

RESEARCH ARTICLE

Evidence of Online Performance Deterioration in User Sessions on Reddit

Philipp Singer^{1,2*}, Emilio Ferrara³, Farshad Kooti³, Markus Strohmaier^{1,2}, Kristina Lerman³

1 GESIS - Leibniz Institute for the Social Sciences, Cologne, Germany, **2** University of Koblenz-Landau, Koblenz, Germany, **3** Information Sciences Institute, University of Southern California, Los Angeles, United States of America

* philipp.singer@gesis.org



OPEN ACCESS

Citation: Singer P, Ferrara E, Kooti F, Strohmaier M, Lerman K (2016) Evidence of Online Performance Deterioration in User Sessions on Reddit. PLoS ONE 11(8): e0161636. doi:10.1371/journal.pone.0161636

Editor: Tobias Preis, University of Warwick, UNITED KINGDOM

Received: April 23, 2016

Accepted: August 9, 2016

Published: August 25, 2016

Copyright: © 2016 Singer et al. This is an open access article distributed under the terms of the [Creative Commons Attribution License](#), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Data Availability Statement: Data is publicly available online. Information can be found at https://www.reddit.com/r/datasets/comments/3bxlg7/i_have_every_publicly_available_reddit_comment and direct access is e.g., possible in Google's BigQuery engine at https://bigquery.cloud.google.com/table/fh-bigquery:reddit_comments.2015_04. Additionally, for reproducibility and own inference, we make our code and experimental steps available online in the form of jupyter notebooks using R kernels in the supplementary material. The utilized Reddit data was not crawled by ourself, but is publicly available and promoted on Reddit. In this work, we only use aggregated data of the whole dataset and thus, the

Abstract

This article presents evidence of performance deterioration in online user sessions quantified by studying a massive dataset containing over 55 million comments posted on Reddit in April 2015. After segmenting the sessions (i.e., periods of activity without a prolonged break) depending on their intensity (i.e., how many posts users produced during sessions), we observe a general decrease in the quality of comments produced by users over the course of sessions. We propose mixed-effects models that capture the impact of session intensity on comments, including their length, quality, and the responses they generate from the community. Our findings suggest performance deterioration: Sessions of increasing intensity are associated with the production of shorter, progressively less complex comments, which receive declining quality scores (as rated by other users), and are less and less engaging (i.e., they attract fewer responses). Our contribution evokes a connection between cognitive and attention dynamics and the usage of online social peer production platforms, specifically the effects of deterioration of user performance.

Introduction

Performance deterioration following a period of sustained mental effort has been documented in settings that include student performance [1], driving [2], data entry [3], and exerting self-control [4]. Although the mechanisms for deteriorating performance are still debated [5–7], deterioration has been shown to be accompanied by physiological brain changes [8–10], suggesting a cognitive origin, whether due to mental fatigue, boredom, or strategic choices to limit attention. Outside of vigilance tasks, however, relatively little is known about whether and how this phenomenon affects online behavior. As our society becomes increasingly interconnected and people spend more time interacting through various online platforms, analyzing online performance is important for understanding how content is produced and consumed [11], how information spreads [12–14], and how people decide what and who to pay attention to [15, 16].

In this work, situated under the broad umbrella of user behavior modeling [17], we study online performance on Reddit, a popular peer production and social news platform. We

data used adheres to the terms and conditions of Reddit.

Funding: The authors received no specific funding for this work.

Competing Interests: The authors have declared that no competing interests exist.

measure online peer production performance as the quality of comments produced by Reddit users over the course of a session, defined as a period of activity without a prolonged break. The dataset we study contains over 55 million comments posted on Reddit in April 2015, and includes a variety of related meta-data, such as time stamps, information about the users, and the score attributed by others to each comment. We segment user activity into sessions, defined as periods of commenting without a break longer than 60 minutes, as suggested in [18] (cf. Fig 1). We link an individual's commenting performance over the course of a session to different proxy measures for a comment's quality, such as its length, readability, the score it receives from others, and the number of responses it triggers.

Our analyses uncover deteriorating online performance over the course of user sessions, with a decline in quality of subsequent comments *across* different proxy measures. Fig 2 illustrates the decline in the average score received by comments posted during sessions with ten comments: the data shows that each subsequent comment receives a rating that is on average 0.3 points lower than the preceding one. The size of this effect is quite large: It is equivalent to a 30% probability increase of receiving a downvote to a comment, for each extra comment posted after the first one in the session. Additionally, we observe that users tend to start with higher quality comments the longer the sessions are. To statistically study these effects, we design and implement mixed-effects models—allowing the incorporation of heterogeneous behavioral differences—which model the effect of session duration on the deterioration of online performance.

Our findings may be linked to effects of cognitive depletion: Exerting mental effort to compose a comment may diminish an individual's capacity to continue producing quality

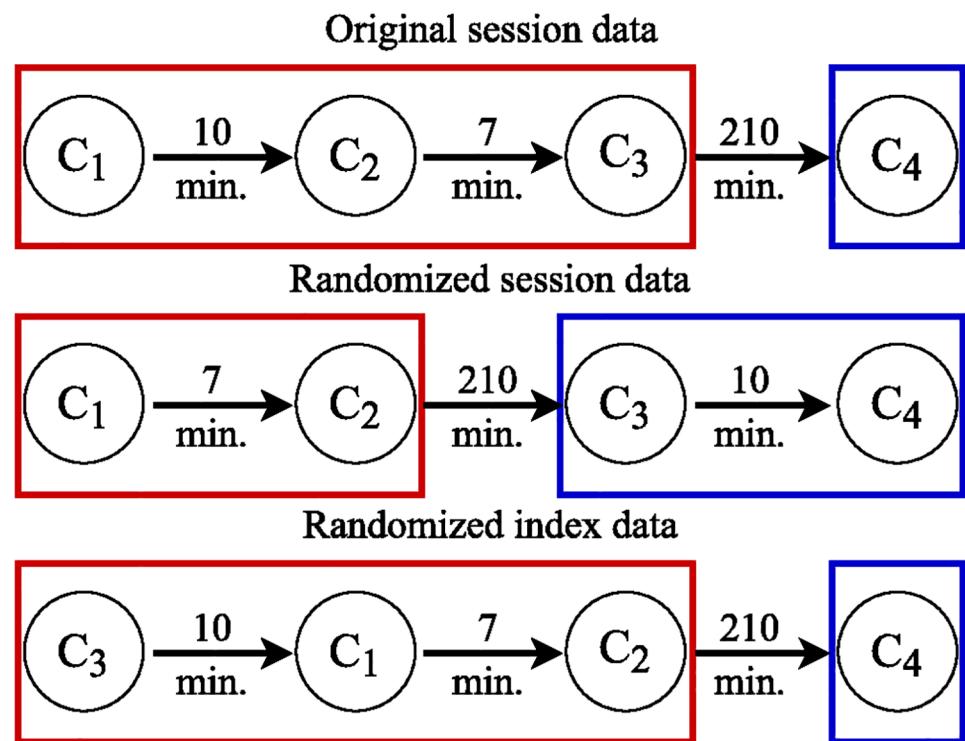


Fig 1. Sessions and randomization. Circles represent comments C_i and arrows depict the time difference $\Delta t_{i,j}$ between subsequent comments C_i and C_j . Sessions are derived by breaking at time differences exceeding 60 min. Original data sessions are shown in the first row. The middle row shows randomized sessions where time differences between comments are swapped for deriving new sessions while retaining the original order of comments. The bottom row depicts the randomized index data where sessions are retained but the order of comments within sessions is swapped.

doi:10.1371/journal.pone.0161636.g001

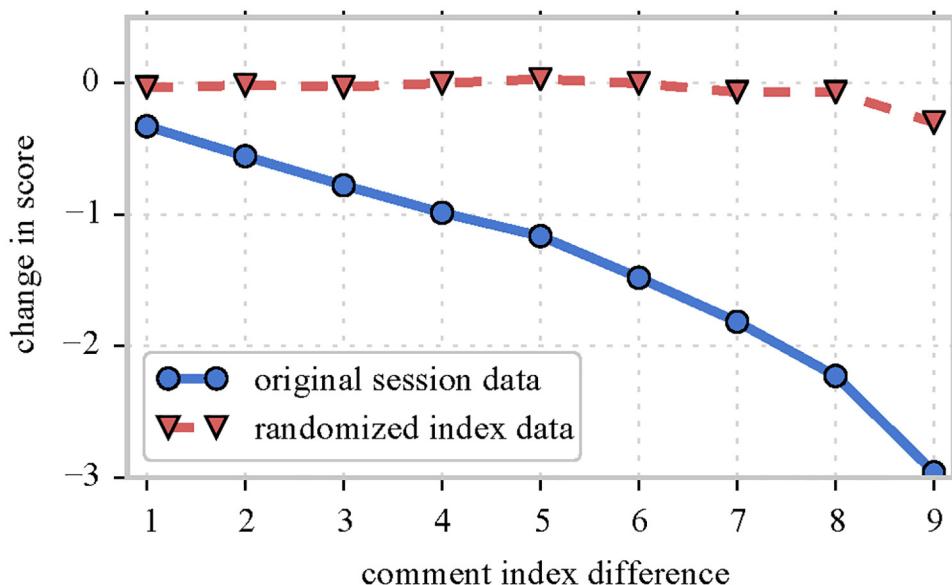


Fig 2. Performance of comments within sessions. We show the average Reddit score for comments in sessions of length 10 (original session data, blue solid line). The average rating of each comment decreases starkly, by about 0.3 points for each comment after the first one in the session. This suggests the presence of (super linear) performance deterioration throughout user sessions. The effect disappears in randomized data having shuffled comments within sessions (red dashed line).

doi:10.1371/journal.pone.0161636.g002

comments, whether through the loss of attention, mental fatigue, or simply the onset of boredom. Evidence also suggests that people, and other primates, have finite cognitive capacity for managing interpersonal relationships [19] limiting their amount of social interaction [20, 21]. Only recently, our research community started investigating the possible relationship between cognitive limits and online interactions, showing the impact of information overload on user behavior [20, 22–24]. Possibly, within-session deterioration of performance could explain the difficulty for users to continue exerting effort to discover information deeper in their social stream [15, 16, 25]. Also, deterioration might be influenced by the passive content consumption within a session, e.g., replies by other users to own comments maybe being toxic or hateful leading to flame wars [26]. The relation between the session length (i.e., number of comments) and the session's first comment's quality might also be explained by different starting capacities to make quality contributions, or, that the perceived quality of the first comment encourages users to produce more follow-up comments.

Although unveiling the mechanism(s) behind observed phenomena goes well beyond the scope of the current study, performance deterioration occurs throughout various critical daily activities, including learning (e.g., prolonged study sessions) and self-regulation (e.g., coping with stress, inhibition, refraining from behaving, or sticking to dietary restrictions). We believe that shedding light on the complex interplay between cognitive limits and individual performance can further our understanding of human behavior in many contexts. Thus, showing initial evidence of online performance deterioration is important and we expect this work to have implications for both computer and cognitive sciences communities.

Results

Next, we present our findings on studying effects of session dynamics on online performance focusing on (*i*) empirical observations, as well as by utilizing mixed-effects models on (*ii*)

performance at session start and (iii) performance over the course of sessions. We study, after pre-processing, around 40 million Reddit comments posted in April 2015. We derive user sessions as periods of commenting activity without breaks longer than 60 minutes as suggested in [18] (cf. [Materials and Methods](#) section and [Fig 1](#) for further details). For measuring performance, we look at four proxies of comment quality: text length, readability, the score a comment receives from others, and the number of responses it triggers. For comparison, we also study effects on two randomized session datasets as described in [Fig 1](#).

Empirical Observations

[Fig 3](#) visualizes changes in online performance over the course of user sessions with respect to our quality features (comment text length, number of responses, score, and readability). Different colors and markers distinguish sessions of distinct length (i.e., number of comments written during the session) of up to a length of 5. The x-axis shows the session index of a comment, the y-axis shows the (population-wide) average of respective feature (with error bars). For example, in the first plot of [Fig 3\(a\)](#), the red triangle at $x = 2$ refers to the average text length of all comments written in second position of all sessions of length 3.

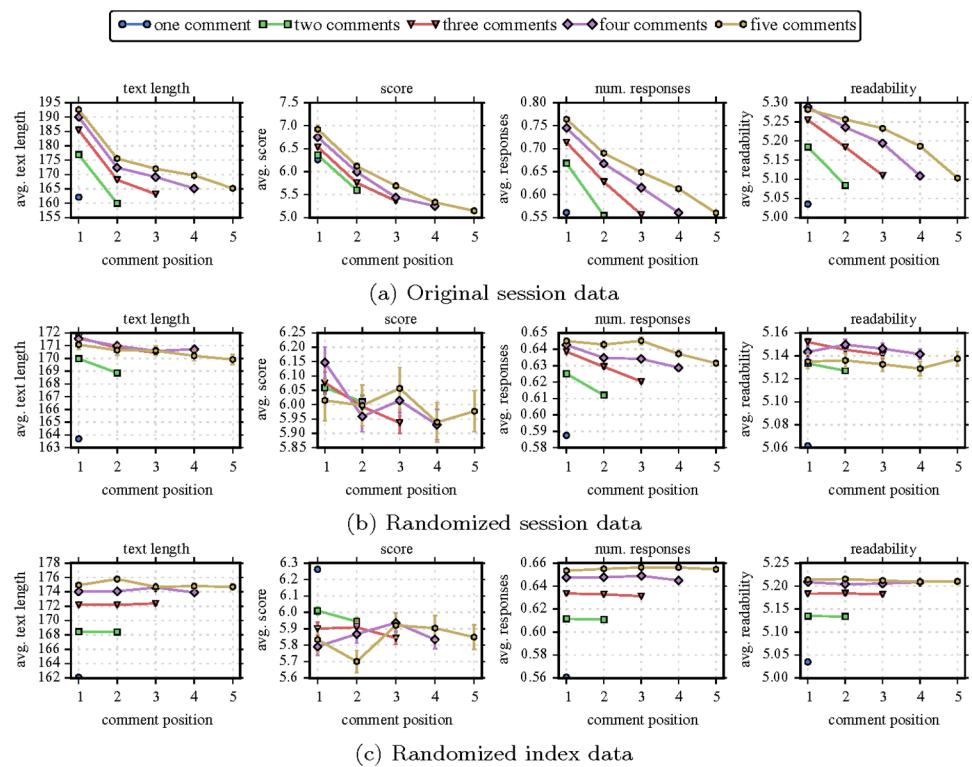


Fig 3. Empirical observations. This figure visualizes the average of all four quality features of interest at their respective position in a session. The colors (different markers) indicate different session lengths (number of comments written in a session, 1 up to a length of 5). The x-axis depicts a comment's index within the session, and the y-axis gives the average feature value (standard error bars of the mean are also depicted, but usually are very small and fall inside the data markers). The first row (a) depicts the original session data while the second (randomized session data) and third row (randomized index data) visualize results for the randomized data. The results indicate that earlier comments in a session tend to be of higher quality than later ones. Additionally, there appears to be a relation between the session length and the performance of the first comment in a session (stacking of lines). These clear patterns for the original data (a) mostly disappear for both of our randomized datasets (b,c). Overall, these empirical insights suggest performance deterioration over the course of sessions.

doi:10.1371/journal.pone.0161636.g003

[Fig 3\(a\)](#) depicts the original session data of interest and suggests interesting dynamics in user behavior. First, all lines are stacked: The first comment of a longer session also starts out with a longer text, a higher score, more responses, and more complex text (evidenced by higher readability score). Second, all feature values decline throughout the course of a session hinting towards some form of performance deterioration. On average, the last comment of a session is shorter, receives a lower score and fewer responses, and is easier to read.

In contrast, these trends largely disappear in our randomized data—i.e., randomized session data shown in [Fig 3\(b\)](#) and randomized index data shown as in [Fig 3\(c\)](#). There is no clear decline in feature values of later comments in comparison to earlier comments in sessions. The reason why some lines (e.g., number of responses) in [Fig 3\(b\)](#) and [3\(c\)](#) are still slightly stacked can be explained by our way of randomizing—see middle and bottom rows of [Fig 1](#). Some sessions (especially in the randomized session data) still stay partly, or sometimes even fully, intact, preserving the original session data. However, the effects are much reduced, for example, the average number of responses in the original data ranges between 0.55 and 0.77, while in the two randomized sets it ranges in the intervals [0.58–0.65] and [0.56–0.66] respectively.

Several considerations limit the conclusions we can draw from these empirical results. First, the population-wide average feature value may not be fully indicative of user performance because some distributions (length, score, responses) are heavy-tailed. Second, we have only visualized sessions up to a length of 5. While visualizations including all lengths up to 10 show similar trends (not shown here), more detailed analyses are necessary. Third, and most importantly, we have ignored the fact that our samples are not independent of each other as we repeatedly measure comments for individual redditors. Each user's behavior may be different, for example, one user may tend to write very long comments, while another one may prefer making shorter ones; mixing these different behavioral aspects in one analysis does not allow for specific inference about performance deterioration. We resolve some of these issues by using mixed-effects models incorporating individual differences (cf. [Materials and Methods](#) section). We start with an (*i*) analysis of the performance on the first comment in sessions, based on our observation of a potential stacking effect, and continue with (*ii*) experiments on performance deterioration over the course of sessions with respect to a potential decline in quality. In the main article, we only report the most appropriate models and the significant fixed effect coefficients. However, we make more extensive experimental results and model analytics available online [27] as well as in the supplementary material. The provided R notebooks ([S1–S8 Notebooks](#)) give insights about the experimental steps taken to come up with the appropriate models utilized in this work by examining a sample of 1 million data points. Additionally, we provide the complete regression output of the final models reported in this article in the supplementary material ([S1–S8 Tables](#)).

Performance at session start

We hypothesized a relation between the length of sessions and their comments' respective quality; readily apparent in the stacking of lines in [Fig 3\(a\)](#). We now statistically study this relation by focusing on the simplified question whether the length (number of comments) of a session has an effect on the performance of the very first comment in the session. We model the data with mixed-effects models specified as: feature $\sim 1 + \text{session length} + (1|\text{user})$. The outcome (dependent) variable refers to one of our four quality features. The session length is the main fixed effect of interest. Additionally, we vary the intercept between users (random effect). For this analysis, we limit our data to only consider the very first comment in each session (around 23.5M comments). The detailed model analytics are openly available and can be found online [27] and in the supplementary information ([S1–S4 Notebooks](#)).

Table 1. Mixed-effects model results. In (a), the models study the effect of session length on the quality of the first comment C_1 in a session; i.e., data only contains the first session comments. In (b), the models investigate the effect of the session index i on the quality of respective comment C_i ; data includes all comments in sessions with more than a single comment. Each table highlights the most appropriate models for each quality features based on extensive model analytics—lmer refers to linear mixed-effects models while glmer refers to generalized linear mixed-effects models. All coefficients are strongly significant as derived from model comparisons based on BIC statistics.

(a) Performance at session start		
feature	best model	coeff (session length)
text length	lmer (log-transform)	+0.0342
num. responses	glmer (Poisson, log link)	+0.0685
score	glmer (Poisson, log link, constant added for positivity)	+0.00015
readability	lmer	+0.0478
(b) Performance over the course of sessions		
feature	best model	coeff (session index)
text length	lmer (log-transform)	-0.0205
num. responses	glmer (Poisson, log link)	-0.0640
score	glmer (Poisson, log link, constant added for positivity)	-0.00028
readability	lmer	-0.0410

doi:10.1371/journal.pone.0161636.t001

The results (fixed session length effects) are summarized in [Table 1\(a\)](#) (cf. also [S1–S4 Tables](#)). As hypothesized in our empirical population-wide observations, the results indicate that there is a positive relation between the length of sessions (i.e., the number of comments) and their first comment's quality. This is imminent from resulting positive fixed effects coefficients meaning that an increase in session length leads, on average, to an increase of the first comment's text length, the number of responses it triggers, the score it receives and its Flesch-Kincaid grade level which corresponds to higher complexity of written text.

A potential explanation for the observed effect is that users start with different capacities to make quality contributions depending on how many more comments they plan to compose. Another (opposite) explanation could be that a higher performance of the first comment encourages users to produce more comments leading to longer sessions. While we believe the first explanation is more plausible—text length and readability are not based on external success measures, and responses accumulate at a (somewhat) longer time scale—future studies should aim at answering these causal questions. Without resolving the nature of causality, the identified relation between session length and quality of the first comment has implications for the experiments we report below that model the dynamics of user performance during the sessions. We have now shown, empirically and statistically, a high heterogeneity between different sessions with respect to their length. Accounting for this (e.g., as a nuisance effect) in our models is thus necessary.

Performance over the course of sessions

We now turn our attention to the dynamics of user performance in sessions on Reddit. Our empirical insights so far have suggested a performance decline throughout the course of a session. We statistically study this hypothesis by investigating whether the index of a comment (relative position in session) has an effect on the quality of the respective comment. To that end, we apply mixed-effects models specified as: $\text{feature} \sim 1 + \text{session index} + \text{session length} + (1|\text{user})$. Again, the dependent variable refers to one of our four quality features. The session index is the main fixed effect that we are interested in for studying performance declines. Our models include an additional nuisance effect controlling for individual session lengths as

suggested by our previous experiments—model analytics confirm the importance of this factor. An additional random effect models the variations of the intercept between different authors. For this analysis, we consider all data for sessions having more than a single comment (around 24.5M comments). The detailed model analytics can be found online [27] and in the supplementary information ([S5–S8 Notebooks](#)).

We summarize the main results in [Table 1\(b\)](#) and again focus on the fixed session index effect (cf. also [S5–S8 Tables](#)). The results now indicate a negative effect of the session index on our respective quality features indicated by the four negative coefficients. This means that with duration of a session, the quality of comments decreases on average. The next comment in a session is of shorter text length, triggers less responses and a fewer score, as well has a lower Flesch-Kincaid grade level indicating easier complexity of written text. This argues for performance deterioration throughout the course of user sessions on Reddit.

To further confirm observed effect, we repeated the above experiments on the randomized data. For both the randomized session and index datasets, the session index effect is not significant for all features of interest indicating no performance depletion effect in the randomized data (cf. [S5–S8 Notebooks](#)). This is in contrast to real session data analyzed above and also confirms that the effects do not simply arise as a result of the *order* in which comments are made, but their *order within* a session.

Discussion

Our work presents novel evidence of performance deterioration during prolonged online activity. By analyzing Reddit, a popular online social network that attracts millions of users, we showed that sessions with more activity are significantly associated with production of lower quality content, as measured by the length of the comment posted, its readability score, its average score and the number of responses it receives. In light of these findings, we developed a mixed-effects model that captures online performance deterioration. The code and results for all model analytics are available online [27] and available in the supplementary information ([S1–S8 Notebooks](#) and [S1–S8 Tables](#)).

Our analysis can be expanded in several directions. For example, we have only accounted for the basic differences between distinct Reddit users in the mixed-effects models. Yet, a much more nuanced analysis of heterogeneous effects of online performance deterioration would be warranted. One interesting direction involves understanding whether all individuals exhibit the same levels of performance deterioration, or whether these effects vary from user to user. For example, we might find that all users consistently exhibit deterioration or that different subgroups of users exist, where some users might even show improvements in performance over time. Neuroscience studies found individual differences in working memory and other cognitive activities in the human brain [28]. However, it remains unclear from a physiological standpoint whether capacity to process or produce information varies from person to person [29]. Online performance deterioration may also depend on acquired experience (as a form of cognitive dexterity) with a system. A new, and thus unfamiliar, user in a system may experience faster performance deterioration than an experienced user, because e.g., the cognitive or attention cost associated with the same operations may be experience-dependent (this is particularly true for information discovery and content production activities). A computational study of online performance in this direction could be very valuable.

Additionally, other hypotheses can be studied, such as that performance deterioration depends on the topic (politics vs. funny images), the time of the day, or the intensity of sessions (shorter average time differences between comments). A further aspect to consider is, that we have considered all comments posted to Reddit as equal, meaning that we did not distinguish

between those comments posted at the root of a comment hierarchy and those posted further down the hierarchy. Future research in that direction is necessary to better understand observed deterioration effect. For example, top-level comments might generally be of higher quality than low-level comments, or performance deterioration might be stronger for successive posts in the same submission thread compared to comments across submissions. Also, the position of a comment in the hierarchy also influences its visibility to others which might have an impact on perceived quality. These and similar questions can be studied by our proposed models. They are highly adaptable and fixed and random effects can be utilized to model these potential heterogeneous effects; for example, including a random effect allowing the deteriorating effects to vary between users could already allow us to make further inference about individual differences.

Furthermore, the set of quality features can be extended arbitrarily and also investigated more closely. In this work, we have focused on two features that are static (text length and readability) and two features that express the perception of the content by others (score and number of responses). Specifically the latter category of features warrants future studies, e.g., in light of potential social influence bias (herding) effects [30]. Yet, also other categories of quality features might be of interest, such as the sentiment of the comment.

Although our study was confined to Reddit, performance deterioration may generalize to other online activities. Future studies are needed to identify the mechanisms leading to observed deterioration, whether through the loss of attention, mental fatigue, or simply the onset of boredom. Regardless of the causes, understanding the complex interplay between individual's cognitive limits and dynamic behavior is key to optimizing individual—and collective—performance in peer production and other online systems.

Materials & Methods

Here, we thoroughly describe utilized data, corresponding pre-processing steps, and statistical mixed-effects modeling approach.

Data

For studying performance deterioration we utilized a publicly available dataset ever written on Reddit starting from the first one on October 17, 2007 to the last one at the end of May 2015 containing all comments (nearly 1.7 billion) [31]. For our experiments, we extracted a smaller sample that limits the data to all comments posted in April 2015. An advantage of this limited data is that we do not need to additionally account for changes in Reddit's platform not only in its interface, but also in its voting mechanisms as well as the general usage patterns of users on the site [32]. Our results are robust to sample data from other months showing similar observations.

Quality features

For measure online performance, we studied the following comment quality features.

Text length. This feature counts the number of characters in a comment and is an indicator for its textual length. Each URL in a comment accounts for one additional character. The overall mean of text lengths is $\mu = 168.08$, the median is $m = 86.00$, and the standard deviation is $\sigma = 281.88$.

Score. The score is a measure the perception of other users and is the difference between their up- and downvotes (the starting score is 1). All ratings can be summarized by the mean $\mu = 6.05$, the median $m = 1.00$ and the standard deviation $\sigma = 51.57$.

Number of responses. We see the number of replies a comment triggers as a proxy for engagement and a comment's success. We only count direct replies in the comment hierarchy.

Table 2. Pearson correlation between features.

	text length	readability	responses	score
text length	1.000	0.296	0.072	0.005
readability	0.296	1.000	0.043	0.005
responses	0.072	0.043	1.000	0.558
score	0.005	0.005	0.558	1.000

doi:10.1371/journal.pone.0161636.t002

The mean number of responses is $\mu = 0.61$, the median is $m = 0.00$ and the standard deviation is $\sigma = 1.44$.

Readability. George Klare provided the original definition of readability [33] as “the ease of understanding or comprehension due to the style of writing”. For measuring readability of Reddit comments, we use the so-called *Flesch-Kincaid grade level* [34] representing the readability of a piece of text by the number of years of education needed to understand the text upon first reading; it contrasts the number of words, sentences and syllables. It is defined as follows:

$$0.39 \left(\frac{\text{total words}}{\text{total sentences}} \right) + 11.8 \left(\frac{\text{total syllables}}{\text{total words}} \right) - 15.59$$

The lowest possible grade is -3.4 which e.g., emerges for comments that only contain a sentences having a single syllable such as “OK”, only a single URL or only emoticons. We set the maximum Flesch-Kincaid grade to be 22. Simply put, a higher Flesch-Kincaid grade indicates higher readability complexity of a given comment. The overall mean of the Flesch-Kincaid grade is $\mu = 5.12$, the median is $m = 4.91$ and the standard deviation is $\sigma = 4.61$.

Correlation of features. As shown in Table 2, most of the features are not strongly correlated (Pearson’s ρ) with each other; however, we can identify two special cases. First, readability and text length have a correlation of $\rho = 0.296$, which is not surprising given that shorter texts are easier to read, which is accounted for in the Flesch-Kincaid grade level formula. Second, the two success features score and number of responses have a correlation of $\rho = 0.558$, meaning that comments that get a high score also tend to receive more replies. However, overall, these correlation results indicate that each feature represents interesting aspects on its own. All correlation coefficients are strongly significant (p -values close to zero) for a significance test with the null hypothesis stating no correlation (also accounting for multiple comparison by e.g., Bonferroni adjustment).

Sessions

We decided to take the time differences between consecutive comments as session indicators. To that end, we followed the approach advised in [18] where a strong regularity in how social media users initiate events across several different platforms was identified. Authors argue that a good rule-of-thumb is an inactivity threshold of 60 minutes to separate sessions. However, as postulated, we first visually and analytically inspect the log-scaled histogram of time differences between consecutive comments (after cleaning comments, before filtering sessions) as depicted in Fig 4. Similar to the results presented for other platforms [18, 35], there is a peak for very short time scales (minutes) and a peak for time differences of one day suggesting daily routine. By fitting a Gaussian Mixture Model (using EM-algorithm, log-normal mixture) with two components to the log-transformed data, we end up with the two means $\mu_1 = 6.85\text{min}$. and $\mu_2 = 794\text{min}$. A natural valley is visible between the two peaks and thus, combined with the results from the log-normal mixture fitting, we follow the rule-of-thumb of [18] and pick a time difference $\Delta t_{i,j}$ of

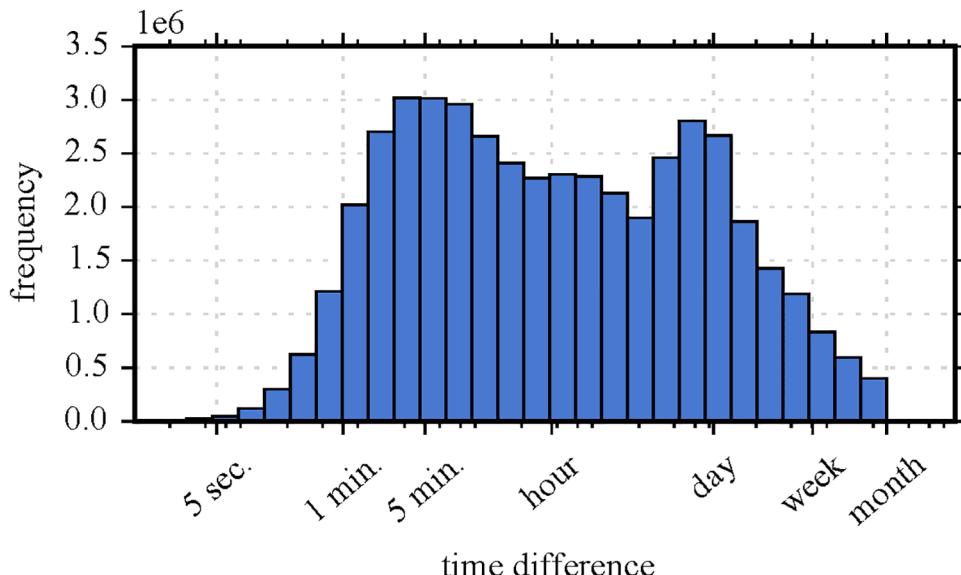


Fig 4. Time differences between consecutive comments of users on Reddit. The x-axis depicts the time differences between consecutive comments (tick labels refer to major ticks) and the y-axis illustrates respective frequency. The log-scaled histogram shows a peak for very short time scales (minutes) and very long ones (1 day) suggesting daily routines. A natural valley emerges between both peaks arguing for the choice of a one hour break between comments for sessions.

doi:10.1371/journal.pone.0161636.g004

one hour between consecutive comments C_i and C_j to separate sessions. Note that other (similar) choices of break time (e.g., 30 or 90 minutes) produce similar inference.

Data pre-processing

We took several steps for pre-processing and cleaning the data. First, we removed users from our data based on these rules: (i) They have posted the exact same comment more than 100 times, (ii) their username is part of an unofficial Reddit bot list [36], or (iii) their account has been deleted; this accounts for around 4.5M comments. Second, we deleted all sessions containing at least one comment (i) that has been deleted, (ii) that is completely empty, or (iii) that contains characters that are not in the ASCII character set (e.g., Chinese characters)—accounting for additional 3M comments. Finally, we removed all sessions containing more than 10 comments accounting for around 7.25M allowing for easier experimental tractability and the removal of further bot accounts. Note though that the inclusion of these sessions into the experiments does not change the main observations of this paper. Our final dataset contains 40,064,930 comments produced by 2,669,969 different users and posted in 47,462 different subreddits.

Randomizing sessions

For comparison, we created two randomized datasets to which we applied our analysis. The first baseline—which we call *randomized session dataset*—attempts to preserve as much information as possible while randomizing the process of deriving user commenting behavior sessions. To do so, we shuffled the time differences $\Delta t_{i,j}$ between consecutive comments made by each user, but preserved all other features, including the temporal order of comments. Then, we simply derived user activity sessions based on shuffled times. An example is provided in Fig 1 (middle row). This baseline dataset is very conservative in terms of randomization and

retains many original sessions. For example, many parts of a session stay intact as only the short time differences are potentially swapped, which does not alter the sessions. The second baseline—which we call *randomized index dataset*—keeps the sessions intact, but randomizes the order of comments inside each session (e.g., exchanging C_1 by C_3). Thus, it does not preserve the original order of comments; see Fig 1 (bottom row). Multiple randomization iterations did not alter the results.

Mixed-effects models

For statistically modeling performance deterioration, we utilized *mixed-effects models* allowing for the incorporation of heterogeneous effects and behavioral differences accounting for the non-independent nature of longitudinal data at hand. Mixed-effects models include both *fixed* and *random* effects; following [37], we refer to fixed effects as effects being constant across levels (e.g., individuals) and random effects as those varying between different levels. An overview of mixed-effects models can be found in [38].

In our setting, the introduction of random effects enabled us to consider variations between different levels; the most important level being different users accounting for the inherent differences between individual Reddit users (e.g., the average quality of their comments). As highlighted in [39], mixed-effects models have further advantages, such as flexibility in handling (i) missing data and (ii) continuous and categorical responses, as well as (iii) the capability of modeling heteroscedasticity. For simplicity, let us specify mixed-effects model equations using the following syntax [40]:

$$\text{outcome} \sim 1 + \text{fixed effect} + (\text{random effect} | \text{level}) \quad (1)$$

This specification describes a model where an outcome (dependent variable) is explained by an intercept 1, one or more fixed effect(s), as well as one or more random effects allowing for variations between levels. For all our experiments, we utilize the *lme4* R package [40] and fit the models with maximum likelihood. Examples about model specifications can be found online [41].

As each of our experiments is conducted on one of our four different features that all exhibit different properties—e.g., count (text length) vs. continuous (readability) data—we performed extensive model analytics to find the most suitable model for each problem setting. Overall, we aimed at finding the most appropriate model for each feature at hand by not only focusing on simple linear mixed-effects models, but also on generalized mixed-effects models such as Poisson or negative Binomial regression suitable for count data. When fitting regression models, several assumptions need to be considered, such as for linear models we need to check for normally distributed residuals and heteroscedasticity. Thus, we performed model diagnostics on the individual models and successively tried to improve our models, for example going from a linear model to a Poisson model. Additionally, we checked for overdispersion and zero-inflation in our count data models (Poisson and negative binomial) and accounted for it. We also tackled problems like multicollinearity, outlier bias, as well as convergence problems. The models reported in this article are the ones that we judged as the most useful ones for each setting at hand after extensive model diagnostics outlined above.

For judging significance of fixed and random effects, we followed an incremental modeling approach starting with the most simple model only explaining the outcome by the intercept and then subsequently adding effects to the model. For comparing the relative fits of these models we used the *Bayesian Information Criterion* (BIC) [42] which balances the likelihood of a model with its complexity. An interpretation table presented by Kass and Raftery [43] can be consulted to determine the strength of the differences between BIC scores. This allows to gain

confidence in the significance of observed effects allowing us to make inference on them. All reported fixed effects in this work are highly significant—except where mentioned (randomized baseline data)—meaning that the differences in BIC scores between the models including the effect and those excluding it are far larger than the maximum threshold of 10 indicating strong evidence as postulated in [43]. For completeness, we also conducted additional significance tests for the fixed effects such as t-tests or F-tests confirming our BIC diagnostics.

In order to enable the reader to follow our individual steps and also allow for personal inference, we provide detailed reports for each experiment—based on a sample of 1 million data points—in the form of jupyter notebooks using R kernels both online [27] and in the supplementary material ([S1–S8 Notebooks](#)). In the main article, we only reported the fixed effects and corresponding inference as those were the main effects we were interested in. However, we make the full regression outputs available in the supplementary information ([S1–S8 Tables](#)). Making our code and all experiments publicly available allows us to carefully document the results, as well as encourage other researchers to make their own inference and further refine our models. At the same time, utilized Reddit data is freely available [31].

Supporting Information

S1 Notebook. Experimental steps for studying effects of session length on first comment's text length.

(HTML)

S2 Notebook. Experimental steps for studying effects of session length on first comment's number of responses.

(HTML)

S3 Notebook. Experimental steps for studying effects of session length on first comment's score.

(HTML)

S4 Notebook. Experimental steps for studying effects of session length on first comment's readability.

(HTML)

S5 Notebook. Experimental steps for studying effects of session index on a comment's text length.

(HTML)

S6 Notebook. Experimental steps for studying effects of session index on a comment's number of responses.

(HTML)

S7 Notebook. Experimental steps for studying effects of session index on a comment's score.

(HTML)

S8 Notebook. Experimental steps for studying effects of session index on a comment's readability.

(HTML)

S1 Table. Mixed-effects model results for effects of session length on first comment's text length. This table presents the detailed mixed-effects model results for studying the effect of session length on the text length of the first comment C_1 in a session; i.e., data only contains the first session comments. The models at hand are linear mixed-effects models (lmer) where

the outcome (text length) has been log-transformed. The baseline model excludes the fixed effect at interest for judging the significance of the effect; comparing the BIC of both models reveals a clear significance. This is confirmed by the AIC as well as the classic t-test on the coefficient.

(PDF)

S2 Table. Mixed-effects model results for effects of session length on first comment's number of responses. This table presents the detailed mixed-effects model results for studying the effect of session length on the number of responses of the first comment C_1 in a session; i.e., data only contains the first session comments. The models at hand are generalized linear Poisson mixed-effects models (glmer) with a log link. The baseline model excludes the fixed effect at interest for judging the significance of the effect; comparing the BIC of both models reveals a clear significance. This is confirmed by the AIC as well as the classic t-test on the coefficient.

(PDF)

S3 Table. Mixed-effects model results for effects of session length on first comment's score. This table presents the detailed mixed-effects model results for studying the effect of session length on the score of the first comment C_1 in a session; i.e., data only contains the first session comments. The models at hand are generalized linear Poisson mixed-effects models (glmer) with a log link—additionally we have added a constant for making the score always positive. The baseline model excludes the fixed effect at interest for judging the significance of the effect; comparing the BIC of both models reveals a clear significance. This is confirmed by the AIC as well as the classic t-test on the coefficient.

(PDF)

S4 Table. Mixed-effects model results for effects of session length on first comment's readability. This table presents the detailed mixed-effects model results for studying the effect of session length on the readability of the first comment C_1 in a session; i.e., data only contains the first session comments. The models at hand are linear mixed-effects models (lmer). The baseline model excludes the fixed effect at interest for judging the significance of the effect; comparing the BIC of both models reveals a clear significance. This is confirmed by the AIC as well as the classic t-test on the coefficient.

(PDF)

S5 Table. Mixed-effects model results for effects of session index on a comment's text length. This table presents the detailed mixed-effects model results for studying the effect of the session index i on the text length of respective comment C_i ; i.e., data includes all session comments. The models at hand are linear mixed-effects models (lmer) where the outcome (text length) has been log-transformed. The baseline model excludes the fixed effect at interest for judging the significance of the effect; comparing the BIC of both models reveals a clear significance. This is confirmed by the AIC as well as the classic t-test on the coefficient.

(PDF)

S6 Table. Mixed-effects model results for effects of session index on a comment's number of responses. This table presents the detailed mixed-effects model results for studying the effect of the session index i on the number of responses of respective comment C_i ; i.e., data includes all session comments. The models at hand are generalized linear Poisson mixed-effects models (glmer) with a log link. The baseline model excludes the fixed effect at interest for judging the significance of the effect; comparing the BIC of both models reveals a clear significance. This is confirmed by the AIC as well as the classic t-test on the coefficient.

(PDF)

S7 Table. Mixed-effects model results for effects of session index on a comment's score.

This table presents the detailed mixed-effects model results for studying the effect of the session index i on the score of respective comment C_i ; i.e., data includes all session comments. The models at hand are generalized linear Poisson mixed-effects models (glmer) with a log link—additionally we have added a constant for making the score always positive. The baseline model excludes the fixed effect at interest for judging the significance of the effect; comparing the BIC of both models reveals a clear significance. This is confirmed by the AIC as well as the classic t-test on the coefficient.

(PDF)

S8 Table. Mixed-effects model results for effects of session index on a comment's readability.

This table presents the detailed mixed-effects model results for studying the effect of the session index i on the readability of respective comment C_i ; i.e., data includes all session comments. The models at hand are linear mixed-effects models (lmer). The baseline model excludes the fixed effect at interest for judging the significance of the effect; comparing the BIC of both models reveals a clear significance. This is confirmed by the AIC as well as the classic t-test on the coefficient.

(PDF)

Author Contributions

Conceived and designed the experiments: PS EF FK KL MS.

Performed the experiments: PS FK.

Analyzed the data: PS EF FK KL.

Contributed reagents/materials/analysis tools: PS FK.

Wrote the paper: PS EF FK KL MS.

References

1. Sievertsen HH, Gino F, Piovesan M. Cognitive fatigue influences students' performance on standardized tests. *Proceedings of the National Academy of Sciences*. 2016; 113(10):2621–2624. doi: [10.1073/pnas.1516947113](https://doi.org/10.1073/pnas.1516947113)
2. Borghini G, Astolfi L, Vecchiato G, Mattia D, Babiloni F. Measuring neurophysiological signals in aircraft pilots and car drivers for the assessment of mental workload, fatigue and drowsiness. *Neuroscience and Biobehavioral Reviews*. 2014; 44:58–75. doi: [10.1016/j.neubiorev.2012.10.003](https://doi.org/10.1016/j.neubiorev.2012.10.003) PMID: [23116991](#)
3. Healy AF, Kole JA, Buck-Gengle CJ, Bourne LE. Effects of prolonged work on data entry speed and accuracy. *Journal of Experimental Psychology Applied*. 2004; 10(3):188–199. doi: [10.1037/1076-898X.10.3.188](https://doi.org/10.1037/1076-898X.10.3.188) PMID: [15462620](#)
4. Muraven M, Baumeister RF. Self-regulation and depletion of limited resources: does self-control resemble a muscle? *Psychological Bulletin*. 2000; 126(2):247–259. doi: [10.1037/0033-2909.126.2.247](https://doi.org/10.1037/0033-2909.126.2.247) PMID: [10748642](#)
5. Boksem MAS, Tops M. Mental fatigue: Costs and benefits. *Brain Research Reviews*. 2008; 59(1):125–139. doi: [10.1016/j.brainresrev.2008.07.001](https://doi.org/10.1016/j.brainresrev.2008.07.001) PMID: [18652844](#)
6. Kurzban R, Duckworth A, Kable JW, Myers J. An opportunity cost model of subjective effort and task performance. *The Behavioral and Brain Sciences*. 2013; 36(6):661–679. doi: [10.1017/S0140525X12003196](https://doi.org/10.1017/S0140525X12003196) PMID: [24304775](#)
7. Vohs KD, Baumeister RF, Schmeichel BJ, Twenge JM, Nelson NM, Tice DM. Making Choices Impairs Subsequent Self-Control: A Limited-Resource Account of Decision Making, Self-Regulation, and Active Initiative. *Journal of Personality and Social Psychology*. 2008; 94(5):883–898. doi: [10.1037/0022-3514.94.5.883](https://doi.org/10.1037/0022-3514.94.5.883) PMID: [18444745](#)

8. Lorist MM, Boksem MAS, Ridderinkhof KR. Impaired cognitive control and reduced cingulate activity during mental fatigue. *Cognitive Brain Research*. 2005; 24(2):199–205. doi: [10.1016/j.cogbrainres.2005.01.018](https://doi.org/10.1016/j.cogbrainres.2005.01.018) PMID: [15993758](#)
9. Lim J, Wu WCC, Wang J, Detre JA, Dinges DF, Rao H. Imaging brain fatigue from sustained mental workload: an ASL perfusion study of the time-on-task effect. *NeuroImage*. 2010; 49(4):3426–3435. doi: [10.1016/j.neuroimage.2009.11.020](https://doi.org/10.1016/j.neuroimage.2009.11.020) PMID: [19925871](#)
10. Pattyn N, Neyt X, Henderickx D, Soetens E. Psychophysiological investigation of vigilance decrement: boredom or cognitive fatigue? *Physiology & Behavior*. 2008; 93(1-2):369–378. doi: [10.1016/j.physbeh.2007.09.016](https://doi.org/10.1016/j.physbeh.2007.09.016)
11. Ciampaglia GL, Flammini A, Menczer F. The production of information in the attention economy. *Scientific Reports*. 2015; 5 (9452). doi: [10.1038/srep09452](https://doi.org/10.1038/srep09452) PMID: [25989177](#)
12. Lerman K, Ghosh R. Information Contagion: An Empirical Study of the Spread of News on Digg and Twitter Social Networks. *International Conference on Web and Social Media*. 2010; 10:90–97.
13. Yang J, Counts S. Predicting the Speed, Scale, and Range of Information Diffusion in Twitter. *International Conference on Web and Social Media*. 2010; 10:355–358.
14. Wu S, Hofman JM, Mason WA, Watts DJ. Who says what to whom on twitter. In: *International Conference on World Wide Web*; 2011.
15. Weng L, Flammini A, Vespignani A, Menczer F. Competition among memes in a world with limited attention. *Scientific Reports*. 2012; 2(335). doi: [10.1038/srep00335](https://doi.org/10.1038/srep00335) PMID: [22461971](#)
16. Hodas NO, Lerman K. How visibility and divided attention constrain social contagion. In: *International Conference on Social Computing*; 2012.
17. Allen RB. User models: theory, method, and practice. *International Journal of Man-Machine Studies*. 1990; 32(5):511–543. doi: [10.1016/S0020-7373\(05\)80032-X](https://doi.org/10.1016/S0020-7373(05)80032-X)
18. Halfaker A, Keyes O, Kluver D, Thebault-Spieker J, Nguyen T, Shores K, et al. User Session Identification Based on Strong Regularities in Inter-activity Time. In: *International Conference on World Wide Web*; 2015.
19. Dunbar RIM. Neocortex size as a constraint on group size in primates. *Journal of Human Evolution*. 1992; 22(6):469–493. doi: [10.1016/0047-2484\(92\)90081-J](https://doi.org/10.1016/0047-2484(92)90081-J)
20. Goncalves B, Perra N, Vespignani A. Modeling Users' Activity on Twitter Networks: Validation of Dunbar's Number. *Plos One*. 2011; 6(8):e22656. doi: [10.1371/journal.pone.0022656](https://doi.org/10.1371/journal.pone.0022656) PMID: [21826200](#)
21. Miritello G, Moro E, Lara R, Martinez-Lopez R, Belchamber J, Roberts SGB, et al. Time as a limited resource: Communication strategy in mobile phone networks. *Social Networks*. 2013; 35(1):89–95. doi: [10.1016/j.socnet.2013.01.003](https://doi.org/10.1016/j.socnet.2013.01.003)
22. Kooti F, Aiello LM, Grbovic M, Lerman K, Mantrach A. Evolution of Conversations in the Age of Email Overload. In: *International Conference on World Wide Web*; 2015.
23. Hodas NO, Lerman K. The Simple Rules of Social Contagion. *Scientific Reports*. 2014; 4 (4343). doi: [10.1038/srep04343](https://doi.org/10.1038/srep04343) PMID: [24614301](#)
24. Gomez Rodriguez M, Gummadi K, Schoelkopf B. Quantifying Information Overload in Social Media and Its Impact on Social Contagions. In: *International Conference on Web and Social Media*; 2014.
25. Ferrara E, Interdonato R, Tagarelli A. Online popularity and topical interests through the lens of Instagram. In: *Conference on Hypertext and Social Media*; 2014.
26. Yasseri T, Sumi R, Rung A, Kornai A, Kertész J. Dynamics of conflicts in Wikipedia. *PloS one*. 2012; 7 (6):e38869. doi: [10.1371/journal.pone.0038869](https://doi.org/10.1371/journal.pone.0038869) PMID: [22745683](#)
27. Singer P, Ferrara E, Kooti F, Strohmaier M, Lerman K. Supplementary Reddit evidence notebooks. https://github.com/psinger/reddit_depletion.
28. Vogel EK, Machizawa MG. Neural activity predicts individual differences in visual working memory capacity. *Nature*. 2004; 428(6984):748–751. doi: [10.1038/nature02447](https://doi.org/10.1038/nature02447) PMID: [15085132](#)
29. Marois R, Ivanoff J. Capacity limits of information processing in the brain. *Trends in Cognitive Sciences*. 2005; 9(6):296–305. doi: [10.1016/j.tics.2005.04.010](https://doi.org/10.1016/j.tics.2005.04.010) PMID: [15925809](#)
30. Muchnik L, Aral S, Taylor SJ. Social influence bias: A randomized experiment. *Science*. 2013; 341 (6146):647–651. doi: [10.1126/science.1240466](https://doi.org/10.1126/science.1240466) PMID: [23929980](#)
31. Baumgartner J. Reddit data. https://www.reddit.com/r/datasets/comments/3bxlg7/i_have_every_publicly_available_reddit_comment/
32. Singer P, Flöck F, Meinhart C, Zeitfogel E, Strohmaier M. Evolution of reddit: from the front page of the internet to a self-referential community? In: *International Conference on World Wide Web Companion*; 2014.
33. Klare GR. Measurement of readability. 1963.

34. Kincaid JP, Fishburne RP Jr, Rogers RL, Chissom BS. Derivation of new readability formulas (automated readability index, fog count and flesch reading ease formula) for navy enlisted personnel. DTIC Document; 1975.
35. Geiger RS, Halfaker A. Using edit sessions to measure participation in Wikipedia. In: Conference on Computer Supported Cooperative Work; 2013. p. 861–870.
36. Reddit bots. <https://www.reddit.com/r/autowikibot/wiki/redditbots>.
37. Gelman A. Analysis of variance—why it is more important than ever. *The Annals of Statistics*. 2005; 33(1):1–53. doi: [10.1214/009053604000001048](https://doi.org/10.1214/009053604000001048)
38. Pinheiro J, Bates D. Mixed-effects models in S and S-PLUS. Springer Science & Business Media; 2006.
39. Baayen RH, Davidson DJ, Bates DM. Mixed-effects modeling with crossed random effects for subjects and items. *Journal of Memory and Language*. 2008; 59(4):390–412. doi: [10.1016/j.jml.2007.12.005](https://doi.org/10.1016/j.jml.2007.12.005)
40. Bates D, Mächler M, Bolker B, Walker S. Fitting Linear Mixed-Effects Models Using lme4. *Journal of Statistical Software*. 2015; 67(1):1–48. doi: [10.18637/jss.v067.i01](https://doi.org/10.18637/jss.v067.i01)
41. R-sig-mixed-models FAQ;. <http://glmm.wikidot.com/faq#toc27>.
42. Schwarz G, et al. Estimating the dimension of a model. *The Annals of Statistics*. 1978; 6(2):461–464. doi: [10.1214/aos/1176344136](https://doi.org/10.1214/aos/1176344136)
43. Kass RE, Raftery AE. Bayes factors. *Journal of the American Statistical Association*. 1995; 90(430):773–795. doi: [10.1080/01621459.1995.10476572](https://doi.org/10.1080/01621459.1995.10476572)