

How Constraints Affect Content: The Case of Twitter’s Switch from 140 to 280 Characters*

Kristina Gligorić
EPFL
kristina.gligoric@epfl.ch

Ashton Anderson
University of Toronto
ashton@cs.toronto.edu

Robert West
EPFL
robert.west@epfl.ch

Abstract

It is often said that constraints affect creative production, both in terms of form and quality. Online social media platforms frequently impose constraints on the content that users can produce, limiting the range of possible contributions. Do these restrictions tend to push creators towards producing more or less successful content? How do creators adapt their contributions to fit the limits imposed by social media platforms? To answer these questions, we conduct an observational study of a recent event: on November 7, 2017, Twitter changed the maximum allowable length of a tweet from 140 to 280 characters, thereby significantly altering its signature constraint. In the first study of this switch, we compare tweets with nearly or exactly 140 characters before the change to tweets of the same length posted after the change. This setup enables us to characterize how users alter their tweets to fit the constraint and how this affects their tweets’ success. We find that in response to a length constraint, users write more tersely, use more abbreviations and contracted forms, and use fewer definite articles. Also, although in general tweet success increases with length, we find initial evidence that tweets made to fit the 140-character constraint tend to be more successful than similar-length tweets written when the constraint was removed, suggesting that the length constraint improved tweet quality.

1 Introduction

The enemy of art is the absence of limitations.—Orson Welles

It is often thought that constraints affect both the form and quality of creative content. For example, there is anecdotal evidence across many domains that imposing length constraints can shape and improve the resulting writing: academic authors edit papers to fit a page limit, poets adhere to a prescribed verse form or rhyme scheme, and journalists edit articles to fit a word count limit (McPhee 2015). More broadly, research in several fields—spanning product design (Joyce 2009; Moreau and Dahl 2005), process management (Fritscher and Pigneur 2009), and education (Hennessey 1989)—suggests that having too much freedom can be paralyzing (epitomized by the feeling of staring at a blank sheet of paper), and that there is a sweet spot with just the right amount of constraints.

Many social media platforms enforce length restrictions for posts, which is why the latter are frequently called *micro-*

posts. For instance, at the time of writing, Instagram captions were limited to 2,200 characters, LinkedIn updates to 700 characters, Pinterest pins to 500 characters, and Twitter posts, *tweets*, to 280 characters. In the context of social media, a less restrictive character limit could conceivably be either positive or negative. In this work, we investigate the relationship between length constraints on post length and the form and quality of the resulting content. Do length constraints steer users towards editing their content in a useful way on social media, or does additional space allow for more creative and engaging content? Answering this question involves resolving two competing hypotheses:

- H1** Imposing constraints has a positive effect on creativity, influencing users to create succinct content that is more likely to be appealing to others (Joyce 2009).
- H2** Relaxing constraints provides users more space for expressing opinion and allows for more potentially interesting content. Also, longer posts occupy a larger portion of the update feeds displayed to users (Rosen 2017), further increasing engagement.

Choosing the right character limit is not trivial, given these conflicting hypotheses. Resolving them would not only help elucidate the relation between content and constraints; given the multi-billion-dollar economy depending on social media, the question is also of significant financial consequence.

The ideal way to answer this question would be through a randomized control study: subject randomly sampled users to various character limits, and observe how the nature and success of their posts depends on the respective limit. Unfortunately, however, only social-media providers themselves have full control over their platforms, including the ability to conduct such A/B tests. Furthermore, even this ideal experiment might suffer from creating an ecosystem with differing content lengths with clashing norms and conventions. Studying the effect of constraints is therefore difficult.

Twitter’s November 2017 switch. In order to circumvent these difficulties, we take advantage of a recent event to conduct an observational study of how length constraints affect microposts: on November 7, 2017, Twitter suddenly and unexpectedly increased its maximum tweet length from 140 to 280 characters.

The switch was reportedly introduced to allow users to express their thoughts without running out of characters, thus preventing them from finishing a thought (Rosen and Ihara

*Full version available at <https://arxiv.org/abs/1804.02318>
Copyright © 2018, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

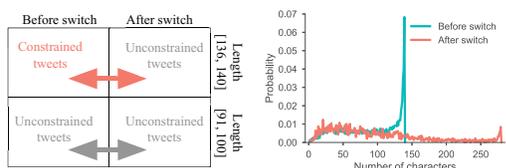


Figure 1: *Left*: Schema of study setup. *Right*: Histogram of tweet lengths, before and after the switch.

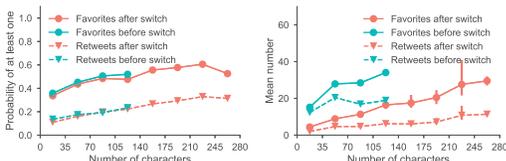


Figure 2: Engagement (retweets and favorites) as a function of length, with 95% confidence intervals. *Left*: Probability of receiving at least one engagement. *Right*: Mean number of engagements (for tweets with at least one engagement).

2017). This change in length constraint—which we henceforth refer to as *the switch*—constitutes an exogenous event that was most likely unexpected to most Twitter users, so we may reasonably assume that user behavior did not change in anticipation of this event. After carefully controlling for certain factors, differences between posts tweeted before *vs.* after the switch can inform us about the effect of different length constraints.

Research questions. Given that our aim is to study the effect of length constraints on the style and success of content in on-line social media platforms, we seek to answer the following questions:

- RQ1** Do length constraints lead to characteristic changes in the writing style of posts?
- RQ2** Do length constraints push users toward creating content that other users are more likely to engage with?

2 Methodology¹

Research design. In this paper, our main methodological contribution is a matched observational study design that allows us to use the exogenous shock to the Twitter ecosystem to investigate the relationship between length constraints and content form and quality, even in the absence of full experimental control of Twitter. Before Twitter’s switch from a 140-character limit to a 280-character limit, tweets that are nearly or exactly 140 characters in length are likely to have been explicitly “squeezed” by the user to comply with the character limit. After the switch, tweets of this length are less likely to have been affected by the character limit, since they are far short of the new 280-character maximum. Our basic design is thus to compare various properties of tweets of a

length between 136 and 140 characters² just before *vs.* right after the switch (Fig. 1, left). The major difference being the presence *vs.* absence of the 140-character limit, comparing tweets of this length from before *vs.* after the switch (corresponding to the upper, red arrow in Fig. 1, left) lets us isolate the impact of the character limit. By comparing content written very close in time, we minimize the likelihood of external factors changing in the time during which our studied tweets were written. In a control study, we also compare tweets with 91 to 100 characters before *vs.* after the switch, as these tweets are unaffected by the character limit during both time periods (lower, gray arrow in Fig. 1, left).

Data. In order to execute this research design, one would ideally analyze all tweets in our character range of interest that were authored right before and right after the switch. In lieu of access to complete Twitter data, we approximate this ideal. We first collect tweets from the 1% sample Twitter supplies via its *Spritzer API*, for the time between April and June 2017. From this set of tweets we sample 100K users, where the probability of being sampled is proportional to the number of 140-character-long tweets they posted. This ensures that we are likely to sample users who disproportionately generate tweets for which the 140-character constraint is relevant. We proceed to collect these users’ timelines using the Twitter *user timeline API*, gathering up to the 3,200 most recent tweets per user. We simultaneously keep collecting new tweets these users post. All tweets we analyze were posted between November 2017 and January 2018.³

Controlling for users and topics. With the goal of mitigating potential confounds in our comparison of tweets before *vs.* after the switch, we perform two analyses. First, we control for users by including, for each user, the same number of tweets written before the switch as after the switch (but this number may vary across users). As we are interested in matching tweets actually written by a particular user, we do not consider retweets of other users’ tweets. Second, we also conduct an analysis in which we control for topics by including, for each unique set of hashtags, the same number of tweets before as after the switch with that set of hashtags.

Controlling for temporal effects. By design, we compare tweets posted before and after the switch, thus introducing a time gap between the observed content. This raises two problems: first, retweet and favorite counts depend on how much time has passed since posting; and second, one might observe differences in the number of retweets and favorites as a side effect of a user’s becoming more popular with time. We address the first issue by making sure to collect a tweet at least 48 hours after it was posted, since by that time 99.99% of retweets happen (Sysomos Retweets Study 2010). The second issue is ruled out by making sure that the number of a user’s followers did not increase by more than 1% between the user’s first and last tweets in our dataset, when controlling for users.

¹More details in appendix of <https://arxiv.org/abs/1804.02318>

²We use a 5-character range because it is not always possible to strictly optimize for the limit: omitting one word to meet the constraint may result in a tweet of fewer than exactly 140 characters.

³Data available at https://github.com/epfl-dlab/140_to_280

Linguistic feature	Matched for user		Matched for topic		Control study	
	Before	After	Before	After	Before	After
Number of hashtags	0.53	0.54	—	—	0.22	0.22
Number of emojis	0.25	0.25	0.37	0.29	0.31	0.32
Number of abbreviations	0.17	0.13	0.09	0.07	0.12	0.12
Number of definite article	0.68	0.71	0.66	0.69	0.48	0.48
Number of indefinite articles	0.55	0.55	0.44	0.45	0.38	0.39
Number of <i>and</i>	0.38	0.42	0.25	0.32	0.32	0.32
Number of <i>&</i>	0.16	0.11	0.18	0.14	0.03	0.03
Number of missing spaces after punctuation	0.59	0.64	0.64	0.68	0.48	0.52
Number of auxiliary verbs and negations in long form	0.73	0.75	0.68	0.66	0.53	0.52
Number of auxiliary verbs and negations in contracted form	0.53	0.44	0.52	0.41	0.39	0.29
Fraction of lexical words (nouns, verbs, adjectives, adverbs)	56.45%	56.39%	55.44%	55.62%	57.34%	57.92%
Lexical variation	0.9089	0.9081	0.9136	0.9121	0.9405	0.9410
Lexical sophistication	0.0014	0.0014	0.0014	0.0014	0.0013	0.0013
Number of words per sentence	17.81	17.86	15.86	15.85	14.82	14.90

	Matched for user	
	Before	After
<i>I am</i>	0.011	0.013
<i>I'm</i>	0.057	0.051
<i>have</i>	0.123	0.128
<i>'ve</i>	0.033	0.026
<i>will</i>	0.065	0.068
<i>'ll</i>	0.027	0.021
<i>would</i>	0.033	0.034
<i>'d</i>	0.014	0.011

Table 1: Results of linguistic analysis (RQ1). *Left*: Linguistic features averaged across tweets, before vs. after switch, when matching for users and topics (Sec. 3.2), and in control study (Sec. 3.3). *Right*: Mean number of occurrences per tweet of expanded and contracted word forms. Bold values are significant with $p < 0.00365$ according to Kolmogorov–Smirnov tests.

	Matched for user				Matched for topic				Control study (tweet length [91, 100])			
	Retweets		Favorites		Retweets		Favorites		Retweets		Favorites	
	Before	After	Before	After	Before	After	Before	After	Before	After	Before	After
Probability of at least one	21.61%	20.38%	45.50%	45.52%	28.78%	29.94%	52.29%	52.17%	19.10%	18.72%	52.04%	52.94%
Quartiles (25%, 50%, 75%)	(1, 2, 7)	(1, 2, 7)	(1, 3, 8)	(1, 3, 8)	(1, 2, 5)	(1, 2, 4)	(1, 2, 6)	(1, 2, 6)	(1, 2, 5)	(1, 2, 5)	(1, 2, 6)	(1, 2, 6)

Table 2: Results of success analysis (RQ2). *Top*: Probability of receiving at least one engagement (retweet or favorite). *Bottom*: Quartiles of numbers of engagements (for tweets with at least one engagement).

Constraint consistency over time. We ensure that all content we analyze was generated under the same character counting policy. Additionally, we are primarily interested in how *text* is altered under a particular restriction. We thus limit ourselves to tweets not containing URLs and @ replies.

Pagination. Twitter users have long been working around the 140-character limit by splitting a long piece of text into a sequence of length-compliant tweets. These tweet threads are usually annotated with the position of the tweet in its sequence, *e.g.*, 2/3. They are usually not intentionally altered in order to fit the character limit requirement using mechanisms we are interested in. Thus, we disregard such tweets.

3 Results

3.1 General impact of length constraints

We begin with an analysis of the general impact of tweet length limits. We follow a data-driven approach and analyze 4M tweets written before, and 1.9M written after, the switch. Tweets are filtered according to our overall study setup (Sec. 2).

Fig. 1 (right) contains a histogram of tweet lengths before and after the switch. The spike at exactly 140 characters before the switch is indicative of users dealing with the character limit constraint. Additionally, the fact that there is no corresponding spike after the switch reflects the fact that the spike was indeed induced by the constraint—once it is lifted, 140 characters is an unexceptional tweet length.

Fig. 2 presents the probability of obtaining at least one engagement (retweet or favorite) and the average number of engagements (given at least one engagement), as a function of tweet length. We observe that, both before and after the switch, longer tweets are more successful on average, as captured by the number of retweets and favorites. This demonstrates the existence of a length effect consistent with hypothesis H2 (Sec. 1). This finding is, however, subject to

several potential confounds; *e.g.*, tweet length might correlate with the importance of what the user has to say, or with the popularity of the user herself, which might in turn be the real cause for the success of the respective tweet. Our matched studies, described next, address these confounds.

3.2 Matched studies

We compare tweets of length [136, 140] posted before, with tweets of a length in the same range posted after, the switch.

Linguistic analysis (RQ1). We are interested in linguistic aspects that might conceivably be altered in an attempt to meet length constraints: hashtags, emojis, abbreviations and acronyms used in the Twitter community (Beal, Vangie 2016), articles, conjunctions, spaces after punctuation, and auxiliary verbs in their long, contracted, or incorrectly contracted forms. We also evaluate stylistic features of text, including the frequency of lexical words, lexical variation, and sophistication as defined by Lu (2012). Moreover, we compute the number of words per sentence, an indicator of readability. Here, we are not interested in content that fails to elicit any engagement at all, as it is unlikely to reflect the optimizing mechanisms of interest. We thus keep only tweets that have received at least one retweet or favorite, resulting in 12K (4K) tweet pairs when matching for users (topics).

Matching for users and for topics, respectively, we compare the above linguistic features before vs. after the switch using Kolmogorov–Smirnov significance tests. Given that we test multiple hypotheses (one per feature), the significance threshold is adjusted via Dunn–Šidák correction (Abdi 2007), where the adjusted significance threshold is calculated as $1 - (1 - \alpha)^{1/m}$. We set $\alpha = 0.05$ and $m = 14$, resulting in an adjusted significance threshold of 0.00365.

Whether we control for users or for topics, constrained tweets systematically differ in linguistic features (Table 1). They contain fewer hashtags, fewer articles, and more words

in their short forms: more abbreviations, more auxiliary verbs in their contracted forms, more *&*, and fewer *and*.

These discrepancies are indicative of edits taking place given a 140-character limit constraint. We do not observe more missing spaces after punctuation in constrained tweets. This signals that users are unwilling to use punctuation in an incorrect way in order to fit a constraint.

Success analysis (RQ2). When quantifying engagement, we evaluate how much other users engage with a tweet in terms of retweets and favorites. Controlling for users and for topics, respectively, we evaluate success before *vs.* after the switch. The results of this study are presented in Table 2. We see that, when controlling for users, constrained tweets are 6.0% (relative) more likely to elicit engagement in the form of at least one retweet, compared to unconstrained tweets. Note that the trend is reversed when controlling for topics (instead of users), but the difference is smaller (−3.9% relative), and controlling for users is arguably a stricter criterion than controlling for topics, which, taken together, indicates that length restrictions slightly improve the success of tweets, in line with hypothesis H1 (Sec. 1).

3.3 Control study for ruling out confounds

As our analysis compares two different time periods to each other, one might object that the observed stylistic changes were in fact caused by a community-wide drift of norms between the two time periods. Although unlikely given the condensed time frame, we aim to rule out this alternative explanation with a control study in which we repeat our analysis (matching for users; 9K tweet pairs), but this time comparing tweets of a length much below the 140-character limit (91 to 100 characters) before *vs.* after the switch (*cf.* Sec. 2).

Such short tweets are unlikely to have been optimized for the length constraint. As a consequence, if any differences are observed for [91, 100] tweets, the same mechanisms might be at work for [136, 140] tweets, and it might be these mechanisms—rather than the 140-character constraint—that also cause the differences among [136, 140] tweets. However, we observe only one statistically significant difference in linguistic features for the [91, 100] range (Table 1), hinting at the length constraint as the cause for the other features.

As for success, we observe that [91, 100] tweets posted before the change are 2.0% (relative) more likely to be retweeted (Table 2), a much smaller difference than for [136, 140] tweets (6.0% relative), which we interpret as support for our above finding that length constraints slightly increase tweet success. However, more studies are necessary to confirm these findings, ideally with more data.

4 Discussion

In this paper, we address the question of how users adapt their contributions to fit the limits imposed by social media platforms and whether these restrictions tend to push users towards producing more or less successful content. We develop a matched observational methodology that involves comparing carefully selected sets of tweets posted just before *vs.* just after the switch.

We find that tweets constrained by a 140-character limit contain fewer hashtags and definite articles, and more words in abridged form: more abbreviations, more contracted auxiliary verbs, more *&*, and fewer *and*, implying that, when subject to a length constraint, users write more tersely.

We also found initial evidence that length-constrained tweets are slightly more successful in terms of the engagement they receive from other users. However, future work needs to develop a better understanding of this phenomenon and determine whether the observed findings are robust. We anticipate that collecting more data from the time after the relaxation of the constraint will be helpful to this end.

In our studies, we assume that tweets in the [136, 140] range posted after the switch are not constrained. We cannot rule out the existence of users adhering to the old constraint despite the switch (*e.g.*, out of sheer nostalgia, or because they are bots generating tweets via scripts written before the switch), but in the absence of such hypothetical users, the amplitude of the observed effects would be even larger.

Future work. This paper presents the first in what we hope will be a series of studies regarding the impact of Twitter’s switch from 140 to 280 characters. In particular, our observational studies should be followed by experimental studies, in order to corroborate the initial findings presented here. Also, our results are from right around the switch, when things were potentially still in flux. It would be interesting to revisit our findings in the future once the Twitter ecosystem has settled into a new steady state. Finally, we hope that future work will generalize our findings to constrained content production beyond Twitter.

References

- Abdi, H. 2007. Bonferroni and Šidák corrections for multiple comparisons. In *Encyclopedia of Measurement and Statistics*, 103–107. Sage Publications.
- Beal, Vangie. 2016. Twitter dictionary: A guide to understanding Twitter lingo. <https://goo.gl/3Wbkkt>.
- Fritscher, B., and Pigneur, Y. 2009. Supporting business model modelling: A compromise between creativity and constraints. In *Proc. International Workshop on Task Models and Diagrams*.
- Hennessey, B. A. 1989. The effect of extrinsic constraints on children’s creativity while using a computer. *Creativity Research Journal* 2(3):151–168.
- Joyce, C. K. 2009. *The Blank Page: Effects of Constraint on Creativity*. PhD thesis, UC Berkeley.
- Lu, X. 2012. The relationship of lexical richness to the quality of ESL learners’ oral narratives. *The Modern Language Journal* 96(2):190–208.
- McPhee, J. 2015. Omission: Choosing what to leave out. *The New Yorker*. September 14, 2015. <https://goo.gl/639Gpr>.
- Moreau, C. P., and Dahl, D. W. 2005. Designing the solution: The impact of constraints on consumers’ creativity. *Journal of Consumer Research* 32(1):13–22.
- Rosen, A., and Ihara, I. 2017. Giving you more characters to express yourself. <https://goo.gl/syjkWs>.
- Rosen, A. 2017. Tweeting made easier. <https://goo.gl/DYzBji>.
- Sysomos Retweets Study. 2010. Replies and retweets on Twitter. <https://goo.gl/hqDv5J>.