

LLM vs Small Model? Large Language Model Based Text Augmentation Enhanced Personality Detection Model

Linmei Hu^{1*}, Hongyu He², Duokang Wang², Ziwang Zhao², Yingxia Shao², Liqiang Nie³

¹School of Computer Science and Technology, Beijing Institute of Technology

²School of Computer Science, Beijing University of Posts and Telecommunications

³School of Computer Science and Technology, Harbin Institute of Technology (Shenzhen)

hulinmei@bit.edu.cn, {hehongyu, wangduokang, zhaoziwang, shaoyx}@bupt.edu.cn, nieliqiang@gmail.com

Abstract

Personality detection aims to detect one’s personality traits underlying in social media posts. One challenge of this task is the scarcity of ground-truth personality traits which are collected from self-report questionnaires. Most existing methods learn post features directly by fine-tuning the pre-trained language models under the supervision of limited personality labels. This leads to inferior quality of post features and consequently affects the performance. In addition, they treat personality traits as one-hot classification labels, overlooking the semantic information within them. In this paper, we propose a large language model (LLM) based text augmentation enhanced personality detection model, which distills the LLM’s knowledge to enhance the small model for personality detection, even when the LLM fails in this task. Specifically, we enable LLM to generate post analyses (augmentations) from the aspects of semantic, sentiment, and linguistic, which are critical for personality detection. By using contrastive learning to pull them together in the embedding space, the post encoder can better capture the psycho-linguistic information within the post representations, thus improving personality detection. Furthermore, we utilize the LLM to enrich the information of personality labels for enhancing the detection performance. Experimental results on the benchmark datasets demonstrate that our model outperforms the state-of-the-art methods on personality detection.

Introduction

Personality is a combination of a person’s internal characteristics, which can be reflected in their behavior (Stajner and Yenikent 2020). Personality traits are defined by long-established personality theories such as the Myers-Briggs Type Indicator (MBTI) taxonomy (Myers-Briggs 1991). Personality detection aims to detect one’s personality traits underlying in social media posts. This emerging task in computational psycho-linguistics can provide specific personality information about the person and has been utilized in various applications, such as dialogue systems (Wen et al. 2021; Yang, Chen, and Narasimhan 2021) and psychological treatments (Bagby et al. 2016).

Traditional personality detection methods mostly rely on manually designed features, such as Linguistic Inquiry and

Word Count (LIWC) (Tausczik and Pennebaker 2010), that bring psychology theories and linguistic methods together. Recent works use end-to-end deep neural network to obtain text representations automatically in a data-driven manner (Jiang, Zhang, and Choi 2020; Keh and Cheng 2019; Lynn, Balasubramanian, and Schwartz 2020; Gjurković et al. 2021). However, extracting effective personality information from online posts is a non-trivial task. One main challenge is the scarcity of ground-truth personality traits, as they are collected from self-report questionnaires, which is time-consuming and often raises concerns related to user privacy. To overcome the issue, pre-trained language models have been applied to learn post representations (Yang et al. 2021b; Xue et al. 2018; Yang et al. 2021b). To further improve the post representations, TrigNet (Yang et al. 2021b) considers the contextual information of a user’s posts as well as the external psycho-linguistic knowledge from LIWC (Xue et al. 2018). However, the learned post representations are still unsatisfactory, resulting in inferior performance. In addition, they treat personality traits as one-hot classification labels, overlooking the semantic information within them.

Recently, Large language models (LLMs) (Brown et al. 2020; Sun et al. 2021; Ouyang et al. 2022; Zhang et al. 2022) have demonstrated significant capability in various natural language processing tasks under a generative format in zero-shot or few-shot scenarios. However, it demonstrates unsatisfactory performance in personality detection compared to task-specific small model (Ji et al. 2023). Inferring people’s personality traits from online posts is a complex and difficult task. Despite the ineffectiveness of LLMs in this task, previous studies have demonstrated that LLMs exhibit strong language abilities, such as text comprehension, summarization, and sentiment analysis (Zhang et al. 2023; Wang et al. 2023b; Hsieh et al. 2023), which can be used to distill useful knowledge for enhancing small models. On the one hand, we can leverage LLMs to generate post analyses (augmentations) from the aspects of semantic, sentiment, and linguistic, which are key factors for personality detection (Fang et al. 2023). In this way, we can use the augmented post information to learn better post embeddings for personality detection. On the other hand, we can also use LLMs to generate explanations of the complex personality labels to enrich the label information for improving personality de-

*Corresponding author.

tection.

In this paper, we propose a large language model based text augmentation enhanced personality detection model, which distills useful knowledge from LLMs to enhance a small model’s personality detection capabilities, both on the data side and the label side. On the data side, we follow the methodology of contrastive sentence representation works (Gao, Yao, and Chen 2021; Cheng et al. 2023a), using the LLM to generate *knowledgeable* post augmentations from semantic, sentiment and linguistic aspects to provide personality-related knowledge. With the post augmentations, our contrastive post encoder can capture more psycho-linguistic information within the post representations. It is noteworthy that this method does not introduce extra costs during inference. On the label side, considering the personality labels are highly complex in their implications, which makes it difficult for detection, we also use the LLMs to generate additional explanations for the labels to enrich the label information for improving the detection performance.

In summary, our main contributions are as follows:

- We propose a novel LLM based text augmentation enhanced personality detection model, which distills useful knowledge from LLMs to the small model, alleviating the issue of data scarcity and improving personality detection.
- Our model enables LLM to generate post augmentations for contrastive post representation learning from three specially designed aspects: semantic, sentiment, and linguistic, which are key factors for personality detection. Additionally, we use LLM to generate explanations of the personality labels for further improving the detection performance.
- Experimental results on benchmark datasets have demonstrated that our model outperforms the state-of-art baselines, which shows the effectiveness of our method.

Related Work

Personality Detection

Traditional personality detection relies on manual statistical feature engineering (Yarkoni 2010; Schwartz et al. 2013), such as extracting psycho-linguistic features from Linguistic Inquiry and Word Count (LIWC) (Tausczik and Pennebaker 2010) or statistical text features from the bag-of-words (Zhang, Jin, and Zhou 2010) model. As deep learning rapidly advances, a variety of Deep Neural Networks (DNNs) have been employed for personality detection tasks, resulting in significant success, including CNN (Xue et al. 2018), LSTM (Tandera et al. 2017), GRU (Lynn, Balasubramanian, and Schwartz 2020), etc. Recently, personality detection has benefited from pre-trained language models. Jiang, Zhang, and Choi (2020) achieves promising performance by simply concatenating all the utterances from a single user into a document and encoding it with BERT (Devlin et al. 2019) and RoBERTa (Liu et al. 2019). Some works improve it by leveraging the contextual information and external psycho-linguistic knowledge from LIWC (Yang et al.

2021a,b; Zhu et al. 2022). For example, Transformer-MD (Yang et al. 2021a) stores posts’ hidden states in the memory of Transformer-XL (Dai et al. 2019) in order to avoid introducing post-order bias. TrigNet (Yang et al. 2021b) constructs a heterogeneous graph between posts for each user based on the psycho-linguistic knowledge in LIWC and aggregates useful information with a GAT. D-DGCN (Yang et al. 2023) builds a dynamic graph, which enables the model to learn the connections between the posts, and employs DGCN to integrate the information. Although these methods achieve promising performance, they still suffer from the limited supervision of insufficient personality labels, which leads to inferior quality of post embeddings that consequently affect the model performance.

Contrastive Sentence Representation Learning

Contrastive learning was initially proposed by Hadsell, Chopra, and LeCun (2006) and has been widely used for self-supervised representation learning in various domains. In the realm of NLP, a fundamental application of contrastive learning lies in sentence representation learning, applicable to both self-supervised and supervised scenarios. Yan et al. (2021) and Gao, Yao, and Chen (2021) propose different data augmentation strategies for contrastive learning using unlabeled data. In the supervised contrastive learning scenario, SimCSE (Gao, Yao, and Chen 2021) demonstrated that supervised Natural Language Inference (NLI) datasets (Bowman et al. 2015; Williams, Nangia, and Bowman 2018) are effective for learning sentence embeddings by taking sentences with entailment labels as positive samples. Capitalizing on the strong capability of large language models, CLAIF (Cheng et al. 2023b) utilizes LLM to generate similar text and similarity score, using them as positive samples and weights in info-NCE loss, thereby achieving improved performance. Different from these approaches, we generate data augmentations from the LLM, which contains task-specific information, to enhance personality detection performance.

Knowledge Distillation from Large Language Models

Knowledge distillation is used to transfer knowledge from a larger, more powerful teacher model, into smaller student models, enhancing their performance in practical applications (Hinton, Vinyals, and Dean 2015; Iliopoulos et al. 2022; Wang et al. 2021). Recently, depending on the strong language abilities of LLMs, generating additional training data from LLMs for improving smaller models has become a new knowledge distillation trend (Wang et al. 2023a). For example, Self-instruct (Wang et al. 2023a) distills instructional data to enhance the instruction-following capabilities of pre-trained language models. Hsieh et al. (2023) generates rational text from LLMs to enhance the inference abilities of smaller models. Similarly, CLAIF (Cheng et al. 2023b) leverages LLMs to produce similar texts, aiming to learn a better sentence representations. In this paper, we propose to leverage LLMs to generate post analyses (augmentations) from specially-designed perspectives to enhance the small model for personality detection.

Prompt	Here are some posts from a user posting on the internet, Please infer his/her MBTI type based on these posts. Your response should clearly provide one MBTI label and give reasons for each dimension in MBTI.
Input	Posts: “Thanks for all the input guys! And talking a bit more with him after calming down did help. I did reiterate how I ...”
Output	<p>ISFJ</p> <p>Introverted (I): The user seems to prefer spending time alone or with a small group of close friends rather than being in large social situations. They also mention feeling overwhelmed by personal stress and overthinking.</p> <p>Sensing (S): The user is detail-oriented and practical, preferring concrete examples and experiences over abstract concepts. They mention enjoying logic puzzles and disliking math when it is connected to grades.</p> <p>Feeling (F): The user is empathetic and sensitive to criticism, expressing hurt when criticized by others. They also prioritize their emotions and the emotions of others in their decision-making, such as in their desire to work in domestic violence awareness advocacy.</p> <p>Judging (J): The user prefers structure and organization, expressing discomfort with uncertainty and change. They also mention feeling the need to apologize for double-posting and being absent from the internet for a week.</p>

Figure 1: An example of personality detection using the LLM: The result provided by the LLM is ISFJ, while the actual ground truth is ENFP.

Method

Personality detection can be defined as a multi-document multi-label classification task (Yang et al. 2021b, 2023). Formally, given a set $P = \{p_1, p_2 \dots p_n\}$ of N posts from a user, where $p_i = [w_{i1}, w_{i2}, \dots, w_{im}]$ is i -th post with m tokens. This task aims to predict t -dimensional personality traits from the trait-specific label space $Y = \{\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_T\}$ based for one user based on P . For MBTI taxonomy, $T = 4$ and \mathbf{y}_t is one-hot vector. In this paper, we instruct the LLM to generate summarizations from semantical, sentimental and linguistic analysis for post augmentation. The augmented data represented as $X = \{P, P^s, P^e, P^l\}$, where P^s, P^e, P^l correspond to the analysis texts for semantic, sentiment and linguistic, respectively. For the personality labels, we also utilize the LLM to generate explanations of the labels from semantic, sentiment and linguistic perspectives, thereby enriching the label information. For dimension t , the label descriptions are represented as $\hat{\mathbf{y}}_t = \{L_{y_t,0}, L_{y_t,1}\}$, where each $L_{y_t,j} = \{l_{y_t,j}^s, l_{y_t,j}^e, l_{y_t,j}^l\}$, corresponding to the semantic, sentiment and linguistic descriptions.

Analysis of LLM Performance on Personality Detection

LLMs show strong capabilities in many downstream tasks. However, their performance on the personality detection task is unsatisfying, as discussed in Section 4.6. Figure 1 shows a failure case of LLMs, where it wrongly classified ENFP as ISFJ. When we prompt the LLM to elucidate

the analysis process underlying its classification, Figure 1 shows that the LLM primarily infers the personality traits based on only the post semantics. However, previous studies (Tausczik and Pennebaker 2010; Stajner and Yenikent 2020) demonstrate that the way people communicate and their sentiments often reveal more about their psychological state than the semantics of their communication content. LLM fails to capture the sentiment and linguistic patterns for personality detection.

Generating Knowledgeable Post Augmentations from LLMs

Personality detection is a complex and difficult task, since the personality traits are latent theoretical variables that cannot be directly or objectively observed (Fang et al. 2023). Previous research has demonstrated a strong link between personality traits and a person’s sentiments, words, and opinions (Kishima et al. 2021; Johnson et al. 2023). Therefore, despite of the ineffectiveness of the LLM in this task, we propose to fully utilize the LLM’s abilities to distill personality-related knowledge, aiming to enhance smaller models for improving personality detection performance.

Specifically, we empirically instruct LLM (chatGPT) to generate post analyses from three aspects: semantic, sentiment, and linguistic, serving as data augmentations to original post. One analyses example is shown in Figure 2 generated with the instruction as follows:

Your task is to analyze the characteristics of a user based on a piece of text published by the user on the Internet. You are required to analyze it from the perspectives of semantic, sentiments, and linguistics. Note that if the text is incomplete and ends with an ellipsis, it may have been truncated due to external reasons, in which case you should ignore it. post:...

Contrastive Post Encoder

Contrastive learning aims to learn efficient representations by aligning semantically similar entities closer and distancing the dissimilar ones (Hadsell, Chopra, and LeCun 2006; Gao, Yao, and Chen 2021). One key component of contrastive learning is the selection of positive pairs.

With the obtained post augmentations from the LLM as positive samples, we learn the post embeddings with the contrastive post encoder that able to capture more personality-related feature.

Formally, for i -th post p_i and it’s analysis from three aspect $P_{pos} = \{p_i^s, p_i^e, p_i^l\}$, we use the final hidden state of “[CLS]” as the sentence representation:

$$\mathbf{h}_i = \text{BERT}(p_i) \in \mathbb{R}^{1 \times d}, \quad (1)$$

$$\mathbf{h}_i^+ = \text{BERT}(p_i^+), \text{ where } p_i^+ \in \{p_i^s, p_i^e, p_i^l\}. \quad (2)$$

Due to the distribution difference between the original post text and the analysis text, we add an extra MLP same to original BERT implementation as a projection head to mitigate this discrepancy:

$$z_i = \delta(\mathbf{W}\mathbf{h}_i + \mathbf{b}), \quad (3)$$

$$z_i^+ = \delta(\mathbf{W}\mathbf{h}_i^+ + \mathbf{b}), \quad (4)$$

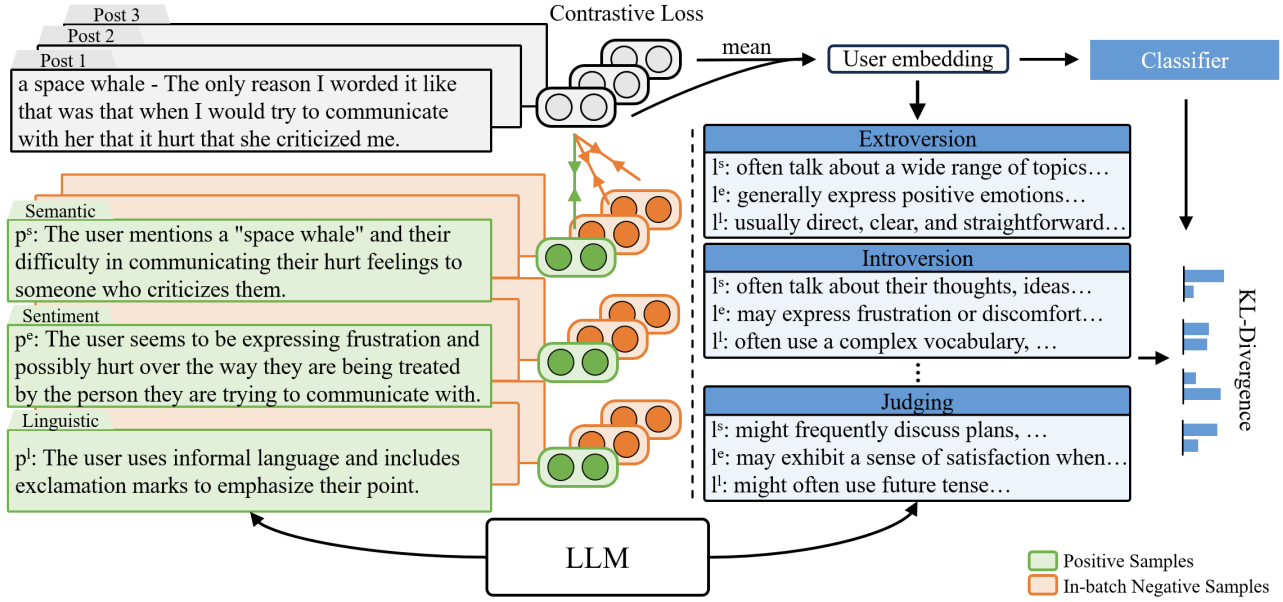


Figure 2: An overview of our TAE.

where δ is Tanh function. We adopt the contrastive framework presented by Chen et al. (2020), utilizing a post-wise info-NCE loss with in-batch negatives, for i -th post in mini-batch:

$$\mathcal{L}_{cl} = \frac{1}{M} \sum_{i=1}^M \ell_i^{cl} \quad (5)$$

$$\ell_i^{cl} = -\log \sum_{(z_i, z_i^+) \in P_{pos}} \frac{e^{\text{sim}(z_i, z_i^+)/\tau}}{\sum_{j=1}^M \sum_{(z_j, z_j^+) \in P_{pos}} e^{\text{sim}(z_i, z_j^+)/\tau}},$$

where M is total number of post in mini-batch, τ is temperature hyperparameter and $\text{sim}(z_1, z_2)$ is cosine similarity $\frac{z_1^\top z_2}{\|z_1\| \cdot \|z_2\|}$.

After obtaining posts representation of a user, we simply use average pooling method to produce the user representation:

$$\mathbf{u} = \text{mean}([\mathbf{h}_1, \mathbf{h}_2, \dots, \mathbf{h}_N]). \quad (6)$$

LLM Enriching Label Information

Previous methods directly use one-hot label for personality detection, overlooking the information in personality traits. Since personality traits are complicated to comprehension, we also turn to the LLM to generate explanations of each trait from semantic, sentiment and linguistic perspectives, as the same to post augmentations.

Specifically, given a personality label $L_{y_t, j}$, containing three descriptions from semantic, sentiment and linguistic aspects $L_{y_t, j} = \{l_{y_t, j}^s, l_{y_t, j}^e, l_{y_t, j}^l\}$, the label representation is obtained by:

$$\mathbf{v}_{y_t, j} = \text{mean}([\text{BERT}(l_{y_t, j})]), \text{ where } l_{y_t, j} \in L_{y_t, j}, \quad (7)$$

where we also use the embedding of "[CLS]" in last layer as the sentence embedding.

Based on label representations, we generate soft labels to overcome the over-confident issue, which is primarily attributed to the information loss in the MBTI taxonomy's binary classification, which neglects the position on the scale (Stajner and Yenikent 2020). Dataset label noise introduced by the questionnaire measurement errors (Anusic et al. 2009) could also lead to the over-confident issue. We leverage the generated label information to assign a soft label to a user based on the similarity between the user embedding and the label embedding.

Formally, to generate soft label, we first calculate the dimension-wise label similarity distribution with cosine similarity and softmax function. Then, we combine the original one-hot vector with a controlling parameter α and apply an additional softmax function:

$$\mathbf{y}_t^s = \text{softmax}(\text{sim}(\mathbf{u}, \mathbf{V}_{y_t})), \quad (8)$$

$$\mathbf{y}_t^c = \text{softmax}(\alpha \mathbf{y}_t + \mathbf{y}_t^s), \quad (9)$$

In this way, we can assign a softer label instead of one-hot label when a user is relatively neutral within a dimension, improving model's generalization ability.

Finally, we employ T softmax-normalized linear transformations to predict personality traits, and apply a KL-divergence as the loss function:

$$\hat{\mathbf{y}}^t = \text{softmax}(\mathbf{u} \mathbf{W}_u^t + \mathbf{b}_u^t), \quad (10)$$

$$\mathcal{L}_{det} = -\frac{1}{B} \sum_{i=1}^B \ell_i^{det},$$

Dataset	Types	Train	Validation	Test
Kaggle	I / E	4011/1194	1326/409	1339/396
	S / N	727/4478	222/1513	248/1487
	T / F	2410/2795	791/944	780/955
	P / J	3096/2109	1063/672	1082/653
Pandora	I / E	4278/1162	1427/386	1437/377
	S / N	610/4830	208/1605	210/1604
	T / F	3549/1891	1120/693	1182/632
	P / J	3211/2229	1043/770	1056/758

Table 1: Statistics of the Kaggle and Pandora datasets.

$$\begin{aligned}
\ell_i^{det} &= \sum_{t=1}^T \text{KL-divergence}(\mathbf{y}^c, \hat{\mathbf{y}}^t) \\
&= \sum_{t=1}^T \sum_{j=0,1} \mathbf{y}_j^c \log\left(\frac{\mathbf{y}_j^c}{\hat{\mathbf{y}}_j^t}\right)
\end{aligned} \tag{11}$$

where B denotes number of samples in batch. Overall, the whole training objective of TAE is formulated as follows:

$$\mathcal{L} = \mathcal{L}_{det} + \lambda \mathcal{L}_{cl}, \tag{12}$$

where λ is a trad-off parameter to balance the two losses.

Experiments

Datasets

Following previous studies (Yang et al. 2021a,b, 2023), we choose two widely used datasets in personality detection, Kaggle¹ and Pandora² (Gjurković et al. 2021). The Kaggle dataset is collected from PersonalityCafe³, a platform where people share their personality types and engage in daily communications. This dataset contains a total of 8675 users, with each user contributing 45-50 posts. Pandora is another dataset collected from Reddit⁴, where personality labels are extracted from users’ self-introductions that contain MBTI types. This dataset contains dozens to hundreds of posts from each of the 9067 users. Both datasets are based on the MBTI taxonomy, which divides people’s personality into four dimensions, each containing two aspects: **I**ntroversion vs. **E**xtraversion (I vs. E), **S**ensing vs. **i**ntuition (S vs. N), **T**hinking vs. **F**eeling (T vs. F), and **P**erception vs. **J**udging (P vs. J). Since the two datasets are severely imbalanced, we employ the Macro-F1 metric, and use average Macro-F1 for overall performance. We adopt the same data division as in (Yang et al. 2023), which shuffles the datasets and splits them in a 60-20-20 proportion for training, validation, and testing, respectively. Table 1 shows the statistics of the two datasets.

Baselines

We compare our model with several baselines as follows.

¹<https://www.kaggle.com/datasnaek/mbti-type>

²<https://psy.takelab.fer.hr/datasets/all>

³<http://personalitycafe.com/forum>

⁴<https://www.reddit.com>

SVM (Cui and Qi 2017) and **XGBoost** (Tadesse et al. 2018): These methods concatenate all the posts of a user into a document first, and then utilize SVM or XGBoost for classification based on features extracted using bag-of-words models.

BiLSTM (Tandera et al. 2017): Glove is employed for generating word embeddings. Subsequently, a Bi-directional LSTM is used to encode each post, with the averaged post representation serving as the user representation for personality detection.

BERT_{concat} (Jiang, Zhang, and Choi 2019): It concatenates a user’s posts into an extended document and then employs BERT to encode this composite text for user representation.

BERT_{mean} (Keh and Cheng 2019): It uses BERT to encode each post individually, extracting the CLS embedding for the post representation. Then, it employs mean pooling to derive the user representation.

AttRCNN (Xue et al. 2018): It adopts a hierarchical structure, in which a variant of Inception (Szegedy et al., 2017) is utilized to encode each post and a CNN-based aggregator is employed to obtain the user representation. Besides, it considers psycho-linguistic knowledge by concatenating the LIWC features with the user representation.

SN+Attn (Lynn, Balasubramanian, and Schwartz 2020): It is also a hierarchical network that employs GRU and attention mechanisms to encode sequences from both the word-level and post-level for user representation, using a pre-trained 200-dimensional word2vec model (Mikolov et al., 2013) for word embeddings.

Transformer-MD (Yang et al. 2021a): Transformer-MD draws inspiration from the Transformer XL, first employing a low-level encoder to individually encode each post, storing the CLS embeddings in memory. Then, a high-level encoder aggregates information from all other posts to further encode the post.

TrigNet (Yang et al. 2021b): TrigNet constructs a tri-graph for each user consisting of posts, words, and word categories, based on LIWC dictionary. It then employs a modified GAT to encode the graph and uses average pooling to obtain a user representation.

D-DGCN (Yang et al. 2023): DDGCN employs a dynamic graph to model a user’s posts, allowing the model to learn the connections between posts autonomously, and then uses DGCN to encode the graph and obtains a user representation. It has a variant **D-DGCN+ ℓ_0** representing adding ℓ_0 norm with Hard Concrete distribution.

ChatGPT⁵: We applied the ‘gpt-3.5-turbo-0301’ version of ChatGPT, and set the temperature to 0, making the outputs mostly deterministic for the identical inputs.

Implementation Details

All the deep learning models are implemented in PyTorch, and the optimizer used is Adam (Kingma and Ba 2014). The learning rate for the pre-trained post encoder is set to 1e-5, and for other parameters is set to 1e-3. We employed ‘bert-base-uncased’ from BERT as our post encoder. The

⁵<https://chat.openai.com/>

Methods	Kaggle					Pandora				
	<i>I/E</i>	<i>S/N</i>	<i>T/F</i>	<i>P/J</i>	Average	<i>I/E</i>	<i>S/N</i>	<i>T/F</i>	<i>P/J</i>	Average
SVM	53.34	47.75	76.72	63.03	60.21	44.74	46.92	64.62	56.32	53.15
XGBoost	56.67	52.85	75.42	65.94	62.72	45.99	48.93	63.51	55.55	53.50
BiLSTM	57.82	57.87	69.97	57.01	60.67	48.01	52.01	63.48	56.21	54.93
BERT _{concat}	58.33	53.88	69.36	60.88	60.61	54.22	49.15	58.31	53.14	53.71
BERT _{mean}	64.65	57.12	77.95	65.25	66.24	56.60	48.71	64.70	56.07	56.52
AttRCNN	59.74	64.08	78.77	66.44	67.25	48.55	56.19	64.39	57.26	56.60
SN+Attn	65.43	62.15	78.05	63.92	67.39	56.98	54.78	60.95	54.81	56.88
Transformer-MD	66.08	69.10	79.19	67.50	70.47	55.26	58.77	69.26	60.90	61.05
TrigNet	69.54	67.17	79.06	67.69	70.86	56.69	55.57	66.38	57.27	58.98
D-DGCN	68.41	65.66	79.56	67.22	70.21	61.55	55.46	71.07	59.96	62.01
D-DGCN+ ℓ_0	69.52	67.19	80.53	68.16	71.35	59.98	55.52	70.53	59.56	61.40
ChatGPT	65.86	51.69	78.60	63.93	66.89	55.52	49.79	71.25	60.51	59.27
TAE(our)	70.90	66.21	81.17	70.20	72.07	62.57	61.01	69.28	59.34	63.05

Table 2: Overall results of our TAE and baseline models in Macro-F1 (%) score.

mini-batch size is set to 8. The temperature τ is set to 0.07 and trade-off parameter λ is set to 1. The controlling parameter α is set to 4. Following previous works (Yang et al. 2021b, 2023, 2021a), We limit each user to a maximum of 50 posts, and limit each post and analysis text in both datasets to a maximum length of 70 words. Additionally, we replace words that match any personality label with $\langle \text{type} \rangle$ to avoid information leaks (Yang et al. 2021b). For the LLM, in consideration of cost and efficiency, we use 'gpt-3.5-turbo-0301' model to generate post augmentations.

Overall Results

The overall results are presented in Table 2. We can find that the proposed TAE consistently outperforms all the baselines on Macro-F1 scores. Compared to the best baselines D-DGCN (71.35% on average Macro-F1) and D-DGCN+ ℓ_0 (62.01%) respectively on Kaggle dataset and Pandora dataset, our model improves them by 1.01% and 1.68%, respectively. The result demonstrates the superiority of our model in personality detection. We believe the reasons are two-fold: (1) Our model TAE benefits from the post augmentations generated by the LLM, enabling the contrastive post encoder to extract information that is more conducive to personality detection. (2) The generated additional explanations of personality labels effectively help in accomplishing the detection task. We can also find that compared to BERT_{mean}, our model has a marked improvement, with respective gains (5.83% on the Kaggle dataset and 6.53% on the Pandora dataset). This demonstrates that data augmentations from LLM in data-scarce situations is highly advantageous. Finally, we can observe that methods leveraging external psycho-linguistic knowledge from LIWC, such as AttRCNN and TrigNet achieve relatively good results, validating the effectiveness of introducing external knowledge. Furthermore, our model achieves the best performance, indicating that the generated post analyses and label explanations from large language models can provide effective information for personality detection.

Methods	Kaggle				
	<i>I/E</i>	<i>S/N</i>	<i>T/F</i>	<i>P/J</i>	Average
TAE _{w/o semantic}	71.24	66.34	80.61	67.43	71.40
TAE _{w/o sentiment}	70.14	65.29	80.03	69.55	71.25
TAE _{w/o linguistic}	70.34	65.69	78.89	69.71	71.07
TAE _{w/o label}	70.57	65.89	81.90	69.78	72.02
TAE _{w/o All}	64.65	57.12	77.95	65.25	66.24
TAE _{Concat}	66.60	64.28	78.19	62.55	67.91
TAE _{WS}	64.63	62.62	77.68	64.30	67.31
TAE	70.90	66.21	81.17	70.20	72.07

Table 3: Results of ablation study on Macro-F1 on the Kaggle dataset.

Ablation Study

To verify the importance of each component in our TAE model, we conduct an ablation study on the Kaggle dataset. First, we analyze the contributions of the post augmentations from each aspect. As we can see from Table 3, among the three aspects, linguistic augmentation proves to be the most important one in our method, as the average Macro-F1 score declines most largely when it is removed. Furthermore, the semantic information is the least influential augmentation. This indicates that the semantic information is relatively less important for personality detection, which is consistent with the observations in previous works (Yang et al. 2021b; Stajner and Yenikent 2020). When removing the LLM based label information enrichment, the performance of the model slightly decreases. The ablation studies demonstrate that our model benefits from the LLM generated post augmentations from semantic, sentiment and linguistic aspects, as well as the LLM based label information enrichment.

To further explore the effectiveness of using the LLM-generated analysis texts as data augmentations for contrastive post representation learning, we conducted ablation experiments. We compared this approach with the variants that directly use these LLM-based analysis texts (post augmentations) as additional input. TAE_{Concat} first encodes post and analyses respectively and then concatenates their em-

Methods	Kaggle				
	I/E	S/N	T/F	P/J	Average
ChatGPT	65.86	51.69	78.60	63.93	66.89
ChatGPT _{cot}	65.13	60.35	75.73	59.30	65.12
ChatGPT _{3 shot}	70.61	58.35	76.58	65.43	67.74
TAE	70.90	66.21	81.17	70.20	72.07

Table 4: ChatGPT performances on kaggle testing set.

beddings. TAE_{WS} denotes using weighted sum as the pooling method of the post features and post augmentation features. $TAE_{w/o All}$ denotes removing all components in TAE which is an original BERT with mean pooling. Table 4 shows their performance on the Kaggle dataset, we can observe that models taking the post augmentations as extra input (TAE_{Concat} , TAE_{WS}) outperform the baseline $TAE_{w/o All}$, but they are all inferior to our TAE which is a contrastive model. This demonstrates the effectiveness of data augmentation and the contrastive learning paradigm. Moreover, the models taking analysis texts as extra input requires the LLM to generate analysis texts during inference stage, which is costly. We also visualize the weights learned by TAE_{WS} in Figure 3. It shows that linguistic analysis is more important, which is consistent with our previous observation.

LLM Performance

To analyze LLM’s performance in the personality detection task, we conducted a series of experiments using ChatGPT. We assessed LLM’s capabilities under three settings: zero-shot, CoT, and few-shot, using the Kaggle testing set. For the few-shot setting, we randomly selected 3 examples from the training set. Specifically, we employed the ‘gpt-3.5-turbo-0301’ version but switched to ‘gpt-3.5-turbo-16k-0613’ for the few-shot setting, as the input length exceeded the limit. As shown in Table 4, the performance of ChatGPT is comparable to that of a small model that has been fine-tuned ($BERT_{mean}$) for a specific task on Macro-F1 metrics. Under the CoT setting and few-shot setting, the classification performance slightly decreased. This indicates that the LLM’s reasoning ability fails in personality detection. Directly applying LLM to the personality detection is not appropriate. Thus, in this paper, we consider leveraging the LLM to enhance the small model for personality detection by distilling useful knowledge from LLM to the small model.

Effect of Trade-Off Parameter

For trade-off parameter λ in TAE, we searched in $\{0.5, 1, 1.5, 2, 2.5, 3, 3.5, 4\}$. Figure 4 demonstrates how the model performance changes with the λ on the validation sets of Kaggle and Pandora datasets. We can observe that the Macro-F1 value first grows and reaches the highest value at $\lambda=1$ while it begins to drop when λ is larger than 1. This may be because that initially with the value of λ increases, TAE can benefit more from the contrastive learning with the LLM generated post augmentations, which helps to learn better post representations. However, if λ is set too large (> 1), the contrastive signal out weight the detection loss. Overall, $\lambda=1$ can reach best balance between detection loss


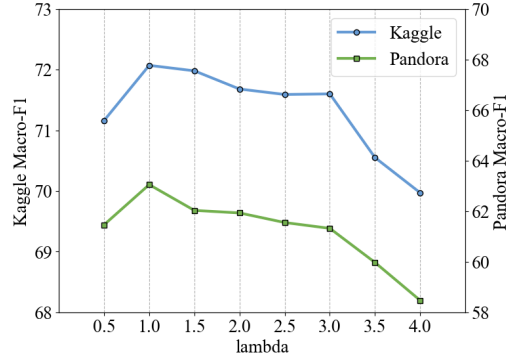
Figure 3: Visualization of learned weight of post and analyses in TAE_{WS} .

Figure 4: Performance curves for different trade-off parameter.

and contrastive loss on both Kaggle and Pandora datasets.

Conclusion

In this paper, we propose a large language model (LLM) based text augmentation enhanced personality detection model, which distills the useful knowledge from the LLM to address the data scarcity issue faced by small models in personality detection, even when the LLM itself struggles with the task. By leveraging the LLM’s abilities in text comprehension, summarization, and sentiment analysis, we instruct it to generate post analyses from three specially-designed perspectives: semantic, sentiment, and linguistic, which play a critical role for personality detection. Taking these analyses as positive samples and using contrastive learning to pull them together in the embedding space enables the post encoder in the small model to better capture the psycholinguistic information within the post representations, thus improving personality detection. Furthermore, we utilize the LLM to generate label descriptions, enriching the semantics of personality labels and utilize label information to generate soft labels to overcome the over-confidence issue, enhancing models generalization ability. Experimental results on two benchmark datasets demonstrate that our model outperforms the state-of-the-art methods on personality detection.

In future work, we will explore how to combine the advantages of the existing knowledge graphs and the LLM in improving personality detection.

Acknowledgments

This work was supported by the National Science Foundation of China (No. 62276029), and CCF-Zhipu. AI Large Model Fund (No. 202217).

References

- Anusic, I.; Schimmack, U.; Pinkus, R.; and Lockwood, P. 2009. The Nature and Structure of Correlations Among Big Five Ratings: The Halo-Alpha-Beta Model. *Journal of personality and social psychology*, 97: 1142–56.
- Bagby, R. M.; Gralnick, T. M.; Al-Dajani, N.; and Uliaszek, A. A. 2016. The Role of the Five-Factor Model in Personality Assessment and Treatment Planning. *Clinical Psychology: Science and Practice*, 23(4): 365–381.
- Bowman, S. R.; Angeli, G.; Potts, C.; and Manning, C. D. 2015. A large annotated corpus for learning natural language inference. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, 632–642.
- Brown, T.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J. D.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; Agarwal, S.; Herbert-Voss, A.; Krueger, G.; Henighan, T.; Child, R.; Ramesh, A.; Ziegler, D.; Wu, J.; Winter, C.; Hesse, C.; Chen, M.; Sigler, E.; Litwin, M.; Gray, S.; Chess, B.; Clark, J.; Berner, C.; McCandlish, S.; Radford, A.; Sutskever, I.; and Amodei, D. 2020. Language Models are Few-Shot Learners. In Larochelle, H.; Ranzato, M.; Hadsell, R.; Balcan, M.; and Lin, H., eds., *Advances in Neural Information Processing Systems*, volume 33, 1877–1901.
- Cheng, Q.; Yang, X.; Sun, T.; Li, L.; and Qiu, X. 2023a. Improving Contrastive Learning of Sentence Embeddings from AI Feedback. In *Findings of the Association for Computational Linguistics: ACL 2023*, 11122–11138.
- Cheng, Q.; Yang, X.; Sun, T.; Li, L.; and Qiu, X. 2023b. Improving Contrastive Learning of Sentence Embeddings from AI Feedback. In *Findings of the Association for Computational Linguistics: ACL 2023*, 11122–11138.
- Cui, B.; and Qi, C. 2017. Survey Analysis of Machine Learning Methods for Natural Language Processing for MBTI Personality Type Prediction. Technical report, Stanford University.
- Dai, Z.; Yang, Z.; Yang, Y.; Carbonell, J.; Le, Q.; and Salakhutdinov, R. 2019. Transformer-XL: Attentive Language Models beyond a Fixed-Length Context. In Korhonen, A.; Traum, D.; and Màrquez, L., eds., *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 2978–2988.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, 4171–4186.
- Fang, Q.; Giachanou, A.; Bagheri, A.; Boeschoten, L.; van Kesteren, E.-J.; Shafiee Kamalabad, M.; and Oberski, D. 2023. On Text-based Personality Computing: Challenges and Future Directions. In *Findings of the Association for Computational Linguistics: ACL 2023*, 10861–10879.
- Gao, T.; Yao, X.; and Chen, D. 2021. SimCSE: Simple Contrastive Learning of Sentence Embeddings. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, 6894–6910.
- Gjurković, M.; Karan, M.; Vukojević, I.; Bošnjak, M.; and Snajder, J. 2021. PANDORA Talks: Personality and Demographics on Reddit. In *Proceedings of the Ninth International Workshop on Natural Language Processing for Social Media*, 138–152.
- Hadsell, R.; Chopra, S.; and LeCun, Y. 2006. Dimensionality Reduction by Learning an Invariant Mapping. In *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'06)*, volume 2, 1735–1742.
- Hinton, G.; Vinyals, O.; and Dean, J. 2015. Distilling the Knowledge in a Neural Network. arXiv:1503.02531.
- Hsieh, C.-Y.; Li, C.-L.; Yeh, C.-k.; Nakhost, H.; Fujii, Y.; Ratner, A.; Krishna, R.; Lee, C.-Y.; and Pfister, T. 2023. Distilling Step-by-Step! Outperforming Larger Language Models with Less Training Data and Smaller Model Sizes. In *Findings of the Association for Computational Linguistics: ACL 2023*, 8003–8017.
- Iliopoulos, F.; Kontonis, V.; Baykal, C.; Menghani, G.; Trinh, K.; and Vee, E. 2022. Weighted Distillation with Unlabeled Examples. In Koyejo, S.; Mohamed, S.; Agarwal, A.; Belgrave, D.; Cho, K.; and Oh, A., eds., *Advances in Neural Information Processing Systems*, volume 35, 7024–7037.
- Ji, Y.; Wu, W.; Zheng, H.; Hu, Y.; Chen, X.; and He, L. 2023. Is ChatGPT a Good Personality Recognizer? A Preliminary Study. arXiv:2307.03952.
- Jiang, H.; Zhang, X.; and Choi, J. D. 2019. Automatic Text-based Personality Recognition on Monologues and Multi-party Dialogues Using Attentive Networks and Contextual Embeddings. arXiv:1911.09304.
- Jiang, H.; Zhang, X.; and Choi, J. D. 2020. Automatic Text-Based Personality Recognition on Monologues and Multi-party Dialogues Using Attentive Networks and Contextual Embeddings (Student Abstract). volume 34, 13821–13822.
- Johnson, R.; Wootten, M.; Spear, A.; and Smolensky, A. 2023. The Relationship Between Personality Traits and the Processing of Emotion Words: Evidence from Eye-Movements in Sentence Reading. *Journal of Psycholinguistic Research*, 1–27.
- Keh, S. S.; and Cheng, I.-T. 2019. Myers-Briggs Personality Classification and Personality-Specific Language Generation Using Pre-trained Language Models. arXiv:1907.06333.
- Kishima, R.; Matsumoto, K.; Yoshida, M.; and Kita, K. 2021. Construction of MBTI Personality Estimation Model Considering Emotional Information. In *Proceedings of the 35th Pacific Asia Conference on Language, Information and Computation*, 262–269. Shanghai, China.
- Liu, Y.; Ott, M.; Goyal, N.; Du, J.; Joshi, M.; Chen, D.; Levy, O.; Lewis, M.; Zettlemoyer, L.; and Stoyanov, V. 2019. RoBERTa: A Robustly Optimized BERT Pretraining Approach. arXiv:1907.11692.
- Lynn, V.; Balasubramanian, N.; and Schwartz, H. A. 2020. Hierarchical Modeling for User Personality Prediction: The Role of Message-Level Attention. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 5306–5316.

- Ouyang, L.; Wu, J.; Jiang, X.; Almeida, D.; Wainwright, C.; Mishkin, P.; Zhang, C.; Agarwal, S.; Slama, K.; Ray, A.; Schulman, J.; Hilton, J.; Kelton, F.; Miller, L.; Simens, M.; Askell, A.; Welinder, P.; Christiano, P. F.; Leike, J.; and Lowe, R. 2022. Training language models to follow instructions with human feedback. In Koyejo, S.; Mohamed, S.; Agarwal, A.; Belgrave, D.; Cho, K.; and Oh, A., eds., *Advances in Neural Information Processing Systems*, volume 35, 27730–27744.
- Schwartz, H. A.; Eichstaedt, J. C.; Kern, M. L.; Dziurzynski, L.; Ramones, S. M.; Agrawal, M.; Shah, A.; Kosinski, M.; Stillwell, D.; Seligman, M. E. P.; and Ungar, L. H. 2013. Personality, Gender, and Age in the Language of Social Media: The Open-Vocabulary Approach. *PLOS ONE*, 8(9): 1–16.
- Stajner, S.; and Yenikent, S. 2020. A Survey of Automatic Personality Detection from Texts. In *Proceedings of the 28th International Conference on Computational Linguistics*, 6284–6295.
- Sun, Y.; Wang, S.; Feng, S.; Ding, S.; Pang, C.; Shang, J.; Liu, J.; Chen, X.; Zhao, Y.; Lu, Y.; Liu, W.; Wu, Z.; Gong, W.; Liang, J.; Shang, Z.; Sun, P.; Liu, W.; Ouyang, X.; Yu, D.; Tian, H.; Wu, H.; and Wang, H. 2021. ERNIE 3.0: Large-scale Knowledge Enhanced Pre-training for Language Understanding and Generation. arXiv:2107.02137.
- Tadesse, M. M.; Lin, H.; Xu, B.; and Yang, L. 2018. Personality Predictions Based on User Behavior on the Facebook Social Media Platform. *IEEE Access*, 6: 61959–61969.
- Tandera, T.; Hendro; Suhartono, D.; Wongso, R.; and Prasetio, Y. L. 2017. Personality Prediction System from Facebook Users. *Procedia Computer Science*, 116: 604–611. Discovery and innovation of computer science technology in artificial intelligence era: The 2nd International Conference on Computer Science and Computational Intelligence (ICCS2017).
- Tausczik, Y. R.; and Pennebaker, J. W. 2010. The Psychological Meaning of Words: LIWC and Computerized Text Analysis Methods. *Journal of Language and Social Psychology*, 29(1): 24–54.
- Wang, S.; Liu, Y.; Xu, Y.; Zhu, C.; and Zeng, M. 2021. Want To Reduce Labeling Cost? GPT-3 Can Help. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, 4195–4205.
- Wang, Y.; Kordi, Y.; Mishra, S.; Liu, A.; Smith, N. A.; Khoshdel, D.; and Hajishirzi, H. 2023a. Self-Instruct: Aligning Language Models with Self-Generated Instructions. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 13484–13508.
- Wang, Z.; Xie, Q.; Ding, Z.; Feng, Y.; and Xia, R. 2023b. Is ChatGPT a Good Sentiment Analyzer? A Preliminary Study. arXiv:2304.04339.
- Wen, Z.; Cao, J.; Yang, R.; Liu, S.; and Shen, J. 2021. Automatically Select Emotion for Response via Personality-affected Emotion Transition. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, 5010–5020.
- Williams, A.; Nangia, N.; and Bowman, S. 2018. A Broad-Coverage Challenge Corpus for Sentence Understanding through Inference. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, 1112–1122.
- Xue, D.; Wu, L.; Hong, Z.; Guo, S.; Gao, L.; Wu, Z.; Zhong, X.; and Sun, J. 2018. Deep learning-based personality recognition from text posts of online social networks. *Applied Intelligence*, 48.
- Yan, Y.; Li, R.; Wang, S.; Zhang, F.; Wu, W.; and Xu, W. 2021. ConSERT: A Contrastive Framework for Self-Supervised Sentence Representation Transfer. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, 5065–5075.
- Yang, F.; Quan, X.; Yang, Y.; and Yu, J. 2021a. Multi-Document Transformer for Personality Detection. volume 35, 14221–14229.
- Yang, R.; Chen, J.; and Narasimhan, K. 2021. Improving Dialog Systems for Negotiation with Personality Modeling. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, 681–693.
- Yang, T.; Deng, J.; Quan, X.; and Wang, Q. 2023. Orders Are Unwanted: Dynamic Deep Graph Convolutional Network for Personality Detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 13896–13904.
- Yang, T.; Yang, F.; Ouyang, H.; and Quan, X. 2021b. Psycholinguistic Tripartite Graph Network for Personality Detection. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, 4229–4239.
- Yarkoni, T. 2010. Personality in 100,000 Words: A large-scale analysis of personality and word use among bloggers. *Journal of Research in Personality*, 44(3): 363–373.
- Zhang, S.; Roller, S.; Goyal, N.; Artetxe, M.; Chen, M.; Chen, S.; Dewan, C.; Diab, M.; Li, X.; Lin, X. V.; Mihaylov, T.; Ott, M.; Shleifer, S.; Shuster, K.; Simig, D.; Koura, P. S.; Sridhar, A.; Wang, T.; and Zettlemoyer, L. 2022. OPT: Open Pre-trained Transformer Language Models. arXiv:2205.01068.
- Zhang, T.; Ladhak, F.; Durmus, E.; Liang, P.; McKeown, K.; and Hashimoto, T. B. 2023. Benchmarking Large Language Models for News Summarization. arXiv:2301.13848.
- Zhang, Y.; Jin, R.; and Zhou, Z.-H. 2010. Understanding bag-of-words model: a statistical framework. *International Journal of Machine Learning and Cybernetics*, 1: 43–52.
- Zhu, Y.; Hu, L.; Ge, X.; Peng, W.; and Wu, B. 2022. Contrastive Graph Transformer Network for Personality Detection. In Raedt, L. D., ed., *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI-22*, 4559–4565.