

WorryWords: Norms of Anxiety Association for over 44k English Words

Saif M. Mohammad

National Research Council Canada

saif.mohammad@nrc-cnrc.gc.ca

Abstract

Anxiety, the anticipatory unease about a potential negative outcome, is a common and beneficial human emotion. However, there is still much that is not known, such as how anxiety relates to our body and how it manifests in language. This is especially pertinent given the increasing impact of anxiety-related disorders. In this work, we introduce *WorryWords*, the first large-scale repository of manually derived word–anxiety associations for over 44,450 English words. We show that the anxiety associations are highly reliable. We use *WorryWords* to study the relationship between anxiety and other emotion constructs, as well as the rate at which children acquire anxiety words with age. Finally, we show that using *WorryWords* alone, one can accurately track the change of anxiety in streams of text. The lexicon enables a wide variety of anxiety-related research in psychology, NLP, public health, and social sciences. *WorryWords* (and its translations to over 100 languages) is freely available.

<http://saifmohammad.com/worrywords.html>

1 Introduction

Anxiety is the apprehensive uneasiness or nervousness due to some anticipated or potential negative outcome. Even though it is often considered to be an unwanted and negative feeling, psychologists have long posited its evolutionary benefits. As a species, we have benefited from feeling anxious: for example, building safe dwellings to avoid being attacked by wild animals and poisonous insects, being respectful and kind to others to avoid social exclusion, and so on (Bateson et al., 2011). Yet, there is much we do not know such as the mechanisms of how anxiety develops in our bodies and mind, how normal and clinical anxiety manifest in language, and effective mechanisms to cope with excessive anxiety. Further, the dramatic technologi-

cal advances over the last few decades have rapidly and substantially changed our environments to such an extent (for instance, the widespread social media usage and human–technology interactions) that there is now a considerable mismatch between the current environment and what our anxiety response slowly evolved to address over millennia (Blease, 2015; Salecl, 2004; Qureshi-Hurst, 2022). This has led some to argue that we are in an ‘age of anxiety’ (Parish, 2000; Blease, 2015).

Anxiety has always been a part of normal human life. Some might feel anxious occasionally and in particular circumstances (state anxiety), and some might have a more diffuse sense of general anxiety over long periods of time (trait anxiety). Some of the major sources of anxiety include tests, sports, technology, and social anxiety (Zeidner and Matthews, 2010). More recently researchers have also identified a growing amount of climate change anxiety (Clayton, 2020). Mental health practitioners clinically diagnose anxiety disorders when the sense of anxiety is so high or so frequent that it impacts the day-to-day functioning of an individual. In the US alone, 6.8 million adults (3.1% of the population) are believed to have Generalized Anxiety Disorder (GAD), 15 million adults (7.1%) have Social Anxiety Disorder (SAD), 7.7 million adults (3.6%) suffer from Post-traumatic Stress Disorder (PTSD), and everyone experiences stress from time to time.¹ Anxiety disorders are highly co-morbid with depression (Kalin, 2020), and together they constitute the *internalizing disorders*.

In addition to impacting well-being and happiness, anxiety (both normal and clinical) has been shown to impact cognitive outcomes such as learning, accurately performing tasks, productivity, as well as our physical bodies (Kim and Gorman, 2005). In contrast, calmness (being in a state of peace and tranquility) has been shown to improve health and cognitive outcomes (Siddaway et al.,

¹<https://adaa.org/understanding-anxiety/facts-statistics>

2018). Siddaway et al. (2018) also show the importance of conceptualizing anxiety as a continuum from high calmness to high anxiety; and that both promoting calmness and decreasing excessive anxiety are important for overall health and well-being. Thus, anxiety and calmness are active areas of research in Psychology, Medicine, Public Health, Social Science, Economics, and even the Humanities.

As is the case with many emotions, people express anxiety through language (consciously and unconsciously). In this work, we study how anxiety and calmness are associated with word meaning. We asked over a thousand respondents to rate 44,450 English words on a scale from -3 (associated with high calmness) to $+3$ (associated with high anxiety). We aggregated the responses to create a word–anxiety association lexicon that has a real-valued anxiety score for each of the words. We refer to this resource by the name *WorryWords*. We show that the scores in *WorryWords* are highly reliable (a split-half reliability score of 0.82, which is in-line with scores for existing resources for other emotions). We use *WorryWords* to study how anxiety is related to other emotion constructs (valence, arousal, dominance, fear, anger, sadness, etc.). We show the rate at which anxiety and calmness words are acquired with age. Finally, we show that *WorryWords* can be used with simple word-counting methods to accurately track anxiety in text streams.

We make *WorryWords*, and its translations in over 100 languages, freely available for research through the project home page.² Some areas where we envisage its use include:

1. Understanding anxiety and the underlying mechanisms; how anxiety relates to other emotions; how it relates to our body; how anxiety changes with age, socio-economic status, weather, green spaces, etc.
2. Determining how anxiety manifests in language; how language shapes anxiety; how culture shapes the language of anxiety; etc.
3. Tracking the degree of anxiety towards targets of interest such as climate change, government policies, biological vectors, etc.
4. Identifying effective coping mechanisms and clinical interventions to manage anxiety.
5. Developing automatic systems for detecting anxiety; developing chat systems that are sensitive to nuances and diverse expressions of anxiety by people from various demographics.

6. Studying anxiety and uneasiness in story telling; its relationship with central elements of narratology such as conflict and resilience.
7. Studying how anxiety impacts social behaviour in physical and virtual environments.

We also created an extensive list of ethical considerations (Section 9) that must be taken into consideration when using *WorryWords*.

2 Related Work

It is generally agreed that the three primary dimensions of affect (our sense of feeling) are valence (positiveness–negativeness / pleasure–displeasure), arousal (active–passive), and dominance (dominant–submissive) (Russell, 1980; Osgood et al., 1957). There is also work showing that certain categorical emotions such as joy, anger, and fear are particularly important (Ekman, 1992; Plutchik, 1980). More recent work shows that we construct and label these emotions (a cognitive process) in our minds from various interoceptive signals (internal signals from our body to the brain) (Barrett, 2017).

Anxiety is often considered a low-valence, high-arousal, and low-dominance emotion, but there is little work empirically exploring the degree to which it is correlated with these affective dimensions. Anxiety is also often compared to fear, but unlike fear which generally tends to be a reaction to a threat from the *present* (something that has already manifested), anxiety is often associated with the *possible* negative outcome of a *future* event (Bateson et al., 2011). Anxiety and stress are sometimes used synonymously and they are generally agreed to be very close emotions (both impact mind and body through excessive worrying, loss of sleep, etc.), although stress is often associated with some physical stimulus and stress can be one of the causes of anxiety.

Manually created word–emotion association lexicons exist for valence, arousal, and dominance (Warriner et al. (2013) and Mohammad (2018a) for English; Moors et al. (2013) for Dutch, and Võ et al. (2009) for German). The largest among these is the NRC VAD lexicon (Mohammad, 2018a, 2024a): version 1 has entries for over 20,000 English words, and version 2 for over 44,000 unigrams and 10,000 bigrams (two-word sequences). The NRC Emotion Lexicon (Mohammad and Turney, 2010) is a popular lexicon for categorical emotions; it has entries for over 14,000 English words.

²<http://saifmohammad.com/worrywords.html>

We do not know of existing large lexicons for anxiety; however, there exist some datasets of social media *posts* manually annotated with anxiety or stress labels. [Rastogi et al. \(2022\)](#) collected and annotated about 5,488 Reddit posts with the labels 0 (no stress) and 1 (stress). [Mitrović et al. \(2024\)](#) collected a series of questions and answers from Quora and Reddit, where the question contained the term ‘anxiety’. [Turcan and McKeown \(2019\)](#) created a dataset of 4000 5-sentence segments from Reddit annotated for presence or absence of stress. These social media datasets were compiled primarily to train and test automatic systems that attempt to detect anxiety from text. For details on automatic anxiety detection, we refer the interested readers to [Mitrović et al. \(2024\)](#); [Shen and Rudzicz \(2017\)](#); [Rastogi et al. \(2022\)](#); [Winata et al. \(2018\)](#). More recent work interviewed volunteers and analyzed their responses to identify language features such as the use of personal pronouns, negations, and emotion words with anxiety ([Stade et al., 2023](#)).

As part of the WorryWords Project, we created a large word–anxiety association lexicon with a focus on reliable annotations.³ We used existing emotion lexicons to examine how anxiety is related to other emotions. We also made use of the [Rastogi et al. \(2022\)](#) dataset to show the usefulness of WorryWords in tracking anxiety in text streams.

3 Obtaining Human Ratings of Anxiety

The keys steps in creating the anxiety dataset were:

1. selecting the terms to be annotated
2. developing the questionnaire
3. developing measures for quality control (QC)
4. annotating terms on a crowdsourcing platform
5. discarding data from outlier annotators (QC)
6. aggregating data from multiple annotators to determine the anxiety association scores

We describe each of these steps below.

1. Term Selection. To allow for the study of anxiety with respect to other affectual and linguistically interesting properties, we included terms from: the NRC Emotion Lexicon ([Mohammad and Turney, 2013, 2010](#)), the NRC VAD Lexicon ([Mohammad, 2018a](#)), and the terms from the Prevalence dataset ([Brysbaert et al., 2019](#)) (specifically, words marked as being known to over 70% of the respondents). Summary details of each of these is provided in the Appendix. The union of the above sets resulted in

44,450 terms that were then annotated for association with anxiety and calmness.

2. Anxiety Questionnaire. The questionnaires used to annotate the data were developed across two rounds of pilot annotations. Detailed directions, including notes directing them to consider predominant word sense (in case the word is ambiguous) and example questions (with suitable responses) were provided. (See Appendix.) The primary instruction and the question presented to annotators is shown below.

Consider anxiety to be a broad category that includes:

jittery, antsy, insecure, nervous, unease, tense, worried, unnerving, nerve-racking, apprehensive, fretful, troubled, self-doubting, discontented, concerned, and keyed up

Consider calmness to be a broad category that includes:

calm, relaxed, comforted, serene, at ease, self-assured, carefree, composed, collected, untroubled, peaceful, contented, unconcerned, indifferent, and uninvolved

Select the options that most English speakers will agree with.

Q1. <term> is often associated with feeling:

- | | |
|--|---------------------|
| 3: very anxious | -1: slightly calm |
| 2: moderately anxious | -2: moderately calm |
| 1: slightly anxious | -3: very calm |
| 0: not associated with feeling anxious or calm | |
-

3. Quality Control Measures. About 2% of the data was annotated beforehand by the authors and interspersed with the rest. These questions are referred to as *gold* (aka *control*) questions.

Half of the gold questions were used to provide immediate feedback to the annotator (in the form of a pop-up on the screen) in case they mark them incorrectly. We refer to these as *popup gold*. This helps prevent the situation where one annotates a large number of instances without realizing that they are doing so incorrectly. It is possible, that some annotators share answers to gold questions with each other (despite this being against the terms of annotation). Thus, the other half of the gold questions were also separately used to track how well an annotator was doing the task, but for these gold questions no popup was displayed in case of errors. We refer to these as *no-popup gold*.

4. Crowdsourcing. We setup the anxiety annotation task on the crowdsourcing platform, *Mechanical Turk*. In the task settings, we specified that we needed annotations from ten people for each word. We obtained annotations from native speakers of English residing around the world. Annotators were free to provide responses to as many terms as they wished. The annotation task was approved by our institution’s review board.

³In follow on work, we also compiled anxiety association norms for over 10K English phrases ([Mohammad, 2024b](#)).

Dataset	#words	Annotators	#Annotators	#Annotations	MAI	SHR (ρ)	SHR (r)
unigrams	44,450	US, India, UK, Canada	1,020	375,796	8.5	0.82	0.89

Table 1: A summary of the WorryWords annotations. MAI = mean annotations per word. SHR, measured through both Spearman rank and Pearson’s correlations (last two columns), indicate high reliability.

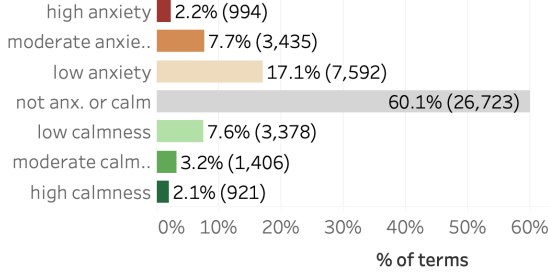


Figure 1: Distribution of terms in WorryWords: percentage of terms associated with each class. (The total number is shown in parenthesis.)

Demographics: About 87% of the respondents who annotated the words live in USA. The rest were from India, United Kingdom, and Canada. The average age of the respondents was 38.9 years. Among those that disclosed their gender, about 54% were female and about 45% were male.⁴

5. Filtering. If an annotator’s accuracy on the gold questions (popup or non-popup) fell below 80%, then they were refused further annotation, and all of their annotations were discarded (despite being paid for). See Table 1 for summary statistics.

6. Aggregation. Every response was mapped to an integer from -3 (very calm) to 3 (very anxious). The final anxiety score for each term is simply the average score it received from each of the annotators. We also created a categorical version of the lexicon by labeling all words that got a score ≥ 2.5 as associated with *high anxiety*, ≥ 1.5 and < 2.5 as *moderate anxiety*, ≥ 0.5 and < 1.5 as *slight anxiety*, > -0.5 and < 0.5 as *neither anxiety nor calmness*, and so on. We refer to the list of words along with their real-valued final scores and categorical labels as the *WorryWords Lexicon*.

Table 6 in the Appendix shows example entries from the lexicon. Figure 1 shows the distribution of the different classes. As expected, most words are associated with neither anxiety nor calmness ($\sim 60\%$), but it is worth noting that about 27% of the words are associated with anxiety (to some degree) and about 13% of the words are associated with calmness (to some degree).

⁴Respondents were shown optional text boxes to disclose their demographic information as they choose; especially important for social constructs such as gender, in order to give agency to the respondents and to avoid binary language.

4 Reliability of the Annotations

A useful measure of quality is the reproducibility of the end result—repeated independent manual annotations from multiple respondents should result in similar scores. To assess this reproducibility, we calculate average *split-half reliability* (SHR) over 1000 trials.⁵ All annotations for an item are randomly split into two halves. Two separate sets of scores are aggregated, just as described in Section 3 (bullet 6), from the two halves. Then we determine how close the two sets of scores are (using a metric of correlation). This is repeated 1000 times and the correlations are averaged. The last two columns in Table 1 show the results (split half-reliabilities). Scores of 0.82 (Spearman rank correlation) and 0.89 (Pearson correlation) indicate high reliability of the real-valued scores obtained from the annotations. (For reference, if the annotations were random, then repeat annotations would have led to an SHR of 0. Perfectly consistent repeated annotations lead to an SHR of 1.)

If the real-valued scores are converted to discrete categories (for example, into 3, 7, or 10, classes), then we can make use of another measure of reliability that factors in the classes. For example, if we just want three classes: *associated with anxiety*, *associated with calmness*, and *neither*. Then one could simply split the -3 to 3 range into three equal partitions of -3 to -1 (calmness), -1 to 1 (neither), and 1 to 3 (anxiousness). To determine reliability of the discrete labels obtained from annotations, we first randomly split all annotations into two halves (just as before). Now, if scores for a term x obtained from the two halves place x in different classes (e.g., if the two scores are 0.36 and 1.72), then this contributes to a lower overall reliability. In contrast, even if the scores for x from the two halves vary a bit, but both scores place x in the same class (e.g., if the two scores are 1.82 and 1.79) then this contributes to a higher overall reliability. Thus we can compute reliability from the extent to which the two halves of annotations result in the same discrete class. We will refer to this as the *Split-Half Class Match Percentage* (SHCMP).

⁵SHR is a common way to determine reliability of responses to generate scores on an ordinal scale (Weir, 2005).

#Bins	Bin size	Random	WorryWords	
		SHCMP	SHCMP	SHCloseP
10	0.60	10.0	0.62	84.9
7	0.86	14.3	0.71	91.6
5	1.20	20.0	0.78	96.8
4	1.50	25.0	0.79	98.4
3	2.00	33.0	0.87	99.4
2	3.00	50.0	0.89	100.0

Table 2: Split-Half Class Match Percentage (SHCMP) and Split-Half Closeness Percentage (SHCloseP) obtained from the anxiety annotations when splitting the -3 to 3 score range into equal-sized bins (partitions) pertaining to ordinal classes. SD: standard deviation.

Note, however, that using discrete classes means that many items that have similar real-valued scores from two independent sets of annotations may be on different sides of a class boundary. Thus even if the original annotations are reliable, using discrete classes adds some unreliability inherently. To get a sense of this, we also calculate the degree to which the split-half scores are close to each other (without regard to class boundaries): e.g., the percentage of times two independent sets of annotations for an item lead to scores that are closer than a pre-chosen threshold (0.1, 0.25, etc.). We will call this the *Split-Half Closeness Percentage (SHCloseP)*.

Table 2 shows the SHCMP and SHCloseP for the WorryWords annotations for various levels of granularity (#classes). The second column shows the class size—obtained by dividing the -3 to 3 range equally by #classes.⁶ The third column shows the expected SHCMP when terms are given scores at random. The fourth column shows SHCMP when using the actual WorryWords annotations. The fifth column shows SHCloseP when using different bin sizes as threshold.

Observe that the SHCMP scores are markedly higher than what one would get with just random annotations. Further, SHCMP scores are fairly high in case of a small number of discrete classes (bins), but drops off markedly as the number of classes is increased. Note the considerably higher SHCloseP scores which indicate that much of the unreliability in discrete classes is due to items falling on either side of class boundaries, and that when there is a class mismatch the involved classes are close to each other (smaller error) as opposed to distant classes (bigger error).

⁶SHCMP can also be determined from custom ranges that are not of the same size. Here we show SHCMP with equal size bins for simplicity.

5 WorryWords: Research & Applications

WorryWords has a wide variety of uses and applications as listed in the Introduction. We explore:

- Two diverse examples where lexicographic analysis of the words helps answer research questions about anxiety (Section 6).
- An example (with evaluation) of how the WorryWords can be used to accurately track anxiety in text streams (Section 7). This has both commercial and research applications.

6 Using WorryWords to Answer Research Questions about Anxiety

6.1 From a lexical standpoint, how is anxiety related to other emotions?

An important question in understanding any emotion is how it relates to other emotions. With the creation of the first large-scale lexicon of word–anxiety associations, we can now study how these norms correlate with other key emotion constructs such as valence (V), arousal (A), and dominance (D), as well as categorical emotions such as sadness, fear, and anger. Since anxiety is commonly understood to be a negative (low-valence emotion), we expect it to have a negative correlation with valence (negative situations are often associated with anxiety). However, we have no prior evidence to indicate how strong this negative correlation is expected to be; especially since, there can be many term pairs w_1 and w_2 such that even though w_1 has lower valence than w_2 , w_2 is associated with a higher degree of anxiety than w_1 . For example, *hopeless* and *frenzy*, *lethargic* and *public speaking*, *weak* and *test*, etc.

Similarly, we expect anxiety to have a positive correlation with arousal (since high anxiety terms are also expected to be high arousal terms), and a negative correlation with dominance (since higher anxiety tends to co-occur with situations of loss of control). Yet, again, it is unclear how strong these correlations truly are across a vast vocabulary of words. Understanding these relationships has implications in psycholinguistics and lexical semantics: e.g., whether anxiety moderates word recognition even when controlling for valence. Further, in the context of lexicon-creation: If the magnitudes of these correlations are very high, then one can simply take an existing VAD lexicon, and keep only the low valence, high arousal, and/or low dominance terms to create an anxiety lexicon.

	anxiety	valence	arousal	domin.
anxiety	1.000	-0.471	0.237	-0.218
valence	-0.471	1.000	-0.014	0.568
arousal	0.237	-0.014	1.000	0.286
domin.	-0.218	0.568	0.286	1.000

Table 3: Spearman rank correlations of association scores between various emotion pairs, for 44k terms.

VAD: The first row (and also the first column) of Table 3 shows the correlations of anxiety association scores of terms in WorryWords with the valence, arousal, and dominance scores of terms in the NRC VAD Lexicon v2 (Mohammad, 2024a). Observe that anxiety has a moderate negative correlation with valence, a mild correlation with arousal, and a mild negative correlation with dominance. For additional context, the bottom 3 rows (and last 3 columns) of Table 3 show the pair-wise correlations between valence, arousal, and dominance. It shows for example, that valence and dominance tend to be more correlated than any of the other pairs (a result also seen in other VAD lexicons such as the Warriner Lexicon (2013)); and in terms of the magnitude, the degree of (negative) correlation between valence and anxiety is lower than that of valence and dominance.

We also created an interactive visualization that shows the WorryWords terms on valence–arousal and valence–dominance spaces (Figure 8 in the Appendix shows a screenshot). (Hovering over individual points shows the corresponding term as well as anxiety, V, A, and D scores.) We observe that, as expected, a majority of the anxiety words (score > 0.5) occur in the top-left quadrant of the V–A space: low V and high A. Nonetheless, there are still non-negligible number of terms that occur in the top-right quadrant (such as *adrenalin*, *revolutionist*, and *whizzing*) and in the bottom-left quadrant (such as *joylessness*, *disheartenment*, and *disreputableness*). Comparatively much fewer anxiety terms occur in the bottom right (high V and low A) quadrants. Examples include *seatbelt* and *post-operative*. In contrast, the majority of the calmness terms (score < -0.5) occur in the high V quadrants. Examples, of low V terms associated with calmness include *pity* and *aloof*. The V–D plots show that anxiety terms tend to be concentrated in the bottom-left quadrant, whereas calmness terms are most commonly found in the top-right quadrant. Anxiety terms associated with high D include terms such as *skyrocket*, *canny*, *gunfight*, and *tyrant*; anxiety terms associated with low D include *persecuted* and *gutlessness*.

Categorical Emotions: Table 4 shows the correlations of anxiety scores of terms in WorryWords with the anger, anticipation, disgust, fear, joy, sadness, surprise, and trust scores of terms in the NRC Emotion Intensity Lexicon v1 (Mohammad, 2018b). Observe that anxiety seems to essentially have no correlation with anticipation, joy, and surprise (correlation magnitudes less than 0.2); and low correlation with the negative emotions: anger, disgust, fear, and sadness. Among these emotions, anxiety has the highest correlation with fear.

Discussion: Overall, we see that while anxiety associations follow some expected trends such as pertaining to low valence, the magnitudes of correlations are moderate with valence and low with arousal and dominance. Thus anxiety presents itself in many different ways and cannot be characterized simply with VAD. The fact that anxiety has the highest correlation with fear (among the four negative categorical emotions) is in line with the belief that fear is a key aspect of anxiety (Rosen and Schulkin, 1998; Shin and Liberzon, 2010).

6.2 At what rate do children acquire words associated with anxiety? How does this change with the Age of Acquisition?

Feeling anxious is a normal part of childhood and growing up (Beesdo et al., 2009). The Cleveland Clinic describes how very young children (6-month to 3-year olds) can have separation anxiety from parents and caregivers; pre-schoolers often develop phobias towards specific things such as insect, darkness, thunder, heights, etc.; and school age brings many sources of anxiety such as going to a new school, exams, social interactions, and managing school work.⁷ Over the course of development, the way children express, experience, and communicate their emotions, including anxiety, changes (Bailen et al., 2019; Thompson, 1991). Understanding these changes is instrumental in promoting healthy socio-emotional functioning across all stages of development. Thus, the associations of emotions with different ages is a key aspect studied by researchers in child development.

With WorryWords we can now look specifically at the relationship between anxiety-associated words and age. Specifically, in our work, we examine the correlation between the age at which words are acquired by children in their vocabulary

⁷<https://my.clevelandclinic.org/health/diseases/anxiety-in-children>

	anxiety	anger	anticipn.	disgust	fear	joy	sadness	surprise	trust
anxiety	1.000	0.336	-0.011	0.279	0.357	-0.132	0.272	0.082	-0.137

Table 4: Spearman rank correlation between the word–emotion association scores for anxiety and 8 other emotions.




Figure 2: Percent of WorryWords terms by AoA.

(the age they learn the meanings of words) and the anxiety associated with those words.

We do so by examining the words with entries in both WorryWords and [Kuperman et al. \(2012\)](#) age of acquisition dataset (which includes the age at which $\sim 30K$ English words are commonly acquired by children). Figure 2 shows the percentage of words in WorryWords that are acquired at various ages. Subsets of words corresponding to the three anxiety classes are shown in shades of orange, corresponding to the three calmness classes are shown in shades of green, and the amount of words neither associated with anxiousness nor calmness are shown in grey.

Observe that the percentage of words acquired increases steadily with age, peaking at 11, and then decreasing rapidly after that. The percentage of anxiety and calmness words also largely follows the same pattern, but some notable additional observations can be made. At a very early age (1 to 4 years) the percentage of calmness words acquired is higher than the percentage of anxiety words. However, from age 5 onwards, children learn more anxiety words than calmness words.

Figure 3 shows the percentage of words acquired each year for each of the anxiety and calmness classes; i.e., after disregarding the *neither anxious nor calm* words, what percentage of words acquired in a year correspond to each of the six other classes. (For every x-axis value, the six percentages sum up to 100%.) Observe that from age 3 onwards, the percentage of anxiety words increases steadily with age (for all three anxiety classes), whereas, the percentage of calmness words decreases with age (for all three calmness classes).




Figure 3: For all of the non-neutral WorryWords terms acquired at age x : the percentage corresponding to each of the non-neutral anxiety–calmness bins. (The six percentages for each age sum up to 100%.)

Discussion: The notable trends in the relative rates of acquisition of anxiety and calmness terms with age have implications in developmental psychology and evolutionary linguistics ([Price, 2013](#); [Breggin, 2015](#)). They may also be relevant in the understanding of and coping with childhood anxieties ([Wright et al., 2010](#); [Legerstee et al., 2010](#)). We hope that WorryWords will spur further research in these areas.

7 Capturing Anxiety Arcs in Text Streams using WorryWords

A common use of emotion lexicons is to determine emotion arcs and trends in how emotions change in a stream of textual data. For example, determining changes in the degree of anger towards a product in social media over time ([Breeze, 2020](#)), tracking happiness and well-being of different cities and counties through geo-located posts on X (formerly Twitter) in public health research ([Frank et al., 2013](#)), generating emotion arcs of stories in digital humanities ([Mohammad, 2011](#); [Ohman et al., 2024](#)), etc. The lexicons may be used in combination with supervised systems and large language models (where such resources are available), or they may be used on their own where interpretability, low-carbon foot print, and simpler systems are

valued — which is often the case in psychology and public health research. Lexicon-only methods are often markedly less accurate at determining emotions of *individual sentences*. However, [Teodorescu and Mohammad \(2023\)](#) have shown that on the task of generating emotion arcs (where emotion scores of hundreds, if not thousands, of individual sentences are averaged to determine points on the emotion arc) lexicon-only methods are competitive with more sophisticated neural and large language models.⁸ In this section we explore:

1. How accurate are anxiety arcs generated using WorryWords alone?
2. Given the close relationship between anxiety and other emotion categories such as fear, anger, and low valence, how good are lexicons for other emotions at capturing anxiety arcs?

Together these results tell us about the usefulness of WorryWords in the task of generating emotion arcs, as well as how anxiety-associated and other-emotion-associated words present themselves in sentences that have been marked for anxiety.

7.1 Creating the Anxiety Arcs

First we quickly summarize how emotion arcs are created. The user (anyone interested in tracking anxiety for a text stream of interest) provides the:

- Text Stream: A series of instances (sentences, posts, etc.) in order of interest (e.g., posts on X that mention *#climatechange* ordered by timestamps, sentences from a novel).
- Bins: A pre-chosen grouping (day, month, year, k sentences, etc. — based on granularity of interest)⁹. An additional parameter relevant when the bin is specified in terms of number of sentences (say, k) is the step size. For example, a step size of s instances means that after including sentences 1 to k in the first bin, the next bin includes instances $1 + s$ to $1 + s + k$, then $1 + 2s$ to $1 + 2s + k$, and so on. Commonly used step sizes are 1 (highly overlapping bins — leading to smoother arcs) and k (non-overlapping bins).

⁸Even though an individual lexicon entry only provides the emotion association with the predominant sense (and does not take context into account), by tracking the emotions associated with hundreds of co-occurring words in a text stream we make use of context. If for example, an anger-inducing event occurs, then the corresponding bin is likely include many words that will have an association with anger.

⁹For e.g., anxiety in tweets about climate change where one wants to track changes month by month, each bin includes the tweets posted in the corresponding *month*.

We generate the predicted arc for this text stream by simply taking the average anxiety association scores of words (taken from WorryWords) for each bin. Terms not found in the lexicon are disregarded when determining these scores. The gold arc can be created by taking the averages of the manually labeled anxiety scores (labels) of each of the posts in the bins. The closer the predicted arc to the gold arc, the more accurate it is.

However, there exists no cohesive text streams dataset with annotations for anxiety. That said, there exist datasets where independent sentences are annotated for anxiety, such as the stress dataset by [Rastogi et al. \(2022\)](#). Since this dataset is not associated with a particular event where there might be ups and downs of emotions, a completely random ordering of the data or using timestamps to order the data into a text stream, produces of a flat line anxiety arc. Thus, we created new text streams by sampling posts from the [Rastogi et al. \(2022\)](#) dataset (5,488 Reddit posts), with replacement, such that the gold anxiety arc had portions with various random amplitudes and slopes (spikiness).¹⁰ (Note that sampling is a way to create a new ordering of the posts.) We continued sampling until we created text streams that were 10,000 sentences long. We then repeated the procedure 1000 times to create 1000 different text streams (and corresponding gold arcs). These text streams can now be used to evaluate automatic methods of determining anxiety arcs. The yellow line in Figure 4 (a) shows the beginning portion of one of the gold anxiety arcs (the 1000th arc) for bin size of 50.

7.2 Results

The green line in Figure 4 (a) shows the beginning portion of the predicted anxiety arc when using the lexicon-only method, word associations from WorryWords, and a rolling window with bin size 50.¹¹ Observe that the green line traces the gold line quite closely. There are also occasions where it overshoots or fails to reach the gold arc. The column on the right in the figure lists the average Spearman correlation between the predicted and

¹⁰Sampling without replacement led to the same results, but sampling with replacement allowed reuse of the limited labeled data to produce longer waves. The random arc sampling algorithm is described in the Appendix.

¹¹Note that bin size is a user specified parameter. Bin size of 50 posts is relatively small and often only used when data is sparse, e.g., tracking character arcs in novels. Results with bigger bin sizes are often better (as more data is used to infer change of emotion across bins) — as shown ahead.



Figure 4: Predicted (green) anxiety arcs when applying four different types of association scores to one of the test data text streams: a. anxiety association from WorryWords, b. valence (from NRC VAD), c. arousal (from NRC VAD), and dominance (from NRC VAD). The gold (yellow) arc pertaining to the same text stream is repeated in all four figures for reference. Bin size used to create the arcs: 50 posts. On the right is the average correlation between 1000 pairs of gold and predicted arcs for each case.

Window Size	Spearman	RMSE
1	0.554	1.011
10	0.767	0.517
50	0.903	0.204
100	0.932	0.136
200	0.946	0.099
500	0.955	0.075
1000	0.955	0.073

Table 5: Avg. Spearman corn. and RMSE of 1K pairs of predicted and gold anxiety arcs for different bin sizes.

gold arcs for all 1000 text streams (for the full length of the text streams: 10,000 posts). We note that the arc generated using the anxiety lexicon has a high correlation with the gold arc (greater than 0.9) – sufficient for most applications.

Figure 4 (b), (c), and (d) show the same gold anxiety arc but show predicted arcs using valence, arousal, and dominance scores, respectively (taken from the NRC VAD Lexicon). These figures act as control arcs (for comparison with the WorryWords arc). The column on the right shows the average correlations for the 1000 pairs of gold and predicted arcs. We see that the gold arc is highly negatively correlated with valence (-0.85). Just in terms of the magnitude of correlation, arousal is closer to anxiety than dominance with anxiety.

Categorical Emotions: Figure 9 in the Appendix shows similar gold and predicted arcs (and overall average correlations) for the gold anxiety arc and predicted arcs generated using lexicons for eight other categorical emotions. The magnitudes of cor-

relation obtained using lexicons for these emotions is lower than that obtained using WorryWords.

Window Size Effect: We repeated the experiments described above with many window sizes and found similar overall trends. However, using larger window sizes generally led to higher correlations. Table 5 shows these results. In addition to Spearman rank correlation the table also shows another metric of arc–arc closeness: root mean square error (RMSE). The closer the two arcs, the smaller the RMSE. Observe that there is a plateauing effect in closeness after a window size of a few hundred posts (around $\rho = 0.955$ and $RMSE = 0.075$).

8 Conclusion

We created WorryWords — a large lexicon of word–anxiety association scores with the help of responses to a simple questionnaire, by over 1000 respondents (mainly from the US, but also from India, UK, and Canada). We showed that the responses led to highly reliable scores. We compared how anxiety association scores correlate with the scores for various other emotions. We determined the rate at which anxiety words are acquired with age; specifically, how from ages 1 to 4 children learn more calmness words than anxiety words, and that trend is reversed age 5 onwards. Finally, we showed how WorryWords can be used to track anxiety in large text streams. We make WorryWords freely available to enable a wide variety of anxiety associated research.

9 Limitations

For limitations associated with the dataset, please refer to the Ethics Statement (§10).

The lexicographic experiments in Section 6 make use of existing emotion and Age of Acquisition datasets. The labels for these datasets reflect their annotators, who were mainly US and European English speakers. Thus further experiments with datasets pertaining to other locations and demographics are needed to determine how widely applicable or universal some of these conclusions are. The conclusions from our work should largely be seen as applicable to the demographics of its annotators (and not beyond).

The emotion arc experiments were done on an existing stress-labeled dataset of Reddit posts. However, the dataset does not correspond to a natural textstream such as tweets about a new government policy or utterances of a character in a novel. (No such dataset has human annotations for anxiety.) However, by creating text streams corresponding to a thousand different emotion arcs from an existing Reddit dataset, we created a robust anxiety arc gold dataset.

10 Ethics Statement

Emotions are complex and nuanced mental states. Additionally, each individual expresses emotions, including anxiety, differently through language, which results in large amounts of variation. See [Mohammad \(2023\)](#) for a discussion of good practices and ethical considerations when using emotion lexicons. See [Mohammad \(2022\)](#) for a broader discussion of ethical considerations relevant to automatic emotion recognition. We discuss below notable ethical considerations when computationally analyzing anxiety.

Importantly, WorryWords should not be used as a standalone tool for detecting anxiety disorders. We hope future research, led by clinicians and medical practitioners, will assess the usefulness of WorryWords in mental health applications, in combination with various other sources of information.

The crowd-sourced task was approved by our Institutional Research Ethics Board. Our annotation process stored no information about annotator identity and as such there is no privacy risk to them. The individual words selected did not pose any risks beyond occasionally reading text on the internet. The annotators were free to do as many

word annotations as they wished. The instructions included a brief description of the purpose of the task (Figure 5).

Any lexical dataset of emotion-association norms entails several ethical considerations.

- *Coverage:* We sampled a large number of English words from many sources (which themselves sample from many sources). Yet, it is possible that the words included do not cover all domains, genres, and people of different locations, socio-economic strata, etc. equally. It likely includes more of the vocabulary common in the United States with a socio-economic and educational backgrounds that allow for technology access.
- *Word Senses and Dominant Sense Priors:* Words when used in different senses and contexts may be associated with different degrees of anxiety. The entries in WorryWords are indicative of the anxiety associated with the predominant senses of the words. This is usually not too problematic because most words have a highly dominant main sense (which occurs much more frequently than the other senses). In specialized domains, some terms might have a different dominant sense than in general usage. Entries in the lexicon for such terms should be appropriately updated or removed.
- *Not Immutable:* The anxiety scores do not indicate an inherent unchangeable attribute. The associations can change with time (e.g., the increase in anxiety associated with *coughing* during a world-wide pandemic), but the dataset entries are largely fixed. They pertain to the time they are created.
- *Socio-Cultural Biases:* The annotations for anxiety capture various human biases. These biases may be systematically different for different socio-cultural groups. Our data was annotated by mostly US and Indian English speakers, but even within the US and India there are many diverse socio-cultural groups. Notably, crowd annotators on Amazon Mechanical Turk do not reflect populations at large. In the US for example, they tend to skew towards male, white, and younger people. However, compared to studies that involve just a handful of annotators, crowd annotations benefit from drawing on hundreds and

thousands of annotators (such as this work).

- *Inappropriate Biases*: Our biases impact how we view the world, and some of the biases of an individual may be inappropriate. For example, one may have race or gender-related biases that may percolate subtly into one's notions of anxiety associated with words. Our dataset curation was careful to avoid words from problematic sources. We also ask people to annotate terms based on what most English speakers think (as opposed to what they themselves think). This helps to some extent, but it can still capture historical anxiety associations with certain identity groups. This can be useful for some socio-cultural studies; but we also caution that anxiety association with identity groups be carefully contextualized to avoid false conclusions.
- *Perceptions (not "right" or "correct" labels)*: Our goal here was to identify common perceptions of anxiety association. These are not meant to be "correct" or "right" answers, but rather what the majority of the annotators believe based on their intuitions of the English language.
- It is more appropriate to make claims about anxiety word usage rather than anxiety of the speakers. For example, '*the use of anxiety words grew by 20%*' rather than '*anxiety grew by 20%*'. A marked increase in anxiety words is likely an indication that anxiety increased, but there is no evidence that anxiety increased by 20%.
- Comparative analyses can be much more useful than stand-alone analyses. Often, anxiety word counts on their own are not useful. For example, '*the use of anxiety words grew by 20% when compared to [data from last year, data from a different person, etc.]*' is more useful than saying '*on average, 5 anxiety words were used in every 100 words*'.
- Inferences drawn from larger amounts of text are often more reliable than those drawn from small amounts of text. For example, '*the use of anxiety words grew by 20%*' is informative when determined from hundreds, thousands, tens of thousands, or more instances. Do not draw inferences about a single sentence or utterance from the anxiety associations of its constituent words.

We recommend careful reflection of ethical considerations relevant for the specific context of deployment when using WorryWords.

Acknowledgments

Many thanks to Sophie Wu and Tara Small for helpful discussions.

References

- Lisa Feldman Barrett. 2017. The theory of constructed emotion: an active inference account of interoception and categorization. *Social cognitive and affective neuroscience*, 12(1):1–23.
- Melissa Bateson, Ben Brilot, and Daniel Nettle. 2011. Anxiety: an evolutionary approach. *The Canadian Journal of Psychiatry*, 56(12):707–715.
- Katja Beesdo, Susanne Knappe, and Daniel S Pine. 2009. Anxiety and anxiety disorders in children and adolescents: developmental issues and implications for dsm-v. *Psychiatric Clinics*, 32(3):483–524.
- CR Blease. 2015. Too many 'friends,' too few 'likes'? evolutionary psychology and 'facebook depression'. *Review of General Psychology*, 19(1):1–13.
- Ruth Breeze. 2020. Angry tweets: A corpus-assisted study of anger in populist political discourse. *Journal of Language Aggression and Conflict*, 8(1):118–145.
- Peter R Breggin. 2015. The biological evolution of guilt, shame and anxiety: A new theory of negative legacy emotions. *Medical hypotheses*, 85(1):17–24.
- Marc Brysbaert, Paweł Mandera, Samantha F McCormick, and Emmanuel Keuleers. 2019. Word prevalence norms for 62,000 english lemmas. *Behavior research methods*, 51:467–479.
- Susan Clayton. 2020. Climate anxiety: Psychological responses to climate change. *Journal of anxiety disorders*, 74:102263.
- Paul Ekman. 1992. An argument for basic emotions. *Cognition and Emotion*, 6(3):169–200.
- Morgan R Frank, Lewis Mitchell, Peter Sheridan Dodds, and Christopher M Danforth. 2013. Happiness and the patterns of life: A study of geolocated tweets. *Scientific reports*, 3(1):2625.
- Ned H Kalin. 2020. The critical relationship between anxiety and depression.
- Jean Kim and Jack Gorman. 2005. The psychobiology of anxiety. *Clinical Neuroscience Research*, 4(5-6):335–347.
- Victor Kuperman, Hans Stadthagen-Gonzalez, and Marc Brysbaert. 2012. Age-of-acquisition ratings for 30,000 english words. *Behavior research methods*, 44:978–990.

- Jeroen S Legerstee, Nadia Garnefski, Francine C Jellesma, Frank C Verhulst, and Elisabeth MWJ Utens. 2010. Cognitive coping and childhood anxiety disorders. *European child & adolescent psychiatry*, 19:143–150.
- Sandra Mitrović, Oscar William Lithgow-Serrano, and Carlo Schillaci. 2024. [Comparing panic and anxiety on a dataset collected from social media](#). In *Proceedings of the 9th Workshop on Computational Linguistics and Clinical Psychology (CLPsych 2024)*, pages 153–165, St. Julians, Malta. Association for Computational Linguistics.
- Saif Mohammad. 2011. [From once upon a time to happily ever after: Tracking emotions in novels and fairy tales](#). In *Proceedings of the 5th ACL-HLT Workshop on Language Technology for Cultural Heritage, Social Sciences, and Humanities*, pages 105–114, Portland, OR, USA. Association for Computational Linguistics.
- Saif Mohammad. 2018a. [Obtaining reliable human ratings of valence, arousal, and dominance for 20,000 English words](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 174–184, Melbourne, Australia. Association for Computational Linguistics.
- Saif M. Mohammad. 2018b. Word affect intensities. In *Proceedings of the 11th Edition of the Language Resources and Evaluation Conference (LREC-2018)*, Miyazaki, Japan.
- Saif M. Mohammad. 2022. Ethics sheet for automatic emotion recognition and sentiment analysis. *Computational Linguistics*, 48(2):239–278.
- Saif M. Mohammad. 2023. Best practices in the creation and use of emotion lexicons. In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, Dubrovnik, Croatia. Association for Computational Linguistics.
- Saif M. Mohammad. 2024a. NRC VAD Lexicon v2: Norms for valence, arousal, and dominance for over 55k English terms. In *ArXiv*.
- Saif M. Mohammad. 2024b. Worrywords: A massive lexicon of anxiety associations for English words and phrases. In *ArXiv*.
- Saif M. Mohammad and Svetlana Kiritchenko. 2015. [Using hashtags to capture fine emotion categories from tweets](#). *Computational Intelligence*, 31(2):301–326.
- Saif M. Mohammad and Peter D. Turney. 2010. Emotions evoked by common words and phrases: Using Mechanical Turk to create an emotion lexicon. In *Proceedings of the NAACL-HLT Workshop on Computational Approaches to Analysis and Generation of Emotion in Text*, LA, California.
- Saif M. Mohammad and Peter D. Turney. 2013. Crowdsourcing a word–emotion association lexicon. *Computational Intelligence*, 29(3):436–465.
- Agnes Moors, Jan De Houwer, Dirk Hermans, Sabine Wanmaker, Kevin Van Schie, Anne-Laura Van Harmelen, Maarten De Schryver, Jeffrey De Winne, and Marc Brysbaert. 2013. Norms of valence, arousal, dominance, and age of acquisition for 4,300 dutch words. *Behavior research methods*, 45(1):169–177.
- Emily Ohman, Yuri Bizzoni, Pascale Feldkamp Moreira, and Kristoffer Nielbo. 2024. [EmotionArcs: Emotion arcs for 9,000 literary texts](#). In *Proceedings of the 8th Joint SIGHUM Workshop on Computational Linguistics for Cultural Heritage, Social Sciences, Humanities and Literature (LaTeCH-CLfL 2024)*, pages 51–66, St. Julians, Malta. Association for Computational Linguistics.
- C.E. Osgood, Suci G., and P. Tannenbaum. 1957. *The measurement of meaning*. University of Illinois Press.
- Jane Parish. 2000. The age of anxiety. *The Sociological Review*, 48(2_suppl):1–16.
- Robert Plutchik. 1980. A general psychoevolutionary theory of emotion. *Emotion: Theory, research, and experience*, 1(3):3–33.
- John Scott Price. 2013. An evolutionary perspective on anxiety and anxiety disorders. *New insights into anxiety disorders*, 3:20.
- Emily Qureshi-Hurst. 2022. Anxiety, alienation, and estrangement in the context of social media. *Religious Studies*, 58(3):522–533.
- Aryan Rastogi, Liu Qian, and Erik Cambria. 2022. Stress detection from social media articles: New dataset benchmark and analytical study. In *2022 IEEE World Conference of Computational Intelligence (WCCI)*.
- Jeffrey B Rosen and Jay Schulkin. 1998. From normal fear to pathological anxiety. *Psychological review*, 105(2):325.
- James A Russell. 1980. A circumplex model of affect. *Journal of personality and social psychology*, 39(6):1161.
- Renata Salecl. 2004. *On anxiety*. Routledge.
- Judy Hanwen Shen and Frank Rudzicz. 2017. [Detecting anxiety through Reddit](#). In *Proceedings of the Fourth Workshop on Computational Linguistics and Clinical Psychology — From Linguistic Signal to Clinical Reality*, pages 58–65, Vancouver, BC. Association for Computational Linguistics.
- Lisa M Shin and Israel Liberzon. 2010. The neurocircuitry of fear, stress, and anxiety disorders. *Neuropsychopharmacology*, 35(1):169–191.

Andy P Siddaway, Peter J Taylor, and Alex M Wood. 2018. Reconceptualizing anxiety as a continuum that ranges from high calmness to high anxiety: The joint importance of reducing distress and increasing well-being. *Journal of personality and social psychology*, 114(2):e1.

Elizabeth C Stadel, Lyle Ungar, Johannes C Eichstaedt, Garrick Sherman, and Ayelet Meron Ruscio. 2023. Depression and anxiety have distinct and overlapping language patterns: Results from a clinical interview. *Journal of psychopathology and clinical science*.

Daniela Teodorescu and Saif Mohammad. 2023. [Evaluating emotion arcs across languages: Bridging the global divide in sentiment analysis](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 4124–4137, Singapore. Association for Computational Linguistics.

Elsbeth Turcan and Kathy McKeown. 2019. [Dreaddit: A Reddit dataset for stress analysis in social media](#). In *Proceedings of the Tenth International Workshop on Health Text Mining and Information Analysis (LOUHI 2019)*, pages 97–107, Hong Kong. Association for Computational Linguistics.

Melissa LH Võ, Markus Conrad, Lars Kuchinke, Karolina Urton, Markus J Hofmann, and Arthur M Jacobs. 2009. The berlin affective word list reloaded (bawl-r). *Behavior research methods*, 41(2):534–538.

Amy Beth Warriner, Victor Kuperman, and Marc Brysbaert. 2013. Norms of valence, arousal, and dominance for 13,915 English lemmas. *Behavior Research Methods*, 45(4):1191–1207.

Joseph P Weir. 2005. Quantifying test-retest reliability using the intraclass correlation coefficient and the sem. *The Journal of Strength & Conditioning Research*, 19(1):231–240.

Genta Indra Winata, Onno Pepijn Kampman, and Pascale Fung. 2018. Attention-based lstm for psychological stress detection from spoken language using distant supervision. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6204–6208. IEEE.

Mark Wright, Robin Banerjee, Willemijn Hoek, Carolien Rieffe, and Sheida Novin. 2010. Depression and social anxiety in children: Differential links with coping strategies. *Journal of abnormal child psychology*, 38:405–419.

Moshe Zeidner and Gerald Matthews. 2010. *Anxiety 101*. Springer Publishing Company.

- All terms in the NRC Emotion Lexicon v1 ([Mohammad and Turney, 2013, 2010](#)): It includes about 14,000 common English words with labels indicating whether they are associated with any of the eight notable discrete emotions identified by [Plutchik \(1980\)](#): anger, anticipation, disgust, fear, joy, sadness, surprise, and trust.
- All terms in the NRC VAD Lexicon v1 ([Mohammad, 2018a](#)): It includes about 20,000 common English terms. Notably, it also includes terms common in X (formerly Twitter). Posts on X include non-standard language such as creatively spelled words (*happee*), hashtags (*#takingastand*, *#lonely*) and conjoined words (*loveumom*). These were taken from the Twitter Emotion Corpus ([Mohammad and Kiritchenko, 2015](#)).
- Terms in the Prevalence dataset ([Brysbaert et al., 2019](#)) that are marked as known to 70% or more respondents: This dataset has prevalence scores (how widely a word is known by English speakers), determined directly by asking people, for 62,000 lemmas. We included a term if it was marked as known to at least 70% of the people who provided responses for the term. From this set we removed terms that are common person names or city names.

A.2 AMT Questionnaire

Screenshots of the detailed instructions, examples, and sample question presented to the annotators are shown in Figures 5, 6, and 7, respectively.

A.3 Distribution of WorryWords

WorryWords is made freely available on the project website as a compressed file. Terms of use will require that users not re-distribute the file and not post any form of the lexicon on the web. This is to prevent the resource being included in the data scrape fed to a large language model. See full list of terms of use at the project home page. Table 6 shows entries for a random sample of words from WorryWords.

A APPENDIX

A.1 Sources of Terms Selected for Annotation

WorryWords includes terms selected from these resources:

Word	Anxiety Score
suffocative	3
smolder	3
exploder	3
exterminate	2.71
ghastly	2.66
haunted	2.62
castrate	2.62
abusive	2.5
manic	2.4
knifelike	2.4
coercion	2.33
conniving	2.16
affliction	2.14
frustrate	2.05
unforeseeable	2
grueling	2
snarling	1.88
obtrusion	1.88
hysterectomy	1.87
scum	1.85
reinjure	1.83
riskily	1.72
downheartedly	1.71
jailhouse	1.71
frostbitten	1.66
stifled	1.62
disturbed	1.57
projectile	1.55
disquieted	1.55
centaur	1.33
brusquely	1.25
psychosomatic	1.12
escapist	1
discordance	0.8
propagandize	0.75
brandy	0.14
ceramics	0.12
photosphere	0
rhinoceros	0
paraphrase	0
forkless	0
witting	0
gemini	-0.1
swordswoman	-0.13
canto	-0.41
roller	-0.5
motorcycle	-0.71
semiprofessional	-0.8
socialization	-0.8
cleanness	-1.55
coldblooded	-1.57
fantail	-1.6
conformed	-1.71
unscratched	-1.75
journalistic	-1.8
fireless	-1.85
reaffirmation	-2
aristocratic	-2.1
limberness	-2.5
unblemished	-2.62
lullaby	-2.79
delightfully	-3
generous	-3

Table 6: Randomly sampled terms and their anxiety-association score from WorryWords.

Introduction:

1. Attempt these questions only if you are fluent in English.
2. Your responses are confidential.

Task:

Words can be associated with different degrees of anxiousness or calmness. While there is some variation from person to person, there is also a fair amount of consensus. For example, most people will agree that the term:

- *nail biting* is often associated with **very high anxiousness**
- *stroll* is often associated **moderately calmness** feelings
- *serene* is often associated with **high calmness** feelings
- *table* is often **not associated** with calmness or anxiousness
- *test* is often associated with **moderate or high anxiousness**
- *waiting* is often associated with **moderate anxiousness**
- *problem solved* is often associated with **moderate calmness**

In this multiple choice task, you will be given common English terms and you have to select the options that best describe the degree of anxiousness or calmness associated with them.

Consider **anxiety** to be a broad category that includes:

- jittery, antsy, insecure, nervous, unease, tense, worried, unnerving, nerve-racking, apprehensive, fretful, self-doubting, troubled, discontented, concerned, and keyed up

Consider **calmness** to be a broad category that includes:

- calm, relaxed, comforted, serene, at ease, self-assured, carefree, composed, collected, untroubled, peaceful, contented, nonchalant, unconcerned, indifferent, uninterested, and uninvolved

Give answers that capture what most English speakers would agree.

If you do not know the meaning of a word or are unsure, you can look it up in a dictionary (e.g., the [Merriam Webster](#)) or on the internet.

Purpose of the task:

Your responses will be used in a research study to better understand how anxiety manifests in language.

Quality Control:

- Responses that are not in accordance with the instructions will not be paid for.
- Some questions have pre-determined correct answers. If you mark these questions incorrectly, we will give you immediate feedback in a pop-up box. We will keep track of your answers for these gold questions. **If you mark too many of these incorrectly, it will lead to the rejection of *all* your HITs.**
- If you see that you are getting quite a few of the gold questions wrong (e.g. more than 2 in every 10 HITs), then do not accept more HITs.
- If you disagree with the answer for a gold HIT, include the correct response in the Feedback textbox. Note that missing an occasional gold question will not lead to the rejection of your responses.
- This quality control measure promotes fairness for those who do the task responsibly.

Notes:

- If a term has more than one meaning, consider the most common meaning.
- A rule of thumb is that a term associated with more anxiousness tends to often occur in sentences that convey anxiousness, whereas a term associated with more calmness tends to often occur in sentences that convey calmness.
- Try not to overthink the answer. **Let your instinct guide you.**

Figure 5: Questionnaire: Detailed instructions.

Very anxious:

- jittery, flooding, heart attack, ship sinking, etc.

Moderately anxious:

- argument, sibling rivalry, elevated pulse, commotion, first day, confront, etc.

Slightly anxious:

- unclear, doubt, meeting, secluded, tricky, etc.

Not associated with anxiousness or calmness:

- furniture, envelope, utensil, fry, garage, profession, very, same, percent, etc.

Slightly calm:

- greenery, done, okay, good to know, in control, etc.

Moderately calm:

- stroll, soothed, fishing, fixed, smooth sailing, assured, confidently, etc.

Very calm:

- heaven, serenity, meditate, sleep, massage, etc.

Figure 6: Questionnaire: Examples.

[View instructions](#)

Summary Instructions

Consider **anxiety** to be a broad category that includes:

- jittery, antsy, insecure, nervous, unease, tense, worried, unnerving, nerve-racking, apprehensive, fretful, self-doubting, troubled, discontented, concerned, and keyed up

Consider **calmness** to be a broad category that includes:

- calm, relaxed, comforted, serene, at ease, self-assured, carefree, composed, collected, untroubled, peaceful, contented, nonchalant, unconcerned, indifferent, uninterested, and uninvolved

If you do not know the meaning of a word or are unsure, you can look it up in a dictionary (e.g., the [Merriam Webster](#)) or on the internet.

Quality Control

Some questions have pre-determined correct answers. If you mark these questions incorrectly, we will give you immediate feedback in a pop-up box. An occasional misanswer is okay. However, if the rate of misanswering is high (e.g., >20%), then all of one's HITs may be rejected

Select the options that **most English speakers** will agree with.

Q1. **disease** is often associated with feeling:

- ☐ 3: very anxious
- ☐ 2: moderately anxious
- ☐ 1: slightly anxious
- ☐ 0: not associated with feeling anxious or calm
- ☐ -1: slightly calm
- ☐ -2: moderately calm
- ☐ -3: very calm

Figure 7: Questionnaire: Sample Question.

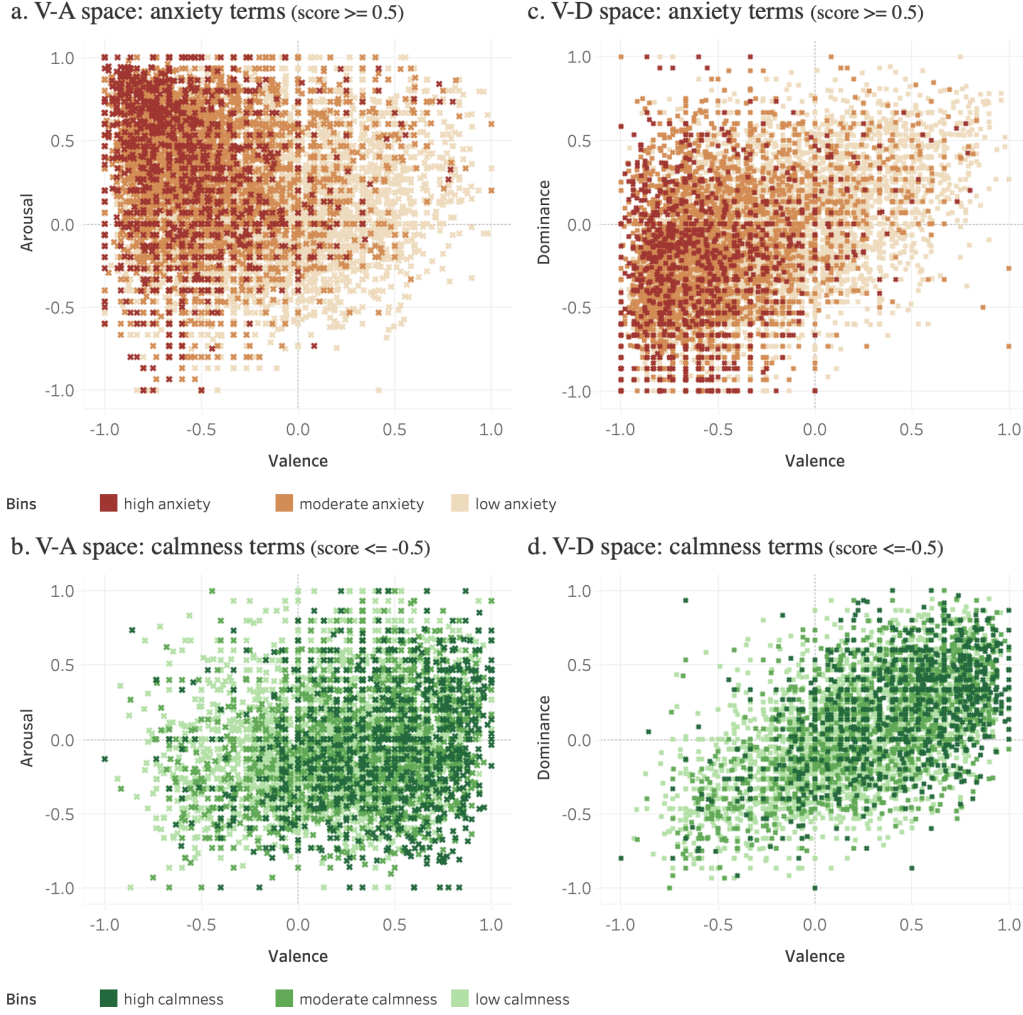


Figure 8: WorryWords terms in the Valence–Arousal (a and b) and Valence–Dominance (c and d) spaces.

A.4 Random Arc Sampling Algorithm

In order to create a text stream (described in Section 6) such that the gold arc has various components of random ups and downs of varying slopes, we sample posts from a dataset of anxiety (1) and no-anxiety (0) posts, with replacement, in short bursts as follows. For each burst (sub-arc), we first randomly determine a slope and number of posts to make up the sub-arc. Then we sample posts from the source datasets accordingly. For example, sampling only anxiety posts will lead to a slope of 90 degrees, sampling only no-anxiety posts will lead to a slope of -90 degrees, and sampling a mixture will lead to a slope between -90 and +90 (depending on the proportion of anxiety and no-anxiety posts). We continue sampling until the randomly chosen length of the subarc is reached. Then the whole procedure is repeated with new randomly generated slopes and lengths of subarcs until the text stream includes 10,000 posts.

A.5 Supplementary Figures and Tables

Figure 8 shows plots of the WorryWords terms on valence–arousal and valence–dominance spaces (described in Section 6.1). Table 7 shows the Spearman rank correlation between the word–emotion association scores for various emotion pairs (discussed in Section 6.1).

Figure 9 shows predicted (green) anxiety arcs when applying nine different emotion association scores to one of the test data text streams (described in Section 7).

A.6 Computational Resources and Carbon Footprint

A nice advantage of using simple lexicon-based approaches is the low carbon footprint and computational resources required. All of the experiments described in the paper were conducted on a regular personal laptop.

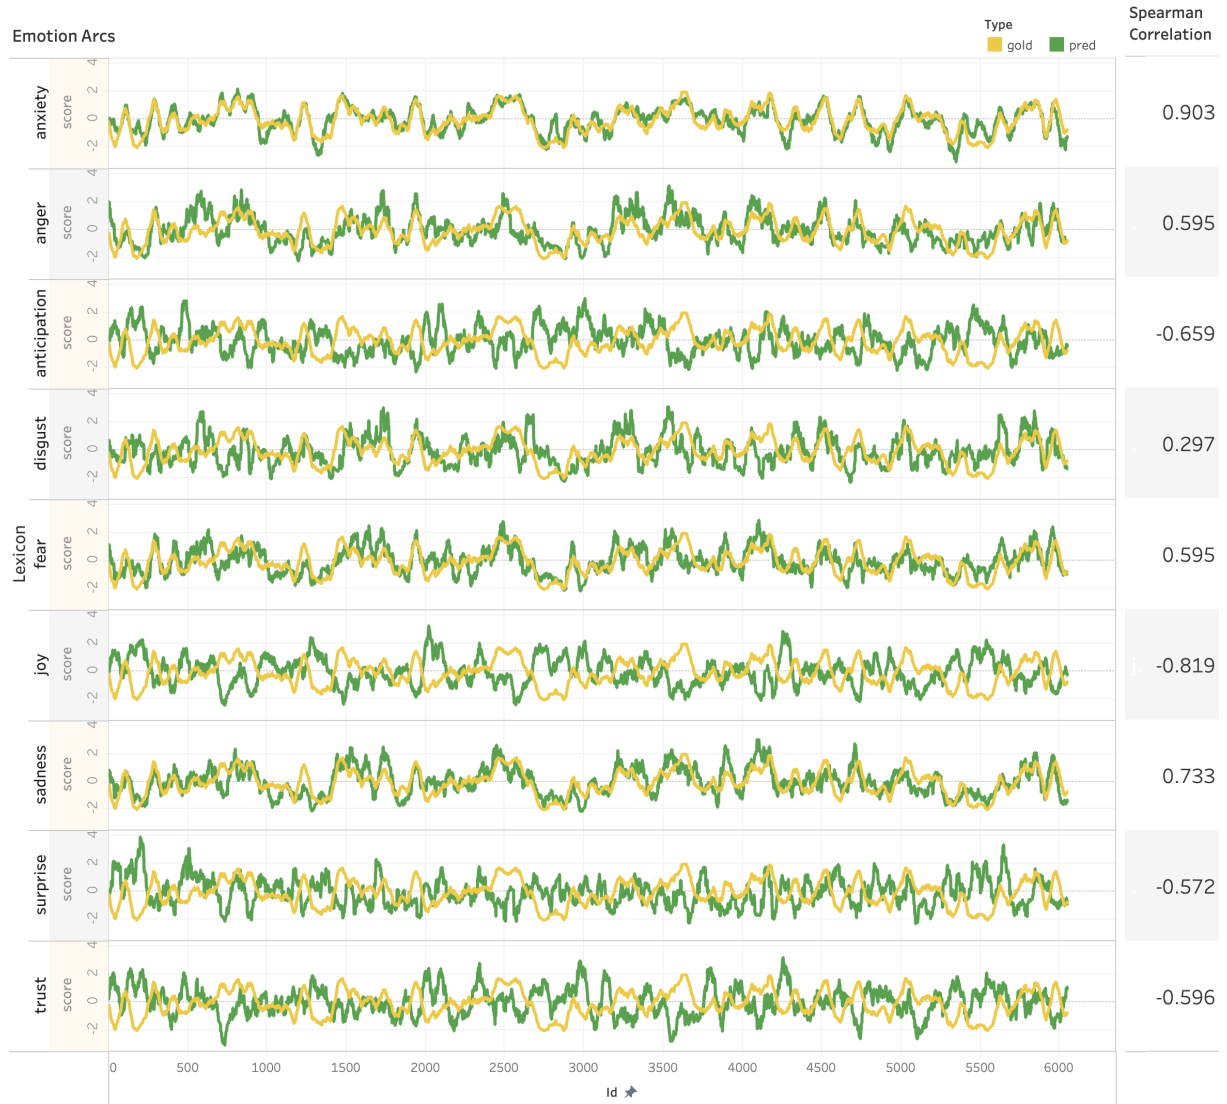


Figure 9: Predicted (green) anxiety arcs when applying nine different types of association scores to one of the test data text streams: 1. associations with anxiety (from WorryWords); 2–9. associations with each of the eight Plutchik emotions (from NRC Emotion Intensity Lexicon). The gold (yellow) arc pertaining to the same text stream is repeated in all nine figures for reference. Bin size used to create the arcs: 50 posts. On the right is the average correlation between 1000 pairs of gold and predicted arcs for each case.

	anxiety	anger	anticipn.	disgust	fear	joy	sadness	surprise	trust
anxiety	1.000	0.336	-0.011	0.279	0.357	-0.132	0.272	0.082	-0.137
anger	0.336	1.000	0.025	0.432	0.461	-0.023	0.372	0.117	-0.048
anticipn.	-0.011	0.025	1.000	-0.014	0.074	0.432	0.020	0.276	0.252
disgust	0.279	0.432	-0.014	1.000	0.328	-0.036	0.323	0.059	-0.054
fear	0.357	0.461	0.074	0.328	1.000	-0.017	0.439	0.167	-0.035
joy	-0.132	-0.023	0.432	-0.036	-0.017	1.000	-0.014	0.283	0.386
sadness	0.272	0.372	0.020	0.323	0.439	-0.014	1.000	0.088	-0.047
surprise	0.082	0.117	0.276	0.059	0.167	0.283	0.088	1.000	0.106
trust	-0.137	-0.048	0.252	-0.054	-0.035	0.386	-0.047	0.106	1.000

Table 7: Spearman rank correlation between the word–emotion association scores for various emotion pairs. For a given anxiety–emotion pair, the correlation is computed for terms that are included in corresponding emotion lexicons.