

Vector Poetics: Parallel Couplet Detection in Classical Chinese Poetry

Maciej Kurzynski

Advanced Institute for Global Chinese Studies
Lingnan University, Hong Kong
maciej.kurzynski@ln.edu.hk

Xiaotong Xu

Lingnan University, Hong Kong
xiaotongxu@ln.hk

Yu Feng

Lingnan University, Hong Kong
yufeng2@ln.hk

Abstract

This paper explores computational approaches for detecting parallelism in classical Chinese poetry, a rhetorical device where two verses mirror each other in syntax, meaning, tone, and rhythm. We experiment with five classification methods: (1) verb position matching, (2) integrated semantic, syntactic, and word-segmentation analysis, (3) difference-based character embeddings, (4) structured examples (inner/outer couplets), and (5) GPT-guided classification. We use a manually annotated dataset, containing 6,125 pentasyllabic couplets, to evaluate performance. The results indicate that parallelism detection poses a significant challenge even for powerful LLMs such as GPT-4o, with the highest F1 score below 0.72. Nevertheless, each method contributes valuable insights into the art of parallelism in Chinese poetry, suggesting a new understanding of parallelism as a verbal expression of principal components in a culturally defined vector space.

1 Introduction

Parallelism is a rhetorical and structural device where two poetic verses correspond with each other in terms of syntax, meaning, tone, and rhythm. The most widely recognized form of parallelism in Chinese culture are the red couplets, known as *fai chun* or *chun lian*, traditionally placed on either side of the doors during Lunar New Year, but the most exquisite ones can be found in classical poetry. Parallelism requires careful selection of words to maintain symmetry while expressing complementary or contrasting ideas, making it a hallmark of linguistic artistry in Chinese poetic craft.

Consider the following, classic example of a pentasyllabic parallel couplet:

山光悦鸟性，
潭影空人心。

*The mountain light pleases the bird's nature;
The pond's reflection empties the human heart.*

Each character (syllable) in the first line corresponds syntactically and semantically with a character at the same position in the second line: 山 (“mountain”) and 潭 (“pond”) both describe natural elements, while 光 (“light”) and 影 (“reflection”) refer to visual phenomena. Similarly, 悦 (“pleases”) and 空 (“empties”) are verbs that describe the effects of these phenomena on 鸟性 (“bird’s nature”) and 人心 (“human heart”), respectively. This correspondence creates a balanced and harmonious structure, parallel in both form and meaning.

While not all Chinese poetry contains parallel structures, one of the most rewarding aspects of reading regulated verses like the one shown above lies in uncovering the intricate correspondences between characters, words, and entire lines. These may be readily apparent, as in our example, or more subtly concealed, drawing upon obscure references and tropes from classical literature that only the most knowledgeable readers will fully appreciate.

In this paper, we report on our attempts to develop a computational method for detecting parallelisms in poetic verses. This complex task involves numerous linguistic dimensions, such as word segmentation, semantic comparison, and part-of-speech alignment, stretching to the limit the abilities of language models and their emergent properties. Yet beyond the technical challenges, computational studies of parallelism can also shed new light on the broader questions of how meaning is generated through language and how humans perceive, structure, and understand the world through patterns and relationships.

2 Previous Work

Modern scholars and traditional *literati* alike have long debated what constitutes a parallel couplet, identifying various categories of parallelisms. In the chapter “Parallel Phrasing” from his *Literary*

Mind and the Carving of Dragons, Liu Xie (460-522) distinguishes four kinds of couplets: verbal (言对), material (事对), antithetical (反对), and direct (正对) (Liu, 2015). A Japanese scholar monk, Kūkai (774-835), proficient in Chinese classical literary criticism, identifies as many as 29 different types of couplets (Kūkai, 1975). Wang Li, a prominent modern Chinese linguist, argues that the classification of words—pairing nouns with nouns, verbs with verbs, etc.—is foundational to understanding parallelism (Wang, 1979). Andrew Plaks views parallelism not merely as a stylistic tool but as a fundamental mode of textual organization and argumentation (Plaks, 1990). Zhang Longxi offers a comparative East-West view on parallelism as a way of reasoning common to all humans as embodied beings (Zhang, 2021). Cai Zong-qi traces the development of parallel patterning in Chinese literature beginning with pre-Qin writings and becoming prominent during the Six Dynasties (Cai, 2022).

From the natural language processing (NLP) perspective, research on parallelism in Chinese poetry has drawn relatively little attention. Lee et al. have applied POS-tagging to detect syntactic parallelism in Tang poetry, confirming traditional views that the middle couplets (2nd and 3rd) in a 4-couplet poem are more often syntactically parallel than the outer couplets (1st and 4th) (Lee et al., 2018). Few other studies focus on the automatic generation of parallel couplets (Yuan et al., 2019; Song, 2022; Qu et al., 2022). Projects like the *Chinese Classical Poetry Matching Dataset (CCPM)* provide structured data that facilitates the study of semantic relationships in poetry, enabling models to better understand the nuances of parallelism by matching ancient Chinese lines with their modern translations (Li et al., 2021). Another prominent approach is the use of transformer-based models, such as the *TransCouplet* model, which leverages a fusion of embeddings, including glyph, pinyin, and part-of-speech embeddings, to capture the syntactic and phonetic nuances of Chinese couplets. The model employs both a transformer encoder and decoder to generate the second line of a couplet given the first, ensuring that the generated line adheres to the linguistic rules of parallelism (Chiang et al., 2021). Yet another innovative method is the *CoupGAN*, which utilizes a generative adversarial network (GAN) for couplet generation. This model focuses on learning the intricate balance between semantic meaning and structural symmetry

that characterizes high-quality couplets (Qu et al., 2022). All of these studies contribute to the ongoing efforts to apply advanced machine learning techniques to the generation and analysis of Chinese parallel couplets, pushing the boundaries of how well machines can replicate this sophisticated form of poetic expression.

3 Dataset

Our manually-annotated test dataset, used to compare different classification methods, contains 6,125 pentasyllabic couplets divided into two classes: parallel (2,139 samples) and non-parallel (3,986 samples). All couplets come from the so-called “Six Dynasties” (222-589), a transformative period in the history of Chinese literature which witnessed the development of parallel poetry, ultimately leading to the elegant regulated verses of the Tang (618-907) and the following eras. The manual annotation ensures a good amount of reliable data to conduct evaluation. Since each method has been using a different training dataset, we provide detailed descriptions of the training process in the relevant subsections.

4 Experiments

4.1 Baseline: Verb Matching

Our baseline method focused on the syntactic positioning of verbs. For training, we have used a manually-annotated dataset of pentasyllabic poetry containing 7,733 verses, with each verse labeled with binary code indicating a verb (1) or non-verb (0) at each position. For example, the line 翩翩吹我衣 (“blowing fluttringly on my clothes”) would be labeled “00100,” with “1” indicating the verb 吹 in the third position. The *SikuBERT* model¹ (109M parameters), pre-trained on a large collection of pre-modern Chinese texts, was then fine-tuned for one epoch on this labeled dataset. We used Adam optimizer, batch size of 16, and learning rate 2e-5. During inference, if verbs were found at corresponding positions in both lines, we considered them as parallel. For example, in the couplet 人来水鸟没, 楫度岸花沈 (“as people arrive, water birds vanish; as oars pass, riverside flowers submerge”), the model correctly identified the verbs that appear at the second and fifth positions in both lines (“01001”).

¹<https://huggingface.co/SikuBERT/sikubert>

4.2 Integrated Semantic and Syntactic Analysis

The second method took a more holistic approach, integrating three distinct models: word segmentation, part-of-speech tagging, and character-to-character semantic matching.

The **word segmentation** model split poetic lines into words, which is essential for Chinese, where word boundaries are not explicitly marked, and where a word can consist of one or more characters. We fine-tuned *SikuBERT* on a dataset of pentasyllabic poetry containing 540k couplets, which had been pre-processed by segmenting each couplet into individual words. To do that, we have built a list of words from the *Hanyu Da Cidian*, or the *Grand Chinese Dictionary*. For example, the couplet 猛虎凭林啸, 玄猿临岸欢 (“Fierce tigers roar from within the forest, mysterious monkeys rejoice by the riverbank”) would be segmented into (first line) 猛虎 “fierce tigers” | 凭 “leaning” | 林 “the forest” | 啸 “roar” and (second line) 玄猿 “mysterious monkeys” | 临 “approaching” | 岸 “the riverbank” | 欢 “rejoice.” We labeled each character as 1 (beginning of the word), 0 (middle or end of the word) or 2 (*BERT* special tokens [CLS] and [SEP]). In the above example, the two lines align perfectly, segmented into 2+1+1+1 structures (or “2101112101112” using our notation). We fine-tuned *SikuBERT* on thus pre-processed couplets (including both parallel and non-parallel ones) for one epoch, with batch size of 16, Adam optimizer, and the learning rate $2e-5$. In inference, we segmented each test couplet and measured the segmentation alignment between the two lines; since the first character is always the beginning of a word, each correct matching after the first character would increase the segmentation score by 0.25.

The **part-of-speech** (POS) tagging model attached tags to each character, providing structural information about the sentence. For example, in the couplet 挥汗辞中宇, 登城临清池 (“Wiping sweat as he leaves the central palace, ascending the ramparts to gaze upon the clear pond”), both lines match perfectly as V+N+V+ADJ+N. For this task, we used a third-party *RoBERTa* model pre-trained on Classical Chinese texts for POS-tagging and dependency-parsing.² Each correct pair-wise matching increased the POS score of a couplet by 0.2.

The **semantic matching** model, finally, com-

pared characters between the two lines to assess semantic relationships. For instance, in the couplet 共矜红颜日, 俱忘白发年 (“Together we cherished the days of youth, equally we forgot the years of grey hair”), all characters in the first line match semantically with those in the second line: adverbs with adverbs, colors with colors, time markers with time markers, etc. Conversely, in 常恐寒节至, 凝气结为霜 (“Always fearing the arrival of cold seasons; the air thickens and turns to frost”), there are no semantically matching characters.

The semantic-matching model required more work than the other two. First, we used *SikuBERT* to generate embeddings (last hidden states, which contain the deep semantic representations) of each character in a couplet. A smaller model, placed on top of the *SikuBERT* encoder, then constructed a rich feature set by concatenating the embeddings of the two corresponding characters as well as their difference. This combined vector (of dimensionality $768 * 3 = 2304$) was passed through a classification layer that predicted whether the character pair matched semantically. By repeating this process across all corresponding character pairs in the couplet, the model generated a final score (the average of pairwise scores) that determined the overall semantic parallelism. The positive training examples were derived from the parallel couplets (the second and third line pairs, or the middle couplets) gathered from the regulated pentasyllabic verses composed in the Tang and the following dynasties, including 270k examples. An equal number of negative examples was generated synthetically by randomly redistributing from 2 to 5 characters within either line of the parallel couplet. For instance, the positive example 青溪通小市, 香稻满平畴 (“A blue stream leads to a small market; Fragrant rice fills the flat fields”), where all characters match semantically (“11111”), would be transformed into a negative example 通青溪小市, 香稻满平畴 (lit. “Leads blue stream small market; Fragrant rice fills the flat fields”), labeled as “00011.” In this case, we shuffled the first three characters in the first line.

Having trained the three models, in inference the scores were then averaged, and if the final score exceeded a predefined threshold (in our case, the best results were achieved with threshold of 75%), the couplet was classified as parallel.

4.3 Difference-Based Character Embeddings

This method focused on the idea of character-to-character differences in embedding space, us-

²KoichiYasuoka/roberta-classical-chinese-large-upos

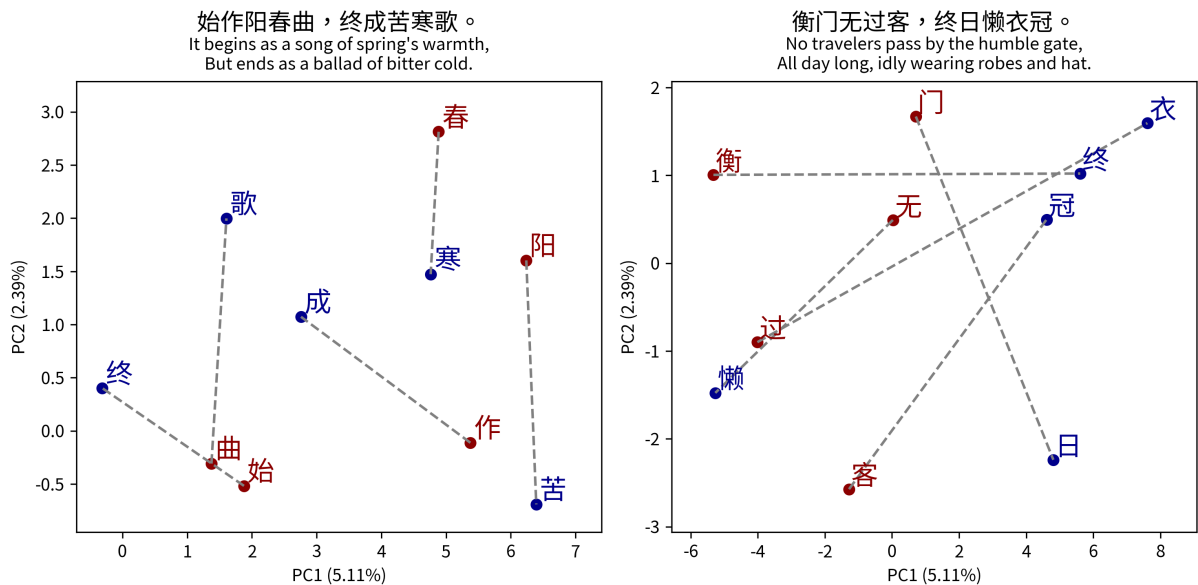


Figure 1: 2-dimensional PCA projections of character embeddings of parallel (left) and non-parallel (right) couplets. Embeddings (last hidden states, $dim = 768$) were generated by *SikuBERT*. We randomly selected 600 couplets to find principal components. Although the first two PCs explain relatively little variance, some parallel relationships can be recognized. For example, in the parallel couplet (left), 阳 (“warmth”) is to 苦 (“bitter”) as 春 (“spring”) is to 寒 (“cold”), and 作 (“to make”) is to 成 (“to become”) as 始 (“to begin”) is to 终 (“to end”).

ing them to distinguish between parallel and non-parallel couplets. For each training couplet, contextual embeddings were generated for all characters with the help of *SikuBERT*. The differences between the contextual embeddings of corresponding characters from the two lines were then calculated and stored in a list as “parallel differences.”

To introduce contrast, the characters in one of the lines were shifted by one or more positions, and the same embedding difference calculation was performed. These results were labeled as “non-parallel differences,” as shifting disrupts the proper alignment necessary for parallelism while preserving semantic connectedness between the lines. A one-layer classifier with RELU non-linearity was then trained for one epoch using both parallel and non-parallel differences (270k examples each).

Once trained, the classifier was applied to our test couplets. If the couplet consisted only of “parallel differences,” we labeled the couplet as parallel. This method aimed to capture the essence of parallelism by focusing on the relative positions and differences between characters in the embedding space (Figure 1).

4.4 Structured Examples

The next method leveraged the structured nature of regulated poetry from the Tang Dynasty onwards,

particularly the *wulü* (five-character eight-line regulated verse). As mentioned before, these poetic forms prescribe parallel lines at specific positions, typically in the second and third couplets. To harness this structural consistency for classification, we collected couplets from publicly available datasets of classical Chinese poetry. Positive examples (270k) were taken from the second and third couplets, while the first and fourth couplets served as negative examples (270k).³ We fine-tuned *SikuBERT* on this dataset for one epoch, with a batch size of 16, learning rate of $2e-5$, and Adam optimizer.

4.5 GPT-Guided Evaluation

The final method leveraged the capabilities of state-of-the-art AI in the form of GPT-4o model (gpt-4o-2024-05-13) accessed over OpenAI API. GPT-4o is a large, multilingual, and multimodal generative transformer model which can be requested to perform various tasks through prompting. This method involved combining a list of calibration examples and a target couplet within a prompt:

Few-Shot Learning: The GPT-4o model was first provided with examples of couplets: in one version, we provided only two examples, one par-

³Some of the first and fourth couplets *are* parallel, a problem explored in the Discussion section below.

allel and one non-parallel, in another as many as ten examples (five parallel and five non-parallel). All examples included a simple analysis and the final judgment based on this analysis. This part of the prompt calibrated the model’s understanding of parallelism in classical Chinese poetry, particularly in terms of grammatical, syntactic, and semantic matching. We provide the 2-shot prompt (original and translation) in the appendix.

Couplet Evaluation: The second part of the prompt included the target couplet. For each couplet, the model was supposed to follow the instructions from the first part of the prompt by providing the couplet’s structural and semantic analysis and then deciding whether it exhibited parallelism. For example, given the lines 驰波催永夜，零露逼短晨 (“Rushing waves hasten the endless night, Scattered dew pressures the brief morning”), the model provided the following explanation (translated here into English, original Chinese below):

“Rushing waves” and “scattered dew” match: both are natural scenes; “hasten” and “pressure” match: both are verbs that imply driving or prompting; “endless night” and “brief morning” also match: one represents a long night and the other a short morning, creating a contrast. Words at each corresponding position in the two lines match semantically, syntactically, and grammatically, thus forming a neatly balanced couplet.

Decision: parallel.

「驰波」与「零露」匹配：都是自然景象。「催」和「逼」匹配：都是动词，表示推动、促使的意思。「永夜」和「短晨」也匹配：一个表示长夜，一个表示短晨，形成对比。两句在每个位置上的词语在语义、句法和语法上都相对应，所以构成工整对仗的对联。

判断：是对联。

5 Results

Despite expectations, the GPT-guided evaluation did not significantly outperform more traditional methods (Table 1). While GPT-4o boasts extensive training data, reinforcement learning from human feedback (RLHF), and a vast parameter count, its performance was comparable to simpler, structured methods. We encountered a number of confusing cases where the final classification seemingly contradicted the preceding analysis, or where the

analysis was hallucinated. As the model is non-deterministic, in some cases GPT-4o generated two contradictory interpretations. The most common errors occurred when the model did not know how to parse a poetic line. Consider the following example of a hallucinated reasoning generated in response to the couplet 淳于前行酒，雍门坐相和 (“Chunyu presents wine as he walks ahead, Yongmen sits and joins in harmony”):

Analysis: “Chunyu” and “Yongmen” are both names of people, matching each other. “Walking ahead” [前行] and “sitting” [坐] express different actions, one walking and the other sitting, forming a neat antithetical movement. “Wine” [酒] and “harmonizing” [相和] represent two scenes related to drinking and music, basically corresponding to each other.

「淳于」和「雍门」都是人名，匹配。「前行」和「坐」表达了不同的动作，一个是行走，一个是坐着，动作对仗工整。「酒」和「相和」表示的是与饮酒和音乐相关的两个场景，基本对应。

While the model correctly identified the proper names, it considered the two verses as parallel despite the syntactic mismatch in the last three characters, comparing words of different lengths. We believe that further prompt engineering could mitigate such errors and improve performance.

The verb-based method, on the other hand, struggled with lines where verbs were either absent or not central to the couplet’s structure, leading to misclassifications. Consider the couplet 春草郁青青，桑柘何奕奕 (“spring grass flourishes, how lush and vibrant; mulberry and elm trees, how splendid and bright”), which features no verbs (“00000”) but still forms a parallel couplet. Moreover, even when verbs were present at identical positions, semantic mismatches could occur, such as in 边地无高木，萧萧多白杨 (“There are no tall trees in the borderland; there are many rustling poplars”), where the verbs align, but the initial disyllabic words (边地 “borderland” and 萧萧 “rustling”) do not match semantically.

The integrated approach remained heavily dependent on the availability and quality of training data, and the compounded errors from individual models led to numerous misclassifications. For example, the word-segmentation model alone could not account for semantic mismatches, as seen in 误落尘网中，一去三十年 (“Mistakenly caught in the

Table 1: Results of Classification Methods. The methods include Verb Matching (VM), which focuses on the syntactic positioning of verbs; Semantic-Syntactic Integration (ISSA), which combines semantic matching with syntactic analysis (POS + word segmentation); Embedding Differences (ED), which utilizes character embedding differences; Structured Examples (SE), which distinguishes between 2nd and 3rd (mostly parallel) couplets and 1st and 4th (mostly non-parallel) couplets from regulated poetry; and GPT Evaluation (GPT-4o), which leverages a large conversational transformer model.

Method	Accuracy (%)	Precision (%)	Recall (%)	F1 Score (%)
VM (baseline)	76.5	72.3	55.0	62.5
ISSA	72.5	68.2	48.0	56.3
ED	75.9	72.0	53.6	61.4
SE	79.1	69.9	72.7	71.3
GPT-4o (2-shot)	79.0	70.7	68.4	69.5
GPT-4o (10-shot)	79.6	73.7	64.4	68.8

web of the mundane, Gone for thirty years”), where the word segmentation is consistent (1+1+2+1), but semantics differ. Likewise, the part-of-speech model would (wrongly) reduce the couplet score in cases where non-matching individual characters formed larger matching units: in the couplet 君居北海阳，妾在江南阴 (“You reside in the north of the North Sea; I dwell in the south of the Yangtze”), the POS tags of individual characters differ in 北海 (North Sea, labeled as adjective+noun) and 江南 (Yangtze’s South, labeled as noun+adjective), even though the two words match semantically.

Finally, although the structured examples approach should have benefited from the fact that the regulated poetry from the Tang onwards contains parallel lines at prescribed positions, we have found many exceptions in the existing datasets, in particular verses from less known poets that do not conform with the stringent formal rules, such as the inner couplet 隔邻呼酒伴，对坐且开尊 (“Calling a drinking companion from next door; They sit facing each other, opening a jar of wine”). Similar examples might have confused the model during training, leading to lower score. It is worth noting that despite such shortcomings, this approach achieved higher F1 score than GPT-4o, and further filtering of the training data might improve the results. On the other hand, it offers less explainability than GPT.

6 Discussion

6.1 Parallelism and Subjectivity

In our evaluation of parallel couplets, it became evident that the notion of parallelism is not always binary—many couplets do not simply fall into “parallel” or “non-parallel” categories. Instead, paral-

lelism often manifests on a spectrum. We observed instances of “almost parallel” couplets where, despite imperfect POS matching, the semantic correspondence between two lines suggested itself: 昔闻汾水游，今见尘外骖 (“Once I heard of travels by the Fen River, Now I see a horse beyond the dusty realms”).

Likewise, the subjective nature of interpreting poetic texts posed a significant challenge in our study. A creative reader can find parallel relationships between any two randomly chosen verses. We frequently encountered disagreements among annotators regarding the classification of certain couplets. These disagreements highlight the inherent interpretive flexibility in literary analysis, where different readers may perceive varying degrees of parallelism based on their linguistic intuition, cultural background, and personal experience.

6.2 Vector Poetics

The concept of “vector poetics” represents a promising avenue explored in this project, drawing inspiration from foundational work in natural language processing. Seminal studies of word embeddings, such as word2vec (Mikolov et al., 2013) or GloVe (Pennington et al., 2014) papers, demonstrate how parallel relationships between words in a given dataset can be quantitatively modeled (“queen” is to “king” as “woman” is to “man,” e.g.). We suggest that this analogy can be usefully extended to sets of words—or entire poetic lines—in parallel couplets (Figure 1). Consider the transformation:

$$\text{春} - \text{生} + \text{死} \approx \text{秋} \quad (1)$$

Translating as “Spring – Life + Death \approx Autumn,”

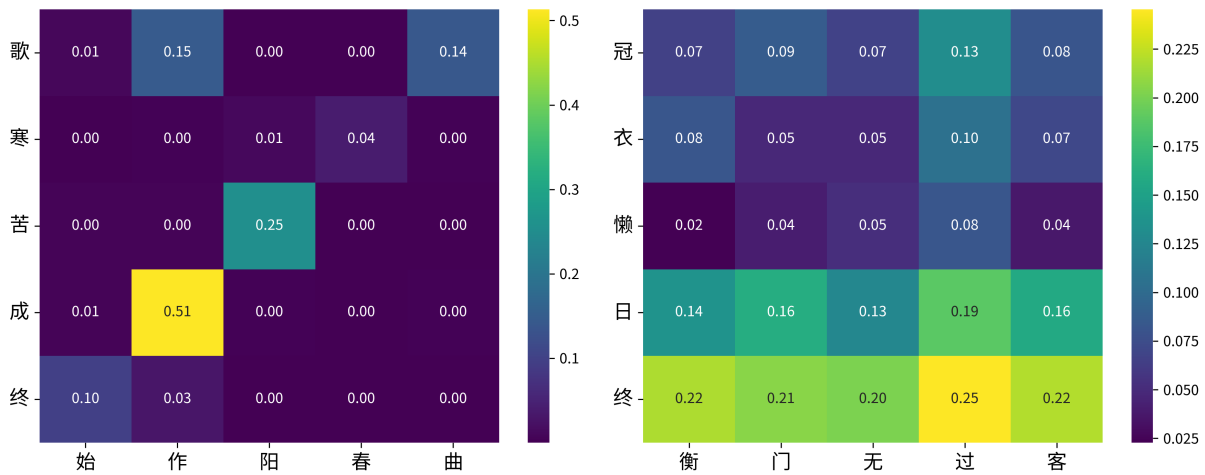


Figure 2: Attention heatmaps comparing parallel and nonparallel couplets (examples identical with Figure 1). Left: Heatmap showing focused attention along the diagonal in a parallel couplet, indicative of strong structural and semantic alignment learned by the model from parallel examples. Right: Heatmap for a non-parallel couplet with diffused attention, highlighting the lack of structured alignment between lines. Attention scores have been averaged (both directions) over all heads from the top layer of *SikuBERT*, which has been fine-tuned on parallel and non-parallel couplets (Structured Examples approach).

the equation conceptualizes the cyclical nature discovered by poets in both seasons and life. Importantly, such equivalences will hold only in cultural contexts where “Spring” is to “Life” as “Autumn” is to “Death,” that is, where these concepts align along the same “principal cultural components.” Vector poetics would thus re-interpret poetic parallelisms as manifestations of cultural dimensions in which two things are (subjectively) perceived as either identical or opposite. This also entails a quantitative re-interpretation of poetic creativity: a search within a historically and culturally defined vector space, and a reshaping of that space through discovery of new similarities and differences between things in the world.

Finally, our analysis points to the innovative use of attention mechanisms within transformer-based models, which can visualize the interaction between poetic lines (Figure 2). Heatmaps generated from the attention scores in a fine-tuned *SikuBERT* reveal distinct patterns: in parallel couplets, there is a noticeable concentration of attention along the diagonal, indicating a structured and direct interaction between corresponding characters. Conversely, in non-parallel couplets, attention is more diffusely spread across the matrix, indicating a weaker alignment and less structured interaction between the lines. Notice that these specific attention patterns have been inferred by the model during training. While more work needs to be done in this direction, currently available results seem to offer further evi-

dence that parallelism is not only a formal device at work within a text but also a cognitive device that guides the reader’s interpretive engagement.

Limitations

Despite the advancements detailed in this study, several limitations remain. Traditional Chinese poetry, particularly from the classical period, adheres to strict patterns of alternating tones, known as *pingze* (平仄), with *ping* (平) tones being level and *ze* (仄) tones oblique. These tonal patterns, crucial for the aesthetic and rhythmic quality of Chinese poetry, add yet another layer of complexity to parallelism detection but are not accounted for in our current computational models, which generally lack the capability to analyze tonal features (Liu et al., 2018). Additionally, our approach overlooks the ancient pronunciation of Chinese characters, which is vital for understanding the original phonetic environment and rhyming schemes of the texts (Tharsen, 2022). This oversight may lead to potential misinterpretations or missed phonetic parallelisms that would have been evident to contemporary audiences. Finally, the pre-trained *SikuBERT* model used in our study is relatively small by contemporary standards; larger models are likely to yield better performance. We are going to address these shortcomings in our future studies.

Acknowledgments

The authors would like to thank the NLP4DH 2024 conference organizers and the two anonymous reviewers for their useful comments and feedback.

References

- Zongqi Cai. 2022. Six dynasties parallel prose: Descriptive and expository. In Zong qi Cai, editor, *How to Read Chinese Prose*, pages 210–225. Columbia University Press, New York.
- Kuan-Yu Chiang, Shihao Lin, Joe Chen, Qian Yin, and Qizhen Jin. 2021. [Transcouplet: Transformer based chinese couplet generation](#). *arXiv preprint arXiv:2112.01707*.
- Kūkai. 1975. *Bunkyo Hifuron*. People’s Literature Publishing House.
- John Lee, Yin Hei Kong, and Mengqi Luo. 2018. [Syntactic patterns in classical chinese poems: A quantitative study](#). *Digital Scholarship in the Humanities*, 33(1):82–95.
- Wenhao Li, Fanchao Qi, Maosong Sun, Xiaoyuan Yi, and Jiarui Zhang. 2021. [CCPM: A Chinese Classical Poetry Matching Dataset](#). *arXiv preprint arXiv:2106.01979*. Accessed: 2024-08-19.
- Chao-Lin Liu, Thomas J Mazanec, and Jeffrey R Tharsen. 2018. Exploring chinese poetry with digital assistance: Examples from linguistic, literary, and historical viewpoints. *Journal of Chinese Literature and Culture*, 5(2):276–321.
- Xie Liu. 2015. *The Literary Mind and the Carving of Dragons*. The Chinese University Press.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, pages 3111–3119.
- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543. Association for Computational Linguistics.
- Andrew H. Plaks. 1990. [Where the lines meet: Parallelism in chinese and western literatures](#). *Poetics Today*, 11(1):105–127.
- Qian Qu, Jiancheng Lv, Dayiheng Liu, and Kexin Yang. 2022. [Coupgan: Chinese couplet generation via encoder–decoder model and adversarial training under global control](#). *Soft Comput.*, 26(15):7423–7433.
- Yan Song. 2022. [Chinese couplet generation with syntactic information](#). In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 6436–6446, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.
- Jeffrey R. Tharsen. 2022. [From form to sound: visual and aural representations of premodern chinese phonology and phonorhetoric with applications for phonetic scripts](#). *International Journal of Digital Humanities*, 4(1-3):115–129.
- Li Wang. 1979. *A Brief Introduction to the Metric Patterns of Chinese Poetry*. Beijing Publishing House.
- Shengqiong Yuan, Luo Zhong, Lin Li, and Rui Zhang. 2019. [Automatic generation of chinese couplets with attention based encoder-decoder model](#). *2019 IEEE Conference on Multimedia Information Processing and Retrieval (MIPR)*, pages 65–70.
- Longxi Zhang. 2021. Parallelism and antithesis: Structural principles in the mind and in literature from a chinese perspective. *European Review*, 29(2):274–284.

A Appendix (2-shot GPT-4o prompt)

[system] You are a professor of Chinese literature and an expert in ancient poetics.

[user] Please determine if the following two five-character lines form a parallel couplet. The requirement for a parallel couplet is that each word in the first line corresponds syntactically, grammatically, and semantically with a word in the second line (there are many methods of correspondence: the two words can respectively represent time and space, morning and evening, life and death, youth and harsh winter, static and dynamic, two different places, place nouns and time nouns, etc.). You need to provide a concise analysis, then make a decision based on this analysis. The decision can only be ‘is parallel’ or ‘is not parallel.’

[user] Couplet: Observing things arouses sorrow, looking after oneself recognizes the source of worries.

[assistant] Analysis: “Observing things” corresponds to “looking after oneself”: one is about observing the outside world, the other is about caring for one’s inner world (both are one-character verb plus one-character noun structures), “arouses” and “recognizes” are both verbs, “sorrow” refers to a sad mood, “source of worries” refers to a sorrowful mood, both match. Decision: Is parallel.

[user] Couplet: Sitting quietly all around, I listen to the song above my hall.

[assistant] Analysis: “Four” and “listen” do not match (one is a number, the other is a verb), “sit” and “I” do not match (one is a verb, the other is a pronoun), “and” and “hall” do not match (one is a

function word, the other is a noun). Decision: Is not parallel.

[system] 你是中国文学的教授和古代诗学的专家。

[user] 请判断下面的两个五言句是不是工整对仗的对联。对联的要求是前句和后句的每一个词语在句法、语法和意义上都相对应（对应的方法有很多：两个词可以分别表达时间与空间、早上与晚上、生与死、春天与冬天、静态与动态、两个不同的地方、地点名词与时间名词等等）。请提供简洁的分析，然后根据这个分析进行判断。判断只能是「是对联」或者「非对联」。

[user] 句对：览物起悲绪，顾己识忧端。

[assistant] 分析：「览物」对应着「顾己」：一个是指观赏外面的世界，一个是指照顾内心的世界（两个都是一个字的动词加一个字的名词），「起」与「识」都是动词，「悲绪」是指悲伤的心情，「忧端」是指忧伤的心情，都匹配。判断：是对联。

[user] 句对：四坐且莫喧，听我堂上歌。

[assistant] 分析：「四」与「听」不匹配（一个是数字，一个是动词），「坐」与「我」不匹配（一个是动词，一个是代名词），「且」和「堂」不匹配（一个是功能词，一个是名词）。判断：非对联。