

Neurosymbolic AI for Personalized Sentiment Analysis

Luyao Zhu¹[0000-0003-4602-0536], Rui Mao¹[0000-0002-1082-8755],
Erik Cambria¹[0000-0002-3030-1280], and Bernard J. Jansen²[0000-0002-6468-6609]

¹ College of Computing and Data Science,
Nanyang Technological University, Singapore

`luyao001@e.ntu.edu.sg`
`{rui.mao,cambria}@ntu.edu.sg`
`https://sentic.net`

² Qatar Computing Research Institute,
Hamad Bin Khalifa University, Doha, Qatar
`bjansen@hbku.edu.qa`

Abstract. Sentiment analysis is crucial in extracting valuable insights from vast amounts of textual data generated across various platforms, such as social media, customer reviews, news articles, etc. Over the years, researchers and business professionals have worked hard to refine sentiment analysis algorithms, but there is a limit to how accurate any algorithm can be without considering personalization. In this work, we propose a framework for personalized sentiment analysis that performs automatic user profiling by modeling users based on different levels of personalization, before performing sentiment analysis. In particular, such framework leverages seven levels of personalization (from bottom to top), namely: Entity, to distinguish between humans and other intelligent agents; Culture, to take into account how different cultures perceive the same concept as positive or negative; Religion, to consider how specific religious beliefs may affect an individual’s opinion about certain topics; Vocation, to better gauge people’s opinion based on their job and education level; Ideology, to take into account political beliefs as well as social, economic, or philosophical viewpoints; Personality, to better classify certain concepts as positive or negative based on personality traits; finally, Subjectivity, to take into account personal preferences and experiences.

Keywords: Sentiment Analysis · Personalization · Persona · Personality · AI · NLP.

1 Introduction

Sentiment analysis has evolved significantly from traditional survey methods, transitioning from structured, manual data collection to automated, unstructured text analysis in the early 2000s. More advanced artificial intelligence (AI) techniques have been applied to the problem of automatically extracting people’s opinions from text [14].




Fig. 1. Conventional sentiment analysis vs. personalized sentiment analysis.

Recent advancements in sentiment analysis are shifting from document-level or sentence-level [8] to finer-grained aspect-level [35]. Most research relied on supervised learning with annotated data. The ground truth labels were commonly identified by recruiting professional annotators and then taking the label agreed upon by majority voting among them [12]. This method works well in domains with consistent sentiment perception, e.g., interpreting the sentiment of product reviews. However, when it comes to understanding the sentimental perception of individuals, the task can be more challenging. For example, many sentiment analysis annotation tasks were achieved with imperfect agreement rates [23, 54]. In other words, not all labels were agreed upon by all annotators. While using majority voting for labeling is acceptable in a probabilistic sense, it can be seen as overlooking the human-centric aspect of sentiment analysis.

We introduce the task of personalized sentiment analysis, which focuses on analyzing individual sentiment perceptions. This approach is motivated by the observation that different individuals may perceive an identical statement differently regarding its sentiment polarity. In contrast, conventional sentiment analysis aims to predict the semantic sentiment of a statement, where the sentiment prediction remains the same for an identical statement. For instance, an introverted person may have a negative sentiment towards performing in front of a large audience, while an extroverted person may view the same situation positively (see Figure 1). This variation in sentiment perception can be attributed to personality traits. While the distinction between introversion and extroversion is from personality theory, the variability in sentiment perception among individuals can also be influenced by other factors. The inconsistency in sentimental perception between individuals may originate from multiple sources. For example, as the adage suggests, “*the enemy of my enemy is my friend*”, the sentimental perception of a person can be driven by the relationship or the context of the situation. In this case, personalized sentiment analysis extends beyond traditional semantic and pragmatic understanding, incorporating a broader range of human subjective factors, such as persona information.




Fig. 2. Personalized Sentiment Analysis Pyramid.

To this end, we propose a novel neurosymbolic AI framework that leverages seven levels of personalization (see Figure 2) for personalized sentiment analysis. In particular, such a framework initially targets identifying whether the user is a human (e.g., male or female) or other intelligent agent, determining whether to consider typical human needs and beliefs. Secondly, the framework aims to identify the user’s cultural background to discern whether a given concept is perceived as positive or negative based on different cultural beliefs. Next, a similar mechanism is applied to Religion. Vocation is considered an important factor in shaping users’ views based on their job and education level. Following, the user’s ideologies are modeled to consider political beliefs and social, economic, or philosophical viewpoints. Next, Personality is detected in order to better classify certain concepts as positive or negative based on the user’s personality traits. Finally, Subjectivity aims to consider specific user preferences learned from historical or training data.

This work aims to evaluate whether the framework of the identified seven levels of personalization can help improve the accuracy of sentiment analysis. We evaluate the framework based on a dialogue dataset [9] that originated from *Harry Potter* novels and large language models (LLMs). The dataset contains conversations between different characters in the *Harry Potter* novels and the associated sentiment perceptions between characters. Considering the advancement of LLMs in diverse domains [33], we leverage ChatGPT and GPT-4 to generate the seven-level persona information for each character by well-designed prompts. Finally, the utility of the persona information is evaluated with a neurosymbolic AI paradigm. The persona information is structured as tailored prompts, feeding into LLMs together with the conversations of the employed dataset. We analyze the variations in the sentiment analysis accuracy of AI after the inclusion or exclusion of different prompts.

We find that in the context of the conversations of *Harry Potter*, Culture, Vocation, and Ideology present the highest utility, resulting in the most accuracy gains in our research domain; All the seven aspects have positive utilities to the neurosymbolic AI system; Integration of personalized neurosymbolic knowledge into LLMs leads to biases towards the object and the subject on sentiment analysis task. It is reasonable that different types of persona information have different utilities in a specific domain because different individuals may prioritize certain aspects of persona information based on the context or the domain they are in. The sentiment perception reasoning may rely on different background knowledge in different scenarios.

The contribution of this work is threefold. (1) It introduces the novel concept of personalized sentiment analysis, designed to enhance AI’s understanding of the varied sentiment perceptions between individuals. (2) It presents a comprehensive framework derived from an analysis of literature, delineating seven levels of personalization for personalized sentiment analysis. (3) It performs an empirical study to assess the effectiveness of various persona information types within a conversational domain.

2 Literature Review

2.1 Sentiment Analysis with AI

Recently, there have been notable advancements in sentiment analysis, characterized by several key trends. Initially, the focus was on lexicon-based approaches to identify sentiment polarities such as positive, negative, and neutral in sentences or documents. The field then evolved by introducing concept-level sentiment analysis systems like pSenti [44] and SenticNet [8]. These systems represented a shift towards more advanced methods that combined lexicon-based and learning-based approaches. They presented greater accuracy in tasks such as sentiment polarity classification and sentiment intensity prediction, surpassing the capabilities of traditional lexicon-based systems [43].

Later, with the development of neural networks, the research focus of sentiment analysis shifts towards developing different learning frameworks to improve accuracy. Convolutional Neural Network-based supervised learning [24], transfer learning [11], adversarial training [15], meta-learning [17], prompt-based [38] algorithms were proposed for sentiment analysis. These research efforts address learning challenges in sentiment analysis, e.g., pattern discovery with label data, efficient learning from few-shot examples, robust representations, and domain adaption. During this period, research in multimodal [55] or cross-lingual [56] sentiment analysis was dynamic because it expanded the scope of sentiment analysis beyond English text. Recently, there has been a significant enrichment in the task setups of sentiment analysis. Researchers are no longer satisfied with simply predicting a sentiment polarity for an input text; they are extending the scope of sentiment analysis to include different levels of granularity and contextual awareness, e.g., aspect-based sentiment analysis [35] and opinion mining [40],

emotion detection [1], conversational sentiment analysis [28], sentiment analysis from electroencephalography (EEG) signals [25], facial expressions [10] or speech [30]. Another trend in sentiment analysis is that researchers paid more attention to the linguistic phenomena that likely affect sentiment analysis, e.g., metaphors [36], sarcasm [53], and ambiguous word senses [58]. Considering the impact of sentiment in broad domains, there are research papers studying sentiment analysis in different science domains, e.g., nature disaster [13], mental health [22], finance [31, 32], legislation [48] and education [2].

To sum up, previous research has addressed sentiment analysis by tackling learning challenges, enhancing sentiment analysis granularity, improving natural language understanding in learning systems, and grounding sentiment analysis in different downstream tasks. However, sentiment perception can vary subjectively in different contexts. There is limited research on personalized sentiment analysis that integrates various types of persona information. This motivates us to bridge this gap by forming a framework to identify the sources of subjectivity in sentiment analysis and developing a neurosymbolic system to process the task of personalized sentiment analysis.

2.2 Sources of Diversity in Sentiment Perception

The theory of mind (ToM) suggests that individuals understand that others may hold beliefs, desires, intentions, emotions, and thoughts that differ from their own [3]. Thus, we believe that multiple factors can influence individual sentiment perceptions. According to the theory of appraisal [41], opinions and sentiments arise not as direct responses to stimuli but as complicated evaluations incorporating subjective judgments across multiple levels. We assume that the factors influencing sentiment perception are hierarchically structured. This hierarchical structure ranges from general factors affecting large populations to specific factors influencing individuals. To explore this idea further in the context of sentiment analysis, we reviewed the following theoretical research.

ToM research is subject to humans and other species or intelligent agents. Early research found that chimpanzees may possess a preliminary form of ToM. The ability allows them to infer the mental states of others, like humans [47]. This ability to attribute intentions, knowledge, and beliefs to others suggests that chimpanzees possess the basic forms of social cognition. However, the following research has shown that there are great differences between humans and animals in terms of the depth and complexity of ToM [7]. For example, while chimpanzees can understand others by a perception-goal psychology, they do not have a fully developed belief-desire psychology like humans. With the development of LLMs, e.g., ChatGPT, researchers also extended ToM tests to AI. ChatGPT and GPT-4 were tasked with difficult questions that required inferring the counterfactual effects of actions on mental states [5]. The findings show that GPT-4 demonstrated strong abilities in these scenarios, possessing an advanced level of ToM.

Cultural norms, values, and beliefs significantly impact how people perceive and understand the world. Researchers compared students from Western (American) and Eastern (Indian) cultural backgrounds and discovered that Western participants had an independent self-construal and saw themselves as considerably more different from others [39]. Indian participants, on the other hand, perceived themselves as somewhat more alike to others, suggesting an interdependent self-perception. Additionally, studies show that different languages have different conceptions of emotion, with variations in the causes, evaluations, outcomes, modes of management and display, and even physiological responses linked to particular notions [45]. Sentiment perception can also be impacted by religious factors. The theory of cognitive dissonance suggests that individuals might adjust their opinions to match their religious beliefs in order to mitigate psychological discomfort [16]. This process can strengthen existing beliefs while eliminating conflicting opinions. Affirming religious beliefs can reduce the negative affect and emotional discomfort that occurs when individuals experience cognitive dissonance [6].

Individuals gain a sense of identity and self-esteem from association with various groups, including those based on occupational and educational backgrounds [50]. To maintain a positive social identity, individuals often adopt perspectives consistent with the norms and values of these groups. Related research also shows that people in high-status occupations tend to show more liberal attitudes toward social issues than people in low-status occupations [29]. Individuals may interpret information in a way that is consistent with their beliefs and values. This bias may lead individuals to view information that aligns with their ideology as more credible and trustworthy, while ignoring conflicting information [26]. Individuals with different ideologies may interpret identical information in divergent ways, thus delivering different opinions on the same target. For example, those with conservative ideologies often emphasize individual responsibility and individual rights, influencing their stance on welfare and healthcare-related policies [49]. In contrast, individuals with liberal ideologies often prioritize social justice and equality, leading to opposing opinions.

The way people construct their ideas is also influenced by certain personality traits [21]. For example, people who are open to new experiences are more likely to be receptive to new ideas and have flexible, open-minded perspectives [42]. On the other hand, those with higher conscientiousness typically base their beliefs on carefully analyzing the available data, leading to more thoughtful viewpoints. Additionally, research indicates that although intuitive thinkers depend more on heuristics and intuition, possibly producing subjective and prejudiced ideas, analytical thinkers often use deliberate, reflective thinking, leading to objective and evidence-driven judgments [46]. Sometimes, sentiment perceptions can be influenced by individuals' subjectivity, e.g., individuals who have had positive interactions with dogs are more likely to view dogs positively, contrasting with those whose experiences have been negative [51]. People tend to focus on information that aligns with their personal preferences and beliefs, potentially distorting their perception of emotions [4].

This tendency, termed confirmation bias, can impact how individuals interpret emotional events. Moreover, varied experiences and subjective feelings can influence using metaphorical language among individuals to express their opinions [27]. For example, financial analysts may employ different metaphors in their reports under different market conditions [34]. The public’s perception of different types of weather disasters is also reflected in their metaphorical expressions [37]. To sum up, theoretical research and empirical studies support that individuals’ sentiment perceptions are subject to multiple factors, including entity diversity that distinguishes between humans and other intelligent agents like animals and AI; culture, religion, vocation, ideology, personality, and subjectivity. These factors may impact personalized sentiment analysis in different scenarios. However, their collective impacts generally represent the complex interplay of individual characteristics and contextual influences on sentiment perception. These factors not only influence how individuals perceive and interpret the sentiment of a target but also alter the language they use to express subjective feelings. In sentiment analysis, understanding these factors is necessary for developing more personalized and context-sensitive systems.

3 Methodology

After reviewing relevant literature in psychology and cognitive science in Section 2.2, we define a hierarchical framework, containing factors that can impact individual sentiment perception and personalized sentiment analysis. This hierarchical framework is termed Personalized Sentiment Analysis Pyramid (see Figure 2), including persona aspects, e.g., entity, culture, religion, vocation, ideology, personality, and subjectivity. Entity refers to the differentiation between human genders and other intelligent agents. Culture represents how various cultures perceive concepts as positive or negative. Religion involves considering how specific religious beliefs can influence an individual’s opinions on certain topics. Vocation aids in understanding people’s opinions based on their occupation and educational background. Ideology involves political beliefs and social, economic, or philosophical viewpoints. Personality assists in categorizing concepts as positive or negative based on personality traits. Finally, subjectivity considers personal preferences and experiences. At the bottom of the pyramid, personalization is more general, e.g., entities of the same gender and species, such as males, females, AI, or other creatures, can share the same persona information. Personalization is more specific at the top layer, e.g., subjectivity level. The persona information can be tailored for individuals. Next, we use an LLM, i.e., GPT-4 Turbo, to analyze the persona information of our subjects related to the above seven aspects. LLMs were trained with broader sources. It has shown superior knowledge in diverse domains, including natural language understanding and generation, multilingual capabilities, commonsense, reasoning, and scientific task processing [33]. It was also suggested as a useful tool for survey research and persona information generation [20]. Thus, it is eligible for analyzing the personalities of a subject from different aspects.

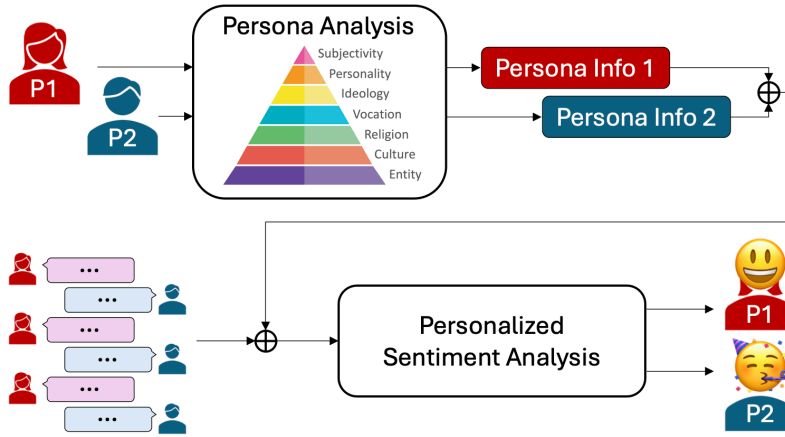


Fig. 3. Personalized sentiment analysis workflow. \oplus denotes textual concatenation.

Finally, we test LLM performance on personalized sentiment analysis tasks. The obtained persona information in the former step is used as symbolic knowledge, guiding the sentiment inference of an LLM. Since we combine the symbolic knowledge and the reasoning ability of neural network-based LLMs together, the methodology is neurosymbolic. The structured symbolic knowledge provides a clear and understandable reasoning basis, enhancing the interpretability and explainability of the system’s decisions. The testing data were sourced from novels, including dialogues with multi-turns. The task is to predict a speaker’s sentiment perceptions towards another speaker involved in the conversation. We hypothesize that the additional persona information has different utilities in this scenario because the intensity of the influence of personal characteristics on sentiment perception changes as the scene changes. We aim to evaluate the utilities of the ensemble and each type of persona information.

The overall workflow of our method can be viewed in Figure 3. Our task setup and computing pipeline represent a novel approach for several reasons. First, we do not focus on analyzing the sentiment of a conversation based solely on its semantic content. Instead, our goal is to analyze how one person perceives another’s sentiments. This means that even in a conversation that may seem neutral, there could still be negative sentiment if the individuals involved in the conversation do not like each other. Second, unlike traditional personalized AI techniques, such as user preference-based dialogue systems [60, 59], personality trait-based recommender systems [52], or annotation-subjectivity-driven sentiment analysis [57], our approach considers persona information from multiple aspects. This allows our system to incorporate a broader range of factors that may be informative for personalized sentiment analysis. Finally, our system prioritizes user subjectivity by generating personalized outputs based on different types of persona information, even when presented with the same dialogue input. This approach to human-computer interaction is more human-centric.

3.1 Persona Information Acquisition

Since we have defined seven aspects for persona analysis and our analytical subjects are characters from *Harry Potter* novels, we can consult the persona information for GPT-4 directly. We formulate the query template as follows.

[Goal]: I want to categorize the given object according to their `{aspect_term}` in the Harry Potter book series. Please suggest a description, one in a line, starting with "-" and surrounded by quotes ". For example: - "`{example_category}`"
 Do not output anything else.
 You may choose only one `{aspect_term}` from the following list: `{category_list}`
 Please categorize `{sample_in_prompt}` into one `{aspect_term}` according to their `{aspect_term}` in the Harry Potter book series.

For each persona aspect, the *Harry Potter* persona analysis-tailored definition can be viewed in Table 1.

Table 1. Persona aspect term for analyzing the characters in Harry Potter novels.

Aspect	Aspect Term
Entity	specie type and gender
Culture	cultural background
Religion	religious beliefs
Vocation	strong feeling of suitability for a particular career or occupation
Ideology	ideologies from aspects of political, social, epistemological, and ethical
Personality	MBTI personality type
Subjectivity	preferences and hobbies

Table 1 shows the aspect terms for the seven aspects we used to analyze the characters in the Harry Potter book series. To prevent overlap among the seven aspects and ensure that the LLM fully understands the meanings of the aspect terms, we provide a list of categories for the first six aspects (shown as follows) for the LLM to reference during inference. For “subjectivity”, we provide examples such as “Quidditch Seeker” and “Painting”, allowing the LLM to generate relevant answers openly.

- Entity: [specie type] Wizards and Witches, Muggles, Werewolves, Dragons, Hippogriffs, Basilisks, Trolls, Hags, Giants, Ghosts, House-elves, Goblins, Centaurs, Veela, Merpeople, Dementors, Vampires. [gender] male, female, or, inapplicable.
- Culture: Gryffindor, Slytherin, Hufflepuff, Ravenclaw, England, Scotland, Wales, Irish, French, Bulgarian, India, African, Romani, Middle Eastern, USA.
- Vocation: Auror, Healer, Transfigurers, Charms Experts, Diviners, Professor, Magizoologist, Potion Master, Curse Breaker, Metamorphmagi, Animagi,

- Occlumens, Legilimens, Runes Experts, Patronus Charm Casters, Unspeakable, Wandmaker, Broom Maker, Quidditch Player, Journalist, Shop Owner, Ministry Official, Librarian, Herbologist, Arithmancer, Servants, Metalworkers, Bankers, Underwater Dwellers, Companions or Pets.
- Religion: Good vs. Evil, Love vs. Indifference, Acceptance_Death vs. Fear_Death vs. Bravery_Death vs. Denial_Death vs. Honor_Death, Sacrifice vs. Selfishness, Redemption vs. Condemnation, Impartiality vs. Prejudice, Tolerance vs. Intolerance, Courage vs. Cowardice, Faith vs. Skepticism, Responsibility vs. Irresponsibility.
 - Ideology: Equality and Inclusivity vs. inequality and exclusivity, Reform vs. Satus Quo, Utilitarianism vs. Moral Absolutism, Knowledge vs. Ignorance, Loyalty and Community vs. disloyalty and individualism, Pragmatism vs. Idealism.
 - Personality: ESTJ, ENTJ, ESFJ, ENFJ, ISTJ, ISFJ, INTJ, INFJ, ESTP, ESFP, ENTP, ENFP, ISTP, ISFP, INTP, INFP.

3.2 Personalized Sentiment Analysis

We conduct personalized sentiment analyses using the information obtained from the persona and an LLM. Given the collection of the persona information (p_i) of a person related to the seven aspects ($p = \{p_1, p_2, \dots, p_i, \dots, p_7\}$), the scene illustration (e) where the conversation happens, the dialogues with multi-turns (d), the interlocutors (m, n), and task description (t), the task aims to predict the mutual sentiment perceptions, e.g., the sentiment perception of m (a perceiving subject) towards n (a perceiving object) ($s^{(m \rightarrow n)}$) and the sentiment perception of n towards m ($s^{(n \rightarrow m)}$). We fit all the aforementioned textual information into a prompt template ($template(\cdot)$), then ask an LLM to predict the sentiment perceptions ($LLM(\cdot)$) from the prompt ($prompt$).

$$prompt^{(m \leftrightarrow n)} = template(e^{(m, n)}, d^{(m, n)}, p^{(m)}, p^{(n)}, m, n, t) \quad (1)$$

$$s^{(m \rightarrow n)}, s^{(n \rightarrow m)} = LLM(prompt^{(m \leftrightarrow n)}) \quad (2)$$

The prompt template ($prompt^{(m \rightarrow n)}$) for inferring the sentiment perception of m towards n can be viewed below. In the prompt box, the content after [Goal] refers to the task description (t) that directs an LLM to deliver desired predictions, following a fixed structure. **scene** (e) is the background illustration in which the conversation takes place. **dialogue_sample** is the dialogue with multi-turns (d). **character_1** and **character_2** correspond to the interlocutors (m, n). **persona_1** and **persona_2** denote their persona information $p^{(m)}$ and $p^{(n)}$ that was obtained in Section 3.1.

[Goal] I want to classify the sentiment scores between the two characters in the Harry Potter book series based on their dialogue and persona.
Please suggest a sentiment score, one in a line, starting with "-" and surrounded by quotes "". For example:
- "<Harry to Hermione> 1"

- "<Hermione to Harry> 2"

The following shows the different meanings of the sentiment scores.
 Please select the sentiment score from the following options: -5: Vendetta; -4: Intentionally inflict harm; -3: Maliciously targeting and harm; -2: Deliberately bullying/deliberately targeting; -1: Rude/Frivolous/Mean characters; 0: Stranger/Neutral; 1: Normal/Polite; 2: Friendly; 3: Kind; 4: Close; 5: Devoted.

[Scene] {scene}
 [Dialogue] {dialogue_sample}
 [Persona] {character_1}: {persona_1}. {character_2}: {persona_2}.

Please classify the sentiment scores between the two characters {character_1} and {character_2} based on the given dialogue and their personas.

4 Experiment

4.1 Research Questions

In this work, we aim to explore the following research questions:

1. What is the utility of using the ensemble of the seven levels of personalization?
2. What is the utility of using individual personalization?
3. How does personalization impact sentiment analysis accuracy across different types of entity and culture factors?

These research questions are explored in Sections 5.1-5.3, respectively.

4.2 Dataset

Our personalized sentiment analysis uses the Harry Potter Dataset (HPD) [9]. It was developed to enhance the alignment of conversation agents with fictional characters from Harry Potter novels. It includes annotating relationships and character attributes that evolve over the storyline. HPD includes background information, such as conversation scenes, speaker identities, and character attributes, to enable dialogue agents to generate replies consistent with the Harry Potter universe. In contrast to our structured persona analysis, the character attributes in the dataset were not derived from the same set of analytical aspects. Thus, we did not use their character attribute descriptions. We leverage their affection labels that indicate the sentiment intensity of a perceiving subject to another perceiving object as our sentiment intensity labels. In our classification task, positive labels correspond to the sentiment intensity, ranging from -5 to -1; negative labels correspond to the sentiment intensity, ranging from 1 to 5. A neutral label corresponds to the sentiment intensity of 0. Our method is evaluated using the English version of the original HPD. The statistics of our employed data and the sentiment intensity distribution are shown in Figure 4.

Item	No. of samples
Dialogues	1,191
Positive sentiment	4,932
Negative sentiment	730
Neutral sentiment	8
Interlocutors	94

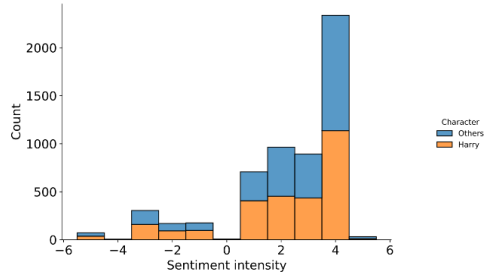


Fig. 4. Dataset statistics and sentiment intensity distribution.

4.3 Large Language Models

The persona information was queried from GPT-4 Turbo. The personalized sentiment analysis was evaluated with GPT-4 Turbo and GPT-3.5 Turbo (i.e., ChatGPT), respectively. GPT-4 Turbo is an upgraded version of GPT-3.5 Turbo. Both LLMs were developed by Open AI, pre-trained with Transformer-based deep neural networks and a large number of corpus resources. By a preliminary test, we found that these LLMs have rich knowledge about *Harry Potter* novels.

4.4 Evaluation Metrics

We evaluate the personalized sentiment analysis performance using two types of metrics. The F1 score and accuracy are used to evaluate the accuracy of sentiment polarity classification. In this task, LLMs aim to predict whether the sentiment perception is positive, negative, or neutral. Mean Squared Error (MSE) is used to evaluate the accuracy of sentiment intensity prediction. Considering the potential failure of LLMs to respond in the required format, the Answer Rate is also integrated into the evaluation. It measures the percentage of queries successfully answered by the LLM. In this task, LLMs aim to predict the sentiment intensity, ranging from -5 to 5. Since the sentiment intensity between the same perceiving subject and perceiving object is identical across all dialogues in the dataset, we combine all dialogues between two identical interlocutors as a united input. A correct prediction is defined as the predicted sentiment label (either a sentiment polarity or a sentiment intensity score) of a perceiving subject towards a perceiving object matching the ground truth label.

5 Results

5.1 The Ensemble Utility of the Seven Levels of Personalization

We compared the sentiment analysis results of GPT-3.5 and GPT-4 with and without the ensembled seven levels of personalization in Table 2. The ensemble utility of seven levels of personalization improved the performance of GPT-3.5 on

sentiment analysis. For GPT-4, the inclusion of the seven levels of personalization seems to contribute to a slight improvement in $F1(p:=pos)$ and accuracy and mean-square error(MSE) or even a marginal decrease in $F1(p:=neg)$ and Macro F1. Taking into consideration the answer rate, however, we found the accuracy of GPT-4 on the whole dataset turns into $0.8977 \times 0.8155 = 0.7321$ while the one of GPT-4 w/ p1:7 becomes $0.9058 \times 0.9492 = 0.8598$. Thus, the adjusted accuracy increase of 0.1277 demonstrates the ensemble effectiveness of the seven levels of personalization on sentiment analysis task.

Table 2. Personalized sentiment analysis. $p:=pos$ means the positive sentiment is defined as positive labels for computing F1; $p:=neg$ means the negative sentiment is defined as positive labels for computing F1.

	F1 (p:=pos)	F1 (p:=neg)	Macro. F1	Acc	MSE	Answer Rate
GPT-3.5	0.9381	0.5340	0.4907	0.8686	0.3586	0.6522
w/ p1:7	0.9484	0.5859	0.5114	0.8936	0.3293	0.6860
Delta	0.0103	0.0519	0.0207	0.0250	-0.2036	0.0338
GPT-4	0.9569	0.7372	0.5704	0.8977	0.1860	0.8155
w/ p1:7	0.9631	0.7092	0.5630	0.9058	0.1812	0.9492
Delta	0.0062	-0.028	-0.0074	0.0081	-0.0048	0.1337

5.2 The Utility Analysis of the Individual Personalization

Table 3 shows the positive influence of each individual personalization on the performance of GPT-3.5 in the personalized sentiment analysis tasks. Among them, Culture, Vocation, Ideology, and Subjectivity strengthened the performance of GPT-3.5 by a significant margin, while Entity, Religion, and Personality contributed to a relatively less improvement. There may be a thought-provoking rationale for such a discrepancy. The Harry Potter book series is deeply rooted in a rich culture, ideology, and subjectivity backdrops, thereby offering a rich tapestry of themes and narratives that resonate deeply with readers’ or even sentiment intensity annotators’ own values. For example, someone who values loyalty and friendship may possibly echo characters like Harry and Ron. Hence, they may easily capture Harry’s negative sentiment towards Peter Pettigrew, who betrayed his friends, James and Lily Potter. Consequently, an LLM knowing these factors (culture, ideology, and subjectivity) may understand characters’ sentiments more precisely by resonating with characters’ values.

5.3 The Personalization Utility Analysis by Entity and Culture

The results presented in Section 5.2 are readily comparable, as the growths in F1, Macro. F1, accuracy, and answer rate are consistent. Unlike Sections 5.1 and 5.2, performance summarized by different categories of entity and culture aspects is more sensitive to the answer rate, since we investigate the results by

Table 3. Personalized sentiment analysis by persona types.

	F1 (pos=pos)	F1 (pos=neg)	Macro. F1	Acc-all	MSE	Answer Rate
GPT-3.5	0.9381	0.5340	0.4907	0.8686	0.3586	0.6522
w/ p1	0.9478	0.5553	0.5010	0.8825	0.3200	0.6769
w/ p2	0.9508	0.5793	0.5138	0.8926	0.2941	0.6801
w/ p3	0.9413	0.5403	0.4939	0.8841	0.3736	0.6713
w/ p4	0.9504	0.5542	0.5015	0.8886	0.3055	0.7171
w/ p5	0.9456	0.5532	0.4996	0.8884	0.3419	0.6938
w/ p6	0.9490	0.5619	0.5073	0.8868	0.3108	0.6667
w/ p7	0.9502	0.5722	0.5074	0.8897	0.3046	0.6780

breaking down the characters into finer-grained groups instead of treating them as a whole. Therefore, we calculated the evaluation metrics based on all the query samples in Tables 4 and 5, by setting the missing result with a fixed value (100 in this paper) out of the scope of ground-truth labels. Moreover, we detailed the sentiment analysis results by presenting the metrics of each category group, both as subjects (conveying sentiments to Harry) and objects (receiving sentiments from Harry).

In general, including seven levels of personalization completely enhanced the performance of groups Ghosts, Acromantula, Veela, Centaurs, and French. For entity breakdowns, the negative effects on GPT-4 of including seven aspects occur only in the cases where the entity groups (Muggles, Wizards and Witches, House-elves, and Goblins) play as an object. For culture breakdowns, the negative impacts of including seven aspects on GPT-4 are primarily observed when the entity groups (England, Gryffindor, Ravenclaw, Slytherin, and Hufflepuff) function as objects. However, effects related to these breakdowns are observed in only two cases (Ravenclaw and Hufflepuff) when they serve as subjects. For both breakdowns (entity and culture), the negative effects on GPT-3.5 of including seven aspects occur relatively more often and mainly when the entity group plays as a subject. The above observation highlights the bias of LLMs (such as GPT-3.5 and GPT-4) towards the subject and the object on sentiment analysis, especially when these models address personalized neurosymbolic knowledge. Additionally, comparing the results of GPT-3.5 and GPT-4, we observed that GPT-3.5, when integrated with personalized neurosymbolic knowledge, achieved comparable or even superior performance to GPT-4. This is validated by the results from several Entity or Culture groups including Muggles, Ghosts, Acromantula, Veela, and England.

6 Conclusion

In this work, we studied the task of personalized sentiment analysis. Personas were widely studied in commercial domains and web research [18, 19]. Unlike conventional sentiment analysis tasks that aim to analyze sentiment by the meanings of the text, personalized sentiment analysis targets to analyzing the individual sentiment perception.

Table 4. Personalized sentiment analysis by entity breakdowns (adjusted by answer rate). * denotes the inclusion of persona information lowers the performance.

Entity	Perc.	Model	F1(p:=pos)	F1(p:=neg)	Macro. F1	Acc
Muggles	Sub.	G3.5	0.16	0.5203	0.2268	0.3208
		w/ p1:7	0.2105	0.6119	0.2742	0.4057
		G4	0	0.4833	0.1611	0.2736
	Obj.	w/ p1:7	0.1111	0.6277	0.2463	0.4151
		G3.5	0.0909	0.0833	0.0581	0.0566
		w/ p1:7	0.32	0.1237	0.1479	0.1321
Giants	Sub.	G4	0.0952	0.198	0.0978	0.1038
		w/ p1:7	0.2069	0.1237*	0.1102	0.0849*
		G3.5	0.875	0	0.2917	0.7778
	Obj.	w/ p1:7	0.8529*	0	0.2843*	0.7436*
		G4	0.9646	0	0.3215	0.9316
		w/ p1:7	0.9957	0	0.3319	0.9915
Wizards and Witches	Sub.	G3.5	0.8641	0	0.288	0.7607
		w/ p1:7	0.8529*	0	0.2843*	0.7436*
		G4	0.9646	0	0.3215	0.9316
	Obj.	w/ p1:7	0.9957	0	0.3319	0.9915
		G3.5	0.7468	0.578	0.4416	0.5941
		w/ p1:7	0.7837	0.6724	0.4853	0.6483
House-elves	Sub.	G4	0.865	0.7664	0.555	0.7644
		w/ p1:7	0.9521	0.82	0.6027	0.9061
		G3.5	0.7497	0.2465	0.3321	0.5713
	Obj.	w/ p1:7	0.7793	0.2189*	0.3328	0.6146
		G4	0.8679	0.6435	0.5079	0.7483
		w/ p1:7	0.9543	0.6323*	0.5331	0.8836
Ghosts	Sub.	G3.5	0.7727	0.5714	0.4481	0.6176
		w/ p1:7	0.6829*	0.375*	0.3526*	0.5*
		G4	0.9412	0.9231	0.6214	0.8824
	Obj.	w/ p1:7	0.9615	1	0.6538	0.9412
		G3.5	0.6818	0.1538	0.2786	0.4706
		w/ p1:7	0.6818	0.1538	0.2786	0.4706
Acromantula	Sub.	G4	0.875	0.5882	0.4877	0.7647
		w/ p1:7	0.8148*	0.1538*	0.3229*	0.6765*
		G3.5	0.88	0	0.2933	0.7857
	Obj.	w/ p1:7	0.9231	0	0.3077	0.8571
		G4	0.9231	0	0.3077	0.8571
		w/ p1:7	0.963	0	0.321	0.9286
Veela	Sub.	G3.5	0.8333	0	0.2778	0.7143
		w/ p1:7	0.88	0	0.2933	0.7857
		G4	0.963	0	0.321	0.9286
	Obj.	w/ p1:7	0.963	0	0.321	0.9286
		G3.5	0.6667	0	0.2222	0.5
		w/ p1:7	0.6667	0	0.2222	0.5
Centraurs	Sub.	G4	0	0	0	0
		w/ p1:7	0	0	0	0
		G3.5	0	0	0	0
	Obj.	w/ p1:7	0	0	0	0
		G4	0	0	0	0
		w/ p1:7	0	0	0	0
Werewolves	Sub.	G3.5	0.5161	0	0.172	0.3478
		w/ p1:7	0.7895	0	0.2632	0.6522
		G4	0.6471	0	0.2157	0.4783
	Obj.	w/ p1:7	0.7568	0	0.2523	0.6087
		G3.5	0.5625	0	0.1875	0.3913
		w/ p1:7	0.7895	0	0.2632	0.6522
Goblins	Sub.	G4	0.6471	0	0.2157	0.4783
		w/ p1:7	0.7568	0	0.2523	0.6086
		G3.5	1	0	0.3333	1
	Obj.	w/ p1:7	1	0	0.3333	1
		G4	1	0	0.3333	1
		w/ p1:7	1	0	0.3333	1
Gobblins	Sub.	G3.5	0.0571	0.5	0.1857	0.0541
		w/ p1:7	0*	0.8	0.2667	0.0541
		G4	0.1111	1	0.3704	0.1351
	Obj.	w/ p1:7	0.1111	1	0.3704	0.1351
		G3.5	0	0	0	0
		w/ p1:7	0	0.3333	0.1111	0.027
Gobblins	Sub.	G4	0.1176	0.5714	0.2297	0.1081
		w/ p1:7	0.1176	0.75	0.2892	0.1351
		G3.5	0.5455	0	0.1818	0.3333
	Obj.	w/ p1:7	0.3333*	0.3333	0.2222	0.2222*
		G4	0.4444	0.4	0.2815	0.3333
		w/ p1:7	0.5714	0.5714	0.381	0.4444
Obj.	G3.5	0.7273	0	0.2424	0.4444	
	w/ p1:7	0.5455*	0	0.1818*	0.3333*	
	G4	0.8	0.6667	0.4889	0.6667	
Gobblins		w/ p1:7	0.7143*	0*	0.3571*	0.5556*

Table 5. Personalized sentiment analysis by culture breakdowns (adjusted by answer rate). * denotes the inclusion of persona information lowers the performance.

Culture	Perc.	Model	F1(p:=pos)	F1(p:=neg)	Macro. F1	Acc
England	Sub.	G3.5	0.5135	0.5362	0.3499	0.3889
		w/ p1:7	0.4762*	0.6104	0.3622	0.4306
		G4	0.6087	0.4776	0.4227	0.3819
	Obj.	w/ p1:7	0.597	0.6133	0.4451	0.4722
		G3.5	0.4752	0.0971	0.1908	0.2014
		w/ p1:7	0.4808	0.1321	0.2043	0.2222
Gryffindor	Sub.	G4	0.5833	0.2456	0.2763	0.2431
		w/ p1:7	0.5854*	0.1651*	0.2502*	0.2292*
		G3.5	0.7609	0.1798	0.3135	0.614
	Obj.	w/ p1:7	0.7985	0.2985	0.3657	0.664
		G4	0.8772	0.2593	0.3788	0.7805
		w/ p1:7	0.9632	0.3478	0.437	0.9276
Ravenclaw	Sub.	G3.5	0.7705	0.1282	0.2296	0.623
		w/ p1:7	0.8042	0.1017*	0.302	0.6694
		G4	0.8802	0.4667	0.4592	0.785
	Obj.	w/ p1:7	0.9652	0.4364*	0.4847	0.9303
		G3.5	0.6443	0	0.2148	0.466
		w/ p1:7	0.5775*	0.069	0.2155	0.4078*
Slytherin	Sub.	G4	0.8047	0	0.2682	0.6602
		w/ p1:7	0.7976*	0.2	0.3325	0.6602
		G3.5	0.5913	0.15	0.2471	0.3592
	Obj.	w/ p1:7	0.6461	0.1081*	0.2514	0.4272
		G4	0.8125	0.3256	0.3794	0.5728
		w/ p1:7	0.8593	0.3111*	0.3901	0.6311
Hufflepuff	Sub.	G3.5	0.5143	0.6981	0.4041	0.543
		w/ p1:7	0.32*	0.7966	0.3722*	0.6471
		G4	0.5143	0.8822	0.4655	0.7692
	Obj.	w/ p1:7	0.6061	0.933	0.513	0.8643
		G3.5	0.3111	0.2797	0.1969	0.2127
		w/ p1:7	0.2281*	0.2712*	0.1664*	0.2036*
Bulgarian	Sub.	G4	0.6667	0.7352	0.4673	0.6018
		w/ p1:7	0.5763*	0.7152*	0.4305*	0.5882*
		G3.5	0.6897	0.8572	0.5156	0.5542
	Obj.	w/ p1:7	0.7333	0.75*	0.4944*	0.6024
		G4	0.7937	0.9333	0.5757	0.6867
		w/ p1:7	0.8905	0.7778*	0.5561*	0.8193
French	Sub.	G3.5	0.6316	0.2857	0.3058	0.4578
		w/ p1:7	0.7097	0.1538*	0.2878*	0.5422
		G4	0.748	0.5	0.416	0.6024
	Obj.	w/ p1:7	0.8467	0.4*	0.4156*	0.7349
		G3.5	0.4286	0	0.1429	0.2727
		w/ p1:7	0.4*	0	0.1333*	0.2273*
Irish	Sub.	G4	0.4	0	0.1333	0.2773
		w/ p1:7	0.6207	0	0.2069	0.4091
		G3.5	0.625	0	0.2083	0.4545
	Obj.	w/ p1:7	0.5806*	0	0.1935*	0.4091*
		G4	0.5806	0	0.1935	0.4091
		w/ p1:7	0.625	0	0.2083	0.4545
G4	Sub.	G3.5	0.5	0	0.1667	0.3333
		w/ p1:7	0.7692	0	0.2564	0.625
		G4	0.6286	0	0.2095	0.4583
	Obj.	w/ p1:7	0.7368	0	0.2456	0.5833
		G3.5	0.5455	0	0.1818	0.375
		w/ p1:7	0.7692	0	0.2564	0.625
G4	Sub.	G4	0.6286	0	0.2095	0.4583
		w/ p1:7	0.7568	0	0.2456	0.5833
		G3.5	0.5455	0.2	0.2485	0.2667
	Obj.	w/ p1:7	0.4*	0.2	0.2*	0.2*
		G4	0.4286	0.2	0.2095	0.2667
		w/ p1:7	0.6316	0.2	0.2772	0.4667
Irish	Sub.	G3.5	0.5455	0.2	0.2485	0.2667
		w/ p1:7	0.4*	0*	0.1333*	0.1333*
		G4	0.4615	0.2	0.2205	0.2667
	Obj.	w/ p1:7	0.6316	0.2	0.2772	0.4667

The difference is that an identical statement can yield the same sentiment prediction by its meaning. In contrast, different people may perceive the message differently based on their own personal preferences, personality traits, beliefs, background, etc. To this end, we devised a framework, termed the Personalized Sentiment Analysis Pyramid, for tackling all these different facets through seven different levels of personalization, namely: Entity, Culture, Religion, Vocation, Ideology, Personality, and Subjectivity.

We evaluated the framework with a dialogue dataset sourced from *Harry Potter* novels. The evaluation showed that personalized neurosymbolic knowledge, i.e., seven levels of personalization, augmented LLMs’ performance on sentiment analysis. We also analyzed the utility of each persona aspect and found that each individual persona aspect can augment sentiment intensity classification results. Finally, we investigated the influence of persona information on several character groups in the *Harry Potter* novels. Results showed that including persona information elevated the performance of groups Ghost, Acromantula, Veela, Centaurs, and French. Furthermore, a bias of LLMs fed with personalized neurosymbolic knowledge towards subject and object groups is observed.

In future work, we plan to develop more robust persona information parsers and classifiers to extract information related to the defined persona aspects from different modalities. With persona information, we will also conduct a wide range of personalized cognitive computing tasks, including investigating how different people use different metaphors to communicate different perspectives, experiences, and emotions, revealing the nuanced ways in which language shapes and reflects cultural, social, and individual identities.

Acknowledgments

This research/project is supported by the Ministry of Education, Singapore under its MOE Academic Research Fund Tier 2 (STEM RIE2025 Award MOE-T2EP20123-0005) and by the RIE2025 Industry Alignment Fund – Industry Collaboration Projects (IAF-ICP) (Award I2301E0026), administered by A*STAR, as well as supported by Alibaba Group and NTU Singapore.

References

1. Muhammad Abdul-Mageed and Lyle Ungar. EmoNet: Fine-grained emotion detection with gated recurrent neural networks. In *Proceedings of ACL*, pages 718–728, Vancouver, Canada, 2017. Association for Computational Linguistics.
2. Nabeela Altrabsheh, Mohamed Medhat Gaber, Mihaela Cocea, et al. SA-E: sentiment analysis for education. *Frontiers in Artificial Intelligence and Applications*, 255:353–362, 2013.
3. Ian A Apperly and Stephen A Butterfill. Do humans have two systems to track beliefs and belief-like states? *Psychological Review*, 116(4):953, 2009.

4. Gordon H Bower. Mood and memory. *American Psychologist*, 36(2):129, 1981.
5. Sébastien Bubeck, Varun Chandrasekaran, Ronen Eldan, Johannes Gehrike, Eric Horvitz, Ece Kamar, Peter Lee, Yin Tat Lee, Yuanzhi Li, Scott Lundberg, Harsha Nori, Hamid Palangi, Marco Tulio Ribeiro, and Yi Zhang. Sparks of artificial general intelligence: Early experiments with GPT-4. *arXiv preprint arXiv:2303.12712*, 2023.
6. Christopher T Burris, Eddie Harmon-Jones, and W Ryan Tarpley. “by faith alone”: Religious agitation and cognitive dissonance. *Basic and Applied Social Psychology*, 19(1):17–31, 1997.
7. Josep Call and Michael Tomasello. Does the chimpanzee have a theory of mind? 30 years later. *Trends in Cognitive Sciences*, 12(5):187–192, 2008.
8. Erik Cambria, Xulang Zhang, Rui Mao, Melvin Chen, and Kenneth Kwok. SenticNet 8: Fusing emotion AI and commonsense AI for interpretable, trustworthy, and explainable affective computing. In *Proceedings of International Conference on Human-Computer Interaction (HCI)*, 2024.
9. Nuo Chen, Yan Wang, Haiyun Jiang, Deng Cai, Yuhan Li, Ziyang Chen, Longyue Wang, and Jia Li. Large language models meet Harry Potter: A dataset for aligning dialogue agents with characters. In *EMNLP Findings*, pages 8506–8520, 2023.
10. Debishree Dagar, Abir Hudait, Hrudaya Kumar Tripathy, and MN Das. Automatic emotion detection model from facial expression. In *ICACCCT*, pages 77–85, 2016.
11. Xin Luna Dong and Gerard De Melo. A helping hand: Transfer learning for deep sentiment analysis. In *Proceedings of ACL*, pages 2524–2534, 2018.
12. Cuc Duong, Qian Liu, Rui Mao, and Erik Cambria. Saving earth one tweet at a time through the lens of artificial intelligence. In *IJCNN*, pages 1–9, 2022.
13. Cuc Duong, Vethavikashini Chithrara Raghuram, Amos Lee, Rui Mao, Gianmarco Mengaldo, and Erik Cambria. Neurosymbolic AI for mining public opinions about wildfires. *Cognitive Computation*, 2024.
14. Chunxiao Fan, Jie Lin, Rui Mao, and Erik Cambria. Fusing pairwise modalities for emotion recognition in conversations. *Information Fusion*, 106:102306, 2024.
15. William Fedus, Ian Goodfellow, and Andrew M Dai. MaskGAN: Better text generation via filling in the _ . In *ICLR*, 2018.
16. Leon Festinger. *A Theory of Cognitive Dissonance*. Stanford University Press, 1957.
17. Kai He, Rui Mao, Tieliang Gong, Chen Li, and Erik Cambria. Meta-based self-training and re-weighting for aspect-based sentiment analysis. *IEEE Transactions on Affective Computing*, 14(3):1731–1742, 2023.
18. Bernard J Jansen. *Understanding user-web interactions via web analytics*. Springer Nature, 2022.
19. Bernard J Jansen, Soon-Gyo Jung, Lene Nielsen, Kathleen W Guan, and Joni Salminen. How to create personas: Three persona creation methodologies with implications for practical employment. *Pacific Asia Journal of the Association for Information Systems*, 14(3):1, 2022.
20. Bernard J Jansen, Soon-gyo Jung, and Joni Salminen. Employing large language models in survey research. *Natural Language Processing Journal*, 4:100020, 2023.
21. Bernard J Jansen, Joni Salminen, Soon-gyo Jung, and Kathleen Guan. *Data-driven personas*. Springer Nature, 2022.
22. Shaoxiong Ji, Xue Li, Zi Huang, and Erik Cambria. Suicidal ideation and mental disorder detection with attentive relation networks. *Neural Computing and Applications*, 34(13):10309–10319, 2022.

23. Kian Kenyon-Dean, Eisha Ahmed, Scott Fujimoto, Jeremy Georges-Filteau, Christopher Glasz, Barleen Kaur, Auguste Lalande, Shruti Bhandari, Robert Belfer, Nirmal Kanagasabai, et al. Sentiment analysis: It’s complicated! In *Proceedings of ACL*, pages 1886–1895, 2018.
24. Yoon Kim. Convolutional neural networks for sentence classification. In *Proceedings of EMNLP*, pages 1746–1751, 2014.
25. Sudhanshu Kumar, Mahendra Yadava, and Partha Pratim Roy. Fusion of EEG response and sentiment analysis of products review to predict customer satisfaction. *information Fusion*, 52:41–52, 2019.
26. Ziva Kunda. The case for motivated reasoning. *Psychological Bulletin*, 108(3):480, 1990.
27. George Lakoff and Mark Johnson. *Metaphors We Live by*. University of Chicago press, 1980.
28. Wei Li, Liuyao Zhu, Rui Mao, and Erik Cambria. SKIER: A symbolic knowledge integrated model for conversational emotion recognition. *Proceedings of AAAI*, 37(11):13121–13129, 2023.
29. Seymour Martin Lipset. *The Confidence Gap: Business, Labor, and Government in the Public Mind*. The Johns Hopkins University Press, 1987.
30. Zhiyun Lu, Liangliang Cao, Yu Zhang, Chung-Cheng Chiu, and James Fan. Speech sentiment analysis via pre-trained features from end-to-end ASR models. In *ICASSP*, pages 7149–7153, 2020.
31. Yu Ma, Rui Mao, Qika Lin, Peng Wu, and Erik Cambria. Multi-source aggregated classification for stock price movement prediction. *Information Fusion*, 91:515–528, 2023.
32. Yu Ma, Rui Mao, Qika Lin, Peng Wu, and Erik Cambria. Quantitative stock portfolio optimization by multi-task learning risk and return. *Information Fusion*, 104:102165, 2024.
33. Rui Mao, Guanyi Chen, Xulang Zhang, Frank Guerin, and Erik Cambria. GPTEval: A survey on assessments of ChatGPT and GPT-4. In *Proceedings of LREC-COLING*, pages 7844–7866, 2024.
34. Rui Mao, Kelvin Du, Yu Ma, Luyao Zhu, and Erik Cambria. Discovering the cognition behind language: Financial metaphor analysis with MetaPro. In *ICDM*, pages 1211–1216, 2023.
35. Rui Mao and Xiao Li. Bridging towers of multi-task learning with a gating mechanism for aspect-based sentiment analysis and sequential metaphor identification. In *Proceedings of AAAI*, volume 35, pages 13534–13542, 2021.
36. Rui Mao, Xiao Li, Kai He, Mengshi Ge, and Erik Cambria. MetaPro Online: A computational metaphor processing online system. In *Proceedings of ACL*, volume 3, pages 127–135, 2023.
37. Rui Mao, Qika Lin, Qiawen Liu, Gianmarco Mengaldo, and Erik Cambria. Understanding public perception towards weather disasters through the lens of metaphor. In *Proceedings of IJCAI*, 2024.
38. Rui Mao, Qian Liu, Kai He, Wei Li, and Erik Cambria. The biases of pre-trained language models: An empirical study on prompt-based sentiment analysis and emotion detection. *IEEE Transactions on Affective Computing*, 14(3):1743–1753, 2023.
39. Hazel Rose Markus and Shinobu Kitayama. Culture and the self: Implications for cognition, emotion, and motivation. In *College student development and academic life*, pages 264–293. Routledge, 2014.
40. Edison Marrese-Taylor, Juan D Velásquez, and Felipe Bravo-Marquez. A novel deterministic approach for aspect-based opinion mining in tourism products reviews. *Expert Systems with Applications*, 41(17):7764–7775, 2014.

41. James R Martin and Peter R White. *The Language of Evaluation*, volume 2. Springer, 2003.
42. Robert R McCrae and Paul T Costa. Validation of the five-factor model of personality across instruments and observers. *Journal of Personality and Social Psychology*, 52(1):81, 1987.
43. Walaa Medhat, Ahmed Hassan, and Hoda Korashy. Sentiment analysis algorithms and applications: A survey. *Ain Shams Engineering Journal*, 5(4):1093–1113, 2014.
44. Andrius Mudinas, Dell Zhang, and Mark Levene. Combining lexicon and learning based approaches for concept-level sentiment analysis. In *ISDOMW*, 2012.
45. Aneta Pavlenko. Emotion and emotion-laden words in the bilingual lexicon. *Bilingualism: Language and cognition*, 11(2):147–164, 2008.
46. Gordon Pennycook, James Allan Cheyne, Nathaniel Barr, Derek J Koehler, and Jonathan A Fugelsang. On the reception and detection of pseudo-profound bullshit. *Judgment and Decision Making*, 10(6):549–563, 2015.
47. David Premack and Guy Woodruff. Does the chimpanzee have a theory of mind? *Behavioral and Brain Sciences*, 1(4):515–526, 1978.
48. Sven-Oliver Proksch, Will Lowe, Jens Wäckerle, and Stuart Soroka. Multilingual sentiment analysis: A new approach to measuring conflict in legislative speeches. *Legislative Studies Quarterly*, 44(1):97–131, 2019.
49. Charles S Taber and Milton Lodge. Motivated skepticism in the evaluation of political beliefs. *American Journal of Political Science*, 50(3):755–769, 2006.
50. Henri Tajfel, John C Turner, William G Austin, and Stephen Worchel. An integrative theory of intergroup conflict. *Organizational Identity: A Reader*, 56(65):9780203505984–16, 1979.
51. Peter Tyrer, Geoffrey M Reed, and Mike J Crawford. Classification, assessment, prevalence, and effect of personality disorder. *The Lancet*, 385(9969):717–726, 2015.
52. Hsin-Chang Yang and Zi-Rui Huang. Mining personality traits from social messages for game recommender systems. *Knowledge-Based Systems*, 165:157–168, 2019.
53. Tan Yue, Rui Mao, Heng Wang, Zonghai Hu, and Erik Cambria. KnowleNet: Knowledge fusion network for multimodal sarcasm detection. *Information Fusion*, 100:101921, 2023.
54. Tan Yue, Xuzhao Shi, Rui Mao, Zonghai Hu, and Erik Cambria. SarcNet: A multilingual multimodal sarcasm detection dataset. In *Proceedings of LREC-COLING*, pages 14325–14335, 2024.
55. Amir Zadeh, Minghai Chen, Soujanya Poria, Erik Cambria, and Louis-Philippe Morency. Tensor fusion network for multimodal sentiment analysis. In *Proceedings of EMNLP*, pages 1103–1114, 2017.
56. Xulang Zhang, Rui Mao, and Erik Cambria. Multilingual emotion recognition: Discovering the variations of lexical semantics between languages. In *IJCNN*, 2024.
57. Xulang Zhang, Rui Mao, and Erik Cambria. SenticVec: Toward robust and human-centric neurosymbolic sentiment analysis. In *ACL Findings*, 2024.
58. Xulang Zhang, Rui Mao, Kai He, and Erik Cambria. Neurosymbolic sentiment analysis with dynamic word sense disambiguation. In *EMNLP Findings*, pages 8772–8783, 2023.
59. Luyao Zhu, Wei Li, Rui Mao, and Erik Cambria. HIPPL: Hierarchical intent-infering pointer network with pseudo labeling for consistent persona-driven dialogue generation. *IEEE Computational Intelligence Magazine*, 2024.
60. Luyao Zhu, Wei Li, Rui Mao, Vlad Pandealea, and Erik Cambria. PAED: Zero-shot persona attribute extraction in dialogues. In *Proceedings of ACL*, volume 1, pages 9771–9787, 2023.