

PsihoRo: Depression and Anxiety Romanian Text Corpus

Alexandra Ciobotaru^{1,2}, Ana-Maria Bucur³, Liviu P. Dinu^{1,4}

¹University of Bucharest, ²Druid AI, ³Università della Svizzera italiana,

⁴Human Language Technologies Research Center

^{1,4}90 Panduri Street, Bucharest, Romania

²1-5 Costache Negri Street, Bucharest, Romania,

³Via Buffi 13, Lugano, Switzerland

alexandra.zbaganu@s.unibuc.ro, ana-maria.bucur-cosma@usi.ch, ldinu@fmi.unibuc.ro

Abstract

Psychological corpora in NLP are collections of texts used to analyze human psychology, emotions, and mental health. These texts allow researchers to study psychological constructs, identify patterns related to mental health problems and analyze emotional language. However, collecting accurate mental health data from social media can be challenging due to the assumptions made by data collectors. A more effective approach involves gathering data through open-ended questions and then assessing participants' mental health status using self-report screening surveys. This method was successfully employed for English, a language with a lot of psychological NLP resources. However, the same cannot be stated for Romanian, which currently has no open-source mental health corpus. To address this gap, we have collected the first open-source corpus focused on depression and anxiety in Romanian, by utilizing a form with 6 open-ended questions along with the standardized PHQ-9 and GAD-7 screening questionnaires. Although the PsihoRo corpus contains texts from only 205 respondents, it represents an important first step toward understanding and analyzing mental health issues within the Romanian population. We employ statistical analysis, text analysis using Romanian LIWC, emotion detection, and topic modeling to identify the most important features of this newly introduced resource for the NLP community. The data is publicly available at <https://huggingface.co/datasets/Alegzandra/PsihoRo>.

Keywords: mental health, Romanian, emotion detection, LIWC, topic modeling

1. Introduction

Mental health research has increasingly used data from social media platforms, which often present a curated view of individuals' personal lives. While social media can provide valuable insights for mental health studies (Liu et al., 2022; Rai et al., 2024; Bucur et al., 2025), there is a risk of positivity bias, where users tend to share mostly positive aspects of their experiences (Salopek and Eastin, 2024). Another important source for mental health research is clinical data, which typically contains sensitive information and is often not publicly accessible (Puspo et al., 2026).

Most research on mental health is conducted using data in English, which may overlook important nuances present in other languages. Previous studies suggest that features predictive of mental health issues in English may not translate effectively to other languages or cultures (Rai et al., 2024; Bucur et al., 2026). For example, the use of the first-person pronoun "I" has been identified as a strong predictor of depression (Rude et al., 2004), and this feature is frequently analyzed in English-based research. However, the relationship between pronoun usage and depression severity can vary significantly across different demographic groups and languages (Rai et al., 2024). While the pronoun "I" is often used as an indicator of depres-

sion in English, its use in other languages requires special consideration due to linguistic differences. For instance, previous research has shown that this association has not been observed among Romanian speakers (Trifu et al., 2024). Romanian is a pro-drop language, which allows for ambiguity in subject pronoun usage and may contribute to a lower frequency of the personal pronoun "I" in everyday communication.

Despite the availability of data in the Romanian language related to emotions, such as the RED dataset (Ciobotaru and Dinu, 2021), REDv2 (Ciobotaru et al., 2022), and the BRIGHTER multilingual dataset (Muhammad et al., 2025), which includes Romanian data, there is a scarcity of mental health-related resources. Very few clinical data sources (Trifu et al., 2024) and posts from medical forums (Briciu and Lupea, 2018; Dudău and Sava, 2022) exist, but unfortunately are not publicly accessible.

This study aims to address this gap by creating the first publicly available dataset in Romanian, specifically designed to analyze emotional expressions in the context of depression and anxiety. By analyzing the linguistic characteristics and emotional patterns in Romanian, our research aims to contribute to a deeper understanding of mental health expressions in languages other than English. Through this initiative, we seek to promote more inclusive research in mental health, paving the way

for targeted interventions and support across diverse linguistic communities.

2. Related Work

Previous research has gathered open-ended responses from participants to explore how expressions of mental health influence language. [Yoon et al. \(2024\)](#) created open-ended personality questions to explore how natural language processing can aid in predicting neuroticism from people's responses. The authors gathered over 30,000 sentences written by 425 adults in Korean and used KoBERT to see which types of questions best revealed neuroticism and related traits like depression and dependency. [Van der Vegt and Kleinberg \(2023\)](#) created the multi-modal Real World Worry Waves Dataset (RW3D), in English, which tracks the psychological effects of the COVID-19 pandemic in the UK over three consecutive years, 2020, 2021 and 2022, through repeated surveys containing also free-text responses. [Dai et al. \(2022\)](#) used open-ended interview responses from more than 58,000 job candidates along with self-assessed personality (HEXACO) and employed "InterviewBERT" to deduce personality traits from textual in English. [Sikström et al. \(2023\)](#) used open-ended questionnaires to study young and middle-aged adults' responses regarding their own mental health, by employing NLP computational methods. The authors analyzed 1,106 responses and found semantic differences between the two study groups.

In terms of combining emotion detection and LIWC term analysis in mental health texts, several studies employ such analysis in English: [Rahman et al. \(2024\)](#), [Zhang et al. \(2023\)](#), [Alshouha et al. \(2024\)](#). [Beech et al. \(2025\)](#) combined LIWC features ([Boyd et al., 2022](#)), sentiment analysis and NLP to explore the relationship between language and mental health issues such as depression, anxiety, and stress during the COVID-19 pandemic. They analyzed written narratives from 883 participants, alongside their DASS-21 scores, and found an association between increased symptom severity and greater use of first-person pronouns and negative emotion words. [Kjell et al. \(2021\)](#) collected responses in English from 411 participants who answered word-response questions related to their mental health and completed the GAD-7 and PHQ-9 rating scales. Their analyses show that individuals' language reflects their mental health status, and combining verbal responses with quantitative measures can enhance the comprehension and assessment of the risk of depressive and anxiety symptoms in both clinical practice and research contexts.

While several datasets have been created using text gathered from open-ended questionnaires,

combined with screening tools to assess the risk of mental health issues, most of these resources are focused on the English language. Currently, there are no open-source psychological corpora available in Romanian, and this was the main driving force behind our study, along with our interest in applying emotional analysis to such a corpus.

3. Data Collection

As opposed to social media collected data, which is collected without participants' consent, we collected data from participants who provided their informed consent for data collection ([Benton et al., 2017](#)). The data collection methodology consists of a completely anonymous survey, organized in three parts. The first part of the survey consists of six open-ended questions, designed to gather extensive written responses. Out of these six questions, three touch on positive topics and three on negative topics. The survey was written in Romanian for Romanian speakers only. We further present the six open-ended questions, in both English and Romanian:

1. To what extent do you think social media affects your emotional state, especially when you see negative news? "În ce măsură crezi că rețelele sociale îți afectează starea emoțională, mai ales când vezi știri negative?" (negative)
2. What is an important lesson about resilience and adaptation that you have learned recently? "Care este o lecție importantă despre reziliență și adaptare pe care ai învățat-o recent?" (positive)
3. How much do you think politics affects people's mental health? Can you describe a time when you felt stressed because of political decisions? "Cât de mult crezi că politica afectează sănătatea mentală a oamenilor? Poți să detaliezi un moment în care ai simțit stres din cauza deciziilor politice?" (negative)
4. What role do you think community or social support plays in maintaining mental health during difficult times? "Ce rol crezi că are comunitatea sau sprijinul social în menținerea sănătății mentale în vremuri dificile?" (positive)
5. How were you affected by discussions about war, the economic crisis, elections, or other tense topics? "Cum ai fost afectat de discuțiile despre război, criza economică, alegeri sau alte subiecte tensionate?" (negative)
6. How do you recharge your batteries after a stressful period? "Cum îți reîncarci bateriile după o perioadă stresantă?" (positive)

Usually, psychological corpora for NLP are gathered from social media through searching for different keywords in online posts that could indicate depression and anxiety. However, after consulting with medical professionals in the psychiatric field, we have concluded that this method of collecting psychological data may not be entirely accurate. Thus, we added two more parts to our survey: one assessing respondents' risk of depression using the PHQ-9 (Kroenke et al., 2001), and another assessing their risk of anxiety using the GAD-7 (Spitzer et al., 2006). Both GAD-7 and PHQ-9 contain questions (7 for GAD-7 and 9 for PHQ-9) that ask the participants how often they have experienced specific problems, and each item is answered on a 4-point Likert scale: Not at all = 0, Several days = 1, More than half the days = 2, Nearly every day = 3. The severity cut-offs for depression based on PHQ-9 are: 0-4: Minimal, 5-9: Mild, 10-14: Moderate, 15-19: Moderately severe, 20-27: Severe, while the severity cut-offs for anxiety based on GAD-7 are: 0-4: Minimal, 5-9: Mild, 10-14: Moderate, 15-19: Moderately severe, 20-27: Severe.

The resulting questionnaire is completely anonymous, and no personal data is requested or stored (such as name, surname, email address, or IP). We did not collect demographic data for the same reason: to build trust with respondents and keep the time required to complete the form short.

The resulting data consists of 205 responses, collected over an eight-month period from March 2025 to October 2025. This time frame includes two major stressful events for Romanians—presidential elections and the tax increase burden generated by the country's new leadership.

4. Methodology

Our research consists of gathering insights from PsihoRo using descriptive statistics to analyze the correlation between depression and anxiety levels among respondents across three main population groups, computed based on severity cut-offs. According to the literature in psychology (Kroenke et al., 2001), a PHQ-9 score of 10 or higher indicates the potential presence of major depressive disorder and prompts further assessment. For the GAD-7 questionnaire, a cut-off score of 10 has also been proposed as a threshold for identifying potential generalized anxiety disorder, warranting further evaluation (Spitzer et al., 2006).

Based on these cut-off scores, three population groups were derived from our dataset:

- Depression Risk Group: 37 respondents scored higher than 10 on PHQ-9 (18%);
- Anxiety Risk Group: 36 respondents scored higher than 10 on GAD-7 (17.56%);

- No Risk Group: 28 respondents scored 0 for at least one out of the two questionnaires (13.66%).

Linguistic Inquiry and Word Count (LIWC) (Boyd et al., 2022) is a lexicon and a text analysis software that provides psychologically meaningful categories and is widely used in computational mental health research (Rude et al., 2004; Liu et al., 2022; Hur et al., 2024; Trifu et al., 2024). The official LIWC software is available for analyzing textual data in the English language; however, validated translations of the lexicon are also available in several languages. In our experiments, we use the Romanian version of LIWC (Ro-LIWC2015) (Dudău and Sava, 2022). The 2015 version of LIWC includes 98 categories that relate to linguistic dimensions and grammar, as well as psychological processes, including affective, cognitive, social, perceptual, and biological processes.

Previous research on the English language has indicated that individuals with mental health issues tend to exhibit specific language patterns. These patterns include a higher frequency of using the personal pronoun “I” (Rude et al., 2004), increased use of words with negative connotations (Fekete, 2002), and a tendency to use words with absolute meanings (Al-Mosaiwi and Johnstone, 2018). These findings highlight the importance of language analysis for understanding mental health and motivate us to use LIWC to examine participants' language patterns. After checking for differences in LIWC categories across PHQ-9 and GAD-7 scores, we wanted to test their discriminatory power for classification. To this end, we employed a LightGBM classifier (Ke et al., 2017) and trained it on LIWC results to predict a binary label for depression and anxiety, using a cut-off score of 10. The model achieved accuracies of 0.85 for depression and 0.83 for anxiety. To determine which LIWC features have the greatest influence on predictions, we used SHapley Additive exPlanations (SHAP) (Lundberg and Lee, 2017).

To complement these analyses, we have created an emotion detection model by fine-tuning the Romanian BERT (Dumitrescu et al., 2020) on the REDv2 dataset (Ciobotaru et al., 2022), for multi-label text classification, using the Huggingface transformers pipeline.¹ We used this model to predict the emotions in PsihoRo, thus gathering emotional information for the text written by each population group mentioned above. The resulting model achieved an F1 score of 66.85% on the REDv2 test set.

We also performed topic modeling to identify the main themes in the PsihoRo corpus, using the `stm` R package. First, the texts were tokenized, lower-

¹<https://huggingface.co/docs/transformers>

	PHQ-9	GAD-7
Count	205	205
Mean	6.00	5.78
SD	4.83	4.95
Min	0	0
25%	3	2
50% (Median)	5	5
75%	8	7
Max	23	21

Table 1: Descriptive statistics for PsihoRo

cased, and cleaned by removing punctuation, numbers, symbols and stopwords. Next, we created a Document-Feature Matrix (DFM) to represent term frequencies across texts. To enhance model performance and reduce noise, we excluded rare terms that appeared in fewer than 3 documents and the very common terms that appeared in more than 90% of the texts. Further, the cleaned DFM was transformed into the format required by the Structural Topic Model (STM), and we trained the model with $K = 5$ topics using the Spectral initialization method.

5. Experiments and Results

5.1. PsihoRo Descriptive Statistics

In Table 1, we show mean, median (50% percentile), min, max, 25% and 75% percentiles, and the standard deviation for both anxiety and depression scores. In the collected data, the maximum value for GAD-7 (anxiety) is 21, which is the maximum that can be achieved, and the maximum score for PHQ-9 (depression) is 23, out of the maximum possible of 27.

We analyzed the length of responses to determine whether individuals at risk of depression or anxiety tend to write less. Our findings indicate that this is not the case for Romanian. The computed T-statistic was 1.55 when comparing the depression risk group with the group not at risk, and 2.26 when comparing the anxiety risk group with the group not at risk. These results suggest that there is no significant difference in response length between these populations.

In Figure 1, we present the distribution plots of both PHQ-9 and GAD-7 scores. The data indicate that the distributions are skewed to the left, with most observations clustered at lower values and a longer tail extending toward the higher end of the scale. This suggests that the majority of respondents experience minimal to mild risk of depression and anxiety.

The Pearson correlation coefficient (r) between the PHQ-9 and GAD-7 total scores is 0.761

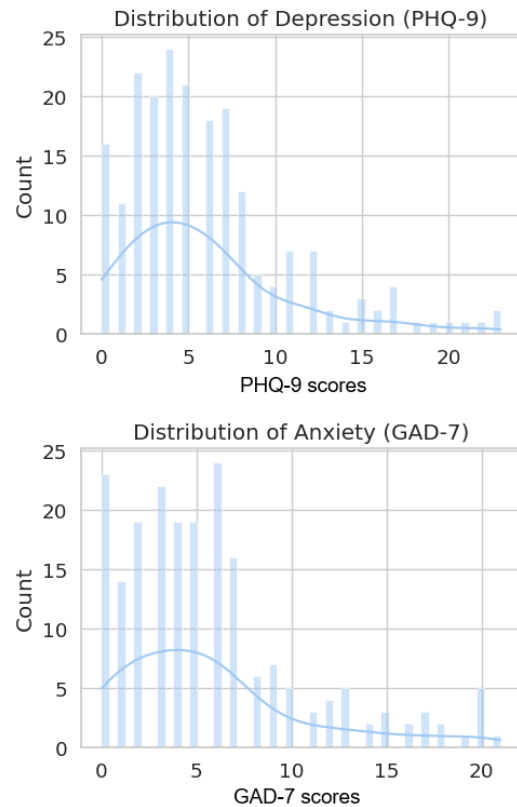


Figure 1: Distributions for GAD-7 and PHQ-9




Figure 2: Correlation between Anxiety and Depression Scores.

($p < 0.001$), indicating that anxiety and depression are strongly correlated in our data. This finding is illustrated in Figure 2 and aligns with the results of several other studies (Kalin, 2020; Hoying et al., 2020; Jacobson and Newman, 2014).

5.2. LIWC

We employ LIWC to analyze language patterns in participants' open-ended responses and correlate them with scores from the PHQ-9 and GAD-7. In Table 2, we present the Pearson correlation coefficients between LIWC categories and the symptom scores. Positive values indicate that higher symp-

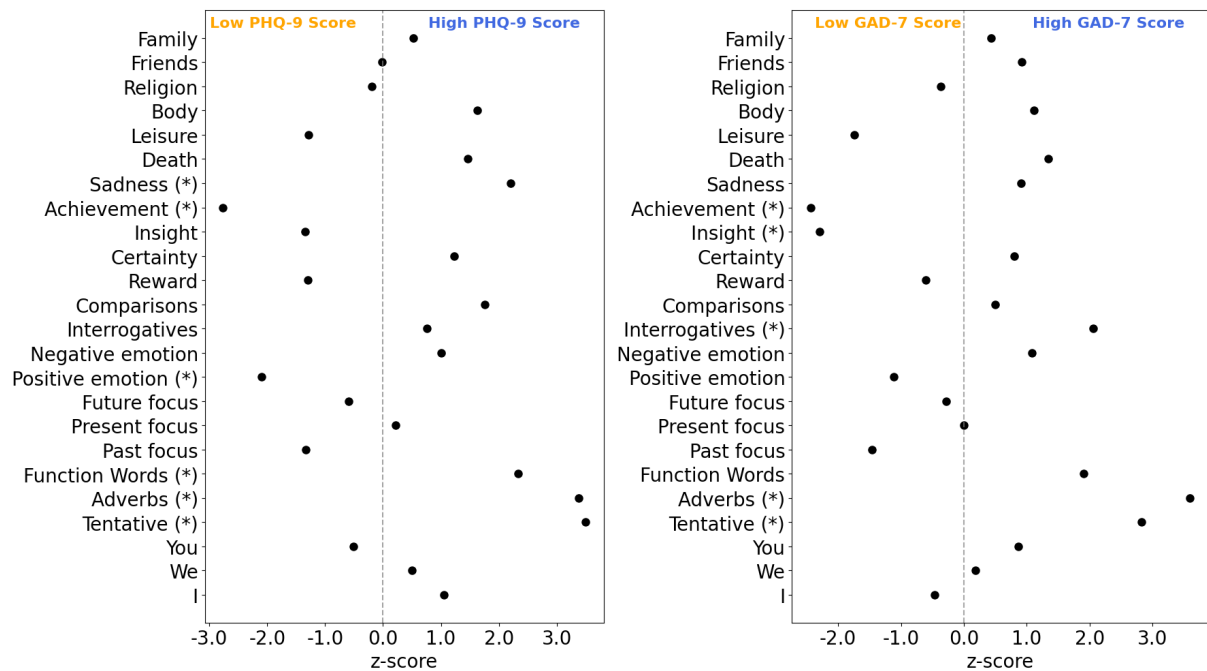


Figure 3: Differences in selected LIWC linguistic categories between participants with low vs. high PHQ-9 (left) and GAD-7 (right) scores. * denotes statistically significant differences ($p < 0.05$) measured using the Mann-Whitney U test. Positive values indicate greater relative use among participants with higher symptom scores.

LIWC Dimension	Examples	PHQ-9	GAD-7
I	I, me, mine	.018	-.046
We	we, us, our	.013	-.039
You	you, your	.006	.084
Verb	eat, come, carry	-.087	-.091
Past focus	ago, did, talked	-.064	-.075
Present focus	today, is, now	-.019	-.034
Future focus	may, will, soon	-.027	-.055
Adjectives	free, happy, long	.118	.120
Negations	no, not, never	-.156*	-.167*
Positive emotion	love, nice, sweet	-.122	-.057
Negative emotion	hurt, ugly, nasty	.028	.028
Anxiety	worried, fearful	-.042	.010
Anger	hate, kill, annoyed	.107	.105
Sadness	crying, grief, sad	.074	.045
Family	daughter, dad, aunt	.079	.145*
Friends	buddy, neighbor	-.026	.043
Female references	girl, her, mom	.061	.074
Male references	boy, his, dad	.156*	.177*
Comparatives	greater, best, after	.167*	.081
Causation	because, effect	-.186**	-.180**
Tentative	maybe, perhaps	.204**	.213**
Insight	think, know	-.075	-.101
Assent	agree, OK, yes	.009	.065
Perceptual processes	look, heard, feeling	.009	.033
Body	cheek, hands, spit	.200**	.180**
Health	clinic, flu, pill	-.051	.035
Achievement	win, success, better	-.163*	-.192**
Leisure	cook, chat, movie	-.145*	-.169*
Death	bury, coffin, kill	.138*	.029

Note. * $p < .05$, ** $p < .01$.

Table 2: Correlations between PHQ-9, GAD-7, and LIWC dimensions.

tom scores are associated with greater use of that linguistic category, while negative values indicate the opposite.

Surprisingly, the use of negations is negatively correlated with both PHQ-9 and GAD-7 scores, suggesting that participants with higher depression and anxiety scores tend to use fewer negations in their responses. Individuals who score higher on the PHQ-9 and GAD-7 tend to use more words in the Tentative category, which indicates uncertainty, as evidenced by correlation coefficients of .204** for PHQ-9 and .213** for GAD-7. Comparatives also show a positive correlation with the depression score (.167*), indicating comparative thinking, which is often linked to depressive rumination (Nolen-Hoeksema et al., 2008). However, another marker for rumination, the high frequency of Past focus verbs, is not observed in our data. There is a positive correlation between the use of Body-related words and PHQ-9 (.200**) and GAD-7 (.180**) scores, indicating that participants with higher scores frequently use more language related to bodily sensations. This may reflect increased somatic awareness, a common trait in depression and anxiety (Penninx et al., 2013). Participants with higher PHQ-9 and GAD-7 scores discuss Achievement less frequently (-.163* and -.192**, respectively), which could indicate lower motivation. In addition, the negative correlation with Leisure-related words (-.145* and -.169*, respectively) may stem from anhedonia, a common symptom of depression and anxiety (Grillo, 2016). Words in the Death category show a significant positive correlation with depression

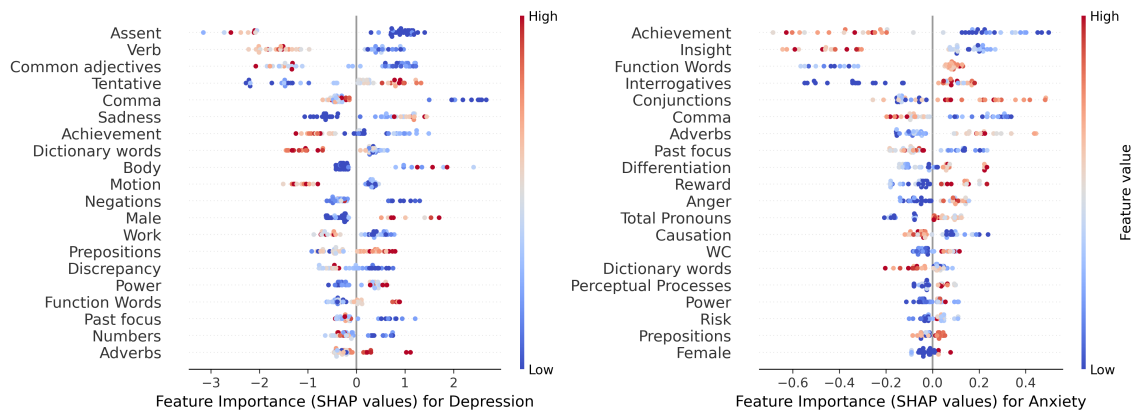


Figure 4: Feature importance for LIWC categories analyzed using SHAP.

(.138*), but not with anxiety, indicating a greater preoccupation with mortality in participants with higher depression scores. In contrast, the Family category is positively correlated with anxiety scores ($r = .145^*$), suggesting that individuals with higher anxiety are more focused on family-related issues.

To better understand the differences between higher and lower symptom scores for depression and anxiety, we conducted an additional analysis of the LIWC categories. According to the cut-off values for both the PHQ-9 and GAD-7, we define a low PHQ-9 score as less than 10, while a score of 10 or higher is considered high. In Figure 3, the goal was to identify linguistic markers that differentiate individuals with higher vs. lower symptom scores. Statistically significant differences are indicated with an asterisk (*), as determined by the Mann-Whitney U test. Positive values indicate greater relative use among participants with higher symptom scores, while negative values indicate greater use among low-symptom participants. The analysis produced results consistent with the correlation analysis, particularly regarding Achievement and tentativeness in relation to depression and anxiety. However, for depression, we observed a decreased use of Positive emotion words and an increased usage of Sadness words, which aligns with existing literature in psychology (Trifu et al., 2024). In addition, the language of participants with high GAD-7 scores showed a higher frequency of words in the Interrogatives category, indicating a tendency to question or seek clarity.

To determine which LIWC features have the greatest influence on predictions, we used SHapley Additive exPlanations (SHAP) (Lundberg and Lee, 2017), as explained in Section 4, with the results presented in Figure 4. The importance of the top 20 features aligns with our correlation analysis and highlights the differences between low and high scores for both depression and anxiety. For predicting depression, low values in the LIWC features for

Assent, Verb, and Common Adjectives, along with high values in Tentative words, contribute to the model's positive predictions. In the case of anxiety, low values in LIWC features related to Achievement and Insight, combined with high values in Function words, Interrogatives, and Conjunctions, also have a higher contribution to a positive prediction.

The LIWC analyses presented in this section demonstrate that lexical indicators can effectively capture underlying psychological patterns related to depression and anxiety. However, the most important feature used in mental health research, the pronoun “I”, appears to be less useful for the Romanian language, likely because it is a pro-drop language. This finding is consistent with the work of Trifu et al. (2024), which also indicated that this marker does not generalize well in a sample of Romanian patients with major depression.

5.3. Emotion Analysis

Using the emotion detection model described in Section 4, we predicted emotions for all PsihoRo texts, then aggregated the results using radar charts to visualize overall emotional patterns. In Figure 5, we show the emotion predictions across all PsihoRo texts, as well as for the negative-topic (pessimistic) questions (1, 3, 5) and the positive-topic (optimistic) questions (2, 4, 6).

Surprise is an emotion that is completely absent from the PsihoRo data. The following emotions are more common, in order of appearance: anger, joy, fear, sadness, trust, and, most commonly, neutral. The emotional pattern for the more pessimistic questions is displayed in the center of Figure 5, while the emotional pattern for the more optimistic questions is displayed on the right. Emotions such as sadness and anger are on the rise on the Pessimistic Questions, while joy and trust are more prevalent on the Optimistic Questions.

The emotions present in the texts written by the




Figure 5: Emotion predictions across: all PsihoRo texts (left), negative-topic (pessimistic) questions (middle), and positive-topic (optimistic) questions (right).

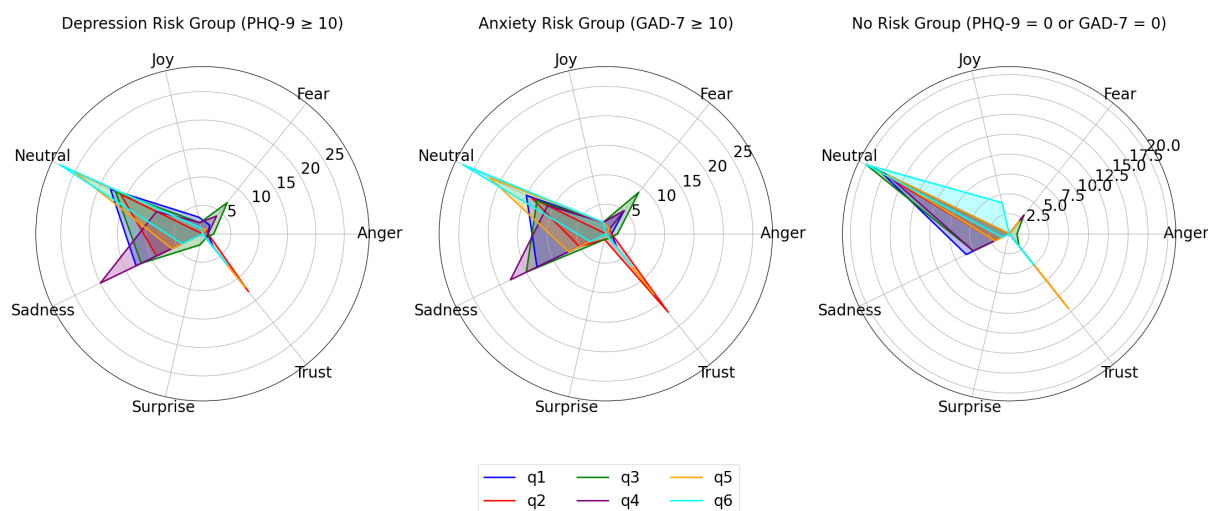


Figure 6: Emotion predictions across the three studied population groups.

three primary demographic groups mentioned at the beginning of Section 4 are presented in Figure 6. While the neutral emotion is strongly represented across all groups, the population at risk for depression shows a high prevalence of sadness. In addition, the group at risk for anxiety exhibits emotions of sadness, fear and trust.

5.4. Topic Modeling

We applied the topic modeling analysis described in Section 4 for each of the three population groups studied in this work. The main topics are derived from the six open questions asked in the form. However, some differences appear among the three groups of respondents, as discussed in the following subsections.

5.4.1. No Risk Group

Topic 1 highlights how individuals perceive the current political and social period in a critical and sometimes negative way. The words “*negative*”, “*decisions*”, “*politics*”, “*measure*”, “*news*” and “*period*” indicate a concern for the effects of government decisions and how they are reflected in the media. The presence of the term “*sport*” suggests a possible comparison between different areas of public life.

Topic 2 reflects the complex interaction between information conveyed through news, political life, and how they influence the general state of the population. Terms such as “*politics*”, “*news*”, “*state*”, “*people*” and “*politics*” indicate an increased attention to how media discourse shapes collective perceptions and emotions.

Topic 3 addresses how individuals interpret and

process information related to elections and political decisions. The terms “*measure*”, “*depends*”, “*people*”, “*time*”, “*elections*” and “*example*” suggest a personal analysis of the electoral context and an attempt to understand the relationship between information, time and decision.

Topic 4 highlights the importance of authentic human interaction and the balance between the inner and social dimensions of the individual. The words “*listen*”, “*outer*”, “*inner*”, “*community*” and “*matter*” suggest an openness to empathetic communication and highlight the value of social belonging.

Topic 5 focuses on the emotional dimension of everyday experiences, indicating how individuals manage their affective states and reactions in common situations. Words such as “*affect*”, “*try*”, “*day*”, “*go*” and “*enough*” express a constant effort to adapt and maintain psychological balance in the face of daily pressures by self-regulation and personal reflection.

5.4.2. Depression Risk Group

Topic 1 captures the perception of fatigue, daily routine and the constant effort that individuals put into their personal and professional lives. Associated terms such as “*sleep*”, “*hard*”, “*walking*” and “*people*” indicate a general state of exhaustion and the need for balance between personal and social life.

Topic 2 highlights the importance of psychological resilience and the ability to adapt to stressful or unpredictable situations. Words such as “*resilience*”, “*situation*”, “*time*”, “*community*” and “*important*” suggest a solution-oriented approach and maintaining a state of balance in contexts of uncertainty.

Topic 3 focuses on how individuals relate to political processes and the electoral climate. The terms “*politics*”, “*elections*”, “*people*”, “*moment*” and “*listening*” indicate both interest in political events and their influence on social attitudes and conversations.

Topic 4 captures the potentially harmful impact of social networks on interpersonal relationships and community cohesion. The presence of the words “*negative*”, “*social*”, “*networks*”, “*community*” and “*socialization*” underlines the perception of a degradation of authentic communication and well-being caused by the digital environment.

Topic 5 reflects the concern for affective states and how social and economic factors contribute to the intensification of collective anxiety. The terms “*feel*”, “*affect*”, “*cause*”, “*anxiety*”, “*country*” and “*family*” suggest an emotional perspective on everyday experiences and the general social context.

5.4.3. Anxiety Risk Group

Topic 1 captures how major events, such as war, economic crises or socio-political changes, affect the perception of daily life. Words such as “*affect*”, “*crisis*”, “*war*”, “*elections*”, “*life*” and “*time*” indicate the concern for external instability and its influence on personal balance.

Topic 2 highlights the close connection between political factors and individual emotional state. Terms such as “*politics*”, “*stress*”, “*affect*”, “*feel*”, “*cause*” and “*family*” suggest a direct relationship between socio-political events and the level of anxiety or psychological tension felt.

Topic 3 addresses the role of community and collective resilience in the face of the constant flow of negative information from the media. The words “*resilience*”, “*news*”, “*negative*”, “*community*”, “*important*” and “*trying*” indicate an attempt to find the balance between information and protecting one’s mental state.

Topic 4 explores the effects of social media on psychological state, with a focus on anxiety and social pressure. Terms such as “*social*”, “*anxiety*”, “*state*”, “*hard*”, “*networks*” and “*present*” suggest a connection between the intensive use of social networks and the degradation of emotional state. The topic reflects how virtual interactions can accentuate psychological discomfort.

Topic 5 emphasizes the connection between people, community and the sense of belonging. The words “*people*”, “*community*”, “*listen*”, “*dear*” and “*happen*” highlight the importance of close relationships and solidarity in maintaining emotional balance.

6. Conclusion and Future Work

In this paper, we presented PsihoRo, the first open-source dataset for mental health in Romanian. The dataset was collected from 205 participants who filled in a comprehensive form that included both open-ended questions for text collection and validated questionnaires, specifically the PHQ-9 and GAD-7, to evaluate their risk of depression and anxiety.

We analyzed the collected data using the psychologically meaningful categories from LIWC, emotion analysis and topic modeling. We observed differences in LIWC categories related to body-related words, which may indicate somatic awareness, and tentative words, which indicate uncertainty, that correlated positively with both depression and anxiety scores. A higher preoccupation with mortality was also observed in relation to depression, while anxiety scores were positively correlated with a preoccupation with family. While LIWC analyses showed that lexical indicators can effectively capture underlying psychological patterns related to

depression and anxiety, we did not find any association between the linguistic marker “I” and depression or anxiety in the Romanian language, consistent with previous research (Trifu et al., 2024). The emotion analysis revealed a prevalence of sadness in the population group at risk for depression, as well as fear in the population group at risk for anxiety, which aligns with previous research (Low et al., 2020; Atapattu et al., 2022; Varastehnezhad et al., 2025). Topic modeling showed that although some perceptions and emotions are captured in textual responses, a multi-perspective approach that goes beyond self-expressed text is necessary to fully comprehend these patterns. By offering high-quality data with precise labels for anxiety and depression risk, PsihoRo seeks to supplement this intricate research.

Ethical Considerations and Limitations

The data collection received approval from the Ethics Committee of the University of Bucharest. All respondents provided their informed consent before participating in the survey, which was completely anonymous. No personal information, such as names, surnames, email addresses, or IP addresses, was requested or stored.

We believe that increasing the number of respondents through crowdsourcing or similar platforms could significantly enhance PsihoRo. We aim to implement this in future versions of this resource.

Acknowledgments

The authors would like to thank psychologist Virginia Neacșu for her valuable feedback and suggestions throughout the development of this work.

This work was partially supported by Druid AI, University of Bucharest and the Ministry of Education and Research, CNCS-UEFISCDI, project SIROLA, number PN-IV-P1- PCE-2023-1701, within PNCDI IV, and by the project "Romanian Hub for Artificial Intelligence - HRIA", Smart Growth, Digitization, and Financial Instruments Program, 2021-2027, MySMIS no. 334906.

Bibliographical References

Mohammed Al-Mosaiwi and Tom Johnstone. 2018. In an absolute state: Elevated use of absolutist words is a marker specific to anxiety, depression, and suicidal ideation. *Clinical psychological science*, 6(4):529–542.

Bashar Alshouha, Jesús Serrano-Guerrero, Francisco Chiclana, Francisco P. Romero, and José A.

Olivas. 2024. [Bioemodetector: A flexible platform for detecting emotions from health narratives](#). *SoftwareX*, page 101670.

Thushari Atapattu, Mahen Herath, Charitha Elvitigala, Piyanjali de Zoysa, Kasun Gunawardana, Menasha Thilakaratne, Kasun de Zoysa, and Katrina Falkner. 2022. [EmoMent: An emotion annotated mental health corpus from two South Asian countries](#). In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 6991–7001, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.

Abigail Beech, Haoxue Fan, Jocelyn Shu, Javiera Oyarzun, Peter Nadel, Olivia T. Karaman, Sophia Vranos, Elizabeth A. Phelps, and M. Alexandra Kredlow. 2025. [Using natural language processing to identify patterns associated with depression, anxiety, and stress symptoms during the covid-19 pandemic](#). *Journal of Affective Disorders*, 376:113–121.

Adrian Benton, Glen Coppersmith, and Mark Dredze. 2017. Ethical research protocols for social media health research. In *Proceedings of the first ACL workshop on ethics in natural language processing*, pages 94–102.

Ryan L Boyd, Ashwini Ashokkumar, Sarah Seraj, and James W Pennebaker. 2022. The development and psychometric properties of liwc-22. *Austin, TX: University of Texas at Austin*, 10:1–47.

Anamaria Briciu and Mihaiela Lupea. 2018. Studying the language of mental illness in romanian social media. In *2018 IEEE 14th International Conference on Intelligent Computer Communication and Processing (ICCP)*, pages 21–28. IEEE.

Ana-Maria Bucur, Andreea-Codrina Moldovan, Krutika Parvatikar, Marcos Zampieri, Ashiqur R. KhudaBukhsh, and Liviu P. Dinu. 2025. [On the state of nlp approaches to modeling depression in social media: A post-covid-19 outlook](#). *IEEE Journal of Biomedical and Health Informatics*, pages 1–13.

Ana-Maria Bucur, Marcos Zampieri, Tharindu Ranasinghe, and Fabio Crestani. 2026. A survey on multilingual mental disorders detection from social media data. In *Proceedings of the 19th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*.

Alexandra Ciobotaru, Mihai Vlad Constantinescu, Liviu P. Dinu, and Stefan Dumitrescu. 2022. [RED v2: Enhancing RED dataset for multi-label emotion detection](#). In *Proceedings of the Thirteenth*

- Language Resources and Evaluation Conference*, pages 1392–1399, Marseille, France. European Language Resources Association.
- Alexandra Ciobotaru and Liviu P Dinu. 2021. Red: A novel dataset for romanian emotion detection from tweets. In *Proceedings of the international conference on recent advances in natural language processing (RANLP 2021)*, pages 291–300.
- Yimeng Dai, Madhura Jayaratne, and Buddhi Jayatilleke. 2022. [Explainable personality prediction using answers to open-ended interview questions](#). *Frontiers in Psychology*, 13.
- Diana Paula Dudău and Florin Alin Sava. 2022. The development and validation of the romanian version of linguistic inquiry and word count 2015 (ro-liwc2015). *Current Psychology*, 41(6):3597–3614.
- Stefan Dumitrescu, Andrei-Marius Avram, and Sampo Pyysalo. 2020. [The birth of Romanian BERT](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 4324–4328, Online. Association for Computational Linguistics.
- Sandor Fekete. 2002. The internet-a new source of data on suicide, depression and anxiety: a preliminary study. *Archives of Suicide Research*, 6(4):351–361.
- Lorenzo Grillo. 2016. [A possible role of anhedonia as common substrate for depression and anxiety](#). *Depression Research and Treatment*, 2016:1598130.
- Jacqueline Hoying, Bernadette Mazurek Melnyk, Elizabeth Hutson, and Angela Tan. 2020. [Prevalence and correlates of depression, anxiety, stress, healthy beliefs, and lifestyle behaviors in first-year graduate health sciences students](#). *Worldviews on Evidence-Based Nursing*, 17(1):49–59. Epub 2020 Jan 8.
- Jihyun K Hur, Joseph Heffner, Gloria W Feng, Jutta Joormann, and Robb B Rutledge. 2024. Language sentiment predicts changes in depressive symptoms. *Proceedings of the National Academy of Sciences*, 121(39):e2321321121.
- Nicholas C. Jacobson and Michelle G. Newman. 2014. [Avoidance mediates the relationship between anxiety and depression over a decade later](#). *Journal of Anxiety Disorders*, 28(5):437–445. Epub 2014 Apr 16.
- Ned H. Kalin. 2020. [The critical relationship between anxiety and depression](#). *American Journal of Psychiatry*, 177(5):365–367. PMID: 32354270.
- Guolin Ke, Qi Meng, Thomas Finley, Taifeng Wang, Wei Chen, Weidong Ma, Qiwei Ye, and Tie-Yan Liu. 2017. Lightgbm: A highly efficient gradient boosting decision tree. *Advances in neural information processing systems*, 30.
- Oscar N. E. Kjell, Pernilla Johnsson, and Sverker Sikström. 2021. [Freely generated word responses analyzed with artificial intelligence predict self-reported symptoms of depression, anxiety, and worry](#). *Frontiers in Psychology*, 12:602581.
- Kurt Kroenke, Robert L. Spitzer, and Janet B. W. Williams. 2001. [The phq-9: Validity of a brief depression severity measure](#). *Journal of General Internal Medicine*, 16(9):606–613.
- Tingting Liu, Lyle H Ungar, Brenda Curtis, Garrick Sherman, Kenna Yadeta, Louis Tay, Johannes C Eichstaedt, and Sharath Chandra Guntuku. 2022. Head versus heart: social media reveals differential language of loneliness from depression. *Npj Mental Health Research*, 1(1):16.
- Daniel M. Low, Laurie Rumker, Tanya Talkar, John Torous, Guillermo Cecchi, and Satrajit S. Ghosh. 2020. [Natural language processing reveals vulnerable mental health support groups and heightened health anxiety on reddit during covid-19: Observational study](#). *Journal of Medical Internet Research*, 22(10):e22635.
- Scott M Lundberg and Su-In Lee. 2017. A unified approach to interpreting model predictions. *Advances in neural information processing systems*, 30.
- Shamsuddeen Hassan Muhammad, Nedjma Ousidhoum, Idris Abdulmumin, Jan Philip Wahle, Terry Ruas, et al. 2025. [BRIGHTER: BRIdging the gap in human-annotated textual emotion recognition datasets for 28 languages](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8895–8916, Vienna, Austria. Association for Computational Linguistics.
- Susan Nolen-Hoeksema, Blair E Wisco, and Sonja Lyubomirsky. 2008. Rethinking rumination. *Perspectives on psychological science*, 3(5):400–424.
- Brenda WJH Penninx, Yuri Milaneschi, Femke Lamers, and Nicole Vogelzangs. 2013. Understanding the somatic consequences of depression: biological mechanisms and the role of depression symptom profile. *BMC medicine*, 11(1):129.

- Sadiya Sayara Chowdhury Puspo, Ana-Maria Bucur, Stevie Chancellor, Ozlem Uzuner, and Marcos Zampieri. 2026. Mental health disorder detection beyond social media: A systematic review of available datasets. In *Proceedings of LREC*. European Language Resources Association.
- Abu Bakar Siddiqur Rahman, Hoang-Thang Ta, Lotfollah Najjar, Azad Azadmanesh, and Ali Saffet Gönül. 2024. [Depressionemo: A novel dataset for multilabel classification of depression emotions](#). *arXiv preprint arXiv:2401.04655*. Preprint.
- Sunny Rai, Elizabeth C Stade, Salvatore Giorgi, Ashley Francisco, Lyle H Ungar, Brenda Curtis, and Sharath C Guntuku. 2024. Key language markers of depression on social media depend on race. *Proceedings of the National Academy of Sciences*, 121(14):e2319837121.
- Stephanie Rude, Eva-Maria Gortner, and James Pennebaker. 2004. Language use of depressed and depression-vulnerable college students. *Cognition & Emotion*, 18(8):1121–1133.
- Abby H Salopek and Matthew S Eastin. 2024. Toxic positivity intentions: an image management approach to upward social comparison and false self-presentation. *Journal of Computer-Mediated Communication*, 29(3):zmae003.
- Sverker Sikström, Blerta Kelmendi, and Niklas Persson. 2023. [Assessment of depression and anxiety in young and old with a question-based computational language approach](#). *npj Mental Health Research*, 2:11.
- Robert L. Spitzer, Kurt Kroenke, Janet B. W. Williams, and Bernd Löwe. 2006. [A brief measure for assessing generalized anxiety disorder: The gad-7](#). *Archives of Internal Medicine*, 166(10):1092–1097.
- Raluca Nicoleta Trifu, Bogdan Nemes, Dana Cristina Herta, Carolina Bodea-Hategan, Dorina Anca Talaş, and Horia Coman. 2024. Linguistic markers for major depressive disorder: a cross-sectional study using an automated procedure. *Frontiers in Psychology*, 15:1355734.
- Isabelle Van der Vegt and Bennett Kleinberg. 2023. [A multi-modal panel dataset to understand the psychological impact of the pandemic](#). *Scientific Data*, 10(1):537.
- Arya Varastehnezhad, Reza Tavasoli, Soroush Elyasi, MohammadHossein LotfiNia, and Hamed Farbeh. 2025. [Ai in mental health: Emotional and sentiment analysis of large language models' responses to depression, anxiety, and stress queries](#). *ArXiv*, abs/2508.11285.
- Seowon Yoon, Jihee Jang, Gaeun Son, Soohyun Park, Jueun Hwang, Joon Yeon Choeh, and Kee-Hong Choi. 2024. [Predicting neuroticism with open-ended response using natural language processing](#). *Frontiers in Psychiatry*, 15:1437569.
- Tianlin Zhang, Kailai Yang, Shaoxiong Ji, and Sophia Ananiadou. 2023. [Emotion fusion for mental illness detection from social media: A survey](#). *Information Fusion*, 92:231–246.