The anatomy of Reddit: An overview of academic research

Alexey N. Medvedev, Renaud Lambiotte, Jean-Charles Delvenne

Abstract Online forums provide rich environments where users may post questions and comments about different topics. Understanding how people behave in online forums may shed light on the fundamental mechanisms by which collective thinking emerges in a group of individuals, but it has also important practical applications, for instance to improve user experience, increase engagement or automatically identify bullying. Importantly, the datasets generated by the activity of the users are often openly available for researchers, in contrast to other sources of data in computational social science. In this survey, we map the main research directions that arose in recent years and focus primarily on the most popular platform, Reddit. We distinguish and categorise research depending on their focus on the posts or on the users, and point to different types of methodologies to extract information from the structure and dynamics of the system. We emphasize the diversity and richness of the research in terms of questions and methods, and suggest future avenues of research.

1 Introduction

Understanding the dynamics and structure of human communication is a central research theme in computational social science. The increasing availability of digital traces of human interactions has allowed to quantify, at a large scale, a variety

——————————————

Alexey N. Medvedev

naXys, Université de Namur, ICTEAM, Université catholique de Louvain, Belgium e-mail: `an_medvedev@yahoo.com`

Renaud Lambiotte

Mathematical Institute, University of Oxford, UK e-mail: `renaud.lambiotte@maths.ox.ac.uk`

Jean-Charles Delvenne

ICTEAM and CORE, Université catholique de Louvain, Belgium e-mail: `jean-charles.delvenne@uclouvain.be`

of phenomena. For instance, phone call logs led to the identification of the bursti-ness of human communication, typically organised into "periods" of short intensive communication followed by long periods of silence [27]; Facebook and email data helped to confirm the smallness of the world i.e. the typical network distance be-tween people is disproportionally small as compared to its size [4]; tweet messages led to studies to uncover the mechanisms leading to information cascades [56]; etc. If early works initially focused on one-to-one communication, the emergence of new communication channels, such as Twitter or online forums, has opened the possibility to study collective discussions [1].

Collective discussions have not been invented by new media. Think of demo-cratic debates in the Greek Senate. As such, they have been and remain a major way for exchanging opinions, and for producing collective decisions. Online forums pro-vide a venue where Internet-goers post questions or comments, which may, or may not, trigger discussions from other members of the community. Understanding how people behave in online forums has important theoretical implications, to improve our understanding of collective thinking, but also practical applications, to improve user experience, increase engagement or facilitate the democratic process [2]. The purpose of this Chapter is to provide an overview of the academic research on online discussion platforms, or online forums, and to bring together the variety of research questions considered the literature. Most of our attention is dedicated to the self-proclaimed "front page of the Internet" [43] – the website Reddit (REDDIT.COM) – which is the largest online discussion forum in the world as of today. Note that several other online discussion platforms have a similar architecture and have also been studied, for instance, in comparative studies; they include Digg, Hacker News, Slashdot, Epinions, Meneame, Barrapunto and even Wikipedia.

The rest of this paper is organised as follows. Section 2 presents the datasets that can be extracted from Reddit and have been widely used by researchers. Academic studies are then divided according to their primary focus on the post or the users, and presented in Sections 3 and 4 respectively. We conclude with a discussion and perspectives for future research.

2 The Reddit dataset

Reddit (launched in 2005) is a social news aggregation, web content rating, and discussion website, ranked as #6 most visited website in the world with 234 million unique users (as of February 2018)[1]. A schematic structure of Reddit is illustrated in Figure 1. Registered users submit posts that contain a title, an external link or a self-written piece of content, which immediately become available to the whole audience of Reddit for voting and commenting. The voting system permits only registered users to upvote (give a positive +1 vote) or to downvote (give a negative -1 vote) on posts and comments. Comments form a *discussion tree*, which can be

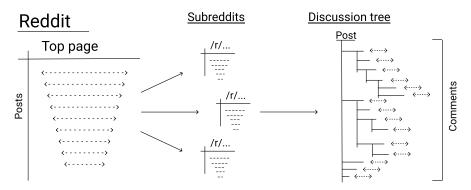---

[1] https://en.wikipedia.org/wiki/Reddit

Fig. 1 The schematic structure of the Reddit platform. The entry point is the top page of Reddit, which is fed of posts from the subreddits followed by a registered user (or from all subreddits, for an anonymous user), and ranked according to votes and posts' age. The user may further proceed to top page of a specific subreddit, where the feed narrows down to only posts from a chosen subreddit. Each post can be upvoted or downvoted and has an attached section of comments. The comments are structured as a rooted tree by the reply-to relation to other comments or the post itself.

described as a rooted tree, where the root is a designated node representing the post itself and each other node represents a comment. There is a link between two nodes if there is a 'reply-to' relation between them.

The huge posting space of Reddit is divided into subreddits – *self-created* communities of users, united by a certain topic. Every submitted post has subreddit name as an intangible attribute. Each subreddit and Reddit itself has a so-called "top page" – the feed where post titles with voting and commenting links are delivered to users. Two factors influence the post's ranking position there: 1) time and 2) voting *score*, or otherwise called *karma*, which is basically the difference between upvotes and downvotes. High score posts have a higher chance of appearing at the top page. However with time, newer information replaces the older in the feed. Users can follow subreddits, but not other users, which constitutes the main distinction with social network platforms, like Facebook or Twitter, where users follow a person and not a content. Other platforms have a similar structure. For example, Slashdot (launched in 1997) is made of news stories, together with comments moderated by selected users, but not by an open voting system[2]. Only a fixed number of topic-based subsections is available. Hacker News (launched in 2007) is an online community very similar to Reddit but with only two pre-made subreddits[3]. Digg (launched in 2004) acts currently as a news aggregator, but it formerly was a socially curated platform with a post submission, commenting and voting system like Reddit[4]. Meneame is a Spanish analogue of Digg, Barrapunto is also a Spanish version of Slashdot [15]. Etc.

---

[2] https://en.wikipedia.org/wiki/Slashdot

[3] https://en.wikipedia.org/wiki/Hacker_News

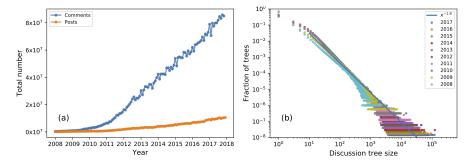[4] https://en.wikipedia.org/wiki/Digg

Fig. 2 The evolution of Reddit from Jan 2008 till Jan 2018: a) monthly counts of posts and comments, b) distribution of discussion sizes. One may notice an exponential increase in the activity counts, but the discussion size distribution follows a similar shape, close to the power-law with exponent α ranging from around 1.7 for early years to 1.9 for later.

Reddit has gained a central place in the scientific literature thanks to the openness, richness and quality of its data, which allows to perform longitudinal studies of the whole system and, critically, to ensure reproducibility of the results. Jason Baumgartner, under the Reddit name Stuck_In_The_Matrix, did a tremendous amount of work when attempted to collect a full dataset of posts and comments, going back to the creation of the site [46]. The figures of this Chapter have all been prepared from this dataset. For instance, basic numbers on the growth of the site and the total sizes of discussions are found in Figure 2. His data repository also contains the data from the platform Hacker News [45].

Despite its recognized quality, the dataset is not flawless either. Gaffney and Matias [10] report several inconsistencies in the data. For example, comment and post data before 2008 appears to be hugely corrupted, having around 80% of posts missing, as well as 90% posts information of one month data at the interface between 2009 and 2010. In total, across the time interval Jan 2006 and Feb 2016, the authors report 0.043% missing comments and 0.65% missing posts. The risks of missampled data are obvious, but in large scale studies they may be safely disregarded due to their smallness. The system has sustained exponential growth [43], thus data volume in early years is negligibly small compared to today's numbers. The newly published rescraped data must be free from the errors found in the data before 2008 [47].

Although missing data causes reasonable inconsistensies, reported comments from authors, whose accounts were deleted, constitute around 25% of the data according to the letter appendix to [10] which imposes a greater obstacle on user-centric studies. In such case, the name of the author turns to default name "[deleted]", but all the other information on the submitted post or comment remains in the system.

3 From the perspective of posts

Posts are at the heart of the platform structure and dynamics. Once posted, they may gain attention and receive feedback in the form of votes and comments, thereby obtaining a good ranking, and even more attention. They may also go quickly unnoticed in the avalanche of newer posts.

Popularity prediction. Anyone who has ever used online social networks is familiar with the concept of "likes" and "dislikes" – the way of expressing attitude towards a piece of content on a binary scale. If "page views" have long been the dominant measure of the success of a content, more and more platforms have moved to voting systems where the number of positive votes is the measure of popularity. Discussion platforms use systems of upvotes and downvotes for different purposes, ranging from the automatic discovery of appreciated items and its delivery them to a wider audience, to the moderation of discussions to protect from spam or malicious content. For these reasons, good models of popularity prediction are of interest for both content creators and platform curators.

In general, the problem of popularity prediction has been considered in various online social systems. Early studies, e.g. in YouTube and Digg, found a direct relation between content's initial popularity (in terms of views and upvotes) and its future counts [48], but more sophisticated models have been proposed since then. For instance, Lee et al. [31] have modelled the lifetime of discussions on Myspace with the Cox proportional hazard regression model. They selected the number of "risk" factors, fitted from the data for each thread, which were further used as a predictor of threads hitting a threshold number of comments. Mishne and Glance [34] analysed the corpus of comments in weblogs, and the relation between weblog popularity and commenting patterns in it. Tsagkias et al. [51] analyse the corpus of comments under news stories in regional Internet news agents. The authors propose a model that predicts the commenting popularity prior to article publication in two cases: first if there is a potential to receive comments, and second if the article receives "low" or "high" comment volume. Bandari et al. [5] also investigated if popularity of news articles can be estimated even before their posting online.

The Reddit dataset shows a proportional relation between the score of a post and the size of its discussion tree *on average*, as shown in Figure 3. This may shed some light on general aspects of the posts' popularity, however, in each particular case, it is more important to make a more tailored estimation of submissions' score. A number of works has been dedicated to the prediction of scores on Reddit. Horne et al. [24] found a number of textual and temporal features of high score comments by considering the discussion threads from 11 popular subreddits appearing in a 6 month period of 2013. The authors proposed a machine learning model for predicting comments' score and pointed the differences in users' preferences in subreddits. In particular, they claimed that timing of the comment, its relevancy and novelty have positive impact, but stale memes or high user ranking (overall number of positive comments in a user's history) does not affect or even pushes down the average comment score. It was observed that moderation does not always impact proper behavior in the community and may shorten the life of a discussion thread.
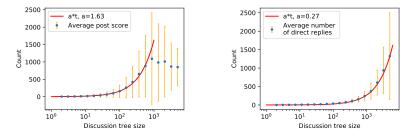
Fig. 3 Average posts' score versus discussion size (left figure) and number of direct replies to a post versus discussion size (right figure) in Reddit. The data shows that average values may be fitted with a linear trend up until a certain values. Dots represent average counts, bars show standard deviation. The figures are based on data from year 2009, but other years show similar results.

Recurrent neural networks (RNN) have also been employed to measure community endorsement. Fang et al. [9] constructed an RNN trained to predict comment scores. Instead of controlling for the submission context, the model learns latent modes of submission context and examines how the context relates to different levels of community endorsement. On a dataset of three popular subreddits, they achieve a good performance, on average, and show that high score comments are usually harder to predict than lower ones. High scoring comments tend to be submitted early in the discussion, and the number of direct replies is not smaller than the height of its hanging discussion subtree. Low and medium score comments have a number of direct replies less than the height of a discussion subtree, indicating the presence of a further discussion. Low score comments tend to come later in the discussion overall, but also later in terms of the group of responses to a parent comment.

While structural features have been shown to be good predictors, researchers have started working on extracting textual linguistic features in order to gain predictive power. In this direction, Jaech et al. [25] have reported general improvement of machine learning classifiers in the problem of ranking comments that appear in a fixed time window in a discussion thread. Note, however, that the gain was reported to be marginal. Later, Zayats and Ostendorf [54] constructed a specific type of RNN called LSTM (long short term memory) for the same purpose of predicting comment scores. The proposed model uses structural and temporal comment features, as well as textual linguistic features of the comments. The authors achieved a slightly better performance (increase in average F1 score from 50 to 54 on average) on a dataset of three subreddits studied earlier in [9]. They found that controversial comments (that further generate a wide discussion in terms of a discussion tree) tend to be overpredicted (with lower score than predicted) and jokes and funny comments, on the contrary, were mostly underpredicted (with a higher score than predicted). Linguistic context was found to be helpful in prediction tasks, and words of underpredicted comments were aligned with comments of positive score, but words associated with overpredicted comments did not show any significant correlation. In another work,

Article	Task	Dataset	Methods
Horne et al. [24]	Predict high scoring comments, assess the impact of thread moderation	Reddit dataset [46], 11 top subreddits	Linear regression, sentiment analysis
Fang et al. [9]	Predict final score of comments	Reddit, three chosen subreddits	Recurrent neural networks (RNN)
Zayats, Ostendorf [54]	Predict final score of comments	Reddit, three chosen subreddits	RNN with long short term memory (LSTM)
Hessel et al. [20]	Given a pair of submissions, predict the one with higher final score	Reddit dataset [46], 6 image-based subreddits	Image description (convolutional neural networks), LSTM
Stoddard [44]	Determine inherent quality of posts and to predict high-scoring posts	Hacker News; Reddit dataset [46], 5 top subreddits	Poisson processes
Lakkaraju et al. [29]	Predict popularity of resubmitted content	Reddit, unique dataset of resubmitted images	Poisson regression
Aragón et al. [3]	Review of the models of discussion trees	Reddit, Slashdot, Meneame, Barrapunto, etc.	Review
Medvedev et al. [33]	Model structure and predict dynamics of discussion trees	Reddit, dataset [46]	Stochastic Hawkes processes

Table 1 Short summary of the articles with studies on Reddit, presented in Section 3

Hessel et al. [20] considered pairs of posts, submitted within a very short time interval into the same communities (to exclude timing bias) and predicted the more popular posts in those pairs. Six primarily image sharing subreddits were selected for the study, with a set of features including textual and temporal information, but also image features assessed by the deep neural networks. The authors concluded that user-centric characteristics, e.g. previous popular submissions, and content-specific features, e.g. more complicated images and simpler titles, make a good predictor of popularity of the submission. The authors also reported an accuracy comparable to the that of human classification.

The popularity of a content is influenced by various factors, and public endorsement may not properly reflect its inherent quality [41]. This phenomenon has been studied by Stoddard [44] by means of a Poisson regression model that infers the intrinsic quality of posts from voting activity of the users. The author collected a unique dataset of users' voting time series by tracking a number of top posts on front pages of 5 subreddits and the front page of Hacker News, and used data to predict final score of posts. A variable of quality was then introduced in the model parameters, and shown to correlate with the total post scores, although there were several situations when similar quality posts had different scores and vice versa. Amongst others, the mechanism of making popular more visible than less popu-

lar ones leads to a multiplicative process that increases the variance of popularity, with the effect of making a substantial fraction of the posts ignored. According to Gilbert [11], Reddit overlooks 52% of the most popular links the first time they were submitted. Lakkaraju, McAuley and Leskovec [29] explored this idea and showed that resubmissions of the same piece of content may gain more popularity than an original submission. The authors collected a dataset of image submissions between 2008 and 2013, where each image was resubmitted roughly 7.9 times[5]. Same pictures can in principle be resubmitted to different subreddits, thus the authors employed a success metric compared to the average post score in the community. The authors also proposed a statistical model that predicts the expected score of a resubmitted picture. The model parameters include the inherent content popularity, penalties from previous success and previous submissions to other communities or to the same community twice. Overall, the study supports the hypothesis that a high quality content 'speaks for itself' and determines its score. The choice of subreddit plays an important role – the model shows that the content, resubmitted to the same subreddit, in general was unlikely to be popular, as well as whether the content was previously highly rated in a popular subreddit (with a high number of visitors or subscribers). This effect gradually disappears with time, indicating forgetfulness of the audience. Resubmission titles were also found indicative: if the title is novel, written using subreddit-specific words and sentiment orientation, the submission has higher chances to receive positive feedback. Similarly, Glenski et al. [12, 13] found that users mostly vote on posts only after glancing over the title, without proper reading of the content or the discussion. The authors of [29] also performed an *in situ* experiment: they manually chose and resubmitted 85 images from the dataset, select a "good" and a "bad" title according to the model for each picture and post them in two different subreddits. The post scores, gathered after one day, show that submissions with a "good" title generated scores three times higher than the "bad" ones.

Generative models for discussion trees. As mentioned above, comments under the post form a rooted tree. Such trees have dynamic nature and their temporal growth reveals the dynamic of attention to the post (see an example of a tree and a histogram of comment arrival on Figure 4). Generative models for discussion trees mostly question the tree structure of a discussion while disregarding the comments' or posts' textual features and exact timings. The reader may refer to the extensive review on generative models given by Aragón et al. [1], and here we give only a brief overview of the representative contributions.

Gomez et al. [15] considered discussion trees in four large Internet boards (Slashdot, Barrapunto, Meneame and Wikipedia) and proposed a generating model based on preferential attachment mechanism (PA model) with respect to the comment degree and the root bias. Later Gomez et al. [16] enriched this model by incorporating a notion of novelty of comments, which is represented by an exponentially decaying function of attractiveness. The model showed better results in likelihood of representing the tree structure and reproduced well the width/depth relation for discus-

---

[5] The authors used `karmadecay.com` – the reverse image search tool specifically designed for Reddit.
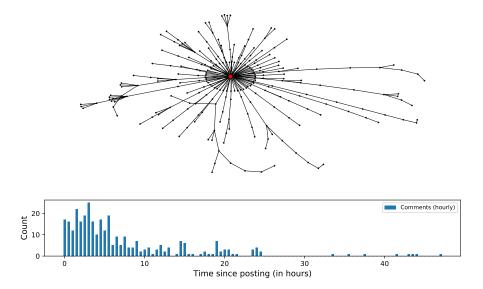
Fig. 4 Sample discussion tree of a post on Reddit with a histogram of comment arrival. Central large red node depicts the post, comments depicted in black. The histogram presents hourly aggregation of comments arrival.

sion trees. Lumbreras et al. [32] proposed to enrich the PA model with the notion of *roles*, which are latent functions of community members. The PA model defines a set of parameters that regulate the place of attachment of a new comment. The authors suggest that this set of parameters is different for different users and propose to group them into the role sets, which are inferred accordingly. Despite of this extra natural assumption, the gain in model likelihood is marginal.

Aragón et al. [3] considered the social system Meneame, where the change in discussion representation happened in 2015 from a plain list to a structured threaded view. This change was observed to have an impact on the structure of discussion trees, and the authors further enrich the PA model, with a reciprocity term that captures the tendency of posting authors to reply back in the discussion. It was observed that change of the platform interface had a positive effect on reciprocity, as well as on other parameters in general.

The above mentioned models focus exclusively on the structure of discussion trees while leaving out the continuous time dynamics of the comment attraction process. Kaltenbrunner et al. [26] found that comment arrival time in Slashdot discussions fits well by double lognormal distribution, although the fitting quality depends on a circadian rhythm of the site. Based on this finding, the authors propose the prediction model, that predicts a total number of comments in a discussion thread. The dynamical aspect of tree generation was first studied by Wang et al. [52], where the authors introduce a merely theoretical model for the structural and temporal

evolution of discussions. The temporal evolution was described as a Lévy process with power-law interevent time distribution, when newly arriving comments were assumed to attach to the existing tree under the simple PA rule. The model was inspired by empirical observations of such discussion boards, like Reddit, Digg and Epinions, however the mean-field nature of it limits its calibration with real-world datasets. Medvedev et al. [33] used a Hawkes process along with its branching tree interpretation to jointly model structure and dynamics of discussion trees in Reddit. The model was further used for prediction of the discussion flow, performing better than contemporary models of cascade dynamics.

Other. Zannettou et al. [53] perform a study of meme evolution and propagation across different platforms, e.g. Twitter, Reddit, 4chan and Gab, where the last two are image boards structurally similar to Reddit. The authors' analysis helps to reveal the influential actors in meme ecosystem, both in terms of creation and propagation, the authors build clusters of similar memes and make an analysis of reciprocal influence between the observed communities using Hawkes processes.

4 From the perspective of users

So far we emphasised the posts as the central pieces of information driving the dynamics of the platform. We now focus on the person that hides behind each post, comment, like or dislike, and review studies on the behavioral features of users.

Activity patterns. Observing the actions of users on a website can lead to interesting conclusions. Glenski et al. [12, 13] have studied a dataset with all recorded activity of 309 Reddit users within one year. The activity log included the information on all clicks, pageloads and votes made within the REDDIT.COM domain. As expected, the majority of users prefer passive browsing and rarely interact with the content (only 16% of users produce more than 50% of interactions). Users mostly vote on posts on average after only browsing the title (73 % of posts), although a non-negligible fraction (17%) of participants follow the link of the post and browse the section of comments before giving a vote. It was noted that users' probability of interaction with a given post decreases with the ranking of the post on the top page of Reddit as well as on subreddits. Text analysis of post titles shows that the probability of interaction increases with the reading ease of the title, i.e. as they use shorter words, and smaller sentences. The authors used the concept of activity sessions, which are the periods of user activity starting from an interaction and finishing after 1 hour without consecutive interactions. This terminology and a 1-hour threshold were adopted from Halfhaker et al. [17] and Singer et al. [42]. The authors reported a mean session length of 53 minutes. However, most participants had much shorter sessions prevailing (>3 mins). Singer et al. [42] also studied performance deterioration within sessions of active commenting on Reddit, where sessions of increasing intensity, i.e. how many posts users produced during sessions, are associated with the production of shorter, progressively less complex comments, which

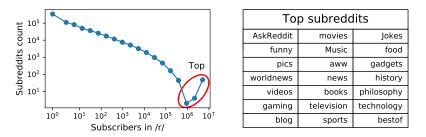| Top subreddits		
AskReddit	movies	Jokes
funny	Music	food
pics	aww	gadgets
worldnews	news	history
videos	books	philosophy
gaming	television	technology
blog	sports	bestof

Fig. 5 Distribution of the number of subscribers of subreddits. The largest subreddits in the rising tail of this distribution are shown along in the table.

receive declining score. In this work, the authors found a similar prevalence of short sessions, and sessions presented a daily circadian rhythms.

Community loyalty. While some of the previous results also applied to other online platforms, this section is devoted specifically to Reddit, a platform where the content is by default submitted to thematic communities, or subreddits. These subreddits are in principle open to anyone, but users can follow particular subreddits of their interest in order to customizing their feed. Interestingly, subreddit sizes are close to have heavy-tailed distribution, and there exists a fraction of subreddits-outliers with a huge number of subscribers 5. Looking closer at those we find that many of these "top" subreddits are those which are proposed to sign up for by default at the registration of a new user.

Tan and Lee [50] studied the posts of users across the subreddits and found that users on average tend to explore and continuously post in new communities, moreover they tend over time to share their activity evenly between a small number of communities with diverse interests. Differences in posting patterns of users may be used for prediction of the users' future settlement status in a community. Vagrant users, on average, post to more similar communities in comparison with the settling users, they use different language patterns from those existing in a community and their posts receive less attention in terms of score. Score of the first post may generally act as a predictor of further postings. The posting activity rate alone showed to be a bad predictor of the future settlement status in a community. An interesting finding is that the very same users tend to use different vocabulary when posting in different communities, therefore adapting to the community language. Hamilton et al. [18] defined loyal communities as the ones that retain their loyal users over time and find that such communities have smaller, but denser user interaction networks – with users as nodes, connected if there is a reply-to comment between them. These networks were found to be less assortative and less clustered, and thus show less fragmentation into groups. The authors then predicted whether a user will be loyal to the community, using a machine learning classifier with linguistic features of user's posts and comments, and achieve on average 63.6% classification accuracy.

Reddit allows users to self-organize into interest communities, which leads to interesting dynamics of communities. Hessel et al. [21] focused on communities

Article	Task	Dataset	Methods
Glenski et al. [12, 13]	Collect and assess the dataset of tracks of user actions	Reddit, unique dataset of user interactions	Statistical analysis
Singer et al. [42]	Assess user performance deterioration during activity sessions	Reddit, all comments made in April 2015	Statistical analysis, negative binomial and Poisson regression
Tan and Lee [50]	Study explorers and exploring phenomena of new communities	Reddit, dataset [46]	Statistical analysis, regression, linear classification
Hamilton et al. [18]	Loyalty prediction for newcoming users, patterns of loyal communities	Reddit, all comments made in 2014	User interaction networks, random forest classifiers
Hessel et al. [21]	Study the dynamic of arise of highly-related communities	Reddit, dataset [46]	Statistical analysis
Newell et al. [39]	Study the user migration across platforms during externally caused unrest period	Reddit, dataset [46]	Statistical analysis
Zhang et al. [55]	Classify subreddits along "niche" and "volatile" dimensions, study user retention	Reddit, dataset [46]	Statistical analysis
Das and Lavoie [7]	Model users posting strategies with respect to community feedback	Self-collected Reddit dataset	Machine learning, reinforcement learning, Hierarchical Dirichlet Process
Kumar et al. [28]	Mobilization and attacks between communities	Reddit, dataset [46]	Reply networks, lexical analysis, LSTM, Mechanical Turk
Tan [49]	Genealogy of subreddits	Reddit, dataset [46]	Relational networks

Table 2 Short summary of the articles with studies on Reddit, presented in Section 4

sharing name affixes, for example affix "ask" (science, askscience), "true" (atheism, trueatheism), "help" (tech, techhelp), etc. A curious finding is that when such a highly-related community is created, users engaging in the newer community tend to be more active than in their original one. However, in a prevailing number of cases, newer related subreddits do not detach out of their old partners (in terms of the user base), but in about 25% of the cases, the newer subreddit overtakes its counterpart in participation rate. Some reasons for this behaviour may include the absence of moderation in the new community, a more general scope, or simply a more appealing name spelling. The authors note an interesting result that users who explore the newer created communities generally become relatively more active in their home communities instead of being distracted.

Users may also migrate under external pressure. In 2015, a series of external events triggered closure of several popular subreddits. Newell et al. [39] studied the

history of this unrest period and observed that users migrating to new subreddits increase their level of participation with respect to their previous community. The authors followed users when they migrated to other discussion platforms and observed that although alternative platforms deliver a space for a broad audience, Reddit users value its advantage of hosting niche communities. In a similar vein, Zhang et al. [55] created a scalable framework for typing subreddits along the "niche" and "volatile" dimensions and used these types to understand the user retention and assimilation in subreddits. Finally, Muchnik et al. [38] performed a large-scale experiment on a Reddit-like platform to study the herding effect of social influence and how the system reacts to the manipulation of comment scores. They observed that users tend to correct artificially down-voted comments. However, comments that are artificially up-voted received an enhanced number of positive votes, thereby increasing the initial bias. Similar herding effects were found in other social systems as well [19, 40]. Das and Lavoie [7] also used a self-collected Reddit dataset of users posts and comments to train a reinforcement-learning model for how users select subreddits to post in reaction to community feedback.

Inherent networks of communities. Kumar et al. [28] considered interactions between communities in the form of mobilization by users of a community (the source of the 'attack') for hateful comments on posts from another community (the target of the attack). Such mobilization happens when a user in source community posts a link to a post in a target community and titles it with the intention of mobilizing a subset of users, who further write hateful comments on the target post. Such interactions may cause users of a target community to leave. By analyzing reply networks in target discussions, the authors found the effect of echo chambers, i.e. attackers preferentially interact with other attackers and defenders with other defenders. When a direct interaction happen, the attackers "gang-up" on defenders and only a small part of the defenders is involved into interactions with the attackers. The authors propose an LSTM neural network model that uses textual and social features in order to identify whether a given cross-linked post will produce a mobilization.

Gomez et al. [14] constructed and analysed the inherent social network of Slashdot, generated by replies in the discussion threads. The network exhibits neutral mixing by degree, almost identical in and out degree distributions, only moderated reciprocity and an absence of a community structure. The authors conjectured that users are more inclined to be linked to people who express different points of view, and that the network may help to identify users with a high diversity in opinions. The authors also proposed a measure for the controversy of a discussion, based on the h-index [22].

Tan [49] considered the genealogy of communities in Reddit. The author builds a weighted directed network of communities, where community A is linked to a community B if a substantial fraction of first 100 posting users in B had their posts in A. The weight of a link (A, B) is simply the fraction of posting users. The network shows user migration across the communities and is useful for predicting growth of communities. One finds that the diverse portfolio of memberships is the most important characteristic of early adopters, whereas community feedback and language similarity does not seem to matter.

Other. Discussion platforms are a tested for a broad spectrum of possible research questions. In addition to the topics covered above, we give now some other research directions. One phenomenon which frequently happens in online discussions is *trolling* – provocative, offensive or menacing messaging [6]. Mojica [35] collected and studied an annotated dataset of trolling comments in discussions on Reddit using a variety of language features. Derczynski and Rowe [8] used Reddit comments to create an annotated corpus of *named entities* – proper nouns representing a person, place or an organisation. Horne and Adali [23] studied how posting news articles on subreddit */r/worldnews* influences their popularity and conclude that changing the article titles results in greater popularity comparing to leaving the original one. In a similar way, Moyer et al. [37] studied how posts on the subreddit */r/todayilearned* influenced the pageviews of Wikipedia.

5 Discussion

This survey does not aim at giving a comprehensive listing of all Reddit-related works, but rather at providing a representative sampling that illustrates the richness of datasets related to online discussion platforms, with Reddit as a dominant example. The richness is clear in terms of quantity as well as quality — in principle the entire dataset can be harvested, while other platforms, e.g. Facebook may offer exhaustivity only at the cost of selecting a small sample of volunteers [30], and studies of Twitter are known to be limited by the volume and bias of their API [36]. The richness also arises from the diversity of the data, featuring an inherent social network between users, texts constitutive of posts and comments, social appreciation (score), tree structure of posts and comments, all that unfolding in time for years.

The richness of the data translates in a variety of topics of investigation. On the theoretical level, it allows to observe an ecosystem of users discussing, agreeing or not, and organizing in communities. Besides fundamental sociological questions, it also allows to investigate a range of Internet-specific questions, such as trolling, echo chambers, polarization, social manipulation, etc. The data also offers material to shape solutions for applied questions. From a platform designer viewpoint, it could help to improve the experience of a user, but also to design more efficient algorithms for the identification of high-quality posts. The design of the commenting system is also expected to affect the dynamics and structure of conversations. In this direction, important problems include the detection and automatic removal of trolling or attacks, as well as ways to stimulate the activity of a forum.

The richness of the data and problems calls for a range of computational methods, which may be explicit statistical models, or black-box machine learning tools, in order to classify or predict the behaviors of users, posts, communities. Overall we observe that the structure of discussion trees is relatively well understood. However, mixing the dynamics and structure with textual features is an important step that has only been studied by means of black box machine learning, such as neural networks techniques, showing a good performance in predicting community appreciation. A
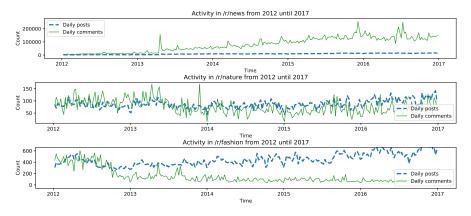
Fig. 6 Daily counts of submission of posts and comments in three selected subreddits. Three possible scenarios of participation dynamics are shown: 1) increase in comment rate exceeds posting rate (top figure); 2) rates are similar (middle figure); 3) comments eventually disappear, while number of posts increase.

challenge remains the fact that many basic statistics of the data (activity of users, popularity of communities, success of a post, etc.) exhibit heavy tails, which may introduce sampling issues as, for instance, a random sampling may fail to observe extreme points (e.g. high activity users) while they carry a large influence in the structure and dynamics of the system. This caveat must be kept in mind when using techniques such as neural networks.

This review is a testimony of the richness and dynamism of academic research on social platforms, in general, and Reddit, in particular. Despite the many progresses overviewed above, we would like to conclude with a list of what we believe to be promising research directions. In our opinion, fruitful avenues of research include a more detailed study of the activity patterns of users, of the influence of platform structure on user experience, and of the inherent networked effects behind posts, comments and votes. To be more precise, the dynamics of posting and commenting shape and define the life of a community. The two processes are coupled, but not necessarily proportional, as can be seen in Figure 6. Important questions include the identification of dynamical and structural features that ensure the growth or resilience of online communities over time. The dynamics of discussions is another interesting, yet mostly unexplored, aspect of research, especially the possible relation between the structure and the dynamics of discussion trees. This question could be explored by means of machine learning. Finally, the huge volume of new posts and comments makes the design of efficient ranking and recommendation algorithms vital, in order to allow users to identify relevant information and improve their online experience.

References

1. Aragón, P., Gómez, V., García, D., Kaltenbrunner, A.: Generative models of online discussion threads: state of the art and research challenges. Journal of Internet Services and Applications **8**(1), 15 (2017)
2. Aragón, P., Gómez, V., Kaltenbrunner, A.: Visualization tool for collective awareness in a platform of citizen proposals. In: Proceedings of the International AAAI Conference on Weblogs and Social Media, pp. 756–757 (2016)
3. Aragón, P., Gómez, V., Kaltenbrunner, A.: To thread or not to thread: The impact of conversation threading on online discussion. In: International AAAI Conference on Web and Social Media (2017)
4. Backstrom, L., Boldi, P., Rosa, M., Ugander, J., Vigna, S.: Four degrees of separation. In: Proceedings of the 4th Annual ACM Web Science Conference, pp. 33–42. ACM (2012)
5. Bandari, R., Asur, S., Huberman, B.A.: The pulse of news in social media: Forecasting popularity. ICWSM **12**, 26–33 (2012)
6. Bishop, J.: The effect of de-individuation of the internet troller on criminal procedure implementation: An interview with a hater. International Journal of Cyber Criminology **7**(1) (2013)
7. Das, S., Lavoie, A.: The effects of feedback on human behavior in social media: An inverse reinforcement learning model. In: Proceedings of the 2014 international conference on Autonomous agents and multi-agent systems, pp. 653–660. International Foundation for Autonomous Agents and Multiagent Systems (2014)
8. Derczynski, L., Rowe, M.: Tracking the diffusion of named entities. arXiv preprint arXiv:1712.08349 (2017)
9. Fang, H., Cheng, H., Ostendorf, M.: Learning latent local conversation modes for predicting comment endorsement in online discussions. In: Proceedings of The Fourth International Workshop on Natural Language Processing for Social Media, pp. 55–64 (2016)
10. Gaffney, D., Matias, J.N.: Caveat emptor, computational social science: Large-scale missing data in a widely-published reddit corpus. arXiv preprint arXiv:1803.05046 (2018)
11. Gilbert, E.: Widespread underprovision on reddit. In: Proceedings of the 2013 conference on Computer supported cooperative work, pp. 803–808. ACM (2013)
12. Glenski, M., Pennycuff, C., Weninger, T.: Consumers and curators: Browsing and voting patterns on reddit. IEEE Transactions on Computational Social Systems **4**(4), 196–206 (2017)
13. Glenski, M., Weninger, T.: Predicting user-interactions on reddit. In: Proceedings of the 2017 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining 2017, pp. 609–612. ACM (2017)
14. Gómez, V., Kaltenbrunner, A., López, V.: Statistical analysis of the social network and discussion threads in slashdot. In: Proceedings of the 17th international conference on World Wide Web, pp. 645–654. ACM (2008)
15. Gómez, V., Kappen, H.J., Kaltenbrunner, A.: Modeling the structure and evolution of discussion cascades. In: Proceedings of the 22Nd ACM Conference on Hypertext and Hypermedia, pp. 181–190 (2011)
16. Gómez, V., Kappen, H.J., Litvak, N., Kaltenbrunner, A.: A likelihood-based framework for the analysis of discussion threads. World Wide Web **16**(5-6), 645–675 (2013)
17. Halfaker, A., Keyes, O., Kluver, D., Thebault-Spieker, J., Nguyen, T., Shores, K., Uduwage, A., Warncke-Wang, M.: User session identification based on strong regularities in inter-activity time. In: Proceedings of the 24th International Conference on World Wide Web, pp. 410–418. International World Wide Web Conferences Steering Committee (2015)

18. Hamilton, W.L., Zhang, J., Danescu-Niculescu-Mizil, C., Jurafsky, D., Leskovec, J.: Loyalty in online communities. In: Proceedings of the International AAAI Conference on Weblogs and Social Media, vol. 2017, p. 540. NIH Public Access (2017)

19. Hanson, W.A., Putler, D.S.: Hits and misses: Herd behavior and online product popularity. Marketing letters **7**(4), 297–305 (1996)

20. Hessel, J., Lee, L., Mimno, D.: Cats and captions vs. creators and the clock: Comparing multimodal content to context in predicting relative popularity. In: Proceedings of the 26th International Conference on World Wide Web, pp. 927–936. International World Wide Web Conferences Steering Committee (2017)

21. Hessel, J., Tan, C., Lee, L.: Science, askscience, and badscience: On the coexistence of highly related communities. In: ICWSM, pp. 171–180 (2016)

22. Hirsch, J.E.: An index to quantify an individual's scientific research output. Proceedings of the National academy of Sciences **102**(46), 16,569–16,572 (2005)

23. Horne, B.D., Adali, S.: The impact of crowds on news engagement: A reddit case study. arXiv preprint arXiv:1703.10570 (2017)

24. Horne, B.D., Adali, S., Sikdar, S.: Identifying the social signals that drive online discussions: A case study of reddit communities. In: 26th International Conference on Computer Communication and Networks (ICCCN), pp. 1–9 (2017). DOI 10.1109/ICCCN.2017.8038388

25. Jaech, A., Zayats, V., Fang, H., Ostendorf, M., Hajishirzi, H.: Talking to the crowd: What do people react to in online discussions? In: Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing, pp. 2026–2031 (2015)

26. Kaltenbrunner, A., Gomez, V., Lopez, V.: Description and prediction of slashdot activity. In: Web Conference, 2007. LA-WEB 2007. Latin American, pp. 57–66. IEEE (2007)

27. Karsai, M., Kivelä, M., Pan, R.K., Kaski, K., Kertész, J., Barabási, A.L., Saramäki, J.: Small but slow world: How network topology and burstiness slow down spreading. Physical Review E **83**(2), 025,102 (2011)

28. Kumar, S., Hamilton, W.L., Leskovec, J., Jurafsky, D.: Community interaction and conflict on the web. In: Proceedings of the 2018 World Wide Web Conference on World Wide Web, pp. 933–943. International World Wide Web Conferences Steering Committee (2018)

29. Lakkaraju, H., McAuley, J.J., Leskovec, J.: What's in a name? understanding the interplay between titles, content, and communities in social media. ICWSM **1**(2), 3 (2013)

30. Lambiotte, R., Kosinski, M.: Tracking the digital footprints of personality. Proceedings of the IEEE **102**(12), 1934–1939 (2014)

31. Lee, J.G., Moon, S., Salamatian, K.: Modeling and predicting the popularity of online contents with cox proportional hazard regression model. Neurocomputing **76**(1), 134–145 (2012)

32. Lumbreras, A., Jouve, B., Velcin, J., Guégan, M.: Role detection in online forums based on growth models for trees. Social Network Analysis and Mining **7**(1), 49 (2017)

33. Medvedev, A.N., Delvenne, J.C., Lambiotte, R.: Modelling structure and predicting dynamics of discussion threads in online boards. Journal of Complex Networks p. cny010 (2018). DOI 10.1093/comnet/cny010

34. Mishne, G., Glance, N.: Leave a reply: An analysis of weblog comments. In: Proc. 3rd Annual Workshop on the Weblogging Ecosystem at the 15th International World Wide Web Conference, 2006 (2006)

35. Mojica, L.G.: Modeling trolling in social media conversations. arXiv preprint arXiv:1612.05310 (2016)

36. Morstatter, F., Pfeffer, J., Liu, H., Carley, K.M.: Is the sample good enough? comparing data from twitter's streaming api with twitter's firehose. In: ICWSM (2013)

37. Moyer, D., Carson, S.L., Dye, T.K., Carson, R.T., Goldbaum, D.: Determining the influence of reddit posts on wikipedia pageviews. In: Proceedings of the Ninth International AAAI Conference on Web and Social Media (2015)

38. Muchnik, L., Aral, S., Taylor, S.J.: Social influence bias: A randomized experiment. Science **341**(6146), 647–651 (2013)

39. Newell, E., Jurgens, D., Saleem, H.M., Vala, H., Sassine, J., Armstrong, C., Ruths, D.: User migration in online social networks: A case study on reddit during a period of community unrest. In: ICWSM, pp. 279–288 (2016)

40. Salganik, M.J., Watts, D.J.: Leading the herd astray: An experimental study of self-fulfilling prophecies in an artificial cultural market. Social psychology quarterly **71**(4), 338–355 (2008)
41. Sinatra, R., Lambiotte, R.: Topical issue-quantifying success. Advances in Complex Systems **21**, 3–4 (2018)
42. Singer, P., Ferrara, E., Kooti, F., Strohmaier, M., Lerman, K.: Evidence of online performance deterioration in user sessions on reddit. PloS one **11**(8), e0161,636 (2016)
43. 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: Proceedings of the 23rd international conference on world wide web, pp. 517–522. ACM (2014)
44. Stoddard, G.: Popularity dynamics and intrinsic quality in reddit and hacker news. In: ICWSM, pp. 416–425 (2015)
45. Stuck_In_the_Matrix: Dataset is available on the following webpage. https://files.pushshift.io/reddit/ (Query: 2017-06-01)
46. Stuck_In_the_Matrix: I have every publicly available reddit comment for research. approx. 1.7 billion comments @ 250 gb compressed. any interest in this? https://redd.it/3bxlg7 (Query: 2017-07-14)
47. Stuck_In_the_Matrix: Update for the reddit corpus. https://redd.it/8aen5g (Query: 2018-09-27)
48. Szabo, G., Huberman, B.A.: Predicting the popularity of online content. Communications of the ACM **53**(8), 80–88 (2010)
49. Tan, C.: Tracing community genealogy: How new communities emerge from the old. arXiv preprint arXiv:1804.01990 (2018)
50. Tan, C., Lee, L.: All who wander: On the prevalence and characteristics of multi-community engagement. In: Proceedings of the 24th International Conference on World Wide Web, WWW '15, pp. 1056–1066. International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, Switzerland (2015)
51. Tsagkias, M., Weerkamp, W., De Rijke, M.: Predicting the volume of comments on online news stories. In: Proceedings of the 18th ACM conference on Information and knowledge management, pp. 1765–1768. ACM (2009)
52. Wang, C., Ye, M., Huberman, B.A.: From user comments to on-line conversations. In: Proceedings of the 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '12, pp. 244–252 (2012)
53. Zannettou, S., Caulfield, T., Blackburn, J., De Cristofaro, E., Sirivianos, M., Stringhini, G., Suarez-Tangil, G.: On the origins of memes by means of fringe web communities. arXiv preprint arXiv:1805.12512 (2018)
54. Zayats, V., Ostendorf, M.: Conversation modeling on reddit using a graph-structured lstm. Transactions of the Association of Computational Linguistics **6**, 121–132 (2018)
55. Zhang, J., Hamilton, W.L., Danescu-Niculescu-Mizil, C., Jurafsky, D., Leskovec, J.: Community identity and user engagement in a multi-community landscape. In: Proceedings of the International AAAI Conference on Weblogs and Social Media, vol. 2017, p. 377. NIH Public Access (2017)
56. Zhao, Q., Erdogdu, M.A., He, H.Y., Rajaraman, A., Leskovec, J.: Seismic: A self-exciting point process model for predicting tweet popularity. In: Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pp. 1513–1522 (2015)