

# Detecting Inspiring Content on Social Media

Oana Ignat  
University of Michigan  
MI, USA  
oignat@umich.edu

Y-Lan Boureau  
Facebook AI  
New York City, USA  
ylan@fb.com

Jane A. Yu  
Facebook AI  
Menlo Park, USA  
janeyu@fb.com

Alon Halevy  
Facebook AI  
Menlo Park, USA  
ayh@fb.com

**Abstract**—Inspiration moves a person to see new possibilities and transforms the way they perceive their own potential. Inspiration has received little attention in psychology, and has not been researched before in the NLP community. To the best of our knowledge, this work is the first to study inspiration through machine learning methods. We aim to automatically detect inspiring content from social media data. To this end, we analyze social media posts to tease out what makes a post inspiring and what topics are inspiring. We release a dataset of 5,800 inspiring and 5,800 non-inspiring English-language public post unique ids collected from a dump of Reddit public posts made available by a third party and use linguistic heuristics to automatically detect which social media English-language posts are inspiring.

**Index Terms**—inspiration, social media data, natural language processing, emotions, sentiment

## I. INTRODUCTION

Inspiration is a distinct psychological construct [1], conceptualized by Thrash and Elliot as possessing three core characteristics: evocation (i.e., it is triggered rather than willed), transcendence (i.e., it orients towards things outside of and greater than the self), and approach motivation (i.e., it energizes approach rather than avoidance [1]–[3]). Inspiration has two distinct stages: one an activation state that is more akin to feeling and emotion, the second an urge to act. Inspiration is also related to “other-praising emotions” such as gratitude, elevation, and admiration [4], in that those emotions often result in inspiration. Inspiration, however, is a broader concept and more focused on the resulting increased motivation to act.

### A. The appeal of inspiration

Inspiration is a promising target of study because it is richly associated with a range of positive outcomes, while its absence is described by lay people as something akin to depression [5]. It can help humans access more creativity [6], productivity, and happiness [7], facilitates progress towards goals [8], and promotes both hedonic (pleasure-oriented) and eudaimonic (growth-oriented) well-being [9]–[11]. In quantitative analyses of inspiration, Thrash et al. showed that inspiration positively correlates with positive affect, the work-mastery component of the need for achievement (but negatively with its competitiveness component), and intrinsic motivation (but negatively with extrinsic motivation) [1]. They also documented that inspiration appears to cause rather than simply follow absorption (a state of focused attention on object qualities such as beauty, rather than diffuse arousal), work mastery, creativity, perceived

competence, self-esteem, and optimism. Inspiration can also predict measurable behavioral outcomes such as exploration and purchase [12], and correlates with outcomes such as higher number of patents [1]. Recall of inspiring content has also been shown to result in significantly more frequent persistence in an unrelated physical endurance task, compared to mere recall of an amusing event [13], with people recalling an amusing event being 5 times more likely to give up than those recalling an inspiring one.

### B. Scope of this work

Despite the attractive promise of inspiration, there has been little work on automatically detecting content that is specifically inspiring, rather than merely engaging or positive. Inspiration requires an “encounter” [5], an event “in which new or better possibilities are revealed by, or revealed in an evocative stimulus object” [2]. Our work aims to facilitate such encounters by providing tools for automatic identification of text content likely to be judged inspiring. We focus on inspiration in everyday content as judged by lay people, similar in spirit to early work by Hart who attempted to capture the experience of inspiration in ordinary life [5], rather than “as if it were reserved for the gifted artist, the breakthrough scientist, or the extraordinary mystic.” This motivates our focus on social media content, as well as our use of crowdsourcing, rather than expert labeling.

The contribution of this work is threefold. (1) We introduce the novel task of automatic detection of inspiring content from social media data. To facilitate the task, we offer a novel dataset of inspiring and non-inspiring post ids from Reddit, sourced from a dump of public posts made available by a third party, pushshift.io [14]. (2) We combine experiments comparing several weak labeling techniques to retrieve English-language inspirational content with a classifier trained on human labels to provide a strong baseline to determine if an English-language post is inspiring or not. (3) We provide a detailed analysis of data labeled as inspiring to gain insights on which topics are inspiring, and how they influence the readers.

## II. RELATED WORK

### A. Characterizing inspiring content and its role

Several works have analyzed characteristics of inspiring content. Interviews in the domain of sports have revealed features of coach speeches to athletes that were deemed inspiring, both in terms of content and delivery [15], for example

showing that inspiring speeches provide an illustration of a path towards success or an increase in the sense of belonging, but the work is limited to a sample of young, male athletes.

There have also been some recent papers on inspiring or self-transcendent content (i.e., orienting people towards things outside of or greater than the self) in relation to social media. Oliver et al. [16] argues for the importance of studying self-transcendent content to promote well-being and other-oriented connectedness. Meier et al. showed that benign envy triggered by social media content could lead to inspiration and increased positive affect [17], and that exposure to enhanced positive content on Instagram led to higher feelings of inspiration and brief improvements in well-being [18]. This is in contrast with findings of negative consequences when inspiration was not involved [19]. According to [20], social media is an important source of inspiring content, especially for younger audiences. Dale et al. studied 100 videos on YouTube tagged with “inspiration,” and identified a variety of specific features in media eliciting self-transcendent emotion, such as triggers associated with hope (e.g., “underdog narratives”), appreciation of beauty and excellence (e.g., nature, art, vastness) [21]. Similar triggers were found in “#inspirational” and “#meaningful” Tumblr memes [22], and Facebook posts [23]. [21] also reported that 76.5% of surveyed users had been inspired at some point by social media content, with 34.9% reporting this being the case within the past week, and the most commonly reported inspiring themes being kindness, overcoming obstacles/perseverance, human connection, and gratitude or thankfulness.

### B. Quantifying word use in inspiring content

Ji et al. explored textual aspects of inspiring New York Times articles reshared on social media [24], [25]. They constructed a Self-Transcendent Emotion Dictionary (STED) to detect the presence of 370 English words, word stems, and phrases corresponding to six self-transcendent emotional experiences (awe, admiration, elevation, gratitude, hope, and general inspiration), and validated how it could be used to quantify self-transcendent qualities in text material without resorting to sophisticated natural language processing techniques. However, [25] reports that “the observed effects sizes appear to be relatively small,” making counts of STED words an insufficient tool to detect inspiring material, all the more so for social media content where word counts are considerably smaller than in news articles and thus extremely noisy. The STED could nonetheless have been useful as one of the filtering heuristics to select posts over which to train our classifier, however it was not yet available at the time we conducted this work. Another difference consists in the target of the word pattern use: in [24], [25], the word patterns apply only to the content itself, rather than also to comments about the content, as we do in this work to capture the fact that inspiration involves elicitation of a response in someone else.

### C. Machine learning methods for affective computing

Despite the fact that inspiration is not, strictly speaking, an emotion, the work closest to ours in terms of the techniques used is in the broader context of emotion recognition (e.g., [26], [27]). Emotion recognition is a sub-field of Affective Computing whose goal is to enable machines to understand and emulate affect and emotion [28]. The first benchmark for emotion recognition, Affective Text [26], [29] opened the field to several emotion datasets that vary in size, domain, taxonomy and applications: news headlines [26], tweets [30]–[32], Reddit posts [33], narrative sequences [34] and dialog [35], [36]. One of the key aspects distinguishing our work from the above is that while other works try to detect the emotion of the *publisher* of the content, our goal is to detect whether the *consumer* of the content would be inspired or not.

To the best of our knowledge, there hasn’t been any work using machine learning techniques for textual content, except for Inspirobot [37], which aims to automatically generate inspiring content, but appears more geared towards satire and entertainment than a serious study of inspiration.

## III. DATA COLLECTION AND ANNOTATION

In order to develop models for recognizing inspiring content, we need test data annotated with human annotations. This section describes the details of our dataset.

### A. Problem definition

The goal of this research is to develop models that can recognize whether a post on social media is likely to inspire someone who reads that post. We focus on text, but recent work on recognizing human intent of images [38] could form the basis for considering multi-modal content in future work.

Whether a social media post is inspiring to a user depends on the user and their current state and motivations in life, nevertheless the models we discuss in this paper are not intended to make predictions for specific users. Rather, we aim to predict whether a piece of content would be generally rated as inspiring by human annotators. We rely on a lay understanding of inspiration, following Thrash and Elliot, who confirmed the relative consistency of how inspiration is interpreted by observing measurement invariance through time and populations [1], and cited interview research by Hart: “Despite the great many shades of meaning the term has in common usage it appears to represent a clear and consistent event” [5], [39].

### B. Data collection

We annotate two datasets. The first, which we refer to as pR, is taken from a dump of public posts made available by a third party (pushshift.io Reddit [14]). The second dataset, which we refer to as S, consists of public social media posts from Facebook. The dataset we make available is based on pR.

The first challenge in creating the dataset for annotation is that the vast majority of posts on social media are not inspiring. Therefore, in order to ensure that the annotators are exposed

to a sizeable amount of inspiring posts, we filter the data using the following heuristics: (1) public posts with at least one comment that contains the substrings “inspir” or “uplift” (pR & S), (2) public posts that authors mark as “feeling inspired” or “feeling up” (S), (3) public posts that are shared at least 10 times (S), (4) public posts from the subreddits that contain the substrings “inspir” or “uplift” (pR), and (5) comments to the following four questions from the “AskReddit” subreddit: “When was the last time you felt inspired?”, “Who or what inspired you?”, “Who inspired you and how?”, “What is the most inspiring thing you have ever seen or heard?” (pR). As control, we also collect random posts: (1) posts with no comment that contains the substrings “inspir” or “uplift” (pR & S), (2) posts from random subreddits that do not contain the substrings “inspir” or “uplift” (pR).

### C. Data pre-processing

We filter out posts that are not in English using the language classifier from fastText [40]. We remove links and double spaces and we filter out posts that contain less than 10 words or more than 200 words. We remove posts that contain offensive language or profanity using an SVM model [41] trained on two datasets of hate speech in social media [42], [43]. This basic process of removing offensive language could be improved, for example by using careful human labeling to decide what should not appear in the dataset. However, there is not yet a fully agreed upon method, and best practices for diverse and inclusive data collection are still very much a matter of active research [44]–[48].

For data analysis, we further pre-process the data: we remove stop words and specialized words (*subscribe*, *comment*, *like*, *follow*, *link*), special characters and digits, hashtags, punctuation, and emojis, whose prevalence in the dataset was not high enough for analysis. We also convert the words to lowercase and lemmatize them.

### D. Data annotation

The resulting posts are annotated by crowdsourced workers to determine: (1) whether the post is inspiring or not; (2) if the post is inspiring, what influence it has on the reader; (3) what emotions it evokes; (4) the annotator’s confidence in the answer. The instructions and interface for annotation are shown in Figure 1. We provide the annotators with a short definition of what inspiration is, one example of an inspiring post and one example of a non-inspiring post. The two example posts from the guidelines are selected by the paper authors, among the filtered posts.

Each post is labeled by three different annotators. We compute the agreement between the annotators using Fleiss Kappa measure [49] and we obtain 0.26, indicating fair agreement yet showing that the task remains subjective.

Due to the challenge of finding inspiring posts at scale, our annotation proceeded in several phases. In the first phase, we annotated 1000 pR posts: 500 from the subreddits that contain the substrings “inspir” or “uplift” and 500 from random subreddits. We find that only 167 posts are labeled as inspiring

# NON-INSPIRING	# INSPIRING		
	1/3	2/3	3/3
5,796	1,517	1,936	2,343

TABLE I  
NUMBER OF HUMAN ANNOTATED POSTS: INSPIRING AND NON-INSPIRING, GROUPED BY AGREEMENT SCORE.

by at least two annotators, and most of these posts are of a particular genre of inspiring quotes. We also find very few false positive posts (9 /500), which suggests that non-inspiring posts are much easier to collect. Therefore, in the next annotation phases, we focus on collecting inspiring posts.

In order to obtain a bigger fraction of inspiring posts, we changed the filtering heuristic in the second phase to consider the content of comments on posts. Specifically, we collect posts with at least one comment that contains the substrings “inspir” or “uplift” and responses to posts asking questions that contain the substring “inspir”, from the “AskReddit” subreddit (pR): “When was the last time you felt inspired?”, “Who or what inspired you?”, “Who inspired you and how?”, “What is the most inspiring thing you have ever seen or heard?”. These posts turn out to be much more diverse and do not focus on inspiring quotes. From a total number of 3,415 annotated posts, 1,225 posts are labeled as inspiring by at least two annotators - considerably more than the first phase.

In the third phase we use the previously annotated posts to fine-tune a pre-trained RoBERTa model [50] to the task of predicting inspiring content. The trained classifier is then used as additional filtering to select the posts to be annotated from the ones that fit the second heuristic. A third set of posts selected in that process is annotated by workers. From 3,210 annotated posts, 3,054 are labeled as inspiring by at least two annotators. This final selection pipeline thus seems to be predictive of subsequent human annotation as “inspiring.”<sup>1</sup>

Non-inspiring posts are posts collected from random subreddits and posts that match the heuristics yet were not labeled inspiring. A post is labeled as non-inspiring if the trained classifier labeled it as non-inspiring or if no annotator labeled it as inspiring. After balancing the number of inspiring and non-inspiring posts, we obtain 11,592 total posts: 5,796 inspiring and 5,796 non-inspiring. A post is inspiring if at least one annotator labeled it as inspiring. The inspiring post is assigned an agreement score corresponding to the number of annotators who labeled it as inspiring (3/3 is the maximum agreement score and 1/3 is the minimum). Note that the agreement score of the post is different from the confidence of the annotator, which can be low or high. We find that the annotators are highly confident in their answers 91% of the time, and thus we decided to not use the confidence of the annotators in our experiments. We show the annotation results in Table I.

<sup>1</sup>Note that despite our efforts to obtain a wide variety of inspiring posts, we do not have a measure of the inspiring posts that we missed, only results showing that the posts we select are largely deemed inspiring.

<b>Instructions</b>	\$(text)
Please read the text provided (the post) and answer if it is inspirational or not and why. We are NOT asking if the poster is inspired.	Do you agree that the above text is inspiring? <input type="radio"/> Yes <input type="radio"/> No
<b>Definitions</b>	If so, why is the text inspirational? <input type="radio"/> It makes me feel good <input type="radio"/> It may influence me to take action <input type="radio"/> Both <input type="radio"/> Not inspirational Type any other reasons here: _____
<b>Inspirational</b>	Something that fills people with an animating or exalting influence, or produces rousing feelings, novel thoughts, or a motivation to act. If a link, photo or video seems missing, please use your best judgement and indicate low confidence as necessary.
Example: "Never regret your decision, every mistake makes u smarter and stronger"	What emotions does the post invoke in you? (Check all that apply) <input type="checkbox"/> Admiration <input type="checkbox"/> Gratitude <input type="checkbox"/> Curiosity <input type="checkbox"/> Optimism <input type="checkbox"/> Neutral Type any other emotions here: _____
<b>Not inspirational</b>	What is your confidence in your answer? <input type="radio"/> Low Confidence <input type="radio"/> High Confidence
Example: "Three months going on one year. He is such a light in the midst of all the chaos. We are so thankful for our little boo."	

Fig. 1. User Interface for the Amazon Mechanical Turk annotation platform.

#### IV. CLASSIFICATION OF INSPIRING POSTS

We now describe our baseline classification methods for the task of detecting inspiring posts on social media. We experimented with fastText [51], and fine-tuned pre-trained language models BERT [52] and RoBERTa [50]. In each case, we fine-tuned the model for our task using our annotated training data. We balance the training data to have an equal number of inspiring and non-inspiring posts. The statistics for this experimental split is shown in Table II. In the experiments, we used early stopping, 5 training epochs and a batch size of 64. More training epochs (10 and 20) did not improve the results.

The results, shown in Table III, show that the model has a high accuracy when distinguishing inspiring from non-inspiring posts. In the experiments we train the model on each dataset separately and the best performing model is RoBERTa. In addition, we tried a transfer learning experiment where we trained the classifier on the dataset S and tested it on the data in pR (last line in Table III). This results in lower performance than training specifically on pR, as expected for data from a different distribution, but there is still some transfer of performance. We hypothesize that the transfer learning was not stronger because of the different nature of the posts in the two datasets: while the posts in S tend to be more personal short stories or announcements, the inspiring posts in pR tend to be more general and broad, such as news, discussions, and asking or answering questions.

During error analysis, we find that the most common errors our best model makes are false negatives: posts that are collected as inspiring, but labeled by the model as non-inspiring. We find that in most cases, the posts that confuse the model are the ones annotated as inspiring by one annotator.

#### V. ANALYZING INSPIRING POSTS

An important goal of our research is to gain insights into what makes a social media post inspiring. Extensive analysis

	S		pR	
	inspiring	non-inspiring	inspiring	non-inspiring
Train	642,517	642,517	5,216	5,216
Test	71,390	71,390	580	580

TABLE II  
NUMBER OF POSTS FOR THE EXPERIMENTAL DATA SPLIT.

Method	S		pR	
	Acc.	F1	Acc.	F1
fastText	74.66	74.58	76.20	76.20
BERT	75.18	76.22	83.53	84.23
RoBERTa	79.81	80.73	<b>87.24</b>	<b>88.08</b>
RoBERTa*	79.81	80.73	67.93	71.73

TABLE III  
RESULTS FROM BASELINES AND MAIN METHOD ON S AND pR TEST DATA.  
\* - MODEL WAS FINE-TUNED ON S DATA. S IS WEAKLY LABELED AND pR IS HUMAN ANNOTATED.

of our data through topics and sentiment analyses confirm the following about inspiring content on social media: *inspiration is a complex state which can contain both positive and negative emotions, and can be found across a diverse range of every-day moments and activities.*

We analyse two different types of inspiring posts: posts that *make people feel good* by changing their mood or their perspective (e.g., feel gratitude, admiration, curiosity), and posts that *make people act* on their thoughts and wishes (e.g., start a new hobby or a new work routine). Table IV shows the split among these categories from our annotated data. Table V shows example posts from each category. Our annotation also allowed workers to add other motives for why they find the post inspiring. Examples from the “other” category include: “its practical life”, “its real truth”, and comments that relate the post to their previous life experience, such as “I was the second fastest sprinter female in my school”, or “people are willing to go through so much for their animals”.

From our data observations, the inspiring posts can also be categorized as *general* or *personal*. Posts like inspiring news or quotes apply to most people and circumstances as they promote popular human values such as gratitude, determination or kindness (e.g., “Don’t be content everyday to do no wrong, be prepared every day to do good” @6:50 in

# INSPIRING POSTS (5,796 IN TOTAL)				
	MOTIVATION TO ACT	FEEL GOOD	NO EFFECT	OTHER
≥ 1	3,333	4,382	3,470	564
≥ 2	1,052	1,711	1,476	13
≥ 3	149	317	20	0

TABLE IV  
NUMBER OF INSPIRING POSTS LABELED WITH THEIR EFFECTS ON THE READERS BY AT LEAST ONE, TWO AND THREE (ALL) ANNOTATORS.

MOTIVATION TO ACT	FEEL GOOD
<p>“This is very inspiring to me, as I will be returning to college soon to begin schooling in Physics. Being schizophrenic can (and will) make you think you won’t get far in life, and its very easy to fall into a non-productive spiral really quick. People like John Forbes Nash Jr. and this fellow can really inspire without even meaning to, and that gives me a lot of hope for my future. Good Post.”</p>	<p>“I say follow your heart. Play what you’re inspired to play. If you’re inspired to polish a piece, work on that. If you’re inspired to just relax and play pieces you already know, do that. The most productive practicing happens when your mind is most concentrated, and when you’re inspired by something, your mind is automatically attracted to that thing. I think it’s best to take advantage to that phenomenon.”</p>
<p>“Just saying you inspired my dad to become a bodybuilder, and inspire me to work out and stay fit. Don’t let anyone get you down, you’re the terminator! Your a good example about you can do anything if you set your mind to it!”</p>	<p>“As simple as this sounds, I’m happy that I have a roof to live under, food to eat, and people that care about me. We all take that for granted until its all gone.”</p>
<p>“Do something to help or enrich the lives of others. Could be volunteering at a youth/vulnerable peoples centre, getting involved in a local art project, or making a set of instructional videos on youtube about something you know a lot about. If you inspire just a single person you’re still making the world that little bit better.”</p>	<p>“Motivation is key. Some mornings I look out, it’s cold raining wet and miserable, I just pop in my headphones and play some uplifting music, stuff to get you pumped up. Failing that, run with others. That way you have the motivation of not letting others down. Join a crew so if you cancel it’s not yourself you’re letting down, but the crew.”</p>

TABLE V  
EXAMPLES OF INSPIRING POSTS FROM PR WITH DIFFERENT INFLUENCE ON THE READER.

that video, definitely advice to live by”). These are different from the more personal posts such as someone sharing how they overcame a personal struggle (e.g., “Over the past year I’ve lost some weight, but have been feeling pretty down and hate how most of my clothes look. You have inspired me to donate what I hate and go get some damn clothes that fit. Thank you!”).

#### A. Topic analysis

Inspiration is present across multiple topics: *spiritual* (e.g., religion, mindfulness), *physical* (e.g., hobbies, travel) and *artistic* (e.g., dance, painting). In what follows, we analyze the distribution of these topics.

**TF-IDF score:** We compute the tf-idf [53] scores of all the words in a document composed of all the inspiring and non-inspiring posts. We show the top 20 words from inspiring posts, sorted by their tf-idf score in Table VI.

**Word frequency:** Table VI shows the words that are most frequent in the inspiring posts, compared to the non-inspiring ones, the top 20 frequent bi-grams and tri-grams and the top frequent hashtags in the inspiring posts. The hashtags act like keywords, summarizing the content of the S and pR posts. The words can be grouped into 4 main categories: (a) names of famous figures (e.g., philosophers, authors, athletes, activists, scientists); (b) human values; (c) spirituality; (d) art and beauty.

**Clustering topics:** We represent the words using GloVe [54] embeddings, the 200 dimensional model trained on two billion Tweets. We apply k-means [55] clustering on the most important (sorted by their tf-idf score) 200 words from inspiring posts. Using the elbow method, we determine that the best number of clusters is 10. After applying t-distributed Stochastic Neighbor Embedding (t-SNE) [56], the two dimensional plots can be seen in fig. 2 and fig. 3.

From the clusters, we observe words related to spiritual domain: *heavenly, faithful, cherish, conscious, psalm, hymn*; words related to human values: *generosity, philanthropy, willingness, willpower, devotion, charismatic, defy, minimalism*;

and words related to artistic or creative pursuits: *poetic, writing, cello, recite, teaching, cuisine*

From the clusters, we can also observe groups of words that express positive and negative emotions: *precious, thrive, generosity, willingness, poetic, philanthropic, majestic, heavenly, cherish, motivates, delighted, scholarship, phd, teaching, firefighter* vs. *misery, adversity, spite, cruelty, arrest, defy, sorrow, starve, endure, WWII, burial, refugee, debt, palsy*. This reflects the fact that inspiration is often found in circumstances where someone rises above adversity.

Examples of these pairs can be seen in posts about overcoming difficulties, which is a popular theme in inspiring posts (e.g., “That a kid with a younger brother with cerebral palsy runs triathlons with his brother attached to him”, “Man, I gotta start running again! The fact that your post has the picture of Emi that I was inspired to draw after beating her arc has me even more motivated!”, “96 Year-Old Latino Man Graduates with Top Honor, Survives WWII, and Now the Pandemic”)

#### B. Emotion and Sentiment Detection

We also analyze the sentiment and emotions expressed in inspiring posts using automatic classification models and human annotation.

a) *Sentiment Analysis:* We measure the sentiment of the inspiring posts using polarity<sup>2</sup>: a score between -1 and 1 which corresponds to how negative or how positive the sentiment of the post is. The results can be seen in Figure 7. We observe that the inspiring posts with positive sentiment are the most frequent.

b) *Emotions:* We also asked annotators to select emotions they felt while reading the posts, among a set of candidates (admiration, gratitude, curiosity, optimism and neutral) that looked likely from previous research (e.g., [4]). The annotators also have the option to type in any other emotion they consider relevant (Figure 4).

For the emotion classifier, we use the GoEmotions dataset [33], the largest manually annotated dataset of 58k English

<sup>2</sup>Polarity is measured using <https://textblob.readthedocs.io/en/dev/>.




Fig. 2. K-means [55] clustering of most important (sorted by their tf-idf score) 200 word embeddings from inspiring pR posts.




Fig. 3. K-means [55] clustering of most important (sorted by their tf-idf score) 200 word embeddings from inspiring S posts.

Reddit comments, labeled for 27 emotion categories or Neutral. We fine-tune the RoBERTa [50] multi-class classification model on the GoEmotions dataset and apply the model to our pR posts. Admiration, optimism, sadness and curiosity are selected frequently by the emotion classifier, although gratitude isn't (Figure 5).

Figure 6 shows the difference in frequency of the predicted emotions in inspiring and non-inspiring pR posts. As can be seen, *admiration* and *relief* are the most frequent emotions in inspiring posts, compared to the non-inspiring posts, while *sadness* and *remorse* are more present in the non-inspiring ones. Note that many non-inspiring posts were selected among posts that matched our heuristics, so may not be representative of a purely random set of negative examples.

## VI. CONCLUDING REMARKS

In this paper, we introduced the novel task of identifying inspiring posts in social media data. We performed extensive analyses on Reddit (pR) data and other social-media data in order to gain insight into what topics are inspiring, what words are specific to inspiration and what emotions the inspiring posts transmit to readers. To facilitate research in that domain, we release a new dataset consisting of 5,796 inspiring and 5,796 non-inspiring post ids collected from pR. We also annotate these posts with the effect they have on the reader and the emotions they transmit.

We have already noted in the introduction the numerous benefits of inspiration as a promoter of wellness, creativity, and motivation. Another potential benefit for inspiring social media content is that it might crowd out less desirable content. A lot of attention has focused on the AI challenges in mitigating negative aspects of social media, due to their potential role in spreading misinformation, hate speech and other forms




Fig. 4. Most frequent emotions labeled by the Amazon Mechanical Turk annotators in the inspiring posts: curiosity, neutral, gratitude, optimism and admiration are solicited in the annotation task, while the others are not.

of policy violating content [57]. An indirect method for mitigating such harms is for social network recommendation systems to find more positive and inspiring content to show to users. This paper takes a first step towards identifying inspiring content that can be shown to users.

Avenues for future work include extending these methods to identify multi-modal inspiring content, analyzing in more detail the effects of inspiring content (including unintended consequences) on users, and analyzing the style of inspiring posts. More powerful classifiers based on larger language models could also be developed to detect inspiring content directly from random posts without first resorting to heuristics.

S	pR
<b>Top words with highest tf-idf scores</b>	
moses, kamala, jemima ; savior, positive, generosity, brutality, radical ; adoration, kneel ; poet, tailor, diy, choir	lincoln, lennon, swartz, thompson ; guilt, hardship, conscious nurture, marathon, teaching ; choir, grave, psalm ; poetic, backyard
<b>Top words more frequent in inspiring than in other posts</b>	
share, people, time, make, work, world, support, thank, help , feel, journey, live , story, video, watch	hope, amaze, love, great, thank, learn, world, others ; feel, dream ; music, beautiful, write
<b>Most frequent bi-grams</b>	
black life, life matter, father day, family friend, thank everyone, happy birthday, hard work, feel free ; god bless, thank god	david brin, carl sagan ; awe inspire, uplift series, thank much make happy, feel well, lose weight, inspire work, inspire hope
<b>Most frequent tri-grams</b>	
sushant singh rajput ; black life matter, happy father day , stand chance win, wish happy birthday, weight loss journey , dream come true, thank family friend ; new york city	david brin uplift, martin luther king ; keep good work, make feel well want say thank, inspire many people, rise uplift war
<b>Most frequent hashtags</b>	
#georgefloyd ; #blacklivesmatter, #love, #motivation, #mentalhealth, #fathersday, #success, #goals, #giveaway, #fitness, #contest, #grateful, #womensupportingwomen, #entrepreneur, #staysafe ; #blessed, #catholic, #transformation, #life	#thursdaytruthful, #standup, #guidelines, #wecanbe, #thelift ; #balance, #inrecovery ; #art, #artist, #artistdiscovery
<b>Personalities ; Human values ; Spirituality ; Art and beauty</b>	

TABLE VI  
STATISTICS FROM INSPIRING POSTS IN S AND pR DATA. THE DATA IS LEMMATIZED.

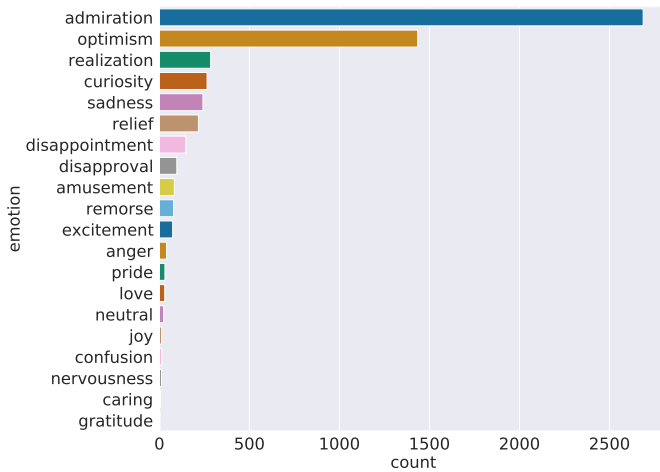


Fig. 5. Emotions predicted by fine-tuned RoBERTa [50] in inspiring pR posts.

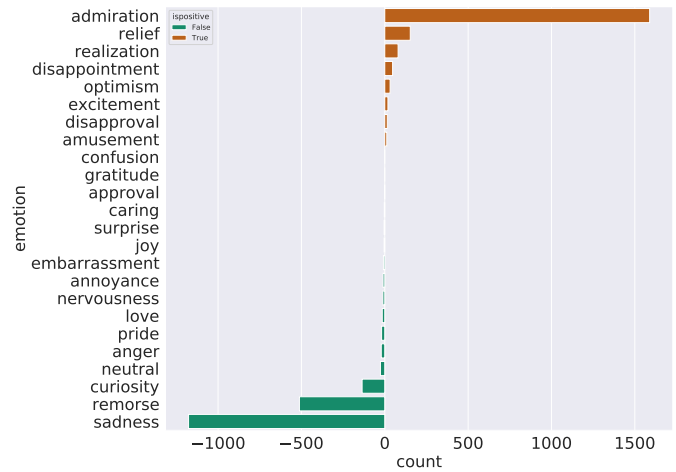


Fig. 6. Difference in frequency between emotions detected by RoBERTa [50] model in inspiring and non-inspiring pR posts. Red: emotions that are more prevalent in inspiring posts. Green: emotions that are more prevalent in non-inspiring posts.

## REFERENCES

- [1] T. M. Thrash and A. J. Elliot, "Inspiration as a psychological construct.," *Journal of personality and social psychology*, vol. 84 4, pp. 871–89, 2003.
- [2] T. M. Thrash and A. J. Elliot, "Inspiration: core characteristics, component processes, antecedents, and function.," *Journal of personality and social psychology*, vol. 87 6, pp. 957–73, 2004.
- [3] A. J. Elliot and T. M. Thrash, "Approach-avoidance motivation in personality: approach and avoidance temperaments and goals.," *Journal of personality and social psychology*, vol. 82, no. 5, p. 804, 2002.
- [4] S. B. Algoe and J. Haidt, "Witnessing excellence in action: The 'other-praising' emotions of elevation, gratitude, and admiration," *The journal of positive psychology*, vol. 4, no. 2, pp. 105–127, 2009.
- [5] T. Hart, "Inspiration: Exploring the experience and its meaning," *Journal of Humanistic Psychology*, vol. 38, no. 3, pp. 7–35, 1998.
- [6] V. C. Oleynick, T. M. Thrash, M. C. LeFev, E. G. Moldovan, and P. D. Kieffaber, "The scientific study of inspiration in the creative process: challenges and opportunities," *Frontiers in Human Neuroscience*, vol. 8, 2014.
- [7] T. M. Thrash, E. G. Moldovan, V. C. Oleynick, and L. A. Maruskin, "The psychology of inspiration," *Social and Personality Psychology Compass*,






Fig. 7. The polarity of all the inspiring posts, grouped into bins by sentiment score: from negative, to neutral and positive sentiment.

- vol. 8, no. 9, pp. 495–510, 2014.
- [8] M. Milyavskaya, I. Ianakieva, E. Foxen-Craft, A. Colantuoni, and R. Koestner, “Inspired to get there: The effects of trait and goal inspiration on goal progress,” 2012.
- [9] T. M. Thrash, L. A. Maruskin, S. E. Cassidy, J. W. Fryer, and R. M. Ryan, “Mediating between the muse and the masses: inspiration and the actualization of creative ideas,” *Journal of personality and social psychology*, vol. 98, no. 3, pp. 469–87, 2010.
- [10] T. M. Thrash, A. J. Elliot, L. A. Maruskin, and S. E. Cassidy, “Inspiration and the promotion of well-being: tests of causality and mediation,” *Journal of personality and social psychology*, vol. 98, no. 3, p. 488, 2010.
- [11] W. C. Belzak, T. M. Thrash, Y. Y. Sim, and L. M. Wadsworth, “Beyond hedonic and eudaimonic well-being: Inspiration and the self-transcendence tradition,” in *The happy mind: Cognitive contributions to well-being*, pp. 117–138, Springer, 2017.
- [12] T. Böttger, T. Rudolph, H. Evanschitzky, and T. Pfrang, “Customer inspiration: Conceptualization, scale development, and validation,” *Journal of Marketing*, vol. 81, no. 6, pp. 116–131, 2017.
- [13] J. W. Klein, T. I. Case, and J. Fitness, “Can the positive effects of inspiration be extended to different domains?,” *Journal of Applied Social Psychology*, vol. 48, no. 1, pp. 28–34, 2018.
- [14] J. Baumgartner, S. Zannettou, B. Keegan, M. Squire, and J. Blackburn, “The pushshift reddit dataset,” in *Proceedings of the International AAAI Conference on Web and Social Media*, vol. 14, pp. 830–839, 2020.
- [15] M. J. Smith, S. G. Figgins, M. Jewiss, and P. E. Kearney, “Investigating inspirational leader communication in an elite team sport context,” *International journal of sports science & coaching*, vol. 13, no. 2, pp. 213–224, 2018.
- [16] M. B. Oliver, A. A. Raney, M. D. Slater, M. Appel, T. Hartmann, A. Bartsch, F. M. Schneider, S. H. Janicke-Bowles, N. Krämer, M.-L. Mares, et al., “Self-transcendent media experiences: Taking meaningful media to a higher level,” *Journal of Communication*, vol. 68, no. 2, pp. 380–389, 2018.
- [17] A. Meier and S. Schäfer, “The positive side of social comparison on social network sites: How envy can drive inspiration on instagram,” *Cyberpsychology, Behavior, and Social Networking*, vol. 21, no. 7, pp. 411–417, 2018.
- [18] A. Meier, A. Gilbert, S. Börner, and D. Possler, “Instagram inspiration: How upward comparison on social network sites can contribute to well-being,” *Journal of Communication*, vol. 70, no. 5, pp. 721–743, 2020.
- [19] E. J. Noon and A. Meier, “Inspired by friends: Adolescents’ network homophily moderates the relationship between social comparison, envy, and inspiration on instagram,” *Cyberpsychology, Behavior, and Social Networking*, vol. 22, no. 12, pp. 787–793, 2019.
- [20] A. A. Raney, S. H. Janicke, M. B. Oliver, K. R. Dale, R. P. Jones, and D. Cox, “Profiling the audience for self-transcendent media: A national survey,” *Mass Communication and Society*, vol. 21, no. 3, pp. 296–319, 2018.
- [21] K. R. Dale, A. A. Raney, S. H. Janicke, M. S. Sanders, and M. B. Oliver, “Youtube for good: A content analysis and examination of elicitors of self-transcendent media,” *Journal of Communication*, vol. 67, no. 6, pp. 897–919, 2017.
- [22] D. Rieger and C. Klimmt, “The daily dose of digital inspiration: A multi-method exploration of meaningful communication in social media,” *New Media & Society*, vol. 21, no. 1, pp. 97–118, 2019.
- [23] K. R. Dale, A. A. Raney, Q. Ji, S. H. Janicke-Bowles, J. Baldwin, J. T. Rowlett, C. Wang, and M. B. Oliver, “Self-transcendent emotions and social media: Exploring the content and consumers of inspirational facebook posts,” *New Media & Society*, vol. 22, no. 3, pp. 507–527, 2020.
- [24] Q. Ji, A. A. Raney, S. H. Janicke-Bowles, K. R. Dale, M. B. Oliver, A. Reed, J. Seibert, and A. A. Raney, “Spreading the good news: Analyzing socially shared inspirational news content,” *Journalism & Mass Communication Quarterly*, vol. 96, no. 3, pp. 872–893, 2019.
- [25] Q. Ji and A. A. Raney, “Developing and validating the self-transcendent emotion dictionary for text analysis,” *PLoS one*, vol. 15, no. 9, p. e0239050, 2020.
- [26] C. Strapparava and R. Mihalcea, “Semeval-2007 task 14: Affective text,” in *SemEval@ACL*, 2007.
- [27] C. Strapparava and R. Mihalcea, “Learning to identify emotions in text,” in *SAC ’08*, 2008.
- [28] R. W. Picard, “Affective computing,” 1997.
- [29] L. A. M. Bostan and R. Klinger, “An analysis of annotated corpora for emotion classification in text,” in *COLING*, 2018.
- [30] “Crowdfunder,” <https://data.world/crowdfunder>.
- [31] V. A. Kharde and S. Sonawane, “Sentiment analysis of twitter data : A survey of techniques,” *ArXiv*, vol. abs/1601.06971, 2016.
- [32] S. M. Mohammad, F. Bravo-Marquez, M. Salameh, and S. Kiritchenko, “Semeval-2018 task 1: Affect in tweets,” in *SemEval@NAACL-HLT*, 2018.
- [33] D. Demszky, D. Movshovitz-Attias, J. Ko, A. S. Cowen, G. Nemade, and S. Ravi, “Goemotions: A dataset of fine-grained emotions,” *ArXiv*, vol. abs/2005.00547, 2020.
- [34] C. Liu, M. Osama, and A. de Andrade, “Dens: A dataset for multi-class emotion analysis,” *ArXiv*, vol. abs/1910.11769, 2019.
- [35] Y. Li, H. Su, X. Shen, W. Li, Z. Cao, and S. Niu, “Dailydialog: A manually labelled multi-turn dialogue dataset,” *ArXiv*, vol. abs/1710.03957, 2017.
- [36] A. H. Miller, W. Feng, D. Batra, A. Bordes, A. Fisch, J. Lu, D. Parikh, and J. Weston, “Parlai: A dialog research software platform,” *ArXiv*, vol. abs/1705.06476, 2017.
- [37] “Inspirobot,” <https://inspirobot.me/>.
- [38] M. Jia, Z. Wu, A. Reiter, C. Cardie, S. J. Belongie, and S. Lim, “Intentionomy: a dataset and study towards human intent understanding,” *CoRR*, vol. abs/2011.05558, 2020.
- [39] T. R. Hart, *Inspiration: An exploration of the experience and its role in healthy functioning*. PhD thesis, ProQuest Information & Learning, 1993.
- [40] P. Bojanowski, E. Grave, A. Joulin, and T. Mikolov, “Enriching word vectors with subword information,” *arXiv preprint arXiv:1607.04606*, 2016.
- [41] C. Cortes and V. Vapnik, “Support-vector networks,” *Machine Learning*, vol. 20, pp. 273–297, 2004.
- [42] A. Schofield and T. Davidson, “Identifying hate speech in social media,” *XRDS: Crossroads, The ACM Magazine for Students*, vol. 24, pp. 56–59, 2017.
- [43] “Toxic comment classification challenge,” <https://www.kaggle.com/c/jigsaw-toxic-comment-classification-challenge/data>.
- [44] T. Davidson, D. Warmesley, M. Macy, and I. Weber, “Automated hate speech detection and the problem of offensive language,” in *ICWSM*, 2017.
- [45] M. Sap, D. Card, S. Gabriel, Y. Choi, and N. A. Smith, “The risk of racial bias in hate speech detection,” in *ACL*, 2019.
- [46] T. Davidson, D. Bhattacharya, and I. Weber, “Racial bias in hate speech and abusive language detection datasets,” *ArXiv*, vol. abs/1905.12516, 2019.
- [47] M. Mozafari, R. Farahbakhsh, and N. Crespi, “Hate speech detection and racial bias mitigation in social media based on bert model,” *PLoS ONE*, vol. 15, 2020.
- [48] X. Zhou, M. Sap, S. Swayamdipta, N. A. Smith, and Y. Choi, “Challenges in automated debiasing for toxic language detection,” *ArXiv*, vol. abs/2102.00086, 2021.
- [49] J. L. Fleiss and J. Cohen, “The equivalence of weighted kappa and the intraclass correlation coefficient as measures of reliability,” *Educational and psychological measurement*, vol. 33, no. 3, pp. 613–619, 1973.



- [50] Y. Liu, M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer, and V. Stoyanov, "Roberta: A robustly optimized bert pretraining approach," *ArXiv*, vol. abs/1907.11692, 2019.
- [51] A. Joulin, E. Grave, P. Bojanowski, and T. Mikolov, "Bag of tricks for efficient text classification," in *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pp. 427–431, Association for Computational Linguistics, April 2017.
- [52] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "Bert: Pre-training of deep bidirectional transformers for language understanding," in *NAACL-HLT*, 2019.
- [53] J. E. Ramos, "Using tf-idf to determine word relevance in document queries," 2003.
- [54] J. Pennington, R. Socher, and C. D. Manning, "Glove: Global vectors for word representation," in *EMNLP*, 2014.
- [55] J. A. Hartigan and M. A. Wong, "A k-means clustering algorithm," 1979.
- [56] L. van der Maaten and G. E. Hinton, "Visualizing data using t-sne," 2008.
- [57] A. Y. Halevy, C. Canton-Ferrer, H. Ma, U. Ozertem, P. Pantel, M. Saeidi, F. Silvestri, and V. Stoyanov, "Preserving integrity in online social networks," *CoRR*, vol. abs/2009.10311, 2020.