Quantifying the Digital Phenotype of Loneliness on Twitter

Ege Otenen (eotenen@iu.edu)[1,2]
Ozgur Can Seckin (oseckin@iu.edu)[1]
Danny Valdez (danvald@iu.edu)[3]
Johan Bollen (jbollen@indiana.edu)[1,2]
[1]Luddy School of Informatics, Computing, and Science, Indiana University, Bloomington, IN 47408, USA
[2]Cognitive Science Department, Indiana University, Bloomington, IN 47408, USA
[3]School of Public Health, Indiana University, Bloomington, IN 47408, USA

Abstract

Social media promotes social connectedness, but social media users can still be lonely which is an important preceding condition to various mental health disorders such as anxiety and depression. Here we aim to describe online loneliness in individuals from the linguistic and social features of their platform use. We define a sample of Twitter users who explicitly report being lonely and compare their language to a matching random control sample. For each user, we create a text embedding - a numerical representation of the content of their online posts, excluding terms and expressions related to loneliness. We utilize principal component analysis on the resulting embeddings to condense the data into a smaller number of variables, while still retaining the majority of the variance. By doing so, we are able to position each user within a two-dimensional space, defined by the first two principal components, which capture the most significant amount of variation in the data. Lonely individuals are spatially separated from the control sample, indicating that lonely individuals exhibit distinct language patterns that is often self-referential, e.g. "I should" and "but I". Indicators of online social relations, such as the number of online friends, favorites, mentions, show that lonely individuals have fewer social relations, while a sentiment analysis demonstrates that their posts have lower valence. Our results provide insights into the lexical, social, and affective markers that characterize loneliness online, providing a starting point for the development of diagnostics and prevention.

Keywords: social media; mental health; online behaviors: natural language processing; correlation

Introduction

Loneliness, defined as periodic feelings of isolation, social disconnection, or perceived dissatisfaction with one's degree of social connection (S. Cacioppo et al., 2014; Eichstaedt et al., 2015), will affect at least one-third of the US population at some point in their life course according to a recent report of Harvard University (2021). Prior research strongly supports that experiencing loneliness, even if periodic or brief, can result in negative mental and physical health outcomes. For example, people who are lonely may experience cognitive impairments (Hawkley & Capitanio, 2015) and they are at a higher risk for developing internalizing disorders such as depression (J. T. Cacioppo et al., 2006) and anxiety (Beutel et al., 2017). Negative mental health outcomes attributed to loneliness, likewise may result in changes to physical health including poor sleep (J. T. Cacioppo et al., 2002), impaired immunity (Hawkley & Capitanio, 2015), increased risk of cardiovascular disease (Golaszewski et al., 2022) and death (Holt-Lunstad et al., 2015).

In spite of prevalence and implications of loneliness for individual and collective health, loneliness remains a rather opaque and possibly under-reported phenomenon. Indeed, lonely people who are conscious about their anxiety in social context (Narchal & McDavitt, 2017), may lack social confidence skills to seek friendships or social connection (Solano & Koester, 1989), and struggle to form and maintain relationships (Wittenberg & Reis, 1986). These self-conscious domains of loneliness may be leading to lower levels of self-disclosure in lonely individuals compared with others who are not feeling lonely.

Social media provides a platform for billions of people to openly discuss their social, affective, cognitive, and behavioral states (Chen et al., 2022; Guntuku, Schneider, et al., 2019). Lonely people in particular may prefer online communication to in-person interactions (Caplan, 2003; Kim et al., 2009; Morahan-Martin & Schumacher, 2000) due to the sense of anonymity provided in online settings (Morahan-Martin & Schumacher, 2000). Prior studies suggest that people who are lonely report that they "feel more like themselves" in online communication compared to face-to-face interactions (Lee et al., 2013) and consequently prefer to communicate and disclose themselves online (Leung, 2011). Additionally, loneliness is not a disorder that can be detected through a diagnosis. As such, constructs to measure loneliness are limited and inconsistently applied.

To date, 72% of the US adult population uses at least one social media platform daily; and almost 40% of them use it daily or more frequently (Pew Research Center, 2021). Since majorities of the US population actively use social media for a variety of purposes, including self-disclosures about health status, social media is increasingly becoming a desirable platform to study mental health outcomes at population level scales. Previous studies resoundingly support that loneliness can be studied via social media. For example, Guntuku, Schneider, et al. (2019) examined tweets posted by a sample of Pennsylvania residents (N=25,966) whose tweets included the terms "alone" or "lonely." A lexical analysis using LIWC categorization (Pennebaker et al., 2015) indicated that lonely individuals tweeted more about their emotions and cognitive processes in general. They tended to tweet at night, and mention drugs, sleep difficulties, psychosomatic symptoms, interpersonal relationships, self-reflection, and confusing emotions at greater levels than the control group. Sim-

In M. Goldwater, F. K. Anggoro, B. K. Hayes, & D. C. Ong (Eds.), *Proceedings of the 45th Annual Conference of the Cognitive Science Society*.

ilarly, Kivran-Swaine et al. (2014) developed a qualitative coding scheme to understand how loneliness-related tweets are categorized and examined the changes in types of loneliness in different demographics. They categorized tweets into four categories which are social, physical, romantic, and somatic. The social loneliness category involves texts that mention loneliness in social gatherings, physical loneliness includes an indication of physical space, somatic loneliness indicates the bodily reaction or situation due to loneliness, and romantic loneliness is related to the aspiration of a romantic partner.

Mahoney et al. (2019) applied a more general typology and criteria. In their study, they examined 22k tweets that contained the word "lonely" and generated twenty-three subcategories embedded in the three main categories of tweets which are "loneliness and self", "loneliness and others", and "irrelevant". They also examined the temporal distribution of tweets and LIWC categorization of the words. They found that tweets including "lonely" have more negative words, pronouns, verbs, health-related words with the more use of past, future, and present tenses, and that word categories can vary depending on whether the tweet was self-directed or not. Negative emotions and sadness increased on weekends and were found to be at the highest levels from 12 to 4 a.m. Therefore, considering the fact that measuring loneliness is very challenging task, and lexical analyses of social media previously provided great insights about loneliness (Guntuku et al., 2017; Kivran-Swaine et al., 2014), using social media may be a better instrument to enable the detection of loneliness.

Present Study

The aforementioned studies yield a description of the lexical and thematic properties of tweets that contain terms related to loneliness but do not specifically focus on loneliness as a personal state or experience. Here, we investigate differences in the lexical properties of tweets posted by individuals who are lonely vs. a random control group, such that we can characterize the state of loneliness from the language people use online. Specifically, we explore (1) whether loneliness can be detected from the language of individuals who self-disclose feeling lonely in their tweets using computational methods, (2) which lexical factors are most informative with respect to such detection, and (3) whether our findings correlate with social and affective indicators evident in other internalizing disorders.

Method

Data Collection

To identify a sample of lonely individuals, we collected 100,000 tweets posted between July 16 to July 31 of 2022 from the Twitter Academic API which returns all tweets that match a given search criterion. In this case we searched for individuals making an explicit self-referential statement of being lonely by searching for tweets containing lexical variations of the terms ("I am", "feel", and "lonely") in conjunction, e.g. "im feeling lonelyyy". We excluded tweets containing the same set of words, but not as a part of an explicit self-referential statement of loneliness, e.g. "I am convinced he feels lonely". We further excluded retweets and quotes from advertisements and songs. We also carefully eliminated negations from the dataset (e.g., "I am not lonely"), but allowed for adverbs and adjectives in the statement such as "I feel so lonely" or "I am a very lonely person". Next, we manually validated 2,500 tweets from the remaining dataset to verify they indeed contained an explicit self-referential statement of an individual experiencing loneliness.

Second, we created a random control group by retrieving the same number of tweets from the same observation period (July 16 to July 31, 2022). We sampled uniformly from this date range to obtain a random sample comparable in size to the set of lonely individuals that also tracks trends in Twitter volume over time. We achieved this sample by randomly selecting two adjacent seconds from each hour and each day for the entire observation period and acquiring all tweets posted in that 2s time interval. We excluded tweets which matched the term "lonely" and its variations (e.g., "lonelyy", "lonel") to exclude individuals referring to loneliness or experiencing loneliness. After applying these exclusions, we retained a sample of 9,424 randomly selected users for the control group.

Third, to prevent the inclusion of chatbots, and institutional or fake accounts, we rate all users in both groups with respect to the likelihood they are "bot-like" with Botometer V4 (Yang et al., 2020). We exclude all accounts that score higher than 2.5 out of 5 in any of the Botometer categories.

Finally, we matched the demographics of the random control sample to the sample of lonely individuals by pairwise matching to avoid the confounding effects of these variables. So that, we matched each individual from the lonely group to a person from control group who has the same age, gender, and Twitter user account creation date. First, we inferred the age and gender of our wider control sample with the M3 Inference library, a deep learning model trained on Twitter data to predict the demographic information of users with given profile information (e.g., name, bio, picture) (Wang et al., 2019).

Following the selection of individuals into either the "Lonely" (L) or "Control" (C) sample, we extracted the Twitter timelines from June 1, 2022 to October 3, 2022 for all individuals in either group, i.e. up to the 3,200 most recent tweets posted by each individual. We only selected individuals whose tweets were at least 90% in English. One thing to note is that we removed the tweets posted in the observation period (July 16 to July 31) from all subsequent analyses since they were used as the sample inclusion criterion ("tweets expressing loneliness") .

This procedure resulted in two equally sized samples of (1) 722 Lonely individuals (L) and (2) a random control sample of 722 individuals. The sample was predominantly female, with 55.82% inferred as female and 44.18% inferred as male.

The majority of individuals (52.77%) were younger than 20 years. For further information, please refer to Table 1.

Creating User Embeddings

To compare the semantic and lexical features of individuals' language between the Lonely and Control groups, we converted each user's tweet timelines into a single embedding vector. We relied on a procedure common to Natural Language Processing (NLP) models which generates vector representations for chunks of texts (Cer et al., 2018; Conneau et al., 2017; Le & Mikolov, 2014) based on patterns learned from very large text corpora.

We employed the Sentence Transformers library provided in Python and used the pre-trained model "distilbert-base-uncased" to create a text embedding for each tweet. Next, we computed the mean text embedding for each user by averaging the embedding tweets for all tweets posted by an individual, resulting in a mean "user embedding" vector.

Here, we obtained such embedding vectors with DistilBert (Sanh et al., 2019), a faster and smaller version of Bidirectional Encoder Representations from Transformers (BERT) model (Devlin et al., 2018). BERT models are masked language models that preserve semantic and syntactic context in text data, provide numerical embeddings that affords a quantitative pairwise comparison of the semantic and syntactical features of texts.

These embedding vectors are 768-dimensional, hence all user embeddings are 768-length. To visualize the similarities or differences between individual user embedding vectors, and to remove factors that are irrelevant to the distinction between Lonely and Control users, we applied a Principal Component Analysis (PCA) (Abdi & Williams, 2010) to the covariance between the user embeddings. We retained PCA Components 1 and 2 which explain respectively 37% and 15% of the variance (see Figure 1) to position each user on a two-dimensional plane spanned by the first and second PCA component. We also calculated Shannon Entropy in the PCA representations to understand the differences between the language of these two groups. The resulting two-dimensional mapping allows a visual inspection of the differences between individuals in the Lonely and Control group according to the content of their Twitter timelines.

To determine the degree to which Lonely and Control group individuals can be separated on the basis of their user embeddings, in addition to the mentioned visualization, we trained a Support Vector Machine (SVM) (Boser et al., 1992) by using the dimensionality reduction applied user embeddings as independent variable to predict whether they belong to the Lonely or Control group. We followed a five-fold cross validation process to find the best hyperparameter setting, which led us to use a linear kernel.

Language Features In addition to the PCA analysis, we examined the usage of words from different grammatical categories such as pronouns, nouns, adjectives, adverbs. Previous studies Guntuku, Schneider, et al. (2019), Kivran-Swaine et al. (2014), and Mahoney et al. (2019) consistently found the increased usage of pronouns in lonely populations. Here, we explored nominative person pronouns within the scope of this study. To achieve this, we implemented the Stanza library (Qi et al., 2020) provided in Python to find the part of speech tags of each word in tweets. A part of speech tag is a label assigned to a word in a text identifying its grammatical role, such as noun, verb, adjective or adverb. Next, for each user, we counted the number of pronouns (I, you, he, she, it, we, they) they used in their tweets, and normalized these numbers by dividing by the total number of words.

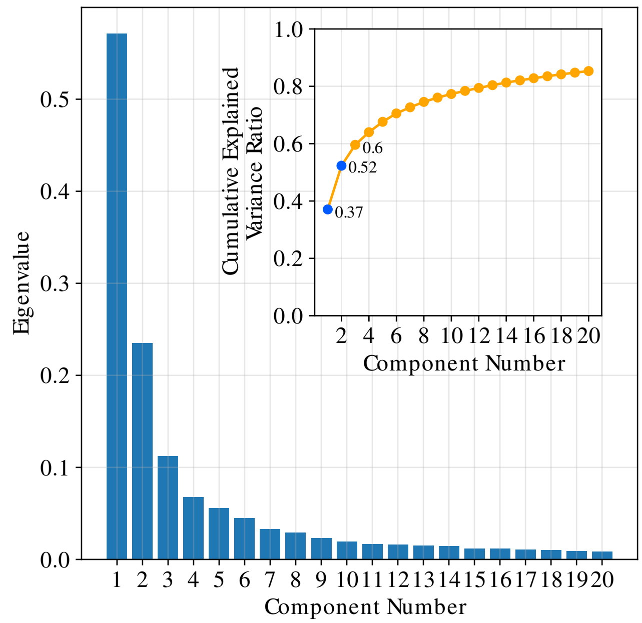

Figure 1: PCA Results

Sentiment Analysis We assessed the emotional polarity of each tweet by a well-vetted, off-the-shelf, open-source sentiment analysis tool, i.e. Valence Aware Dictionary and sEntiment Reasoner (VADER, Hutto & Gilbert, 2014), which operates on the basis of an extensive lexicon and a sophisticated set of syntactical and grammatical rules to assess the semantic polarity (valence) of a given tweet. VADER was specifically designed and validated to recognize features of online language (tweets), slang, emojis, acronyms, and common vernacular, as well as negations and hedging. In this current study, we used the compound score which indicates an overall Valence score of on a range of -1 (negative) to 1 (positive).

Social Network Features To assess an individual's online connectedness, we retrieved the number of their Twitter followers and friends. We also gathered the overall number of retweets and favorites for each tweet, and the number of mentions each user made to other users using the "@" tag. This allows us to cross-validate the generated user embeddings with

Table 1: Descriptives of Sample Dataset

		Lonely Group		Control Group	
		Number of accounts	Number of tweets	Number of accounts	Number of tweets
N		722 (100.00%)		722 (100.00%)	
Gender	Male	319 (44.18%)	256,830	319 (44.18%)	253,545
	Female	403 (55.82%)	309,538	403 (55.82%)	277,934
Age (years)	≤ 18	381 (52.77%)	293,226	381 (52.77%)	278,367
	19–29	260 (36.01%)	206,524	260 (36.01%)	180,735
	30–39	51 (7.06%)	46,305	51 (7.06%)	49,352
	≥ 40	30 (4.16%)	20,313	30 (4.16%)	23,025

the social connectedness of each user.

Results

User Embeddings

Figure 2 shows a transition from lonely to control group individuals as one move from left to right on the X axis, which represents the PCA component 1. Separation of density curves between groups suggests that lonely users are linguistically distinguishable from a random sample of users. SVM results also supported that the lonely and control groups can be separable from each other with a .70 accuracy and .66 F1 score.

To assess how individual language changes as we move along Component 1, we divided PCA Component 1 into multiple bins and aggregated the tweets of the users that belong to each bin. Next, we calculated the Shannon Entropy shift in bigrams for each bin against the opposite hemisphere of the PCA Component 1 to observe the word couples that characterize the discourse. Namely, if the given bin is lower than the median Component 1, it would be compared with the highest 50% of the bins of Component 1. In the same way, if the given bin is higher than the median Component 1, it would be compared with the lowest 50%. Then, we ranked the entropy shifts and showed randomly selected words in Fig. 2.b. As indicated by the Shannon Entropy shift, as Component 1 increases, the discourse changes towards terms indicative of subjective and internalizing thoughts rather than factual information such as “I feel”, “I wish”, “I just”, “so hard” or “makes me”. In the following sections, we discuss a quantitative understanding of how the language of these two groups changes as we move along Component 1.

Language Features Considering the relation between Component 1 and pronoun usage, we conducted Spearman’s rank correlation tests to examine the usage of each nominative person pronoun. The relation between Component 1 and “I” was positive, $r(1437) = .57$, $p < .001$. This suggests that usage of “I” increased from the left (control group) to the right (lonely group) along Component 1. However, the use of all other pronouns, except for “she” ($r(1437) = .01$, $p > .05$), increased when moving to the control group from the lonely group along Component 1. Specifically, usage of “you” $r(1437) = -.28$, $p < .001$., “he” $r(1437) = -.17$, $p < .001$., “we” $r(1437) = -.33$, $p < .001$., “they” $r(1437) = -.19$, $p < .001$., and “it” $r(1437) = -.11$, $p < .001$ were negatively related with Component 1.

Sentiment Analysis As we move from left (control group) to right (lonely group) along PCA Component 1, the compound VADER sentiment of tweets decreases $r(1437) = -.37$, $p < .001$ indicating that lonely individuals tend to use more negative language compared to the control group.

Social Network Features With respect to indicators of social relations, we observed that as Component 1 increases, the total number of Favorites received by the user decreases $r(1437) = -.06$, $p = .03$ as well. This statement holds for the total number of mentions made by the user $r(1437) = -.41$, $p < .001$, and the number of friends $r(1437) = -.18$, $p < .001$. These results suggest that lonely users, as expected, have a lower number of online friends and interact less with others, providing a degree of cross-validation to our observation.

Discussion

This study explores whether lonely individuals can be detected from their online language and if this is the case, which online factors are most strongly associated with loneliness. We demonstrated three main findings in this study which we briefly discuss below in light of previous literature. First, we demonstrated that at least one of the PCA components generated from BERT embeddings robustly distinguishes the lonely group from the control group. The robustness and reliability of this distinction is supported with an accurate, predictive SVM model. Although previous studies (Guntuku, Schneider, et al., 2019; Kivran-Swaine et al., 2014; Mahoney et al., 2019) were able to separate tweets of lonely individuals from others based on qualitative indicators of language use and themes, this is the first study as far as we know that demonstrates the capability to reliably and accurately separate lonely individuals from a random control group on the basis of their language use. Given that social media enables lonely people to express themselves (Leung, 2011), being able to identify lonely individuals based on individuals’ tweets is a valuable capability since it allows a more accurate characterization of loneliness from online data sources and

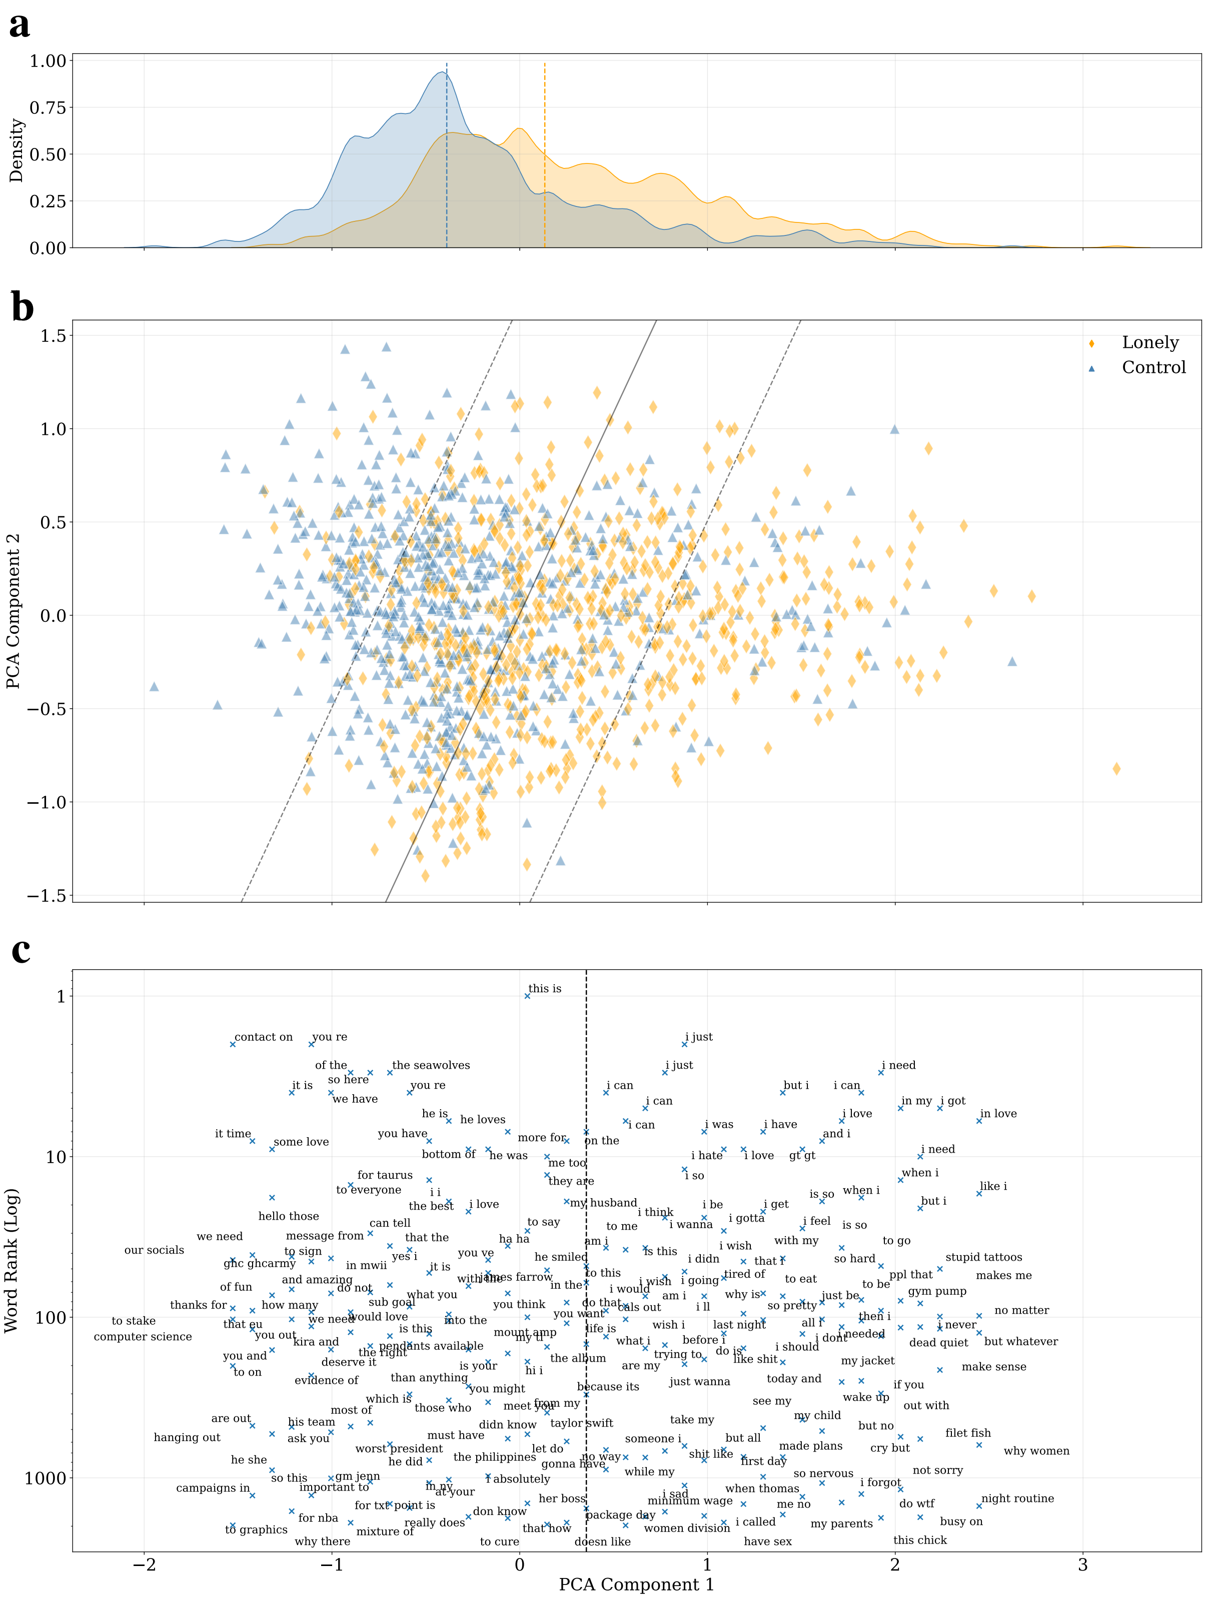

Figure 2: The x-axis for all graphs reflects PCA component.(a) Distribution of the PCA component 1 for the lonely and control groups (b) 2-dimensional representation of lonely and control group users. Solid and dashed lines represent the decision boundary and support vectors of SVM classifier, respectively. (c) Bi-grams according to their entropy shift ranks compared to the opposite 50% of the PCA Component 1 .

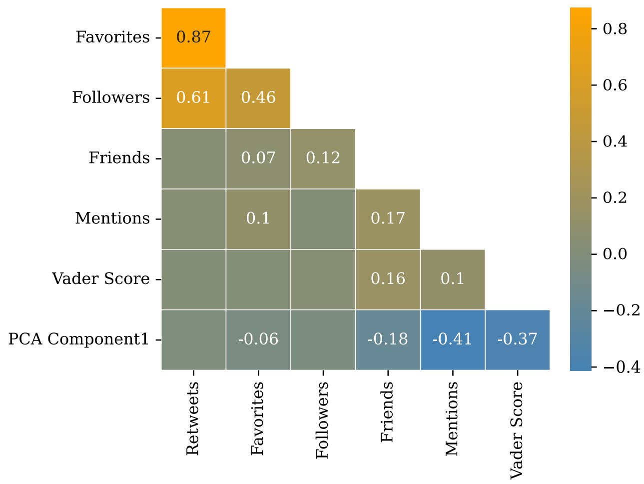

Figure 3: Correlation among network attributes and PCA Component 1

thereby underpins possible mitigation efforts of a significant public health concern.

Since loneliness is a precursor of several mental disorders (e.g., depression, alcohol abuse, personality disorders, Mushtaq et al., 2014), early detection of loneliness may contribute to the prevention of these disorders. Even though language models are still actively being developed toward better performance and accuracy, their performance may presently be sufficient to assist fast screenings of individuals for risk of loneliness at nearly societal scales. Future studies could benefit from fast computational methods to develop further intervention and prevention methods.

Second, our findings indicate that loneliness sharply increased prevalence of the “I” pronoun in language, whereas it decreases the usage of all other pronouns except for “she”. These findings are in line with previous studies that show the significant usage of first-person pronouns among lonely tweets (Guntuku, Schneider, et al., 2019; Mahoney et al., 2019). As previously observed (Guntuku, Schneider, et al., 2019; Jones et al., 1981; Kivran-Swaine et al., 2014; Mahoney et al., 2019), lonely individuals use more negative affect in their language. Considering these consistent findings and the fact that lonely people have a negative and self-oriented focus (Jones et al., 1981), these language patterns underpin the capability of NLP and AI methods to detect loneliness. Considering the increased prevalence of first-person pronouns in the language of depressed and stressed individuals (Guntuku, Buffone, et al., 2019; Guntuku et al., 2017; Holtzman et al., 2017), we caution that our results may be confounded by the comorbidity of loneliness and internalizing disorders such as depression. Although we only focused on people who self-identify themselves as lonely to eliminate loneliness related to boredom, or social isolation, we did not control for other mental health issues. Future studies may advance our study by controlling for these other sample characteristics.

Third, as we transition from the control group to the lonely group in Component 1, the number of friends, mentions and received likes also declined. Even though lonely people may not be particularly skilled in maintaining and forming relationships in a person (Wittenberg & Reis, 1986), activity on online platforms does not seem to compensate for this. The alignment of the individuals through the Component 1 was mostly related to the number of people they mentioned. Therefore, lonely people might be finding it difficult to reach out to other people. However, lonely individuals might still derive considerable benefits from their online activities, rather than face-to-face communication. Follow-up studies should investigate whether the in-person social networks of lonely individuals are less densely structured than their online networks. These results may indicate whether online behaviors conform to the individual’s social relations outside of online social media platforms.

Limitations

This study is limited in several important points which can be addressed in future studies. First, even though our lonely and control samples match in demographics, our sample is dominated by younger females. Considering that loneliness is most prevalent among young adults (age between 18-22, Cigna, 2020) and the fact that young adults are better represented in Twitter than the U.S. population (Wojcik & Hughes, 2019), this distribution is not surprising. However, future studies should examine a wider sample with various characteristics since the effects and patterns of loneliness differ among various demographics (Andy et al., 2022). Also, we only focused on language, sentiment, and social network features. In addition to these features, the real network of these Twitter users beyond the social media platform should be assessed as a basis of comparison. Understanding friendship networks may also be a very simple and exploratory feature to distinguish lonely users from others. The scope of this study is restricted to understanding the online lexical, semantic, emotional, and social features that mark the presence of loneliness for online individuals. The drivers and mechanics of the development and progress of loneliness among these people are not within the scope of this study but can be addressed with similar methods as we demonstrated in this study. An important distinction with our data is that we were only to capture the various features of loneliness among people who self-disclosed their lonely feelings on social media. However, not all lonely people may be willing to disclose their sentiments on these platforms, or even have social media accounts. As such, while we stand by our findings, we acknowledge they may not be generalizable to, or representative of, all lonely people. To meaningfully reach lonely people online, even without disclosures, improved models relying on other metrics than self-declaration are needed.

References

Abdi, H., & Williams, L. J. (2010). Principal component analysis. *Wiley Interdisciplinary Reviews: Computational Statistics*, *2*(4), 433–459.

Andy, A., Sherman, G., & Guntuku, S. C. (2022). Understanding the expression of loneliness on twitter across age groups and genders. *Plos ONE*, *17*(9), e0273636.

Beutel, M. E., Klein, E. M., Brähler, E., Reiner, I., Jünger, C., Michal, M., Wiltink, J., Wild, P. S., Münzel, T., Lackner, K. J., et al. (2017). Loneliness in the general population: Prevalence, determinants and relations to mental health. *BMC Psychiatry*, *17*(1), 1–7.

Boser, B. E., Guyon, I. M., & Vapnik, V. N. (1992). A training algorithm for optimal margin classifiers. *Proceedings of the fifth annual workshop on Computational learning theory*, 144–152.

Cacioppo, J. T., Hawkley, L. C., Berntson, G. G., Ernst, J. M., Gibbs, A. C., Stickgold, R., & Hobson, J. A. (2002). Do lonely days invade the nights? potential social modulation of sleep efficiency. *Psychological Science*, *13*(4), 384–387.

Cacioppo, J. T., Hughes, M. E., Waite, L. J., Hawkley, L. C., & Thisted, R. A. (2006). Loneliness as a specific risk factor for depressive symptoms: Cross-sectional and longitudinal analyses. *Psychology and aging*, *21*(1), 140.

Cacioppo, S., Capitanio, J. P., & Cacioppo, J. T. (2014). Toward a neurology of loneliness. *Psychological Bulletin*, *140*(6), 1464.

Caplan, S. E. (2003). Preference for online social interaction: A theory of problematic internet use and psychosocial wellbeing. *Communication Research*, *30*(6), 625–648.

Cer, D., Yang, Y., Kong, S.-y., Hua, N., Limtiaco, N., John, R. S., Constant, N., Guajardo-Cespedes, M., Yuan, S., Tar, C., et al. (2018). Universal sentence encoder. *arXiv preprint arXiv:1803.11175*.

Chen, J., Yan, Y., & Leach, J. (2022). Are emotion-expressing messages more shared on social media? a meta-analytic review. *Review of Communication Research*, *10*.

Conneau, A., Kiela, D., Schwenk, H., Barrault, L., & Bordes, A. (2017). Supervised learning of universal sentence representations from natural language inference data. *arXiv preprint arXiv:1705.02364*.

Devlin, J., Chang, M.-W., Lee, K., & Toutanova, K. (2018). Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.

Eichstaedt, J. C., Schwartz, H. A., Kern, M. L., Park, G., Labarthe, D. R., Merchant, R. M., Jha, S., Agrawal, M., Dziurzynski, L. A., Sap, M., et al. (2015). Psychological language on twitter predicts county-level heart disease mortality. *Psychological Science*, *26*(2), 159–169.

Golaszewski, N. M., LaCroix, A. Z., Godino, J. G., Allison, M. A., Manson, J. E., King, J. J., Weitlauf, J. C., Bea, J. W., Garcia, L., Kroenke, C. H., et al. (2022). Evaluation of social isolation, loneliness, and cardiovascular disease among older women in the us. *JAMA Network Open*, *5*(2), e2146461–e2146461.

Guntuku, S. C., Buffone, A., Jaidka, K., Eichstaedt, J. C., & Ungar, L. H. (2019). Understanding and measuring psychological stress using social media. *Proceedings of the international AAAI conference on Web and Social Media*, *13*, 214–225.

Guntuku, S. C., Schneider, R., Pelullo, A., Young, J., Wong, V., Ungar, L., Polsky, D., Volpp, K. G., & Merchant, R. (2019). Studying expressions of loneliness in individuals using twitter: An observational study. *BMJ Open*, *9*(11), e030355.

Guntuku, S. C., Yaden, D. B., Kern, M. L., Ungar, L. H., & Eichstaedt, J. C. (2017). Detecting depression and mental illness on social media: An integrative review. *Current Opinion in Behavioral Sciences*, *18*, 43–49.

Hawkley, L. C., & Capitanio, J. P. (2015). Perceived social isolation, evolutionary fitness and health outcomes: A lifespan approach. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *370*(1669), 20140114.

Holt-Lunstad, J., Smith, T. B., Baker, M., Harris, T., & Stephenson, D. (2015). Loneliness and social isolation as risk factors for mortality: A meta-analytic review. *Perspectives on Psychological Science*, *10*(2), 227–237.

Holtzman, N. S., et al. (2017). A meta-analysis of correlations between depression and first person singular pronoun use. *Journal of Research in Personality*, *68*, 63–68.

Hutto, C., & Gilbert, E. (2014). Vader: A parsimonious rule-based model for sentiment analysis of social media text. *Proceedings of the international AAAI Conference on Web and Social Media*, *8*(1), 216–225.

Jones, W. H., Freemon, J., & Goswick, R. A. (1981). The persistence of loneliness: Self and other determinants 1. *Journal of Personality*, *49*(1), 27–48.

Kim, J., LaRose, R., & Peng, W. (2009). Loneliness as the cause and the effect of problematic internet use: The relationship between internet use and psychological wellbeing. *Cyberpsychology & Behavior*, *12*(4), 451–455.

Kivran-Swaine, F., Ting, J., Brubaker, J., Teodoro, R., & Naaman, M. (2014). Understanding loneliness in social awareness streams: Expressions and responses. *Proceedings of the International AAAI Conference on Web and Social Media*, *8*(1), 256–265.

Le, Q., & Mikolov, T. (2014). Distributed representations of sentences and documents. *International conference on machine learning*, 1188–1196.

Lee, K.-T., Noh, M.-J., & Koo, D.-M. (2013). Lonely people are no longer lonely on social networking sites: The mediating role of self-disclosure and social support. *Cyberpsychology, Behavior, and Social Networking*, *16*(6), 413–418.

Leung, L. (2011). Loneliness, social support, and preference for online social interaction: The mediating effects of identity experimentation online among children and adolescents. *Chinese Journal of Communication*, *4*(4), 381–399.

Mahoney, J., Le Moignan, E., Long, K., Wilson, M., Barnett, J., Vines, J., & Lawson, S. (2019). Feeling alone among

317 million others: Disclosures of loneliness on twitter. *Computers in Human Behavior*, *98*, 20–30.

Morahan-Martin, J., & Schumacher, P. (2000). Incidence and correlates of pathological internet use among college students. *Computers in Human Behavior*, *16*(1), 13–29.

Mushtaq, R., Shoib, S., Shah, T., & Mushtaq, S. (2014). Relationship between loneliness, psychiatric disorders and physical health? a review on the psychological aspects of loneliness. *Journal of Clinical and Diagnostic Research: JCDR*, *8*(9), WE01.

Narchal, R., & McDavitt, S. A. (2017). Loneliness and anxiety sensitivity: Understanding behavioural avoidance in the lonely. *Acta Psychopathologica*.

Pennebaker, J. W., Boyd, R. L., Jordan, K., & Blackburn, K. (2015). *The development and psychometric properties of liwc2015* (tech. rep.).

Qi, P., Zhang, Y., Zhang, Y., Bolton, J., & Manning, C. D. (2020). Stanza: A python natural language processing toolkit for many human languages. *arXiv preprint arXiv:2003.07082*.

Sanh, V., Debut, L., Chaumond, J., & Wolf, T. (2019). Distilbert, a distilled version of bert: Smaller, faster, cheaper and lighter. *arXiv preprint arXiv:1910.01108*.

Solano, C. H., & Koester, N. H. (1989). Loneliness and communication problems: Subjective anxiety or objective skills? *Personality and Social Psychology Bulletin*, *15*(1), 126–133.

Wang, Z., Hale, S., Adelani, D. I., Grabowicz, P., Hartman, T., Flöck, F., & Jurgens, D. (2019). Demographic inference and representative population estimates from multilingual social media data. *The World Wide Web Conference*, 2056–2067.

Wittenberg, M. T., & Reis, H. T. (1986). Loneliness, social skills, and social perception. *Personality and Social Psychology Bulletin*, *12*(1), 121–130.

Wojcik, S., & Hughes, A. (2019). Sizing up twitter users. *PEW research center*, *24*, 1–23.

Yang, K.-C., Varol, O., Hui, P.-M., & Menczer, F. (2020). Scalable and generalizable social bot detection through data selection. *Proceedings of the AAAI conference on artificial intelligence*, *34*(01), 1096–1103.