

SOLO: A Corpus of Tweets for Examining the State of Being Alone

Svetlana Kiritchenko¹, Will E. Hipson², Robert J. Coplan², Saif M. Mohammad¹

¹National Research Council Canada, ²Carleton University, Ottawa, Canada

svetlana.kiritchenko@nrc-cnrc.gc.ca, williamhipson@cmail.carleton.ca,

robert.coplan@carleton.ca, saif.mohammad@nrc-cnrc.gc.ca

Abstract

The state of being alone can have a substantial impact on our lives, though experiences with time alone diverge significantly among individuals. Psychologists distinguish between the concept of *solitude*, a positive state of voluntary aloneness, and the concept of *loneliness*, a negative state of dissatisfaction with the quality of one’s social interactions. Here, for the first time, we conduct a large-scale computational analysis to explore how the terms associated with the state of being alone are used in online language. We present SOLO (*State of Being Alone*), a corpus of over 4 million tweets collected with query terms *solitude*, *lonely*, and *loneliness*. We use SOLO to analyze the language and emotions associated with the state of being alone. We show that the term *solitude* tends to co-occur with more positive, high-dominance words (e.g., *enjoy*, *bliss*) while the terms *lonely* and *loneliness* frequently co-occur with negative, low-dominance words (e.g., *scared*, *depressed*), which confirms the conceptual distinctions made in psychology. We also show that women are more likely to report on negative feelings of being lonely as compared to men, and there are more teenagers among the tweeters that use the word *lonely* than among the tweeters that use the word *solitude*.

Keywords: solitude, lonely, mental health, well-being

1. Introduction

We have all experienced the state of being alone one time or another: perhaps, a loved one was away, or our Instagram post did not stir up a barrage of likes, or we enjoyed a quiet hike, or we felt disconnected from those around us. Further, older people and young adults experience loneliness at markedly higher rates than others (Luhmann and Hawkey, 2016; Hawkey and Capitanio, 2015).

The state of being alone can have a substantial impact on our lives. On the one hand, *loneliness*—a negative and unwanted state of being alone—has been shown to be correlated with increased cognitive decline, dementia, depression, suicide ideation, self harm, and even death (Gerst-Emerson and Jayawardhana, 2015; Hawkey and Capitanio, 2015; Luo et al., 2012; Endo et al., 2017).¹ On the other hand, *solitude*—a positive and self-driven state of being alone—has been shown to improve autonomy, creativity, and well-being (Long et al., 2003; Knafo, 2012; Coplan and Bowker, 2017; Coplan et al., 2019a). Loneliness and solitude have also been shown to play a role in the adaptive fitness of our species (Hawkey and Capitanio, 2015; Larson, 1990). Thus loneliness and solitude are starting to receive substantial amounts of attention from the medical and psychological research. Yet, there is no large-scale computational work on analyzing the language of being alone.

Here, for the first time, we present a large corpus of tweets associated with the state of being alone. We will refer to it as the *State of Being Alone* corpus, or *SOLO* for short. SOLO includes over 4 million tweets, each of which consists of at least one of the following tokens: *solitude*, *lonely*, and *loneliness*. We use SOLO to analyze the language and emotions associated with the state of being alone. Specifically, we explore the following questions:

- When people use terms such as *solitude*, *alone*, and *loneliness* in tweets, how often are they referring to the state of being alone as opposed to some other sense of those words?
- Do we find evidence from the text that solitude is indeed more self-driven than loneliness (as theorized by psychologists)?
- Do we find evidence from the text that the speakers view solitude as a more positive concept than loneliness (as theorized by psychologists)?
- Which words are associated with solitude, and which words are associated with loneliness?
- Do different demographic groups (e.g., different genders, age groups, etc.) perceive solitude and loneliness differently?

Most of the past studies exploring such questions come from Psychology (see next section). They involve self-reports from a small number of people. Here, for the first time, we computationally examine millions of tweets associated with the state of being alone for the language used, and especially the emotion associations. We also make SOLO freely available for research.² We hope that this new dataset will bring fresh attention to the relationship between the state of being alone and our well-being.

2. Related Work

Time spent alone can have varying emotional effects. For instance, time alone is experienced negatively in those cases when we are unable to fulfill our needs for social interaction (Baumeister and Leary, 1995), but positively when we are exhausted from long periods of social interaction and desire time for relaxation and reflection (Nguyen et al., 2018; Long et al., 2003). Given that an estimated 25–33% of waking time is spent being alone (Larson et

¹The negative public health impacts of loneliness are so great that in 2018 the UK appointed a minister for loneliness.

²<https://svkir.com/projects/solo.html>

al., 1982), identifying and distinguishing between ‘positive’ and ‘negative’ instances of being alone has substantial implications for improving people’s well-being.

Many theoretical perspectives have emerged to explain these divergent experiences of being alone. Proponents of self-determination theory (Deci and Ryan, 2010) postulate that time alone that is intrinsically motivated (i.e., *choosing* to spend time alone) is better for one’s well-being than time alone that arises for external reasons (e.g., one who is alone due to the nature of their work) (Chua and Koestner, 2008; Nguyen et al., 2018).

However, the experience of being alone may also differ as a result of when this state arises. Someone who spends a lot of time alone may come to feel lonely because they perceive their social network as deficient (Hawkley and Cacioppo, 2010), in which case subsequent moments in solitude are likely to diminish in pleasantness. Conversely, someone who is inundated with social activity may become dissatisfied with the amount of time they get to spend alone (Coplan et al., 2019a), in which case being alone would be experienced as even more pleasant than usual. As far as we know, there has been no large-scale computational work examining text associated with the state of being alone.

Even though emotions are central to human experience and they have been studied for centuries, there are still many unknowns about their inner workings. Two prominent models of emotions are the dimensional model and the basic emotions model. As per the dimensional model (Osgood et al., 1957; Russell, 1980; Russell, 2003), emotions are points in a three-dimensional space of valence (positive–negative), arousal (active–passive), and dominance (dominant–submissive). Thus, when comparing the meanings of two words, we can compare their degrees of valence, arousal, or dominance. For example, the word *party* indicates more positiveness than the word *crying*; *terrible* indicates more arousal than *conversation*; and *hike* indicates more dominance than *abandoned*.

According to the basic emotions model (aka discrete model) (Ekman, 1992; Plutchik, 1980; Frijda, 1988), some emotions, such as joy, sadness, fear, etc., are more basic than others, and these emotions are each to be treated as separate categories.

We use the NRC Valence, Arousal, and Dominance (NRC VAD) lexicon (Mohammad, 2018a) and the NRC Emotion lexicon (Mohammad and Turney, 2013; Mohammad and Turney, 2010) to determine the emotion associations of the words in SOLO. These lexicons were created by manual annotation. The NRC VAD lexicon has valence, arousal, and dominance scores for over twenty thousand English terms, and it was created using a comparative annotation technique called Best-Worst Scaling (BWS) (Louviere, 1991; Louviere et al., 2015; Kiritchenko and Mohammad, 2016). It has been shown to have high reliability (repeated annotations produce similar association scores). The NRC Emotion lexicon has binary (associated or not associated) scores for about fourteen thousand English terms (a subset of terms in the VAD lexicon) with eight basic emotions (joy, sadness, fear, anger, surprise, anticipation, disgust, and trust) as well as positive and negative sentiment.

3. Creating the SOLO Corpus

We now describe how we collected tweets related to the state of being alone and created the SOLO corpus.

3.1. Query Term Selection

After consulting with psychologists on our team and utilizing different thesauri, we created a list of words and short phrases related to the state of being alone: *alone*, *alone time*, *aloneness*, *confinement*, *desert*, *detachment*, *get away from it all*, *get away from people*, *hermit*, *isolation*, *loneliness*, *lonely*, *lonesomeness*, *me time*, *peace and quiet*, *privacy*, *quarantine*, *reclusiveness*, *retirement*, *seclusion*, *separateness*, *serenity*, *silence*, *solitariness*, *solitude*, *tranquility*, *undisturbed*, *wilderness*, *withdrawal*. We collected tweets using these query terms for a few weeks, and then manually checked the relevance of the obtained tweets. Some query terms (e.g., *solitariness*, *reclusiveness*, *lonesomeness*, *aloneness*, *get away from it all*) were rarely used on Twitter and, therefore, were discarded. Some terms (e.g., *silence*, *privacy*, *retirement*, *desert*) were often used in other senses, not related to the state of being alone. Even for the query word *alone*, only about half of the collected tweets related to the concept of being alone. In many tweets, *alone* was used for emphasis (e.g., “*only you and you alone can thrill me like you do*”, “*I barely like Christmas music on Christmas lol, let alone in early November*”). After this manual inspection, we decided to keep three terms: *solitude* and *loneliness* (nouns), and *lonely* (adjective).

3.2. Collecting Tweets

SOLO Corpus: Tweets related to the state of being alone were collected by polling the Twitter API from August 28, 2018 to July 10, 2019 with the following query terms: *loneliness*, *lonely*, and *solitude*. We discarded duplicate tweets, short tweets (containing less than three words), and tweets with external URLs. Further, we kept only up to three tweets per user. This minimizes the impact of prolific tweeters and bots on the corpus. We refer to the combined set of the remaining tweets as the *State of being Alone* corpus, or *SOLO* for short. We refer to the individual sets of tweets as the *loneliness sub-corpus*, the *lonely sub-corpus*, and the *solitude sub-corpus*, respectively. Table 1 shows the number of tweets in each sub-corpus. In total, the SOLO Corpus contains over four million tweets.

General Tweets: As a control corpus, we collected tweets by polling the Twitter API from May 16, 2019 until June 12, 2019 using English function words (e.g., *is*, *on*, *they*, etc.) as query terms. Again, we discarded duplicate tweets, short tweets (containing less than three words), tweets with external URLs, and kept only up to three tweets per user. We will refer to this set of tweets as the *General Tweet Corpus*. It includes over 21 million tweets.

3.3. Tweet Volume

For the same time period (about a year), we were able to collect seventeen times more tweets with the word *lonely* and two-and-a-half times more tweets with the word *loneliness* than tweets with the word *solitude*. This suggests that

Corpus	# of tweets	# of users
SOLO Corpus:		
loneliness	489,264	408,659
lonely	3,339,166	2,443,210
solitude	191,643	158,878
All	4,020,073	3,010,747
General Tweet Corpus	21,719,409	12,096,240
Total	25,739,482	15,106,987

Table 1: The number of tweets for each query term.

Corpus	Percentage of relevant tweets
loneliness	93%
lonely	96%
solitude	92%
Average	94%

Table 2: Percentage of relevant tweets for each query term.

most users refer to the state of being alone through the use of words *lonely* and *loneliness*, and rarely using the word *solitude*. In a period of one year, close to three million users posted at least one tweet with the words *lonely* or *loneliness*, which reflects the magnitude of the loneliness problem.

4. Assessing Relevance of the SOLO Tweets to the State of Being Alone

A tweet may include the term *loneliness*, *lonely*, or *solitude* and yet may not be relevant to the state of being alone. Thus we manually examined a small sample of SOLO to determine the percentage of relevant tweets. We considered a tweet to be relevant if it directly referred to the state of being alone. This included (but was not limited to):

- a personal statement about being alone,
- a statement about other people being alone,
- a general statement about aspects of being alone,
- a message of support (e.g., “*you are not alone*”),
- a quote from literature about being alone.

We considered tweets to be irrelevant if the query word (*loneliness*, *lonely*, *solitude*) was used as part of a title (of a book, song, etc.) or a name (of a place, a stadium, etc.). Tweets containing advertisements were also considered irrelevant.

For each query term, we randomly selected 100 tweets with that term and counted the percentage of relevant tweets. Table 2 shows the results. Observe that for all the query terms, over 90% of examined tweets were relevant to the state of being alone. This confirms the suitability of the SOLO Corpus for studying the everyday language associated with the state of being alone.

Categories	loneliness	lonely	solitude
first-hand experience	0.35	0.62	0.47
other people’s experience	0.15	0.16	0.09
general statement	0.30	0.09	0.21
literary quote	0.19	0.06	0.16
offering support	0.00	0.01	0.05
other	0.01	0.06	0.02

Table 3: Different types of SOLO tweets and their relative frequency in each sub-corpus.

5. Analyzing the Language and Emotions Associated with the State of Being Alone

We examine the language of the SOLO tweets to determine if the concept words *loneliness*, *lonely*, and *solitude* tend to be used in different emotional contexts. In particular, we explore the question of whether Twitter users perceive the concept of solitude as more positive and self-driven and the concept of loneliness as more negative and externally imposed as suggested by psychology literature. For this, in Section 5.1, we manually analyze a sample of tweets for the types of contexts in which people use the words *loneliness*, *lonely*, and *solitude*. We also computationally identify and compare words strongly associated with each of these terms. In Section 5.2, we examine the words occurring in SOLO for their emotional associations.

5.1. Language Associated with Being Alone

First, we look at how people use the terms *loneliness*, *lonely*, and *solitude* in everyday language of tweets. Do people often describe their own feelings and experiences or offer support to other people? Do they just make general statements about different aspects of being alone? Which words are most likely to co-occur with these terms?

Manual Examination of the SOLO Tweets: We manually examined randomly selected samples of 100 tweets from the loneliness, lonely, and solitude sub-corpora to identify the types of messages users are likely to post using these terms. Table 3 shows the results.

In tweets with the word *solitude*, people often describe their own experiences and attitudes (e.g., “*I fell in love with my solitude.. everything changed after that.*”), provide general statements about positive or negative aspects of being alone (e.g., “*Solitude can be either comforting or really painful.*”), and cite relevant quotes from notable people and literary sources (e.g., “*The monotony and solitude of a quiet life stimulates the creative mind - Albert Einstein*”). They less often discuss other people’s experiences (e.g., “*It seems like they hate everything that isn’t profitable - whether it’s wolves, wild horses, stunning landscapes, solitude...*”) or offer support (e.g., “*that is saaaaaddd, but don’t worry, solitude is a nice friend*”).

When people use the word *lonely*, they mostly report on their own feelings (e.g., “*Feeling lonely and forgotten :/*”) and those of other people (e.g., “*Well you were clearly very lonely.*”). In tweets with the word *loneliness*, users less often describe their own experiences (e.g., “*The level of loneliness I’ve reached is at an all time high*”), and more often make general statements (e.g., “*we don’t know how*

SOLO term	Words associated with the term
loneliness	<i>alone, feeling, lonely, depression, pain, sadness, isolation, fear, killing, feelings, anxiety, happiness, cure, solitude, hurts, emptiness, crippling, anger, silence, fill, suffering, relationships, empty, darkness, boredom</i>
lonely	<i>feel, sad, feeling, alone, friends, sometimes, single, felt, bored, feels, nights, scared, depressed, af, cold, island, christmas, empty, hearts, loneliness, miserable, surrounded, horny, asf, desperate</i>
solitude	<i>alone, enjoy, peace, silence, loneliness, fortress, quiet, hundred, lonely, enjoying, comfort, prefer, nature, isolation, comfortable, bliss, moments, sea, presence, peaceful, seek, embrace, darkness, gabriel, inner</i>

Table 4: The most frequent words strongly associated with the terms *loneliness*, *lonely*, and *solitude*.

to appreciate loneliness”) and quote celebrities and literary sources (e.g., “If you are afraid of loneliness, don’t marry. -Anton Chekhov”), than in tweets with the words *lonely* and *solitude*.

Notably, in 14% of tweets from the solitude sample, tweeters explicitly assert their need to spend some time alone to reflect, heal, or focus on important tasks (e.g., “It’s funny how the universe works...this moment of solitude was unplanned but definitely needed.”).

Words Associated with Loneliness, Lonely, and Solitude: We identify words that are associated with the SOLO query terms, *loneliness*, *lonely*, and *solitude*, i.e., words that tend to appear in tweets with these query terms more often than they do in the General Tweet Corpus. For this, we calculate an association score of a word w with the target sub-corpus C_{target} ($target \in \{loneliness, lonely, solitude\}$) as compared to the corpus of general tweets (the reference corpus, $C_{reference}$):

$$Score(w) = PMI(w, C_{target}) - PMI(w, C_{reference}) \quad (1)$$

PMI stands for pointwise mutual information:

$$PMI(w, C_{target}) = \log_2 \frac{freq(w, C_{target}) * N}{freq(w) * freq(C_{target})} \quad (2)$$

where $freq(w, C_{target})$ is the number of times the word w occurs in the target corpus, $freq(w)$ is the total frequency of the word w in the two corpora (target and reference), $freq(C_{target})$ is the total number of words in the target corpus, and N is the total number of words in the two corpora. $PMI(w, C_{reference})$ is calculated in a similar way. Thus, Equation 1 is simplified to:

$$Score(w) = \log_2 \frac{freq(w, C_{target}) * freq(C_{reference})}{freq(w, C_{reference}) * freq(C_{target})} \quad (3)$$

Since PMI is known to be a poor estimator of association for low-frequency events, we ignore terms that occur less than 25 times in total in both corpora.

Association scores can range from $-\infty$ to $+\infty$; in practice, however, they usually range from around -6 to 6 . A positive score indicates a greater overall association with the target corpus, that is the word appears at a higher rate (more occurrences per 100 words) in the target corpus than in the reference corpus. A negative score indicates that a word appears at a lower rate in the target corpus than in the

reference corpus. The magnitude is indicative of the degree of association. Note that there exist numerous other methods to estimate the degree of association of a word with a category (e.g., cross entropy, Chi-squared test, and information gain). We have chosen PMI because it is simple and robust and has been successfully applied in a number of NLP tasks (Clark et al., 2016; Kiritchenko et al., 2014).

We calculate association scores with the loneliness, lonely, and solitude sub-corpora for all words in the SOLO corpus. We say that a word is **strongly associated** with a sub-corpus if the corresponding association score is greater than or equal to 1.5.³ Table 4 shows 25 most frequent words in the loneliness, lonely, and solitude sub-corpora that are strongly associated with them. Observe that the words strongly associated with *solitude* are mostly positive. Tweets in the solitude sub-corpus tend to describe peaceful, enjoyable moments, often in the natural surroundings. The presence of high-dominance words, such as *enjoy*, *prefer*, and *comfort*, indicate that the person most likely feels in control over a situation, that the time alone was self imposed and desirable. Words strongly associated with *lonely* and *loneliness*, on the other hand, are mostly negative and low in dominance. These tweets often refer to the feelings of sadness, anxiety, depression, and boredom. Words like *friends*, *relationships*, and *Christmas* probably reflect the unfulfilled need for social interaction that is often felt more strongly during traditional family holidays like Christmas.

Solitude–Loneliness Dimension of Word Association:

We can use the *solitude* corpus to study how people talk about solitude. Similarly, we can use the *lonely* and *loneliness* corpora (jointly) to study how people talk about loneliness.⁴ In the sub-section above, we explored each of the query term sub-corpora in comparison with the General Tweets Corpus. Here, in order to determine the extent to which words are associated with solitude as opposed to loneliness, we calculate the solitude–loneliness association score as shown below:

$$Score(w) = PMI(w, C_{solitude}) - PMI(w, C_{loneliness}) \quad (4)$$

Using this score we can place words along the solitude–loneliness dimension, where words strongly associated with solitude but not with loneliness are towards one end

³The threshold of 1.5 is somewhat arbitrary, but reasonable.

⁴We use the italicized term (e.g., *loneliness*) to refer to the query term, and the non-italicized form (e.g., loneliness) to refer to the mental/physical state.

SOLO term	Words associated with the term
solitude	<i>enjoy, peace, silence, fortress, quiet, hundred, enjoying, prefer, nature, bliss, complete, presence, peaceful, seek, embrace, gabriel, inner, marquez, value, spiritual, noise, superman, competing, recharge, prayer</i>
lonely & loneliness	<i>feeling, im, sad, girl, ass, nights, lmao, bitch, boy, baby, scared, bored, girls, hi, cuz, somebody, depressed, hearts, sucks, broke, club, af, pls, hurts, cute</i>

Table 5: The most frequent words strongly associated with *solitude* as opposed to *lonely* and *loneliness*.

Sentiment	loneliness	lonely	solitude
positive	0.05	0.03	0.71
negative	0.71	0.84	0.11
mixed	0.14	0.06	0.18
unclear	0.10	0.07	0

Table 6: Proportions of the SOLO tweets with different sentiments towards the state of being alone.

and words strongly associated with loneliness but not with solitude are towards the other end.

Table 5 shows 25 most frequent words that are more strongly associated with solitude than with loneliness (solitude–loneliness association score ≥ 1.5), and 25 most frequent words that are more strongly associated with loneliness than with solitude (solitude–loneliness association score ≤ -1.5). Observe that words that are more strongly associated with solitude than with loneliness are positive and high dominance words. These are words referring to peaceful and spiritual activities of being with oneself, recharging, and enjoying the present moment. In contrast, the words more strongly associated with loneliness than with solitude refer to negative personal experiences of being sad, scared, bored, hurt, and broken-hearted.

5.2. Emotions Associated with Being Alone

In this section, we measure the emotional context in which the SOLO query terms, *loneliness*, *lonely*, and *solitude*, occur. In particular, we investigate whether people use these terms in different emotional contexts and whether they are associated with the qualities suggested in the psychology literature. We analyze a sample of the SOLO corpus manually and the full corpus computationally using existing word–emotion association lexicons.

Manual Examination of Sentiment in the SOLO Corpus: We randomly sampled 100 tweets each from the loneliness, lonely, and solitude sub-corpora, and manually examined each of these tweets to determine whether they express positive, negative, or mixed attitudes towards the state of being alone. Table 6 shows the results.

Observe that tweeters that use the term *solitude* mostly have a positive attitude towards being alone (e.g., “*Have you ever: felt lonely? No, I love my solitude.*”), yet sometimes mixed (e.g., “*What is the balance for those of us that love the solitude but wanna have companionship ?*”) or even negative (e.g., “*Some people prefer to live in solitude, but no one can withstand it*”) sentiments can be expressed. On the other hand, the vast majority of tweeters that use the words *lonely* and *loneliness* have a negative attitude to-

wards being alone (e.g., “*i’m really lonely and really sad*”). Only rarely do people include the words *lonely* and *loneliness* when they express positive sentiments in the SOLO tweets (e.g., “*Loneliness is designed to help you discover who you are ... and to stop looking outside yourself for your worth. — Mandy Hal*”).

Basic Emotions Associated with Words in SOLO: Next, we look at the whole SOLO Corpus and analyze emotions associated with words occurring in the SOLO tweets. We use the NRC Word–Emotion Association Lexicon (Mohammad and Turney, 2013; Mohammad and Turney, 2010) which has entries for over 14,000 English common words.⁵ It provides labels for eight basic emotions (anger, fear, sadness, disgust, joy, anticipation, surprise, and trust) and two sentiments (positive and negative). The labels are binary indicating whether a word is associated with an emotion (or sentiment) or not. The lexicon was created by crowdsourcing the annotations. We consider only those words in SOLO that appear in the lexicon, and count the percentage of words associated with each emotion (i.e., out of every 100 words, how many are associated with sadness, joy, etc.). (The SOLO query words are excluded from the analysis.)

Figure 1 shows the results for the different sub-corpora of the SOLO corpus. For comparison, we also show the results for the General Tweets Corpus. For each emotion, the differences between the word percentages for the sub-corpora are statistically significant (Chi-squared test, $p < 0.0001$). Observe that tweets in the solitude sub-corpus contain more words associated with the positive sentiment and more words associated with the emotions of joy, anticipation, and trust than the tweets in other sub-corpora, including the general tweets. There are 25–30% more positive words in the solitude tweets than in the lonely and loneliness tweets. On the other hand, tweets with the words *lonely* and *loneliness* have more words associated with the negative sentiment and more words associated with the emotions of anger, fear, sadness, and disgust. There are 60% more negative words in the loneliness tweets than in the solitude or the general tweets. Somewhat surprisingly, tweets in the loneliness sub-corpus have significantly more (20–40%) words associated with the negative sentiment and the negative emotions of anger, fear, and sadness than tweets in the lonely sub-corpus.

Valence-Arousal-Dominance of Words in SOLO: To analyze the SOLO corpus with regard to the dimensional theory of emotions, we use the NRC Valence, Arousal,

⁵<http://saifmohammad.com/WebPages/NRC-Emotion-Lexicon.htm>




Figure 1: The percentage of words associated with eight basic emotions in different sub-corpora.

and Dominance (VAD) Lexicon (Mohammad, 2018a).⁶ The VAD lexicon provides real-valued ratings of valence, arousal, and dominance for over 20,000 English words. The scores range from 0 to 1 along each of the three dimensions: valence (from maximally unpleasant to extremely pleasant), arousal (from maximally calm, sleepy to maximally active, intense), and dominance (from maximally weak to maximally powerful). The annotations were obtained through crowd-sourcing.

We consider words that appear in the VAD lexicon, and count the percentage of words that have high/low valence, arousal, and dominance scores. (The SOLO query words are excluded from the analysis.) For all three dimensions, we consider scores greater than or equal to 0.75 as high scores, and scores lower than or equal to 0.25 as low scores. Table 7 shows the percentage of words in the different sub-corpora with high/low valence, arousal, and dominance scores. Within each row, all the differences are statistically significant (Chi-squared test, $p < 0.0001$).

Dimension	general	loneliness	lonely	solitude
Valence				
low	9.3	15.8	12.3	9.2
high	29.4	30.2	30.3	33.7
Arousal				
low	9.1	10.9	11.5	14.4
high	8.3	8.6	7.2	6.2
Dominance				
low	4.8	8.3	8.5	7.1
high	11.9	9.9	7.9	12.3

Table 7: The percentage of words with high/low valence, arousal, and dominance scores in the SOLO corpus. ‘general’ stands for ‘General Tweets Corpus’. The highest numbers in each row are in bold. Within each row, all the differences are statistically significant (Chi-squared test, $p < 0.0001$).

We can see again that the solitude tweets have the highest number of strongly positive words (high valence), and the lonely and loneliness tweets have the most strongly negative words (low valence). The loneliness corpus has the highest number of negative words, 72% more than the solitude corpus. The lonely and loneliness sub-corpora also have more high-arousal words than the solitude corpus, while the solitude corpus has the highest amount of low-arousal words. The solitude tweets tend to describe quiet and relaxing moments, in natural surroundings, with no agenda to follow. When lonely, people can feel scared and anxious, showing more arousal. Also, loneliness is associated with both momentary and chronic stress, which may explain why *lonely* occurs among higher arousal words (Seeman, 1996). The solitude corpus has the most high-dominance words, 56% more than the lonely corpus and 24% more than the loneliness corpus. This is consistent with the conceptual definition of *solitude* as a positive, voluntary state of being alone. In contrast, when feeling lonely, people usually perceive the situation as undesirable, they feel scared, depressed, miserable, and powerless.

VAD Trends Along the Solitude–Loneliness Dimension of Word Association: We analyze the trends in valence, arousal, and dominance scores along the solitude–loneliness dimension. We use the solitude–loneliness association scores for words computed as described in Section 5.1. We order the words by their solitude–loneliness association scores from smallest to largest, bin the scores with a 0.5 step, and average the valence, arousal, and dominance scores for all words that fall in each bin. For example, for bin with the score of 1 we average the VAD scores of all the words whose association scores fall in the range [0.75, 1.25). The VAD scores for words are taken from the NRC VAD Lexicon. Figure 2 shows the trends in the average VAD scores along the solitude–loneliness dimension. (Only bins with at least 100 words are shown.) Recall that words with positive association scores occur at a higher rate in the solitude sub-corpora and at a lower rate in the lonely and loneliness sub-corpora while the words with negative association scores occur at a higher rate in the lonely and loneliness sub-corpora and at a lower rate in the solitude

⁶<http://saifmohammad.com/WebPages/nrc-vad.html>




Figure 2: Trends in average valence (V), arousal (A), and dominance (D) scores along the solitude–loneliness dimension of association. Positive association scores indicate the word’s stronger association with solitude than with loneliness; negative association scores indicate stronger association with loneliness than with solitude.

sub-corpus. Along all three dimensions (valence, arousal, and dominance), the trends are very consistent: the more the word is associated with *solitude*, the higher its valence and dominance scores are, and the lower its arousal score is. While the range of the average arousal scores is relatively small (from 0.48 to 0.55), the differences in the average valence and dominance scores are substantial (from 0.45 to 0.59 for valence, and from 0.43 to 0.56 for dominance). This once again supports the hypothesis that solitude is often viewed as positive, intrinsically motivated state of being alone, and loneliness is viewed as negative, externally imposed state of being alone.

6. Demographic Differences in the Language Associated with the State of Being Alone

In this section, we examine the differences in the language and emotions associated with the state of being alone between genders (male vs. female) and age groups (adolescents vs. adults). Researchers have long been interested in exploring differences in language use between genders in different communication media and sociocultural contexts (Park et al., 2016; Coates, 2015). Here, we continue this line of work and investigate whether men and women tend to use the SOLO concepts, *loneliness*, *lonely*, and *solitude*, in different emotional contexts. Psychologists are also interested in identifying developmental differences in the perception and experiences with the state of being alone

Corpus	Total tweets	Tweets with inferred gender
General Tweets	21,719,409	8,355,543 (38%)
SOLO Corpus:		
loneliness	489,264	169,305 (35%)
lonely	3,339,166	1,131,935 (34%)
solitude	191,643	68,721 (36%)

Table 8: The total number of tweets with inferred gender of the tweeter.

Corpus	Tweets written by	
	Females	Males
General Tweets	3,730,986 (45%)	4,624,557 (55%)
SOLO Corpus:		
loneliness	87,228 (52%)	82,077 (48%)
lonely	636,388 (56%)	495,547 (44%)
solitude	33,000 (48%)	35,721 (52%)

Table 9: The number of tweets written by (inferred) female and male users.

(Coplan et al., 2019b). Using the large amounts of tweets in the SOLO Corpus and an existing word–age association lexicon, we analyze the tendency of different age groups to describe their experiences of being alone as solitude or loneliness states.

6.1. Gender Differences in the Language Associated with the State of Being Alone

To infer the gender of the tweeters, we use the US Social Security Administration database⁷. From the database, we select first names that occur more than 100 times in total over the years from 1940 until 2017 and that were used for males (females) at least 95% of the times. In total, we found 19,714 female and 10,909 male such names. We split the user names of the tweeters by punctuation marks and match the first token against the selected first names. If the first token matches one of the female (male) first names, the user is considered female (male).

Table 8 shows the number of tweets with inferred tweeter gender for each sub-corpora. We are able to infer gender of the tweeter in 34%–38% of the tweets. Table 9 shows the percentage of tweets written by female and male users. Notice that in the General Tweets, the majority of the tweets with inferred gender is from male users (55%). Similar percentage of male users is inferred in the solitude sub-corpora (52%). However, in the lonely and loneliness sub-corpora the majority of the inferred users are female (56% and 52%, respectively). This suggests that women have and/or report their negative experiences of being alone more often than men.

⁷<https://www.ssa.gov/oact/babynames/limits.html>. We acknowledge that users may identify their gender as non-binary, but we did not have the data to explore this. We also acknowledge that US Social Security information is not representative of the names from around the world. Thus, the gender analysis is mostly representative of US residents.

emotion	general	loneliness	lonely	solitude
negative	-0.1	-0.2	0.5	-0.6
positive	1.7	0.9	0.5	1.6
anger	-0.2	-0.1	0.2	-0.3
anticipation	1.0	0.1	0.5	1.3
disgust	0.4	-0.1	0.4	-0.2
fear	-0.2	-0.1	0.1	-1.2
joy	3.0	0.9	1.4	2.6
sadness	0.4	0.4	0.8	-0.6
surprise	0.4	0.2	0.3	0.3
trust	0.0	0.5	0.1	0.3

Figure 3: The differences in percentages of words associated with eight basic emotions in tweets written by female and male users. Positive scores (shown in red) indicate that females tend to use more words associated with this emotion than males do. Negative scores (shown in blue) indicate that males tend to use more words associated with this emotion than females do. Darker shades of red/blue highlight differences with larger absolute values. ‘general’ stands for ‘General Tweets Corpus’.

To examine the differences in emotional content of tweets written by different genders, we perform analyses of basic emotions and valence, arousal, and dominance in a similar manner as described in Section 5.2. The analyses are performed separately on the tweets written by male users and on the tweets written by female users. Figure 3 shows the differences in percentages of words associated with eight basic emotions in tweets written by female and male users. Observe that in the General Tweets Corpus the differences are minor, most of them are below 1%. The only differences that are 1% or larger are for the emotions of joy (3% more in text written by women) and anticipation (1% more in text written by women) as well as for positive sentiment (1.7% more in text written by women). We see similar trends in the solitude sub-corpus: the only differences that are larger than 1% in absolute values are for the emotions of joy, anticipation, fear, and for positive sentiment. In the lonely and loneliness sub-corpora, the differences across genders are even smaller—below 1% for all, except for the emotion of joy in the lonely sub-corpus (1.4%). The results for valence, arousal, and dominance are also similar (numbers not shown here). Overall, within tweets associated with the state of being alone, the differences in emotional content across the two genders are small.

6.2. Age Differences in the Language Associated with the State of Being Alone

Since we do not have age information for the tweeters in our corpus, we use an available Word–Age Association Lexicon (Schwartz et al., 2013). This lexicon provides association scores and the corresponding p-values for common words and phrases (1-grams, 2-grams, and 3-grams) with four age groups: 13 to 18 years old, 19 to 22 years old, 23 to 29 years old, and 30 and over years old. Schwartz

Corpus	Percentage of words associated with an age group:			
	13 to 18	19 to 22	23 to 29	30+
General Tweets	31.0	5.1	10.5	55.3
SOLO Corpus:				
loneliness	29.8	5.4	9.8	54.4
lonely	37.5	7.4	8.8	48.2
solitude	27.4	4.6	11.0	57.2

Table 10: Percentage of words associated with different age groups. Within each age group (column), all the differences are statistically significant (Chi-squared test, $p < 0.0001$).

et al. (2013) collected Facebook messages of 75,000 volunteers, along with the information on their age and gender. Then, they calculated the association scores by fitting a linear function between the target variable (word’s relative frequency) and the dependent variable (age), adjusted for gender. The lexicon includes only those words and phrases that were used by at least 1% of all subjects. From the lexicon, for each age group, we collect single, alpha-numeric tokens that are significantly positively associated with the age group ($p \leq 0.05$). Out of 8,093 single, alpha-numeric tokens in the lexicon, 1,921 were significantly positively associated with the 13 to 18 years old group, 845 were significantly positively associated with the 19 to 22 years old group, 1,130 were significantly positively associated with the 23 to 29 years old group, and 3,055 were significantly positively associated with the 30 and over years old group.

Using the Word–Age Association Lexicon, we calculate the percentage of words associated with each age group in each sub-corpus (loneliness, lonely, solitude, and general tweets). For this, we divide the number of occurrences of words associated with a particular age group by the total number of occurrences of all the words in the lexicon. Table 10 shows the results. Within each age group, all the differences between the numbers for each sub-corpus (loneliness, lonely, and solitude) and the general tweets are statistically significant (Chi-squared test, $p < 0.0001$). Observe that the lonely sub-corpus has higher percentages of words associated with the two younger groups (as compared with the general tweets) and lower percentage of words associated with the two older groups. The differences for the teenage group and the older adults (30+ years old) are particularly large (21% increase for the teenage group, 13% decrease for the 30 and over group). The solitude sub-corpus shows the opposite pattern with lower percentage of words associated with the two younger groups and higher percentage of words associated with the two older groups. The differences between the numbers for the loneliness corpus and the general tweets are relatively small for all four age groups. These results suggest that there are more younger people (especially teenagers) among the tweeters that use the word *lonely* when talking about being alone and, therefore, have more negative experiences when alone, than among the tweeters that use the word *solitude* and have more positive attitudes to the state of being alone. This finding does not support the psychology literature that proposes that adolescence may be a time when being alone is adap-

tive and enjoyable (Coplan et al., 2019b). It is possible, however, that adolescents may use Twitter to vent or share feelings about loneliness more often than other age groups.

7. Applications

In this section, we list the potential applications and the directions for future work using the resources created as part of this project: the SOLO Corpus, the lexicons of words associated with the SOLO concept terms, and the list of search terms related to the concept of being alone.

SOLO Corpus: The corpus can be used to further study how people understand and experience the state of being alone, and how these vary across situations, individuals, and development. For example, the following research questions can be addressed:

- How do people understand different experiences that could be considered ‘solitary’? Do people distinguish between different degrees of solitude? For example, is someone more ‘alone’ if they are away from their Smartphone?
- What are the different motivations (intrinsic and extrinsic) for people to spend time alone?
- Do people recognize some solitary experiences as being more beneficial or costly than others? What are the different benefits that might arise from being alone (e.g., creativity, relaxation, productivity)?
- Can we identify developmental differences in experiences and attitudes towards being alone?

Words associated with the SOLO concept words: Words highly associated with the terms *loneliness*, *lonely*, and *solitude* can be used to identify pieces of text that do not necessarily mention either of these three words, but nevertheless discuss the experiences of being alone. This can apply to tweets, but also to other types of text (blogs, emails, novels, etc.). For example, texts rich in words highly associated with *lonely* have a high probability of discussing feelings of being lonely and the related issues even if the word *lonely* itself is not mentioned.

Search terms: We have shown that by using the search terms *loneliness*, *lonely*, and *solitude* we can collect voluminous corpora of tweets highly related to the state of being alone. Therefore, this search strategy over the Twitter stream can be used to monitor the positive and negative aspects of being alone and their relation to well-being over the entire population across time, geographical regions, and demographic groups.

Building Other SOLO Corpora: The approach presented in this paper can also be used to create other more focused corpora pertaining to specific demographics for whom solitude and loneliness are particularly relevant, such as the elderly and teenagers (Luhmann and Hawkey, 2016; Hawkey and Capitanio, 2015).

8. Conclusion

We presented the *SOLO (State of Being Alone)* corpus—a large corpus of tweets associated with the state of being alone. SOLO includes over 4 million tweets collected using one of the three terms *loneliness*, *lonely*, and *solitude*. Manual examination showed that the corpus contains over 94% of the tweets related to the concept of being alone.

We used the SOLO Corpus to examine the language and emotions associated with the state of being alone. We found evidence that Twitter users tend to use the word *solitude* to describe more positive and self-imposed states of being alone, and tend to use the words *lonely* and *loneliness* when their experiences are negative and undesirable, which is consistent with conceptual definitions proposed in psychology literature. Furthermore, we found that the word *loneliness* tends to be used in more negative contexts than the word *lonely*.

Over the same period of time, the term *lonely* triggered 17 times more tweets than the term *solitude*. There were 12% more tweets with the word *lonely* written by female users than tweets written by males, even though in the General Tweet Corpus (used as control) there were 10% more tweets written by male users. However, the emotional content in the SOLO tweets written by male and female users was strikingly similar. We also found more words associated with the adolescent age group (especially, teenagers) and less words associated with the adult age group in the lonely corpus as compared to the solitude corpus, which suggests a higher vulnerability of teenagers to the negative experiences of feeling lonely.

We make SOLO and other resources created in this project freely available to encourage further research on health, economical, and other issues related to people’s experiences of being alone and how these issues affect the population’s well-being.

The current study focused on English-language social media, in particular tweets. In future work, texts from other genres, such as blogs, news, poetry, and fiction, can be analyzed in a similar manner. While this study examined the percentage of basic emotion words, one can also use lexica such as the NRC Emotion Intensity Lexicon (Mohammad, 2018b) to examine the use of high and low intensity emotion words in expressions of solitude.⁸ By comparing sources from different time periods, we can track how people’s perception of solitude and loneliness change over time. Furthermore, parallel studies in other languages can shed light on cultural differences in people’s attitudes towards and experiences with the state of being alone.

Finally, we are exploring the creation of corpora similar to SOLO with a focus on text generated by specific demographics such as teenagers, elderly, as well as, those coping with disabilities, stress, or other mental and physical conditions. We believe that a better understanding of people’s attitudes towards solitude and loneliness will help identify new ways to improve their well-being.

Acknowledgments

We thank Samuel Larkin for help in collecting tweets.

⁸<http://saifmohammad.com/WebPages/AffectIntensity.htm>

9. Bibliographical References

- Baumeister, R. F. and Leary, M. R. (1995). The need to belong: desire for interpersonal attachments as a fundamental human motivation. *Psychological Bulletin*, 117(3):497.
- Chua, S. N. and Koestner, R. (2008). A self-determination theory perspective on the role of autonomy in solitary behavior. *The Journal of Social Psychology*, 148(5):645–648.
- Clark, P., Etzioni, O., Khot, T., Sabharwal, A., Tafjord, O., Turney, P., and Khashabi, D. (2016). Combining retrieval, statistics, and inference to answer elementary science questions. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Coates, J. (2015). *Women, men and language: A sociolinguistic account of gender differences in language*. Routledge.
- Coplan, R. J. and Bowker, J. C. (2017). “Should we be left alone?” psychological perspectives on the costs and benefits of solitude. In *Cultures of Solitude: Loneliness, Limitation, and Liberation*. Peter Lang Frankfurt.
- Coplan, R. J., Hipson, W. E., Archbell, K. A., Ooi, L. L., Baldwin, D., and Bowker, J. C. (2019a). Seeking more solitude: Conceptualization, assessment, and implications of aloneliness. *Personality and Individual Differences*, 148:17–26.
- Coplan, R. J., Ooi, L. L., and Baldwin, D. (2019b). Does it matter when we want to be alone? Exploring developmental timing effects in the implications of unsociability. *New Ideas in Psychology*, 53:47–57.
- Deci, E. L. and Ryan, R. M. (2010). Self-determination. *The Corsini Encyclopedia of Psychology*, pages 1–2.
- Ekman, P. (1992). An argument for basic emotions. *Cognition and Emotion*, 6(3):169–200.
- Endo, K., Ando, S., Shimodera, S., Yamasaki, S., Usami, S., Okazaki, Y., Sasaki, T., Richards, M., Hatch, S., and Nishida, A. (2017). Preference for solitude, social isolation, suicidal ideation, and self-harm in adolescents. *Journal of Adolescent Health*, 61(2):187–191.
- Frijda, N. H. (1988). The laws of emotion. *American Psychologist*, 43(5):349.
- Gerst-Emerson, K. and Jayawardhana, J. (2015). Loneliness as a public health issue: the impact of loneliness on health care utilization among older adults. *American Journal of Public Health*, 105(5):1013–1019.
- Hawley, L. C. and Cacioppo, J. T. (2010). Loneliness matters: A theoretical and empirical review of consequences and mechanisms. *Annals of Behavioral Medicine*, 40(2):218–227.
- Hawley, L. C. and Capitano, 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.
- Kiritchenko, S. and Mohammad, S. M. (2016). Capturing reliable fine-grained sentiment associations by crowdsourcing and best–worst scaling. In *Proceedings of the Annual Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL)*, pages 811–817, San Diego, California.
- Kiritchenko, S., Zhu, X., and Mohammad, S. M. (2014). Sentiment analysis of short informal texts. *Journal of Artificial Intelligence Research*, 50:723–762.
- Knafo, D. (2012). Alone together: Solitude and the creative encounter in art and psychoanalysis. *Psychoanalytic Dialogues*, 22(1):54–71.
- Larson, R., Csikszentmihalyi, M., and Graef, R. (1982). Time alone in daily experience: Loneliness or renewal. *Loneliness: A sourcebook of current theory, research and therapy*, pages 40–53.
- Larson, R. W. (1990). The solitary side of life: An examination of the time people spend alone from childhood to old age. *Developmental Review*, 10(2):155–183.
- Long, C. R., Seburn, M., Averill, J. R., and More, T. A. (2003). Solitude experiences: Varieties, settings, and individual differences. *Personality and Social Psychology Bulletin*, 29(5):578–583.
- Louviere, J. J., Flynn, T. N., and Marley, A. A. J. (2015). *Best-Worst Scaling: Theory, Methods and Applications*. Cambridge University Press.
- Louviere, J. J. (1991). Best-worst scaling: A model for the largest difference judgments. Working Paper.
- Luhmann, M. and Hawkey, L. C. (2016). Age differences in loneliness from late adolescence to oldest old age. *Developmental Psychology*, 52(6):943.
- Luo, Y., Hawkey, L. C., Waite, L. J., and Cacioppo, J. T. (2012). Loneliness, health, and mortality in old age: A national longitudinal study. *Social Science & Medicine*, 74(6):907–914.
- Mohammad, S. M. and Turney, P. D. (2010). Emotions evoked by common words and phrases: Using mechanical turk to create an emotion lexicon. In *Proceedings of the NAACL HLT 2010 Workshop on Computational Approaches to Analysis and Generation of Emotion in Text*, pages 26–34.
- Mohammad, S. M. and Turney, P. D. (2013). Crowdsourcing a word-emotion association lexicon. *Computational Intelligence*, 29(3):436–465.
- Mohammad, S. M. (2018a). Obtaining reliable human ratings of valence, arousal, and dominance for 20,000 english words. In *Proceedings of The Annual Conference of the Association for Computational Linguistics (ACL)*, Melbourne, Australia.
- Mohammad, S. M. (2018b). Word affect intensities. In *Proceedings of the 11th Edition of the Language Resources and Evaluation Conference (LREC-2018)*, Miyazaki, Japan.
- Nguyen, T.-v. T., Ryan, R. M., and Deci, E. L. (2018). Solitude as an approach to affective self-regulation. *Personality and Social Psychology Bulletin*, 44(1):92–106.
- Osgood, C., G., S., and Tannenbaum, P. (1957). *The measurement of meaning*. University of Illinois Press.
- Park, G., Yaden, D. B., Schwartz, H. A., Kern, M. L., Eichstaedt, J. C., Kosinski, M., Stillwell, D., Ungar, L. H., and Seligman, M. E. (2016). Women are warmer but no less assertive than men: Gender and language on facebook. *PLoS One*, 11(5):e0155885.

- Plutchik, R. (1980). A general psychoevolutionary theory of emotion. *Emotion: Theory, research, and experience*, 1(3):3–33.
- Russell, J. A. (1980). A circumplex model of affect. *Journal of Personality and Social Psychology*, 39(6):1161.
- Russell, J. A. (2003). Core affect and the psychological construction of emotion. *Psychological Review*, 110(1):145.
- Schwartz, H. A., Eichstaedt, J. C., Kern, M. L., Dziurzynski, L., Ramones, S. M., Agrawal, M., Shah, A., Kosinski, M., Stillwell, D., Seligman, M. E. P., et al. (2013). Personality, gender, and age in the language of social media: The open-vocabulary approach. *PloS One*, 8(9):e73791.
- Seeman, T. E. (1996). Social ties and health: The benefits of social integration. *Annals of Epidemiology*, 6(5):442–451.