Orthography in social media: Pragmatic and prosodic interpretations of caps lock

Maria Heath[*]

Abstract. Orthography in social media is largely understudied, but rich in pragmatic potential. This study examines the use of "caps lock" on Twitter, which has been claimed to function as an emotive strengthener. In a survey asking participants to rate tweets on gradient scales of emotion, I show that this claim does not account for all the data. I instead propose that caps lock should be understood as an indicator of prosody in text. I support this theory by drawing on Twitter corpus data to show how users employ single-word capitalization in positions indicative of emphatic stress and semantic focus. A prosodic interpretation of capitalization accounts for all the data in a unified way.

Keywords. orthography; capitalization; Twitter; social media; prosody; corpus; computer mediated communication

1. Introduction. Non-standard orthography on social media is a frequently referenced but largely under-studied phenomenon. Social media users deviate from the norms of English orthography in numerous ways, particularly in the areas of capitalization, punctuation, spelling, and use of emoji. Capitalization alone exhibits several recognizably distinct patterns of use, a few of which are exemplified in (1) through (4) below.

(1) CUTEST DOG EVER[1]
(2) my skin has been SO disrespectful lately, I feel so betrayed.
(3) How Are You So Fucking Hawwwwt!!!??
(4) also i'M S oOO exCITE FOR TAEYeon aND dean

I postulate that non-standard orthography is meaningful and aim to characterize a few of its uses.

Sociolinguistic research has become a popular topic in internet linguistics in recent years (e.g., Baym 2015, Seargeant & Tagg 2014, Zappavigna 2012), but it often takes common assumptions about orthographic patterns as given and builds upon them without further examination. There is also a range of research in computational linguistics which deals with internet language (e.g., Carvalho et al 2009, Gimpel et al 2011), and this research attempts to categorize and analyze social media language in purely quantifiable terms, making it computationally readable. But just as in sociolinguistics, this field makes no attempt at a rigorous analysis of the pragmatic uses of innovative orthography. It is clear that such an analysis needs to be developed, as a foundation upon which sociolinguists and computational linguists alike may build with confidence, knowing that they have a firm understanding of the full range of pragmatic information that may be carried by non-standard orthographic patterns.

As a preliminary step toward filling in this gap, this study looks at two related patterns of non-standard capitalization found on Twitter and examines how they are used. The first pattern in question is that of using capitals for the entire or a large portion of the tweet, as in example (1) above, which I term "all caps". The stereotyped meaning for this pattern is one of anger or

[1] Tweets taken from my corpus have been anonymized by removal of tags and usernames. I have also removed retweet markers and hyperlinks for clarity of data in this paper.

yelling, but some descriptive work on Twitter suggested that this pattern actually serves to upscale the emotive content of the message (Zappavigna 2012). My initial survey eliciting judgements on the emotive strength of tweets with and without the capitalization pattern in question indicates that there is indeed an upscaled strength judgement in happiness, but no such boost in anger or sadness. This data cannot be explained by previous theories, so I propose that this upscaling function derives instead from writers' attempts to write upscaled prosody. I further test this theory by looking at other linguistic uses of prosody, such as focus. The capitalization pattern found in example (2) above, which I refer to as "isolated" or "single-word" capitalization, is an example of how writers use capitalization in contexts which would be marked with intonational focus in speech.[2] Assuming this prosodic explanation, the data for both patterns can be clearly motivated, and we can make predictions about the behavior of other related patterns.

In the next section, I will lay the theoretical background and introduce the prosodic framework necessary to my analysis. Section 3 will outline my initial hypothesis about the interpretation of capitalization and describe the reader judgement experiment aimed at testing it. Since the results do not uphold the initial hypothesis, I turn to the prosodic analysis, which I supplement in section 4 with a qualitative corpus study on focus. Section 5 is a discussion of how the prosodic analysis accounts for both sets of data, and addresses additional variables for future consideration, and section 6 concludes.

2. Background. According to Crystal (2011), the systematic nature of most orthographic deviations on social media sites suggests that the observed range of variation is not due to carelessness or typos, but should rather be analyzed as stylistic (and, I will argue, linguistic) choices. In other words, these deviations are intentional, and therefore meaningful in some way. One lens through which to understand this idea is through Sperber and Wilson's Relevance Theory (2004), which states that it is human instinct to search for relevance in communication, and that communicators can and do take advantage of this. They posit that:

> "as a result of constant selection pressure towards increasing efficiency, the human cognitive system has developed in such a way that our perceptual mechanisms tend automatically to pick out potentially relevant stimuli, our memory retrieval mechanisms tend automatically to activate potentially relevant assumptions, and our inferential mechanisms tend spontaneously to process them in the most productive way" (Sperber & Wilson 2004:254).

In this way, a writer may utilize non-standard orthography as a stimulus, knowing that it will attract the attention of the reader, who will then pick it out as relevant and assign it a meaning. This assigned meaning will be strongly influenced by contextual clues and may additionally require one to draw on other assumptions and competencies.

In this paper, I propose that the interpretation that is assigned to the all caps pattern is prosodic in nature, and suggest that the specifics of the prosodic interpretation are determined by both the orthographic and pragmatic context. I further suggest that the meaning contributed by this pattern can be either paralinguistic or semantic in nature. To understand this claim, we must first understand some of the ways that prosody can be used in English speech and which prosodic features are manipulated to achieve these meanings.

Prosodic studies have shown that heightened states of emotion typically involve physiological changes which result in prosodic changes in speech. Heightened emotional states

---

[2] Analysis of patterns such as those found in examples (3) and (4), word-initial capitalization and alternating capitalization respectively, will be left to future work.

will, for example, lead to increases in pitch, volume, or speed of speech, changes which cue the listener in to the change in emotional state of the speaker (Piotrovskaya 2003, Cruttenden 1997). Though largely unconscious, these changes are widely recognized and understood, and can be manipulated by speakers to affect a false emotional state, as in acting or sarcasm. In addition to the literal content of the utterance, these prosodic features give the listener more information about the emotion being expressed and its degree.

In addition to heightened emotion, another common use of prosody in English is to indicate linguistic stress or emphasis. Stress can be indicated prosodically in English with either increases in pitch, length, or volume, pitch being the most prominently used in English (Cruttenden 1997). A word that is intended to be stressed should have at least one of these three features. Prosodic stress, unlike heightened emotion, is often a linguistic use of prosody, encoding semantic focus. Nevertheless, there seems to be a lot of overlap between stress prosody and the prosodic realization of heightened emotional states, pitch and volume in particular being prominently used in both cases.

It is generally accepted that when reading, many (though not all) people experience a "voice in their head" which applies a default prosody to the text. Fodor (2002) shows that, although this is not a phenomenon which can be studied directly, through its observable effects on language processing and comprehension we can get a glimpse of what people experience "phonetically" in their head when reading.

If a writer wishes to orthographically indicate that a deviation from the default prosody should be applied to a text, such as heightened pitch or volume, how can this be achieved? Relevance Theory suggests any deviations from the standard or most basic form will trigger an assumption of relevance, and that our "inferential mechanisms" will process these deviations using contextual evidence to build the most productive possible interpretation (Sperber and Wilson 2004). Since, according to Cruttenden (1997), pitch and volume are suprasegmental prosodic features, a deviating orthographic feature with the ability to trigger such a prosodic representation would have to be able to span multiple segments.

Capitalization is a great candidate, since it is both suprasegmental and easily accessible on all platforms (unlike bolding, for example). Context can serve to further disambiguate the prosodic interpretation. For example, pitch indicating heightened emotion is often suprasegmental across multiple words, but focus pitch is only heightened on one word or syllable, though phrasal focus interpretations can follow. So, while a longer string of all caps may indicate the former, a single word in all caps seems more likely to be associated with the latter.

The kind of orthography-to-prosody correspondence which I am proposing has precedence in other patterns with a more intuitive one-to-one relationship between written and spoken form, such as lengthening. In example (5) below, the reduplication of letters is analogous to increased prosodic length of the corresponding vowels (Lamontagne & McCulloch 2017).

(5) Im sooooooo hungryyy

The analogy between capitalization and increased pitch or volume may not be as intuitive, but even arbitrary choices may standardize over time. English spelling conventions were shown in Berg and Aranoff (2017) to have undergone self-standardization through popular usage trends, and the internet offers ample opportunity for increasingly rapid standardization. For example, Lamontagne and McCulloch (2017) have shown that even orthographic lengthening does not strictly follow phonological rules, but rather written conventions. Consider example (6), in which

the duplicated letter is silent, but follows the standard convention of final-letter duplication, regardless of pronounceability.

(6)   RT @user: "You're the bombbbbbb" haha

With the popularization of social media and the sheer number of users participating in online conversation daily, social media has become a mecca of rapid language change and standardization. We should expect to see innovative new orthographic trends arising from this cultural phenomenon, but we need not look far to discover the motivation for change when prosody is such a ripe candidate for incorporation into written language. With these observations in mind, we can now turn to looking at explicit evidence for the prosodic role of capitalization in social media.

3. All caps reader judgment experiment. For my first study, I set out to test reader interpretations of the all caps pattern in tweets.[3] This pattern specifically is noteworthy for having prevalent stereotyped meanings associated with it. Typing in all caps may get you accused of yelling, or "flaming",[4] and may be interpreted as an indicator that you are angry. Example (5) below shows that this interpretation fits some cases of all caps use. A detailed examination of Twitter corpus data tells a different story, however. Consider examples (8)-(11):

(7)   THERE ARE ALREADY LAWS ON THE BOOKS CONCERNING GUNS BUT THEY
        ARE NOT BEING ENFORCED! MORONS!
(8)   IM OFFICIALLY IN CALIFORNIA AND IT IS BEAUTIFUL
(9)   I NEED TO DO THIS
(10)   YOU GUYS ARE TOO NICE
(11)   LISTENING TO THE SMITHS HAS MADE ME MOODY AND EMO.

Based on the literal content of these tweets, it would be difficult to interpret their messages as being angry or as yelling. A range of possible emotions are represented, from happiness to desire to something akin to dramatic overacting. Clearly anger is not the only meaning that can be associated with the use of this pattern, nor can they all be described as yelling, though this analysis seems slightly closer to the mark.

In her sociolinguistic book on Twitter language, Zappavigna (2012) notes how the all caps pattern can be used for affective meaning, but offers a more nuanced explanation than stereotyped anger. She suggests that both the "caps lock" pattern and the potentially related pattern of reduplicated exclamation points are related to what she terms "upscaled graduation". She specifies that reduplication of exclamation points serves to upscale the "interpersonal meaning" of a written utterance, and that such an upscaling can be applied to both positive and negative meanings. The content of the tweet itself determines which meaning will be intensified. Though she makes a more explicit connection between punctuation and prosody, she largely simplifies capitalization to this kind of emotive upscaling, which her few examples (as well as my own above) seem to intuitively support. The aim of this initial experiment was to reformat Zappavigna's observations as a concrete hypothesis and to test it in a controlled experiment.

---

[3] Following other research dealing with language and social media (Zappavigna 2012, Baym 2015, Lamontagne & McCulloch 2017), Twitter is used in favor of other social media platforms due to its public nature, character limit, and accessibility through the Twitter API.

[4] Flaming is defined by Baym (2015) as messages which "include swearing, insults, name calling, negative affect, and typographic energy".

3.1. EXPERIMENTAL DESIGN. The experimental hypothesis is that the presence or absence of the caps lock condition in a tweet will influence readers' interpretations of the tweet's meaning. Specifically, I predicted, based on initial observations of data items and following Zappavigna (2012), that naturally occurring tweets including the caps lock pattern which are also independently judged as being exemplary for a particular emotion should be judged as having a lower degree of that emotion if the capitalization pattern is standardized. I use the term 'standardized' in this section to mean 'following the formal standard of current written English', not to indicate a process of creating new standards, as I use it elsewhere. In the case of capitalization, I consider the formal standard to be the capitalization of proper nouns, the first letter of sentences, and acronyms, but nothing else.

For this experiment, I downloaded a sample of about 100,000 English tweets via the Twitter streaming application programming interface (API), run on June 23rd, 2016. From this sample, I extracted every example which included the caps lock pattern and, after discarding duplicates and spam messages, I was left with a capitalization corpus of about 144 tweets.

From this corpus, a number of all caps tweets in a variety of possible emotions were selected for inclusion in the test data.[5] Items containing hyperlinks and emoticons[6] were dispreferred, and tweets were selected based on overall clarity of content out of context and possibility of exemplifying one of the test emotions. Identifying information, such as usernames, was removed from the test items. Each tweet was given a minimal pair, which maintained all of the spelling, punctuation, and other orthographic choices, but reduced the capitalized letters to occurring only sentence- and proper-noun-initially. An example of an original tweet and its standardized minimal pair can be seen in examples (10) and (11) below.

(12) RT @user1: @user2 HI ITS MY BIRTHDAY!!!
(13) RT @user1: @user2 Hi its my birthday!!!

The test was run in a survey format via Qualtrics, and test participants were recruited and compensated via Amazon Mechanical Turk. A total of 23 participants[7] took the survey, 12 male and 11 female. Ages of participants ranged from 18-60, with most of the participants being in their 20s. 17 participants reported using Twitter, and all participants reported using some social media platform for at least an hour a week. Participants were compensated $1.00 for completing the survey, as per standard rates for a ~20 minute survey on Amazon Mechanical Turk.

Each survey contained 30 content items and 6 filler items. For each item, the participant would see either the original version of the tweet (containing the relevant capitalization pattern) or the standardized version (without the capitalization), and no participant saw both versions of the same tweet. The question formatting as seen by the participants can be seen in Figure 1 below.

---

[5] A pilot study gathered readers' ratings of a variety of emotions for a selection of tweets from the corpus. This norming study helped to narrow the test emotions down to three, and these initial judgments allowed me to select the highest rated tweets for each emotion to include in the study, as a way of balancing out the data as much as possible.

[6] Emoticons are interpreted as paralinguistic non-verbal cues according to Zappavigna (2012) and others, and were thus excluded from this research, though future studies may benefit from examining the interrelation between emoticons and pragmatic orthography.

[7] 24 total participants were recruited, but one response set had to be discarded due to the participant not answering any of the content questions.
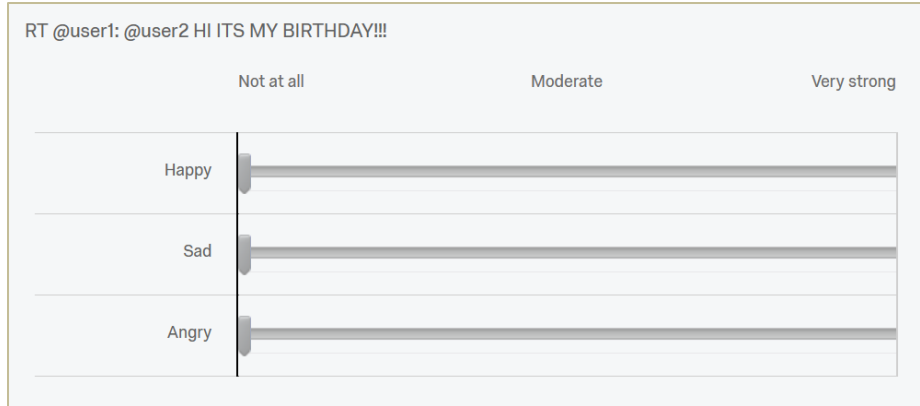
Figure 1. Question format

Three options of emotions were given for each item - happy, sad, and angry - and the participants were asked to slide the bars for an emotion if they deemed it to be present in the given item, to whatever degree they deemed it present. Participants were allowed to slide more than one bar if they chose, or to slide none of the bars if they judged that none of the given emotions were present. No exact numbers were shown to the participants to quantify the degree to which they slid the bar, but Qualtrics recorded their scores on a scale of 0-100 for the purposes of analysis. Participants saw only three degree markings: "Not at all" at the lowest end of the scale, "Moderate" in the middle, and "Very strong" at the high end.

3.2. RESULTS. In order to analyze the results, first I needed to determine which tweets were independently rated as being a good example of a tweet fitting one of the three emotions. On the 0-100 scale, if the original version of the tweet was rated above 60 for an emotion (i.e. falling above the "Moderate" marking on the slider) on average for all participants, I selected it as a strong example of that emotion. 10 of the 30 tweets were rated as happy by these standards, 4 were rated as angry, and none were rated as sad.

The results of the experiment are summarized in Table 1 and presented graphically in Figure 2. The estimated means out of 100 for the happiness rating of the original condition was $\mu=75.06$ and the standardized condition was $\mu=67.75$. A pair of by-participants (F1) and by-items (F2) repeated-measures ANOVAs revealed a significant effect of standardization on the happiness ratings ($F1(1, 22) = 7.69, p < 0.05$; $F2(1, 9) = 6.49, p < 0.05$). The difference between the estimates for original ($\mu=69.27$) and standardized ($\mu=68.77$) angry tweets was not significant.

	Original	Standardized
Happy	75.06 (2.32)	67.75 (2.56)
Angry	69.27 (3.41)	68.77 (3.87)

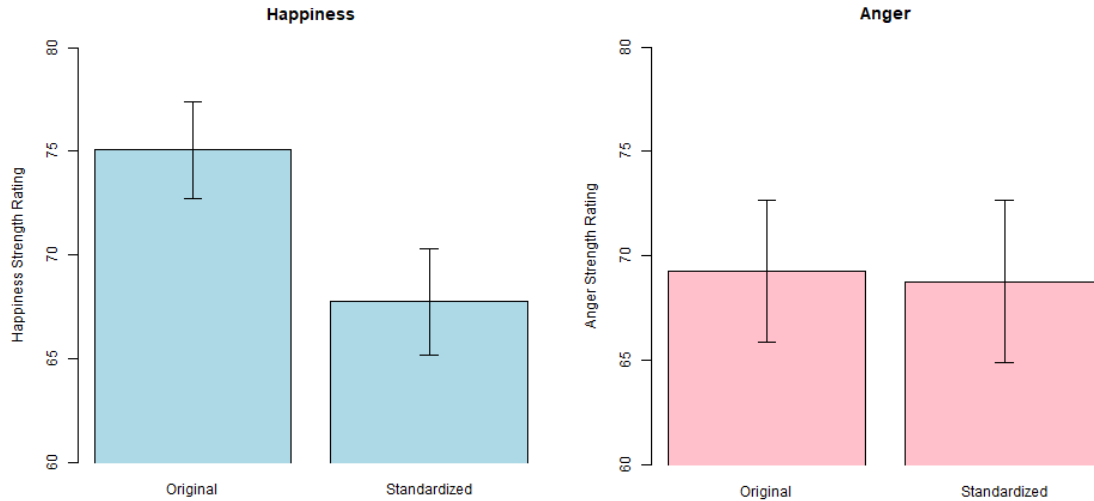Table 1: Mean happiness ratings with standard error in parentheses

Figure 2. Mean happiness and anger ratings by capitalization pattern

3.3. INTERPRETATION AND DISCUSSION. The results show that, for happiness, the realization of a tweet's orthography as "caps lock" or standardized influences the reader's interpretation of it. Tweets with the "caps lock" pattern are rated as happier than those without, for tweets that are rated overall as being happy. This difference supports the hypothesis that capitalization plays a pragmatic role in interpretation. Additionally, the simple fact that so many tweets were rated on average as happy disproves the stereotyped assumption that capitalization correlates to anger. But importantly, the lack of a significant difference in the ratings of the angry tweets also failed to support the broader theory that capitalization plays a role of general emotive strengthening in a text. If that were the case, we would expect to see the same pattern for all three emotions, when in fact we see three different patterns: happiness ratings strengthened, anger ratings unchanged, and sad ratings not attested for the pattern.

Given the failure of this experiment to support the original hypothesis and the surprising experimental results, I propose that rather than serving as a general emotive strengthener, the all caps pattern instead serves to implicate a heightening of prosodic features (pitch, volume, speed, etc.). A prosodic analysis explains the data more thoroughly, because it predicts that all caps would only be used to indicate emotions that, when indicated prosodically, evidence an increase in pitch or volume. Thus, an emotion like sadness, for example, which is considered to be a low arousal emotion which physiologically corresponds to lower pitch and volume (Bänziger & Scherer 2005, Burkhardt & Sendlmeier 2000), shouldn't regularly correspond with the all caps pattern. And in fact, though several the tweets in our survey were rated higher for sadness than happiness or anger, the average rating for original versions of the 10 tweets with the highest sadness ratings was only 29.4, while for happiness that average was 75.1. These results hint that capitalization is not normally used to heighten the degree of sadness, because it doesn't correspond to the prosodic features of stronger sadness.

A more interesting story can be told about the results for the tweets rated as angry. The lack of significant difference in ratings between the original and standardized versions of the four angry tweets may seem puzzling if we rely on the basic assumption that capitalization heightens emotions, but consider examples (14) and (15) below:

(14)   @user fiGHT ME
(15)   @user fight me

7

These two examples were both rated highly for anger in the survey. As it turns out, there are two different widely accepted types of anger with different prosodic realizations that may correlate to the two examples (Bänziger & Scherer 2005, Burkhardt & Sendlmeier 2000). Prosodic studies generally recognize hot anger as a high-arousal version of anger that is marked by heightened prosodic features such as pitch and volume. Cold anger, on the other hand, is the low arousal counterpart, which can have correspondingly lower prosodic feature realizations.[8] If we analyze capitalization as a marker of prosodic intensity, then it would follow that the example in (14) would be interpreted as hot anger, and (15) as cold anger. This would explain both the similarity in ratings, due to the survey allowing for no distinction between the two, and for overall rating strength.

 Since the prosodic analysis fits the experimental data, and the original emotive strengthening hypothesis was not supported by the reader judgments, we have reason to adopt a new prosodic hypothesis. But since the judgment study was not explicitly testing for prosody, further support will be necessary before we confidently adopt this theory. In the next study, we will look at a different type of evidence for the prosodic function of capitalization: an analysis of grammatical stress.

4. Prosodic corpus study. The intuition that orthographic tools such as punctuation, spelling, and capitalization can be used to compensate for a lack of prosody, intonation, and paralinguistic cues in writing can be seen in early work on discourse patterns in online communities. Werry (1996) observes chatroom-users writing in all-caps and proposes that it is being used for emphasis. Nearly two decades later, Baym (2015) makes similar claims, pointing out that a variety of orthographic tools can be used to set certain words or phrases apart from other text visually, thus indicating that they should be emphasized. In this section, we will turn to this kind of emphatic capitalization of isolated words, which I propose has the same underlying correspondence to prosody, but with a different linguistic function.

 In English, linguistic focus and emphasis are typically marked with increased pitch, volume, or length, with pitch being the most salient (Cruttenden 1997). Recall that this set of features overlaps significantly with the prosodic features of heightened emotions, though in this case, we are not talking about suprasegmental increases across multiple words, but rather increases on the scale of a single syllable and, by extension, a single word. Therefore, we might expect that if the all caps pattern is isolated to a single word, then its analogously increased prosodic features may be indicating stress or emphasis. In fact, this use of capitalization exactly mirrors the use of capitalization found in linguistic texts on intonation, especially to indicate focus. Is this simply because linguists needed to develop a formalism and arbitrarily chose capitalization, or is there a broader instinct that capitalization can be used prosodically in this way? To test this, we will be looking at examples from Twitter of the contexts in which people tend to use single-word capitalization to see if they line up with contexts which would be verbally focused or emphasized in speech.

---

[8] It is worth noting that most prosodic studies also differentiate two different levels of happiness (elation vs. happiness) and sadness (despair vs. sadness)(Bänziger & Scherer 2005; Burkhardt & Sendlmeier 2000). For the two levels of happiness, though, both tend to have prosodic features at or above base levels, so we don't expect to find a contrast between the two, but rather a gradient of intensity like what we see in the survey data. Sadness and despair do fall on different ends of the arousal scale, much like anger, but despair examples were uncommon overall in the corpus data. If we included despair example, we would expect to see them pattern with happiness in increased intensity with capitalization.

For this corpus study, I collected approximately 600,000 non-repeating English tweets between October and December of 2017, from which I extracted examples which included the single-word capitalization pattern.

For my initial test, I looked at work on focus sensitive expressions. These expressions have a relationship with focus such that some word in their scope must be focused, and where the focus falls affects its truth conditions or presuppositions (Beaver and Clark 2008). If capitalization can mark focus, we would expect to find instances of capitalized words appearing in the scope of focus sensitive expressions. And in fact, though not all instances of focus sensitive expressions in the Twitter corpus data included a capitalized element in their scope, there were examples that did, such as the following, in which the focus sensitive expression is underlined:

(16) I <u>just</u> need ONE photo of Tony and Kerry okay.
(17) <u>only</u> a fucking MORON wld feel hv to leave b/c ppl speak a diff language
(18) but don't <u>even</u> THINK about bringing ya shampoo & conditioner

As previously mentioned, it is generally understood that readers will assign default prosody to sentences with no explicit indication, so we should expect that most of the time, when default prosody is appropriate, these expressions which require focus in their scope don't need to be explicitly marked. But in cases where there may be ambiguity (such as between a focus on "one" versus on "photo" in example (16)) or a desire for extra emphasis (as in example (17) under the emphatic "fucking"), it is clear that writers are intuitively using capitalization to mark the linguistic focus.

Since not every example of linguistic focus falls in the scope of a focus-sensitive operator, the most frequent words in the corpus that were capitalized in isolation were examined. This frequency data offered an easy method of checking whether writers use capitalization in various focus environments. (19) shows some of the most frequent words capitalized in isolation and their frequencies.

(19) 'NOT', 409   'LIKES', 94   'THIS', 50
      'ALL', 238   'LIVE', 89   'BEST', 50
      'FREE', 187   'REALLY', 65   'THE', 49
      'YOU', 180   'ONLY', 62   'ME', 48
      'SO', 161   'STILL', 60   'IS', 48
      'AND', 150   'MY', 60   'PLEASE', 48
      'NO', 149   'YOUR', 59   'VERY', 47
      'NEVER', 123   'REAL', 56   'ANY', 46
      'THAT', 109   'ONE', 56   'ALWAYS', 46
      'LOVE', 103   'NOW', 53   'EVER', 43
      'WIN', 102   'AMAZING', 52

The words "free", "likes", "live", and "win" are all on the list because of viral advertising language patterns, which we will not be concerned with in this study. The words "love", "amazing", and "please" seem to be more emotive or emphatic in nature, as seen in examples (20)-(22) below. Though these examples are largely subjective, the use of capitals seems to be playing a role like that of the longer all caps examples, but on a smaller scale. They could easily be interpreted as pragmatic emphasis.

(20) i LOVE seeing photographer kim in action

(21) Thank you to everyone who came to our shows in the USA. We had an AMAZING time.
(22) If y'all do not do anything else today PLEASE watch this shit

There are also several pronouns which appear high on the list, namely "you", "my", "your", and "me". When we look at the contexts in which they occur, it becomes clear why they made the list. In example (23) below, "me" is the focused answer to the current question under discussion, "who gave Grayson a freaking hickey?".[9] This is an example of informative focus, wherein the focused element is the new, non-presupposed information. Examples (24) and (25), on the other hand, are examples of contrastive focus, where the pronouns are contrastive with a set of other possible referents, in this case someone else's party and someone else's Twitter.

(23) For anyone asking who gave Grayson a freaking hickey, it's ME madaforker
(24) Our own party cheering for the loss of a seat in Congress. How utterly despicable. You don't represent MY party!
(25) your Twitter is YOUR twitter and fuck it if ur tweets don't get a lot of retweets or likes that's not YOUR problem

Both informative and contrastive focus in English typically carry pitch accents, so these capitalized words are appearing in positions known to be verbally represented by a pitch accent. It is very likely that the single-word capitalization pattern is frequently used to represent these genres of focus on Twitter, but that only the pronouns were common enough words in general to show up on the frequency list. This trend is good evidence that writers have an intuitive sense that pitch accent can be marked with capitalization, and that the intended interpretation will be easily recoverable.

In addition to pronouns, there are also several items in the frequency list that have in common being an extreme point on some scale, whether that be a degree scale or an entailment-based "Horn scale" (Horn 1972). Words that fall into this broad category include "all", "never", "best", and "really", for example. These words all evoke a relevant scale of related alternatives. For example, "best" is on a scale with "better", "good", "bad", etc. as varying degrees, and "never" entails the weaker alternatives "rarely" and "sometimes". Focus on scalar words of this kind would evoke a contrastive interpretation, with the relevant alternatives being the other degrees on the scale. Additionally, these high-frequency scalar words happen to all be extreme high or low degrees on the scale, such that any alternatives would result in weaker assertions. For example, "best" is the highest degree positive superlative, and all the possible alternatives are weaker. Evoking these alternatives by focusing the stronger version may indicate that there was an expectation or presupposition of the weaker alternative, and the stronger alternative is surprising. This contrastive focus of the strongest scalar alternative leads to predictable interpretive effects, which we will explore below.

In addition to "best", other examples of scalar extremes which were frequently capitalized include the degree modifiers "very", "really", and "so". Again, these are all contrastive on a relevant scale of degrees, and examples which include these words presuppose some degree. In (26), for example, it seems to be contextually given that the writer wants to be the smoke bomb person, but the degree to which this is the case needs to be specified, and is contrastive with weaker possible degrees, such as "kind of". And the preface of "I'm not going to lie" indicates that we might not have expected them to make such a strong claim as they do with a focused "really". Likewise, in (27) there seems to be a presupposed scale on which a work day can be

---

[9] The capitalization in this example also happens to be in a cleft sentence focus position.

interesting or uninteresting, and it is surprising to find that it is an extreme degree of interesting in this particular case.

(26)   okay, I'm not going to lie, I REALLY want to be the smoke bomb person.
(27)   today was a VERY interesting day at work

In fact, Paradis (1997) proposes that these kinds of extreme, or what she terms "non-limiting" degrees are typically marked with different pitch contours than their limiting counterparts. Strong degree modifiers like "very" emphatically support (or reject) the assertion, and generally take falling tone. On the other hand, limiting modifiers like "pretty" or "kind of" serve to tone down the assertion being made or accept it with reservation, and are more typically associated with a fall-rise contour, indicating uncertainty or reluctance to commit to the assertion. If capitalization generally encodes high pitch accent on individual words, the contour will naturally evoke a falling contour from the capitalized word to the following lowercase word. This effect is another possible explanation for why strong modifiers frequently get capitalized and predicts that limiting modifiers may cause a prosodic clash if capitalized. And in fact, we find pairs like "always" and "never" in the frequency data, but not "sometimes" or "often", and we find "all" and "no", but not "a few" or even "many".

In additional support for the analysis of capitalization encoding prosodic emphasis on degree modifiers, note that for the word "so", which can be used in multiple senses, only ever appears in all caps in its degree modifier sense, as seen in (28).

(28)   my skin has been SO disrespectful lately, I feel so betrayed.

In other words, there are no examples in the corpus of "so" being capitalized when it is being used as a discourse marker or conjunction, only as a degree modifier. And in fact, we would expect these senses to rarely be the recipient of any kind of prosodic focus, whereas the degree modifier sense can take the scalar contrastive focus previously mentioned.

Although this preliminary research offers only a qualitative analysis, the previous examples indicate that many instances of single-word capitalization can be accounted for in the framework of common focus patterns, especially contrastive focus. Due to individual variation and other noise (such as the advertisement tweets), not every example can be perfectly described under this account, but it seems clear that, in a large number of cases, writers are sensitive to linguistic focus and have chosen to use capitalization to represent it.

5. Discussion. We have seen now that all caps patterns get used both in emotive contexts where they seem to be indicating a heightening of certain emotions through prosodic analogy, and that they can also be used to indicate linguistic focus and emphatic stress. A prosodic analysis, in which capitalization corresponds to heightened pitch, volume, or other related features, offers a cohesive explanation for both data sets.

A clear possible explanation for the emergence of the prosodic capitalization pattern is the growing need to enhance the communicative ability of written text. In an attempt to make writing more speech-like, writers may be drawing on readers' prosodic competence by writing in a visually stylistic way that triggers knowledge of the most contextually appropriate prosodic realization of an utterance. Speech-like features have been attested in writing frequency since the advent of instant text communication, and with the rapid evolution permitted by the size and popularity of online textual communication, it is unsurprising that an orthographic convention for marking prosody has emerged (Baym 2015, Crystal 2011, Werry 1996, Zappavigna 2012, Ferrara et al 1991).

There are a few additional factors which may play into the interpretation of capitalization on social media which I want to briefly acknowledge. First, different social media platforms make different subsets of orthographic variability available to users, and so differences in usage patterns may arise due to a reader's social media use patterns. Relatedly, different social groups across different platforms tend to develop their own conventions over time, which may also play into a reader's expectations. Because of the prolific and accessible nature of capitalization, I don't believe this to be a major issue, but there will always be individual stylistic differences to account for in research of this nature.

Additionally, demographics may play a large role. As Crystal (2011) points out, there are attested differences in language usage online between users of different age groups, men versus women, users of different educational backgrounds, and so forth. This kind of data is typically unavailable in corpus work, but large-scale studies may be able to tease out some of these variations.

Additionally, although I've chosen two specific capitalization patterns to examine here, there are additional levels of distinction that can be made within this category. For example, we have categorized tweets like (26) below as fitting the caps lock pattern, since a large portion of the tweet is capitalized.

(29) Can you guys tell that I switched markers halfway thROUGH BECAUSE THE DAMN MARKER RAN OUT OF INK

But although this study found no significant difference in judgements between tweets in which the entire utterance was capitalized and tweets like the (26), this may simply be because the distinction in their meanings is more fine-grained that this test was designed to determine. Further work is necessary to make these finer-grained distinctions.

Social media text has endless potential as corpus data, but in its raw form it is messy and difficult to analyze using traditional methods. Non-standard orthography and spelling in particular make it difficult for computational analysis, though work is being done to develop new models of computational analysis that take the unique features of internet language into account (Zappavigna 2012, Gimpel et al 2011). To such an end, developing a more complete model of pragmatic orthography would be a significant step towards making internet language usable as corpus data. Such a model would be an indispensable aid for both linguists (sociolinguists, psycholinguists, and computational linguists in particular) working on corpuses of social media text and to those working in sentiment analysis or marketing.

6. Conclusion. In this study, we have observed that internet users tend to deviate from the orthographic norms of English, and under the assumption that these deviations must be linguistically meaningful in some way, have set out to determine how. Specifically, we examined the caps lock pattern on Twitter, hypothesizing that it might correlate to heightened emotional content in a tweet. This hypothesis held up in the initial experiment for happiness, but not for anger or sadness, which led us to revise our hypothesis. Adopting the theory that certain capitalization patterns are used to convey prosody allows us to explain the data for all three emotions by forcing us to consider the prosodic realizations of different emotions and how they might be represented orthographically.

To further test whether there is a prosodic correlate to capitalization, we looked at twitter corpus data that included the single-word capitalization pattern. We observed that most examples of single-word capitalization occurred in positions which are typically associated with informational or contrastive focus, both of which are prosodically realized with stress features

such as heightened pitch and volume. These examples support the theory that capitalization corresponds to an abstract prosodic increase which can be interpreted in the context of other discourse information for a variety of linguistic and paralinguistic effects.

Ultimately, there is much work yet to be done in the area of pragmatic orthography in social media, but this study has taken the first steps toward systematically determining what meanings orthography may contribute to a text.

References

Bänziger, Tanja & Klaus Scherer. 2005. The role of intonation in emotional expressions. *Speech Communication* 46. 252–267. https://doi.org/10.1016/j.specom.2005.02.016.

Beaver, David and Brady Clark. 2008. *Sense and sensitivity.* Wiley-Blackwell. https://doi.org/10.1002/9781444304176.

Berg, Kristian and Mark Aranoff. 2017. Self-organization in the spelling of English suffixes: The emergence of culture out of anarchy. *Language* 93: 37-64. https://doi.org/10.1353/lan.2017.0000.

Burkhardt, Felix & Walter Sendlmeier. 2000. Verification of acoustical correlates of emotional speech using formant synthesis. ISCA tutorial and research workshop (ITRW) on speech and emotion. September 5-7, 2000. Newcastle, Northern Ireland, UK.

Carvalho, Paula et al. 2009. Clues for detecting irony in user-generated contents: Oh…!! It's "so easy" ;-). In *TSA '09: proceeding of the 1st international CIKM workshop on topic-sentiment analysis for mass opinion*. 53-56. New York: AMC.

Cruttenden, Alan. 1997. *Intonation.* Melbourne: Cambridge University Press.

Fodor, Janet D. (2002). Prosodic disambiguation in silent reading. In *Proceedings of the North East Linguistic Society* 32, M. Hirotani (ed.), GSLA, University of Massachusetts, Amherst.

Gimpel, Kevin et al. 2011. Part-of-speech tagging for Twitter: Annotations, features, and experiments. *Proceedings of ACL 2011*.

Heath, Maria. 2017. Interpretations of non-standard capitalization on Twitter. *LSO Working Papers in Linguistics*. Proceedings from WIGL 14. 15-29. http://vanhise.lss.wisc.edu/ling/?q=node/167.

Horn, Laurence. 1972. On the semantic properties of logical operations in English. Doctoral dissertation. UCLA.

Lamontagne, Jeff & Gretchen McCulloch. 2017. Wayyy longgg: Orthotactics and phonology in lengthening on Twitter. *LSA 2017 Annual Meeting*. January 5-8, Austin TX

Paradis, Carita. 1997. Degree modifiers of adjectives in spoken British English. *Lund Studies in English* Vol. 92. Lund University Press. https://doi.org/10.1017/S1360674301250187.

Piotrovskaya, L. 2003. Emotional prosody and emotive intonation. *International Society of Phonetic Sciences.*

Potts, Christopher. 2007. The expressive dimension. *Theoretical Linguistics* 33: 165-198. https://doi.org/10.1515/TL.2007.011.

Werry, Christopher C. 1996. Linguistic and interactional features of internet relay chat. In Susan C. Herring, ed., *Computer-mediated Communication: Linguistic, Social and Cross-Cultural Perspectives*: 47-64. https://doi.org/10.1075/pbns.39.

Wilson, Deirdre & Dan Sperber. 2004. Relevance theory. In L. Horn and G. Ward (eds.) *Handbook of Pragmatics*. Oxford: Blackwell.

Zappavigna, Michele. 2012. *Discourse of Twitter and social media*. London: Bloomsbury.