

Surveying the Dead Minds: Historical-Psychological Text Analysis with Contextualized Construct Representation (CCR) for Classical Chinese

Yuqi Chen
Peking University
cyq0722@pku.edu.cn

Ying Li
Peking University
yingliclaire@pku.edu.cn

Sixuan Li
Xiaoying AI Lab
lisixuan@xiaoyingai.com

Mohammad Atari
University of Massachusetts Amherst
matari@umass.edu

Abstract

In this work, we develop a pipeline for historical-psychological text analysis in classical Chinese. Humans have produced texts in various languages for thousands of years; however, most of the computational literature is focused on contemporary languages and corpora. The emerging field of historical psychology relies on computational techniques to extract aspects of psychology from historical corpora using new methods developed in natural language processing (NLP). The present pipeline, called Contextualized Construct Representations (CCR), combines expert knowledge in psychometrics (i.e., psychological surveys) with text representations generated via Transformer-based language models to measure psychological constructs such as traditionalism, norm strength, and collectivism in classical Chinese corpora. Considering the scarcity of available data, we propose an indirect supervised contrastive learning approach and build the first Chinese historical psychology corpus (C-HI-PSY) to fine-tune pre-trained models. We evaluate the pipeline to demonstrate its superior performance compared with other approaches. The CCR method outperforms word-embedding-based approaches across all of our tasks and exceeds prompting with GPT-4 in most tasks. Finally, we benchmark the pipeline against objective, external data to further verify its validity.¹

1 Introduction

Humans have been producing written language for thousands of years. Historical populations have expressed their norms, values, stories, songs, and more in these texts. Such historical corpora represent a rich yet underexplored source of psychological data that contains the thoughts, feelings, and actions of people who lived in the past (Jackson

¹Our code and data are available at <https://github.com/yukiyuqichen/His-Psy>.

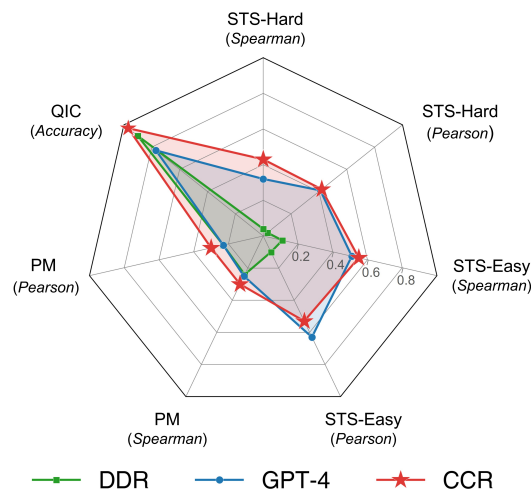


Figure 1: Comparison of the best performance among the DDR, CCR, and prompting methods on three tasks in the C-HI-PSY test set. (STS: Semantic Textual Similarity, PM: Psychological Measure, QIC: Questionnaire Item Classification)

et al., 2021). The emerging field of “historical psychology” has been developed to understand how different aspects of psychology vary over historical time and how the origins of our contemporary psychology are rooted in historical processes (Atari and Henrich, 2023; Muthukrishna et al., 2021; Baumard et al., 2024). Since we cannot access “dead minds” directly but can access their textual remains, natural language processing (NLP) is the primary method to extract aspects of psychology from historical corpora. Previous works, however, are often monolingual and in English (Blasi et al., 2022). In addition, much of the literature at the intersection of psychology and NLP has relied on bag-of-words or word embedding models, focusing on non-contextual word meanings rather than a holistic approach to language modeling.

Recently, more research attention in the NLP community has been directed to historical and ancient languages (Johnson et al., 2021), including

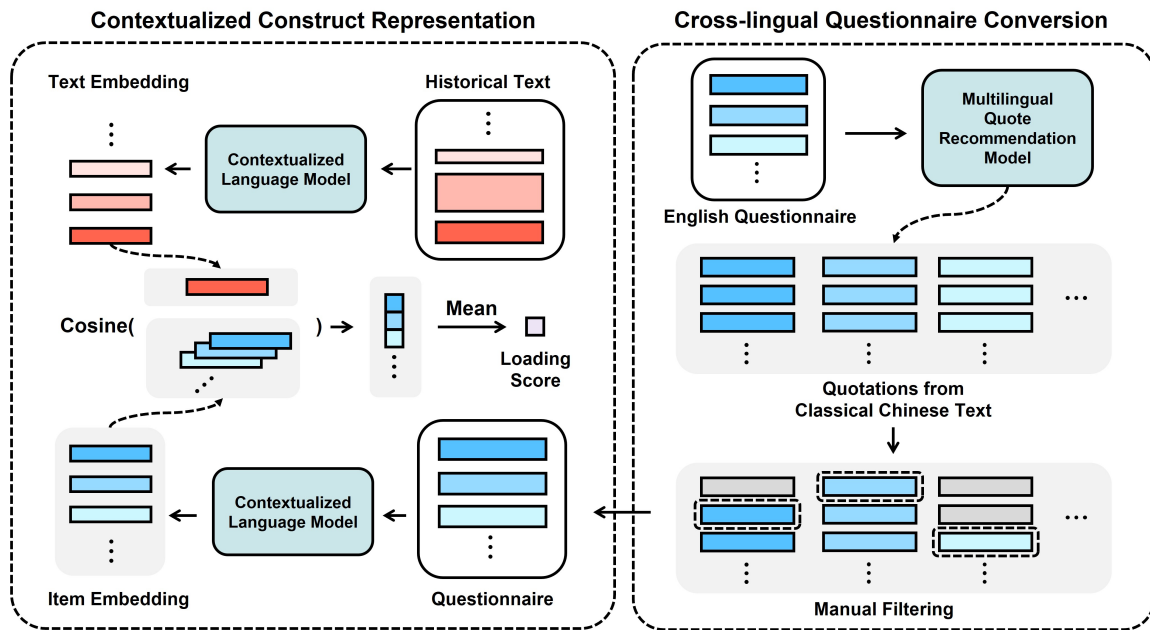


Figure 2: Pipeline for cross-lingual questionnaire conversion and contextualized construct representation.

but not limited to English (Manjavacas Arevalo and Fonteyn, 2021), Latin (Bamman and Burns, 2020), ancient Greek (Yousef et al., 2022), and ancient Hebrew (Swanson and Tyers, 2022). While all these languages have historical significance, classical Chinese is particularly important in the quantitative study of history. China has a long history spanning thousands of years, largely recorded in classical Chinese. The language served as a medium for expressing and disseminating influential philosophical and religious ideas. Confucianism, Daoism, and later Buddhism (through translations from Sanskrit) all found expression in classical Chinese, profoundly shaping Chinese thought, ethics, governance, and norms. As more resources become readily available for classical Chinese, scholars of ancient China can test more specific hypotheses using computational methods (Liu et al., 2023; Slingerland, 2013; Slingerland et al., 2017).

Due to its historical significance and geographical coverage, classical Chinese represents one of the most important languages in historical psychology (Atari and Henrich, 2023). Prior work in social science has often relied on bag-of-words approaches (Zhong et al., 2023) or bottom-up techniques such as topic modeling (Slingerland et al., 2017). In the NLP community, various Transformer-based models for classical Chinese have been developed (Tian et al., 2021; Wang and Ren, 2022; Yan and Chi, 2020; Wang et al., 2023a), primarily for tasks like punctuation pre-

diction (Zhou et al., 2023), poem generation (Tian et al., 2021), and translation (Wang et al., 2023b). However, they have not been applied to theory-driven psychological text analysis for extracting psychological constructs (e.g., moral values, norms, cultural orientation, mental health, religiosity, emotions, and thinking styles) from historical data.

Transformer-based language models (Vaswani et al., 2017) are crucial for psychological text analysis because psychological constructs are often complex, and sentence-level semantics (and above) will more effectively capture psychological meanings than isolated words (Demszky et al., 2023) or non-contextual word embedding models (Kennedy et al., 2021).

Here, we create a pipeline called Contextualized Construct Representation (CCR) for historical-psychological text analysis in classical Chinese. Although similar method has recently been developed for contemporary psychological text analysis (Atari et al., 2023b), it can be adapted for historical NLP. Taking advantage of contextual language models, CCR does not require selecting a priori lists of words to represent a psychological construct (e.g., the popular Linguistic Inquiry and Word Count program, Boyd et al., 2022); instead, making use of thousands of existing questionnaires (which typically include face-valid declarative sentences with which participants agree or disagree) that have been developed and validated in psychology over the last century.

The pipeline of CCR for classical Chinese proceeds in five steps: (1) selecting a questionnaire for the psychological construct of interest; (2) converting the questionnaire, usually in English, into classical Chinese; (3) representing questionnaire items as embeddings using a contextual language model; (4) generating the embedding of the target text using a contextual language model; (5) computing the cosine similarity between the item and text embeddings. This straightforward pipeline is particularly useful for social science, wherein researchers are interested in interpretability and hypothesis testing.

There are two main challenges of using the CCR pipeline in analyzing Chinese historical texts: (1) popular self-report questionnaires, widely accepted by psychologists, are often in English, making it difficult to align them with classical Chinese texts; (2) there is a lack of psychology-specific Transformer-based models for classical Chinese, making it difficult to obtain high-quality representations of Chinese historical texts. To address the first challenge, we propose a pipeline that uses a multilingual quotation recommendation model (Qi et al., 2022) to convert contemporary English questionnaires into contextually meaningful classical Chinese sentences (§3.1). To tackle the second challenge, we build the first Chinese historical psychology corpus (C-HI-PSY) and introduce an approach based on indirect supervision (He et al., 2021; Yin et al., 2023; Xu et al., 2023a) and contrastive learning (Chopra et al., 2005; Schroff et al., 2015; Gao et al., 2021; Chuang et al., 2022) to fine-tune pre-trained models (§3.2).

2 Related Work

Psychological Text Analysis Textual corpora contain imprints of historical biases that can be captured by machine-learning models (Caliskan et al., 2017; Garg et al., 2018). The changes in these biases can also serve as clues for understanding the shifts in human sociality. Given the increasing amount of online textual data, many social scientists are turning to NLP to test their theories. Unlike in some computational fields, social scientists traditionally give primacy to “theory” rather than prediction (Yarkoni and Westfall, 2017). Hence, theory-driven text analysis is the first methodological choice in social sciences, including psychology (Jackson et al., 2021; Wilkerson and Casas, 2017; Boyd and Schwartz, 2021). Given the impor-

tance of theory development and hypothesis testing, many social scientists have developed dictionaries to assess psychological constructs as diverse as moral values (Graham et al., 2009), stereotypes (Nicolas et al., 2021), polarization (Simchon et al., 2022), and threat (Choi et al., 2022).

Distributed Dictionary Representation (DDR)

Aiming to integrate psychological theories with the capabilities of word embeddings, Garten et al. (2018) proposed the Distributed Dictionary Representation (DDR) as a top-down psychological text-analytic method. This method involves (a) defining a concise list of words by experts to capture a specific concept, (b) using a word-embedding model to represent these individual words, (c) computing the centroid of these word representations to define the dictionary’s representation, (d) determining the centroid of the word embeddings within a given document, and (e) assessing the cosine similarity between the dictionary’s representation and that of the document. DDR has been a useful approach in measuring moral rhetoric (Wang and Inbar, 2021), temporal trends in politics (Xu et al., 2023b), and situational empathy (Zhou et al., 2021).

Sentence Transformer While BERT (Devlin et al., 2018) can identify sentences with similar semantic meanings, this process can be resource-intensive. To enhance the performance of BERT-like models for tasks like semantic similarity assessments, Reimers and Gurevych (2019) developed Sentence Transformer with a Siamese network structure, which outperforms conventional models in tasks related to sentences and significantly reduces the time needed for computations. It is engineered to generate sentence embeddings that capture the core semantic content, ensuring that sentences with comparable meanings are represented by closely positioned embeddings in the vector space. Therefore, Sentence Transformer provides an efficient and less computationally demanding method for evaluating semantic similarities between sentences, making it particularly useful in fields such as psychology (Juhng et al., 2023; Sen et al., 2022).

3 Methodology

Employing the CCR pipeline for historical-psychological text analysis necessitates the use of valid questionnaires and appropriate contextual language models that can effectively represent sen-

tences or paragraphs. We propose two distinct pipelines: (1) a cross-lingual questionnaire conversion pipeline to obtain psychological questionnaires in classical Chinese; (2) an indirect supervised contrastive learning pipeline to fine-tune pre-trained Transformer-based models using a historical psychological corpus.

3.1 Cross-lingual Questionnaire Conversion

In order to calculate semantic similarities between questionnaires, typically in English, and the Chinese historical texts to be measured, typically in classical Chinese, we introduce a novel workflow for Cross-lingual Questionnaire Conversion (CQC). Considering the gaps between the two languages and cultures, we employ quotations from authentic historical texts, as they can integrate more naturally within the context of classical Chinese, instead of relying on translations.

The process of converting a contemporary English questionnaire Q into a classical Chinese questionnaire \tilde{Q} is illustrated in the right panel of Figure 2. For each questionnaire item ($q_i \in Q$), the multilingual quote recommendation model, “QuoteR” (Qi et al., 2022), which is trained on a dataset that includes English, modern Standard Chinese, and classical Chinese, can identify a set of quotations $\{\tilde{q}\}_i$ in classical Chinese that are semantically similar to the English sentence q_i . All the items are entered into the model for each questionnaire, resulting in a pool of corresponding quotations. Then, manual filtering is followed to eliminate quotations of low quality, which can be either inappropriate or not explicitly relevant to the psychological construct. Ultimately, the most similar quotations \tilde{q}_i are selected, substituting for every English q_i to construct \tilde{Q} in classical Chinese.

3.2 Indirect Supervised Contrastive Learning

To obtain better psychology-specific representations for CCR in Chinese historical texts, we introduce an indirect supervised contrastive learning approach to finetune pre-trained Transformer-based models, as shown in Figure 3. A detailed example is shown in Figure 9 in the Appendix.

Historical Psychology Corpus We assemble a refined corpus named Chinese Historical Psychology Corpus (C-HI-PSY), which is comprised of 21,539 paragraphs (\mathcal{S}) from 667 distinct historical articles and book chapters in classical Chinese (Table 3 in the Appendix) that meet specific criteria.

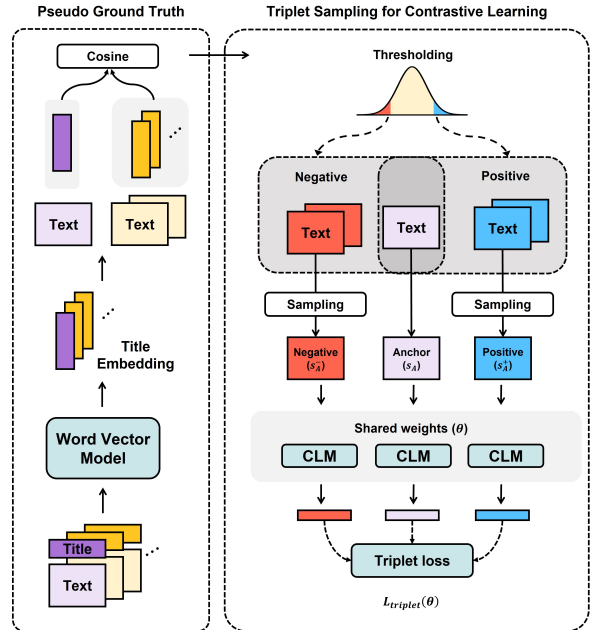


Figure 3: Pipeline of triplet sampling and contrastive learning. CLM stands for contextual language model.

Our selection criteria included (a) whether the article revolved around a theme related to psychology and (b) whether the article’s title revealed the main theme. The titles of these works (\mathcal{T} , $|\mathcal{T}| \ll |\mathcal{S}|$), each carefully selected for their relevance to moral values, were labeled by the ancients for us and thus can serve as pseudo for their topics, including “節義” (moral integrity), “孝弟” (filial piety and fraternal duty), “盡忠” (utmost loyalty), “廉恥” (sense of shame), “清介” (pure and incorruptible), and “愛己” (love oneself), among others.

We divide our data into training, validation, and testing sets, allocating 60%, 20%, and 20% of the data to each set, respectively. The distribution of paragraph lengths across different sets is consistent, as shown in Figure 7 in the Appendix.

Pseudo Ground Truth from Titles Since the title ($t_i \in \mathcal{T}$) of a paragraph ($s_i \in \mathcal{S}$) is a concise summary of the moral values reflected in the paragraph, the semantic similarity between titles, $\mathbf{sim}(t_i, t_j)$, can be considered as the pseudo ground truth for the semantic similarity between corresponding paragraphs, $\mathbf{sim}(s_i, s_j)$. The similarity between titles can be obtained by embedding the titles via $E_T(\cdot)$ and calculating their cosine similarity $\mathbf{cos}(E_T(t_i), E_T(t_j))$. To perform word embedding on the titles, we trained five word vector models on a large classical Chinese corpus containing over a billion tokens using different architectures, and picked the best-performing one (Appendix B).

Positive and Negative Sampling We calculate the cosine similarities between the title embeddings $\mathbf{cos}(E_T(t_i), E_T(t_j))$, obtained through the word vector model, of all title pairs (the Cartesian product $\mathcal{T} \times \mathcal{T}$) in the corpus. The distribution of title similarities is illustrated in Figure 8 in the Appendix. We obtain positive and negative paragraph pairs by thresholding the similarities of title pairs. Paragraphs whose titles have similarities exceeding the upper threshold δ^+ , as well as those with identical titles, were identified as positive pairs $(\mathcal{S} \times \mathcal{S})^+$, that is,

$$\{(s_i, s_j)^+ \mid \mathbf{sim}(E_T(t_i), E_T(t_j)) > \delta^+\}$$

Conversely, those with titles having similarities below the lower threshold δ^- were designated as negative pairs $(\mathcal{S} \times \mathcal{S})^-$, that is,

$$\{(s_i, s_j)^- \mid \mathbf{sim}(E_T(t_i), E_T(t_j)) < \delta^-\}$$

We experiment with several threshold settings, including 0.5th/99.5th, 1st/99th, 10th/90th, and 25th/75th percentiles, on the C-HI-PSY validation set using the base model “bert-ancient-chinese” (Wang and Ren, 2022). Our findings demonstrate that the 10th/90th percentile threshold yields the best performance, see Figure 4. Hence, for the following experiments, if not specified, the threshold setting has been taken as 10th/90th.

Triplet Sampling We implement two strategies, random sampling and hard sampling, to construct triplets of anchor-positive-negative paragraphs (s_A, s_A^+, s_A^-) from the training set. In random sampling, we select one positive instance s_A^+ and one negative instance s_A^- randomly from the respective positive pairs $(s_A \times \mathcal{S})^+$ and negative pairs $(s_A \times \mathcal{S})^-$ of the anchor s_A . In hard sampling, we utilize the pre-trained model $f_\theta(\cdot)$, which is later fine-tuned on these triplets, to embed paragraphs and calculate cosine similarities between the positive and negative pairs as $\mathbf{cos}(f_\theta(s_A), f_\theta(s_A^{+/-}))$. For the positive instance, we choose the paragraph with the lowest similarity to the anchor from its positive pairs, that is,

$$s_A^+ = \operatorname{argmin}_s \{\mathbf{cos}(f_\theta(s_A), f_\theta(s)) \mid (s_A, s) \in (s_A \times \mathcal{S})^+\}$$

Conversely, for the negative instance, we select the paragraph with the highest similarity to the anchor

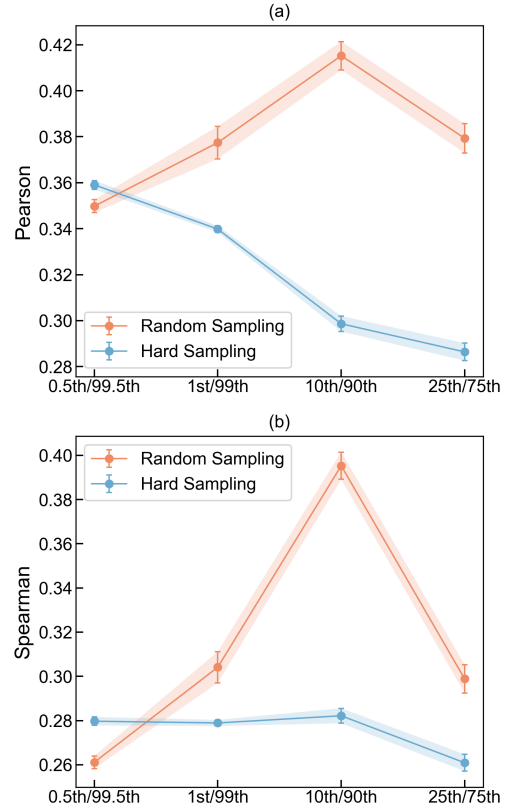


Figure 4: Performance variation with sampling methods and thresholds.

from its negative pairs, that is,

$$s_A^- = \operatorname{argmax}_s \{\mathbf{cos}(f_\theta(s_A), f_\theta(s)) \mid (s_A, s) \in (s_A \times \mathcal{S})^-\}$$

To prevent the model from over-fitting, we ensure that each paragraph is used as an anchor only once, applying this rule across both random and hard sampling strategies. We also compare the two sampling procedures in Figure 4 with respect to each positive-negative splitting threshold. We find that the random sampling procedure is better than hard sampling ever since the threshold is higher/lower than 0.5th/99.5th; we note that the case could be due to the noise inevitably caused by the indirect supervised learning approach, which drove the hard sampling procedure to fail at finding helpful instances (see Limitation).

Fine-tuning with Contrastive Learning We fine-tune several pre-trained Transformer-based models (Wang and Ren, 2022; Yan and Chi, 2020; Reimers and Gurevych, 2019; Xu, 2023) on the C-HI-PSY training set, using a triplet loss function (Schroff et al., 2015),

$$L_{\text{triplet}}(\theta) = \sum_{s_A \in S} \max\{\mathcal{D}^+ - \mathcal{D}^- + \alpha, 0\}$$

where θ denotes the pre-trained weights to be fine-tuned, \mathcal{D}^+ denotes the distance between the positive pair, i.e. $\|f_{\theta}(s_A) - f_{\theta}(s_A^+)\|_2$, \mathcal{D}^- denotes the distance between the negative pair, i.e. $\|f_{\theta}(s_A) - f_{\theta}(s_A^-)\|_2$, and α stands for a margin between positive and negative pairs to ensure that the model does not trivially satisfy the condition by making the embeddings of the anchor, positive, and negative samples equal to each other. This loss function aims to minimize the squared Euclidean norm between the anchor and positive, and maximize the squared Euclidean norm between the anchor and negative.

We construct paragraph pairs from the C-HI-PSY validation set through random sampling to validate the models during training, using the similarities between titles as pseudo ground truth to gauge the similarities between paragraphs. We perform a hyperparameter sweep (Table 5), to select the best-performing configuration for each model, as shown in Table 6 in the Appendix.




Figure 5: Comparison of model performance using the CCR method on the three tasks in the C-HI-PSY test set before and after fine-tuning. (Model A: bert-ancient-chinese, B: guwenbert-base, C: guwenbert-large, D: paraphrase-multilingual-MiniLM-L12-v2, E: text2vec-base-chinese, F: text2vec-base-chinese-paraphrase, G: text2vec-large-chinese)

4 Evaluation and Results

We set up different levels of tasks to evaluate the CCR method (using Sentence Transformer models),

and compare it with the DDR method (using word embedding models) and the prompting method (using generative LLMs). The results are shown in Table 1.

4.1 Semantic Understanding

Understanding of Historical Text: Semantic Textual Similarity For the CCR method, we embed whole paragraphs with Sentence Transformer models, and then calculate the cosine similarity between each pair of paragraphs. For the DDR method, we average the word vectors of all the words in the paragraph, and then calculate the cosine similarity between each pair of paragraphs. For the LLM-prompting method, we craft a few-shot prompt (Brown et al., 2020; Si et al., 2023) (Figure 10 in the Appendix) asking for a similarity score, ranging from 0 to 1, between each pair of paragraphs. As mentioned, similarities between the titles of each pair of paragraphs are used as the pseudo ground truth.

We construct paragraph pairs for evaluation from the C-HI-PSY test set using two sampling methods: (1) random sampling, where paragraphs are randomly paired, and (2) threshold sampling, which pairs paragraphs with either positive or negative samples based on a specific threshold (10th/90th). Threshold sampling produces distinctly positive or negative pairs; thus, we refer to it as the Easy Task. Conversely, random sampling can result in ambiguous pairs, making for a more challenging Hard Task.

Understanding of Questionnaire Item: Text Classification We convert several broadly accepted questionnaires from English into classical Chinese, including Collectivism, Individualism (Oyserman, 1993), Norm Tightness and Norm Looseness (Gelfand et al., 2011), by employing the CQC approach described in §3.1. A good classification of items from different questionnaires, e.g. “I accept the decisions made by my group” from Collectivism questionnaire, and “The best decisions are the ones I make on my own” from Individualism questionnaire, can demonstrate the model’s ability to understand and distinguish the psychological constructs represented by these questionnaires.

For both the CCR and DDR methods, all the items from these questionnaires are embedded. Then we conduct 10-fold cross-validation, using Support Vector Machines (SVM) as the classifier, and text embeddings or averaged word vectors as

features. For the prompting method, we craft a few-shot prompt (Figure 12 in the Appendix) directly asking for classification.

4.2 Psychological Measure

For the CCR method, we calculate the average cosine similarity between each paragraph in the C-HI-PSY test set and all the items in the questionnaire, representing the “loading score” of the paragraph on the questionnaire. For the DDR method, we build a corresponding dictionary for each psychological construct (Appendix C), and calculate the cosine similarity between the centroid of words in each paragraph and the centroid of words in the dictionary. For the prompting method, we craft a few-shot prompt (Figure 11 in the Appendix) to request a score from 0 to 1 for each paragraph on each questionnaire. Items in each questionnaire are provided in the prompt. Average similarities between the title of each paragraph and all the words in the dictionary, calculated by the word vector model, are used as the pseudo ground truth.

4.3 Results

For the Semantic Textual Similarity (STS) task, we evaluate the DDR and CCR methods through a rigorous process involving 20 rounds of random sampling. In each round, 4,308 random paragraph pairs are constructed from the C-HI-PSY test set. After completing these 20 evaluations, we calculate the average scores along with standard errors. When evaluating the prompting method, due to the high costs, we only conduct a single round of random sampling. For the Questionnaire Item Classification (QIC) task, we utilize 60 items from questionnaires on Collectivism, Individualism (Oyserman, 1993), Norm Tightness and Looseness (Gelfand et al., 2011), selecting 15 items from each questionnaire. For the Psychological Measure (PM) task, we measure the loading scores of all 4,308 paragraphs in the C-HI-PSY test set across the questionnaires mentioned above, and report the average scores along with standard errors.

Figure 5 illustrates that the performance metrics of almost all pre-trained models (Wang and Ren, 2022; Yan and Chi, 2020; Reimers and Gurevych, 2019; Xu, 2023) in the CCR baseline have substantially improved after fine-tuning. As shown in Table 1, the CCR method outperforms the DDR method across all tasks and surpasses the prompting method with GPT-4 (version 2024-04-09) in

most tasks, demonstrating its superiority in effectively extracting psychological variables from text.

5 Benchmarking: Traditionalism, Authority, and Attitude toward Reform

To address the lack of benchmark datasets related to psychological measurement in classical Chinese, we further validate the effectiveness of the CCR method using externally annotated data.

Officials’ Attitudes toward Reforms Moral values and political orientations are closely intertwined (Federico et al., 2013; Kivikangas et al., 2021). For example, the attitude of individuals toward reforms, policy changes, and new legislation often reflects traditionalism, conservatism, and respect for authority (Hackenburg et al., 2023; Koleva et al., 2012). Those with stronger traditionalist views are more likely to identify with the existing social order and resist changes to the status quo (Osborne et al., 2023; Jost and Hunyady, 2005).

Throughout Chinese history, there have been numerous instances of significant reforms, one of the most notable of which being the Wang Anshi’s New Policies (Anderson, 2001) in the 11th century, which faced mixed reactions from officials. We draw upon a dataset manually compiled by Wang (2022), who annotated the attitudes of 137 major officials toward the reform. Unlike data easily influenced by annotators’ subjectivity (Davani et al., 2022), officials’ attitudes toward reforms are objective and less controversial, making them highly suitable for benchmarking our pipeline.

Measure of Traditionalism and Authority We extract writings of officials documented in the *Complete Prose of the Song Dynasty* (Zeng and Liu, 2006). Questionnaires of traditionalism (Samore et al., 2023) and authority (Atari et al., 2023a) are converted from English into classical Chinese, by employing the CQC approach (§3.1). Employing the best-performing fine-tuned model, we use our CCR pipeline to measure the levels of traditionalism and authority expressed in texts. For each individual official, results are aggregated by calculating the average score of corresponding writings.

Results We find a significant negative correlation (Table 2) between officials’ attitudes toward the reforms and the levels of traditionalism and authority measured through CCR. Officials with greater traditionalism and respect for existing authority are

Framework	Base Model	Semantic Textual Similarity (Easy Task)		Semantic Textual Similarity (Hard Task)		Questionnaire Item Classification	Psychological Measure	
		Pears.	Spear.	Pears.	Spear.	Accuracy	Pears.	Spear.
		(a) DDR						
Word2Vec (CBOW)	/	.02±.11	.02±.10	-.03±.02	-.02±.01	.80±.16	.22±.07	.23±.05
Word2Vec (Skip-gram)	/	.08±.11	.09±.11	.02±.02	.02±.01	.87±.15	.18±.07	.18±.06
FastText (CBOW)	/	.05±.11	.04±.10	-.01±.01	.01±.01	.90±.13	.23±.08	.24±.06
FastText (Skip-gram)	/	.10±.10	.11±.10	.03±.02	.04±.01	.85±.16	.20±.07	.20±.05
GloVe	/	.07±.10	.09±.11	.01±.02	.01±.01	.83±.15	.16±.09	.19±.05
(b) Prompting								
GPT	GPT-3.5-turbo-0125	.08	.04	.26	.28	.63	.05±.08	.08±.10
GPT	GPT-4-turbo-2024-04-09	.63	.51	.41	.32	.77	.23±.16	.25±.15
(c) CCR (ours)								
BERT	Bert-ancient-chinese	.53±.07	.55±.07	.42±.01	.43±.01	.93±.11	.30±.04	.30±.04
RoBERTa	Guwenbert-base	.29±.07	.46±.09	.25±.01	.40±.01	.90±.11	.20±.06	.23±.09
RoBERTa	Guwenbert-large	.41±.05	.44±.07	.28±.01	.31±.01	.83±.13	.22±.04	.20±.05
SBERT	Paraphrase-multilingual-MiniLM-L12-v2	.20±.15	.21±.14	.18±.01	.19±.01	.82±.19	.15±.04	.14±.05
MacBERT+CoSENT	text2vec-base-chinese	.41±.09	.40±.09	.32±.01	.31±.01	.95±.08	.21±.10	.20±.10
ERNIE+CoSENT	text2vec-base-chinese-paraphrase	.45±.09	.45±.09	.38±.01	.37±.01	.93±.11	.21±.03	.20±.04
LERT+CoSENT	text2vec-large-chinese	.46±.12	.47±.08	.36±.01	.38±.01	.97±.07	.28±.05	.27±.05

Table 1: Performance on the test set across three tasks using three methods: DDR, LLM Prompting, and CCR.

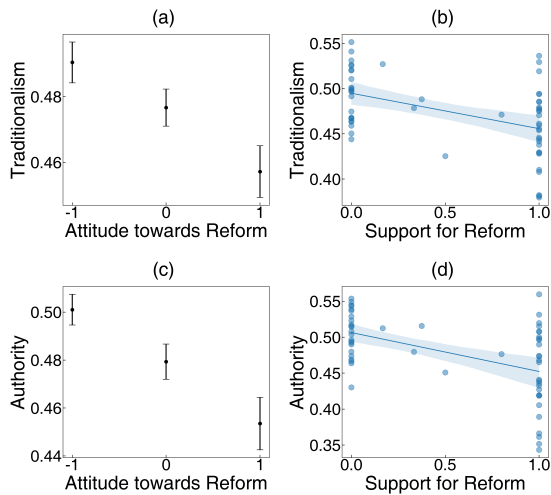


Figure 6: Correlation between traditionalism, authority and officials' attitudes toward reforms. (a) and (c) present the average psychological measure scores with standard errors, using an ordinal variable where -1 signifies opposition to the reform, 0 indicates a neutral or no explicit attitude, and 1 denotes support for the reform ($N = 108$). (b) and (d) depict the linear regression lines accompanied by 95% confidence intervals, employing a continuous variable that ranges from 0 to 1 to quantify officials' degree of support for the reform ($N = 56$).

	Support for Reform	Attitude toward Reform
Traditionalism	-0.441***	-0.279**
Authority	-0.472***	-0.310**

Table 2: Spearman correlation between CCR-based measure of moral values and actual attitude toward reform of officials. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

more likely to oppose reform (Figure 6), which is in line with the theoretical assumptions.

We also conduct Ordinary Least Squares (OLS) regressions of officials' attitudes toward reforms on traditionalism and authority. To ensure robustness, we include the officials' rank as a control variable and the fixed effects of the officials' hometowns. The results (Table 7 in the Appendix) show that psychological indicators of traditionalism and authority measured through CCR, can be used to predict the corresponding officials' attitudes toward reforms, proving the effectiveness of our method.

This benchmarking against historically verified data supports the validity of CCR as a valid computational pipeline to extract meaningful psychological information from classical Chinese corpora.

6 Discussion and Conclusion

Historical-psychological text analysis is a new line of research focused on extracting different aspects of psychology from historical corpora using state-of-the-art computational methods (Atari and Henrich, 2023). Here, we create a new pipeline, CCR, as a helpful tool for historical-psychological text analysis. Evaluating our model against word embedding models (e.g., DDR) and more recent LLMs (e.g., GPT-4), we demonstrate that CCR performs better than these alternatives while keeping its high level of interpretability and flexibility. Classical Chinese is of great historical significance, and the proposed approach can be particularly helpful in testing new insights about the “dead minds” who lived centuries or even millennia prior. We hope our tool motivates future work at the intersections of psychology, quantitative history, and NLP. Importantly, benchmarking historical-psychological tools, especially in ancient languages, is difficult because obtaining ground truth is challenging and dependent upon the quality of historical data. That said, we validate CCR against a historically verified knowledge base about attitudes toward reform and traditionalism.

Limitation

Due to the lack of fine-grained data available for training in the context of classical Chinese and with historical-psychological texts, we propose an indirect supervised learning approach where the similarities between titles are used as the pseudo ground truth for similarities between paragraphs. This approach may lead to the model learning some noise from the data, negatively affecting the model’s performance in downstream tasks, but it can save a significant amount of resources, especially for low-resource ancient languages.

Our experiments show that hard sampling is counterintuitively worse than random sampling on our dataset (Figure 4). This is the case because although the title of a text represents the main idea of most of the content, there may still be parts of the text that are unrelated to the title. For example, in a pair of paragraphs that are identified as positive samples due to their highly similar titles, one paragraph might be irrelevant to the title. Consequently, the text similarity calculated after embedding by a pre-trained model might not be high for this pair of paragraphs. The difference between the similarity prediction made by the pre-trained model and

the pseudo ground truth based on title similarity may result in these paragraph pairs being identified as hard samples. However, in such cases, the pre-trained model’s prediction could be more accurate than the pseudo ground truth derived from title similarity. It is the noise caused by the indirect supervised approach that makes the hard sampling fail to find helpful instances.

Additionally, it is worth mentioning that, due to the unavoidable survivorship bias in existing historical texts, we may only have access to texts that are physically accessible for analysis. Moreover, the psychological indicators measured from historical corpora may not represent the psychological state of the majority of people living in a specific period and location, but rather the psychological state of the specific class who wrote these documents, often the elite and intellectuals. We propose an objective measurement approach, but the results obtained from this method need to be subjected to more rigorous benchmarking and statistical analysis, taking into account population characteristics and recording efforts. The findings should be interpreted based on the domain knowledge of historians and psychologists.

Our future efforts will be directed toward assembling datasets with expert annotations to address these issues. Moreover, we aim to contribute to both historical psychology and NLP by compiling new open-source datasets for benchmarking purposes.

References

- Greg Anderson. 2001. To change china: A tale of three reformers. *Asia Pacific Perspectives*, 1(1):1–19.
- Mohammad Atari, Jonathan Haidt, Jesse Graham, Sena Koleva, Sean T. Stevens, and Morteza Dehghani. 2023a. [Morality beyond the weird: How the nomological network of morality varies across cultures](#). *Journal of Personality and Social Psychology*, 125(5):1157–1188.
- Mohammad Atari and Joseph Henrich. 2023. [Historical psychology](#). *Current Directions in Psychological Science*, 32(2):176–183.
- Mohammad Atari, Ali Omrani, and Morteza Dehghani. 2023b. [Contextualized construct representation: Leveraging psychometric scales to advance theory-driven text analysis](#).
- David Bamman and Patrick J. Burns. 2020. [Latin bert: A contextual language model for classical philology](#). *ArXiv*, abs/2009.10053.

- Nicolas Baumard, Lou Safra, Mauricio Martins, and Coralie Chevallier. 2024. [Cognitive fossils: using cultural artifacts to reconstruct psychological changes throughout history](#). *Trends in Cognitive Sciences*, 28(2):172–186.
- Damián E. Blasi, Joseph Henrich, Evangelia Adamou, David Kemmerer, and Asifa Majid. 2022. [Over-reliance on english hinders cognitive science](#). *Trends in Cognitive Sciences*, 26(12):1153–1170.
- Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. 2017. Enriching word vectors with subword information. *Transactions of the Association for Computational Linguistics*, 5:135–146.
- Ryan L Boyd, Ashwini Ashokkumar, Sarah Seraj, and James W Pennebaker. 2022. The development and psychometric properties of liwc-22. *Austin, TX: University of Texas at Austin*, pages 1–47.
- Ryan L Boyd and H Andrew Schwartz. 2021. Natural language analysis and the psychology of verbal behavior: The past, present, and future states of the field. *Journal of Language and Social Psychology*, 40(1):21–41.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.
- Aylin Caliskan, Joanna J Bryson, and Arvind Narayanan. 2017. Semantics derived automatically from language corpora contain human-like biases. *Science*, 356(6334):183–186.
- Virginia K Choi, Snehash Shrestha, Xinyue Pan, and Michele J Gelfand. 2022. When danger strikes: A linguistic tool for tracking america’s collective response to threats. *Proceedings of the National Academy of Sciences*, 119(4):e2113891119.
- S. Chopra, R. Hadsell, and Y. LeCun. 2005. [Learning a similarity metric discriminatively, with application to face verification](#). In *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR’05)*, volume 1, pages 539–546 vol. 1.
- Yung-Sung Chuang, Rumen Dangovski, Hongyin Luo, Yang Zhang, Shiyu Chang, Marin Soljagic, Shang-Wen Li, Scott Yih, Yoon Kim, and James Glass. 2022. [DiffCSE: Difference-based contrastive learning for sentence embeddings](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4207–4218, Seattle, United States. Association for Computational Linguistics.
- Aida Mostafazadeh Davani, Mark Díaz, and Vinodkumar Prabhakaran. 2022. Dealing with disagreements: Looking beyond the majority vote in subjective annotations. *Transactions of the Association for Computational Linguistics*, 10:92–110.
- Dorottya Demszky, Diyi Yang, David S Yeager, Christopher J Bryan, Margaret Clapper, Susannah Chandhok, Johannes C Eichstaedt, Cameron Hecht, Jeremy Jamieson, Meghann Johnson, et al. 2023. Using large language models in psychology. *Nature Reviews Psychology*, 2(11):688–701.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *CoRR*, abs/1810.04805.
- Christopher M. Federico, Christopher R. Weber, Damla Ergun, and Corrie Hunt. 2013. [Mapping the connections between politics and morality: The multiple sociopolitical orientations involved in moral intuition](#). *Political Psychology*, 34(4):589–610.
- Tianyu Gao, Xingcheng Yao, and Danqi Chen. 2021. [SimCSE: Simple contrastive learning of sentence embeddings](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6894–6910, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Nikhil Garg, Londa Schiebinger, Dan Jurafsky, and James Zou. 2018. Word embeddings quantify 100 years of gender and ethnic stereotypes. *Proceedings of the National Academy of Sciences*, 115(16):E3635–E3644.
- Justin Garten, Joe Hoover, Kate M Johnson, Reihane Boghrati, Carol Iskiwitch, and Morteza Dehghani. 2018. Dictionaries and distributions: Combining expert knowledge and large scale textual data content analysis: Distributed dictionary representation. *Behavior research methods*, 50:344–361.
- Michele J. Gelfand, Jana L. Raver, Lisa Nishii, Lisa M. Leslie, Janetta Lun, Beng Chong Lim, Lili Duan, Assaf Almaliach, Soon Ang, Jakobina Arndottir, Zeynep Aycan, Klaus Boehnke, Pawel Boski, Rosa Cabecinhas, Darius Chan, Jagdeep Chhokar, Alessia D’Amato, Montserrat Subirats Ferrer, Iris C. Fischlmayr, Ronald Fischer, Marta Fülöp, James Georgas, Emiko S. Kashima, Yoshishima Kashima, Kibum Kim, Alain Lempereur, Patricia Marquez, Rozhan Othman, Bert Overlaet, Penny Panagiotopoulou, Karl Peltzer, Lorena R. Perez-Florizno, Larisa Ponomarenko, Anu Realo, Vidar Schei, Manfred Schmitt, Peter B. Smith, Nazar Soomro, Erna Szabo, Naline Taveesin, Midori Toyama, Evert Van de Vliert, Naharika Vohra,

- Colleen Ward, and Susumu Yamaguchi. 2011. [Differences between tight and loose cultures: A 33-nation study](#). *Science*, 332(6033):1100–1104.
- Jesse Graham, Jonathan Haidt, and Brian A Nosek. 2009. Liberals and conservatives rely on different sets of moral foundations. *Journal of personality and social psychology*, 96(5):1029.
- Kobi Hackenburg, William J Brady, and Manos Tsakiris. 2023. Mapping moral language on us presidential primary campaigns reveals rhetorical networks of political division and unity. *PNAS nexus*, page pgad189.
- Hangfeng He, Mingyuan Zhang, Qiang Ning, and Dan Roth. 2021. [Foreseeing the benefits of incidental supervision](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 1782–1800, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Joshua Conrad Jackson, Joseph Watts, Johann-Mattis List, Curtis Puryear, Ryan Drabble, and Kristen A. Lindquist. 2021. [From text to thought: How analyzing language can advance psychological science](#). *Perspectives on Psychological Science*, 17(3):805–826.
- Kyle P. Johnson, Patrick J. Burns, John Stewart, Todd Cook, Clément Besnier, and William J. B. Mattingly. 2021. [The Classical Language Toolkit: An NLP framework for pre-modern languages](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing: System Demonstrations*, pages 20–29, Online. Association for Computational Linguistics.
- John T Jost and Orsolya Hunyady. 2005. Antecedents and consequences of system-justifying ideologies. *Current directions in psychological science*, 14(5):260–265.
- Swanie Juhng, Matthew Matero, Vasudha Varadarajan, Johannes Eichstaedt, Adithya V Ganesan, and H Andrew Schwartz. 2023. Discourse-level representations can improve prediction of degree of anxiety. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 1500–1511.
- Brendan Kennedy, Mohammad Atari, Aida Mostafazadeh Davani, Joe Hoover, Ali Omrani, Jesse Graham, and Morteza Dehghani. 2021. Moral concerns are differentially observable in language. *Cognition*, 212:104696.
- J Matias Kivikangas, Belén Fernández-Castilla, Simo Järvelä, Niklas Ravaja, and Jan-Erik Lönnqvist. 2021. Moral foundations and political orientation: Systematic review and meta-analysis. *Psychological Bulletin*, 147(1):55.
- Spassena P Koleva, Jesse Graham, Ravi Iyer, Peter H Ditto, and Jonathan Haidt. 2012. Tracing the threads: How five moral concerns (especially purity) help explain culture war attitudes. *Journal of research in personality*, 46(2):184–194.
- Zhou Liu, Hongsu Wang, and Peter K Bol. 2023. Automatic biographical information extraction from local gazetteers with bi-lstm-crf model and bert. *International Journal of Digital Humanities*, 4(1-3):195–212.
- Enrique Manjavacas Arevalo and Lauren Fonteyn. 2021. [MacBERTh: Development and evaluation of a historically pre-trained language model for English \(1450-1950\)](#). In *Proceedings of the Workshop on Natural Language Processing for Digital Humanities*, pages 23–36, NIT Silchar, India. NLP Association of India (NLPAl).
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. [Efficient estimation of word representations in vector space](#). In *1st International Conference on Learning Representations, ICLR 2013, Scottsdale, Arizona, USA, May 2-4, 2013, Workshop Track Proceedings*.
- Michael Muthukrishna, Joseph Henrich, and Edward Slingerland. 2021. [Psychology as a historical science](#). *Annual Review of Psychology*, 72(1):717–749.
- Gandalf Nicolas, Xuechunzi Bai, and Susan T Fiske. 2021. Comprehensive stereotype content dictionaries using a semi-automated method. *European Journal of Social Psychology*, 51(1):178–196.
- Danny Osborne, Thomas H. Costello, John Duckitt, and Chris G. Sibley. 2023. [The psychological causes and societal consequences of authoritarianism](#). *Nature Reviews Psychology*, 2(4):220–232.
- Daphna Oyserman. 1993. [The lens of personhood: Viewing the self and others in a multicultural society](#). *Journal of Personality and Social Psychology*, 65(5):993–1009.
- 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, Doha, Qatar. Association for Computational Linguistics.
- Fanchao Qi, Yanhui Yang, Jing Yi, Zhili Cheng, Zhiyuan Liu, and Maosong Sun. 2022. [QuoteR: A benchmark of quote recommendation for writing](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 336–348, Dublin, Ireland. Association for Computational Linguistics.
- Fanchao Qi, Lei Zhang, Yanhui Yang, Zhiyuan Liu, and Maosong Sun. 2020. Wantwords: An open-source online reverse dictionary system. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 175–181.

- 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.
- Theodore Samore, Daniel M. T. Fessler, Adam Maxwell Sparks, Colin Holbrook, Lene Aarøe, Carmen Gloria Baeza, María Teresa Barbato, Pat Barclay, Renatas Berniūnas, Jorge Contreras-Garduño, Bernardo Costa-Neves, Maria del Pilar Grazioso, Pınar Elmas, Peter Fedor, Ana Maria Fernandez, Regina Fernández-Morales, Leonel Garcia-Marques, Paulina Giraldo-Perez, Pelin Gul, Fanny Habacht, Youssef Hasan, Earl John Hernandez, Tomasz Jarmakowski, Shanmukh Kamble, Tatsuya Kameda, Bia Kim, Tom R. Kupfer, Maho Kurita, Norman P. Li, Junsong Lu, Francesca R. Luberti, María Andréa Maegli, Marínés Mejía, Coby Morvinski, Aoi Naito, Alice Ng’ang’a, Angélica Nascimento de Oliveira, Daniel N. Posner, Pavol Prokop, Yaniv Shani, Walter Omar Paniagua Solorzano, Stefan Stieger, Angela Oktavia Suryani, Lynn K. L. Tan, Joshua M. Tybur, Hugo Viciana, Amandine Visine, Jin Wang, and Xiao-Tian Wang. 2023. [Greater traditionalism predicts covid-19 precautionary behaviors across 27 societies](#). *Scientific Reports*, 13(1).
- Florian Schroff, Dmitry Kalenichenko, and James Philbin. 2015. Facenet: A unified embedding for face recognition and clustering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 815–823.
- Indira Sen, Daniele Quercia, Marios Constantinides, Matteo Montecchi, Licia Capra, Sanja Scepanovic, and Renzo Bianchi. 2022. Depression at work: exploring depression in major us companies from online reviews. *Proceedings of the ACM on Human-Computer Interaction*, 6(CSCW2):1–21.
- Chenglei Si, Zhe Gan, Zhengyuan Yang, Shuohang Wang, Jianfeng Wang, Jordan Boyd-Graber, and Lijuan Wang. 2023. [Prompting gpt-3 to be reliable](#). In *International Conference on Learning Representations (ICLR)*.
- Almog Simchon, William J Brady, and Jay J Van Bavel. 2022. Troll and divide: the language of online polarization. *PNAS nexus*, 1(1):pgac019.
- Edward Slingerland. 2013. Body and mind in early china: An integrated humanities–science approach. *Journal of the American Academy of Religion*, 81(1):6–55.
- Edward Slingerland, Ryan Nichols, Kristoffer Neilbo, and Carson Logan. 2017. The distant reading of religious texts: A “big data” approach to mind-body concepts in early china. *Journal of the American Academy of Religion*, 85(4):985–1016.
- Daniel Swanson and Francis Tyers. 2022. [Handling stress in finite-state morphological analyzers for Ancient Greek and Ancient Hebrew](#). In *Proceedings of the Second Workshop on Language Technologies for Historical and Ancient Languages*, pages 108–113, Marseille, France. European Language Resources Association.
- Huishuang Tian, Kexin Yang, Dayiheng Liu, and Jiancheng Lv. 2021. Anchibert: A pre-trained model for ancient chinese language understanding and generation. In *2021 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8. IEEE.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Dongbo Wang, Chang Liu, Zhixiao Zhao, Si Shen, Liu Liu, Bin Li, Haotian Hu, Mengcheng Wu, Litao Lin, Xue Zhao, and Xiyu Wang. 2023a. [Gujibert and gujigpt: Construction of intelligent information processing foundation language models for ancient texts](#). *Preprint*, arXiv:2307.05354.
- Jiahui Wang, Xuqin Zhang, Jiahuan Li, and Shujian Huang. 2023b. [Pre-trained model in Ancient-Chinese-to-Modern-Chinese machine translation](#). In *Proceedings of ALT2023: Ancient Language Translation Workshop*, pages 23–28, Macau SAR, China. Asia-Pacific Association for Machine Translation.
- Pengyu Wang and Zhichen Ren. 2022. The uncertainty-based retrieval framework for ancient chinese cws and pos. In *Proceedings of the Second Workshop on Language Technologies for Historical and Ancient Languages*, pages 164–168.
- Sze-Yuh Nina Wang and Yoel Inbar. 2021. Moral-language use by us political elites. *Psychological Science*, 32(1):14–26.
- Yuhua Wang. 2022. [Blood is thicker than water: Elite kinship networks and state building in imperial china](#). *American Political Science Review*, 116(3):896–910.
- John Wilkerson and Andreu Casas. 2017. Large-scale computerized text analysis in political science: Opportunities and challenges. *Annual Review of Political Science*, 20:529–544.
- Jiashu Xu, Mingyu Derek Ma, and Muhao Chen. 2023a. [Can NLI provide proper indirect supervision for low-resource biomedical relation extraction?](#) In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2450–2467, Toronto, Canada. Association for Computational Linguistics.
- Mengyao Xu, Lingshu Hu, and Glen T Cameron. 2023b. Tracking moral divergence with ddr in presidential debates over 60 years. *Journal of Computational Social Science*, 6(1):339–357.

- Ming Xu. 2023. Text2vec: Text to vector toolkit. <https://github.com/shibing624/text2vec>.
- Tan Yan and Zewen Chi. 2020. Guwenbert. <https://github.com/ethan-yt/guwenbert>.
- Tal Yarkoni and Jacob Westfall. 2017. Choosing prediction over explanation in psychology: Lessons from machine learning. *Perspectives on Psychological Science*, 12(6):1100–1122.
- Wenpeng Yin, Muhao Chen, Ben Zhou, Qiang Ning, Kai-Wei Chang, and Dan Roth. 2023. [Indirectly supervised natural language processing](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 6: Tutorial Abstracts)*, pages 32–40, Toronto, Canada. Association for Computational Linguistics.
- Tariq Yousef, Chiara Palladino, David J. Wright, and Monica Berti. 2022. [Automatic translation alignment for Ancient Greek and Latin](#). In *Proceedings of the Second Workshop on Language Technologies for Historical and Ancient Languages*, pages 101–107, Marseille, France. European Language Resources Association.
- Zaozhuang Zeng and Lin Liu, editors. 2006. *Complete Prose of the Song Dynasty*, volume 360. Shanghai cishu chubanshe and Anhui jiaoyu chubanshe, Shanghai and Hefei. In Chinese.
- Lei Zhang, Fanchao Qi, Zhiyuan Liu, Yasheng Wang, Qun Liu, and Maosong Sun. 2020. Multi-channel reverse dictionary model. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 312–319.
- Ying Zhong, Valentin Thouzeau, and Nicolas Baumard. 2023. [The evolution of romantic love in chinese fiction in the very long run \(618 - 2022\): A quantitative approach](#). In *Workshop on Computational Humanities Research*.
- Bo Zhou, Qianglong Chen, Tianyu Wang, Xiaomi Zhong, and Yin Zhang. 2023. [WYWEB: A NLP evaluation benchmark for classical Chinese](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 3294–3319, Toronto, Canada. Association for Computational Linguistics.
- Ke Zhou, Luca Maria Aiello, Sanja Scepanovic, Daniele Quercia, and Sara Konrath. 2021. The language of situational empathy. *Proceedings of the ACM on Human-Computer Interaction*, 5(CSCW1):1–19.

A Historical Psychology Corpus Details

A.1 Distribution of Paragraph Lengths

To ensure the inclusion of sufficient semantic information, paragraphs containing fewer than 50 characters have been merged with the preceding paragraph of the article or chapter, wherever possible. To accommodate the token limitations of models such as BERT, paragraphs that exceed 500 characters have been divided into segments with fewer than 500 characters each, while maintaining the integrity of the original sentence structure as much as possible. The average length of paragraphs is 195 characters.




Figure 7: Distributions of paragraph lengths in different sets.

A.2 Distribution of Title Similarities




Figure 8: Distribution of title similarities with thresholds.

B Word Embedding Model Details

B.1 Pre-processing

Before training the word vector model, we conducted word segmentation on the corpus, employing the pre-trained tokenizer “COARSE_ELECTRA_SMALL_ZH” from HanLP (<https://hanlp.hankcs.com/docs/api/hanlp/pretrained/tok.html>).

After word segmentation, the corpus consists of 1.04 billion word tokens and an initial vocabulary containing 15.55 million unique words. By truncating the vocabulary at a minimum word count threshold of 10, the final vocabulary size is reduced to 1.27 million words.

B.2 Training Hyperparameters

We train our word vector models on the same corpus using various frameworks and architectures, such as Word2Vec (with CBOW and Skip-gram) (Mikolov et al., 2013), FastText (with CBOW and Skip-gram) (Bojanowski et al., 2017), and GloVe (Pennington et al., 2014). The hyperparameters are presented in Table 4.

C Dictionary Details

We build a dictionary for each classical Chinese questionnaire by using an open-source dictionary system named “WantWords” (Qi et al., 2020), which is based on a multi-channel reverse dictionary model (MRDM) (Zhang et al., 2020) and takes sentences (descriptions of words) as input and yields words semantically matching the input sentences.

The process involves three steps: (1) we employ the “WantWords” model to obtain the top n most similar words to each sentence in the questionnaire; (2) a process of deduplication is then conducted; (3) the words are labeled manually by a native Chinese speaker with “relevant” or “irrelevant” to the corresponding topic, after which all irrelevant words are discarded.

Text	Title	Source
<p>臣之所聞曰：「臣事君，子事父，妻事夫，三者順則天下治，三者逆則天下亂，此天下之常道也，明王賢臣而弗易也。」則人主雖不肖，臣不敢侵也。</p> <p>It is said: "When a subject serves his ruler, a son serves his father, and a wife serves her husband, if all these are in harmony, then the world is in order. If they are contrary, the world is in chaos. This is the eternal way of the world, and wise kings and virtuous ministers do not deviate from it." Even if a ruler is unworthy, a subject dares not encroach.</p>	忠孝 Loyalty and Filial Piety	韓非子 <i>Han Feizi</i>
<p>君子謂節女仁孝厚於恩義也。夫重仁義輕死亡，行之高者也。《論語》曰：「君子殺身以成仁，無求生以害仁。」此之謂也。頌曰：京師節女，夫讎劫父，要女間之，不敢不許，期處既成，乃易其所，殺身成仁，義冠天下。</p> <p>A gentleman considers a woman of virtue and righteousness to hold loyalty and kindness above all else. To value righteousness and morality over life and death is the act of the noble. Confucius said: "A gentleman would die to achieve righteousness, without seeking to live at the expense of righteousness." This is what it means. The ode says: "In the capital, a woman of principle, when bandits threatened her father, demanded her in exchange, she could not refuse. When the arrangement was made, she died to uphold righteousness, her virtue renowned under the heavens."</p>	節義 Virtue and Righteousness	列女傳 <i>Biographies of Exemplary Women</i>
<p>康僧淵在豫章，去郭數十里，立精舍。旁連嶺，帶長川，芳林列於軒庭，清流激於堂宇。乃閒居研講，希心理味，庾公諸人多往看之。觀其運用吐納，風流轉佳。加已處之怡然，亦有以自得，聲名乃興。後不堪，遂出。</p> <p>Kang Sengyuan in Yuzhang, some tens of li away from the settlement, established a hermitage. Beside a ridge, along a long river, fragrant forests lined the pavilion, and clear streams rushed by the hall. He then lived in seclusion to study and lecture, seeking the essence of philosophy, attracting many, including Yu Gong, to visit. Observing his breathing techniques and the transformation of his demeanor, his grace and elegance stood out. Living in such contentment, he found fulfillment.</p>	棲逸 Hermitage	世說新語 <i>A New Account of the Tales of the World</i>

Table 3: Sample entries from the Chinese Historical Psychology Corpus (C-HI-PSY).

Framework	Architecture	Vector Size	Epoch	Window Size	Other Parameters
Word2Vec	CBOW	300	5	5	negative=5
	Skip-gram	300	5	5	negative=5
FastText	CBOW	300	5	5	negative=5, min_n=1, max_n=4
	Skip-gram	300	5	5	negative=5, min_n=1, max_n=4
GloVe		300	15	5	x_max=100, alpha=0.75

Table 4: Word vector model training hyperparameters and evaluation results.

Sampling	Positive/Negative Sampling Thresholds	{(10th, 90th)}
	Triplet Sampling Option	{random}
	Sampling Seed	{42}
Training	Batch Size	{16, 32}
	Epochs	{3}
	Warmup Epochs	{1, 2, 3}
	Learning Rate	{1e-6, 1e-5, 2e-5}
	Optimizer	{Adam}

Table 5: Hyperparameter sweep for triplet sampling and validation for fine-tuned models.

Framework	Base Model	Specific to Classical Chinese	Batch Size	Warmup Epochs	Learning Rate	Pearson	Spearman
BERT	bert-ancient-chinese	✓	32	3	1.0e-05	.43	.42
RoBERTa	guwenbert-base	✓	32	2	2.0e-05	.30	.37
	guwenbert-large	✓	16	1	2.0e-05	.29	.30
SBERT	paraphrase-multilingual-MiniLM-L12-v2	✗	32	1	2.0e-05	.19	.19
MacBERT+CoSENT	text2vec-base-chinese	✗	32	2	2.0e-05	.34	.32
ERNIE+CoSENT	text2vec-base-chinese-paraphrase	✗	32	2	2.0e-05	.40	.40
LERT+CoSENT	text2vec-large-chinese	✗	16	2	2.0e-05	.36	.37

Table 6: Fine-tuned models’ performance on the validation set. We show the best performing configuration which is also the final configuration used to report each models’ performance on the test test.

	Model 1	Model 2	Model 3	Model 4
Traditionalism	-5.290** (1.743)	-6.170*** (1.568)		
Authority			-4.518** (1.320)	-5.207*** (1.194)
Official Rank	0.073 (0.074)	0.011 (0.065)	0.102 (0.074)	0.031 (0.062)
Province Fixed Effects	Yes	Yes	Yes	Yes
R-squared	0.292	0.542	0.31	0.571
Observations	108	56	108	56

Table 7: Ordinary Least Squares (OLS) regressions of officials towards reforms on CCR-based measures of traditionalism and authority, with control variable of official rank and fixed effects for officials’ hometown provinces. Models 1 and 3 use an ordinal variable for officials’ attitudes towards reform, where -1 signifies opposition to the reform, 0 indicates a neutral or no explicit stance, and 1 denotes support for the reform. Models 2 and 4 use a continuous variable that ranges from 0 to 1 to quantify the degree of officials’ support for the reform. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

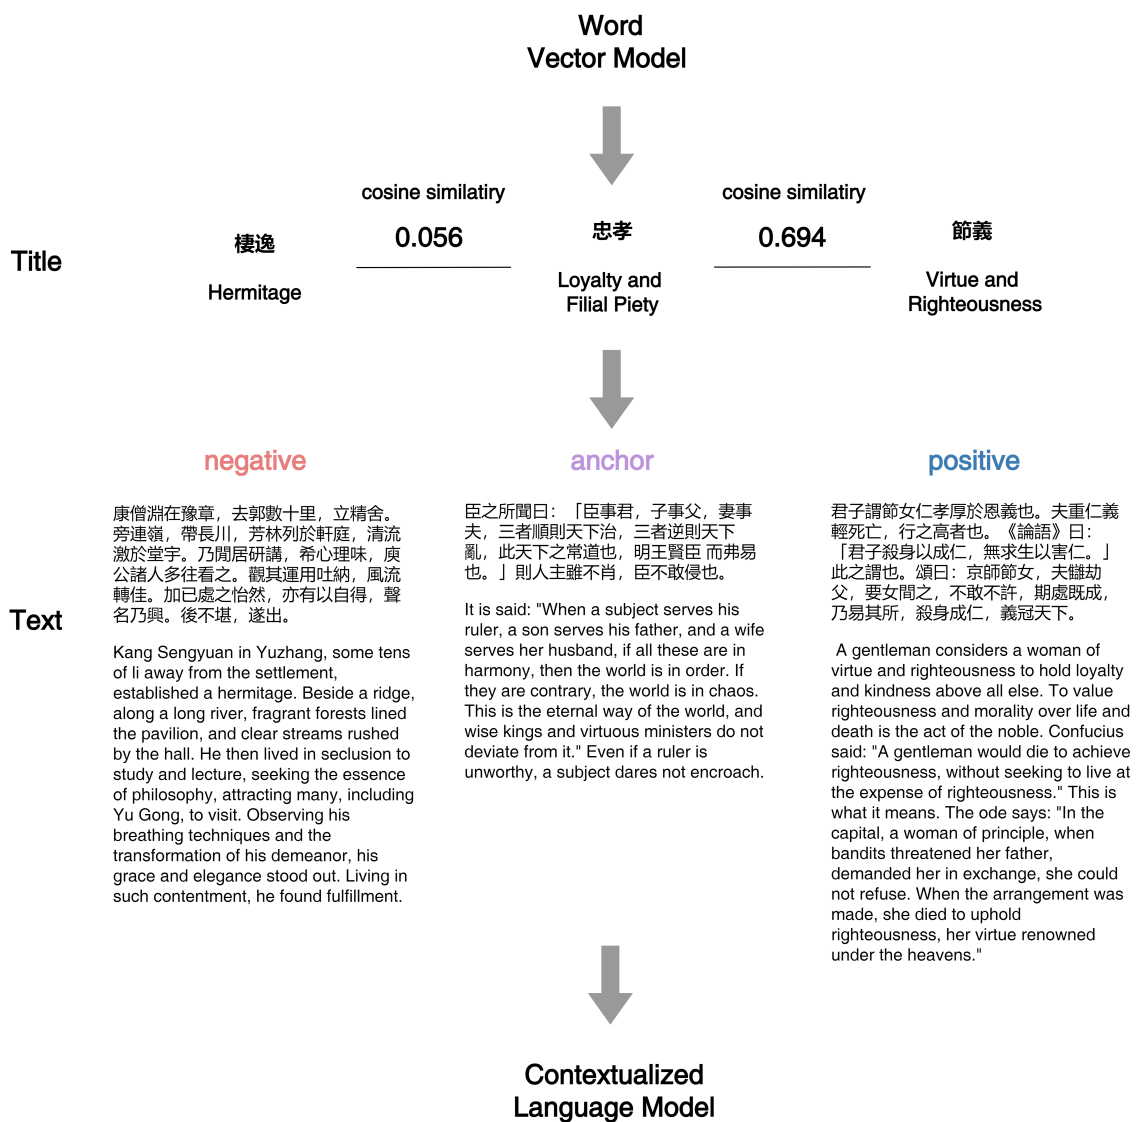


Figure 9: An example of the indirect supervised contrastive learning.

Instructions

On a scale ranging from 0 to 1, how similar are the following two paragraphs with respect to moral values? Respond only with a score ranging from 0 to 1.

Examples

Here are two examples:

Input:
Paragraph 1: 蘇應喜，正安州人，年十八，母劉氏。道光乙酉，東街火延及西街，民居殆盡。喜方在書院肄業，聞報奔回，不問物，惟尋母耗。不得，既而聞火中哭聲，喜知是母，急入救。眾以火猛，入必死，挽之，喜哭曰：「天下豈有無母之子哉？」奮身入救，死之。後灰燼中見喜覆母，母通身焦黑，而喜面如生。(Su Yingxi, from Zheng'an County, was eighteen years old and his mother was Mrs. Liu. During the year of Yiyou in the Daoguang era, a fire that started on the East Street spread to the West Street, nearly destroying all the homes there. Yingxi was studying at the academy at the time and rushed home upon hearing the news. He didn't bother with his belongings; his only concern was to find his mother. Unable to find her initially, he then heard crying from within the flames and knew it was his mother. Despite the crowd warning him that entering the fierce fire would lead to certain death, Yingxi, in tears, exclaimed, "How can there be a son without a mother in this world?" He bravely entered the fire to save her, sacrificing his own life in the process. Later, amidst the ashes, Yingxi was found covering his mother; her body was completely charred, yet Yingxi's face appeared as if he were still alive.)
Paragraph 2: 十歲喪父，事母克盡孝道。母病，藥餌無效，振乃焚香鑪天，割股肉調羹以進，母病遂痊。事聞，縣令旌曰真孝格天。(At the age of ten, after losing his father, he devoted himself fully to serving his mother, fulfilling his filial duties. When his mother fell ill and no medicine proved effective, he burned incense to pray to the heavens and then cut flesh from his own leg to make a soup for her. After consuming this, his mother's health miraculously recovered. This act of filial piety was so remarkable that it was reported to the local authorities, who commended him, declaring that his genuine filial piety moved the heavens.)

Output:
0.74346477

Input:
Paragraph 1: 蘇應喜，正安州人，年十八，母劉氏。道光乙酉，東街火延及西街，民居殆盡。喜方在書院肄業，聞報奔回，不問物，惟尋母耗。不得，既而聞火中哭聲，喜知是母，急入救。眾以火猛，入必死，挽之，喜哭曰：「天下豈有無母之子哉？」奮身入救，死之。後灰燼中見喜覆母，母通身焦黑，而喜面如生。(Su Yingxi, from Zheng'an County, was eighteen years old and his mother was Mrs. Liu. During the year of Yiyou in the Daoguang era, a fire that started on the East Street spread to the West Street, nearly destroying all the homes there. Yingxi was studying at the academy at the time and rushed home upon hearing the news. He didn't bother with his belongings; his only concern was to find his mother. Unable to find her initially, he then heard crying from within the flames and knew it was his mother. Despite the crowd warning him that entering the fierce fire would lead to certain death, Yingxi, in tears, exclaimed, "How can there be a son without a mother in this world?" He bravely entered the fire to save her, sacrificing his own life in the process. Later, amidst the ashes, Yingxi was found covering his mother; her body was completely charred, yet Yingxi's face appeared as if he were still alive.)
Paragraph 2: 東坡在儋耳，語其子過曰：「我決不為海外人。近日頗覺有還中州氣象。」乃澹研焚香，寫平日所作八賦，當不脫誤一字以卜之。寫畢，大喜曰：「吾歸無疑矣。」後數日，廣州之命至。八賦鏗勁初歸梁師成，後入禁中。(While Su Dongpo was in Danzhou, he told his son that he was determined not to remain a man of the overseas (remote regions) forever, feeling a strong premonition of returning to the Central Plains. To ascertain this, he meticulously cleaned and burned incense before copying eight essays he had written, believing that not making a single mistake in transcription would be an auspicious sign. Upon completing the task without error, he joyfully proclaimed his certain return. A few days later, an official appointment from Lianzhou arrived. The manuscripts of the eight essays were initially given to Liang Shicheng, and later they found their way into the imperial collection.)

Output:
0.043439843

Figure 10: Few-shot prompt for the semantic textual similarity task.

Instructions

The Collectivism questionnaire is as follows: 1.人人親其親，長其長，而天下平。2.教以慈睦，而民貴有親。3.忠孝友悌，正己化人，矜孤恤寡，敬老懷幼。4.君仁臣忠，父慈子孝，兄愛弟敬，夫和妻柔，姑慈婦聽，禮之至也。5.事父母者莫善於順，宜兄弟者莫善於讓，故順，孝之貴也；讓，友之本也。(1.When everyone cherishes their own relatives and respects their elders, peace will prevail throughout the world. 2.Teach with kindness and harmony, and people will value kinship. 3.Be loyal and show filial respect, maintain integrity and influence others positively, care for orphans and widows, and respect the elderly while cherishing the young. 4.When rulers are benevolent and ministers are loyal, fathers are kind and sons are filial, brothers love and younger siblings respect, husbands and wives harmonize, and mothers-in-law are kind and daughters-in-law are obedient, such is the epitome of proper conduct. 5.Serving one's parents, nothing is better than obedience; among siblings, nothing is better than yielding. Therefore, obedience is the essence of filial piety; yielding is the foundation of brotherly love.)

The Individualism questionnaire is as follows: 1.餘凡事獨出己見，不問他人是非。2.人生何必同？要在有所立。3.冤頭獨宿今不與取同。4.人須有自信之能力，當從自己良心上認定是非，不可以眾人之是非為從違。5.凡事須先求自己，臨場何必仰他人。(1.I generally prefer to express my own opinions on everything, disdainful of following others in matters of right and wrong. 2.Why must life be uniform? What matters is establishing oneself. 3.Alone I sleep with a graceful neck, not conforming to the crowd. 4.One must have the ability to be self-confident, to discern right from wrong based on one's own conscience, not swayed by the opinions of the masses. 5.In all things, one should prioritize oneself; why should one look to others in the moment?)

The Tightness questionnaire is as follows: 1.不設規矩，無以順人；不切刑罰，無以息暴。2.法律政令者，吏民規矩繩墨也。3.峻法，所以凌過遊外私也；嚴刑，所以遂令慙下也。4.刑一而正百，殺一而慎萬。5.以刑止刑，以殺止殺。(1.Without rules and regulations, there can be no guiding of people; without strict punishments, there can be no quelling of violence. 2.Laws and decrees serve as the guidelines and standards for officials and the people. 3.Harsh laws are in place to curb excessive and private misconduct; severe punishments are to ensure that orders are obeyed and the lower ranks are disciplined. 4.Punish one to correct a hundred, execute one to caution ten thousand. 5.Use punishment to stop further offenses, use execution to deter further killings.)

The Looseness questionnaire is as follows: 1.人不擾，政己和，世慮寡，山情多。2.我清靜則民自正。3.無為，無我，無欲，居下，清虛，自然。4.思想自由，為凡百自由之母。5.我無為，人自寧。(1.When the people are undisturbed, governance is harmonious. With few worldly concerns, there is an abundance of natural affection. 2.When I am calm and tranquil, the people will correct themselves. 3.By taking no action, having no ego, desiring nothing, staying humble, embracing emptiness, and being natural. 4.Freedom of thought is the mother of freedom for all. 5.By my non-action, the people will naturally find peace.)

On a scale ranging from 0 to 1, what score would the following text receive in these four questionnaires respectively? Respond only with four scores, ranging from 0 to 1, in the same order as the questionnaires.

Examples

Here are two examples:

Input:
 龔詠，字橫道，子述之後，隱居好水，不求聞達。以道誼相尚，朝野鉅公，聞人多就。見稱，或勸其出，答曰：吾讀祖父書，樂堯舜道，平生之志漸將矣。出處之際，非吾所固必也。一舉不合，遂自隱，後辟，皆不就。所著有《昌江風土記》、《道水經》、《楊子史》、《尤長於春秋》、《致明放情》、《有知後學》、《我於嚴穴》、《可謂博雅君子》、《伴七世孫》、《性孝友》、《好學不爭利》、《放情山水》、《有司以寬良鳴，辭弗就》。(Nie Zhu, also known as Ji Dao, was a descendant of Ji Shu who chose seclusion by the Xu River over public life, valuing moral integrity over fame. He was revered for his adherence to the virtuous paths of ancient sages Yao and Shun and preferred a life of simplicity and study to holding office. Despite being sought after by officials and scholars for his profound understanding of the classics, especially the "Spring and Autumn Annals," and his contributions to scholarship, Nie Zhu declined all offers of official positions. He authored "The Geography and Customs of Changjiang," and lived his later years as a hermit, embodying the qualities of a learned and noble gentleman without pursuing wealth or status.)

Output:
 0.23639056, 0.33721533, 0.12832133, 0.4139657

Input:
 濟南西關有約承基者，年十五，父歿，祖母尚存，年八十矣。家貧，自知祖孫難以存活，因備於修造工程處，日得工錢三百文以養祖母。(In the Xiguan area of Jinan, there was a young man named Yue Chengji, who was fifteen years old when his father passed away. His grandmother, still living, was eighty years old. Facing poverty and aware of the difficulty in supporting himself and his grandmother, he took up work at a construction site, earning 300 Wen a day to provide for his grandmother.)

Output:
 0.37603232, 0.28392938, 0.1424624, 0.0230375

Figure 11: Few-shot prompt for the psychological measure task.

Instructions

There are four topics that are related to moral values: "Collectivism", "Individualism", "Tightness", and "Looseness". Respond with the topic that best fits the following paragraph.

Examples

Here are two examples:

Input:

忠孝友悌，正己化人，矜孤恤寡，敬老懷幼。

(Be loyal and show filial respect, maintain integrity and influence others positively, care for orphans and widows, and respect the elderly while cherishing the young.)

Output: Collectivism

Input:

禮者禁於將然之前，而法者禁於已然之後。

(Rituals prevent misconduct before it happens, while laws address misconduct after it has occurred.)

Output: Tightness

Figure 12: Few-shot prompt for the questionnaire item classification task.

Traditionalism

規矩者，方圓之正也。

The square and the compass are the standards of straightness and roundness.

禮之于正國也，猶衡之於輕重也；繩墨之于曲直也，規矩之於方圓也。

In governing a state properly, rites are like a scale in measuring weight, a plumb line in discerning straightness or crookedness, and a square or compass in defining squareness or roundness.

國藉十世之基，家承百年之業，士食舊德之名氏，農服先疇之畝畝，商循族世之所鬻，工用高曾之規矩。

A nation relies on the foundations built over decades, a family inherits centuries of enterprise, scholars are known by the virtues of their ancestors, farmers work the fields of their forefathers, merchants follow the trade of their clan, and craftsmen adhere to the standards passed down from their predecessors.

法古無過，循禮無邪。

Ancient laws have no faults, following rites has no evil.

傳前世之法則，作後世之規矩。

Transmit the laws of the former generation as the standards for the next.

教之而信，必先有令人愛者；禁之而畏，必先有令人敬者。

To teach and be believed, one must first present what is dependable; to forbid and be feared, one must first present what is respectable.

Authority

國將興，必貴師而重傳；貴師而重傳，則法度存。

For a nation to prosper, it must value teachers and honor mentors; by valuing teachers and honoring mentors, the norms remain intact.

稱尊長，勿呼名。對尊長，勿見能。

Respect elders, do not call them by their names. When facing elders, do not show off your abilities.

祖宗雖遠，祭祀不可不誠；子孫雖愚，經書不可不讀。

Though ancestors are long gone, sacrifices to them must be conducted sincerely; though descendants may be unwise, the classics must not go unread.

先師遺訓，餘豈之墜？

How can I neglect the teachings left by my former teacher?

傳前世之法則，作後世之規矩。

Transmit the laws of the former generation as the standards for the next.

古者人君持權於上，而內外莫敢不肅。

In ancient times, sovereigns held authority supreme, and none dared to be disrespectful.

Figure 14: Authority questionnaire in classical Chinese.

Figure 13: Traditionalism questionnaire in classical Chinese.