

Automatic Identification and Classification of Bragging in Social Media

Mali Jin¹, Daniel Preoțiu-Pietro², A. Seza Doğruöz³, Nikolaos Aletras¹

¹University of Sheffield, ²Bloomberg, ³Ghent University

{mjin6, n.aletras}@sheffield.ac.uk

dpreotiucpie@bloomberg.net

as.dogruoz@ugent.be

Abstract

Bragging is a speech act employed with the goal of constructing a favorable self-image through positive statements about oneself. It is widespread in daily communication and especially popular in social media, where users aim to build a positive image of their persona directly or indirectly. In this paper, we present the first large scale study of bragging in computational linguistics, building on previous research in linguistics and pragmatics. To facilitate this, we introduce a new publicly available data set of tweets annotated for bragging and their types. We empirically evaluate different transformer-based models injected with linguistic information in (a) binary bragging classification, i.e., if tweets contain bragging statements or not; and (b) multi-class bragging type prediction including not bragging. Our results show that our models can predict bragging with macro F1 up to 72.42 and 35.95 in the binary and multi-class classification tasks respectively. Finally, we present an extensive linguistic and error analysis of bragging prediction to guide future research on this topic.¹

1 Introduction

The desire to be viewed positively is a key driver of human behavior (Baumeister, 1982; Leary and Kowalski, 1990; Sedikides, 1993; Tetlock, 2002) and creating a positive image often leads to personal rewards (Gilmore and Ferris, 1989; Hogan, 1982; Schlenker, 1980). Self-presentation strategies are means for individuals to build and establish this positive social image to meet their goals (Goffman et al., 1978; Jones et al., 1982; Jones, 1990; Bak et al., 2014a). Bragging (or self-praise) is one of the most common strategies and involves disclosing a positively valued attribute about the speaker or their in-group (Dayter, 2014, 2018).

Social media platforms tend to promote self-presentation tendencies (Chen et al., 2016) and allow users to craft an idealized self-image of themselves (Chou and Edge, 2012; Michikyan et al., 2015; Halpern et al., 2017). Self-presentation online is predominantly positive (Chou and Edge, 2012; Lee-Won et al., 2014; Matley, 2018). Furthermore, self-promotion is acceptable and even desired in certain online contexts (Dayter, 2018). This is also amplified by social media platforms through the presence of likes or positive reactions to users' posts (Reinecke and Trepte, 2014) which often are used to quantify impact on the platform (Lampos et al., 2014). Bragging in particular was found to be more frequent on social media than face-to-face interactions (Ren and Guo, 2020).

However, bragging is considered a high risk act (Brown and Levinson, 1987; Holtgraves, 1990; Van Damme et al., 2017) and can lead to the opposite effect than intended, such as dislike or decreased perceived competence (Jones et al., 1982; Sezer et al., 2018; Matley, 2018). It is, thus, paramount to understand the types of bragging and strategies to mitigate the face-threat introduced by bragging as well as how effective the self-presentation attempt is (Herbert, 1990). Table 1 shows examples of a non-bragging and bragging statements grouped in six types under a taxonomy that we propose in this paper based on previous linguistic research (Dayter, 2018; Matley, 2018).

Despite its pervasiveness and importance in online communication, bragging has yet to be studied at scale in computational (socio) linguistics. The ability to identify bragging automatically is important for: (a) linguists to better understand the context and types of bragging through empirical studies (Dayter, 2014; Ren and Guo, 2020); (b) social scientists to analyze the relationship between bragging and personality traits, online behavior and communication strategies (Miller et al., 1992; Van Damme et al., 2017; Sezer et al., 2018); (c) on-

¹Data is available here: https://archive.org/details/bragging_data

Type	Definition	Tweet
Achievement	Concrete outcome obtained as a result of the tweet author's actions. These may include accomplished goals, awards and/or positive change in a situation or status (individually or as part of a group).	<i>Finally got the offer! Whoop!!</i>
Action	Past, current or upcoming action of the user that does not have a concrete outcome.	<i>Guess what! I met Matt Damon today!</i>
Feeling	Feeling that is expressed by the user for a particular situation.	<i>Im so excited that I am back on my consistent schedule. I am so excited for a routine so I can achieve my goals!!</i>
Trait	A personal trait, skill or ability of the user.	<i>To be honest, I have a better memory than my siblings</i>
Possession	A tangible object belonging to the user.	<i>Look at our Christmas tree! I kinda just wanna keep it up all year!</i>
Affiliation	Being part of a group (e.g. family, fanclub, university, team, company etc.) and/or a certain location including living in a city, neighborhood or country.	<i>My daughter got first place in the final exam, so proud of her!</i>
Not Bragging	The tweet is not about bragging or (a) there is not enough information to determine that the tweet is about bragging; (b) the bragging statements belong to someone other than the author of the tweet; (c) the relationship between the author and people or things mentioned in the tweet are unknown.	<i>Glad to hear that! Well done Jim!</i>

Table 1: Bragging taxonomy together with type definitions and examples of tweets.

line users to enhance their self-presentation strategies (Miller et al., 1992; Dayter, 2018); (d) enhancing NLP applications such as intent identification (Wen et al., 2017) and conversation modeling.

In this paper, we aim to bridge the gap between previous work in pragmatics and the computational study of speech acts. Our contributions are:

- A new publicly available data set containing a total of 6,696 English tweets annotated with bragging and their types;
- Experiments with transformer-based models combined with linguistic features for bragging identification (binary classification) and bragging type classification (seven classes);
- A qualitative linguistic analysis of markers of bragging in tweets and the model behavior in predicting bragging.

2 Related Work

Bragging as a Speech Act Bragging as a speech act is considered a face-threatening act to positive face (i.e. the desire to be liked) under politeness theory (Brown and Levinson, 1987). It is directly oriented to the speaker and may threaten their likeability if the bragging is perceived negatively, while also may affect hearer's face by implying that their feelings are not valued by the speaker (Matley, 2018). Bragging online plays an important role in self-presentation and its pervasiveness challenges classic politeness theories, such as the modesty maxim (Leech, 2016) and the self-denigration

maxim (Gu, 1990). Thus, research in social psychology and linguistics has mostly focused on identifying the pragmatic strategies for bragging that mitigate face threat and their impact of likeability and perceived competence, which the speakers aim to increase with this self-presentation strategy.

Bragging Strategies Modest and sincere self-presentation styles are more likely to be perceived positively (Sedikides et al., 2007). Bragging framed as mere information-sharing, but with positive connotation to the speaker, can make the speaker be perceived as more likeable (Miller et al., 1992). It can also be perceived negatively and causes greater aggression when it involves boasting, elements of competitiveness, use of superlatives and explicit comparisons to others (Miller et al., 1992; Hoorens et al., 2012; Scopelliti et al., 2015; Matley, 2018). In addition, competence related statements are more likely to be negatively perceived than those based on warmth (e.g. the ability to form connections with others) (Van Damme et al., 2017). Common mitigation strategies include speaker's attempts to deny compliments, shifting focus to persons closely related to them, reframing bragging as praise from a third party, admitting the bragging act through disclaimers (e.g. using #brag) or expressing it as a complaint (Wittels, 2011; Sezer et al., 2018), question, narration or sharing (Dayter, 2018; Matley, 2018; Ren and Guo, 2020). The success of self-presentation strategies are also impacted by the social context (Tice et al., 1995) or speaker identity (Paramita and Septianto, 2021).

Analysis of Bragging Bragging has been studied in the context of a small ballet community (Dayter, 2014), a pick-up artist forum (Rüdiger and Dayter, 2020) and a small set of WhatsApp conversations (Dayter, 2018). On social media, Matley (2018) studied the functional use of hashtags (e.g. #brag, #humblebrag) in Instagram posts, Tobback (2019) examined bragging strategies on LinkedIn, Ren and Guo (2020) investigated bragging and its pragmatic functions in Chinese social media and Matley (2020) studied impact of mitigating bragging through irony showing that bragging was negatively perceived. However, all these studies rely on manual analyses of small data sets (e.g. <300 posts).

Speech Acts in NLP Speech acts have been studied in NLP with examples including politeness (Danescu-Niculescu-Mizil et al., 2013), complaints (Preoțiuc-Pietro et al., 2019; Jin and Aletras, 2020, 2021), humor (Yang et al., 2021), parody (Maronikolakis et al., 2020), irony (Bamman and Smith, 2015), deception (Chen et al., 2020) and self-disclosure (Bak et al., 2012; Levontin and Yom-Tov, 2017; Ravichander and Black, 2018). Self-disclosure is closer to bragging as it is related to revealing personal information about oneself. It is usually employed to improve or maintain relationships (Bak et al., 2012) as measured through conversation frequency (Bak et al., 2014b). On the other hand, bragging is about aspects that are positively valued by the audience with the goal of improving the speaker’s self-image. Bak et al. (2014a) aim to predict different levels of self-disclosure statements, from general to sensitive; while Wang et al. (2021) examine gender differences in self-promotion by Congress members on Twitter. Bragging also involves in some cases possessions (Chinnappa and Blanco, 2018).

3 Bragging Data

3.1 Bragging Definition & Types

Definition Bragging is a speech act which explicitly or implicitly attributes credit to the speaker for some *good* (e.g. possession, skill) that is positively valued by the speaker and their audience (Dayter, 2014). A bragging statement should clearly express what the author is bragging about.

Types We generalize and extend the bragging types based on the definitions by Dayter (2018) and Matley (2018). The former summarizes them as accomplishments and some aspects of self; while the

latter includes everyday achievements (e.g. cooking) and personal qualities. We divide the ‘some aspects of self’ category into two categories, namely ‘Possession’ and ‘Trait’ respectively. We also add an ‘Affiliation’ category for bragging involving a group to which the speaker belongs. In total, we consider six bragging types and a non-bragging category. Table 1 shows the definitions of each type.

Classification Tasks Given the taxonomy above, we define two classification tasks: (i) *binary* bragging prediction (i.e. if a tweet contains a bragging statement or not); and (ii) seven-way *multiclass* classification for predicting if a tweet contains one of the six bragging types or no bragging at all.

3.2 Data Collection

To the best of our knowledge, there is no other data set available for our study. We use Twitter for data collection as tweets are openly available for research and widely used in other related tasks, e.g. predicting sentiment (Rosenthal et al., 2017), affect (Mohammad et al., 2018), sarcasm (Bamman and Smith, 2015), stance (Mohammad et al., 2016).

Random Sampling We select tweets for annotation by randomly sampling from the 1% Twitter feed one day per month from January 2019 to December 2020 (approximately 10k tweets per day) to ensure diversity using the Premium Twitter Search API for academic research.²

Keyword-based Sampling To give a model access to more positive examples of bragging statements for training, we use a keyword-based sampling method that increases the hit rate of bragging, following previous work on labeling infrequent linguistic phenomena, e.g. irony (Mohammad et al., 2018) or hate speech (Waseem and Hovy, 2016).

We build queries based on indicators of positive self-disclosure (e.g. *I, just*) (Dayter, 2018) and stylistic indicators, e.g. positive emotion words, present tense verbs (Bazarova et al., 2013). As the frequency of these keywords is high, we construct multi-word queries consisting of a personal pronoun and an indicator. In addition, we use a short list of curated bragging-related hashtags.³ After annotating 1,000 tweets, we compute the percentage

²<https://tinyurl.com/2p8wnure>

³The queries are: {[*I, proud*], [*I, glad*], [*I, happy*], [*I, best*], [*I, amazed*], [*I, amazing*], [*I, excellent*], [*I, just*], [*I'm, proud*], [*I'm, glad*], [*I'm, happy*], [*I'm, best*], [*I'm, amazed*], [*I'm, amazing*], [*I'm, excellent*], [*me, proud*], [*my, best*], #brag, #bragging, #humblebrag, #humble, #braggingrights}.

Label	Training set	Dev/Test set	All
	(Keyword sampling)	(Random sampling)	
Binary			3382 (3.54%)
Bragging	544 (16.09%)	237 (7.15%)	781 (11.66%)
Not Bragging	2838 (83.91%)	3077 (92.85%)	5915 (88.34%)
Multi-class			
Achievement	166 (4.91%)	71 (2.14%)	237 (3.54%)
Action	127 (3.76%)	58 (1.72%)	185 (2.76%)
Feeling	39 (1.15%)	27 (0.82%)	66 (0.99%)
Trait	91 (2.69%)	48 (1.45%)	139 (2.08%)
Possession	58 (1.72%)	28 (0.84%)	86 (1.28%)
Affiliation	63 (1.86%)	5 (0.15%)	68 (1.01%)
Not Bragging	2838 (83.91%)	3077 (92.85%)	5915 (88.34%)
Total	3382	3314	6696

Table 2: Bragging data set statistics.

of bragging tweets for each keyword and remove from sampling tweets with less than 5% (i.e. [*I, amazed*], [*I’m, amazing*], [*I’m, best*], [*my, best*], [*I, excellent*], *#humble*).

We initially collected around 6K and 368K tweets using hashtags and multi-word queries respectively. We obtain over 9k tweets by keeping all tweets collected using hashtags and sample 1% from those collected using multi-word queries to balance the two types.

Data Filtering After collecting tweets, we exclude those with duplicate or no meaningful textual content (e.g. only @-mentions or images). We only focus on English posts and filter out non-English ones using the language code provided by Twitter. We also exclude retweets and quoted tweets, as these do not typically express the thoughts of the user who retweeted them. Moreover, we exclude 131 tweets containing a URL in the text because these were related to advertisements based on initial results from our annotation calibration rounds. This resulted in a total of 6,696 tweets which is of similar size with data sets recently released for social NLP (Oprea and Magdy, 2020; Chung et al., 2019; Beck et al., 2021; Mendelsohn et al., 2021).

3.3 Annotation and Quality Control Process

We manually annotate tweets for providing a solid benchmark and foster future research. All authors of the paper have significant experience in linguistic annotation. We run three calibration rounds of 100 tweets each, where all annotated all tweets and discussed disagreements, until a Krippendorff’s Alpha above 0.80 in the seven-class task was reached.

To monitor quality, a subset of 1,564 tweets were annotated by two annotators or more in case of disagreements. If a tweet fits into multiple bragging types, we assign the more prominent one.⁴ The

⁴For example, we annotate “*New car*✓*New crib*✓*New*

Class	Self-disclosure (%)	Non-self-disclosure (%)
Bragging	31.63	68.37
Non-bragging	24.04	75.96
Achievement	31.65	68.35
Action	27.57	72.43
Feeling	31.82	68.18
Trait	36.69	63.31
Possession	29.07	70.93
Affiliation	35.29	64.71
Non-bragging	24.04	75.96
Total	24.93	75.07

Table 3: Percentages of self-disclosure class across bragging classes

annotation is based only on the actual text of the tweet without considering additional modalities (e.g. images), context or replies. This is similar to the information available to predictive models during training. We selected the final label as the majority vote and a final label was assigned after consensus in cases of three different votes.⁵ The full task guidelines, examples and interface are presented in Appendix B.

The inter-annotator agreement between two annotations of all tweets is: (a) percentage agreement: 89.03; (b) Krippendorff’s Alpha (Krippendorff, 2011) (7-class): 0.840; (c) Krippendorff’s Alpha (binary): 0.786. Agreement values are between the upper part of the *substantial* agreement band and the *perfect* agreement band (Artstein and Poesio, 2008). The final data set consists of 6,696 tweets with one of the seven classes. Before annotation, the keyword-based and randomly sampled tweets were shuffled to not induce frequency bias. Data set statistics are shown in Table 2, including statistics across the two sampling strategies. The model performance curve by varying the training set size indicates that annotating more data is not likely to lead in substantial improvements in bragging prediction (see Figure 3 in Appendix).

3.4 Self-disclosure in Bragging

We conduct an analysis of the relationship between self-disclosure and bragging as they are closely related. We use self-disclosure lexicon by Bak et al. (2014a) to assign each tweet in our data set a label (i.e. self-disclosure or non-self-disclosure). The percentages of self-disclosure across each bragging type are shown in Table 3. We also used self-disclosure models as a predictor for bragging in

barbershop✓*20 years young*” as ‘Possession’ because bragging is mostly about possessions (*crib, car, barbershop*).

⁵We experimented on training models using the subset annotated by a single annotator compared to multiple annotators and find no significant differences (see Appendix A).

early experimentation but the results are omitted due to the low performance.

3.5 Data Splits

We use the keyword sampled data for training and the random data for development and testing (in the ratio of 2:8) because the latter is representative of the real distribution of tweets (see Table 2).

4 Predictive Models

We evaluate vanilla transformer-based models (Vaswani et al., 2017) and further leverage external linguistic information to improve them.

BERT, RoBERTa and BERTweet We experiment with Bidirectional Encoder Representations from Transformers (BERT; Devlin et al. (2019)), RoBERTa (Liu et al., 2019) and BERTweet (Nguyen et al., 2020). RoBERTa is a more robust variant of BERT that obtains better results on a wide range of tasks. BERTweet is pre-trained on English tweets using RoBERTa as basis and achieves better performance on Twitter tasks (Nguyen et al., 2020). We fine-tune BERT, RoBERTa and BERTweet for binary and multiclass bragging prediction by adding a classification layer that takes the [CLS] token as input.

BERTweet with Linguistic Features We inject linguistic knowledge that could be related to bragging to the BERTweet model with a similar method proposed by Jin and Aletras (2021),⁶ that was found to be effective on complaint severity classification, a related pragmatics task. The method is adapted from Rahman et al. (2020), which integrates multimodal information (e.g. audio, visual) in transformers using a fusion mechanism called Multimodal Adaption Gate (MAG). MAG integrates multimodal information to text representations in transformer layers using an attention gating mechanism for modality influence controlling. We first expand vectors of linguistic information to a comparable size to the embeddings fed to the pre-trained transformer. We, then, use MAG to concatenate contextual and linguistic representations after the embedding layer of the transformer similar to Rahman et al. (2020). The output is sent to a pre-trained BERTweet encoder for fine-tuning followed by an output layer.

We experiment with these linguistic features:

⁶Early experimentation with simply concatenating or applying attention resulted in lower performance.

- **NRC:** The NRC word-emotion lexicon contains a list of English words mapped to ten categories related to emotions and sentiment (Mohammad and Turney, 2013). We represent each tweet as a 10-dimensional vector where each element is the proportion of tokens belonging to each category.
- **LIWC:** Linguistic Inquiry and Word Count (Pennebaker et al., 2001) is a dictionary-based approach to count words in linguistic, psychological and topical categories. We use LIWC 2015 to represent each tweet as a 93-dimensional vector.
- **Clusters:** We use Word2Vec clusters proposed by PreoŃiu-Pietro et al. (2015) to represent each tweet as a 200-dimensional vector over thematic subjects.

5 Experimental Setup

Text Processing We pre-process text by lower-casing, replacing all username mentions with placeholder tokens @USER and emojis with words using demojize.⁷ We also remove hashtags that are used as keywords (e.g. #brag) in data collection. Finally, we tokenize the text using TweetTokenizer.⁸

Baselines

Majority Class: As a first baseline, we label all tweets with the label of the majority class.

LR-BOW: We train a Logistic Regression with bag-of-words using L2 regularization.

BiGRU-Att: We also train a bidirectional Gated Recurrent Unit (GRU) network (Cho et al., 2014) with self-attention (Tian et al., 2018). Tokens are first mapped to GloVe embeddings (Pennington et al., 2014) and then passed to a bidirectional GRU. Subsequently, its output is passed to a self-attention layer and an output layer for classification.

Hyperparameters For **BiGRU-Att**, we use 200-dimensional GloVe embeddings (Pennington et al., 2014) pre-trained on Twitter data. The hidden size is $h = 128$ where $h \in \{64, 128, 256, 512\}$ with dropout $d = .2$, $d \in \{.2, .5\}$. We use Adam optimizer (Kingma and Adam, 2015) with learning rate $l = 1e-2$, $l \in \{1e-3, 1e-2, 1e-1\}$. For **BERT**,

⁷<https://pypi.org/project/emojize/>

⁸<https://www.nltk.org/api/nltk.tokenize.html>

Model	Precision	Recall	Macro-F1	Precision	Recall	Macro-F1
	Bragging Classification (Binary)			Bragging and Type Classification (7 class)		
Majority Class	46.42	50.00	48.15	13.26	14.29	13.76
LR-BOW	54.53	63.16	52.68	18.52	20.02	18.59
BiGRU-Att	55.93 ± 1.53	51.41 ± 0.47	51.29 ± 1.40	18.32 ± 0.10	26.16 ± 3.41	19.19 ± 0.31
BERT	64.24 ± 1.40	65.91 ± 3.32	64.58 ± 0.80	24.16 ± 1.15	39.66 ± 4.84	26.85 ± 0.81
RoBERTa	66.53 ± 0.29	68.43 ± 2.05	67.34 ± 1.02	28.99 ± 0.61	45.90 ± 3.59	32.82 ± 0.65
BERTweet	70.43 ± 0.16	72.62 ± 0.89	71.44 ± 0.43	30.82 ± 0.75	47.25 ± 2.68	34.86 ± 0.79
BERTweet-NRC	72.89 ± 1.26	70.95 ± 0.96	71.80 ± 0.49	30.95 ± 0.54	47.98 ± 1.12	34.36 ± 0.19
BERTweet-LIWC	72.65 ± 0.20	72.21 ± 0.43	72.42 [†] ± 0.31	32.06 ± 2.42	46.68 ± 7.45	34.83 ± 0.79
BERTweet-Clusters	71.26 ± 2.27	72.53 ± 1.91	71.60 ± 0.21	32.51 ± 1.36	46.97 ± 2.36	35.95 ± 0.54

Table 4: Macro precision, recall and F1-Score (\pm std. dev. for 3 runs) for bragging prediction (binary and multiclass). Best results are in bold. [†] indicates significant improvement over BERTweet (t-test, $p < 0.05$).

RoBERTa and **BERTweet**, we use the base cased model (12 layers and 109M parameters, 12 layers and 125M parameters and 12 layers and 135M parameters accordingly) and fine-tune them with learning rate $l = 3e-6$, $l \in \{1e-4, 1e-5, 5e-6, 3e-6, 1e-6\}$. For **BERTweet with linguistic features**, we project these to vectors of size $l_{NRC} = 200$, $l_{LIWC} = 400$, $l_{Clusters} = 768$, $l \in \{10, 93, 200, 400, 600, 768\}$. For MAG, we use the default parameters from [Rahman et al. \(2020\)](#). For *multi-class classification*, we apply class weighting due to the imbalanced data and set the training epoch to $n = 40$, $n \in \{15, 20, 25, 30, 35, 40, 45, 50, 55, 60, \}$. The maximum sequence length is set to 50 covering 95% of tweets in the training set. We use a batch size of 32.

Training and Evaluation We train each model three times using different random seeds and report the mean Precision, Recall and F1 (macro). We apply early stopping during training based on the dev loss. The experiments with linguistic features are performed with the best pre-trained transformer in each of the two classification tasks.

6 Results

Binary Bragging Classification Table 4 (left) shows the predictive performance of all models on predicting bragging (i.e. binary classification). Overall, BERTweet models with linguistic information achieve better overall performance. Transformer models perform substantially above the *majority class baseline* (+23.29 F1) and above *Logistic Regression* (+18.76). *BERTweet* (71.44 F1) performs better than *BERT* (64.58 F1) and *RoBERTa* (67.34 F1), which illustrates the advantage of pre-training on English tweets for this task.

Performance is further improved (+0.98 F1) by using *LIWC features* alongside BERTweet, which indicates that injecting extra linguistic information

benefits bragging identification. We speculate that this is because a bragging statement usually contains particular terms (e.g. personal pronouns, positive terms) or involves at least one certain aspect or theme (e.g. reward or property), which can be captured by linguistic features (e.g. feature *I* and *ACHIEVE* in LIWC). Combining lexicons lead to worse results than using a single one, so we refrain from reporting these results for clarity.

Multi-class Bragging Classification Table 4 (right) shows the predictive performance of all models on multiclass bragging type prediction including not bragging. We again find that pre-trained transformers substantially outperform the *majority class baseline* (+21.1 F1) and *logistic regression* (+16.27 F1). In line with the binary results, we find that *BERTweet* (34.86 F1) performs best out of all transformers. *BERTweet-Clusters* outperforms all models (35.95 F1), which indicates that topical information helps to identify different types of bragging. Each bragging type might be particularly specialized to certain topics (e.g. *weight loss* in ‘Achievement’ category).

7 Analysis

Linguistic Feature Analysis We analyze the linguistic features i.e. unigrams, LIWC and part-of speech (POS) tags associated with bragging and its types in all tweets of our data set. For this purpose, we first tag all tweets using the Twitter POS Tagger ([Derczynski et al., 2013](#)). Each tweet is represented as a bag-of-words distribution over POS unigrams and bigrams to reveal distinctive syntactic patterns of bragging and their types. For each unigram, LIWC and POS feature, we compute correlations between its distribution across posts and the label of the post. Then, we use the method introduced by [Schwartz et al. \(2013\)](#) to rank the features using univariate Pearson correlation with

Bragging		Non-Bragging		Bragging type											
				Achievement		Action		Feeling		Trait		Possession		Affiliation	
Feature	r	Feature	r	Feature	r	Feature	r	Feature	r	Feature	r	Feature	r	Feature	r
Unigrams and LIWC															
AUTHENTIC	0.149	CLOUT	0.109	FOCUSPAST	0.200	get	0.146	happy	0.228	APOSTRO	0.197	own	0.211	FAMILY	0.276
my	0.127	YOU	0.089	Number	0.157	trip	0.128	POSEMOE	0.218	COGPROC	0.181	buy	0.175	CLOUT	0.271
I	0.122	DISCREP	0.078	Analytic	0.153	RELATIV	0.119	♥	0.191	FOCUSPRESENT	0.179	bought	0.149	proud	0.263
TONE	0.104	NEGEMO	0.077	finished	0.150	ready	0.114	blessed	0.190	cute	0.159	car	0.146	rights	0.215
FOCUSPAST	0.102	SOCIAL	0.076	3	0.133	him	0.114	AFFECT	0.184	PRONOUN	0.157	bedroom	0.144	SOCIAL	0.209
WC	0.100	FOCUSPRESENT	0.070	WORK	0.132	happen	0.105	feels	0.176	take	0.143	extra	0.144	amazing	0.205
RELATIV	0.090	INFORMAL	0.056	managed	0.130	FOCUSFUTURE	0.105	love	0.169	COMPARE	0.143	xr	0.142	♥	0.197
TIME	0.081	COGPROC	0.056	over	0.129	fun	0.102	sunrise	0.166	ANGER	0.138	macbook	0.055	law	0.185
during	0.078	ANGER	0.056	under	0.119	gave	0.097	weighted	0.162	I	0.137	new	0.139	team	0.182
ACHIEVE	0.075	just	0.054	beat	0.112	hours	0.096	july	0.159	if	0.137	afford	0.139	OTHERP	0.181
PREP	0.073	your	0.052	race	0.104	before	0.095	time	0.159	SWEAR	0.134	PERIOD	0.106	words	0.164
managed	0.072	IPRON	0.051	office	0.103	sitting	0.095	truly	0.156	am	0.133	HOME	0.105	teams	0.164
REWARD	0.069	?	0.043	possible	0.103	VERB	0.094	BIO	0.147	PPRON	0.132	DASH	0.084	#baseball	0.164
row	0.068	not	0.038	5	0.101	PREP	0.089	CERTAIN	0.143	me	0.130	I	0.077	fan	0.163
got	0.067	why	0.037	SIXLTR	0.100	INGEST	0.085	TONE	0.140	look	0.122	DISCREP	0.071	MALE	0.160
POS (Unigrams and Bigrams)															
PRP_VBD	0.104	NNP	0.081	CD_NNS	0.198	DT_NNP	0.139	RB_JJ	0.183	VBP	0.252	\$_CD	0.161	FW_.	0.164
VBD	0.093	VB	0.061	VBD	0.171	VBP_TO	0.124	VBP_IN	0.174	PRP	0.193	\$	0.130	VB_VBD	0.161
CD_NNS	0.077	RB_VB	0.056	CD	0.164	IN_.	0.117	VB_RBR	0.161	PRP_VBP	0.191	NN_PDT	0.130	CC_UH	0.159
PRPS	0.074	NNP_NNP	0.049	NNS	0.145	VBP_WP	0.116	JJR_WRB	0.161	VBP_JJ	0.162	NNS_UH	0.122	VBZ_DT	0.151
VBD_DT	0.062	VBP_PRP	0.048	VBD_DT	0.141	NNP_UH	0.116	RB_VBZ	0.146	UH_DT	0.150	SYM_.	0.114	DT_RBS	0.146
NN_IN	0.061	VBZ	0.039	PRP_VBD	0.132	NFP_NNP	0.116	CC_JJ	0.143	VBP_DT	0.150	VBZ_JJ	0.110	UH_NNP	0.145
IN_CD	0.060	MD	0.035	NN_IN	0.132	NNP	0.116	VBD_.	0.131	RB_VB	0.149	VB_PRPS	0.109	_SYM	0.138
IN_PRPS	0.060	NNP_VBZ	0.033	IN_CD	0.130	NNP_NNS	0.114	_VBG	0.123	MD	0.149	PRPS_JJ	0.109	NFP_CC	0.137
PRPS_NN	0.058	RB_RB	0.031	VBN	0.129	TO_VB	0.109	UH_WP	0.118	MD_VB	0.134	_VBD	0.109	PRP_PRPS	0.136
VBD_PRPS	0.057	MD_PRP	0.031	VB_JJR	0.109	TO	0.107	POS_RB	0.118	CC_WP	0.131	NN_PRPS	0.106	NN_NN	0.135

Table 5: Feature correlations including unigrams (lowercase), LIWC (uppercase), part-of-speech (POS) unigrams and bigrams with bragging and non-bragging tweets (left) and bragging tweets grouped in six types (right), sorted by Pearson correlation (r). All correlations are significant at $p < .01$, two-tailed t-test.

words normalized to sum up to unit for each tweet.

Table 5 (left) presents the top 15 features from unigrams (lowercase) and LIWC (uppercase) and top 10 features from POS unigrams and bigrams correlated with bragging and non-bragging tweets. We notice that the top words in the bragging category can be classified into (a) personal pronouns (e.g. *my*, *I*) that usually indicate the author of the bragging statement; (b) words related to time (e.g. *FOCUSPAST*, *TIME*, *during*); and (c) words related to a specific bragging target (e.g. *RELATIV*, *ACHIEVE*, *REWARD*, *managed*). These findings are in line with the indicators of positive self-disclosure by Dayter (2018) and Bazarova et al. (2013). Furthermore, personal pronouns followed by a verb in past tense (*PRP_VBD*) is common in bragging (e.g. *I forgot what it’s like to be good at school. Today I finished a thing we were doing so fast that everyone around me started asking ME for help instead of the prof :’)*)

Table 5 (right) presents the top 15 features from unigrams (lowercase) and LIWC (uppercase) correlated with bragging tweets grouped in six types. We observe that **Achievement** statements usually involve verbs that are in past tense or indicate a result (e.g. *FOCUSPAST*, *finished*, *beat*). A POS pattern common in **Achievement** statements is a cardinal number followed by nouns in plural (*CD_NNS*), similar to its unigram and LIWC features (*NUMBER*, *3*, *5*) (e.g. *I made a total of 5 dollars from*

online surveys wooo). It is worth noting that one of the prevalent LIWC features for **Action** is *FOCUS-FUTURE*. This is because the user may brag about a planned action (e.g. *@USER You know what? I’m going to make some PizzaRolls Brag*). Most of the top words in **Feeling** express emotion or sensitivity (e.g. *happy*, *blessed*), which is consistent with the top POS feature, *RB_JJ* (e.g. *absolutely chuffed, so happy*). In **Trait** category, words are mostly pronouns (e.g. *I*, *PRP*, *PRP_VBP*) and verbs (e.g. *VBP*, *VBP_JJ*). Words appear frequently in **Possession** category are actions related to purchase (e.g. *own*, *buy*) and nouns related to a tangible object (e.g. *car*, *bedroom*). In addition, users usually show off the value of their possessions using statements that involve currency signs (\$) or currency signs followed by a number (*\$_CD*) (e.g. *I just signed a new three-year contract and I’ll be getting 235 any-time minutes per month. Plus, the company is going to throw in a phone for just \$ 49 per month. I’ll bet you can’t beat that deal!*). Finally, top words in **Affiliation** category involve positive feeling towards belonging to a group (e.g. *proud*, *amazing*) and nouns related to it (e.g. *FAMILY*, *team*).

Bragging and Post Popularity We also analyze the association between bragging posts and the number of favorites/retweets they receive by other users. Similar to the previous linguistic feature analysis, we use univariate Pearson correlation to compute the correlations between the log-scaled fa-

Class	Mean	Median
Achievement	3.06	3.00
Action	0.91	0
Feeling	0.50	0
Trait	2.38	2.00
Possession	2.00	0.50
Affiliation	5.50	2.00

Table 6: Mean and median Twitter favorites across bragging classes on a sample set of the data.

avorites/retweets number of each tweet and its label (i.e. bragging or non-bragging) by controlling the numbers of followers and friends of the user who post the tweet. Our results show that the number of favorites is positively correlated with bragging (see Appendix Figure 5) while there is no correlation between bragging and the number of retweets.

We further explore the popularity of different bragging types. We randomly analyze a set of 443 tweets containing 56 bragging statements, where the follower and friend number of users are within a similar range: from 100 to 500 followers and from 500 to 1000 friends ($r = 0.19$, $p < .01$). We compute the mean and median Twitter favorites across the six bragging classes (see Table 6). We observe that bragging statements about *Affiliation* such as family members or sports teams are more likely to receive considerable amount of favorites with the mean of 5.5. For example, 14 users favorite the tweet *This maybe is a little, but I'm SO proud of my research group. We represent so many different personality types, cultures, ways of thinking, etc, and every single member of my lab (all 21 of them)*. We speculate this is because praising the group that one belongs to instead of oneself as a bragging strategy enables users be perceived as more likeable. Furthermore, bragging about *Achievement* is generally marked as favorite by other users with the median of 3, where bigger achievements in the content such as job offers may receive more favorites (e.g. tweet *Scored 80 % on my thesis. Rather proud of that given the circumstances: new baby; pandemic; late topic change due to lockdown; minimal uni support because of furloughs; and an international move.* was marked as favorite 15 times).

Class Confusion Analysis Figure 1 presents the confusion matrix of human agreement on seven classes normalized over the actual values (rows). We observe that *Non-bragging* (97%), *Achievement* (81%) and *Action* (78%) have high agreement, consistent with the class frequency. *Affiliation* (77%), *Possession* (76%) and *Trait* (72%) have compara-

ble percentages as these are easily associated with a bragging target or group. The *Feeling* category has the lowest percentage mostly caused by misclassification to the *Action* category. This is due to the fact that both types are not associated to a concrete outcome by definition, with the *feeling* class linked to a feeling linked to an action. Thus, it makes the boundary between bragging about the action or the feeling associated to the action more challenging to interpret. The next most frequent confusion is between *possession* and *achievement*, which usually arises when a tangible possession is involved and the annotators disagree if the author was bragging about the actual possession or the action that lead to the author obtaining that possession (e.g. *@USER I just got some stealth 300 easily the best headset I've ever had going from astro to turtle beach was a night and day difference*).

Figure 2 presents the confusion matrix between bragging type predictions from the best performing model, BERTweet-Clusters, on the multi-class classification task. First, we observe that the model is more likely to misclassify other classes as the dominant class, *Non-bragging*. Secondly, the most unambiguous classes are *Non-bragging* (87%) and *Achievement* (52%), which are in line with human agreement. Also, the model is good at identifying *Trait* (50%) and *Possession* (46%) due to the particular bragging targets (e.g. personalities, skills or tangible objects). Furthermore, we notice that the percentages of *Action* (31%) and *Feeling* (37%) are low. We speculate this is because they share more similarities with other classes (e.g. involving actions). This might also explain the high percentage of misclassified data points between *Action* and *Achievement*, *Feeling* and *Action*. Lastly, the model often confuses *Affiliation* with *Feeling* likely because the terms that express positive feelings (e.g. 'proud', ❤️) also appear frequently in *Affiliation* (see Table 5).

Error Analysis Finally, we perform an error analysis to examine the behavior and limitations of our best performing model (i.e. BERTweet-LIWC for binary classification and BERTweet-Clusters for multi-class classification) and identify pathways to improve the task modeling.

We first start with the binary bragging classification. We observe that non-bragging tweets containing positive sentiment are easy to be misclassified as bragging and even if such tweets involve something valued positively by authors, the purpose is




Figure 1: Confusion matrix of annotator agreement on seven bragging categories.

usually to express recommendation, compliment or appreciation to others:

T1: @USER paid for my **new bottle of vodka** & **I Love Her with all my heart** ❤️

Another frequent error happens when non-bragging tweets contain popular bragging targets such as achievement-oriented (e.g. weight loss, marathon) or possession-oriented (e.g. car, electronics):

T2: 4 spaces left on my budget **weight loss** program. £ 5 a week!???

Bragging often involves contextual understanding that goes beyond word use and require deep understanding of the context to determine the label. For example, common terms such as *first*, *finally*, *just* often appear in both non-bragging (T3) and bragging (T4) tweets:

T3: **just cleaned** my cats' toilets
 T4: It happened again! **I just completed** 30 minutes of meditation with @USER. **Just** sitting and resting in presence.

Models also fail to detect bragging mainly because it is indirect or there are no typical trigger terms, so they lean on pre-training to contextualize:

T5: 9 hr drives feel like nothing now lol

Some bragging statements use additional mitigation strategies, e.g. re-framing the bragging statement as irony, as a complaint or invoking praise from a third party:

T6: I find it strange how I was always the weird one in school and irl but online people think im cool for some reason

Finally, we highlight some representative examples of model confusion between bragging types. One example is when users' actions lead or not to a concrete result. In this example the model predicted *Action*, but the actual label is *Achievement*:




Figure 2: Confusion matrix of the best performing model on multi-class bragging classification, i.e. BERTweet-Clusters

T7: not to appropriate the gang escapes culture but me n my parents just did an escape room n actually got out?

Another example is an *Action* misclassified as *Possession*. This usually happens when a common phrase indicative of a certain type of bragging (*a new dish*) is invoked as part of an action:

T8: I had **a new dish** "egusi" it's so damn good! Love Nigerian food!

Other errors occur when multiple types of bragging are present (e.g. feeling and action) but the label expresses the more salient type, such as the *feeling* highlighted in this example:

T9: Literally **had the best time** with the girls last night, don't think I've drank that much in my life?

8 Conclusion

We presented the first computational approach to analyzing and modeling bragging as a speech act along with its types in social media. We introduced a publicly available annotated data set in English collected from Twitter. We experimented using transformer models combined with linguistic information on binary bragging and multiclass bragging type prediction. Finally, we presented an extensive analysis of features related to bragging statements and an error analysis of the model predictive behavior. In future work, we plan to study the extent to which bragging is used across various locations (Sánchez Villegas et al., 2020; Sánchez Villegas and Aletras, 2021) and languages and how it is employed by users across contexts.

Acknowledgements

We would like to thank Ari Silburt, Danae Sánchez Villegas, Yida Mu, and all the anonymous reviewers for their valuable feedback.

Ethics Statement

Our work has received approval from the Ethics Committee of the Department of Computer Science at the University of Sheffield (No 037572) and complies with Twitter’s data policy for research.⁹

References

- Ron Artstein and Massimo Poesio. 2008. Inter-coder agreement for Computational Linguistics. *Computational Linguistics*, 34(4):555–596.
- JinYeong Bak, Suin Kim, and Alice Oh. 2012. Self-disclosure and relationship strength in Twitter conversations. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 60–64, Jeju Island, Korea. Association for Computational Linguistics.
- JinYeong Bak, Chin-Yew Lin, and Alice Oh. 2014a. Self-disclosure topic model for classifying and analyzing Twitter conversations. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1986–1996.
- JinYeong Bak, Chin-Yew Lin, and Alice Oh. 2014b. Self-disclosure topic model for Twitter conversations. In *Proceedings of the Joint Workshop on Social Dynamics and Personal Attributes in Social Media*, pages 42–49, Baltimore, Maryland. Association for Computational Linguistics.
- David Bamman and Noah A Smith. 2015. Contextualized Sarcasm Detection on Twitter. In *Proceedings of the 9th International Conference on Weblogs and Social Media, ICWSM*, pages 574–577.
- Roy F Baumeister. 1982. A self-presentational view of social phenomena. *Psychological bulletin*, 91(1):3.
- Natalya N Bazarova, Jessie G Taft, Yoon Hyung Choi, and Dan Cosley. 2013. Managing Impressions and Relationships on Facebook: Self-Presentational and Relational Concerns Revealed Through the Analysis of Language Style. *Journal of Language and Social Psychology*, 32(2):121–141.
- Tilman Beck, Ji-Ung Lee, Christina Viehmann, Marcus Maurer, Oliver Quiring, and Iryna Gurevych. 2021. Investigating label suggestions for opinion mining in German covid-19 social media. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1–13, Online. Association for Computational Linguistics.
- Penelope Brown and Stephen C Levinson. 1987. *Politeness: Some universals in language usage*, volume 4. Cambridge University Press.
- Xi Chen, Gang Li, YunDi Hu, and Yujie Li. 2016. How Anonymity Influence Self-disclosure Tendency on Sina Weibo: An Empirical Study. *The anthropologist*, 26(3):217–226.
- Xi (Leslie) Chen, Sarah Ita Levitan, Michelle Levine, Marko Mandic, and Julia Hirschberg. 2020. Acoustic-prosodic and lexical cues to deception and trust: Deciphering how people detect lies. *Transactions of the Association for Computational Linguistics*, 8:199–214.
- Dhivya Chinnappa and Eduardo Blanco. 2018. Mining possessions: Existence, type and temporal anchors. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 496–505, New Orleans, Louisiana. Association for Computational Linguistics.
- Kyunghyun Cho, Bart van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. Learning phrase representations using RNN encoder–decoder for statistical machine translation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1724–1734, Doha, Qatar. Association for Computational Linguistics.
- Hui-Tzu Grace Chou and Nicholas Edge. 2012. “they are happier and having better lives than i am”: The impact of using facebook on perceptions of others’ lives. *Cyberpsychology, behavior, and social networking*, 15(2):117–121.
- Yi-Ling Chung, Elizaveta Kuzmenko, Serra Sinem Tekiroglu, and Marco Guerini. 2019. CONAN - COUNTER NARRATIVES THROUGH NICHE-SOURCING: A MULTILINGUAL DATASET OF RESPONSES TO FIGHT ONLINE HATE SPEECH. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2819–2829, Florence, Italy. Association for Computational Linguistics.
- Cristian Danescu-Niculescu-Mizil, Moritz Sudhof, Dan Jurafsky, Jure Leskovec, and Christopher Potts. 2013. A computational approach to politeness with application to social factors. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 250–259, Sofia, Bulgaria. Association for Computational Linguistics.
- Daria Dayter. 2014. Self-praise in microblogging. *Journal of Pragmatics*, 61:91–102.
- Daria Dayter. 2018. Self-praise online and offline: The hallmark speech act of social media? *Internet Pragmatics*, 1(1):184–203.
- Leon Derczynski, Alan Ritter, Sam Clark, and Kalina Bontcheva. 2013. Twitter Part-of-Speech Tagging for All: Overcoming Sparse and Noisy Data. In *Proceedings of the international conference recent advances*

⁹<https://developer.twitter.com/en/developer-terms/agreement-and-policy>

- in natural language processing ranlp 2013*, pages 198–206.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- David C Gilmore and Gerald R Ferris. 1989. The effects of applicant impression management tactics on interviewer judgments. *Journal of management*, 15(4):557–564.
- Erving Goffman et al. 1978. *The Presentation of Self in Everyday Life*, volume 21. Harmondsworth London.
- Yueguo Gu. 1990. Politeness phenomena in modern chinese. *Journal of pragmatics*, 14(2):237–257.
- Daniel Halpern, James E Katz, and Camila Carril. 2017. The online ideal persona vs. the jealousy effect: Two explanations of why selfies are associated with lower-quality romantic relationships. *Telematics and Informatics*, 34(1):114–123.
- Robert K Herbert. 1990. Sex-based differences in compliment behavior¹. *Language in society*, 19(2):201–224.
- Robert Hogan. 1982. A socioanalytic theory of personality. In *Nebraska symposium on motivation*. University of Nebraska Press.
- Thomas Holtgraves. 1990. The language of self-disclosure.
- Vera Hoorens, Mario Pandelaere, Frans Oldersma, and Constantine Sedikides. 2012. [The Hubris Hypothesis: You Can Self-Enhance, But You’d Better Not Show It](#). *Journal of personality*, 80(5):1237–1274.
- Mali Jin and Nikolaos Aletras. 2020. [Complaint identification in social media with transformer networks](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 1765–1771, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Mali Jin and Nikolaos Aletras. 2021. [Modeling the severity of complaints in social media](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2264–2274, Online. Association for Computational Linguistics.
- Edward E Jones. 1990. *Interpersonal perception*. WH Freeman/Times Books/Henry Holt & Co.
- Edward E Jones, Thane S Pittman, et al. 1982. Toward a general theory of strategic self-presentation. *Psychological perspectives on the self*, 1(1):231–262.
- Diederik P Kingma and Jimmy Ba Adam. 2015. Adam: A Method for Stochastic Optimization. *Optimization*. In, *ICLR*, 5.
- Klaus Krippendorff. 2011. Computing Krippendorff’s Alpha-Reliability.
- Vasileios Lampos, Nikolaos Aletras, Daniel Preotiuc-Pietro, and Trevor Cohn. 2014. [Predicting and characterising user impact on Twitter](#). In *Proceedings of the 14th Conference of the European Chapter of the Association for Computational Linguistics*, pages 405–413, Gothenburg, Sweden. Association for Computational Linguistics.
- Mark R Leary and Robin M Kowalski. 1990. Impression management: A literature review and two-component model. *Psychological bulletin*, 107(1):34.
- Roselyn J Lee-Won, Minsun Shim, Yeon Kyoung Joo, and Sung Gwan Park. 2014. Who puts the best “face” forward on facebook?: Positive self-presentation in online social networking and the role of self-consciousness, actual-to-total friends ratio, and culture. *Computers in Human Behavior*, 39:413–423.
- Geoffrey Leech. 2016. *Principles of pragmatics*. Routledge.
- Liat Levontin and Elad Yom-Tov. 2017. Negative self-disclosure on the web: the role of guilt relief. *Frontiers in psychology*, 8:1068.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. RoBERTa: A Robustly Optimized BERT Pretraining Approach. *arXiv preprint arXiv:1907.11692*.
- Antonis Maronikolakis, Danae Sánchez Villegas, Daniel Preotiuc-Pietro, and Nikolaos Aletras. 2020. [Analyzing political parody in social media](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4373–4384, Online. Association for Computational Linguistics.
- David Matley. 2018. [“this is Not a# humblebrag, this is just a# brag”](#): The pragmatics of self-praise, hashtags and politeness in Instagram posts. *Discourse, context & media*, 22:30–38.
- David Matley. 2020. [Isn’t working on the weekend the worst?# humblebrag”](#): The impact of irony and hashtag use on the perception of self-praise in instagram posts.
- Julia Mendelsohn, Ceren Budak, and David Jurgens. 2021. [Modeling framing in immigration discourse on social media](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2219–2263, Online. Association for Computational Linguistics.

- Minas Michikyan, Jessica Dennis, and Kaveri Subrahmanyam. 2015. Can You Guess Who I Am? Real, Ideal, and False Self-Presentation on Facebook among Emerging Adults. *Emerging Adulthood*, 3(1):55–64.
- Lynn Carol Miller, Linda Lee Cooke, Jennifer Tsang, and Faith Morgan. 1992. Should I Brag? Nature and Impact of Positive and Boastful Disclosures for Women and Men. *Human Communication Research*, 18(3):364–399.
- Saif Mohammad, Svetlana Kiritchenko, Parinaz Sobhani, Xiaodan Zhu, and Colin Cherry. 2016. Semeval-2016 task 6: Detecting Stance in Tweets. In *Proceedings of the 10th International Workshop on Semantic Evaluation (SemEval-2016)*, *SEM, pages 31–41.
- Saif M. Mohammad, Felipe Bravo-Marquez, Mohammad Salameh, and Svetlana Kiritchenko. 2018. Semeval-2018 Task 1: Affect in tweets. In *Proceedings of International Workshop on Semantic Evaluation (SemEval-2018)*, *SEM, pages 1–17.
- Saif M Mohammad and Peter D Turney. 2013. Crowdsourcing a Word–Emotion Association Lexicon. *Computational intelligence*, 29(3):436–465.
- Dat Quoc Nguyen, Thanh Vu, and Anh Tuan Nguyen. 2020. BERTweet: A pre-trained language model for English tweets. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 9–14, Online. Association for Computational Linguistics.
- Silviu Oprea and Walid Magdy. 2020. iSarcasm: A dataset of intended sarcasm. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1279–1289, Online. Association for Computational Linguistics.
- Widya Paramita and Felix Septianto. 2021. The benefits and pitfalls of humblebragging in social media advertising: the moderating role of the celebrity versus influencer. *International Journal of Advertising*, pages 1–24.
- James W Pennebaker, Martha E Francis, and Roger J Booth. 2001. Linguistic Inquiry and Word Count: LIWC 2001. *Mahway: Lawrence Erlbaum Associates*, 71(2001):2001.
- Jeffrey Pennington, Richard Socher, and Christopher D Manning. 2014. GloVe: Global Vectors for Word Representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 1532–1543.
- Daniel Preoțiu-Pietro, Mihaela Gaman, and Nikolaos Aletras. 2019. Automatically identifying complaints in social media. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5008–5019, Florence, Italy. Association for Computational Linguistics.
- Daniel Preoțiu-Pietro, Svitlana Volkova, Vasileios Lampos, Yoram Bachrach, and Nikolaos Aletras. 2015. Studying User Income through Language, Behaviour and Affect in Social Media. *PLoS one*, 10(9):e0138717.
- Wasifur Rahman, Md Kamrul Hasan, Sangwu Lee, Amir Zadeh, Chengfeng Mao, Louis-Philippe Morency, and Ehsan Hoque. 2020. Integrating Multimodal Information in Large Pretrained Transformers. In *Proceedings of the conference. Association for Computational Linguistics. Meeting*, volume 2020, page 2359. NIH Public Access.
- Abhilasha Ravichander and Alan W. Black. 2018. An empirical study of self-disclosure in spoken dialogue systems. In *Proceedings of the 19th Annual SIGdial Meeting on Discourse and Dialogue*, pages 253–263, Melbourne, Australia. Association for Computational Linguistics.
- Leonard Reinecke and Sabine Trepte. 2014. Authenticity and well-being on social network sites: A two-wave longitudinal study on the effects of online authenticity and the positivity bias in sns communication. *Computers in Human Behavior*, 30:95–102.
- Wei Ren and Yaping Guo. 2020. Self-praise on Chinese social networking sites. *Journal of Pragmatics*, 169:179–189.
- Sara Rosenthal, Noura Farra, and Preslav Nakov. 2017. SemEval-2017 Task 4: Sentiment analysis in Twitter. In *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*, *SEM, pages 502–518.
- Sofia Rüdiger and Daria Dayter. 2020. Manbragging online: Self-praise on pick-up artists’ forums. *Journal of Pragmatics*, 161:16–27.
- Danae Sánchez Villegas and Nikolaos Aletras. 2021. Point-of-interest type prediction using text and images. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 7785–7797, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Danae Sánchez Villegas, Daniel Preoțiu-Pietro, and Nikolaos Aletras. 2020. Point-of-interest type inference from social media text. In *Proceedings of the 1st Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 10th International Joint Conference on Natural Language Processing*, pages 804–810, Suzhou, China. Association for Computational Linguistics.
- Barry R Schlenker. 1980. *Impression management*, volume 222.
- H Andrew Schwartz, Johannes C Eichstaedt, Margaret L Kern, Lukasz Dziurzynski, Stephanie M Ramones, Megha Agrawal, Achal Shah, Michal Kosinski, David Stillwell, and Martin EP Seligman. 2013.

- Personality, Gender, and Age in the Language of Social Media: The Open-vocabulary Approach. *PLoS ONE*, 8(9).
- Irene Scopelliti, George Loewenstein, and Joachim Vosgerau. 2015. You call it “self-exuberance”; i call it “bragging” miscalibrated predictions of emotional responses to self-promotion. *Psychological science*, 26(6):903–914.
- Constantine Sedikides. 1993. Assessment, enhancement, and verification determinants of the self-evaluation process. *Journal of personality and social psychology*, 65(2):317.
- Constantine Sedikides, Aiden P. Gregg, and Claire M. Hart. 2007. The importance of being modest. *The Self: Frontiers in Social Psychology*, edited by Constantine Sedikides and Stephen J. Spencer, 163(84).
- Ovul Sezer, Francesca Gino, and Michael I Norton. 2018. Humblebragging: A distinct—and ineffective—self-presentation strategy. *Journal of Personality and Social Psychology*, 114(1):52.
- Philip E Tetlock. 2002. Social functionalist frameworks for judgment and choice: intuitive politicians, theologians, and prosecutors. *Psychological review*, 109(3):451.
- Zhengxi Tian, Wenge Rong, Libin Shi, Jingshuang Liu, and Zhang Xiong. 2018. Attention Aware Bidirectional Gated Recurrent Unit Based Framework for Sentiment Analysis. In *International Conference on Knowledge Science, Engineering and Management*, pages 67–78. Springer.
- Dianne M Tice, Jennifer L Butler, Mark B Muraven, and Arlene M Stillwell. 1995. When modesty prevails: Differential favorability of self-presentation to friends and strangers. *Journal of personality and social psychology*, 69(6):1120.
- Els Toubback. 2019. Telling the world how skilful you are: Self-praise strategies on LinkedIn. *Discourse & Communication*, 13(6):647–668.
- Carolien Van Damme, Eliane Deschrijver, Eline Van Geert, and Vera Hoorens. 2017. When Praising Yourself Insults Others: Self-Superiority Claims Provoke Aggression. *Personality and Social Psychology Bulletin*, 43(7):1008–1019.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention Is All You Need. In *Advances in neural information processing systems*, pages 5998–6008.
- Jun Wang, Kelly Cui, and Bei Yu. 2021. Self Promotion in US Congressional Tweets. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4893–4899, Online. Association for Computational Linguistics.
- Zeeraq Waseem and Dirk Hovy. 2016. Hateful symbols or hateful people? predictive features for hate speech detection on Twitter. In *Proceedings of the NAACL Student Research Workshop*, pages 88–93, San Diego, California. Association for Computational Linguistics.
- Liyun Wen, Xiaojie Wang, Zhenjiang Dong, and Hong Chen. 2017. Jointly Modeling Intent Identification and Slot Filling with Contextual and Hierarchical Information. In *National CCF Conference on Natural Language Processing and Chinese Computing*, pages 3–15. Springer.
- Harris Wittels. 2011. Humblebrag hall of fame. *Grantland.com*.
- Zixiaofan Yang, Shayan Hooshmand, and Julia Hirschberg. 2021. CHoRaL: Collecting humor reaction labels from millions of social media users. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 4429–4435, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.

A Impact of Multiple Annotations

Table 7 shows the performance of binary bragging classification of the best performing model (BERTweet-LIWC) on two different subsets of the test data: one annotated by a single annotator (2,130 tweets) and the other annotated by two or more annotators until consensus is reached (522 tweets). The results show that the same model tested on the two different subsets of test data lead to similar results. This shows there is no quantitative difference between the data sets annotated by two or more annotators when compared to a single annotator.

Data set	Precision	Recall	Macro-F1
Single Annotation	73.81	71.78	72.74
Multiple Annotations	68.24	83.31	73.23
Entire set	72.92	72.81	72.86

Table 7: Precision, Recall and macro F1-Score obtained by the same best performing model (BERTweet-LIWC) for binary classification on two different subsets of training data, annotated either by a single annotator or by multiple annotators.




Figure 3: Learning curve for performance across each bragging type.

B Guidelines and Annotation Interface

Thank you for your participation in our study. During our experiment, we will ask you to read and evaluate a tweet which may include a bragging or a praisal statement.

Instructions You need to identify whether or not a tweet includes a bragging statement.

Bragging Bragging is a speech act which explicitly or implicitly attributes credit to the speaker for some ‘good’ (possession, accomplishment, skill, etc.) which is positively valued by the speaker and the potential audience. As such, bragging includes announcements of accomplishments, explicit positive evaluations of some aspect of self and other



Figure 4: Screenshot of annotation interface on our platform.

types defined below. A bragging statement should clearly express what the author is bragging about (i.e. the target of bragging).

If the tweet is about bragging, decide on the category where the tweet belongs to from the following categories:

Achievement The act of bragging is about a concrete outcome obtained as a result of the tweet author’s actions. These results may include achievements, awards, products, and/or positive change in a situation or status (individually or as part of a group).

Examples:

- *Finally got that offer! Whoop!!*
- *Our team won the championship*

Action The act of bragging is about a past, current or upcoming action of the user that does not have a concrete outcome

Examples:

- *Hanging at Buffalo Wild Wings with @user for the #ILLvsASU game. #BraggingRights*
- *Guess what! I met Matt Damon today!*

Feeling The act of bragging is about a feeling that is expressed by the user for a particular situation.

Example:

- *Im so excited that I am back on my consistent schedule. I am so excited for a routine so I can achieve my goals!!*

Trait The act of bragging is about a personal trait, skill or ability of the user .

Examples:

- *To be honest, I have a better memory than my siblings*
- *I look great after losing weight*

Possession The act of bragging is about a tangible object belonging to the user.

Example:

- *Look at our Christmas tree! I kinda just wanna keep it up all year!*

Affiliation The act of bragging is about being part of a group (e.g. family, team, org etc.) and/or a certain location including living in a city, neighborhood or country, enrolled into a university, supporting a team, working in a company etc.

Example:

- *My daughter got first place in the final exam, so proud of her!*

Not bragging If the tweet is not about bragging, then select "No. This is not a bragging statement."

Examples:

- *One of the best books I've ever read*
- *hahahahahaha*
- *You gotta admit, that's some mighty awesome aim!*
- *Vote in the poll below for your book of choice!*
- *I think this is great*
- *dear everyone announcing they are at "Friends-giving", we get it, you have friends*
- *In case you didn't know, Adam Silver is in charge*
- *I feel terrible*
- *I don't know why you are celebrating*
- *This is exactly what is going on!*
- *I love you*

Select "No. This is not a bragging statement", also in cases when:

- there is not enough information to determine that the tweet is about bragging
- the bragging statements belong to someone other than the author of the tweet
- the relationship between author and people/things mentioned in the tweet are unknown:
 - *This kid is smart*
 - *That was an amazing stream*
 - *Kudos to mike Dunleavy! It's hard to get a franchise record ANYTHING in Chicago*
- the post is about the act of bragging:
 - *We want to hear you brag!*
 - *Trump isn't Bragging anymore as his trade-war hits the stockmarket hard*
 - *Dudes are getting too cocky these days. Them lil labels and that dar don't impress everyone.*




Figure 5: Pearson correlation between Twitter favorite number and bragging by controlling the number of followers and friends. All correlations are significant at $p < .01$, two-tailed t-test.

brag differently

Not available Finally, if the tweet is not available or displayed, or is in a language other than English, please select the "Not available" option.

Other considerations Please verify the content of hashtags as these may give clues towards the category of the tweet. The judgment should be made only based on the given content of the tweet - please do not search the tweet on Twitter or online in order to identify additional context.