Exploring Optimism and Pessimism in Online Discourse of Individuals with Depression

Anonymous ACL submission

Abstract

The relationship between depression and the concepts of optimism and pessimism has been extensively researched by psychologists. In this paper, we use computational approaches to study how optimism and pessimism are expressed in the online discourse of people diagnosed with depression. Publicly available datasets are used for the development of an optimism/pessimism detection model, as well as for the analyses performed on social media posts of individuals with depression, as measured by BDI-II, a validated questionnaire for assessing depression. To analyze the optimistic and pessimistic posts by individuals with depression, we use LIWC features and perform topic modeling. Our results show that while there might not be significant differences between the amount of optimistic versus pessimistic posts depressed and control individuals have, the content of the posts differ meaningfully, both in terms of linguistic features and approached topics.

1 Introduction

Depression is one of the most prevalent mental disorders and has been extensively researched (Lim et al., 2018; Xu et al., 2021). Many studies focus on understanding how depression manifests and its relationship with mood and emotions (Rottenberg, 2005). In addition to emotions, previous research has also investigated the connection between depression and the concepts of optimism and pessimism. Karhu et al. (2024) demonstrate a bidirectional relationship: optimism not only buffers against depressive symptoms but is also eroded by them, while pessimism both predicts and is intensified by depression. Complementary studies by (Korn et al., 2014) and Hobbs et al. (2022) reveal that, unlike healthy individuals who display an optimistic bias when updating beliefs about the future, those with depression tend to weigh negative information more heavily. In addition, optimism is associated with better psychological well-being and more effective coping (Scheier et al., 2001), as well as better treatment outcomes, including reduced rehospitalization (Tindle et al., 2012). Prior research also highlights a reduced risk of work disability and an enhanced likelihood of returning to work following a depression-related disability (Kronström et al., 2011).

In recent years, computational analyses of social media data have offered significant insights into the interplay between psychological constructs and mental health. Depression detection is a prominent topic in Natural Language Processing, with traditional methods such as Support Vector Machines (SVMs), logistic regression, and random forests being used (Gan et al., 2024). More recently, there has been a transition to modern methods that use attention, deep learning, and pre-trained models (De Santana Correia and Colombini, 2022), demonstrating significant performance increases. However, in addition to identifying mental health disorders, language can offer insights into broader psychological states, such as optimism and pessimism, which are often associated with conditions like depression (Herwig et al., 2009). Previous research from NLP has explored the manifestations of emotions (Uban et al., 2021; Aragon et al., 2021) and even happy moments using social media data from individuals with depression (Bucur et al., 2024). Although research from NLP has focused on developing more effective models for detecting optimism and pessimism (Ruan et al., 2016; Caragea et al., 2018; Alshahrani et al., 2021), to our knowledge, there has been no analysis of optimism and pessimism in the social media language used by individuals with depression.

This work extends current research by examining how expressions of optimism and pessimism on social media correlate with depressive symptoms, a link strongly supported by psychological literature. The primary objectives of this study are to develop optimism-pessimism detection systems using advanced transformer-based architecture and conduct studies based on the relationship between optimism-pessimism and mental health issues, such as depression. As a result, detecting optimism and pessimism in social media is considered a first step toward more accurately understanding and detecting mental health issues. To the best of our knowledge, this study is the first to computationally analyze the correlation between optimistic and pessimistic social media language in people with depression. Thus, we aim to answer the following research questions:

- **RQ1**: In what proportions are optimism and pessimism respectively manifested in the discourse of individuals with depression?
- **RQ2**: How is optimism manifested in the social media language of individuals with depression?

2 Related Work

Though it is still in its early stages, research on detecting optimism and pessimism in social media is expanding, partly because of the COVID-19 epidemic. A deep-learning technique was presented by Blanco and Lourenço (2022) to examine the expression of optimistic and pessimistic sentiments in COVID-19-related Twitter conversations. They examined several network configurations using a pre-trained transformer embedding for semantic feature extraction and found that bi-LSTM systems produced the most successful models. According to the study, optimistic interactions tended to stay positive whereas conversations with strong pessimistic signals showed little emotional change.

In order to improve prediction accuracy for optimism and pessimism, Alshahrani et al. (2020) employed XLNet, a network that combines several auto-regressive language models, to capture semantic relationships and negations. On the benchmark dataset OPT (Ruan et al., 2016), the study's significant 63.32% error reduction increased the state-of-the-art accuracy from 90.32% to 96.45%. Accuracy at the tweet and user levels for two defined thresholds—0 and 1/-1—was one of the assessment measures.

Cobeli et al. (2022) developed a Multi-Task Knowledge Distillation architecture to achieve an accuracy of 86.60% using a best model. They used the same OPT dataset that rates respondents' optimism and pessimism on a scale from -3 (very pessimistic) to 3 (extremely optimistic). The research found that certain POS tags are consistently prevalent throughout all optimism ranges, such as nouns in 80% to 90% of tweets. Other tags, such as hashtags, have been found to be associated with optimism levels. The use of emoticons, punctuation, and user remarks also influenced optimism. As tweets became more positive, first-person singular pronouns were used less frequently, supporting the argument that pessimism and depression may be related. The researchers found that optimizing BERT on the OPT dataset improved performance compared to non-BERT baselines. BERTweet, pre-trained on tweets, performed better with a mean accuracy of 84.58% on the validation set. The attention-based models, such as BERT and BERTweet, fared better than earlier baselines, while MTKD enhanced the results. BERTweet and MTKD outperformed earlier models in demonstrating the best outcomes for the 1/-1 threshold definition of optimism.

The concept of computational analyses in the field of mental health detection correlations in social media speech has been investigated to an extent in the study by Bucur et al. (2021), which looks into the relationship between offensive language and depression by examining how people with depression use offensive speech in their social media posts. According to the authors' data, there is a greater prevalence of derogatory language in the online speech of individuals who have been diagnosed with depression.

In our research, we use computational methods to analyze the online discourse of individuals with depression. We aim to explore the impact of optimism and pessimism, motivated by the existing psychological research and advancements in NLP models designed to detect these two mental attitudes.

3 Data

We use two data collections in our experiments: the OPT dataset (Ruan et al., 2016) with annotations for optimism and pessimism and the eRisk 2021 dataset (Parapar et al., 2021) with social media individuals with depression.

The most popular dataset for optimism/pessimism identification was introduced by Ruan et al. (2016). It contains 7,475 randomly chosen tweets from 500 pessimistic individuals and 500 who were considered optimists. To select the texts, tweets containing optimism or pessimism-related keywords were found, highlighting both optimistic and pessimistic users. Each tweet was evaluated and classified by human annotators using Amazon Mechanical Turk on a scale. To guarantee accuracy, quality control procedures were put in place, such as defining optimism and pessimism precisely, excluding commentators who answered "check" questions incorrectly, and comparing annotations to the average score to spot anomalies. Human annotators rated tweets on a disposition scale from 3 (extremely optimistic) to -3 (very pessimistic); this scale made it possible to distinguish between tweets in a complex way, allowing different levels of optimism and pessimism to be identified within the text. The average of all the evaluations for the acquired annotations is the final score.

In our experiments, we consider the three possible classes: posts with an average annotation below -1 are labeled as pessimistic, those with a score of -1 to 1 belong to the neutral class, and the remaining posts are labeled as optimistic. This three-class setting provides greater granularity and intuitiveness.

Our approach is different from the direction taken in the studies mentioned in the previous section; both works identify the need to address posts with average scores between -1 and 1 separately, as they are the most ambiguous in the given context, even for human interpretation. In one of their approaches, Cobeli et al. (2022) choose to eliminate the specific group of posts, and consider the two classes, optimistic and pessimistic, so as to have a clearer distinction between the two attitudes. Alshahrani et al. (2020) employed the same method of ignoring the respective posts to address the ambiguity, calling it the -1/1 threshold. In both studies, this approach significantly improved model performance, however, for our work we chose not to use a similar technique, but rather keep the ambiguous data and create an additional class for it, for two main reasons:

1. We believe retaining this data ensures preserving the complexity and authenticity of real-world social media posts, as realistically, not all posts are and should be classified either optimistic or pessimistic

2. Eliminating the respective posts would mean reducing the data to almost half of the original size (3,847).

The eRisk 2021 dataset related to depression (Losada and Crestani, 2016; Parapar et al., 2021) contains social media users who were asked to fill in the BDI-II questionnaire (Beck et al., 1996) for the assessment of their depression status. Following this, their Reddit social media data was collected with their consent. The BDI-II questionnaire contains 21 questions related to depression symptoms, and the answers are used to calculate an overall score that indicates the level of depression. The training dataset consists of 90 users with ground truth BDI-II scores and 46,502 posts from Reddit. The test dataset contains 80 users with a total of 32,237 posts. In our experiments, we use the data from all 170 users in the eRisk dataset. Because BDI-II is used by mental health professionals to diagnose depression, we consider users with a score above the established cut-off of 19 (Subica et al., 2014; von Glischinski et al., 2019) as having depression, while those with scores below this threshold are considered control users.

4 Methodology

4.1 Detection of optimism and pessimism

Due to its good downstream performance across a great variety of tasks (Liu et al., 2019; Guo et al., 2022; Amin et al., 2023), we use in our experiments a RoBERTa-based model fine-tuned on the OPT dataset, which is then used to predict optimism, pessimism and neutral labels on the eRisk depression data.

The model, which we will can RoBERTa-OPT-3Labels from now on, was trained using the HuggingFace platform, with twitter-roberta-base-sentiment-latest serving as the base model (Camacho-Collados et al., 2022). The base model was refined for sentiment analysis using the TweetEval benchmark (Barbieri et al., 2020) after being trained on about 124 million tweets. In our training, we set a learning rate of 5e-5, three epochs, a maximum sequence length of 128 characters, an 8-batch size, and a warmup ratio of 0.1. To reduce overfitting, the optimizer employed was AdamW, a variation of the Adam optimizer with weight decay. The learning rate was decreased linearly from the starting value to zero using the "linear" learning rate scheduler. In order to avoid exploding gradient problems, the maximum gradient norm was fixed at 1. To guarantee consistency of outcomes, the seed was set to 42. If after five successive evaluations, there was no progress in the validation metric, early

Optimism		Pessimism		Neutral	
Control	**Depression**	**Control**	**Depression**	**Control**	**Depression**
I'm happy that everything turned out rather well for you in the end, and that gives me a lot of hope for my future.	I graduated [...] and got my driver's license! [...] I know what the next goal to work for is. [...] I honestly value my friendships more.	It is sad to think that the life that we will live in is set for imminent destruction.	Something must always [...] remind me how painful life is and that it will never GENUINELY get better. [...] Everyone would be better off without me [...] I will never be good enough.	Beagles are usually listed as a breed that tends to get along well with cats [...]	I only consume great, but lesser-known media. Are you familiar with Steins;Gate and Morrowind? Thought so.

Table 1: Selected examples that were predicted as optimistic, pessimistic, or neutral from the depression and control groups.

stopping was employed by setting the early stopping patience to 5. The early stopping threshold, which denotes the minimum significant change in the tracked metric needed for it to be deemed an improvement, was set at 0.01.

4.2 LIWC

LIWC 22 (Boyd et al., 2022) is an advanced text analysis tool that categorizes language into different dimensions, including psychologically meaningful ones, enabling the detection of cognitive, emotional, and social cues within the written content. In our study, we focus on the most context-significant LIWC-derived features to analyze optimistic and pessimistic posts by individuals with depressive symptoms. We quantify these differences using z-scores derived from the Mann–Whitney U test, a nonparametric statistical method that assesses whether one group systematically ranks higher or lower than another on a given variable, being particularly suited for analyzing linguistic features that may not follow a normal distribution. Specifically, we use the test to compare how the linguistic features (as categorized by LIWC) differ between the optimistic and pessimistic posts within the depression and control groups. The z-scores reflect the magnitude of these differences, allowing us to quantify how strongly specific language patterns (such as references to future focus, negative emotions, or social behavior) are associated with either optimistic or pessimistic contexts in each group.

4.3 Topic Modeling

We implemented a robust topic modeling framework using BERTopic (Grootendorst, 2022) to uncover themes within social media posts, and to explore their associations with the sentiment and mental health indicators. Our approach leveraged a customized BERTopic pipeline, which integrates text representation, dimensionality reduction, and clustering techniques.

First, we generated dense text embeddings with SentenceTransformer ('all-MiniLM-L6-v2') and reduced dimensionality using UMAP, preserving intrinsic data structure. Clustering was achieved with HDBSCAN, following text preprocessing with a CountVectorizer that included the standard English stopwords, extended with common internet noise words: 'http', 'https', 'amp', 'com', 'www', 'r/'.

To enhance interpretability, topics were refined using a custom representation that leverages KeyBERT, combined with Part-of-Speech filtering and Maximal Marginal Relevance (MMR), yielding high-quality, contextually relevant keywords. The final model assigned topics to each post, which were aggregated by sentiment (optimism, neutral, pessimism) and depression status (depressed vs. control). Chi-squared tests of independence were then employed to statistically assess differences in topic distributions across the target groups.

5 Results and Discussions

5.1 Model Performance

The RoBERTa-OPT-3Labels model shows consistent and competitive performance, with an accuracy of 71.65%, a weighted F1 score of 71.23%, and nearly matching precision and recall values on the test set. The weighted AUC of 0.8452 further underlines its ability to effectively distinguish among the three classes. As this is, to the best of our knowledge, the first work to consider a 3-class approach, it would be interesting to see the results of the state-of-the-art models that interpreted the 1/-1 scenario by eliminating the neutral/ambiguous posts (Caragea et al. (2018); Alshahrani et al. (2020); Alshahrani et al. (2021); Cobeli et al. (2022)). We present selected predicted samples in Table 1.

5.2 General Statistics Interpretation

After running the predictions for optimism and pessimism using the RoBERTa-OPT-3Labels model, we find that users in the depression group have, on average, fewer optimistic posts than the control group, but a similar number of pessimistic posts. In addition, users in the control group have more posts labeled as neutral. The exact descriptive statistics can be found in Appendix A Tables 4 and 5.

To test for statistical significance, we compare the number of optimistic, pessimistic and neutral posts between the two groups, using Mann–Whitney U test, Cohen's d and Pearson correlation (Table 2). The Mann–Whitney U test yields non-significant z-scores and p-values for both optimistic (-1.23, p = 0.22) and pessimistic (-0.20, p = 0.84) posts, suggesting that both groups produce similar amounts of content in these categories. In addition, the small effect sizes (Cohen's d = -0.18 for optimism, 0.06 for pessimism) and weak Pearson correlations further support this lack of meaningful distinction.

However, a more significant difference can be seen in the number of neutral posts for the performed tests, with a small to moderate effect size (d = -0.35). This suggests that individuals with depression post significantly fewer neutral statements than people not diagnosed with depression, potentially reflecting a tendency to engage more with emotionally valenced (optimistic or pessimistic) language rather than neutral discourse (Broome et al., 2015).

	Mann–Whitney U test (z, p)	Cohen's d	Pearson Correlation (r, p)
Optimistic	(-1.23, 0.22)	-0.18	(-0.09, 0.26)
Pessimistic	(-0.20, 0.84)	0.06	(0.03, 0.68)
Neutral	(-2.21, 0.03)	-0.35	(-0.17, 0.03)

Table 2: Statistical Test Results for Optimism and Pessimism

While the statistical tests indicate no significant differences in the number of optimistic or pessimistic posts between depression and control individuals, our subsequent analyses will demonstrate that the content of these posts may vary substantially. We will proceed to show that the way optimism and pessimism are expressed in language differs between depressed and non-depressed users in a meaningful way.

5.3 LIWC Analysis Results

Figure 1 presents a side-by-side comparison of LIWC feature usage across optimistic (left panel) and pessimistic (right panel) posts by individuals with and without depression, measured via z-scores. The categories marked by (*) are statistically significant ($p<0.05$), as measured by the Mann–Whitney U test. By analyzing these scores, we have outlined several key patterns.

First, in the optimistic posts, depressed users still exhibit relatively higher frequencies of negative emotion and mental health references compared to the control group, suggesting that even ostensibly hopeful content may be interlaced with underlying emotional distress (Yang et al., 2023). They also display more tentative and cognitively complex language, indicating ongoing uncertainty and self-reflection, which may highlight core aspects of depressive cognition. On the other hand, depressed users' optimistic posts contain frequent future-oriented words and a significant positive tone, suggesting a forward-looking, positive outlook (Ji et al., 2016). This aligns with prior findings that individuals with depression, despite their condition, often maintain beliefs that their lives will improve in the future. However, research also shows that such expectations do not necessarily serve as a protective factor, instead being linked to an increased risk of recurrent depressive symptoms over time (Busseri and Peck, 2015). Interestingly, even in positive contexts, individuals with depression seem to consistently engage less in cultural and lifestyle/leisure topics. This reduced engagement indicates a persistent disengagement from activities that typically enhance well-being and contribute to a richer quality of life (Eisemann, 1984). This also extends to the usage of language related to achievement and reward, where we see a reduction compared to their control counterparts. This may reflect a diminished sense of agency or efficacy, which is commonly observed in depression (Halahakoon et al., 2020; Winer and Salem, 2015).

When looking at pessimistic posts, the gap between depressed and control users is also pronounced. Depressed individuals exhibit a marked increase in words related to negative affect (e.g., general negative emotion, sadness, anxiety) and self-focused attention (significant usage of first-person pronouns) and health concerns (talking more about general and mental health, illness, death), coupled with all-or-none thinking and dis-

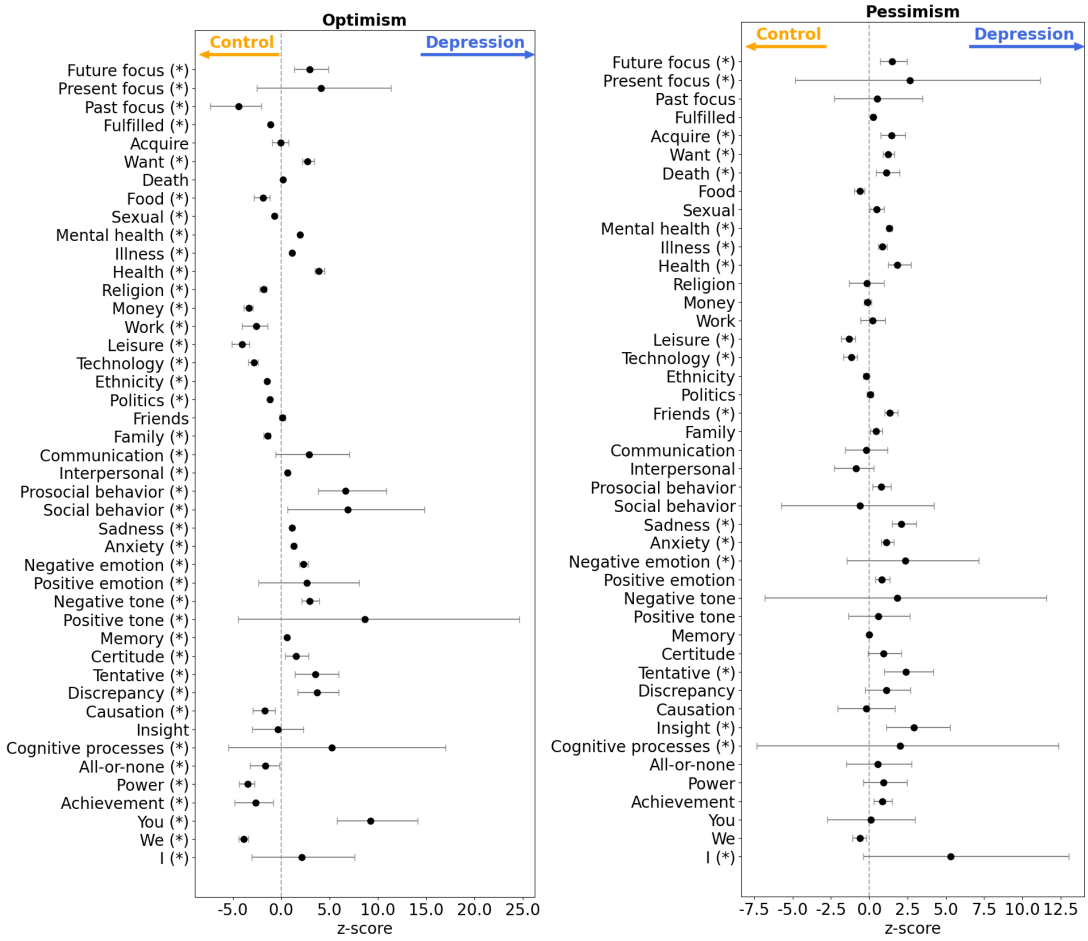

Figure 1: Z-scores for the differences between the depression and control groups for posts labeled as optimistic and pessimistic by the RoBERTa model. Results with (*) are statistically significant (p<0.05, Mann-Whitney U test).

crepancy terms, suggesting a tendency toward more rigid, negative and self-critical thought processes (Mor and Winquist, 2002). Control users, while also expressing negative content in pessimistic posts, tend to do so with fewer markers of pervasive distress and exhibit less dichotomous thinking.

Focusing exclusively on depressed individuals, the figure reveals clear linguistic distinctions between their optimistic and pessimistic posts. In optimistic posts, the language still retains subtle markers of distress—such as moderate levels of negative affect and tentative wording—indicating an underlying cognitive dissonance. Conversely, pessimistic posts are characterized by a significant amplification of negative emotion terms, self-referential language, and rigid, absolutist thinking patterns (Al-Mosaiwi and Johnstone, 2018). However, it is also worth noting that in their optimistic posts, depressed users exhibit a higher frequency of social and prosocial language, which includes terms relating to altruistic behaviors and social engagement. This pattern indicates that when depressed individuals adopt an optimistic tone, they are more likely to express social behaviors and have a desire for connection and support (Carver et al., 1994). In contrast, pessimistic posts are marked by a relative reduction in social language, suggesting that negative emotional states may suppress expressions of social engagement.

The statistical differences measured with the Mann–Whitney U test, and highlighted by z-scores reveal how depressed individuals use their language differently, firstly in comparison to the control group, but also based on the sentiment of the content, with pessimistic posts exhibiting a more pronounced negative linguistic profile.

5.4 Topic Modeling Results

The chi-squared results across the target (depression versus control) groups reveal significant thematic differences in how individuals communicate optimism, pessimism, and neutrality. We will be addressing results for six distinct sub-

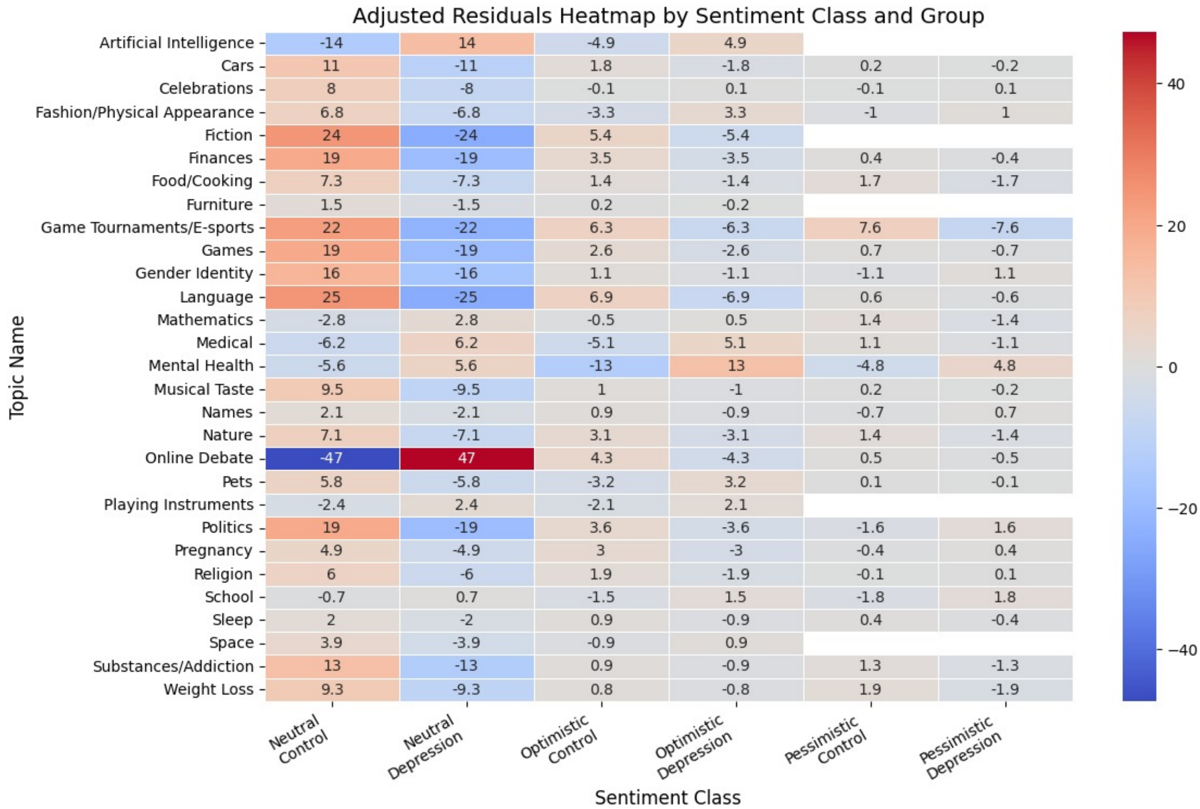

Figure 2: Heatmap of standardized residuals. The colors indicate which topics are significantly overrepresented (red) or underrepresented (blue) in each group.

Group	Category	Overrepresented Topics	Underrepresented Topics
Depressed	**Neutral**	Medical, AI, Online Debate	Fiction, Language, E-sports
	Optimism	Mental Health, Medical, AI	Language, E-sports, Fiction
	Pessimism	Mental Health, School, Politics	Online Debate, AI, Pets
Control	**Neutral**	Fiction, Language, E-sports	Online Debate, AI, Medical
	Optimism	Language, E-sports, Fiction	Mental Health, AI, Medical
	Pessimism	E-sports, Weight Loss, Food	Mental Health, School, Politics

Table 3: Top three topic overrepresentation and underrepresentation across depression and control groups

groups, based on the depression label and the optimism/pessimism/neutral associations, with visualizations available in Figure 2 as a heatmap. We present in Table 3 the most overrepresented and underrepresented topics for each target group. In Appendix A Tables 6 and 7, we present the top 10 topics for the depression and the control group, respectively. Also in Appendix A, Figure 3 displays the standardized residuals, calculated from the observed and expected topic frequencies across the three sentiment classes (neutral, optimism, pessimism).

The disparities suggest that psychological states influence topic preferences in online discourse. The pronounced engagement of Depression-Neutral posts in online debate and artificial intelligence contrasts with the avoidance of these topics in Control-Neutral posts, highlighting a potential association between depression and increased argumentative or analytical engagement when not expressing optimism or pessimism. On the other hand, individuals with depression seem overall more comfortable engaging in mental health-related discourse in all sentiment settings, even in optimistic posts. The significant engagement with e-sports in Control-Pessimism posts may indicate a preference for structured, competitive digital interactions in this category, perhaps as a coping mechanism or an outlet for engagement that does not necessitate personal disclosure. The control group also seems to be engaged in talks about fictional works and general leisure/lifestyle topics, which don't seem as prevalent in the depression group, a theory also supported by literature that suggests reduced engagement in such activities by people diagnosed with depression (Eisemann, 1984). This result is consistent with the observations from the LIWC feature analysis.

To be noted that the missing values seen in the Pessimistic category (both for depression and control groups) were intentionally excluded as there were no posts of the respective topics belonging to that specific subgroup, thus not being statistically significant.

5.5 Revisiting Research Questions

Addressing **RQ1**, our analyses reveal that the overall proportions of optimistic and pessimistic posts among individuals with depression are statistically similar to those of the control group. This indicates that, in terms of frequency, individuals in the depression group do not necessarily exhibit a reduced tendency to express optimism compared to control users, though the control group moderately engages more in neutral content. However, while the quantity of such expressions appears consistent, the qualitative content differs markedly.

In response to **RQ2**, our findings indicate that optimism in the social media language of individuals with depression is manifested in a more nuanced and complex manner. Although optimistic posts are present at comparable rates, the linguistic features and thematic content of these posts suggest a distinct expression of optimism that is intertwined with elements of resilience and coping. Specifically, while their optimistic posts are marked by a significant positive tone and a frequent use of future-oriented terms, a higher frequency of negative emotion words and compared to the control group is still notable. Notably, even within contexts that are ostensibly positive, individuals with depression demonstrate less engagement with cultural, lifestyle, and leisure topics, maintaining a great focus on mental health discussions.

6 Conclusions and Future Work

Our study investigated the expressions of optimism and pessimism in the social media discourse of individuals with depression using computational methods. Although no significant differences were observed in the actual amounts of optimistic versus pessimistic posts between the depression and control groups, our analyses revealed meaningful differences in the linguistic content and thematic focus of these posts. Notably, while pessimistic posts from individuals with depression exhibited a pronounced negative linguistic profile, the expressions of optimism—though subtler—appear to represent a complex interplay of resilience and coping mechanisms. These findings might insights into adaptive strategies within this target group. Overall, our results not only corroborate existing psychological theories regarding language, psychological and depressive states but also highlight the potential of transformer-based models, topic modeling and LIWC features in capturing nuanced variations in online discourse related to mental health.

Subsequent investigations may benefit from a longitudinal approach to examine how expressions of optimism and pessimism evolve over time in relation to depressive symptoms. Additionally, integrating multimodal data—such as images, user interactions, and metadata—may provide a more comprehensive understanding of online expressions of optimism and pessimism.

Limitations

In our experiments, we used the OPT dataset obtained from Twitter/X to train a transformer-based model for predicting optimism and pessimism labels in depression-related content sourced from Reddit. This choice was made due to the limited availability of datasets from the same domain. The OPT dataset is the most commonly used dataset for this specific task (Caragea et al., 2018; Cobeli et al., 2022). Additionally, we selected the eRisk 2021 dataset because it includes social media users who have completed the validated BDI-II questionnaire, which provides a reliable assessment of depression. Prior research suggests that transformer-based models are effective for transfer learning across different platforms (Uban et al., 2022), although future work could explore domain-specific adaptations. To address this limitation, we conduct statistical tests on our results, to strengthen our findings.

Ethical Considerations

This paper uses OPT, a publicly available dataset with annotations for optimism and pessimism. In addition, the eRisk 2021 dataset was made available to us after signing a data usage agreement form. We have adhered to the data agreement, and we did not make any attempt to contact the users or to de-anonymize the data. The sample of posts presented in this paper has been paraphrased to ensure the anonymity of the users. Our primary focus is on quantifying and analyzing optimistic and pessimistic sentiments within the texts of the mental health dataset. We do not aim to predict mental health status or conditions based on this dataset.

References

M. Al-Mosaiwi and T. 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.

A. Alshahrani, M. Ghaffari, K. Amirizirtol, and X. Liu. 2020. Identifying optimism and pessimism in twitter messages using xlnet and deep consensus. In *International Joint Conference on Neural Networks (IJCNN)*.

A. Alshahrani, M. Ghaffari, K. Amirizirtol, and X. Liu. 2021. Optimism/pessimism prediction of twitter messages and users using bert with soft label assignment. In *International Joint Conference on Neural Networks (IJCNN)*.

Mostafa M. Amin, E. Cambria, Björn Schuller, and E. Cambria. 2023. Will affective computing emerge from foundation models and general artificial intelligence? a first evaluation of chatgpt. *IEEE Intelligent Systems*, 38:15–23.

Mario Ezra Aragon, Adrian Pastor Lopez-Monroy, Luis Carlos González-Gurrola, and Manuel Montes-y Gómez. 2021. Detecting mental disorders in social media through emotional patterns-the case of anorexia and depression. *IEEE transactions on affective computing*, 14(1):211–222.

Fabio Barbieri, José Camacho-Collados, Lisa Ellen Anke, and Leandro Neves. 2020. Tweeteval: Unified benchmark and comparative evaluation for tweet classification. In *Findings of the Association for Computational Linguistics: EMNLP 2020*. Association for Computational Linguistics.

Aaron T Beck, Robert A Steer, Gregory K Brown, et al. 1996. Beck depression inventory.

G. Blanco and A. Lourenço. 2022. Optimism and pessimism analysis using deep learning on covid-19 related twitter conversations. *Information Processing and Management*, 59(3):102918.

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.

M. R. Broome, K. E. A. Saunders, P. J. Harrison, and S. Marwaha. 2015. Mood instability: Significance, definition and measurement. *British Journal of Psychiatry*, 207(4):283–285.

A. Bucur, M. Zampieri, and L. P. Dinu. 2021. An exploratory analysis of the relation between offensive language and mental health. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 3600–3606.

Ana-Maria Bucur, Berta Chulvi, Adrian Cosma, and Paolo Rosso. 2024. The expression of happiness in social media of individuals reporting depression. *IEEE Transactions on Affective Computing*.

M. A. Busseri and E. Peck. 2015. Do (even) depressed individuals believe that life gets better and better? the link between depression and subjective trajectories for life satisfaction. *Clinical Psychological Science*, 3(5):715–725.

J. Camacho-Collados, K. Rezaee, T. Riahi, A. Ushio, D. Loureiro, D. Antypas, J. Boisson, L. E. Anke, F. Liu, and E. M. Camara. 2022. Tweetnlp: Cutting-edge natural language processing for social media. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*.

C. Caragea, L. P. Dinu, and B. Dumitru. 2018. Exploring optimism and pessimism in twitter using deep learning. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*.

Charles S. Carver, L. A. Kus, and M. F. Scheier. 1994. Effects of good versus bad mood and optimistic versus pessimistic outlook on social acceptance versus rejection. *Journal of Social and Clinical Psychology*, 13(2):138–151.

Ş. Cobeli, I. Iordache, S. Yadav, C. Caragea, L.P. Dinu, and D. Iliescu. 2022. Detecting optimism in tweets using knowledge distillation and linguistic analysis of optimism. In *International Conference on Language Resources and Evaluation*.

A. De Santana Correia and E. L. Colombini. 2022. Attention, please! a survey of neural attention models in deep learning. *Artificial Intelligence Review*, 55(8):6037–6124.

M. Eisemann. 1984. Leisure activities of depressive patients. *Acta Psychiatrica Scandinavica*, 69(1):45–51.

L. Gan, Y. Guo, and T. Yang. 2024. Machine learning for depression detection on web and social media. *International Journal on Semantic Web and Information Systems*, 20(1):1–28.

Maarten Grootendorst. 2022. Bertopic: Neural topic modeling with a class-based tf-idf procedure.

Yuting Guo, Yao Ge, Yuan-Chi Yang, Mohammed Ali Al-Garadi, and Abeed Sarker. 2022. Comparison of pretraining models and strategies for health-related social media text classification. In *Healthcare*, volume 10, page 1478. MDPI.

D. C. Halahakoon, K. Kieslich, C. O'Driscoll, A. Nair, G. Lewis, and J. P. Roiser. 2020. Reward-processing behavior in depressed participants relative to healthy volunteers. *JAMA Psychiatry*, 77(12):1286.

U. Herwig, A. B. Brühl, T. Kaffenberger, T. Baumgärtner, H. Boeker, and L. Jäncke. 2009. Neural correlates of 'pessimistic' attitude in depression. *Psychological Medicine*, 40(5):789–800.

Catherine Hobbs, Petra Vozarova, Aarushi Sabharwal, Punit Shah, and Katherine Button. 2022. Is depression associated with reduced optimistic belief updating? *Royal Society Open Science*, 9(2):190814.

J. L. Ji, E. A. Holmes, and S. E. Blackwell. 2016. Seeing light at the end of the tunnel: Positive prospective mental imagery and optimism in depression. *Psychiatry Research*, 247:155–162.

Jutta Karhu, Juha Veijola, and Mirka Hintsanen. 2024. The bidirectional relationships of optimism and pessimism with depressive symptoms in adulthood–a 15-year follow-up study from northern finland birth cohorts. *Journal of Affective Disorders*, 362:468–476.

Christoph W Korn, Tali Sharot, Hendrik Walter, Hauke R Heekeren, and Raymond J Dolan. 2014. Depression is related to an absence of optimistically biased belief updating about future life events. *Psychological medicine*, 44(3):579–592.

Kim Kronström, Hasse Karlsson, Hermann Nabi, Tuula Oksanen, Paula Salo, Noora Sjösten, Marianna Virtanen, Jaana Pentti, Mika Kivimäki, and Jussi Vahtera. 2011. Optimism and pessimism as predictors of work disability with a diagnosis of depression: a prospective cohort study of onset and recovery. *Journal of affective disorders*, 130(1-2):294–299.

Grace Y Lim, Wilson W Tam, Yanxia Lu, Cyrus S Ho, Melvyn W Zhang, and Roger C Ho. 2018. Prevalence of depression in the community from 30 countries between 1994 and 2014. *Scientific reports*, 8(1):2861.

Y. Liu, M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer, and V. Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *OpenReview*.

D. E. Losada and F. Crestani. 2016. A test collection for research on depression and language use. In *Lecture notes in computer science*, pages 28–39.

N. Mor and J. Winquist. 2002. Self-focused attention and negative affect: A meta-analysis. *Psychological Bulletin*, 128(4):638–662.

Javier Parapar, Patricia Martín-Rodilla, David E Losada, and Fabio Crestani. 2021. Overview of erisk at clef 2021: Early risk prediction on the internet (extended overview). *CLEF (Working Notes)*, 1:864–887.

Jonathan Rottenberg. 2005. Mood and emotion in major depression. *Current Directions in Psychological Science*, 14(3):167–170.

X. Ruan, S. R. Wilson, and R. Mihalcea. 2016. Finding optimists and pessimists on twitter. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*.

Michael F Scheier, Charles S Carver, and Michael W Bridges. 2001. Optimism, pessimism, and psychological well-being. *Optimism & pessimism: Implications for theory, research, and practice.*, pages 189–216.

Andrew M Subica, J Christopher Fowler, Jon D Elhai, B Christopher Frueh, Carla Sharp, Erin L Kelly, and Jon G Allen. 2014. Factor structure and diagnostic validity of the beck depression inventory–ii with adult clinical inpatients: Comparison to a gold-standard diagnostic interview. *Psychological assessment*, 26(4):1106.

Hilary Tindle, Bea Herbeck Belnap, Patricia R Houck, Sati Mazumdar, Michael F Scheier, Karen A Matthews, Fanyin He, and Bruce L Rollman. 2012. Optimism, response to treatment of depression, and rehospitalization after coronary artery bypass graft surgery. *Psychosomatic medicine*, 74(2):200–207.

Ana-Sabina Uban, Berta Chulvi, and Paolo Rosso. 2021. An emotion and cognitive based analysis of mental health disorders from social media data. *Future Generation Computer Systems*, 124:480–494.

Ana Sabina Uban, Berta Chulvi, and Paolo Rosso. 2022. Multi-aspect transfer learning for detecting low resource mental disorders on social media. In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 3202–3219.

Michael von Glischinski, Ruth von Brachel, and Gerrit Hirschfeld. 2019. How depressed is "depressed"? a systematic review and diagnostic meta-analysis of optimal cut points for the beck depression inventory revised (bdi-ii). *Quality of Life Research*, 28:1111–1118.

E. S. Winer and T. Salem. 2015. Reward devaluation: Dot-probe meta-analytic evidence of avoidance of positive information in depressed persons. *Psychological Bulletin*, 142(1):18–78.

Dong Xu, Yi-Lun Wang, Kun-Tang Wang, Yue Wang, Xin-Ran Dong, Jie Tang, and Yuan-Lu Cui. 2021. A scientometrics analysis and visualization of depressive disorder. *Current neuropharmacology*, 19(6):766–786.

C. Yang, X. Zhang, Y. Chen, Y. Li, S. Yu, B. Zhao, T. Wang, L. Luo, and S. Gao. 2023. Emotion-dependent language featuring depression. *Journal of Behavior Therapy and Experimental Psychiatry*, 81:101883.

A Appendix

	Neutral	Optimistic	Pessimistic	Total
mean	303.59	65.52	12.76	381.88
std	325.32	76.75	16.81	391.66
min	9.00	2.00	0.00	16.00
25%	47.00	14.00	1.00	66.00
50%	150.00	33.00	5.00	199.00
75%	590.00	88.00	18.00	702.00
max	1132.00	416.00	81.00	1208.00

Table 4: Descriptive Statistics for Depression Group

	Neutral	Optimistic	Pessimistic	Total
mean	422.40	79.33	11.68	513.42
std	362.08	75.55	15.79	428.46
min	21.00	2.00	0.00	26.00
25%	57.00	14.75	2.00	66.00
50%	317.50	64.50	6.00	396.00
75%	784.25	117.50	14.00	969.50
max	1258.00	334.00	92.00	1478.00

Table 5: Descriptive Statistics for Control Group

Rank	Topic	Count
1	Online Debate	8,923
2	Mental Health	1,992
3	Politics	1,127
4	Language	822
5	Game Tournaments/E-sports	733
6	Games	590
7	Gender Identity	574
8	Musical Taste	567
9	Pets	525
10	Substances/Addiction	393

Table 6: Control Group – Top 10 Topics by Count

Rank	Topic	Count
1	Online Debate	27,241
2	Mental Health	5,430
3	Politics	1,174
4	Pets	903
5	Musical Taste	722
6	Fashion/Physical Appearance	506
7	Gender Identity	506
8	Artificial Intelligence	499
9	Language	450
10	Game Tournaments/E-sports	402

Table 7: Depressed Group – Top 10 Topics By Count

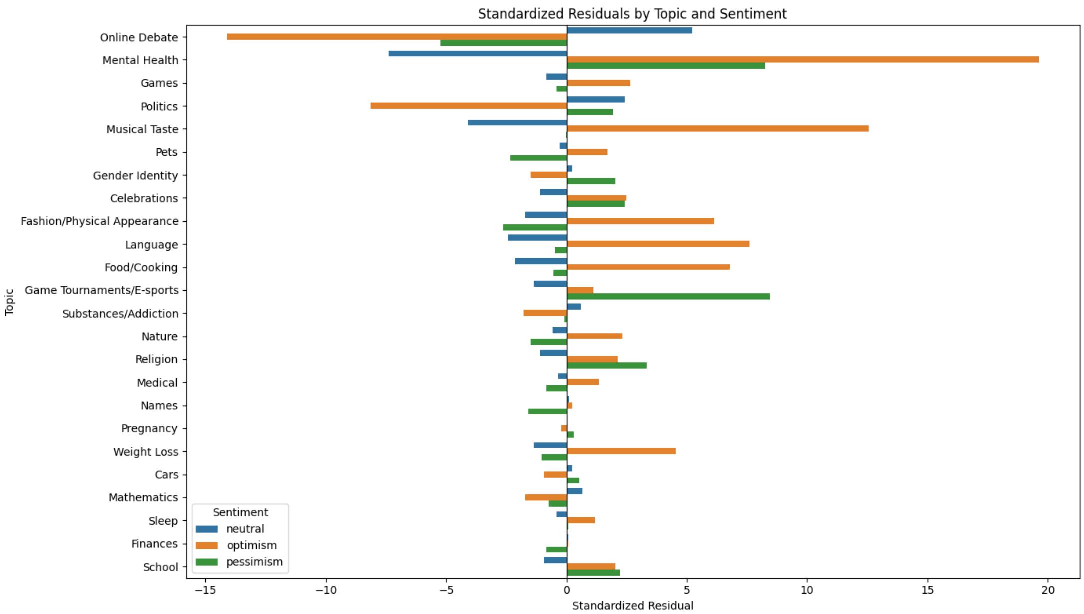

Figure 3: Standardized residuals - observed and expected topic frequencies across the three sentiment classes (neutral, optimism, pessimism).