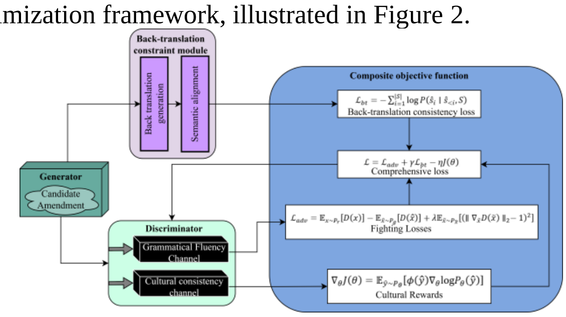

Figure 2: Optimization framework for adversarial training and back-translation constraints

As shown, generator outputs are evaluated by the discriminator on grammatical and cultural dimensions, and the back-translation module enforces semantic alignment. This yields an adversarial loss, a back-translation consistency loss, and a cultural reward signal, jointly applied to both generator and discriminator to optimize grammatical accuracy, semantic fidelity, and cultural compatibility. Inference tests without generator updates confirm that the dual-channel discriminator maintains consistent ranking on unseen samples, demonstrating coherent discrimination behavior beyond the adversarial loop.

2.4 Output reordering

During the generation phase, the decoder outputs multiple candidate revisions, which require a re-ranking mechanism to select the sequence with the best grammatical and cultural fit. The discriminator has already learned a joint representation of grammatical fluency and cultural consistency during training, so its dual-channel output is directly used as the core scoring function during the ranking phase. Let the candidate set be $\{y^{(1)},y^{(2)},\dots,y^{(k)}\}$, the discriminator outputs a grammatical score of $f(y^{(i)})$, and a cultural score of $\phi(y^{(i)})$. The two are weighted together to form a comprehensive score:

$$S(y^{(i)})=\alpha f(y^{(i)})+\beta\phi(y^{(i)})(11)$$

where $\alpha$and $\beta$are adjustment coefficients, controlling the relative importance of grammar and culture. The scoring function performs normalization within the candidate set, ultimately selecting the optimal text by ranking by score.

In order to ensure the stability of the sorting, a length normalization constraint is introduced to eliminate the influence of different sentence lengths on the score.

Assume that the candidate sequence length is $|y^{(i)}|$, and the normalized score is

$$\hat{S}(y^{(i)})=\frac{S(y^{(i)})}{(|y^{(i)}|)^{\gamma}}(12)$$

where γ γ is a balancing factor, controlling the influence of sentence length on the final ranking. This design avoids over-favoring short sentences due to their probability accumulation advantage and prevents long sentences from being penalized by diluting their cultural identity. The length normalization factor γ is set to 0.7. A sensitivity analysis using intervals of 0.5, 0.7, and 1.0 found that a γ of 0.7 achieves the optimal balance between the BLEU score and the manual cultural consistency score, maintaining stability in the ranking of short and long sentences while avoiding the dilution of cultural identity in long sentences. During the reordering process, the internal feature vector of the discriminator is also used as a similarity constraint. $\psi(y^{(i)})$ The cosine similarity between the encoding vector of the candidate sequence and the encoder representation of the original input sentence is calculated and set as: $\psi(y^{*})$

$$R(y^{(i)})=\frac{\psi(y^{(i)})\cdot\psi(y^{*})}{\|\psi(y^{(i)})\|\|\psi(y^{*})\|},(13)$$

This similarity is introduced as an additional term in the comprehensive score to constrain semantic deviation between the corrected output and the original input at the representation level. The encoder representation used for similarity computation is obtained from the same frozen encoder applied to the input sentence during inference, and no external correction references or annotated targets are accessed in this process. The parameters of the discriminator have converged during the training phase and are only used to provide a stable evaluation basis during ranking. The final ranking function after comprehensive optimization is

$$F(y^{(i)})=\hat{S}(y^{(i)})+\lambda R(y^{(i)})(14)$$

where $\lambda$is the semantic constraint weight, which is used to balance the relationship between grammar, culture and semantics. The final score of each candidate is computed as a normalized weighted sum of grammatical score, cultural consistency score, and semantic similarity, and the candidate with the highest score is selected.

3. Experimental Setup

3.1 Dataset Construction

The experimental corpus consists of bilingual parallel data for modeling semantic and syntactic mappings, and manually edited data providing supervision on culture and tone. It includes news, business correspondence, and academic exchanges, spanning everyday and formal domains. All data are preprocessed with word segmentation, part-of-speech tagging, and dependency analysis, and manually annotated for honorifics, greetings, politeness levels, and tone.

Stratified sampling was applied to split the data into training, validation, and test sets, preserving consistent ratios of grammatical and cultural features. The overall split was approximately 8:1:1, with a larger training set to meet generator and discriminator iteration requirements in adversarial training. Table 3 summarizes dataset sizes and cultural feature ratios.

Table 3: Statistics of the size and cultural feature ratio of training, validation and test datasets

Dataset	Number of sentence pairs	Honorifics (%)	Greetings (%)	Politeness level (%)	Tone adjustment (%)	Comprehensive coverage (%)
Training set	48,000	24.5	21.8	27.6	26.1	100
Validation set	6,000	24.7	22.0	27.4	25.9	100
Test set	6,000	24.6	21.9	27.5	26.0	100

Low-frequency cultural categories were balanced via oversampling, allowing the model to learn diverse cultural expressions while maintaining stable discriminator representations. Corpora were checked for consistency, with duplicate or misannotated pairs removed, yielding a structured, standardized dataset for experiments. In a manual consistency check of 1,000 randomly selected sentence pairs, two annotators achieved a Cohen's Kappa of 0.82, confirming high inter-annotator agreement and reliable cultural annotation.

3.2 Experimental parameters

During training, the generator and discriminator are iteratively updated under the same optimization framework. The optimizer uses the Adam architecture, with fixed learning rates and gradient penalties to maintain adversarial training stability. The batch size and iterations are adapted to the data size, while the hidden layer dimensions and number of attention heads remain consistent across the encoder and decoder. The embedding dimensions of the culture vector are combined with grammatical features through a gating mechanism, and reward weights are used to balance culture scoring with semantic preservation. The main parameter configurations are shown in Table 4.

Table 4 : Main parameter configuration table of experimental training

Parameter Type	Value	Parameter Type	Value
Batch Size	64	Learning Rate	1e-4
Epochs	50	Hidden Dimension	512
Attention Heads	8	Cultural Vector Dimension	128
Gradient Penalty	10	Reward Weight	0.3

Under this configuration, the model performed stably in the three-dimensional training of grammatical repair, cultural consistency, and semantic preservation, with a clear parameter convergence trend, providing basic support for the calculation of subsequent evaluation indicators.

3.3 Evaluation metrics

The experiment employs a multi-dimensional quantitative evaluation system. Grammatical effectiveness is measured by alignment accuracy with manual benchmarks, while cultural evaluation combines manual scoring and automatic classifier outputs on honorifics, greetings, politeness levels, and tone. Manual cultural assessment uses a five-point ordinal scale, with three bilingual evaluators independently scoring each sentence; the final score is the arithmetic mean. Evaluators interacted only with system outputs, which were randomly shuffled and anonymized to avoid bias. Fluency is assessed via manual scores and language model perplexity, reflecting formality and naturalness. Semantic preservation is measured by cosine similarity between sentence-level embeddings of the original input and corrected output, derived from the shared encoder representation. This framework ensures complementary evaluation of structural repair, cultural consistency, fluency, and semantic stability. Evaluator consistency, measured by weighted Cohen's Kappa, reached 0.81, confirming reliable and stable human scoring.

4 Results analysis

4.1 Syntax accuracy analysis

Grammatical correction performance was evaluated using the original translation as a baseline and four methods: Method 1, a bidirectional RNN (RNN baseline); Method 2, a Transformer-based correction model; Method 3, a GAN-RNN without cultural features; and Method 4, a complete GAN-RNN jointly optimizing semantic, grammatical, and cultural dimensions. The evaluation considered two aspects: typical error types (verb tense, particle misuse, subject-verb disagreement, modifier errors, and syntactic incompleteness) and sentence lengths (0–10, 11–20, 21–30, 31–40, 40+ tokens), enabling assessment of fine-grained performance and robustness under varying complexity. Figure 3 presents the results.

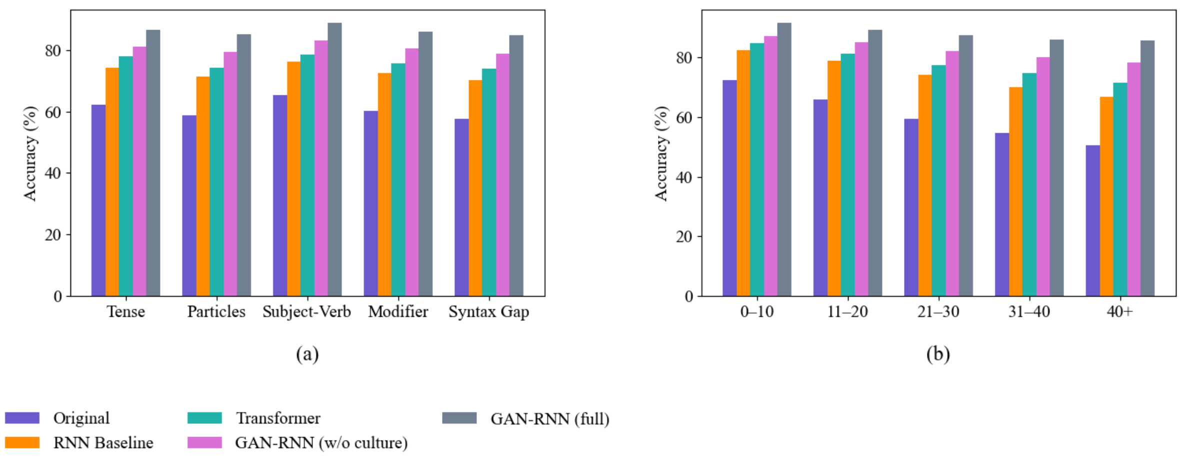

Figure 3: Comparison of the correction accuracy of different models under typical grammatical error types and sentence length ranges
a. Error type dimension, b. Sentence length interval dimension

The original translations achieved 57.6%–65.4% accuracy across error types (average 60.8%), dropping to 50.5% for sentences longer than 40 tokens, indicating numerous uncorrected grammatical defects. The RNN baseline reached 70.1%–76.1% (avg. 72.9%) across error types and 66.8%–82.6% across sentence lengths, showing temporal modeling improves repair but struggles with syntactically incomplete and long sentences. The Transformer baseline achieved 73.8%–78.5% (avg. 76.0%) across error types and 71.5%–84.8% across lengths, demonstrating better long-range dependency handling. GAN-RNN without cultural features corrected 78.9%–83.0% of error types (avg. 80.6%) and 78.2%–87.2% across lengths, showing adversarial training enhances grammatical correction, though performance varied for particle misuse and long sentences. The full GAN-RNN achieved 84.8%–88.7% across error types (avg. 86.2%) and 85.6%–91.5% across sentence lengths, demonstrating robustness in complex contexts. Overall, the model achieves dual optimization for grammatical errors and structural complexity while preserving semantic consistency. Reported accuracy is measured under standard test-time generation without discriminator-based candidate reranking.

To further align the results in Figure 3 with contemporary correction paradigms, Table 5 summarizes the averaged grammatical accuracy of three extended baselines across the same error-type and sentence-length dimensions. This aggregation provides a compact quantitative reference under an identical evaluation setting, facilitating direct comparison with the proposed framework.

Table 5: Average grammatical accuracy across error-type and sentence-length dimensions for extended baselines

Extended Model	Error Type Dimension Avg. Accuracy (%)	Sentence Length Interval Avg. Accuracy (%)	Overall Avg. Accuracy (%)	Significance vs. Proposed Model
Pre-trained Transformer-based GEC	81.4	83.1	82.3	significantly lower
General-purpose Instruction-tuned LLM	83.7	79.8	81.8	significantly lower
Translation-oriented LLM	84.2	85.0	84.6	significantly lower

As reported in Table 5, the pre-trained Transformer-based GEC achieves an average accuracy of 81.4% over error types and 83.1% across sentence length intervals, yielding an overall average of 82.3%, which indicates improved local grammatical modeling but reduced stability on longer structures. The general-purpose instruction-tuned LLM attains 83.7% accuracy on error types but drops to 79.8% across sentence length intervals, resulting in an overall average of 81.8%, reflecting strong correction of frequent grammatical patterns but weaker global structural control in longer sequences. The translation-oriented LLM demonstrates a more balanced performance with 84.2% and 85.0% accuracy on the two dimensions, respectively, and an overall average of 84.6%, yet it remains below the full GAN-RNN model, whose averaged grammatical accuracy exceeds 86% under the same evaluation conditions.

Overall, these results indicate that while contemporary large-scale and pre-trained models provide competitive grammatical correction performance, explicit adversarial optimization over grammatical structure and sequence-level consistency remains advantageous for robust correction across diverse syntactic conditions.

4.2 Cultural consistency analysis

Within the cross-cultural adaptation evaluation framework, this experiment used both subjective manual scoring and rule-based/classifier-based automatic scoring to examine translation differences across five cultural dimensions: use of honorifics (C1) , naturalness of greetings (C2) , politeness (C3) , contextual fit (C4) , and emotional fit (C5) . Model 1 is Baseline-NMT (Neural Machine Translation) . This model, centered on direct sequence-to-sequence mapping, focuses on literal translation and lacks cultural pragmatic modeling. Model 2 is Bi-LSTM. Its bidirectional recurrent structure enhances contextual semantics, improving word order and coherence, but does not explicitly embed cultural embeddings. Model 3 is GAN-only. Adversarial training improves naturalness and fluency, but lacks targeted optimization for cultural markers such as honorifics and greetings. Model 4 is GAN+RNN (the method used in this paper). This model incorporates cultural feature vectors into the recurrent network and combines it with adversarial optimization, aiming to enhance cultural fit while maintaining semantic equivalence. Figure 4 shows the comparison of the above four models under the two types of scoring systems: manual scoring and automatic scoring (on a 100-point scale).

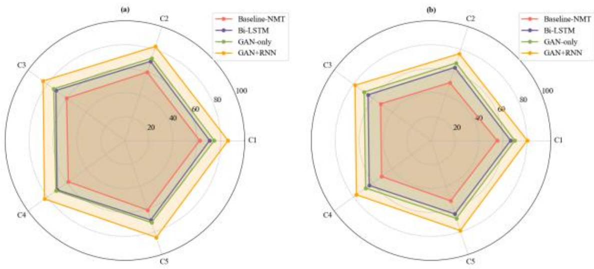

Figure 4 : Multi-dimensional radar chart of cultural consistency of Japanese translations
a: manual scoring , b: automatic scoring

Data analysis shows that the GAN+RNN algorithm scored 86.1 for honorifics and 82.7 for greetings in manual scoring, while the automatic scoring yielded 80.4 for honorifics and 76.2 for greetings. Baseline-NMT scored 62.4 for honorifics and 60.2 for politeness in manual scoring, while the automatic scoring yielded 55.2 for honorifics and 52.1 for politeness. Bi-LSTM and GAN-only algorithms scored 70.8 and 74.6 for honorifics, respectively, in manual scoring, while the corresponding values decreased and the difference between the two approaches widened in automatic scoring. Manual scoring generally outperforms automatic scoring because manual reviewers, guided by a sense of language and overall context, are more tolerant of some politeness or greetings. Automatic scoring relies on rules or classifiers to precisely match honorific markers with sentiment, resulting in more conservative scores. GAN+RNN maintains its lead in both systems because the cultural feature vectors provide the model with explicit pragmatic cues, enabling the generator to optimize fluency and calibrate the formal markers of honorifics and greetings under the pressure of adversarial training. The contextual advantage of Bi-LSTM facilitates semantic coherence, but the lack of adversarial guidance limits the implementation of cultural features. While GAN-only improves naturalness, it fails to bridge the pragmatic gap through explicit cultural embedding. Based on these observations, GAN+RNN, which integrates cultural vectors and adversarial training, demonstrates the most robust effect in improving cultural consistency in Japanese translations.

4.3 Fluency and naturalness analysis

This paper examines translation fluency and naturalness on both subjective and objective scales. The effectiveness of different correction strategies is analyzed alongside human ratings and language model perplexity, aiming to demonstrate the differences in how these methods correct grammar while maintaining cultural fit. Model 1 consists of the original Japanese translation, reflecting the characteristics of unprocessed machine translation output. Model 2 is a system using only a recurrent neural network (RNN-only) for correction, focusing on local correction of sequence dependencies. Model 3 is a system using only a generative adversarial network (GAN-only) for correction, focusing on discriminator-driven naturalness enhancement. Model 4 is the proposed RNN+GAN combined correction method, which builds on Bi- LSTM encoding and attention decoding by incorporating adversarial training and back-translation consistency constraints. Model 5 is a corrector based on a standard sequence-to-sequence Transformer (Transformer-based model), using global self-attention to capture long-range grammatical relationships. The relevant results are shown in Figure 5 below.

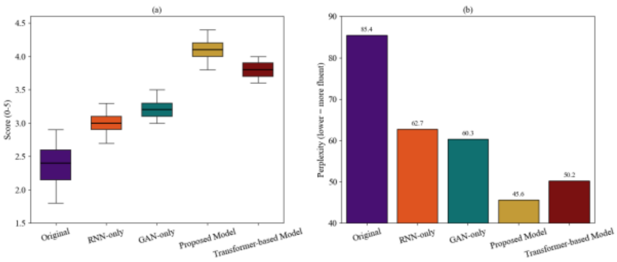

Figure 5: Comparison of translation fluency and confusion
(a) is the distribution of manual ratings, (b) is the comparison of perplexity

Judging from the quantitative results, our method leads in both subjective and objective metrics: the proposed model achieved an average human score of 4.10 and a perplexity of 45.6, significantly outperforming the original text's average score of 2.37 and perplexity of 85.4. The Transformer-based model achieved an average score of 3.81 and a perplexity of 50.2, demonstrating good grammatical grasp but falling short of the naturalness achieved by the joint adversarial training. The GAN-only approach achieved an average score of 3.21 and a perplexity of 60.3, while the RNN-only approach achieved an average score of 3.00 and a perplexity of 62.7. Both approaches contributed to reducing grammatical errors and improving fluency, but to a limited extent. The differences between these approaches stem from their design: the joint model simultaneously optimizes fluency and semantic preservation through sequence modeling and the discriminator's adversarial signal, while back-translation consistency constraints suppress semantic drift, thereby simultaneously reducing language model perplexity and improving human subjective scores. The RNN-only approach lacks a global discriminative constraint, resulting in insufficient processing of subtle naturalness, while the GAN-only approach, without a strong sequence prior, provides weak protection for semantic stability. Conclusion: The collaborative mechanism of RNN and GAN achieves a relatively balanced grammatical repair and cultural expression, making the RNN+GAN joint correction significantly ahead of other control models.

4.4 Semantic preservation analysis

This experiment compares models on semantic preservation, assessing their ability to retain original meaning across sentence lengths and cultural loads. Model 1, a Seq2Seq baseline, uses a standard encoder-decoder without attention or adversarial training. Model 2, GAN-only, optimizes text distribution via a discriminator to enhance naturalness but relies on basic sequence modeling for semantics. Model 3, RNN-only with attention, uses a bidirectional LSTM to capture context and improve semantic continuity. Model 4, the proposed GAN-RNN model, combines RNN context modeling, GAN optimization, and cultural feature vectors to correct grammar and cultural expressions. Evaluation spans five sentence length intervals (1–10, 11–20, 21–30, 31–40, >40 tokens) and five cultural load levels (none to very strong). Figure 6 shows results, with higher similarity values indicating stronger semantic retention; all scores use identical embeddings and normalization for comparability.

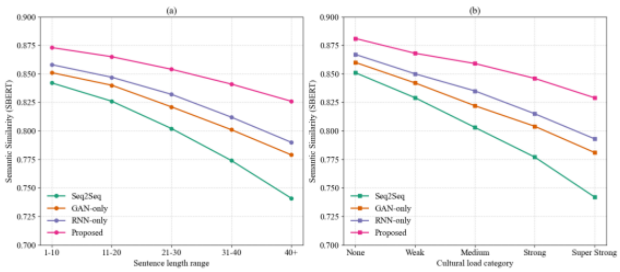

Figure 6: Semantic preservation analysis of the Japanese translation text revision model: comparison of sentence length and cultural load dimensions
(a) Line graph of semantic retention by sentence length , (b) Line graph of semantic retention by cultural load category

The figure shows that the average semantic preservation of the Seq2Seq baseline across all sentence lengths decreases from 0.842 for short sentences to 0.741 for very long sentences, indicating that basic sequence modeling fails to capture the context of long sentences. The GAN-only model slightly outperforms the baseline across all lengths, with a preservation of 0.851 for short sentences and 0.779 for long sentences. This demonstrates that adversarial training improves the naturalness of text but has limited improvement in semantic continuity for long sentences. The RNN-only model maintains stable performance as sentence length increases, with a preservation of 0.858 for short sentences and 0.790 for very long sentences. The attention mechanism effectively mitigates the semantic loss caused by long-distance dependencies. The proposed model consistently leads across all lengths, with a preservation of 0.873 for short sentences and 0.826 for very long sentences. Its fusion mechanism combines context modeling and adversarial optimization, significantly improving semantic preservation for complex sentences and texts with high cultural load. Regarding cultural load, the Seq2Seq baseline drops from 0.851 for texts without cultural features to 0.742 for texts with strong cultural features, demonstrating that pure sequence models lack semantic stability when dealing with rich cultural expressions. The GAN-only model improved to 0.781, demonstrating that naturalness enhancement partially contributes to semantic fidelity. The RNN-only model maintained a retention rate between 0.867 and 0.793 despite varying cultural loads, demonstrating that the attention mechanism helps maintain semantic continuity. The proposed model outperformed other models across all cultural categories, ranging from 0.881 to 0.829, demonstrating its ability to enhance cultural expression while maintaining semantic consistency with the original text, validating the model's effectiveness in cross-cultural semantic correction tasks.

4.5 Ablation experiment analysis

This experiment quantitatively evaluates the contributions of core model components via ablation, analyzing their effects on grammatical correction, cultural fit, fluency, and semantic preservation. Six conditions are tested: full model with cultural vector, discriminator score, and back-translation consistency (A1); without cultural vectors to assess impact on cultural consistency (A2); without discriminator outputs to evaluate its role in grammar and naturalness (A3); without back-translation consistency to examine semantic preservation (A4); combined removal of cultural vectors and discriminator outputs to study interactive effects (A5); and combined removal of cultural vectors and back-translation consistency to observe joint impact on metrics (A6). Figure 7 presents trends of grammatical accuracy, cultural consistency, fluency, and semantic retention across these conditions.

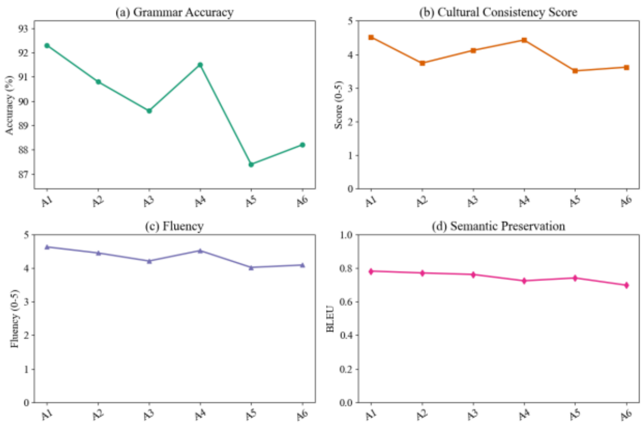

Figure 7: Line chart showing the impact of ablation experiments on the performance of the Japanese translation text correction model.

(a) Changes in grammatical accuracy, (b) Changes in cultural consistency scores, (c) Changes in fluency, and (d) Changes in semantic retention

Figure 7 shows that the full model achieved 92.3% grammatical accuracy, a cultural consistency score of 4.52, fluency of 4.63, and semantic preservation of 0.782, where grammatical accuracy is measured after discriminator-guided candidate selection. After removing the cultural vectors, grammatical accuracy dropped to 90.8%, cultural consistency dropped significantly to 3.74, and fluency and semantic preservation dropped to 4.45 and 0.771, respectively, indicating that cultural characteristics provide key information for the naturalness and cultural fit of the translation. Without the discriminator score, grammatical accuracy dropped to 89.6%, and fluency dropped to 4.21, demonstrating the central role of the discriminator in optimizing syntactic regularity and fluency. Without the back-translation consistency constraint, semantic preservation dropped to 0.724, reflecting the importance of the back-translation consistency signal in maintaining semantic fidelity. Combining the ablation conditions further reduced the levels of each metric. The combination of removing the cultural vectors and removing the discriminator score resulted in a grammatical accuracy of only 87.4%, cultural consistency of 3.51, and semantic preservation of the combination of removing the cultural vectors and removing the back-translation consistency constraint, dropping to 0.698. The results show that the core modules of the model have complementary effects on translation quality, and cultural characteristics, discriminator scores and back-translation consistency jointly support grammatical correction and optimization of cross-cultural expression.

5 Discussion

Experimental results show that the proposed framework consistently outperforms recurrent, Transformer, and adversarial baselines in grammatical accuracy, cultural consistency, fluency, and semantic preservation. Unlike methods focusing solely on sequence modeling or surface fluency, joint optimization of grammar and cultural expression induces systematic improvements. Compared with recurrent and Transformer baselines, the model is more robust on long or complex sentences while maintaining higher cultural alignment, and unlike adversarial-only approaches, it preserves semantic continuity under cultural modification.

These gains stem from cultural vector gating, which modulates the influence of cultural features during encoding and decoding without distorting grammar, and the discriminator, which provides joint feedback on grammatical and cultural quality. Policy-gradient reward propagation reinforces sequence-level correction trajectories that satisfy these constraints.

The framework has limitations. Cultural features rely on domain-specific annotations, limiting transfer to unseen genres or languages without redefining pragmatic categories. Adversarial training requires balanced cultural distributions; sparse or skewed data may reduce stability. Evaluations focus on Japanese translation due to available annotations, but the model is language-agnostic in formulation, governed by abstract cultural representations and sequence-level objectives rather than surface linguistic forms. Cultural representations encode formality and pragmatic function only, not demographic attributes, so corrections reflect normative usage patterns rather than individual identities. Discriminator-driven reward signals favor dominant politeness strategies, emphasizing norm alignment over stylistic diversity. Practical deployment should monitor distributional stability across politeness levels and communicative contexts to avoid narrowing acceptable cultural expressions.

6 Conclusion

This paper presents an automatic Japanese translation correction model integrating GAN and RNN for unified grammatical repair and cultural alignment. Experiments show 86.2% accuracy on typical grammatical errors, outperforming RNN and Transformer baselines. Human ratings on honorifics and greetings reach 86.1 and 82.7, compared with baseline-NMT (62.4 and 60.2), demonstrating that combining cultural feature vectors with adversarial optimization improves cross-cultural naturalness. Fluency analysis shows average human rating of 4.10 and language-model perplexity of 45.6 (vs. 2.37 and 85.4 for original translations), evidencing gains at both syntactic and pragmatic levels. Methodologically, the model employs bi-directional LSTM encoding, attention-based decoding, cultural feature injection, and a convolutional discriminator with back-translation constraints to enhance cultural alignment while preserving semantics. Overall, improvements in grammar, cultural consistency, and fluency highlight the model's practical value for cross-cultural translation correction and scalable automated text repair.

Funding

Innovative Research on the Application-Oriented Undergraduate Foreign Language Translation Talent Training Model Driven by Ideological and Political Education and AI Empowerment" (Provincial General Teaching Reform Project FBJY20250066, University-Level Key Teaching Reform Project 401/XJBJY03).

Research on Japanese translation of Chinese Cultural Documents in the context of Cultural Exchange and mutual Learning" (KY2023FS09).

References

[1] Cao Z. A comparative analysis of Japanese learners' translation bias using neurosemantic analysis[J]. Applied Mathematics and Nonlinear Sciences, 2024, 9(1). https://doi.org/10.2478/amns-2024-0550

[2] Chen Y, Li X. Analysis of linguistic and cultural differences in Japanese translation based on differential equation modeling in the perspective of PMC assessment[J]. Applied Mathematics and Nonlinear Sciences, 2023, 9(1). https://doi.org/10.2478/amns.2023.2.01166

[3] Koizumi R, In'nami Y. Modeling complexity, accuracy, and fluency of Japanese learners of English: A structural equation modeling approach[J]. JALT Journal, 2014, 36(1): 25. https://doi.org/10.37546/jaltjj36.1-2

[4] Povoroznyuk R, Pocheniuk I, Gaidash A, et al. Neuropedagogical guidelines for translation studies: Perceiving the linguistic-cultural markers of the other (foreign) in translation[J]. Revista Romaneasca pentru Educatie Multidimensionala, 2024, 16(4): 185-209. https://doi.org/10.18662/rrem/16.4/912

[5] Cai L. Research on grammatical error correction algorithm in English translation via deep learning[J]. Journal of Intelligent Systems, 2024, 33(1): 20230282. https://doi.org/10.1515/jisys-2023-0282

[6] Bryant C, Yuan Z, Qorib M R, et al. Grammatical error correction: A survey of the state of the art[J]. Computational Linguistics, 2023, 49(3): 643-701. https://doi.org/10.1162/coli_a_00478

[7] Yang M, Li F. Improving machine translation formality with large language models[J]. Computers, Materials & Continua, 2025, 82(2). https://doi.org/10.32604/cmc.2024.058248

[8] Chiba Y, Higashinaka R. Analyzing variations of everyday Japanese conversations based on semantic labels of functional expressions[J]. ACM Transactions on Asian and Low-Resource Language Information Processing, 2023, 22(2): 1-26. https://doi.org/10.1145/3552310

[9] Tarumoto S, Hatagaki K, Miyata R, et al. Evaluating ChatGPT's ability to generate Japanese[J]. Journal of Natural Language Processing, 2024, 31(2): 349-373. https://doi.org/10.5715/jnlp.31.349

[10] Miyamoto T, Nagai N, Satada Y, et al. Risky politeness strategy[J]. Journal of the Artificial Intelligence Society, 2022, 37(3): IDS-G_1-16.

[11] Son J, Kim B. Translation performance from the user's perspective of large language models and neural machine translation systems[J]. Information,

2023, 14(10): 574. https://doi.org/10.3390/info14100574

[12] Boluwatife O S. Cultural nuances in translation: AI vs human translators[J]. 2025.

[13] Zhang Y, Kamigaito H, Okumura M. Bidirectional transformer reranker for grammatical error correction[J]. Journal of Natural Language Processing, 2024, 31(1): 3-46. https://doi.org/10.5715/jnlp.31.3

[14] Al-Sabahi K, Yang K. Supervised copy mechanism for grammatical error correction[J]. IEEE Access, 2023, 11(1): 72374-72383. https://doi.org/10.1109/ACCESS.2023.3294979

[15] Zan C, Ding L, Shen L, et al. Building accurate translation-tailored large language models with language-aware instruction tuning[J]. Frontiers of Information Technology & Electronic Engineering, 2025, 26(8): 1341-1355. https://doi.org/10.1631/FITEE.2400458

[16] Mahsuli M M, Khadivi S, Homayounpour M M. LENM: Improving low-resource neural machine translation using target length modeling[J]. Neural Processing Letters, 2023, 55(7): 9435-9466. https://doi.org/10.1007/s11063-023-11208-1

[17] Lin J C W, Díaz V G Í, Molinera J A M. Introduction to the special issue of recent advances in computational linguistics for Asian languages[J]. ACM Transactions on Asian and Low-Resource Language Information Processing, 2023, 22(3): 1-5. https://doi.org/10.1145/3588316

[18] Katushemererwe F, Caines A, Buttery P. Building natural language processing tools for Runyakitara[J]. Applied Linguistics Review, 2021, 12(4): 585-609. https://doi.org/10.1515/applirev-2020-2004

[19] Al Sharoufi H, Al-Fadhli W S. Bridging the gap: Pragmatic and cultural challenges in machine translation[J]. International Journal of Society, Culture & Language, 2025, 13(2): 29-43.

[20] Ishida T, Murakami Y. Impact of large language models on conversational translation[J]. Authorea Preprints, 2024. https://doi.org/10.36227/techrxiv.173014408.89705418/v1

[21] Ding L, Wang L, Liu S. Recurrent graph encoder for syntax-aware neural machine translation[J]. International Journal of Machine Learning and Cybernetics, 2023, 14(4): 1053-1062. https://doi.org/10.1007/s13042-022-01682-9

[22] Li Z, Parnow K, Zhao H. Incorporating rich syntax information in grammatical error correction[J]. Information Processing & Management, 2022, 59(3): 102891. https://doi.org/10.1016/j.ipm.2022.102891

[23] Huang X. Research on error detection in English translation texts using machine learning algorithms[J]. Intelligent Decision Technologies, 2024, 18(2): 1403-1409. https://doi.org/10.3233/IDT-240111

[24] Xiang Y, Chen Y, Fan W, et al. Enhancing computer-aided translation system with BiLSTM and convolutional neural network using a knowledge graph approach[J]. The Journal of Supercomputing, 2024, 80(5): 5847-5869. https://doi.org/10.1007/s11227-023-05686-2

[25] Chen S, Xiao Y. An intelligent error correction model for English grammar with hybrid attention mechanism and RNN algorithm[J]. Journal of Intelligent Systems, 2024, 33(1): 20230170. https://doi.org/10.1515/jisys-2023-0170

[26] Yirmibeşoğlu Z, Güngör T. Morphologically motivated input variations and data augmentation in Turkish-English neural machine translation[J]. ACM Transactions on Asian and Low-Resource Language Information Processing, 2023, 22(3): 1-31. https://doi.org/10.1145/3571073

[27] Li C, Teney D, Yang L, et al. CulturePark: Boosting cross-cultural understanding in large language models[J]. Advances in Neural Information Processing Systems, 2024, 37: 65183-65216. https://doi.org/10.52202/079017-2082

[28] Jiao Z, Ren F. WRGAN: Improvement of RelGAN with Wasserstein loss for text generation[J]. Electronics, 2021, 10(3): 275. https://doi.org/10.3390/electronics10030275

[29] Pang J, Yang B, Wong D F, et al. Rethinking the exploitation of monolingual data for low-resource neural machine translation[J]. Computational Linguistics, 2024, 50(1): 25-47. https://doi.org/10.1162/coli_a_00496

[30] Hao Y, Liu Y, Mou L. Teacher forcing rewards functions for text generation[J]. Advances in Neural Information Processing Systems, 2022, 35: 12594-12607.