

Analyzing Sentiment in Classical Chinese Poetry

Yufang Hou Anette Frank

Institute for Computational Linguistics, Heidelberg University, Germany

(hou|frank)@cl.uni-heidelberg.de

Abstract

Although sentiment analysis in Chinese social media has attracted a lot of interest in recent years, it has been less explored in traditional Chinese literature (e.g., classical Chinese poetry) due to the lack of sentiment lexicon resources. In this paper, we propose a weakly supervised approach based on Weighted Personalized PageRank (WPPR) to create a sentiment lexicon for classical Chinese poetry. We evaluate our lexicon intrinsically and extrinsically. We show that our graph-based approach outperforms a previous well-known PMI-based approach (Turney and Littman, 2003) on both evaluation settings. On the basis of our sentiment lexicon, we analyze sentiment in the Complete Anthology of Tang Poetry. We extract topics associated with positive (negative) sentiment using a position-aware sentiment-topic model. We further compare sentiment among different poets in Tang Dynasty (AD 618 – 907).

1 Introduction

Classical Chinese poetry is a precious cultural heritage. Among its over 3,000 years of history, the Tang Dynasty (AD 618 – 907) is widely viewed as the zenith of the art of classical Chinese poetry. The Complete Anthology of Tang Poetry, edited during the Qing Dynasty (1644 – 1911), contains over 42,860 poems in 900 volumes by more than 2,500 poets. The collection provides a magnificent insight into all aspects of social life of that period.

Research on sentiment/emotion and imagery analysis of Tang poetry is an active subfield in Chinese philology, with a vast literature (Watson, 1971; Kao and Mei, 1971; Kao and Mei, 1978). In this paper, we seek to analyze the sentiment (i.e.,

positive or negative) of textual elements in Tang poetry from a computational perspective. Specifically, we propose a novel graph-based method to create a sentiment lexicon for classical Chinese poetry. Such a lexicon is a valuable resource for other computational research on classical Chinese poetry, such as semantic analysis (Lee and Taksum, 2012) or poetry generation (He et al., 2012; Zhang and Lapata, 2014).

Turney and Littman (2003) propose a PMI-based algorithm to estimate the *semantic orientation* or *polarity* of a word. The semantic orientation of a given word is calculated by comparing its similarity to positive reference words (e.g., excellent or beautiful) with its similarity to negative reference words (e.g., poor or bad). Instead of calculating the similarity between a given word and each of the positive (negative) reference words separately, we apply Weighted Personalized PageRank (WPPR) to measure the similarity between the given word and all positive (negative) reference words **simultaneously** in a lexical network that we build from a poetry corpus. Our graph-based method is able to find globally optimal solution because the lexical network is analyzed as a whole (Section 3).

We evaluate our poetry sentiment lexicon intrinsically and extrinsically. For the intrinsic evaluation, we compile two test datasets. The first dataset contains 933 words (532 positive and 401 negative) taken from three Chinese sentiment lexicons¹. The second dataset contains 55 words taken from literature of imagery analysis for Tang poetry. These words reflect the common imageries in classical Chinese poetry and have certain fixed emotional connotations. For instance, the character “猿” (ape) often relates to sadness, anxiety and distress, while the character “荷” (lotus)

¹Although these lexicons are for contemporary Chinese, some words keep the same meaning and polarity as in classical Chinese poetry.

is the symbol of beauty, love and rectitude. We show that our method outperforms the very competitive PMI-based approach when evaluating on both datasets (Section 4.1). Our method also outperforms the baseline on an extrinsic evaluation task of predicting sentiment orientation of classical Chinese poetry (Section 4.2).

On the basis of our sentiment lexicon, we analyze sentiment in the Complete Anthology of Tang Poetry. We first analyze topic distributions under positive/negative sentiment in Tang poetry using a position-aware sentiment-topic model (Section 5.1). We then compare sentiment among different poets in Tang Dynasty (Section 5.2).

The main contributions of our work are:

- We propose a graph-based method to build a sentiment lexicon for classical Chinese poetry. Our method is weakly supervised and does not rely on existing lexical resources (e.g., WordNet). It can be easily ported to other domains/languages.
- We evaluate our sentiment lexicon systematically and demonstrate that it can be utilized to analyze sentiment orientation of classical Chinese poetry.
- We analyze sentiment in Tang poetry on the basis of our sentiment lexicon. We apply a *position-aware sentiment-topic model* to extract themes which are tightly associated with positive/negative sentiment. Our model builds in specific assumptions that characterize sentiment expression in classical Chinese poetry. It assumes that lexical items from the same region are generated from a single sentiment-topic pair. We compare sentiment among different famous poets and show that our results are in accordance with studies in Chinese philology.

The poetry sentiment lexicon described in the paper as well as all test datasets are freely available at <http://www.cl.uni-heidelberg.de/~hou/resources.mhtml>.

2 Related Work

Sentiment lexicons. In recent years, considerable attention has been given to the creation of large polarity (positive and negative) lexicons, including various corpus-based approaches (Turney and Littman, 2003; Kanayama and Nasukawa,

2006; Kaji and Kitsuregawa, 2007; Kiritchenko et al., 2014) and dictionary-based approaches (Kamps et al., 2004; Esuli and Sebastiani, 2005; Mohammad et al., 2009; Baccianella et al., 2010). Unlike previous graph-based approaches which create sentiment lexicons based on existing lexical resources (e.g., WordNet, thesauri) (Takamura et al., 2005; Rao and Avichandran, 2009; Hassan et al., 2011), there are no such lexical resources for classical Chinese poetry. Therefore, we choose a corpus-based approach.

While our approach for building sentiment lexicons is domain independent, in this paper we apply it to classical Chinese poetry. This is not a trivial task. There are a variety of reliable resources for English sentiment analysis. However, only a few sentiment lexicons for Chinese are available. In particular, these lexicons are for contemporary Chinese. Moreover, given that these lexicons are developed for contemporary Chinese, they will only have partial coverage for classical Chinese poetry. There might also be divergences due to the change of language over several thousand years. To improve sentiment analysis for Chinese, one line of work seeks to leverage rich English sentiment resources through machine translation (Wan, 2008; Wan, 2009; He et al., 2010). These approaches depend on the quality of machine translation and translation of classical Chinese poetry to English is hard even for professional translators. Our work is similar to Zagibalov and Carroll (2008) in the sense that both approaches are weakly supervised. They build a sentiment lexicon iteratively, starting from a small set of seed items and several lexical patterns (negated adverbial constructions) which can indicate lexical polarity. However, such lexical patterns (e.g., 不 (not) 很 (quite) + 满意 (satisfied) (target word)) are not applicable in classical Chinese poetry.

Computational analysis of classical Chinese poetry. There has been previous work focusing on classical Chinese poetry generation (Zhou et al., 2010; He et al., 2012; Zhang and Lapata, 2014). Lee and Kong (2012) develop a dependency treebank for the Complete Anthology of Tang Poetry. On the basis of this corpus, Lee and Tak-sum (2012) quantitatively analyze the semantic content and word usage in the Complete Anthology of Tang Poetry. Voigt and Jurafsky (2013) find that the classical characters of Chinese poetry decreased across the century by comparing classical

poetry and contemporary prose.

There are only a few works trying to analyze sentiment in classical Chinese poetry. Hu (2001) proposes “similarity search” by using word association measures. For instance, given typical emotional words such as “悲伤 (sadness) 哀 (sorrow)”, the system can find words (e.g., 南浦 (southern shore, a place often used to hold farewell parties in ancient China)) associated with sad emotions. However, he does not analyze sentiment in classical Chinese poetry quantitatively. Based on manually annotated data, Luo (2009) analyzes the sentiment of classical Chinese Song poetry among different poets. To the best of our knowledge, there is no publicly available sentiment lexicon for classical Chinese poetry.

3 Building a Sentiment Lexicon for Classical Chinese Poetry

In this section, we briefly introduce Weighted Personalized PageRank (WPPR). We then detail how we construct a lexical network and how we apply WPPR over the lexical network to build a sentiment lexicon for classical Chinese poetry.

3.1 Weighted Personalized PageRank

The original PageRank algorithm was first introduced by Brin and Page (1998). It is a link-based algorithm for ranking the vertices in a graph. Later, various extensions have been proposed. Weighted PageRank (Xing and Ghorbani, 2004) takes into account the importance of both the in-links and the outlinks of the vertices when distributing rank scores based on the popularity of the vertices. Personalized PageRank (Haveliwala, 2002; White and Smyth, 2003) computes the importance of vertices in a graph relative to one or more root vertices. It has been successfully applied in other NLP tasks, such as word sense disambiguation (Agirre and Soroa, 2009).

Here we combine Weighted PageRank and Personalized PageRank to measure the similarity of lexical items in a lexical network relative to sentiment seeds. Let G be a lexical network with N vertices $v_1, \dots, v_n \in V$ and w_{ij} be the weight associated with the edge from v_i to v_j . Let M be a $N \times N$ transition probability matrix, where $M_{ij} = w_{ij}$ if a link from v_i to v_j exists, and zero otherwise, let S be a set of sentiment seeds where $S \subseteq V$. Then the Weighted Personalized PageR-

ank vector R over G can be calculated as follows:

$$R = \alpha MR + (1 - \alpha)P, \quad (1)$$

where α is the *damping factor* and its value usually set in the $[0.85..0.95]$ range. P is a $N \times 1$ vector, where $P_i = \frac{1}{|S|}$ for $v_i \in S$, and zero otherwise, i.e., all vertices in the sentiment seeds have equal prior probability.

Equation 1 can be viewed as the result of a random walk process starting from the seed nodes, where the random walkers can jump back to the seed nodes S with a given probability $1 - \alpha$. The final rank of vertex v_i , biased towards the set S (the bias is encoded in P), represents the probability of a random walk over the weighted graph (weights associated with edges are encoded in M) ending on vertex v_i , at a sufficiently large time.

3.2 Lexical Network Construction

To create a sentiment lexicon for classical Chinese poetry, we first build a lexical network on the basis of the Complete Anthology of Tang Poetry². Since poetry is imbued with emotions, we assume that: (1) each lexical items in the lexical network bears positive or negative sentiment; and (2) lexical items within a small window are more likely to share the same sentiment. Therefore, by applying WPPR on the basis of a small set of positive (negative) lexical items, we can trace how positive (negative) sentiment information is distributed over the whole lexical network.

The lexical network G is a directed weighted graph, where each vertex v_i is a lexical item. We define a lexical item as a word containing one or two characters. Classical Chinese poetry is typically written in a highly compressed style, where each line normally has a fixed five or seven characters. As a result, each character itself or words containing two characters are expressive and can be used as the main semantic units. Instead of carrying out word segmentation, we simply use a frequency threshold to extract lexical items: a lexical unit is extracted as a lexical item if it appears at least x times in the corpus (x is ten for single-character unit and 50 for two-character unit).

We then create an edge from v_i to v_j , if v_i and v_j co-occur within a window of five characters, i.e., v_i occurs within a window of five characters before or after v_j . Let f_{ij} be the number of times

²The corpus can be downloaded from <http://datatang.com>

that v_i and v_j co-occur in the whole corpus, we set the weight of the edge from v_i to v_j as follows:

$$w_{ij} = \frac{f_{ij}}{\sum_{k=1}^N f_{ik}} \quad (2)$$

Alternatively, w_{ij} can be viewed as the probability of lexical item v_j occurring nearby, given the lexical item v_i .

As a result, we construct a lexical network containing 8656 lexical items (4779 are single-character items, 3877 are two-character items) and 8,832,234 edges. This lexical network contains the word co-occurrence information in the Complete Anthology of Tang Poetry.

3.3 Sentiment Lexicon Creation

We compile a small set of sentiment seeds, which contains six positive lexical items and six negative lexical items (see Table 1). These lexical items are frequent single characters in the Complete Anthology of Tang Poetry and carry strong sentiment. Similar to Turney and Littman (2003) who use 14 sentiment seeds (seven positive words and seven negative words), we only focus on a small number of sentiment seeds to study whether we can build a reasonable sentiment lexicon from weak supervision.

	Characters
positive seeds	香 (fragrant) 爱 (love) 欢 (happy) 贤 (virtuous) 喜 (delight) 瑞 (lucky)
negative seeds	寒 (cold) 愁 (anxiety) 孤 (lonely) 苦 (painful) 悲 (sorrow) 怨 (resentment)

Table 1: Positive and negative sentiment seeds.

We apply WPPR (Section 3.1) twice over the lexical network described in Section 3.2, initialized with the positive seeds and negative seeds respectively. We follow the common practice of setting the *damping factor* to 0.85. Consequently, we get two PageRank vectors Rp and Rn . They can be seen as a measure of similarity of lexical items to all positive seeds and all negative seeds respectively. Finally, we calculate the sentiment vector as follows:

$$Rs = Rp - Rn \quad (3)$$

A lexical item i has a positive sentiment orientation if its corresponding entry in vector Rs (hence Rs_i) is positive, and a negative sentiment orientation if Rs_i is negative. The value of Rs_i can be viewed as the strength of the sentiment orientation associated with the lexical item i .

4 Sentiment Lexicon Evaluation

We evaluate our poetry sentiment lexicon intrinsically and extrinsically. For the intrinsic evaluation, we utilize sentiment lexicons for contemporary Chinese because there is a partial overlap between these lexicons and sentiment expressions in classical poetry. We also evaluate lexical items in our sentiment lexicon appearing only in classical poetry. In the extrinsic evaluation, we test whether our sentiment lexicon can be used to predict sentiment orientation of classical Chinese poetry.

4.1 Intrinsic Evaluation

Test Datasets. To evaluate our approach, we compile two test datasets. The first dataset (SentiLexicon) contains 933 sentiment words taken from three Chinese sentiment lexicons: HowNet³, NTUSD (Ku et al., 2006), and Tsinghua sentiment lexicon⁴. Although these lexicons are for contemporary Chinese, some words keep the same meaning and polarity as in classical Chinese poetry. We merge these three lexicons by removing duplicate or contradictory entries. This yields a big sentiment lexicon containing 12,945 positive words⁵ and 17,114 negative words. We then create SentiLexicon by choosing single-character words and two-character words from the big sentiment lexicon if they do not appear in the set of sentiment seeds (Table 1) and occur at least 50 times⁶ in the Complete Anthology of Tang Poetry. This leads to a dataset containing 532 positive lexical items and 401 negative lexical items.

However, SentiLexicon does not reflect an important aspect of classical Chinese poetry, i.e., emotions are expressed implicitly through imagery. Skilled poets often apply concrete imagery to evoke emotions and sensations. Certain imageries have fixed emotional connotations. For example, the falling autumn leaf (“落叶”) often refers to personal or dynastic decline. We call such words imagery words. We collect 55 typical imagery words (ImageryLexicon) from literature of imagery analysis for Tang poetry. Every word in ImageryLexicon does not appear in SentiLexicon. Table 2 shows some examples of ImageryLexicon.

³http://www.keenage.com/html/e_index.html

⁴<http://nlp.csai.tsinghua.edu.cn/~lj/>

⁵A word can contain one character or several characters.

⁶We carried out preliminary experiments with the thresholds ranging from ten to 50. We found that the accuracy of our method varies in a small range, and our approach outperforms the baseline at all threshold levels.

	Characters
positive	鸳鸯 (mandarin duck) 芙蓉 (hibiscus) 凤凰 (phoenix) 兰 (orchid) 竹 (bamboo)
negative	梧桐 (sycamore) 鸦 (crow) 柳 (willow) 鹧鸪 (partridge) 夕阳 (sunset) 子规 (cuckoo)

Table 2: Examples of ImageryLexicon.

	SentiLexicon	ImageryLexicon
	Accuracy	Accuracy
<i>PMI</i>	60.8	70.2
<i>WPPR</i>	64.4*	74.5

Table 3: Accuracy of WPPR compared to PMI (baseline) for two test datasets. * indicates significant improvement relative to the baseline (McNemar’s test at $p < 0.05$ level).

Baseline. We reimplement a previous PMI-based approach (Turney and Littman, 2003) as the baseline. We use the same sentiment seeds and the same co-occurrence window of five characters as our method. The sentiment orientation of a lexical item (single or two-character) is calculated as follows:

$$SO(w) = \sum_{s \in pSeeds} PMI(w, s) - \sum_{s \in nSeeds} PMI(w, s) \quad (4)$$

Results on test datasets. Table 3 shows the results of our method described in Section 3 (*WPPR*) and the baseline (*PMI*) against two test datasets. Our graph-based approach outperforms the baseline in both cases. Our method is more robust than the baseline because it measures the similarity between the candidate lexical item and the whole positive (negative) sentiment seeds together.

Evaluation on sample data. Our test datasets (SentiLexicon and ImageryLexicon) only cover about 11.5% of lexical items of our sentiment lexicon. To evaluate the lexical items that are not in the test sets, we randomly choose 100 items (50 single and 50 two-character lexical items, both with the equal positive/negative sentiment distribution). They were manually checked by the first author. We obtain an accuracy of 53% in this hard evaluation setting.

4.2 Extrinsic Evaluation

We also carry out an extrinsic evaluation to judge whether our sentiment lexicon can be utilized to analyze sentiment orientation of classical Chinese

poetry. We choose 160 poems from the Tang poetry analysis dictionary (Xiao, 1999), which contains around 1,000 Tang poems paired with professional reviews. We manually annotate the sentiment of each poem as positive or negative according to the reviewers’ analysis. This leads to a dataset (sentiPoetry) containing 83 negative poems and 77 positive poems. For each poem, we predict its sentiment based on the accumulated sentiment orientations of all lexical items (single and two-character) in the poem. Specifically, a poem is predicted as positive if its accumulated sentiment orientation is bigger than a threshold t , and negative otherwise. A subset of sentiPoetry containing 30 positive poems and 30 negative poems is used to tune the threshold t , the remaining 100 poems are reserved as test data. Table 4 shows the accuracy of predicting poetry sentiment on the test dataset using the sentiment lexicon for contemporary Chinese described in Section 4.1, as well as the two lexicons based on the baseline (*PMI*) and our method (*WPPR*) respectively. Using our lexicon achieves an accuracy of 71% on predicting poetry sentiment, which is 14% better than using *PMI Lexicon*. It is obvious that the Out-of-Domain lexicon (*contemporarySenti Lexicon*) performs the worst because of its low coverage of lexical items used in classical Chinese poetry. A closer look at the results indicates that positive poems are hard to predict because happy/joyful emotions are often expressed in a very subtle, implicit way.

	Accuracy
<i>contemporarySenti Lexicon</i>	51.0
<i>PMI Lexicon</i>	57.0
<i>WPPR Lexicon</i>	71.0

Table 4: Results for poetry sentiment prediction.

5 Analyzing Sentiment in Tang Poetry

Poems are saturated with emotions that correlate to positive or negative sentiment. But how are sentiments expressed in different topics? How does sentiment differ between individual poets? We aim to answer these questions in this section.

5.1 Sentiment-based Topic Distribution

Position-aware sentiment-topic model. Traditional topic models like latent Dirichlet allocation (LDA) (Blei et al., 2003) have been explored extensively to discover topics from text. Recently,

LDA has been extended to capture correlations between sentiment and topic from textual data (Mei et al., 2007; Titov and McDonald, 2008; Lin and He, 2009; He et al., 2011; Lazaridou et al., 2013; Li et al., 2013).

Here we modify a joint sentiment-topic model (JST) (Lin and He, 2009) to extract topics associated with positive/negative sentiment. Lin and He (2009) assume that topics are generated dependent on sentiment distributions and words are generated conditioned on the sentiment-topic pairs. JST can detect sentiment and topics simultaneously by encoding word prior sentiment information. However, words in the JST model are position-unaware, i.e., words from the same sentence/clause thus can have different topics or sentiments. We modify the JST model by assuming that lexical items from the same couplet are generated conditioned on the same sentiment-topic pairs. In Chinese poetry, a couplet is a pair of lines which have the same length and express a complete meaning. Lexical items within the same couplet usually relate to the same topic and keep the same polarity. Our position-aware JST model is depicted in Figure 1.

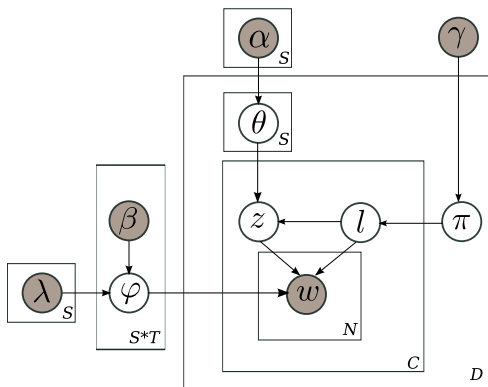


Figure 1: Position-aware JST model.

Assume we have a corpus consisting of D documents; each document is a sequence of C couplets and each word (lexical item) in the couplet is an item from a vocabulary index with V distinct terms; let S be the number of distinct sentiment labels and T the total number of topics. The process of generating a word w in document d under our position-aware JST model is as follows:

- For each sentiment label $l \in S$ and each topic $j \in T$, draw $\varphi_{lj} \sim \text{Dirichlet}(\lambda_l \times \beta_{lj}^T)$
- For each document d ,

- Draw the document’s sentiment distribution $\pi_d \sim \text{Dirichlet}(\gamma)$
- For each sentiment label l , draw a topic distribution $\theta_{d,l} \sim \text{Dirichlet}(\alpha)$
- For each couplet,
 1. choose a sentiment label $l_i \sim \text{Multinomial}(\pi_d)$
 2. choose a topic $z_i \sim \text{Multinomial}(\theta_{d,l_i})$
 3. generate words $w \sim \varphi_{l_i,z_i}$

Model Priors and hyperparameter settings.

We incorporate our sentiment lexicon (described in Section 3) containing 4153 positive lexical items and 4503 negative lexical items as prior knowledge into the position-aware JST model. Specifically, if w is found in the sentiment lexicon, then the element $\lambda_{lw} = 1$, and zero otherwise. Following Lin and He (2009), we set the symmetric $\beta = 0.01, \gamma = (0.05 \times L)/S$, where L is the average document length, S is the total number of sentiment labels⁷. The asymmetric prior α is learned from data.

Topics under different sentiment polarities.

We apply the position-aware sentiment-topic model to the Complete Anthology of Tang Poetry. The corpus contains 42,862 poems written by 2630 authors, with an average of 60 single characters in each poem. We represent each poem as couplets containing single and two-character lexical items⁸. We conduct experiments on $T = 15, 25, 50$ respectively. Table 5 shows the topic examples extracted by the position-aware sentiment-topic model with $T = 25$ under positive and negative sentiment labels respectively. The topics are labeled manually according to the lexical items found in them. Each topic is represented by the top 15 lexical items. These topics reflect common themes associated with positive/negative sentiment in Tang poetry. Moreover, the sentiment-topic distribution provides us with more insight on different aspects of social life in that historical period of China (AD 618 – 907).

For instance, poets wrote poems to praise the beauty of women and flowers (topic PT0 and

⁷In our experiments, the value of γ is around 1.00. This means that we do not assume any prior knowledge of the sentiment distribution of a poem and all possible sentiment distributions are equal.

⁸Similar to lexical network construction (Section 3.3), the two-character lexical items should appear at least 50 times in the whole corpus.

<i>Positive Sentiment</i>	
Topics	Lexical items
PT0: beautiful women	玉 (jade) 金 (gold) 红 (red) 罗 (silk) 香 (perfume) 女 (women) 翠 (green jade) 歌 (song) 舞 (dance) 楼 (building) 画 (painting) 珠 (pearl) 眉 (eyebrow) 双 (pair) 锦 (brocade)
PT1: feast_drinking	酒 (wine) 醉 (drunk) 歌 (song) 杯 (cup) 欢 (happy) 殷勤 (attentive) 客 (guest) 对 (face to face) 饮 (drink) 弦 (chord) 乐 (happy) 劝 (advise) 酌 (drink) 酣 (intoxicated) 筵 (feast)
PT2: war_victory	军 (military) 旌 (banner) 旗 (banner) 马 (horse) 将军 (general) 天 (god) 剑 (sword) 骑 (ride) 弓 (bow) 戎 (army) 功 (achievement) 旌旗 (banner) 箭 (arrow) 战 (battle) 射 (shoot)
PT3: literary	诗 (poetry) 文 (article) 书 (letter) 名 (reputation) 题 (inscribe) 句 (sentence) 君 (gentleman) 字 (character) 章 (chapter) 才 (gift) 篇 (article) 咏 (sing) 赋 (compose) 笔 (pen) 高 (high)
PT4: royal graciousness	德 (moral) 神 (god) 礼 (ritual) 乐 (music) 圣 (holy) 灵 (spirit) 明 (bright) 万 (great number) 惟 (only) 肃 (pay respect) 皇 (emperor) 帝 (emperor) 国 (country) 功 (achievement) 天 (god)
PT5: supernature	仙 (god) 玉 (jade) 丹 (vermilion) 天 (god) 神 (god) 霞 (rosy clouds) 紫 (violet) 灵 (spirit) 龙 (dragon) 烟 (mist) 神仙 (god) 瑶 (fairy) 清 (clean) 蓬 (fairyland)
PT6: Buddhism	香 (perfume) 禅 (meditation) 师 (master) 心 (heart) 世 (world) 法 (dharma) 清 (clean) 真 (truth) 莲 (lotus) 道 (Taoism) 净 (clean) 钟 (bell) 界 (world) 士 (scholar) 佛 (buddhism)
PT7: traveling	楼 (building) 高 (high) 上 (go up) 登 (climb) 陵 (hill) 游 (travel) 州 (province) 台 (terrace) 南 (south) 长安 (place name) 下 (go down) 临 (overlook) 佳 (beautiful) 浮云 (clouds) 武陵 (place name)
PT8: flowers	香 (perfume) 花 (flower) 枝 (twig) 露 (dew) 芙蓉 (hibiscus) 桃 (peach) 光 (light) 芳 (fragrant) 轻 (gentle) 袅 (delicate) 艳 (charming) 团 (round) 新 (fresh) 翠 (green jade) 兰 (orchid)
<i>Negative Sentiment</i>	
Topics	Lexical items
NT0: lovesickness	别 (farewell) 离 (leave) 惆怅 (melancholy) 相思 (lovesick) 忆 (memorize) 望 (hope) 相逢 (meet) 恨 (hate) 送 (farewell) 年 (year) 君 (gentleman) 梦 (dream) 远 (remote) 归 (return) 泪 (tears)
NT1: hardness of life	老 (old) 病 (sickness) 多 (many) 无 (no) 白 (gray) 鬓 (temple hair) 不 (no) 吟 (sing) 愁 (anxious) 衰 (feeble) 贫 (poor) 白头 (old age) 卧 (lie down) 更 (more) 难 (hard)
NT2: hardness of war	塞 (fortress) 边 (border) 城 (city) 河 (river) 胡 (barbarian) 关 (barrier) 征 (attack) 风 (wind) 月 (moon) 北 (north) 战 (battle) 雪 (snow) 虏 (captive) 戍 (garrison) 鼓 (drum)
NT3: traveling by boat	江 (river) 水 (water) 舟 (boat) 湖 (lake) 波 (wave) 岸 (bank) 浪 (wave) 风 (wind) 沙 (sand) 帆 (sail) 孤 (lonely) 海 (sea) 船 (boat) 浦 (riverside) 月 (moon)
NT4: homesickness	堪 (bear) 断 (break) 故 (home) 落 (fall) 乡 (home) 凄 (sorrow) 愁 (anxious) 归 (return) 泪 (tears) 路 (road) 肠 (intestine, often used to describe heartbroken) 伤 (sad) 悲 (sad) 涯 (shore) 音 (news)
NT5: visiting monks	山 (mountain) 僧 (monk) 松 (pine) 寺 (temple) 石 (rock) 林 (forest) 深 (deep) 泉 (fountain) 夜 (night) 寒 (cold) 寻 (search) 客 (guest) 峰 (peak) 云 (cloud) 溪 (stream)
NT6: sad scenery	秋 (autumn) 风 (wind) 叶 (leaf) 寒 (cold) 雨 (rain) 夕阳 (sunset) 晚 (evening) 水 (water) 暮 (twilight) 霜 (frost) 菊 (chrysanthemum) 蝉 (cicada) 山 (mountain) 落 (fall) 凉 (cold)
NT7: sad scenery	雨 (rain) 猿 (ape) 峡 (gorge) 啼 (cry) 江 (river) 云 (clouds) 湘 (river name) 楚 (place name) 山 (mountain) 暮 (twilight) 蹉跎 (waste time) 巫峡 (gorge name) 云雨 (clouds and rain) 梦 (dream) 巫山 (mountain name)
NT8: death and destruction	死 (death) 生 (birth) 苦 (miserable) 骨 (bone) 悴 (sad) 饥 (hungry) 血 (blood) 泥 (mud) 杀 (kill) 枯 (withered) 憔悴 (thin and pallid) 鬼 (ghost) 恶 (evil) 祸 (disaster) 土 (dust)

Table 5: Examples of topics extracted by the position-aware sentiment-topic model.

PT8). They eulogized the Tang empire and were proud of the country's victory in war (topic PT4 and PT2). They enjoyed drinking at banquet (topic PT1), admire others' literary achievements (topic PT3) and praised beautiful landscapes while traveling (topic PT7). In addition, it seems that Buddhism and supernatural beings are the favored topics of poetry in the Tang Dynasty (topic PT6 and PT5). This might reflect that Buddhism was at its peak in the Tang Dynasty and many poets were devout Buddhists.

On the other hand, poets felt sad for women

who separated with their loved ones (topic NT0). They were angry for death and destruction caused by tyranny and turmoil (topic NT8). Poets were tired of war (topic NT2) and they were homesick while traveling alone (topic NT3 and NT4).

It is worth noticing that some topics are associated with both positive and negative sentiment. For instance, poets were happy about the victory of war (topic PT2). At the same time, they were sad about the destruction/pain caused by war (topic NT2), i.e., soldiers were forced to leave their homelands and loved ones. Also, traveling

could involve both sentiments: poets praised the beauty of nature (PT7); but they also felt lonely while traveling alone (NT3, NT4). Specifically, some scenes during traveling become common imageries/symbols which imply sad emotions, such as things related to water (i.e., 江 (river), 舟 (boat), 湖 (lake), 波 (wave), 岸 (bank), 帆 (sail), 浪 (wave), 浦 (riverside) in topic NT3), or 猿 (ape) in topic NT7. This also reflects that poets liked traveling and that traveling by boat was popular in the Tang Dynasty.

5.2 Sentiment of Different Poets

In order to analyze how sentiment differs among poets, we choose four famous poets from the Tang dynasty: 李白 (Li Bai), 杜甫 (Du Fu), 王维 (Wang Wei), and 白居易 (Bai Juyi).

Li Bai enjoys the title of the “Supernatural Being of Poem”. His works are full of passion, imagination and elegance. Du Fu, known as the “Poet Sage”, is known for his anti-war stance and concerns for the poor. Wang Wei, the poet of landscape, has written lots of elegant and exquisite poems. Bai Juyi has been known for his plain and easily comprehensible style of poem, and for his social and political criticism. Table 6 lists the poets’ lifetimes and the number of their poems collected in the Complete Anthology of Tang Poetry.

Poet	Lifetime	Number of Poems
李白 (Li Bai)	701 – 762	891
杜甫 (Du Fu)	712 – 770	1,151
王维 (Wang Wei)	701 – 761	350
白居易 (Bai Juyi)	772 – 846	2,640

Table 6: Number of poems by chosen poets.

On the basis of our sentiment lexicon, we predict the sentiment orientation of each poem using the method described in Section 4.2 with threshold $t = 0^9$. We then compare the percentage of poems with different sentiment for each poet.

Figure 2 shows that the percentage of positive poems written by Li Bai is the highest among our four poets, whereas the percentage of negative poems written by Du Fu is the highest¹⁰. It

⁹We could also use our position-aware sentiment-topic model to predict the sentiment orientation of each poem. However, for the task of predicting sentiment orientation of poems, we found that the position-aware JST model does not perform as well as our simple method described in Section 4.2 on the same test dataset.

¹⁰We find that the comparison under different value of t keeps the same pattern as shown in Figure 2.




Figure 2: Sentiment of different poets.

seems that Du Fu expressed sad emotions more frequently in his poems. This may relate to his frustrating experiences. He aspired a career as a civil servant, but his failure in the examination put an end to his chances to have an official career. During the period of political turmoil, Du Fu fled to the capital. But he was captured and then wandered as a refugee. Most of Du Fu’s life was spent in poverty. One of his sons even died from starvation because of the family’s poverty. Du Fu wrote several poems to express his deep sadness for his son’s death. It is worth noting that although Bai Juyi has been known for his critical political poems, he also wrote a great amount of poems expressing leisurely and comfortable mood, especially in his late years.

6 Conclusions

We propose a novel graph-based method to build a sentiment lexicon for classical Chinese poetry. Our approach is weakly supervised and outperforms a previous PMI-based approach (Turney and Littman, 2003) in different evaluation settings. On the basis of our sentiment lexicon, we analyze sentiment in Tang poetry from different perspectives: which topics are associated with positive and negative sentiment, and how sentiment is distributed among different poets. Our analysis results are in line with the main findings established in classical Chinese literary studies.

The work presented in this paper provides a quantitative means to study sentiment in classical Chinese poetry. We hope it can benefit other research topics, such as poetry generation (He et al., 2012; Zhang and Lapata, 2014) and poetry imagery/style analysis (Fang et al., 2009).

References

- Eneko Agirre and Aitor Soroa. 2009. Personalizing PageRank for word sense disambiguation. In *Proceedings of the 12th Conference of the European Chapter of the Association for Computational Linguistics*, Athens, Greece, 30 March – 3 April 2009, pages 33–41.
- Stefano Baccianella, Andrea Esuli, and Fabrizio Sebastiani. 2010. SentWordNet 3.0: An enhanced lexical resource for sentiment analysis and opinion mining. In *Proceedings of the 7th International Conference on Language Resources and Evaluation*, La Valetta, Malta, 17–23 May 2010, pages 2200–2204.
- David M. Blei, Andrew Y. Ng, and Michael I. Jordan. 2003. Latent dirichlet allocation. *J. Mach. Learn. Res.*, 3(1):993–1022.
- Sergey Brin and Lawrence Page. 1998. The anatomy of a large-scale hypertextual web search engine. *Computer Networks and ISDN Systems*, 30(1–7):107–117.
- Andrea Esuli and Fabrizio Sebastiani. 2005. Determining the semantic orientation of terms through gloss analysis. In *Proceedings of the ACM SIGIR Conference on Information and Knowledge Management (CIKM)*.
- Alex Chengyu Fang, Fengju Lo, and Cheuk Kit Chinn. 2009. Adapting NLP and corpus analysis techniques to structured imagery analysis in classical Chinese poetry. In *Proceedings of the RANLP 2009 Workshop on Adaptation of Language Resources and Technology to New Domains*, Borovets, Bulgaria, 14 June 2009, pages 27–34.
- Ahmed Hassan, Amjad Abu-Jbara, Rahul Jha, and Dragomir Radev. 2011. Identifying the semantic orientation of foreign words. In *Proceedings of the ACL 2011 Conference Short Papers*, Portland, Oregon, 19–24 June 2011, pages 592–597.
- Taher H. Haveliwala. 2002. Topic-sensitive PageRank. In *Proceedings of the 11th World Wide Web Conference*, Honolulu, Hawaii, USA, 7–11 May 2002, pages 517–526.
- Yulan He, Harith Alani, and Deyu Zhou. 2010. Exploring English lexicon knowledge for chinese sentiment analysis. In *CIPS-SIGHAN Joint Conference on Chinese Language Processing*. Beijing, China, 28–29 August 2010, pages 121–128.
- Yulan He, Chengua Lin, and Harith Alani. 2011. Automatically extracting polarity-bearing topics for cross-domain sentiment analysis. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics*, Portland, Oregon, 19–24 June 2011, pages 123–131.
- Jing He, Ming Zhou, and Long Jiang. 2012. Generating chinese classical poems with statistical machine translation models. In *Proceedings of the 26th Conference on the Advancement of Artificial Intelligence*, Toronto, Canada, 22–26 July 2012, page 1650–1656.
- Junfeng Hu. 2001. *The Lexicon Meaning Analysis-based Computer Aided Research Work of Chinese Ancient Poems*. Ph.D. thesis, Peking University. Published by Graduate Linguistics Student Organisation.
- Nobuhiro Kaji and Masaru Kitsuregawa. 2007. Building lexicon for sentiment analysis from massive collection of HTML documents. In *Proceedings of the 2007 Joint Conference on Empirical Methods in Natural Language Processing and Computational Language Learning*, Prague, Czech Republic, 28–30 June 2007, pages 1075–1083.
- Jaap Kamps, Maarten Marx, Robert J. Mokken, and Maarten de Rijke. 2004. Using WordNet to measure semantic orientation of adjectives. In *Proceedings of the 4th International Conference on Language Resources and Evaluation*, Lisbon, Portugal, 26–28 May 2004, pages 1115–1118.
- Hiroshi Kanayama and Tetsuya Nasukawa. 2006. Fully automatic lexicon expansion for domain-oriented sentiment analysis. In *Proceedings of the 2006 Conference on Empirical Methods in Natural Language Processing*, Sydney, Australia, 22–23 July 2006, pages 355–363.
- Yu-kung Kao and Tsu-lin Mei. 1971. Syntax, diction, and imagery in T’ang poetry. *Harvard Journal of Asiatic Studies*, 31(2):49–136.
- Yu-kung Kao and Tsu-lin Mei. 1978. Meaning, metaphor, and allusion in T’ang poetry. *Harvard Journal of Asiatic Studies*, 38(2):281–355.
- Svetlana Kiritchenko, Xiaodan Zhu, and Sarif Mohammad. 2014. Sentiment analysis of short informal texts. *Journal of Artificial Intelligence Research*, 50(1):723–762.
- Lun-Wei Ku, Yu-Ting Liang, and Hsin-Hsi Chen. 2006. Tagging heterogeneous evaluation corpora for opinionated tasks. In *Proceedings of the 5th International Conference on Language Resources and Evaluation*, Genoa, Italy, 22–28 May 2006, pages 667–670.
- Angeliki Lazaridou, Ivan Titov, and Caroline Sporleder. 2013. A Bayesian model for joint unsupervised induction of sentiment, aspect and discourse representations. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics*, Sofia, Bulgaria, 4–9 August 2013, pages 1630–1639.
- John Lee and Yin Hei Kong. 2012. A dependency treebank of classical chinese poems. In *Proceedings of the 2012 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, Montréal, Québec, Canada, 3–8 June 2012, pages 191–199.

- John Lee and WONG Tak-sum. 2012. Glimpses of ancient china from classical chinese poems. In *Proceedings of the 24th International Conference on Computational Linguistics*, Mumbai, India, 8–15 December 2012, pages 621–632.
- Chiwei Li, Myle Ott, and Claire Cardie. 2013. Identifying manipulated offerings on review portals. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, Seattle, Wash., 18–21 October 2013, pages 1933–1942.
- Chenghua Lin and Yulan He. 2009. Joint sentiment/topic model for sentiment analysis. In *Proceedings of the ACM 18th Conference on Information and Knowledge Management (CIKM 2009)*, Hong Kong, 2–6 November 2009, pages 375–387.
- Fengzhu Luo. 2009. The application of affective and semantics computing in literature and history research. *Literature Heritage*, 1(1):138–141.
- Qiaozhu Mei, Xu Ling, Matthew Wondra, Hang Su, and ChengXiang Zhai. 2007. Topic sentiment mixture: Modeling facets and opinions in weblogs. In *Proceedings of the 16th World Wide Web Conference*, Banff, Canada, 8–12 May, 2007, pages 171–180.
- Saif Mohammad, Cody Dunne, and Bonnie Dorr. 2009. Generating high-coverage semantic orientation lexicons from overtly marked words and a thesaurus. In *Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing*, Singapore, 6–7 August 2009, pages 599–608.
- Delip Rao and Deepa avichandran. 2009. Semi-supervised polarity lexicon induction. In *Proceedings of the 12th Conference of the European Chapter of the Association for Computational Linguistics*, Athens, Greece, 30 March – 3 April 2009, pages 675–682.
- Hiroya Takamura, Takashi Inui, and Manabu Okumura. 2005. Extracting semantic orientation of words using spin model. In *Proceedings of the 43rd Annual Meeting of the Association for Computational Linguistics*, Ann Arbor, Mich., 25–30 June 2005, pages 133–140.
- Ivan Titov and Ryan McDonald. 2008. A joint model of text and aspect ratings for sentiment summarization. In *Proceedings of the 46th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, Columbus, Ohio, 15–20 June 2008, pages 308–316.
- Peter D. Turney and Michael L. Littman. 2003. Measuring praise and criticism: Inference of semantic orientation from association. *ACM Transactions on Information Systems (TOIS)*, 21(4):315–346.
- Rob Voigt and Dan Jurafsky. 2013. Tradition and modernity in 20th century Chinese poetry. In *Proceedings of the NAACL-HLT 2013 Workshop on Computational Linguistics for Literature*, Atlanta, USA, 14 June 2013, pages 1–6.
- Xiaojun Wan. 2008. Using bilingual knowledge and ensemble techniques for unsupervised Chinese sentiment analysis. In *Proceedings of the 2008 Conference on Empirical Methods in Natural Language Processing*, Waikiki, Honolulu, Hawaii, 25–27 October 2008, pages 553–561.
- Xiaojun Wan. 2009. Co-training for cross-lingual sentiment classification. In *Proceedings of the Joint Conference of the 47th Annual Meeting of the Association for Computational Linguistics and the 4th International Joint Conference on Natural Language Processing*, Singapore, 2–7 August 2009, pages 235–243.
- Burton Watson. 1971. *Chinese Lyricism: Shih Poetry from the Second to the Twelfth Century*. New York: Columbia University Press.
- Scott White and Padhraic Smyth. 2003. Algorithms for estimating relative importance in networks. In *Proceedings of the 9th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, Washington DC, USA, 24–27 August 2004, pages 266–275.
- Difei Xiao, editor. 1999. *Tang Poetry analysis dictionary*. Shanghai Lexicographical Publishing House, Shanghai, China.
- Wenpu Xing and Ali Ghorbani. 2004. Weighted PageRank algorithm. In *Proceedings of the Second Annual Conference on Communication Networks and Services Research*, Fredericton, N.B., Canada, 19-21 May 2004, pages 305–314.
- Taras Zagibalov and John Carroll. 2008. Automatic seed word selection for unsupervised sentiment classification of Chinese text. In *Proceedings of the 22nd International Conference on Computational Linguistics*, Manchester, U.K., 18–22 August 2008, pages 1073–1080.
- Xingxing Zhang and Mirella Lapata. 2014. Chinese poetry generation with recurrent neural networks. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing*, Doha, Qatar, 25–29 October 2014, pages 670–680.
- Chang-Le Zhou, Wei You, and Xiao-Jun Ding. 2010. Genetic algorithm and its implementation of automatic generation of Chinese SONGCI. *Journal of Software*, 21(3):427–437.