

# Capabilities and Evaluation Biases of Large Language Models in Classical Chinese Poetry Generation: A Case Study on Tang Poetry

Bolei Ma<sup>\*♣,♡</sup> Yina Yao<sup>\*♣</sup> Anna-Carolina Haensch<sup>♣,♡,◇</sup>

♣LMU Munich ♡Munich Center for Machine Learning  
◇University of Maryland, College Park

<sup>\*</sup>Equal contributions.

Correspondence: [bolei.ma@lmu.de](mailto:bolei.ma@lmu.de)

## Abstract

Large Language Models (LLMs) are increasingly applied to creative domains, yet their performance in classical Chinese poetry generation and evaluation remains poorly understood. We propose a three-step evaluation framework that combines computational metrics, LLM-as-a-judge assessment, and human expert validation. Using this framework, we evaluate six state-of-the-art LLMs across multiple dimensions of poetic quality, including themes, emotions, imagery, form, and style, in the context of Tang poetry (唐诗) generation. Our analysis reveals a critical “echo chamber” effect: LLMs systematically overrate machine-generated poems that mimic statistical patterns yet fail strict prosodic rules, diverging significantly from human expert judgments. These findings underscore the limitations of using LLMs as standalone evaluators for culturally complex tasks, highlighting the necessity of hybrid human-model validation frameworks.<sup>1</sup>

## 1 Introduction

Large Language Models (LLMs) have recently shown promising performance in text generation, including applications to creative writing. Among these, classical Chinese poetry poses a unique challenge: it requires adherence to strict prosodic and tonal constraints while achieving semantic richness, emotional resonance, and cultural authenticity (Cai, 2008; He et al., 2021; Yang et al., 2018). Tang poetry (唐诗), in particular, is widely recognized as a pinnacle of Chinese literary tradition, and thus serves as an ideal testbed for evaluating AI creativity in a highly structured cultural domain.

Despite the impressive fluency of modern LLMs, recent studies show that they still struggle with key aspects of poetry generation. Models often fail to maintain coherence across lines, exhibit limited

<sup>1</sup>We release our code and generated poems at <https://github.com/boleima/Tang-Poetry>.

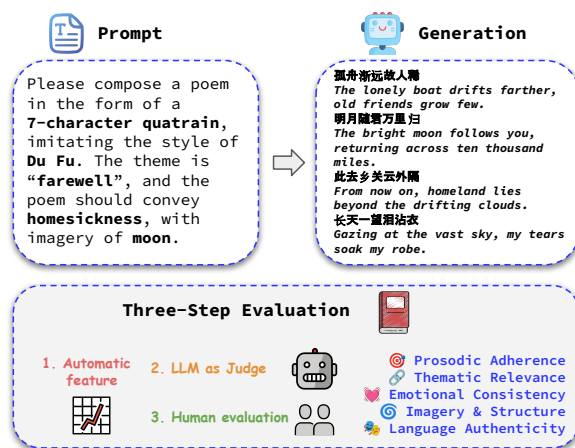


Figure 1: The basic framework of poetry generation and evaluation.

originality in imagery, or collapse into reproducing memorized verses (Chakrabarty et al., 2023; Gómez-Rodríguez and Williams, 2023; Chen et al., 2024). These shortcomings suggest that poetry generation provides not only a challenging benchmark for natural language generation, but also a lens for examining the boundaries of machine creativity.

A further challenge lies in evaluation. Traditional automatic metrics such as BLEU and ROUGE (Papineni et al., 2002) capture surface-level similarity but fail to account for rhythm, imagery, or aesthetic quality (Yang et al., 2018). LLM-based evaluators (“LLM-as-a-judge”) have emerged as a promising alternative, offering richer semantic assessment (Liu et al., 2023; Zhu et al., 2024). However, concerns remain about systematic biases: models may inflate their own outputs, align with peers in an “echo chamber” fashion, or overlook critical dimensions such as prosodic accuracy (Clark et al., 2021). This raises fundamental questions about the reliability of automated evaluation in creative and culturally sensitive domains.

In this work, we present the first systematic study of LLMs in generating and evaluating classical

Chinese poetry, focusing on Tang poetry as a representative form. The basic research framework is shown in Figure 1. We let the models generate Tang poems giving instructions of different dimensions including form, poet, theme, emotion and imagery. We then introduce a three-step evaluation framework that integrates (1) computational feature extraction, (2) LLM-as-a-judge evaluation, and (3) human expert validation. Applying this framework to six state-of-the-art LLMs, we uncover distinct performance tiers in generation quality and reveal systematic evaluation biases that deviate from human judgments.

Our contributions are threefold:

- We propose a multi-dimensional three-step framework for evaluating LLM-generated Chinese poetry.
- We conduct the first empirical study of self- and cross-evaluation biases among LLMs in the classical Chinese poetry domain, identifying “echo chamber” effects in automated assessment.
- We provide comparative benchmarks across six LLMs, highlighting strengths, weaknesses, and the necessity of hybrid collaboration of human and AI validation in evaluating creative AI.

## 2 Related Work

**Chinese Classical Poetry as an NLP Generation Task.** Tang poetry is characterized by strict metrical rules and tonal patterns, making it a demanding test case for text generation. Beyond formal correctness, effective poetry must convey imagery, emotion, and cultural resonance (Cai, 2008; Yang et al., 2018; Ma et al., 2023; Huang and Shen, 2025; Zou, 2025). These qualities distinguish poetry generation from standard text generation tasks and require evaluation frameworks that capture both structural and aesthetic dimensions.

**LLMs for Poetry Generation.** These are existing works on LLMs for generating poetry in Languages such as French (Hämäläinen et al., 2022), Greek (Chatzikyriakidis and Natsina, 2025), and Russian (Koziev and Fenogenova, 2025). While early work on Chinese poetry generation explored template- or retrieval-based systems (He et al., 2021), recent advances in LLMs have greatly improved fluency and stylistic control (Yu et al., 2024; Liu et al., 2025). There are also LLMs trained for the poetry tasks (Yao et al., 2025). Nonetheless, studies have documented persistent challenges:

models often produce repetitive or collage-like outputs, lack originality in imagery, and fail to consistently obey prosodic rules (Chakrabarty et al., 2023; Chen et al., 2024). These limitations motivate systematic evaluation of current models in highly constrained creative tasks such as Tang poetry.

**Evaluation of Generated Poetry.** Automatic evaluation remains a bottleneck. Metrics like BLEU and ROUGE (Papineni et al., 2002) are ill-suited to capturing the cultural and aesthetic value of poetry (Yang et al., 2018). LLM-as-a-judge approaches provide richer assessments (Liu et al., 2023; Zhu et al., 2024), and multi-dimensional evaluation frameworks have been proposed for creative generation (Zhipeng et al., 2019; Deng et al., 2024). Yet, recent work highlights systematic biases in LLM evaluations, including inflated self-assessments and convergence toward flawed standards (Clark et al., 2021). This motivates integrating automated metrics with human expert validation to ensure reliability in evaluating poetry.

## 3 Experimental Design and Methods

This section integrates the experimental design and analytical methods for generating and evaluating the Tang poems. As shown in Figure 1, the overall methodology involves large-scale generation, and a three-step evaluation framework (computational features, LLM cross-evaluation, and human expert validation).

### 3.1 Generation

**Models.** We select six open-source instruction-tuned LLMs as generators and evaluators: DeepSeek-V2-Lite-Chat (DeepSeek-AI et al., 2025), Qwen2.5-7B-Instruct (Yang et al., 2025b), GLM-4-9B-Chat-HF (GLM et al., 2024), Mistral-7B-Instruct-v0.3 (Jiang et al., 2023), Baichuan2-7B-Chat (Yang et al., 2025a), Gemma-2-9B-it (Team et al., 2024).

**Poetry Dimensions.** We cover five important dimensions for Tang poetry: poetic form, target poet style, theme, emotion, and imagery (Table 1). For each dimension, we list representative example elements. For each model, we generate a controlled cross sample set (15,000 poems in total, around 2,500 per model), covering all five dimensions and their elements.

Dim.	Specific Elements
Form	5/7-character quatrains/regulated verse
Poet	Li Bai, Du Fu, Bai Juyi, Wang Wei, Li Shangyin
Theme	landscape, homesickness, nostalgic, pastoral, parting
Emotion	sadness, tranquility, boldness, romance, joy
Imagery	wind, flowers, willows, moon, geese

Table 1: Poetry dimensions and elements for generation. The original Chinese version is in Appendix A.

**Prompt.** Generation is guided by explicit prompts specifying form and imitation target, with temperature set to  $T = 0.4$  for diversity. An example prompt is shown in Appendix B.

### 3.2 Evaluation

We propose a three-step evaluation framework (Figure 1) designed to assess both the generative quality and evaluation biases of LLMs in Tang poetry. The framework integrates (i) automated feature analysis, (ii) LLM-as-judge cross-assessment, and (iii) human expert validation.

#### 3.2.1 Evaluation Pipeline

Evaluation proceeds in three parts. First, automated feature extraction provides objective measures of lexical and semantic properties of generated poems.

Second, in the LLM-as-judge stage, a  $6 \times 6$  evaluation matrix is constructed, where each model acts both as generator and evaluator, scoring all generated samples. This setup enables simultaneous observation of self-assessment and peer-assessment behaviors. Each evaluation is conducted along five dimensions:

- **Prosodic Adherence:** Assesses the harmony of sounds, correctness of tonal patterns, use of rhyme, and overall rhythmic quality.
- **Thematic Relevance:** Judges the depth, clarity, and accuracy with which the poem expressed its designated theme.
- **Emotional Consistency:** Evaluates the authenticity and resonance of the emotion conveyed, ensuring it aligned with the specified emotional tone.
- **Imagery and Structure:** Assesses the vividness and originality of the poem’s imagery, as well as its logical coherence and structural integrity.
- **Language Authenticity:** Judges the precision, conciseness, and authenticity of the language used.

These five dimensions were adopted by synthesizing the core structural and aesthetic criteria defined

in classical Chinese literary theory (Liu, 1962; Cai, 2008). Scores range from 1 to 5.

Finally, in the human validation stage (Step 3), the top three models (ranked by peer evaluation, excluding self-scores) are selected. From each, 50 poems (150 total) are sampled and anonymized. Experts in classical Chinese literature evaluate these poems using the same five dimensions, providing an external calibration of automated evaluation.

#### 3.2.2 Evaluation Methods (three-step analysis)

We detail the analytical methods used in each step.

**Step 1: Computational feature extraction.** We employ three categories of objective metrics to characterize lexical and semantic properties of generated poems.<sup>2</sup>

- 1) **TF-IDF.** We compute TF-IDF to extract lexical fingerprints of poems and models, quantifying characteristic vocabulary usage. For poem  $d \in D$  and term  $t$ :

$$\text{TF-IDF}(t, d, D) = \text{TF}(t, d) \cdot \log \frac{|D|}{|\{d' \in D : t \in d'\}|} \quad (1)$$

enabling comparison of lexical tendencies across models (Sparck Jones, 1988; Manning and Schütze, 1999).

- 2) **Shannon entropy.** To assess lexical diversity, we compute Shannon entropy:

$$H(X) = - \sum_i p(x_i) \log_2 p(x_i) \quad (2)$$

Distributions are compared via one-way ANOVA, with Tukey HSD post-hoc tests when significant, distinguishing between “formulaic” and “diverse” generative styles.

- 3) **Semantic similarity.** To measure thematic relevance and internal coherence, we embed sentences/poems using Chinese sentence embeddings and compute cosine similarity:

$$\text{sim}(A, B) = \frac{A \cdot B}{\|A\| \|B\|} \quad (3)$$

This is applied to poem-prompt alignment and average line-to-line coherence (Reimers and Gurevych, 2019).

<sup>2</sup>These metrics capture diversity and stylistic tendencies but do not directly measure literary quality; they must be interpreted in conjunction with step 2/3.

**Step 2: Descriptive analysis of LLM-as-judge scores.** For the  $6 \times 6$  matrix of cross-assessments, we conduct descriptive and inferential analyses:

- **Generator effect.** Column averages yield consensus rankings of generators.
- **Evaluator effect.** Row averages capture evaluator temperament (leniency vs. harshness).

This analysis provides both performance ranking and diagnostic insight into evaluation biases.

**Step 3: Human-AI alignment.** To measure agreement between automated and human evaluation, we compute Spearman rank correlation between expert scores and mean LLM scores:

$$\rho = 1 - \frac{6 \sum d_i^2}{n(n^2 - 1)} \quad (4)$$

where  $d_i$  is the rank difference for poem  $i$ , and  $n$  is the sample size. High  $\rho$  indicates stronger alignment, while low values reveal systematic blind spots (e.g., insensitivity to metrical violations).

**Sample and statistical strategy.** The human evaluation set comprises 150 anonymized poems (50 per model from the top three peer-ranked models). Each poem is rated by a single expert.

## 4 Experiments and Results

This chapter presents the analysis results of the comparative study, structured according to the three-step evaluation framework.

### 4.1 Computational Feature Analysis

The first analysis employs quantitative methods to objectively characterize the textual properties of the generated poems. Three metrics are considered: lexical fingerprinting (TF-IDF), lexical diversity (Shannon entropy), and semantic coherence (embedding-based similarity).

#### 4.1.1 Lexical Fingerprinting

To identify the characteristic vocabulary of each generative model, we applied the TF-IDF weighting defined in Eq. 1. By calculating scores for all terms across the corpus, a distinct “lexical fingerprint” was derived for each of the six models.

**Shared classical lexicon captured the foundational thematic elements of the genre.** A comparative analysis of the top 10 keywords with the highest TF-IDF scores for each model (see Figure 2) shows that all models successfully reproduce the core lexicon of classical Chinese poetry.

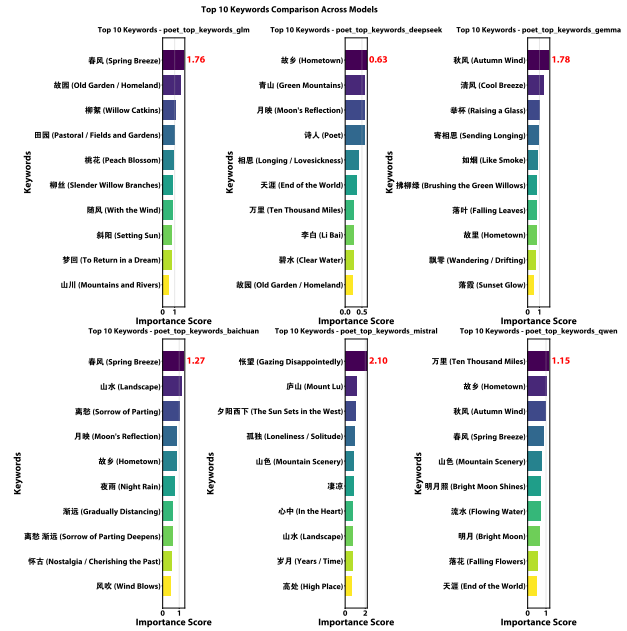


Figure 2: Top 10 keywords for all models, with highest TF-IDF scores marked red for each model.

Traditional poetic imagery such as 春风 (“spring breeze”), 明月 (“bright moon”), 青山 (“green mountains”), and 故乡 (“hometown”) appears frequently across models. This common vocabulary suggests that all generators captured the foundational thematic elements of the genre.

**Distinct lexical preferences reveal model-specific styles.** Closer inspection, however, reveals clear stylistic divergences. Qwen and DeepSeek, for example, show a marked preference for expansive landscape expressions, with terms such as 万里 (“ten thousand miles”) and 天下 (“all under heaven”) scoring prominently. In contrast, Baichuan, Gemma, and Mistral lean toward more personal and emotional expressions, frequently using words such as 离愁 (“sorrow of parting”), 萧瑟 (“desolate”), and 孤独 (“loneliness”). By comparison, GLM favors more joyful and serene vocabulary, such as 春风 (“spring breeze”) and 悠然 (“leisurely”). These differences suggest that models have internalized stylistic patterns from their training corpora.

#### 4.1.2 Lexical Diversity and Predictability

Lexical diversity and unpredictability were quantified using Shannon entropy ( $H$ ), defined in Eq. 2. Entropy measures the average uncertainty of the word distribution in a text: higher values indicate richer lexical variety and lower predictability.

The analysis reveals **three distinct performance tiers among the six models** (see Figure 3). **Class 1**

**(High entropy): GLM and Mistral.** Both produce poems with the highest lexical diversity. **Class 2 (Medium entropy): DeepSeek, Qwen, and Gemma.** These models form a central cluster. **Class 3 (Low-entropy outlier): Baichuan,** which consistently yields the lowest lexical diversity. To formally validate these observations, a one-way ANOVA was conducted, yielding a highly significant result ( $F(5, 14994) = 1986.6, p < .001$ ). Post-hoc comparisons with Tukey’s HSD (see Table 2) confirm the tiered structure. Results confirm three entropy-based classes: Baichuan as an outlier, a mid-tier cluster (DeepSeek, Qwen, Gemma), and a high-entropy group (Mistral, GLM).

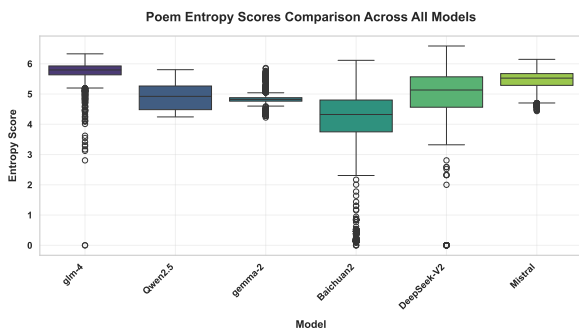


Figure 3: Entropy scores across models.

Model Group 1	Model Group 2	Mean Diff.	p-adj	Reject $H_0$	Conclusion
Baichuan	Other Models	> 1.03	<.001	True	Stat. Unique (Class 3)
DeepSeek	Qwen	0.0197	0.936	False	Stat. Indistinguishable
DeepSeek	Gemma	-0.1865	<.001	True	Sig. Different
Mistral	GLM	0.2652	<.001	True	Sig. Different
Mid-Tier Models	High-Tier Models	> 0.43	<.001	True	Sig. Different

Table 2: Pairwise comparison of mean entropy scores using Tukey’s HSD.

**The strict prosodic rules of regulated poetry naturally constrain lexical choices, reducing entropy.** Hence, Baichuan’s low entropy may reflect either (a) genuine lexical limitation and repetitiveness, or (b) stronger mastery of poetic form. Conversely, the high entropy of GLM and Mistral could signal creative flexibility or, alternatively, a tendency to sacrifice formal correctness for lexical variety. Further analysis of entropy by theme (Figure 4) provides additional nuance. Baichuan’s entropy decreases sharply when tasked with thematically specific prompts (e.g., 怀古 “Nostalgic”), indicating vocabulary restriction and repeated reliance on signature terms such as 万里 (“ten thousand miles”). Combined with TF-IDF results, this suggests a **template-driven generation strategy, in which iconic words are repeatedly reused rather than flexibly adapted to context.**

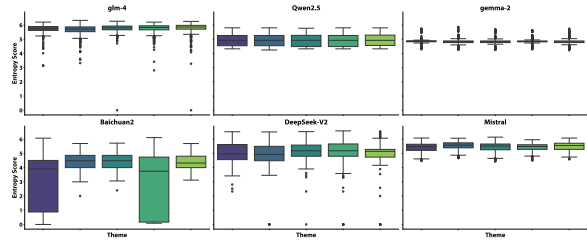


Figure 4: Entropy scores across models by theme.

### 4.1.3 Semantic and Cultural Coherence

To address the ambiguities left by the entropy analysis and to move beyond purely lexical metrics, we conducted a semantic analysis using the [shibing624/text2vec-base-chinese](#) sentence embedding model.

**Semantic Theme Distinction: Lexical diversity does not guarantee semantic or stylistic richness.** This analysis evaluates how well models generate semantically distinct outputs across creative prompts. Lower similarity scores indicate stronger separation. The results (Figure 5) show that models do not separate themes well. Baichuan performs better with lowest similarity score (0.821). This aligns with its low entropy: concentrated thematic focus can naturally yield more predictable vocabulary, suggesting that its low entropy reflects thematic control rather than lexical weakness. Even high-entropy models like GLM and Mistral fail to convert lexical variety into semantic nuance. This reinforces the entropy Paradox: lexical diversity does not guarantee semantic or stylistic richness.

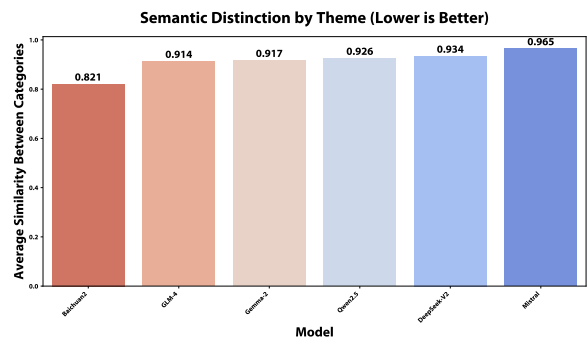


Figure 5: Semantic distinction by theme.

**Generated poems reveal cultural blind spots.** This analysis examines the models’ grasp of classical Chinese “imagery-emotion” pairings (Figure 6). They reliably reproduced canonical links such as Moon -> Tranquility and Wild Goose -> Sadness, with Gemma performing especially well on the latter. Yet every model failed to capture the

quintessential Willow -> Sadness association, a cornerstone of Chinese poetic symbolism where willow branches signify parting and melancholy.<sup>3</sup> This omission highlights a critical deficit in cultural association fidelity that undermines the authenticity of the generated poems. The gap may stem from sparse representation of this motif in broad training data or from deeper architectural limits in modeling nuanced metaphorical relationships.

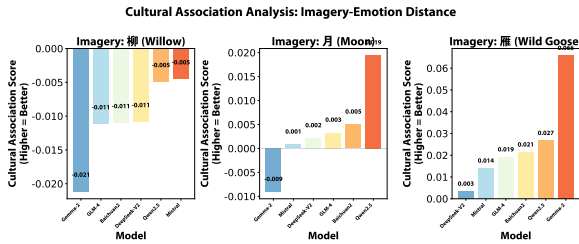


Figure 6: Imagery-Emotion Association in cultural association scores.

**Models exhibit strong semantic convergence.** We measured cross-model semantic similarity (Figures 7a and 7b) and found mean scores above 0.94, revealing a shared representation of poetic meaning likely rooted in overlapping training data. Despite noticeable lexical variation, the underlying semantics are strikingly uniform. We clarify that the alignment paradox refers to the contradictory relationship between a model’s internal semantic knowledge and its actual generative output. We further show the average similarity by poets in Figure 8 (the lower, the more distinct). Similar to the findings from Figure 5, models do not separate poems generated for different poets well. This suggests that while prompting different models with different poets (e.g., 李白 “Li Bai” vs. 杜甫 “Du Fu”), they struggle to manifest these differences in the practice when generating poems. They are essentially trapped in a shared semantic space where they can describe stylistic differences in theory but produce outputs that are semantically indistinguishable from one another.

<sup>3</sup>The willow (柳, liǔ) is a canonical cultural allusion in Chinese poetry for parting and melancholy (due to the homophony with 留, liú, meaning to stay) (Peng et al., 2024). For the target genre and theme, this is a cornerstone of poetic symbolism. In Chinese poetry, there is canonical links between the imagery and emotions, such as Moon -> Tranquility, Wild Goose -> Sadness, Willow -> Sadness, as detailed in Classical Chinese studies such as Cai (2008).

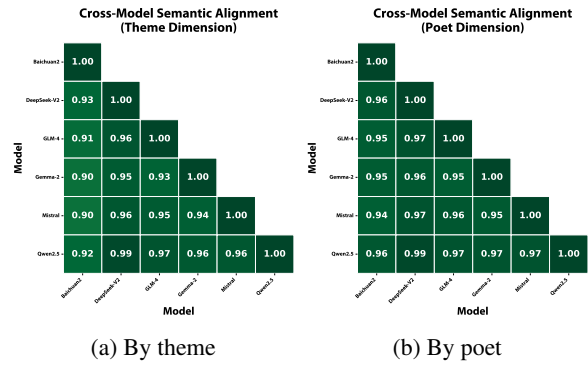


Figure 7: Cross-model semantic alignment.

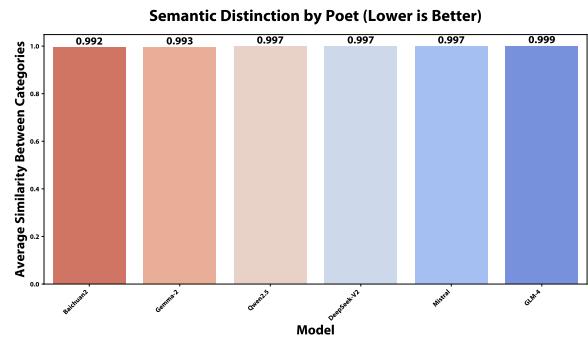


Figure 8: Semantic distinction by poets.

## 4.2 LLMs-as-judges Evaluation

Building upon the objective computational foundations established in analysis 1, the second analysis of our evaluation framework shifts to subjective poetry quality assessment through the LLMs themselves. This LLM-as-a-judge evaluation analysis employs the comprehensive 6 × 6 evaluation matrix, where each model serves as both generator and evaluator, creating a total of 36 combinations. This approach treats the collective judgment of all six models as a proxy for objective quality assessment, similar to how academic peer review aggregates expert opinions. To establish a baseline understanding of generation quality based on “peer consensus”, we analyze the average scores across the five evaluation dimensions received by each generator from each evaluator (see columns in Figure 9).

**Generation Capabilities: Six models vary in generation capabilities, while the overall evaluation scores are high.** The LLM-as-a-judge evaluation heatmap shows that LLMs like Qwen, Gemma, GLM, DeepSeek as generators consistently receive higher scores across, as indicated by the predominantly red coloring, while Baichuan and Mistral receive the lower scores across evaluators, shown by the lighter, less red coloring in its row. However,

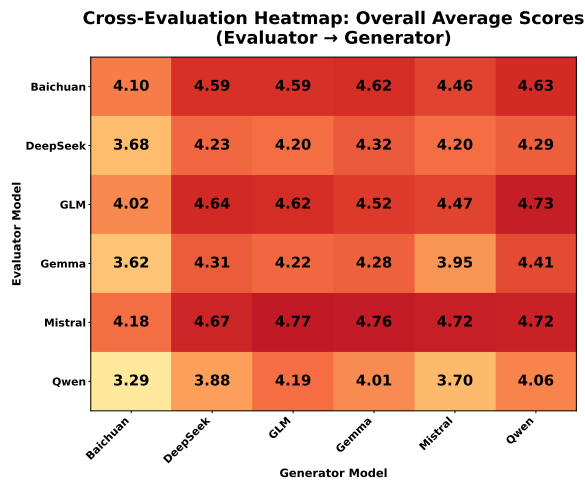


Figure 9: LLM-as-a-judge cross evaluation results. The x axes shows the generation models, while the y axes shows the evaluation models.

we still notice that the overall evaluation scores of all models are above 3, showing the high LLM-generated consensus for the overall quality assessment across models.

**Evaluator Characterization: The evaluators have their scoring tendencies and biases.** We then look into each evaluator model’s evaluation results on different generator models (based on the horizontal view on the y axes on Figure 9). We notice that the models Mistral, GLM, Baichuan evaluate with generally higher scores across all the generator models (showing darker colors on heatmap), while DeepSeek, Gemma, Qwen give lower scores with lighter colors on the heatmap. We further show in Figure 10 the evaluation results of the models for self-evaluation (i.e., the average results of the same model evaluating its generated poems) vs. other-LLM-as-judges evaluation (i.e., the average results of other four LLMs-as-judges evaluating the generated poems). We also show the biases of self – others results. We notice that Baichuan, GLM, Mistral tend to rate self higher than other models, while DeepSeek, Gemma, Qwen tend to rate self lower than other models. Based on these findings, we notice that the models have their own scoring tendencies.

**Interaction and Evidence for the “Echo Chamber”.** Figure 11 plots the mean scores with generator models on the x-axis and evaluator models as separate traces. The pronounced non-parallel lines indicate a strong generator-evaluator interaction, the visual signature of a **methodological echo chamber**. In news and social media, an echo cham-

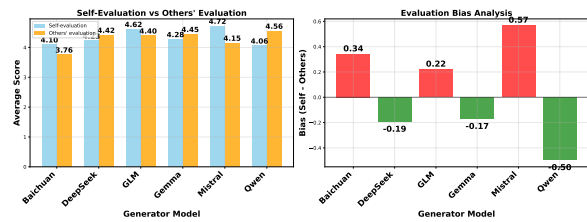


Figure 10: Bias between self and other-LLM-as-a-judge Evaluation. Left: Average scores of self vs. others’ evaluation; Right: Bias (self – others) results.

ber describes an environment in which participants repeatedly encounter and reinforce the same beliefs, insulated from external challenge (Cinelli et al., 2021). We use this term to describe a closed evaluative loop among LLMs in which models, trained on overlapping data and architectures, reinforce one another’s patterns rather than reflecting true poetic quality. This pattern shows that each evaluator adjusts its scoring behavior depending on the generator under review, suggesting collective reinforcement of shared biases rather than independent, objective judgment.

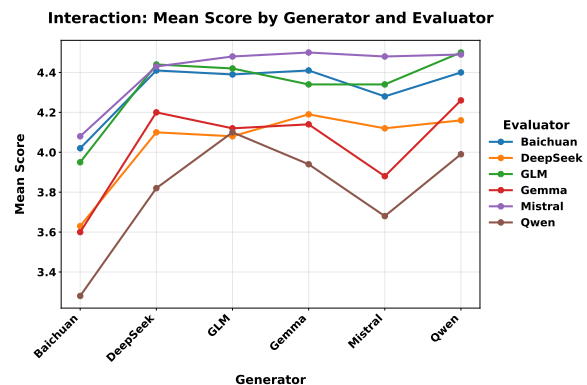


Figure 11: Interaction plot of mean scores by Generator (x-axis) and Evaluator (traces).

**Dimensional Analysis: Bias Patterns Across Evaluation Criteria.** We show the detailed LLM evaluation for each dimension in Figure 12. It reveals that bias patterns are not uniform across evaluation criteria. The *Prosodic Adherence* dimension shows the most consistent scoring patterns across evaluators, suggesting models have developed reliable criteria for formal poetry assessment. Bias variations are minimal, indicating this represents the most “objective” evaluation dimension. The content-focused *Thematic Relevance* and *Emotional Consistency* dimensions exhibit moderate bias variations, with GLM and Mistral showing particular confidence in their own thematic and

emotional expression capabilities. The *Imagery & Structure* and *Classical Language Authenticity* represent the most subjective dimensions, where cultural knowledge and aesthetic judgment are paramount. Here, bias patterns are most pronounced, with Qwen showing notably harsh self-assessment in Classical Language despite peer recognition of its linguistic sophistication.

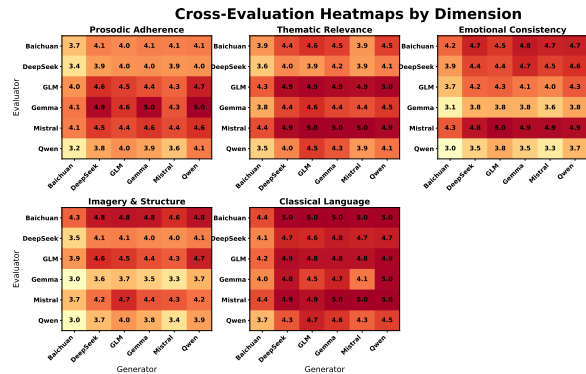


Figure 12: LLM-as-a-judge evaluation scores for each dimension.

### 4.3 Human Evaluation

The final analysis introduces human expert evaluation as the ultimate *gold standard* for poetic quality. While the first two tiers revealed the internal logic and biases of LLM-based assessment, only expert judgment provides a definitive reference for both artistic and technical merit. A stratified random sample of 150 poems: 50 each from Qwen, DeepSeek, and Gemma, the three highest-ranked generators in Tier 2, was scored by 2 professional scholars of classical Chinese poetry using the same five-dimensional rubric applied in earlier tiers.

**Human ratings diverge sharply from machine consensus.** Qwen was judged a clear winner, achieving a mean score of 4.53. Gemma (mean 3.76) and DeepSeek (3.72) lagged far behind, despite their near-parity with Qwen in the automated evaluation. Dimensional analysis explains the discrepancy. Qwen received uniformly high marks, especially in *Imagery & Structure* (4.96) and *Prosodic Adherence* (4.46). Gemma, although strong in *Thematic Relevance* (4.56), failed catastrophically on *Prosodic Adherence* (mean 1.72) due to frequent rule violations and occasional English intrusions in Chinese verses, while this has not been noticed by LLM evaluations. DeepSeek showed moderate competence across dimensions but lacked both technical precision and artistic flair. These

findings highlight limitations of machine consensus, which missed Gemma’s critical prosodic flaws. See the detailed results and visualization of human evaluation scores in Appendix M.

**Statistical alignment of human and LLM scores.** We computed Spearman correlations between human and machine scores across all 150 poems (Table 3). At the aggregate level, modest but significant positive correlations emerge for *Thematic Relevance* ( $\rho = 0.309, p < 0.001$ ) and *Language Authenticity* ( $\rho = 0.297, p < 0.001$ ), indicating a shared but limited sense of quality.

Generator-level analysis reveals sharper divergences. For Qwen, the top human-rated model, no dimension showed significant correlation, demonstrating that machine evaluators failed to capture qualities valued by experts. Gemma provides clear evidence of systematic misalignment: a strong positive correlation on *Thematic Relevance* ( $\rho = 0.622, p < 0.001$ ) but a negative, non-significant correlation on *Prosodic Adherence* ( $\rho = -0.176$ ). This shows that LLM evaluators collectively overlooked Gemma’s prosodic violations.

Generator	Prosody	Thematic	Emotional	Imagery	Language
Overall (N=150)	0.196*	0.309***	0.226**	0.095	0.297***
Qwen (n=50)	-0.033	0.225	0.017	-0.161	0.052
DeepSeek (n=50)	0.097	0.152	0.354**	0.038	0.225
Gemma (n=50)	-0.176	0.622***	0.304*	0.113	0.604***

Significance: \*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$

Table 3: Spearman correlations ( $\rho$ ) between human and LLM judges scores. Detailed visualization of the correlations is shown in Appendix M.

### Validating the “Echo Chamber” hypothesis.

We empirically validate this hypothesis through Gemma’s specific failure case. While human judges identified significant prosodic defects, the LLM evaluator cohort failed to penalize them, demonstrating a collective reinforcement of errors. This phenomenon is rigorously quantified by the Spearman Rank Correlation ( $\rho$ ) in Table 3. A striking example is the negative correlation for Gemma’s Prosodic Adherence ( $\rho = -0.176$ ), which stands in sharp contrast to the high alignment observed in semantic tasks like Thematic Relevance ( $\rho = 0.622, p < 0.001$ ). This disparity suggests that while LLMs align on semantic content, they share a critical blind spot regarding structural constraints. We attribute this to overlapping limitations in pretraining data, i.e., the presence of cross-lingual artifacts and a lack of rigorous, rule-based prosodic representation across the models.

## 5 Discussion

Our findings show that LLMs can produce surface-level fluency in classical Chinese poetry while exhibiting fundamental weaknesses that automated metrics fail to detect. Using a three-step evaluation framework including computational analysis, cross-model scoring, and expert review, we reveal model capabilities in literacy, evaluation biases, and cultural competence, with following insights:

**The most striking result is an echo-chamber effect: models trained on similar data converge on shared yet inaccurate standards for poetic quality.** Gemma’s severe prosodic failures received ratings nearly equivalent to Qwen’s technically correct output when judged by other LLMs, showing that cross-model consensus does not guarantee accuracy. This undermines the common assumption (e.g., Liu et al., 2023; Wang et al., 2024; Chan et al., 2024; Hu et al., 2026) that ensemble evaluation reduces bias; semantic alignment ( $>0.94$ ) suggests that common training data instead amplifies collective blind spots. In creative tasks, human oversight therefore remains indispensable.

**Single metrics such as entropy cannot capture the multifaceted nature of poetic quality.** Entropy and other one-dimensional scores provide useful diagnostics but cannot serve as stand-alone measures. The “entropy paradox”, where low entropy reflects either strong formal control or restricted vocabulary, shows how isolated metrics obscure key distinctions. Combining entropy with single “fingerprints” and semantic-coherence measures proved more informative, revealing, for example, Baichuan’s template-driven generation. Future work should adopt multi-metric diagnostic frameworks (Ghazvininejad et al., 2017; Van de Cruys, 2020) rather than simple performance rankings.

**High semantic similarity masks a lack of cultural and stylistic depth.** Despite striking inter-model agreement, systems failed to capture core cultural associations or to differentiate canonical poet styles. The inability to represent the classic “willow-sadness” motif exemplifies this gap. Such results temper optimistic claims (e.g., Deng et al., 2024; Yu et al., 2024; Qu et al., 2026) that LLM poetry approaches human quality: the alignment paradox shows that models may agree semantically on styles they cannot actually realize.

**Models recombine learned patterns rather than create genuinely new poetry.** High metric overlap of traditional imagery and frequent reuse of

canonical lines indicate recombination rather than genuine innovation (Bender et al., 2021). LLM output largely falls into the so-called “combinatorial creativity” - producing novel combinations of familiar ideas (Boden, 2003), as in humans. Even Qwen’s superior performance reflects precise pattern matching and constraint satisfaction more than novel literary creation, reinforcing critiques of LLMs as advanced auto-completion systems.

**A three-step evaluation framework offers a path toward more reliable creative-AI assessment.** Our framework-integrating computational diagnostics, cross-model evaluation, and expert judgment-provides a transferable methodology for assessing creative AI in culturally rich domains. Future research should (i) design metrics that jointly capture structural fidelity and cultural nuance, (ii) diversify training and evaluation data to reduce shared blind spots, and (iii) emphasize explicit structural constraints over vague stylistic prompts.

**Qualitative analysis.** Appendix J further presents a few qualitative cases of the LLM generated poems.

## 6 Conclusion

Our study demonstrates that while state-of-the-art LLMs can generate fluent Tang-style poetry, they exhibit critical limitations that surface-level metrics and cross-model evaluation often fail to detect. The proposed three-step framework, combining computational analysis, LLM-based assessment, and human expert validation, reveals systematic biases, including echo-chamber effects, shared poem outputs and shallow cultural association fidelity. Models tend to recombine learned patterns rather than create genuinely novel poetic expressions, and high inter-model agreement does not guarantee alignment with human judgments.

These findings underscore the continued importance of human oversight in culturally and structurally complex creative tasks. Future work should focus on multi-metric diagnostic evaluation, integration of explicit structural and cultural constraints, and strategies to mitigate shared blind spots in LLM assessment. Overall, our results highlight both the potential and the current limitations of LLMs in creative language generation, providing a roadmap for more reliable and culturally informed AI evaluation and development.

## Limitations

The focus on classical Chinese poetry, while providing cultural specificity and technical constraints necessary for rigorous evaluation, could have the limits the generalizability of findings to other creative domains in different languages. Future research should extend this methodological approach to other languages in digital literacy domains to test the multilingual and multicultural capabilities of the models, thereby to test the broader applicability of the echo chamber effect and evaluation framework.

The computational metrics we employed, particularly semantic similarity measurements, were limited by their reliance on models trained primarily on prosaic text. Future work should develop evaluation metrics specifically designed for poetic and metaphorical language, potentially incorporating specialized embeddings trained on literary corpora or rule-based systems that can capture formal poetic structures.

A further limitation concerns the study’s scope regarding prompt sensitivity and latent content biases. We did not exhaustively test diverse prompting architectures, such as Chain-of-Thought (Wei et al., 2022) or role-playing protocols, which might enhance poetic reasoning. Additionally, the models’ underlying safety alignments may suppress certain historical themes or archaic expressions. Future work should specifically investigate how to decouple these biases from valid creative constraints to improve generation authenticity.

## Ethical Considerations

We address some potential ethical considerations regarding the paper in this section.

**Risk Concerns.** This work examines LLMs in the culturally sensitive domain of Tang poetry generation. While the generated texts draw on public-domain traditions, they may misrepresent or oversimplify cultural values if interpreted outside research contexts. We caution against deploying LLM-generated poems as authentic cultural artifacts and emphasize that they should only serve as analytical probes of model capabilities. Our study also reveals systematic evaluation biases, underscoring the risks of over-reliance on automated judgments without human validation. By releasing code and generated outputs for research purposes only (the code and the outputs have been submit-

ted together with the paper as supplementary documents and will be publicly released upon paper publication), we aim to promote transparency while minimizing risks of misappropriation.

**Use of AI Assistants.** The authors acknowledge the use of ChatGPT solely for correcting grammatical errors, enhancing the coherence of the final manuscript, and providing assistance with coding.

## References

- Douglas Bates, Martin Mächler, Ben Bolker, and Steve Walker. 2015. [Fitting linear mixed-effects models using lme4](#). *Journal of Statistical Software*, 67(1):1–48.
- Emily M. Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. [On the dangers of stochastic parrots: Can language models be too big?](#) In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, FAccT ’21, pages 610–623, New York, NY, USA. Association for Computing Machinery.
- Margaret A. Boden. 2003. *The Creative Mind: Myths and Mechanisms*, 2nd edition. Routledge, London.
- Zong-Qi Cai, editor. 2008. *How to Read Chinese Poetry: A Guided Anthology*. Columbia University Press.
- Tuhin Chakrabarty, Vishakh Padmakumar, He He, and Nanyun Peng. 2023. [Creative natural language generation](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing: Tutorial Abstracts*, pages 34–40, Singapore. Association for Computational Linguistics.
- Chi-Min Chan, Weize Chen, Yusheng Su, Jianxuan Yu, Wei Xue, Shanghang Zhang, Jie Fu, and Zhiyuan Liu. 2024. [Chateval: Towards better LLM-based evaluators through multi-agent debate](#). In *The Twelfth International Conference on Learning Representations*.
- Stergios Chatzikyriakidis and Anastasia Natsina. 2025. [Poetry in RAGs: Modern Greek interwar poetry generation using RAG and contrastive training](#). In *Proceedings of the 5th International Conference on Natural Language Processing for Digital Humanities*, pages 257–264, Albuquerque, USA. Association for Computational Linguistics.
- Yanran Chen, Hannes Gröner, Sina Zarriß, and Steffen Eger. 2024. [Evaluating diversity in automatic poetry generation](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 19671–19692, Miami, Florida, USA. Association for Computational Linguistics.
- Matteo Cinelli, Gianmarco De Francisci Morales, Alessandro Galeazzi, Walter Quattrocchi, and

- Michele Starnini. 2021. [The echo chamber effect on social media](#). *Proceedings of the National Academy of Sciences of the United States of America*, 118(9):e2023301118.
- Elizabeth Clark, Tal August, Sofia Serrano, Nikita Haduong, Suchin Gururangan, and Noah A. Smith. 2021. [All that’s ‘human’ is not gold: Evaluating human evaluation of generated text](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 7282–7296, Online. Association for Computational Linguistics.
- DeepSeek-AI, Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, Xiaokang Zhang, Xingkai Yu, Yu Wu, Z. F. Wu, Zhibin Gou, Zhihong Shao, Zhuoshu Li, Ziyi Gao, and 181 others. 2025. [Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning](#). *Preprint*, arXiv:2501.12948.
- Zekun Deng, Hao Yang, and Jun Wang. 2024. [Can ai write classical chinese poetry like humans? an empirical study inspired by turing test](#). *Preprint*, arXiv:2401.04952.
- Marjan Ghazvininejad, Xing Shi, Jay Priyadarshi, and Kevin Knight. 2017. [Hafez: an interactive poetry generation system](#). In *Proceedings of ACL 2017, System Demonstrations*, pages 43–48, Vancouver, Canada. Association for Computational Linguistics.
- Team GLM, :, Aohan Zeng, Bin Xu, Bowen Wang, Chenhui Zhang, Da Yin, Dan Zhang, Diego Rojas, Guanyu Feng, Hanlin Zhao, Hanyu Lai, Hao Yu, Hongning Wang, Jiadai Sun, Jiajie Zhang, Jiale Cheng, Jiayi Gui, Jie Tang, and 40 others. 2024. [Chatglm: A family of large language models from glm-130b to glm-4 all tools](#). *Preprint*, arXiv:2406.12793.
- Carlos Gómez-Rodríguez and Paul Williams. 2023. [A confederacy of models: a comprehensive evaluation of LLMs on creative writing](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 14504–14528, Singapore. Association for Computational Linguistics.
- Mika Härmäläinen, Khalid Alnajjar, and Thierry Poibeau. 2022. [Modern french poetry generation with roberta and GPT-2](#). In *Proceedings of the 13th International Conference on Computational Creativity, ICCO 2022, Bozen-Bolzano, Italy, June 27 - July 1, 2022*, pages 12–16. Association for Computational Creativity (ACC).
- Jing He, Ming Zhou, and Long Jiang. 2021. [Generating chinese classical poems with statistical machine translation models](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 26(1):1650–1656.
- Zhengyu Hu, Jieyu Zhang, Zhihan Xiong, Alexander Ratner, Kaize Ding, and Ranjay Krishna. 2026. [Towards acyclic preference evaluation of language models via multiple evaluators](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 40(26):21903–21911.
- Chihan Huang and Xiaobo Shen. 2025. [PoemBERT: A dynamic masking content and ratio based semantic language model for Chinese poem generation](#). In *Proceedings of the 31st International Conference on Computational Linguistics*, pages 50–60, Abu Dhabi, UAE. Association for Computational Linguistics.
- Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, Léo Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timothée Lacroix, and William El Sayed. 2023. [Mistral 7b](#). *Preprint*, arXiv:2310.06825.
- Ilya Koziev and Alena Fenogenova. 2025. [Generation of Russian poetry of different genres and styles using neural networks with character-level tokenization](#). In *Proceedings of the 9th Joint SIGHUM Workshop on Computational Linguistics for Cultural Heritage, Social Sciences, Humanities and Literature (LaTeCH-CLfL 2025)*, pages 47–63, Albuquerque, New Mexico. Association for Computational Linguistics.
- James J. Y. Liu. 1962. *The Art of Chinese Poetry*. University of Chicago Press.
- Yang Liu, Dan Iter, Yichong Xu, Shuhang Wang, Ruochen Xu, and Chenguang Zhu. 2023. [G-eval: NLG evaluation using gpt-4 with better human alignment](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 2511–2522, Singapore. Association for Computational Linguistics.
- Yang Liu, Lan Lan, Jiahuan Cao, Hiuyi Cheng, Kai Ding, and Lianwen Jin. 2025. [Large-scale corpus construction and retrieval-augmented generation for Ancient Chinese poetry: New method and data insights](#). In *Findings of the Association for Computational Linguistics: NAACL 2025*, pages 779–817, Albuquerque, New Mexico. Association for Computational Linguistics.
- Bolei Ma, Berk Yoztyurk, Anna-Carolina Haensch, Xinpeng Wang, Markus Herklotz, Frauke Kreuter, Barbara Plank, and Matthias Aßenmacher. 2025. [Algorithmic fidelity of large language models in generating synthetic German public opinions: A case study](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1785–1809, Vienna, Austria. Association for Computational Linguistics.
- Jingkun Ma, Runzhe Zhan, and Derek F. Wong. 2023. [Yu sheng: Human-in-loop classical Chinese poetry generation system](#). In *Proceedings of the 17th Conference of the European Chapter of the Association*

- for *Computational Linguistics: System Demonstrations*, pages 57–66, Dubrovnik, Croatia. Association for Computational Linguistics.
- Christopher D. Manning and Hinrich Schütze. 1999. *Foundations of Statistical Natural Language Processing*. MIT Press, Cambridge, MA.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [Bleu: a method for automatic evaluation of machine translation](#). In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Xingjuan Peng, Huan Cao, Zhaojing Shi, and Qiming Cao. 2024. [Vegetative imagery and emotional expression: Symbolism and heritage of plant imagery in chinese classical poetry](#). *Frontiers in Humanities and Social Sciences*, 4(3):241–252.
- Zhan Qu, Shuzhou Yuan, and Michael Färber. 2026. [Poetone: A framework for constrained generation of structured chinese songci with llms](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 40(39):32764–32772.
- Nils Reimers and Iryna Gurevych. 2019. [Sentence-BERT: Sentence embeddings using Siamese BERT-networks](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3982–3992, Hong Kong, China. Association for Computational Linguistics.
- Karen Sparck Jones. 1988. *A statistical interpretation of term specificity and its application in retrieval*, pages 132–142. Taylor Graham Publishing, GBR.
- Gemma Team, Thomas Mesnard, Cassidy Hardin, Robert Dadashi, Surya Bhatnagar, Shreya Pathak, Laurent Sifre, Morgane Rivière, Mihir Sanjay Kale, Juliette Love, Pouya Tafti, Léonard Hussenot, Pier Giuseppe Sessa, Aakanksha Chowdhery, Adam Roberts, Aditya Barua, Alex Botev, Alex Castro-Ros, Ambrose Slone, and 89 others. 2024. [Gemma: Open models based on gemini research and technology](#). *Preprint*, arXiv:2403.08295.
- Tim Van de Cruys. 2020. [Automatic poetry generation from prosaic text](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2471–2480, Online. Association for Computational Linguistics.
- Peiyi Wang, Lei Li, Liang Chen, Zefan Cai, Dawei Zhu, Binghuai Lin, Yunbo Cao, Lingpeng Kong, Qi Liu, Tianyu Liu, and Zhifang Sui. 2024. [Large language models are not fair evaluators](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 9440–9450, Bangkok, Thailand. Association for Computational Linguistics.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc V Le, and Denny Zhou. 2022. [Chain-of-thought prompting elicits reasoning in large language models](#). In *Advances in Neural Information Processing Systems*, volume 35, pages 24824–24837. Curran Associates, Inc.
- Aiyuan Yang, Bin Xiao, Bingning Wang, Borong Zhang, Ce Bian, Chao Yin, Chenxu Lv, Da Pan, Dian Wang, Dong Yan, Fan Yang, Fei Deng, Feng Wang, Feng Liu, Guangwei Ai, Guosheng Dong, Haizhou Zhao, Hang Xu, Haoze Sun, and 36 others. 2025a. [Baichuan 2: Open large-scale language models](#). *Preprint*, arXiv:2309.10305.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, and 41 others. 2025b. [Qwen3 technical report](#). *Preprint*, arXiv:2505.09388.
- Cheng Yang, Maosong Sun, Xiaoyuan Yi, and Wenhao Li. 2018. [Stylistic Chinese poetry generation via unsupervised style disentanglement](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3960–3969, Brussels, Belgium. Association for Computational Linguistics.
- Xinyu Yao, Mengdi Wang, Bo Chen, and Xiaobing Zhao. 2025. [Wenyangpt: A large language model for classical chinese tasks](#). In *Proceedings of the Thirty-Fourth International Joint Conference on Artificial Intelligence, IJCAI-25*, pages 8339–8347. International Joint Conferences on Artificial Intelligence Organization. Main Track.
- Chengyue Yu, Lei Zang, Jiaotuan Wang, Chenyi Zhuang, and Jinjie Gu. 2024. [CharPoet: A Chinese classical poetry generation system based on token-free LLM](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 3: System Demonstrations)*, pages 315–325, Bangkok, Thailand. Association for Computational Linguistics.
- Guo Zhipeng, Xiaoyuan Yi, Maosong Sun, Wenhao Li, Cheng Yang, Jiannan Liang, Huimin Chen, Yuhui Zhang, and Ruoyu Li. 2019. [Jiuge: A human-machine collaborative Chinese classical poetry generation system](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 25–30, Florence, Italy. Association for Computational Linguistics.
- Danhao Zhu, Zhixiao Zhao, Yiping Zhang, Guangyao Sun, Chang Liu, Die Hu, and Dongbo Wang. 2024. Research on large language model evaluation for the generation task of natural language processing in classical chinese. *Journal of Information Resources Management*, 14(5):45–58.

Xu Zou. 2025. [BIPro: Zero-shot Chinese poem generation via block inverse prompting constrained generation framework](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1116–1134, Vienna, Austria. Association for Computational Linguistics.

## A Poetry Dimensions for Generation

We cover the five important dimensions in Tang poetry: poetic form (格式), target poet style (诗人), theme (主题), emotion (情感), and imagery (意象), with each including representative specific elements (Table 4).

Dim.	Specific Elements
格式	五/七言绝句、五/七言律诗
诗人	李白、杜甫、白居易、王维、李商隐
主题	山水、思乡、怀古、田园、送别
情感	悲伤、宁静、豪放、浪漫、喜悦
意象	风、花、柳、月、雁

Table 4: Poetry dimensions and elements for generation in Chinese. The English translation is in §3.1.

## B Generation and Evaluation Prompt Templates

**Generation prompts.** For each model, the generation prompts provided explicit instructions regarding the required poetic form (e.g., “five-character quatrain”), the style of a specific poet (e.g., “in the style of Li Bai”), the thematic content (e.g., “farewell”), the emotion (e.g., “melancholy”), and the imaginary (e.g., “willows”). An example prompt (with translation in English) is shown below:

**ZH:** 请模仿杜甫的风格，创作一首关于送别的七言绝句，体现悲伤情感，并包含柳的意象。请只输出诗歌本身。

**EN:** Please imitate the style of Du Fu and compose a seven-character quatrain about farewell, conveying melancholy and including the imagery of willows. Output only the poem itself.

**Evaluation prompts.** Evaluation prompts instructed the model to assess a given poem on five

dimensions: *Prosodic Adherence*, *Thematic Relevance*, *Emotional Consistency*, *Imagery & Structure*, and *Classical Language Authenticity*. Each dimension required an integer score from 1 to 5, returned in structured JSON. The original Chinese prompt and its translation are shown as follows:

**ZH:**

# 角色  
你是一位中国古典文学评价专家。请对以下诗歌进行评分，每个维度给出1-5的整数分数。

# 评分标准  
1分 = 很差  
2分 = 较差  
3分 = 一般  
4分 = 良好  
5分 = 优秀

# 评价维度

1. \*\*韵律格律 (prosodic adherence)\*\*: 音韵、平仄、押韵、节奏  
2. \*\*主题切合度 (thematic relevance)\*\*: 主题明确性与契合度  
3. \*\*情感一致性 (emotional consistency)\*\*: 情感表达的真挚性与一致性  
4. \*\*意象与结构 (imagery structure)\*\*: 意象生动性与结构合理性  
5. \*\*语言经典性 (language authenticity)\*\*: 语言的古典韵味与准确性

# 诗歌信息

- \*\*诗人\*\*: {poet}  
- \*\*主题\*\*: {theme}  
- \*\*情感\*\*: {emotion}  
- \*\*意象\*\*: {imagery}  
- \*\*形式\*\*: {form}  
- \*\*诗歌正文\*\*: {poem\_text}

# 输出要求

请直接输出JSON格式的评分结果，每个维度只需给出1-5的整数分数:

```
{
  "prosodic_adherence": 分数,
```

```

    "thematic_relevance": 分数,
    "emotional_consistency": 分数,
    "imagery_structure": 分数,
    "language_authenticity": 分数
}

```

**EN:**

```

# Role
You are an expert in evaluating
↳ classical Chinese literature.
↳ Please score the following poem,
↳ giving an integer from 1 to 5
↳ for each dimension.

# Scoring Criteria
1 = Very Poor
2 = Poor
3 = Fair
4 = Good
5 = Excellent

# Evaluation Dimensions
1. prosodic adherence: Rhyme,
↳ tonal pattern, rhyme scheme,
↳ rhythm
2. thematic relevance: Clarity
↳ and appropriateness of theme
3. emotional consistency:
↳ Sincerity and consistency of
↳ emotional expression
4. imagery structure: Vividness
↳ of imagery and soundness of
↳ structure
5. language authenticity:
↳ Classical flavor and accuracy of
↳ language

# Poem Information
- Poet: {poet}
- Theme: {theme}
- Emotion: {emotion}
- Imagery: {imagery}
- Form: {form}
- Poem Text: {poem_text}

# Output Requirement
Please output the scoring results
↳ directly in JSON format, giving
↳ an integer 1-5 for each
↳ dimension:

{

```

```

    "prosodic_adherence": score,
    "thematic_relevance": score,
    "emotional_consistency": score,
    "imagery_structure": score,
    "language_authenticity": score
}

```

## C Poem Generation Procedure

For each of the six models, we independently generated poems for every combination of target poet, theme, emotion, imagery, and form. Prompts were derived from a fixed template shown in the previous section (Appendix B) and explicitly instructed the model to satisfy all specified constraints while emulating the designated poet’s style (Figure 13).

The generation loop enumerated all task combinations using Python’s `itertools.product`. For each task, the model produced a poem whose raw output was immediately cleaned to isolate the poetic text.<sup>4</sup> The output was then validated against basic structural and content requirements. If a poem failed validation, the system automatically retried generation up to five times, incrementally increasing the sampling temperature to promote diversity. All valid outputs were written incrementally to JSON Lines (`.jsonl`) files to prevent data loss in the event of interruption.

## D Generated Poem Cleaning Heuristics

Raw outputs often contain prefixes, commentary, or formatting noise. Cleaning follows a two-step procedure: first, a model-driven cleaner (e.g., Qwen2.5) is prompted to extract the text into JSON-formatted poems; if this fails, regex and heuristic rules are applied as fallback to remove explanatory phrases or extraneous lines, while retaining necessary generation traces for later analysis. The cleaned results are stored in a standardized structure for automated evaluation.

In cases where model-driven cleaning (Qwen2.5 extractor) failed to return a valid JSON object, regex-based heuristics were applied:

- Remove leading/trailing commentary or role markers (e.g., “As an AI model, here is your poem:”).
- Collapse multiple empty lines into a single line break.

<sup>4</sup>This initial cleaning was a lightweight check to remove empty or malformed responses and does not duplicate the subsequent validation step.

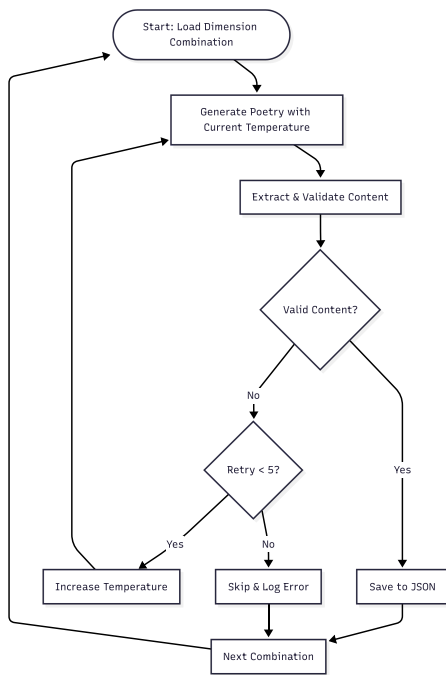


Figure 13: Poem Generation Flow Chart. The experimental loop generated poems for combinations of all dimension elements, after which outputs were cleansed and validated against structural and content criteria. Failed generations triggered adaptive retries with increased temperature, and all results were stored incrementally in JSONL to ensure reliability.

- Retain non-standard artifacts (e.g., repeated punctuation) when directly embedded in poem lines.

The following prompt (original prompt in Chinese and translated prompt in English) was used for poem cleaning:

**ZH:**

"""请从以下文本中提取纯净的诗歌内容。

**\*\*规则\*\*:**

- 保留诗歌的所有诗句和原有的换行分段
- 删除任何前缀后缀（如“好的，这是一首...”、“希望您喜欢”等）
- 删除诗歌标题（如《送别》）
- 删除任何注释、解释或旁白
- 删除与诗歌无关的对话性文字

**\*\*输出格式\*\*:** 请严格按照JSON格式返

回，如果没有诗歌则extracted\_poem为空字符串:

```
{{"status": "success", "extracted_poem": "提取的纯净诗歌文本"}}
```

**\*\*待处理文本\*\*:**

```
{raw_text}
```

**\*\*回答\*\*:** """

**EN:**

"Please extract the pure poetic content from the following text.

Rules:

- Retain all lines of the poem and the original line breaks
- Remove any prefixes or suffixes (e.g., 'This is a poem' or 'I hope you enjoy it')
- Remove the poem title (e.g., 'Farewell')
- Remove any annotations, explanations, or commentary
- Remove any dialogue-related text unrelated to the poem

Output Format: Please return the

output strictly in JSON format.

If no poem is extracted, set

'extracted\_poem' to an empty

string:

```
{'status': 'success',
```

```
'extracted_poem': 'Extracted
```

```
pure poetic text'}
```

Raw Text: {raw\_text}

Response: "

```
%\end{verbatim}
```

The flow for data cleaning could be seen in Figure 14

## E Embedding models

Sentence-level embeddings were obtained using a Chinese sentence transformer (compatible with

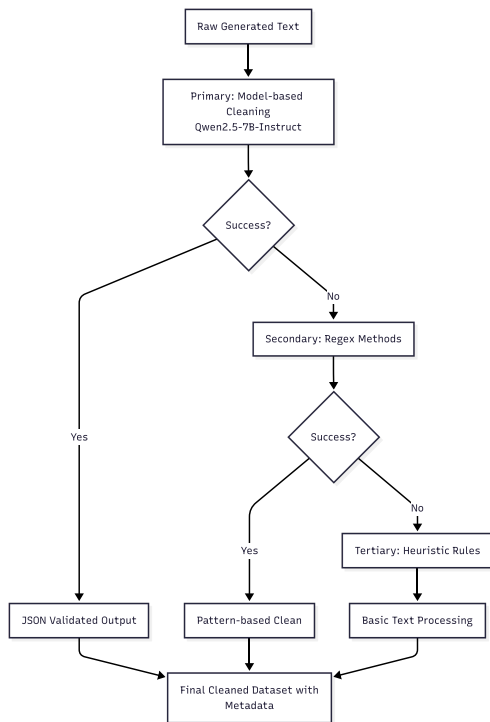


Figure 14: Poem Cleaning Flow Chart. For text cleaning, I used Qwen, which had an excellent performance for generating complete and correctly formatted poems in the initial format consistency review. A structured prompt guided the model to remove extraneous content and output the cleaned text strictly as a JSON object.

Reimers and Gurevych (2019)). For coherence analysis, cosine similarity was computed both across poem-prompt pairs and between adjacent lines.

## F Hyperparameters

- **Generation temperature:**  $T = 0.4$  (with up to 2 retries for diversity).
- **Evaluation temperature:**  $T = 0.2$  (to enforce deterministic scoring).
- **Maximum length:** 128 tokens (covering the longest regulated verse).

## G Data Formats and Storage

All cleaned outputs were stored in structured JSON with metadata including:

- Model identifier
- Prompt specification (form, style, theme, emotion, imagery)
- Generated poem text

- Evaluation scores (when available)

This design ensured automatic parsing for downstream statistical analysis.

## H Statistical Procedures

Entropy comparisons were conducted using one-way ANOVA, with Tukey HSD for post-hoc tests. For evaluator-generator interactions, linear mixed-effects models were implemented in lme4 (Bates et al., 2015) (R package), treating generator/evaluator as fixed and random effects respectively.

## I Human Expert Evaluation

One of the leading authors and another project collaborator were asked to conduct the evaluation of the 150 samples (500 each from Qwen, DeepSeek, and Gemma, the three highest-ranked generators in Step 2 LLM evaluations) as volunteers. There exists disagreement on minor cases that were resolved after discussion. Both annotators are consent about the annotated data use.

## J Qualitative Analysis

While large-scale metrics provide an essential overview of model performance, close reading of individual outputs exposes error patterns that aggregate statistics cannot capture. This section examines 150 poems (50 each) generated by the three top-performing LLMs, Qwen, Gemma, and DeepSeek, highlighting characteristic strengths and weaknesses.

**DeepSeek: Structural Confusion and Collage Effects.** DeepSeek frequently produced text that was descriptive commentary rather than original poetry. For example, when prompted for a five-character quatrain on “homesickness” in the style of Li Bai, it generated a prose-like analysis of Li Bai’s famous poem (Table 5). Other outputs resembled collages of celebrated lines from multiple poets (e.g., Gao Shi, Cen Shen, Liu Zongyuan, and even Mao Zedong), resulting in structurally incoherent verse without a unified voice or theme. Nevertheless, DeepSeek generally respected basic distinctions among quatrains, regulated verse, and line length, suggesting only a partial grasp of formal constraints.

**Gemma: Systematic Format Violations and Cross-Lingual Intrusions.** Gemma’s strongest semantic performance was undermined by persistent failures of form. It routinely defaulted to seven-character quatrains, even when explicitly instructed

Dim.	Elements	English Translation
Poet	李白	Li Bai
Theme	思乡	Homesickness
Emotion	悲伤	Sadness
Imagery	月亮	Moon
Form	五言绝句	Five-character Quatrain

#### Generated Text (Original)

是唐代诗人李白的代表作之一，以简洁明快的语言和深远的意境，表达了作者对故乡的深深思念。诗中“床前明月光”描绘了一个宁静的夜晚，月光如水，洒在床前，营造出一种幽远而神秘的氛围。后两句“疑是地上霜”则巧妙地运用比喻，将月光比作秋霜，既展现了月光的皎洁，又暗示了季节的更替，引发读者对时光流逝的感慨。整首诗...

#### Generated Text (Translation)

This is one of the representative works of the Tang Dynasty poet Li Bai, which uses concise language and profound artistic conception to express the author's deep longing for his hometown. In the poem, 'Bright moonlight before my bed' depicts a quiet night where the moonlight, like water, spills before the bed, creating a distant and mysterious atmosphere. The latter two lines, 'I mistake it for frost on the ground,' cleverly use a metaphor, comparing the moonlight to autumn frost, which not only shows the brightness of the moon but also hints at the changing of seasons, evoking the reader's sigh over the passage of time. The entire poem...

Table 5: Example of a descriptive model output with English translation.

to produce other forms such as five-character regulated verse. A second recurring error was cross-lingual contamination, with English words embedded within Chinese lines (Table 6). Table 7 illustrates another case where Gemma ignored the requested five-character regulated structure. Despite these technical lapses, human evaluators often rated Gemma highly for theme, emotion, and imagery, underscoring a notable gap between semantic quality and formal correctness.

Dim.	Elements	English Translation
Poet	李白	Li Bai
Theme	怀古	Nostalgic
Emotion	豪放	Bold
Imagery	花	Flower
Form	七言绝句	Seven-character Quatrain

**Generated Poem**

Original Chinese	English Translation
落花飘零满地红，	Fallen flowers drift down, covering the ground in red.
旧梦如烟随风blown。	Old dreams, like smoke, are blown with the wind.
举杯邀月寄相思，	I raise my cup to the moon, sending my longing.
何处寻得当年雄。	Where can I find the hero of bygone years?

Table 6: Example of a Model Output with Cross-Lingual Contamination.

**Qwen: Consistent Formal Control with Minor Literary Issues.** Qwen displayed near-perfect adherence to prescribed forms and meter. Format violations were rare, and outputs consistently satisfied structural and stylistic requirements. Remaining issues were subtle: occasional thematic ambiguity or slight rhythmic unevenness that affected aesthetic polish but not correctness. This reliability confirms Qwen as the most technically proficient

Dim.	Elements	English Translation
Poet	白居易	Bai Juyi
Theme	怀古	Nostalgic
Emotion	豪放	Bold
Imagery	月	Moon
Requested Form	五言律诗	Five-character Regulated Verse
Actual Form	七言绝句	Seven-character Quatrain

#### Generated Poem

Original Chinese	English Translation
旧梦如烟消残阳，	Old dreams, like smoke, fade in the setting sun,
明月照破落城墙。	The bright moon shines on the dilapidated city walls.
英雄泪湿青史册，	A hero's tears wet the pages of history,
多少悲欢今何在？	How many joys and sorrows, where are they now?

Table 7: Example of a Model Output Exhibiting a Format Violation.

generator in the study.

## K Additional Results – Information Gain

To quantify how different creative dimensions in the prompts influence the diversity of the generated poems, we analyzed the information gain for each dimension, drawing from Ma et al. (2025). It measures how much information one random variable provides about another. It is calculated as the difference between the entropy of the variable from general group (Group) and its conditional entropy given another variable from a specific subgroup (Subgroup):

$$I(X; Y) = H(Y) - H(Y | X) \quad (5)$$

It indicates how much knowing one variable (e.g.,  $X$ ) reduces uncertainty about another variable (e.g.,  $Y$ ). A higher information gain indicates that knowing one variable reduces uncertainty about another variable. A significant drop in entropy indicates that a given dimension strongly constrains the model's vocabulary, forcing it into a more predictable and focused generative mode. The analysis reveals which dimensions have the most and least impact on the models' outputs.

### Poetic Form as the Strongest Constraint.

Across all five dimensions, Poetic Form provides the highest information gain for nearly all models (see Figure 15). The highly structured forms, such as “7-Char Quatrain” and “5-Char Regulated Verse”, consistently cause the most significant drop in entropy. This is a logical finding, as the strict rules of meter, tone, and rhyme inherent in these forms naturally limit the permissible word choices, thereby reducing the diversity of the output. This demonstrates that all models, to varying degrees,

are sensitive to the explicit structural rules of classical poetry.

**Thematic Prompts Provide Evidence for Baichuan’s Template-Driven Strategy.** The analysis of the Theme dimension provides the strongest quantitative evidence for the “template-driven” hypothesis regarding the Baichuan model. When tasked with a specific theme like 怀古 (“Nostalgic”), Baichuan’s entropy plummets, showing by far the largest information gain of any model (see Figure 16). This aligns perfectly with its TF-IDF fingerprint, which showed a preference for grand, historical terms. This suggests that the “Historical” prompt triggers a rigid, pre-set template in the Baichuan model, leading to a highly predictable and low-diversity output.

**Minimal Impact from Poet, Emotion, and Imagery Dimensions.** In stark contrast to Poetic Form, the dimensions of Poet Style, Emotion, and Imagery show minimal information gain across most models (see Figure 17, Figure 18, Figure 19). The entropy lines for these dimensions are generally flat, indicating that specifying a poet like “Li Bai”, an emotion like “Sadness”, or an image like “Moon” does not significantly constrain the models’ vocabulary. This powerful finding corroborates the results from our semantic distinction analysis (Section 4.1.3), which showed that the models struggle to produce semantically distinct styles. They may be generating poems about Li Bai or sadness, but they are not adopting a specific, constrained vocabulary that reflects a deep understanding of that style or emotion.

**The Baichuan Anomaly: An Indicator of Extreme Rigidity.** An interesting anomaly is observed with the Baichuan model, where for certain prompts like the theme “Homesickness” (思乡) or the imagery “Wild Goose” (雁), the entropy paradoxically increases for the subgroup. A plausible explanation is that Baichuan’s general-purpose (Group) vocabulary is so exceptionally limited that a specific creative prompt forces it to access a different, slightly more varied set of words, thus slightly increasing its entropy relative to its own baseline. This counter-intuitive result further underscores the model’s extreme rigidity and non-flexible generative process.

## L Additional Results – LLM Evaluation

**Comprehensive Evaluator Profiles: Applying the Enhanced Framework.** To better understand

the dynamics of the LLM-as-a-judge evaluation, we analyzed the behavior of each LLM as an evaluator. This revealed four distinct evaluator profiles, each with a unique temperament and scoring bias. Mistral, for instance, acts as a “Generous Optimist”, while Qwen behaves as a “Strict Perfectionist”. A full comparative breakdown of these profiles, detailing their absolute and relative biases, is presented in Table 8.

## M Additional Results – Human Evaluation

Here we show the additional results of the overall score distribution and distributions by the five dimensions of human evaluations on the sampled subset in Figure 20, 21, 22. We notice that all three models were evaluated with moderately high scores, with Deepseek and Gemma lying slightly behind Qwen. We especially notice that for both Deepseek and Gemma models, the score for prosodic adherence is largely lower than average, indicating that the models might be less capable of the prosodic nuance when generating the poems.

We further did correlation analysis of the human evaluation with the LLM judges. Figure 23, 24, 25 visualize the correlations of human vs. LLM evaluation (averaged across 6 LLM judges) for Deepseek, Gemma and Qwen generated poems respectively. Figure 26, 27, 28 further visualize the the correlations of human vs. LLM evaluation by the dimension. We also show the human LLM correlation by the six judge models and different dimensions in Figure 29, 30, 31.

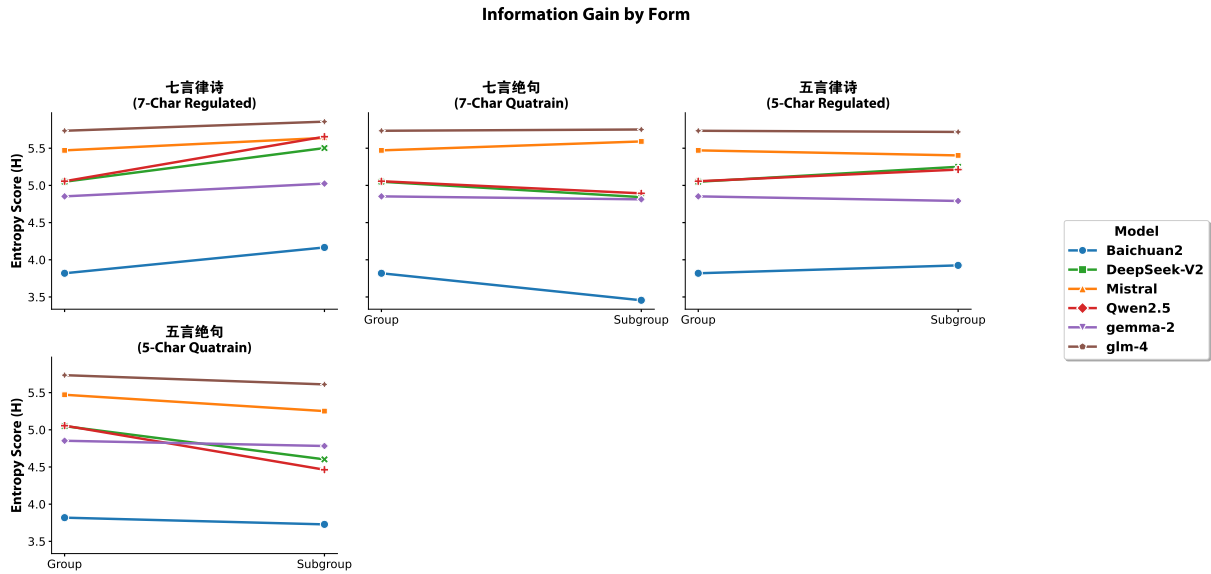


Figure 15: Information Gain by Form. Each subplot compares the overall group entropy ( $H(Y)$ , labeled 'Group') with the subgroup entropy conditioned on a specific poetic form ( $H(Y|X_i)$ , labeled 'Subgroup').

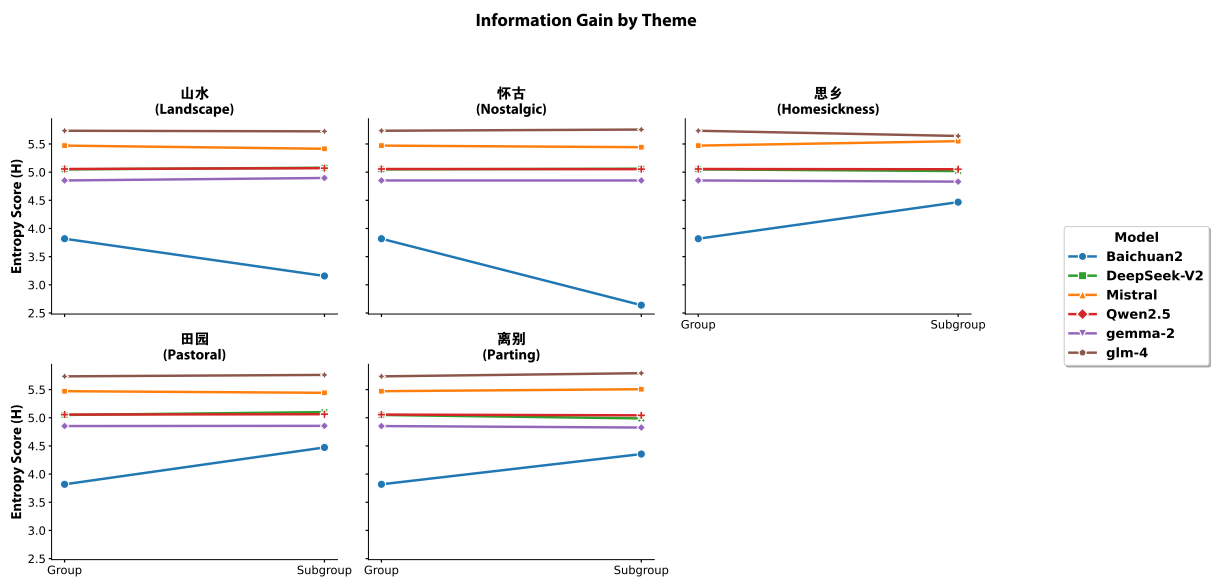


Figure 16: Information Gain by Theme. Each subplot compares the overall group entropy ( $H(Y)$ , labeled 'Group') with the subgroup entropy conditioned on a specific theme ( $H(Y|X_i)$ , labeled 'Subgroup').

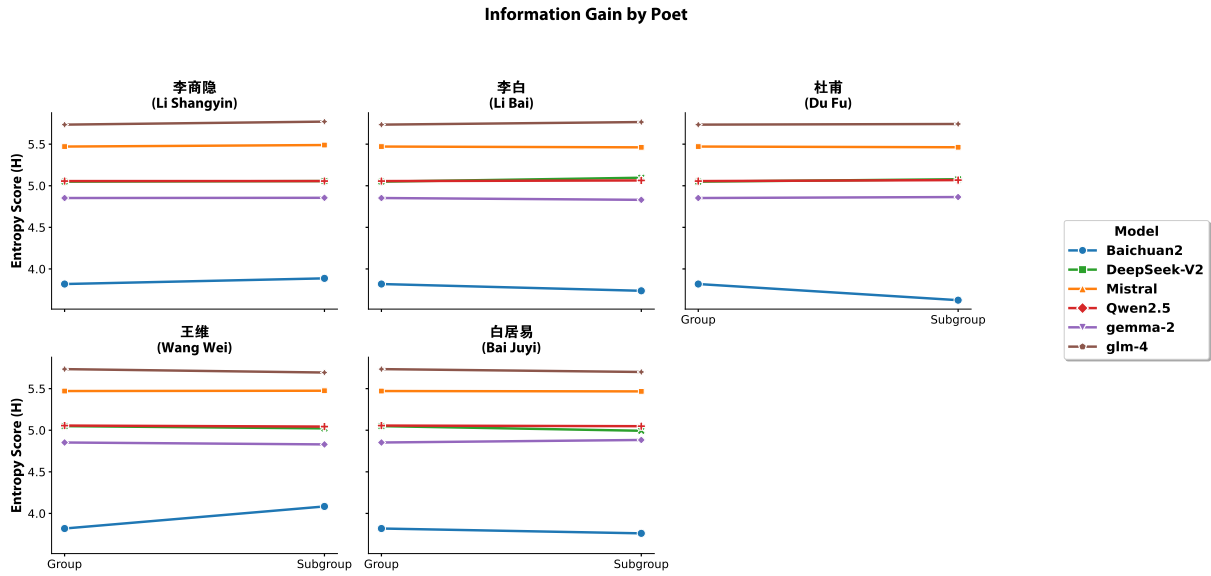


Figure 17: Information Gain by Poet. Each subplot compares the overall group entropy ( $H(Y)$ , labeled 'Group') with the subgroup entropy conditioned on a specific poet style ( $H(Y|X_i)$ , labeled 'Subgroup').

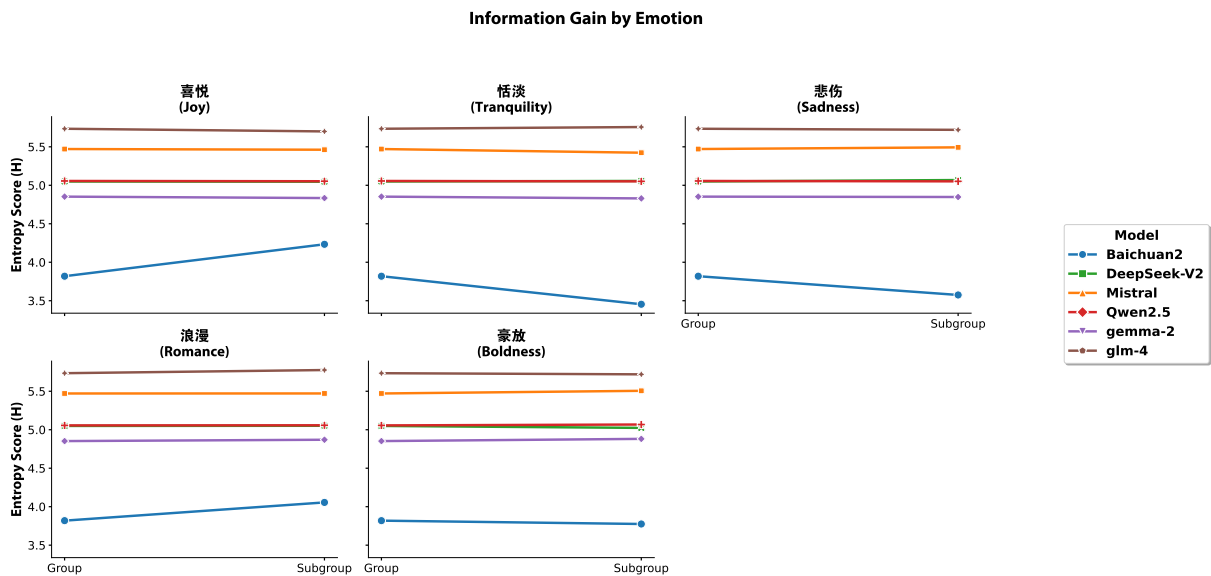


Figure 18: Information Gain by Emotion. Each subplot compares the overall group entropy ( $H(Y)$ , labeled 'Group') with the subgroup entropy conditioned on a specific emotion ( $H(Y|X_i)$ , labeled 'Subgroup').

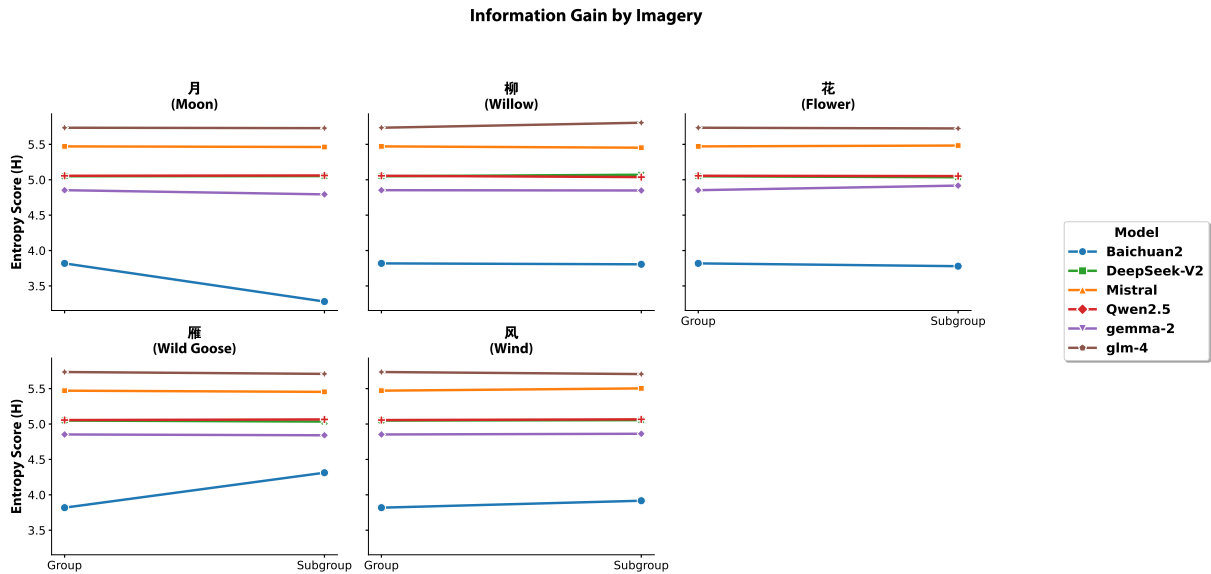


Figure 19: Information Gain by Imagery. Each subplot compares the overall group entropy ( $H(Y)$ , labeled 'Group') with the subgroup entropy conditioned on a specific core image ( $H(Y|X_i)$ , labeled 'Subgroup').

Attribute	Mistral (The Generous Optimist)	GLM (The Balanced Self-Advocate)	Baichuan (The Modest Performer)	Qwen (The Strict Perfectionist)
<b>Temperament</b>	Exceptionally lenient (Avg. Score: 4.64)	Moderately lenient (Avg. Score: 4.60)	Moderately lenient (Avg. Score: 4.50)	Exceptionally strict (Avg. Score: 3.86)
<b>Absolute Bias</b>	+0.08 (Self: 4.72 vs. Peers: 4.64)	+0.13 (Self: 4.73 vs. Peers: 4.60)	-0.13 (Self: 4.63 vs. Peers: 4.50)	+0.20 (Self: 4.06 vs. Peers: 3.86)
<b>Relative Bias</b>	Ranks itself 2nd (4.72) within its own scoring system.	Consistently ranks itself 1st (4.73) in its own evaluations.	Ranks itself highly but not at the top.	Ranks itself 2nd (4.06) within its own harsh framework.
<b>Profile</b>	Confident but not arrogant. Its generosity suggests an optimistic assessment philosophy.	Demonstrates healthy self-confidence with a slight tendency toward self-promotion.	Exhibits balanced self-evaluation with a tendency toward modesty.	Presents complex behavior, maintaining high self-regard within its own stringent standards.

Table 8: **Comparative Analysis of Evaluator Model Profiles.** The table summarizes the distinct “personalities” of the four LLM evaluators based on their scoring behavior and self-assessment tendencies. The data highlights a clear spectrum of evaluation temperaments, from the “Generous Optimist” (Mistral) with an average score of 4.64, to the “Strict Perfectionist” (Qwen) with an average score of 3.86. Each model’s absolute and relative biases in self-evaluation are also detailed.

### Human Evaluation Score Analysis - DeepSeek

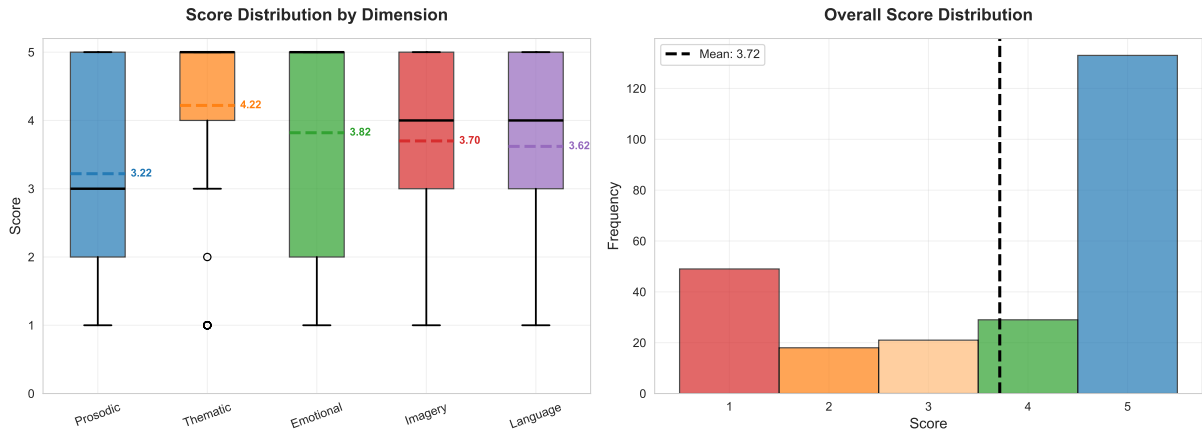


Figure 20: Human evaluation score analysis for Deepseek.

### Human Evaluation Score Analysis - Gemma

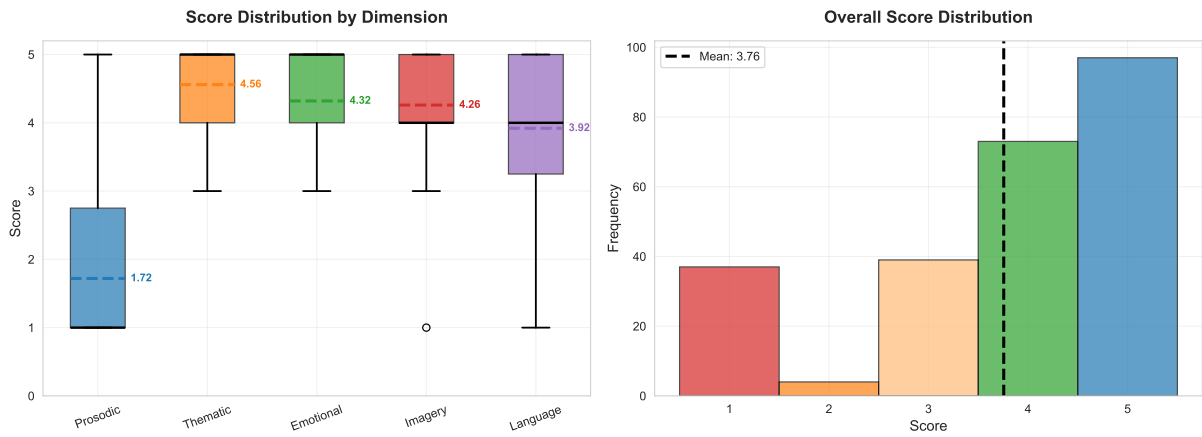


Figure 21: Human evaluation score analysis for Gemma.

### Human Evaluation Score Analysis - Qwen

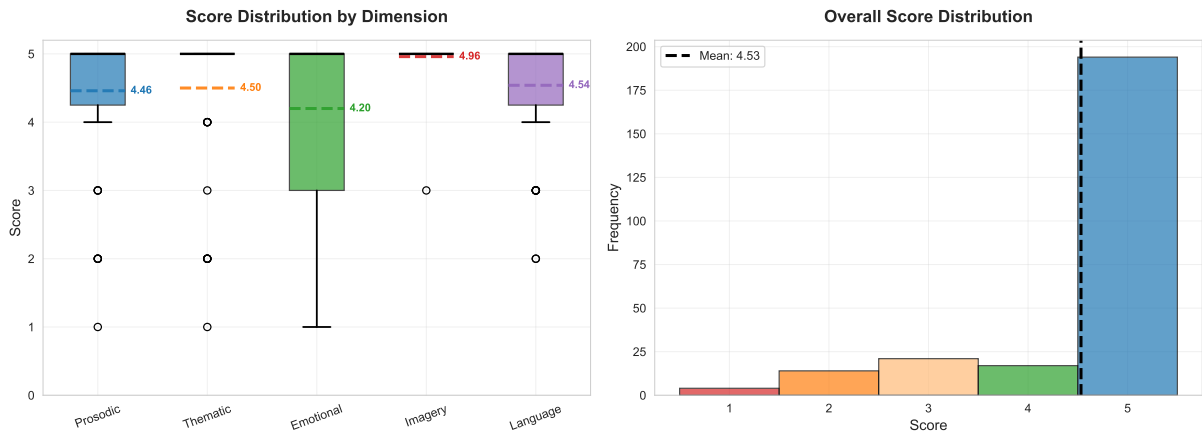


Figure 22: Human evaluation score analysis for Qwen.

Human vs LLM Evaluation Correlation Analysis - DeepSeek Generated Poems  
(Averaged Across 6 LLM Judges)

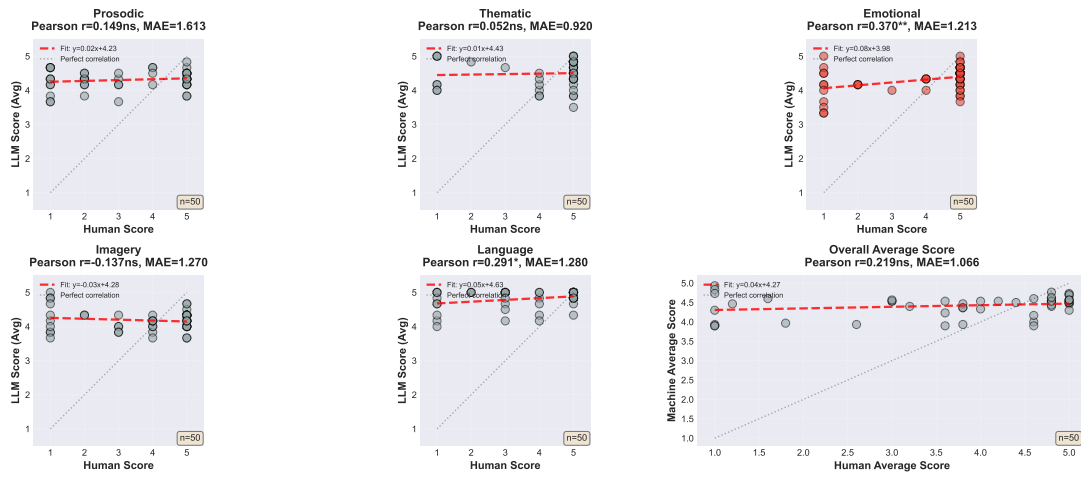


Figure 23: Human vs LLM evaluation correlation for DeepSeek generated poems (averaged across 6 LLM judges).

Human vs LLM Evaluation Correlation Analysis - Gemma Generated Poems  
(Averaged Across 6 LLM Judges)

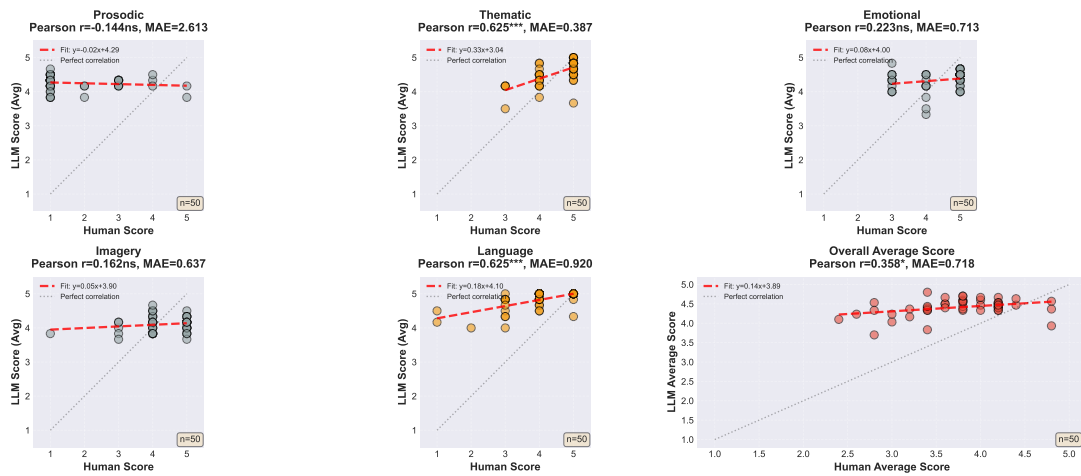


Figure 24: Human vs LLM evaluation correlation for Gemma generated poems (averaged across 6 LLM judges).

Human vs LLM Evaluation Correlation Analysis - Qwen Generated Poems  
(Averaged Across 6 LLM Judges)

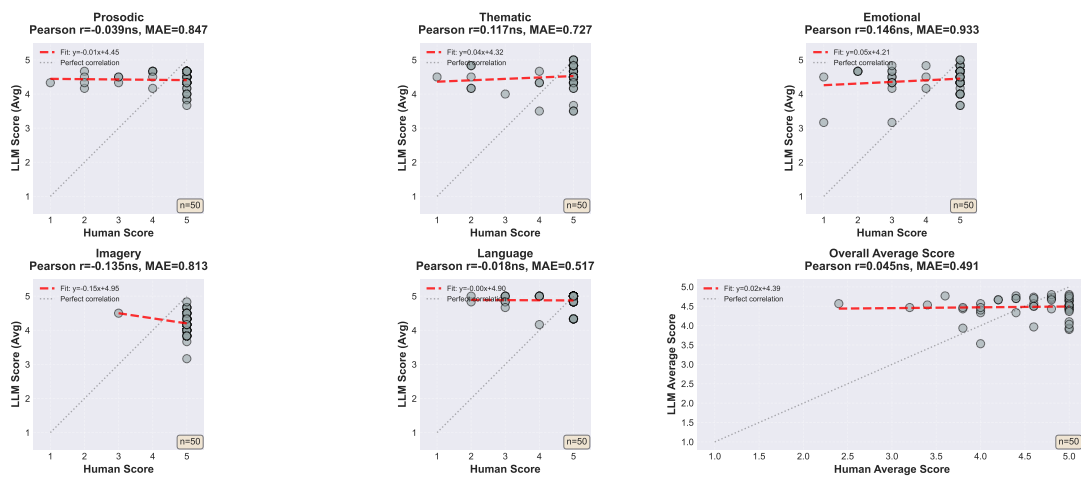


Figure 25: Human vs LLM evaluation correlation for Qwen generated poems (averaged across 6 LLM judges).

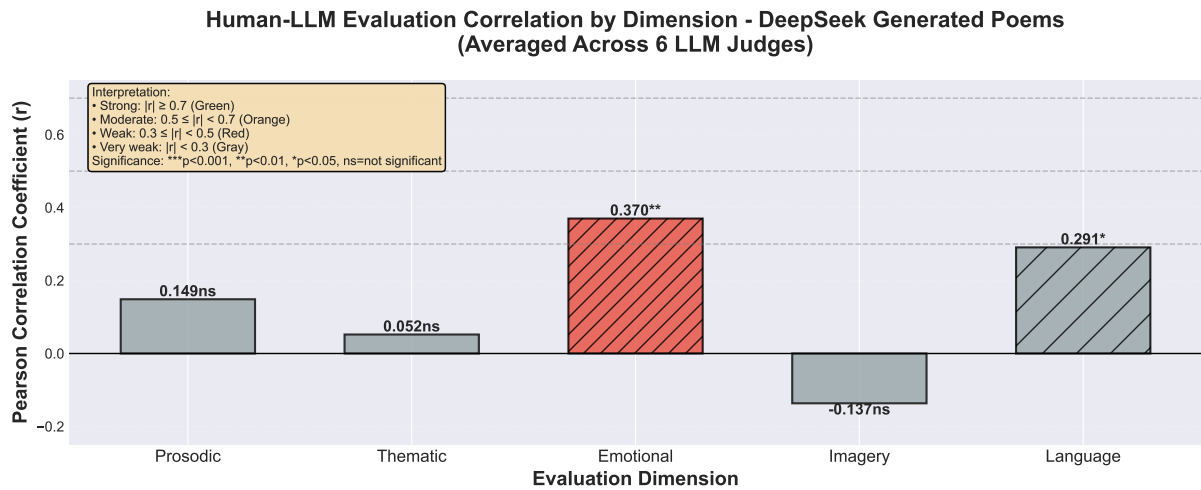


Figure 26: Human vs LLM evaluation correlation by dimension for DeepSeek generated poems (averaged across 6 LLM judges).

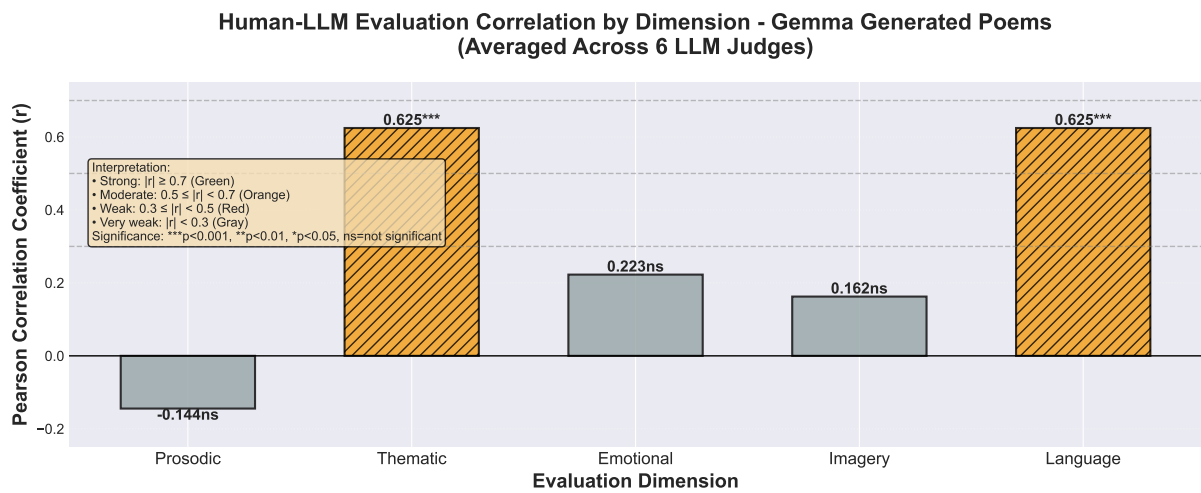


Figure 27: Human vs LLM evaluation correlation by dimension for Gemma generated poems (averaged across 6 LLM judges).

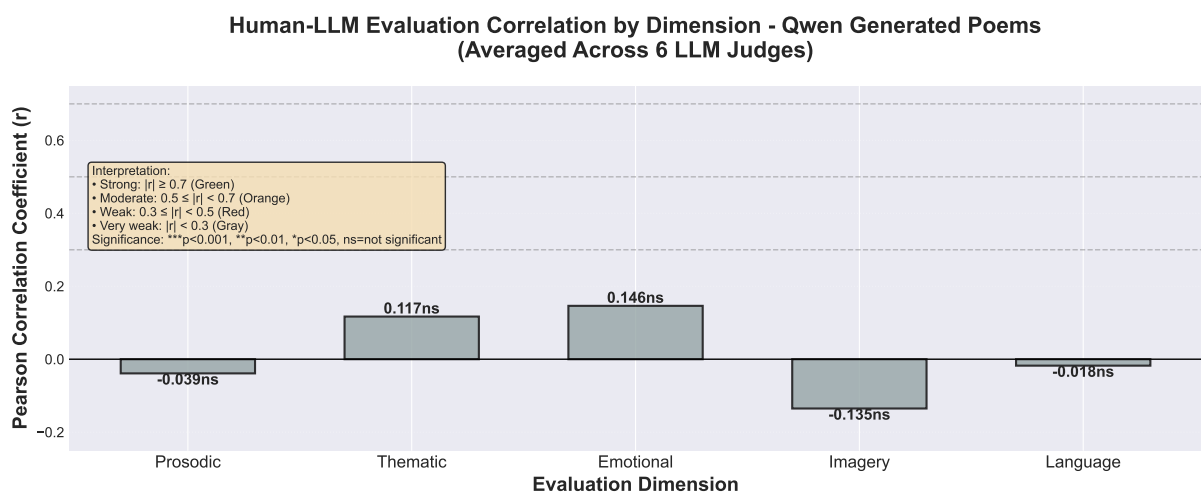


Figure 28: Human vs LLM evaluation correlation by dimension for Qwen generated poems (averaged across 6 LLM judges).

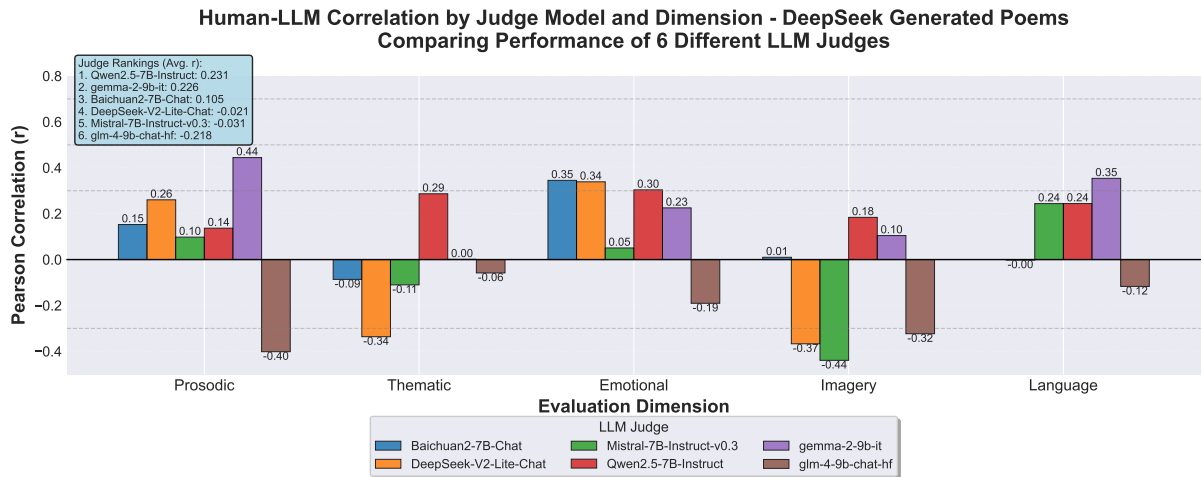


Figure 29: Human-LLM correlation by judge model and dimension for DeepSeek generated poems.

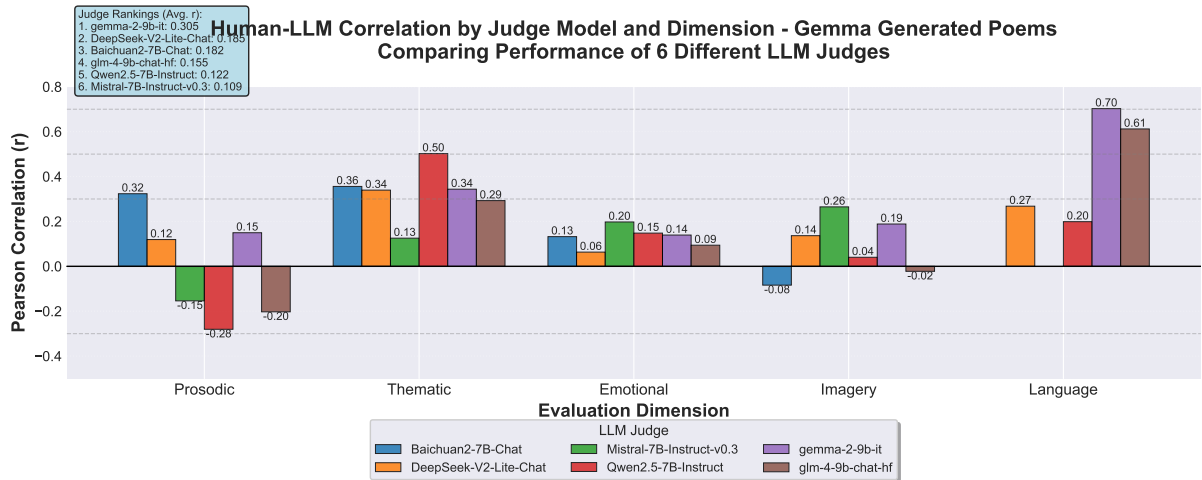


Figure 30: Human-LLM correlation by judge model and dimension for Gemma generated poems.

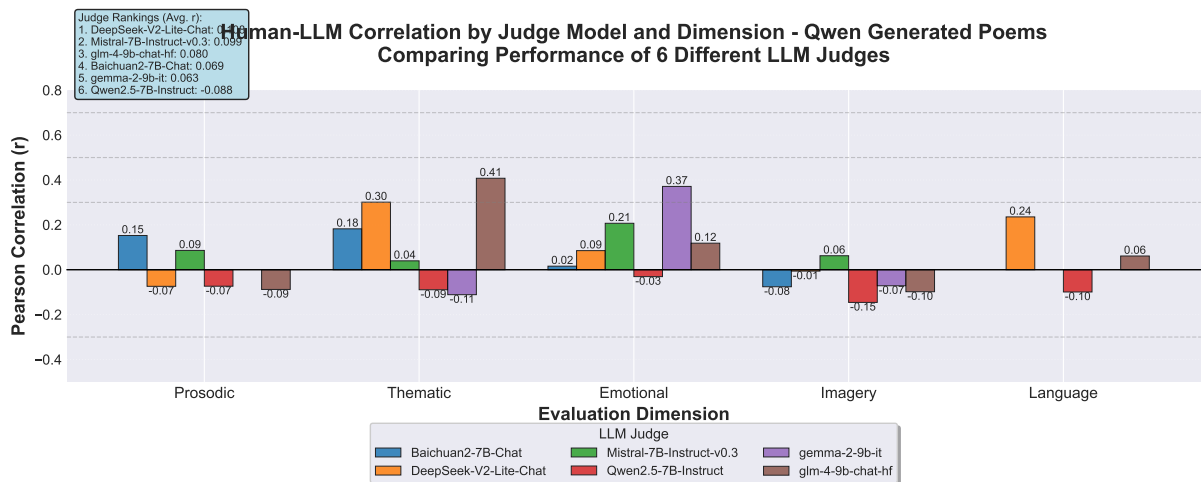


Figure 31: Human-LLM correlation by judge model and dimension for Qwen generated poems.