Spam, Spam, Spam, Spam: Methodological Considerations and Challenges for Studying Educators' Twitter Use

Jeffrey P. Carpenter
Elon University
United States
jcarpenter13@elon.edu

Matthew J. Koehler
K. Bret Staudt Willet
Michigan State University
United States
mkoehler@msu.edu
staudtwi@msu.edu

Spencer P. Greenhalgh
University of Kentucky
United States
spencer.greenhalgh@uky.edu

Abstract: Social media have come to play an important role in the professional lives of many educators. Platforms such as Twitter create new spaces in which collegial contact can occur, opening up various avenues for support and development. These spaces, however, can attract users whose behaviors creates challenges for researchers who try to understand and analyze educators' experiences on social media. In this paper, we describe the kinds of spam that we have encountered in our research specifically on educators' Twitter use. After summarizing the relevant extant literature, we discuss the benefits and drawbacks of various approaches to dealing with the digital noisiness of educator Twitter spaces. We define a number of practical metrics that researchers can utilize to help identify spam, and describe the process used in one study to identify and remove spam. We conclude by considering implications for educational researchers and teacher educators.

Introduction

The advent of the World Wide Web has created opportunities for educators to find online communities and spaces that can prove quite useful to them for professional learning that is personalized to their unique needs and interests (e.g., Dede, 2006). In particular, the microblogging service Twitter has proven to be attractive to a substantial number of educators for its professional applications (Carpenter & Krutka, 2014; Rosenberg, Greenhalgh, Koehler, Hamilton, & Akcaoglu, 2016; Xing & Gao, 2018). Education-related Twitter *hashtags* (i.e., words or phrases preceded by a hash or pound sign that organize content by topics) offer spaces for educators to connect and discuss a wide variety of education topics (Rosenberg, Greenhalgh, Koehler, Hamilton, & Akcaoglu, 2016; Carpenter, Tani, Morrison, & Keene, 2018).

Although online educator professional learning and networking is popular (Greenhow, Campbell, Galvin, & Askari, 2018; Lantz-Andersson, Lundin, & Selwyn, 2018; Macià & García, 2016), many facets of educators' digital activities remain under-researched and at best partially understood. One area that has received only scant attention in the literature is the presence of *spam* -- "unsolicited, repeated actions that negatively impact other people" (Twitter, n.d.) -- in online spaces utilized by educators. Because of the open nature of Twitter, the potential exists for various kinds of *spam* to affect the spaces and communities that educators access via the platform. Popular hashtags can be an attractive target for spammers because these hashtags concentrate attention from a considerable audience—attention that spammers can then attempt to redirect to other messages and content.

Researchers who seek to study educators' Twitter use and education-related hashtags must first consider how spam affects educators and their online spaces; then they must make decisions regarding how they will address the presence of such undesired content. These issues are important because spam content can change the appearance of patterns in the data and thereby obscure what is happening with real, everyday educators on Twitter. Spam content can have different purposes and patterns of interaction than educators; for example, spam content may not engage other users in terms of replies, retweets, or likes in the same way that other content does. As such, leaving spam content in a dataset for analysis may skew analysis of the level of interaction and community present among the "real" users of the hashtag—what researchers might consider to be average, normal, or typical usage. Yet, removing spam from the data being considered runs the risk of misrepresenting what is actually experienced by these everyday educators, and there may be compelling reasons to retain spam content in some studies.

The purpose of this paper, then, is to begin an exploration of the implications of spam data for research in educational hashtags. We offer a broad, yet practical strategy for educational researchers to identify and remove spam from data analyses.

Literature Review

In the following sections, we provide background on the relationship between educator-focused Twitter and spam, the kinds of spam that exist on Twitter, and implications of spam for participants and researchers.

Professional Learning on Twitter as an Affinity Space

Those describing educators' professional use of Twitter have frequently used the language of "communities" to describe the groups that organize themselves using this social networking service (e.g., Britt & Paulus). Brunton (2013) described spam in terms of—or in opposition to—community, making the concept of community important not only for the broad phenomenon of educators' use of Twitter but also for the specific theoretical and methodological issue of spam. Spam potentially has corrosive effects on communities, but its effects likely vary across the diverse landscape of Twitter hashtags and users (Carpenter et al., 2018). Furthermore, what constitutes a "community" in digital spaces is a subject of much debate (Preece & Maloney-Krichmar, 2005), and popular conceptions of community in the educational literature (e.g., the *community of practice*) are not always reflected in educators' Twitter practices (Greenhalgh, 2018).

In this paper, we approach educational hashtags using Gee's (2004) *affinity space* framework, which acknowledges the existence of social interactions in digital spaces without concern for whether users share a single identity or have strong interpersonal relationships. On Twitter, there is a very low barrier for entry into such spaces; for example, a hashtag (e.g., #Edchat) can serve as a *portal* (Rosenberg, Greenhalgh, Koehler, Hamilton, & Akcaoglu, 2016) to everyone else's contributions to the space. This ease of access has positive implications for participants looking for loose affiliation but also makes it easy for off-topic, irrelevant, or even malicious messages —in other words, spam—to enter the space (Greenhalgh, 2018). As a result, in an affinity space mediated by a Twitter hashtag, there are some forms of participation that are core to educators' use of Twitter as professional resource; however, there are also forms of participation (e.g., spam) that contribute little to and potentially detract from the professional purposes for which these spaces exist. To date, spam has received only very limited attention in the extant literature on educators' professional Twitter activities.

Types and Characteristics of Spam

The *existence* of spam in educator-focused, hashtag-mediated professional learning environments may not come as a surprise, but the *diversity* of this spam is noteworthy. Brunton (2013) described spam in broad terms, as "undesirable text, whether repetitive, excessive, or interfering" (p. xxii). This undesirability may take a number of different forms—consider, for example, the following characteristics that may distinguish legitimate messages from spam or more tolerable spam from egregious spam:

- *Offensiveness*: Spam tweets in a hashtag may be judged as such because they reference or include offensive material, such as disparaging remarks, ad hominem attacks, profane language, or pornography.

- *Commercial nature*: The motivations behind spam are often commercial in nature, and educator-focused hashtags might be used to bring participants' attention to particular products or services—some related to education and others not.
- *Relevance to the space:* Although the ostensible purpose of a Twitter hashtag is to index tweets on a specific topic, the open nature of hashtags makes it difficult to filter out tweets that are not related to the purpose of a space.
- *Intentionality*: It is possible for a user to employ an educator-focused Twitter hashtag in a tweet unintentionally, whether because of typos, a shared acronym, or other mistakes or miscommunications.
- *Personality*: It is possible to contribute to Twitter in a variety of impersonal ways, including using fake accounts, using real accounts to automatically tweet material from other sources, or operating automated agents, commonly known as *bots* (Mowbray, 2014).

Individual tweets that might qualify as spam can include various combinations of these characteristics, which will be explained further in the section that follows.

How Participants Experience Spam

Donath (2007) described how spam complicates participant decisions regarding which social ties to forge—in the case of Twitter, who to follow and which hashtag spaces to frequent. Similarly, spam content complicates a discernment of desirable, legitimate, and quality content; sorting through tweets can become more and more time consuming. Furthermore, spam can lead not just to inefficiency, but also to "disillusionment with the entire experience" (Donath, 2007, p. 239). In the case of social media platforms such as Twitter, spam may "change the social atmosphere, causing people to become more suspicious" (Donath, 2007, p. 239). The effects of spam upon user experiences may therefore be an important consideration for researchers investigating educator Twitter use.

Participants in an educator-focused Twitter hashtag space might have varied responses to different kinds of spam based on their evaluation of the previously discussed characteristics of spam. For example, Greenhalgh, Rosenberg, & Wolf (2016) identified a number of tweets using the #MAET hashtag that discussed food rather than the Master of Arts in Educational Technology program that was the ostensible focus of the hashtag—however, it became clear that some of these "spam" tweets were due to the Danish word for "full" being "mæt." While educational technology-focused participants might still find such messages irritating, they might not be as problematic as contributors who post political conspiracy theories to busy hashtags like #Edchat; although irrelevant, the foodie #MAET tweets do not appear to be intentional and may not leave other users feeling that bad actors are significantly detracting from the space. Similarly, Greenhalgh (2018) identified two tweets that used educational hashtags to share commercial information—however, because one was an educational workshop and the other was a handbag being sold on eBay, the first was more relevant to the hashtag spaces in question and likely to be judged differently. Indeed, participants in Twitter spaces that are closely linked to particular books (e.g., Carpenter et al., 2018) or technologies (e.g., #gafe) might appreciate commercial messages rather than consider them an indication of spam.

On the other hand, a tweet might be considered by some users to be inoffensive, non-commercial, relevant to the hashtag topic, correctly using the intended hashtag, and personal—yet still be considered spam because it is still undesirable. In particular, this may occur in the context of the synchronous Twitter chats that are associated with many education-focused hashtags. These chats are typically one-hour long moderated conversations that are structured around pre-determined prompts. The pace of some chats can be quite rapid; for example, participants in Britt and Paulus (2016)'s research on the weekly synchronous #Edchat described their experience of #Edchat as " just absolutely ridiculously fast" and "overwhelming for a new person on Twitter" (p. 56). Furthermore, in the context of such a chat, users may be more likely to find some content undesirable. For example, a tweet composed during a chat that is related to the general topic of the hashtag but not the specific theme of that week's synchronous chat, might be considered an unwelcome distraction by some participants.

How Researchers Respond to Spam

Like participants in educator-focused Twitter hashtags, researchers of these social spaces may respond to spam in a variety of ways, depending on the nature of the content and their research questions. For certain research questions, such as those related to participants' experiences with an education-related Twitter hashtag, spam and those participants producing it might be considered an essential part of the dataset. Indeed, the spam itself may even be of interest—for example, if teachers in a particular hashtag are being criticized through *offensive* messages using their hashtag, researchers may wish to call attention to this abuse to comment on broader social issues (e.g. Veletsianos, Houlden, Hodson, & Gosse, 2018). Similarly, scholars may be interested in the way that educators use social media to sell products of their own design (Carpenter, Abrams, & Dunphy, 2017; Shelton & Archambault, 2019), even if other participants find this self-promotion to be frustrating.

However, spam can also present "serious challenges for the researcher. We can rarely be sure whether our findings are an artifact of a flawed dataset" (Karpf, 2012, p. 642). For research questions seeking to describe "normal" behavior or compare activity between different hashtags, the main concern of the researcher is ultimately the *relevance* of all of the messages in the collection that they are studying to their phenomenon of interest. In these cases, spam skews statistical results in a way that introduces bias and can be described as "noise" (Kwak, Lee, Park, & Moon, 2010).

A Practical Approach to Dealing with Spam in Educational Twitter Research

Throughout the remainder of this paper, we focus on scenarios in which education researchers chiefly face spam as a challenge to their research and must identify and make decisions regarding the removal of spam within their research data. To date, the science of spam detection has guided prior research into automated approaches that identify individual spammers at the user level, or tweets as spam (Wang, Zubiaga, Liakata, & Procter, 2015). Such computational approaches to spam detection seek to solve a *classification* problem—the goal is to create an algorithm that correctly classifies users or tweets as spam or not spam. These approaches are resource intensive, requiring access to specialized training in machine learning, large training datasets of previously classified tweets, and computational resources (e.g., Lin & Huang, 2013; Wang et al., 2015). When considering machine-learning approaches for dealing with spam in educational social media research, two problems present themselves. First, spammers and spam-detection algorithms are in constant state of evolution—the newest spam detection approaches become less effective as spammers devise new approaches to circumvent the latest detection algorithms (Chen, Zhang, Xiang, Zhou, & Oliver, 2016). Second, the machine-learning algorithm approach is beyond what is currently readily available to many educational researchers, who lack access to and training in the latest computer science methods. In contrast to the complex and resource-heavy machine learning approaches to spam detection, we argue here for an approach that focuses on practicality for educational researchers—first we consider metrics, then we focus on using these metrics in holistic decision-making.

Practical Metrics for Educational Researchers

Our practical approach for educational researchers starts with a consideration of some simple, easy-to-compute metrics that are readily available to educational researchers. These metrics, although not direct indicators of spam themselves, can be used in conjunction with each other to help researchers make informed decisions about which data to include or not include in analysis. This list presented below is not meant to be exhaustive, but rather, to serve as a starting point for discussion of the possibilities for practically dealing with spam. These metrics focus primarily on identifying spam at the user level.

- *Volume of tweeting*: One indicator of spam is high-volume tweeting that is often-bot generated. Practical indicators of spam include counts of the raw number of tweets, the percentage of tweets to a hashtag accounted for by a user, or more standardized metrics such as *z*-scores of tweets per user.
- *Level of interaction*: Spam can also be suggested by the lack of interaction with others, as spammers tend to broadcast messages, which others tend to ignore (Lin & Huang, 2013). Practically speaking, a relatively easy metric for researchers to measure interaction is to examine the extent to which a users' tweets result in likes, retweets, and replies.
- *Following vs. followers*: Spammers have often been characterized by their tendency to follow many other users, but have relatively low number of followers in return (e.g., Yang et. al, 2011). Researchers can

quickly measure this phenomenon by calculating the ratio of following to followers for users in their dataset.

- *Level of hyperlinking*: Many spammers love to share hyperlinks (e.g., Lin & Huang, 2013)! For example, a spam tweet might advertise goods for sale and include a hyperlink to the web site where the actual purchase would occur. To identify potential spam accounts, researchers can count the raw number of links, the percentage of tweets that contain a link, or the average number of links per tweet.
- *Bot detection*: Automated accounts (or "bots") are responsible for much of the commercial spam in educational social media. Researchers and Twitter (the company) are constantly trying to develop bot-detection techniques (Chen et al., 2016), while at the same time, bot-creators are racing to change techniques to avoid detection. A relatively simple and accessible tool to educational researchers' is Kearney's (n.d.) *TweetBotOrNot* tool, that examines a users' Twitter activity and gives the probability that the user is a bot.
- *Profile information*: Twitter profiles typically include qualitative information that can be helpful in determining whether users are educators or have education-related purposes. A review of profiles will also reveal if an account is currently suspended, a strong indicator of prior spamming activities.

It is important to note that these metrics do not directly align to the types and characteristics of spam that we outlined earlier. Characteristics such as *offensiveness* and *relevance* are often judgement calls that require interpretation and context not directly detectable in large collections of Twitter data. However, *offensive, irrelevant, commercial,* and *impersonal* spam often leaves behind a footprint that is reflected in the above metrics. In the following section, we describe how researchers can apply these metrics to identify potential sources of spam, but additionally use judgement, reason, and context to make the final determination on what is, or is not, spam.

Holistic Decision-Making with the Metrics

In computational approaches to spam-detection, machine learning algorithms rely upon the collected metrics to make decisions about which content should be considered spam. In contrast, we suggest using human-raters who look across the metrics we have described above to make holistic decisions about what data should be excluded or not from educational datasets.

This decision-making process is not one-size fits all. Whether or not to exclude data from analysis should depend upon the research questions being asked, the nature of the data itself, and level of certainty in the decision being made. For example, research questions that focus on describing the full range of experiences, uses, and purposes of an educational hashtag are likely to exclude very little data from analysis, because spam unfortunately can be part of the day-to-day experience of educational uses of Twitter. In contrast, if the research question is focused more on the professional interactions that can happen on Twitter, spammers who do not actually interact with other users could easily be removed from data analysis, relying primarily upon the aforementioned *level of interaction* metric.

An Example of this Practical Approach in Use

In recent research, the Carpenter and colleagues (2018) conducted an exploratory study of 16 education-related Twitter hashtags. This comparative study sought to map the terrain of educational Twitter spaces by contrasting the traffic associated with these various hashtags. The comparisons across hashtags were complicated by the differing nature and levels of spam associated with the various hashtags. Although some hashtags appeared to feature little undesired content, the traffic on several others included considerable amounts of spam. For example, two hashtags that regularly attracted tweets from more than one thousand educators per month both included an individual user who accounted for more than 20% of all tweet traffic during the data collection period. The presence of such outliers made it challenging to present quantitative comparisons of the hashtag traffic that accurately reflected the nature of educators' use of those hashtags.

Using the practical approach described in the previous section, Carpenter and colleagues (2018) conducted spam removal that used the following metrics: (a) *volume of tweets*, (b) *profile information,* (c) *number of hyperlinks,* and (d) *interaction with others*. Using the *number of tweets*, they identified the 10 most active contributors to each of the 16 hashtags (n = 160) for further scrutiny. For each of these, the researchers read the users' *profile information* (including if the account was currently suspended or not) and the users' recent tweets to assess if the users' Twitter presence appeared to be consonant with the intended purposes of the hashtags to which they posted. It was determined that 28 of these contributors required additional scrutiny to decide if they merited removal from the data set.

For these 28 users, Carpenter and colleagues created a spreadsheet containing the following data: (a) if the account was suspended or not (e.g., seven were suspended); (b) the percentage of traffic accounted for by the user based upon *volume of tweeting*; (c) the level of *interaction with others* based upon the percentage of users' tweets that received retweets (e.g., six users posted more than 1000 tweets to the studied hashtags but had received zero retweets); (d) the *frequency of hyperlinks* (e.g., some users had a link in every tweet); and (e) the probability of being a bot as determined by Kearney's (n.d.) *TweetBotOrNo*t tool. Using these metrics, the research team holistically looked at each user to decide whether or not their data should be excluded from the analysis. Of the 28 users selected for further scrutiny, 24 were ultimately removed from the analysis. These 24 users participated in six of the 16 hashtags—10 of the hashtags data sets were thus unaffected by this spam removal process.

Ultimately, this removal of users allowed for the Carpenter and colleagues to compare and contrast the professional educator traffic on these hashtags in a variety of ways that were less skewed by spam and were aligned with their research questions. For example, an initial analysis suggested that #BFC530—a hashtag that hosts a brief synchronous education-focused chat every Monday to Friday at 5:30 a.m. EST—featured quite different patterns of use from the other hashtags being studied. Prior to spam removal, it appeared that #BFC530 hashtag attracted an average of approximately 26 tweets per month per user, while the other 15 hashtags averaged between 3-10 tweets per user. Such a difference between hashtags could have been interpreted as suggesting that #BFC530 was attracting a level of intensity of posting that was dramatically different from other hashtags. However, after the removal of spam, the tweets per user number for #BFC530 dropped to just over 8.5 per month, no longer outside of the range of other hashtags studied. Spam removal thus allowed for comparisons of the 16 education-focused hashtags that were more accurate considering the study's research questions. We offer this example not with the intention of recommending a static or rigid one-size-fits-all approach to spam removal for educational researchers, but rather as an opportunity for other scholars to consider appropriate strategies to take regarding the spam they might encounter in their own studies.

Discussion

Depending on their research questions and their data, researchers will make varied choices regarding whether to include or remove spam and can be flexible in how holistic decisions are made. Researchers can use some or all of the metrics that we have defined above and may consider additional metrics as well. In time, computational methods of spam detection may become more effective at identifying spam content in educational spaces and may become more accessible to a wide variety of educational researchers. Twitter and other social media services may also become more effective at limiting certain kinds of spam. In the immediate future, however, we see value in many education research contexts in utilizing a combination of metrics and a final holistic human decision to identify spam and make decisions about its removal. The diverse purposes and motivations of users who access education-focused Twitter hashtags (e.g., Carpenter & Krutka, 2014) and the many flavors of spam that can be found in those spaces can confound purely algorithmic approaches to spam identification and removal.

One fundamental challenge in the identification of spam that complicates the endeavor and any decisions regarding the removal of users and/or tweets from a data set is the reality that in some cases spam is in the eye of the beholder. While tweets to an education hashtag that include hyperlinks to pornography or that sell counterfeit handbags are relatively easy to characterize as spam, there are many instances in which it is less straightforward whether content should be considered spam or not. For example, while some educators may see tweets that advertise education products as spam that detracts from their Twitter experience, other educators may find value in some of the products being advertised, or simply accept such tweets as a minor nuisance associated with utilizing a commercial platform. The consequences of spam may thus vary depending on different users' perspectives. Some educators may find it an insignificant burden to block or mute a few prominent spammers, and depending on which hashtags they use, this might be all that is necessary to render spam a non-factor in their Twitter experience. Other users might instead consider the presence of spam to be toxic and therefore choose to disengage from hashtags that attract more spam.

Although online spaces that attract a larger number of participants may offer some benefits to educators, such as access to a larger number of resources, perspectives, and potential new colleagues (Xing & Gao, 2018), bigger may not always be better (Greenhalgh, 2018), particularly in light of concerns related to spam. Indeed, spam could be considered the inevitable abuse of the same technologies that allow such large communities to form in the first place (Brunton, 2013). Educators who are new to Twitter may be discouraged from continuing their use of the platform if they first encounter only the more popular spaces where they must wade through more spam. Furthermore, the quantity and pace of content in some Twitter spaces can be daunting for many users, even if individual messages do not otherwise appear to have the qualities of spam. Twitter hashtags that attract a smaller quantity of tweets and users may in some instances be richer spaces in part because of how they may be less likely to attract spam.

The example of spam removal we offered above focused on identifying spam at the user level, in part because of the large volume of spam tweets we initially noticed coming from a small number of users. It may be that for some education hashtags and/or research questions, identifying spam at the tweet level is more relevant or important. For example, some teachers use social media such as Twitter to engage in *teacherpreneurship* (Shelton & Archambault, 2019). Researchers investigating this phenomenon might find teachers who sometimes tweet content that is oriented towards professional learning and networking, but who at other times tweet advertisements for their products on a lesson marketplace such as TeachersPayTeachers.com. While the former type of tweets is thus unlikely to be interpreted by other users as spam, the latter might not be welcomed by some other educators.

Implications for Teacher Education

Given that various teacher education programs have their pre-service teachers (PSTs) use Twitter for educational and professional purposes (e.g., Carpenter & Morrison, 2018; Damico & Krutka, 2017; Krutka, Nowell, & Whitlock, 2017), teacher educators too should be aware of how the presence of spam affects educational spaces on Twitter. PSTs can fail to make the kinds of connections with other educators intended by the teacher educators who task them with exploring Twitter (Hsieh, 2017). If PSTs and even inservice educators encounter large quantities of spam, they may fail to recognize the benefits of professional networking and collaboration that Twitter can potentially facilitate. Teacher educators may thus want to prepare teachers for navigating the presence of a range of different kinds of spam in educator-focused spaces on Twitter. Given the diversity of education-focused hashtags that exist, teacher educators might also want to be strategic in identifying and recommending certain hashtags that feature less spam to their PSTs. In particular, given that research outside of professional K-12 educator spaces has suggested that Twitter can be rife with misogyny and racism, teacher educators must also consider how to prepare PSTs with critical social media literacy (Nagle, 2018). PSTs could also benefit from discussion of the benefits and drawbacks associated with teacherpreneurs' use of education-focused hashtags to market their wares (e.g., Shelton & Archambault, 2019). PSTs understandably may be drawn to content shared by practicing teachers, and they may thus need guidance regarding how to recognize and evaluate more subtly commercial content. As previously noted, while some offensive and irrelevant content would clearly be considered spam by the overwhelming majority of educators who frequent Twitter, in other cases it may be a more subjective matter whether content is deemed relevant to a particular affinity space.

Future Research

As previously noted, the extant research has paid only very limited attention to the presence and effects of spam in online educator spaces and communities. While this paper aims to further the discussion of spam, the ground is fertile for further work in this area. Future studies could investigate whether the impact of spam on educators may vary according to a variety of user demographics, such as gender, race, age, years of experience teaching, and years of experience on Twitter, among others. Researchers might also choose to explore the effects of spam at the level of the community or space. For example, do some hashtags seem to maintain their character and thrive despite the presence of spam, while other hashtags appear to suffer more harmful consequences? Given how increasingly common it is for individuals to utilize multiple social media, research that compares and contrasts spam across platforms could also be beneficial to the field. Furthermore, we recommend that researchers who study online educator spaces and communities consider at least acknowledging the ways in which spam may affect their data and any methodological steps they employ to mitigate those impacts, even if spam is not a main focus of their research.

Conclusion

With large quantities of educator professional learning, networking, and collaboration occurring online, the spam that can exist in online spaces is a matter of importance for educational researchers who study educators' professional lives in a digital era. Spam likely impacts—at least to some degree—how educators use and experience social media; it also presents challenges to researchers as they collect and analyze data from such media. For many research questions, spam makes the phenomenon of interest less clear; for example, when trying to describe normal or typical behavior or participation in an online space. In this paper, we have made the case that for many research questions, the presence of spam should be addressed. For this purpose, we have provided practical metrics and a holistic decision-making approach that should be accessible to and useful for a broad range of educational researchers. In the end, we advocate for a human approach to decision making about spam in educational research and that researchers should be mindful of the varieties of spam and the diverse ways in which educators may interpret and respond to its existence.

References

Britt, V. G., & Paulus, T. (2016). "Beyond the four walls of my building": A case study of #Edchat as a community of practice. *American Journal of Distance Education*, *30*, 48–59. doi:10.1080/08923647.2016.1119609

Brunton, F. (2013). *Spam: A shadow history of the Internet*. Cambridge, MA: MIT Press.

Carpenter, J., Abrams, A., & Dunphy, M. (2016, March). Educators' professional uses of pinterest. In G. Chamblee & L. Langub (Eds.), *Society for Information Technology & Teacher Education International Conference* (pp. 1925-1930). Chesapeake, VA: Association for the Advancement of Computing in Education (AACE).

Carpenter, J. P., & Krutka, D. G. (2014). How and why educators use Twitter: A survey of the field. *Journal of Research on Technology in Education*, *46*(4), 414-434. doi: 10.1080/15391523.2014.925701

Carpenter, J. P., & Morrison, S. A. (2018). Enhancing teacher education… with Twitter?. *Phi Delta Kappan*, *100*(1), 25-28. doi: 10.1177/0031721718797118

Chen, C., Zhang, J., Xiang, Y., Zhou, W., & Oliver, J. (2016). Spammers are becoming "smarter" on Twitter. *IT Professional*, *18*(2), 66–70. doi:10.1109/MITP.2016.36

Damico, N., & Krutka, D. G. (2018). Social media diaries and fasts: Educating for digital mindfulness with pre-service teachers. *Teaching and Teacher Education*, *73*, 109-119. doi: 10.1016/j.tate.2018.03.009

Dede, C. (2006). *Online professional development for teachers: Emerging models.* Cambridge, MA: Harvard Education Press.

Donath, J. (2007). Signals in social supernets. *Journal of Computer-Mediated Communication*, *13*, 231-251. doi: 10.1111/j.1083-6101.2007.00394.x

Gee, J. P. (2004). *Situated language and learning: A critique of traditional schooling*. New York, NY: Routledge.

Greenhalgh, S. P. (2018). *Spaces and their social frontiers: Using community dimensions to distinguish between teacher-focused hashtags on Twitter* (Doctoral dissertation). Michigan State University, East Lansing MI.

Greenhalgh, S. P., Rosenberg, J. M., & Wolf, L. G. (2016). For all intents and purposes: Twitter as a foundational technology for teachers. E-Learning and Digital Media, 13(1-2), 81-98. doi:10.1177/2042753016672131

Greenhow, C., Campbell, D., Galvin, S., & Askari, E. (2018, March). Social Media in Teacher Professional Development: A Literature Review. In E. Langran & J. Borup (Eds.), *Proceedings of Society for Information Technology & Teacher Education (SITE) International Conference 2018* (pp. 2256-2264). Waynesville, NC: Association for the Advancement of Computing in Education (AACE).

Hsieh, B. (2017). Making and missing connections: Exploring Twitter chats as a learning tool in a preservice teacher education course. *Contemporary Issues in Technology and Teacher Education*, *17*(4), 549-568.

Karpf, D. (2012). Social science research methods in Internet time. *Information, Communication & Society*, *15*, 639-661. doi: 10.1080/1369118X.2012.665468

Kearney, M. W. (n.d.). {TweetBotOrNot} [Web application.] Retrieved from https://mikewk.shinyapps.io/botornot/

Krutka, D. G., Nowell, S., & Whitlock, A. M. (2017). Towards a social media pedagogy: Successes and shortcomings in educative uses of Twitter with teacher candidates. *Journal of Technology and Teacher Education*, *25*(2), 215-240.

Kwak, H., Lee, C., Park, H., & Moon, S. (2010). What is Twitter, a social network or a news media?. In *Proceedings of the 19th International Conference on World Wide Web (WWW 2010)* (pp. 591-600). Association for Computing Machinery (ACM).

Lantz-Andersson, A., Lundin, M., & Selwyn, N. (2018). Twenty years of online teacher communities: A systematic review of formally-organized and informally-developed professional learning groups. *Teaching and Teacher Education*, *75*, 302-315. doi: 10.1016/j.tate.2018.07.008

Lin, P.-C., & Huang, P.-M. (2013). A study of effective features for detecting long-surviving Twitter spam accounts. In *The 15th International Conference on Advanced Communications Technology (ICACT), Technical Proceedings 2013* (pp. 841–846). Piscataway, NJ: IEEE.

Macià, M., & García, I. (2016). Informal online communities and networks as a source of teacher professional development: A review. *Teaching and Teacher Education*, *55*, 291-307. doi: 10.1016/j.tate.2016.01.021

Mowbray, M. (2014). Automated twitter accounts. In K. Weller, A. Bruns, J. Burgess, M. Mahrt, & C. Puschmann (Eds). *Twitter and society* (pp. 183-194). New York, NY: Peter Lang.

Nagle, J. (2018). Twitter, cyber-violence, and the need for a critical social media literacy in teacher education: A review of the literature. *Teaching and Teacher Education, 76*, 86-94. doi: 10.1016/j.tate.2018.08.014

Preece, J., & Maloney-Krichmar, D. (2005). Online communities: Design, theory, and practice. *Journal of Computer-Mediated Communication, 10*(4). doi:10.1111/j.1083- 6101.2005.tb00264.x

Rosenberg, J. M., Greenhalgh, S. P., Koehler, M. J., Hamilton, E. R., & Akcaoglu, M. (2016). An investigation of state educational Twitter hashtags (SETHs) as affinity spaces. *E-learning and Digital Media*, *13*(1-2), 24-44. doi: 10.1177/2042753016672351

Shelton, C., & Archambault, L. (2019). Discovering how teachers build virtual relationships and develop as professionals through online teacherpreneurship. *Journal of Interactive Learning Research*, *29*(4), 579-602.

Twitter (n.d.). Report spam on Twitter. Retrieved from https://help.twitter.com/en/safety-and-security/report-spam

Veletsianos, G., Houlden, S., Hodson, J., & Gosse, C. (2018). Women scholars' experiences with online harassment and abuse: Self-protection, resistance, acceptance, and self-blame. *New Media & Society*, *20*(12), 4689–4708. doi:10.1177/1461444818781324

Wang, B., Zubiaga, A., Liakata, M., & Procter, R. (2015). Making the most of tweet-inherent features for social spam detection on Twitter. *CEUR Workshop Proceedings*, *1395*, 10–16. Retrieved from https://arxiv.org/abs/1503.07405

Xing, W., & Gao, F. (2018). Exploring the relationship between online discourse and commitment in Twitter professional learning communities. *Computers & Education, 126*, 388-398. doi: 10.1016/j.compedu.2018.08.010

Yang, C., Harkreader, R. C., & Gu, G. (2011, September). Die free or live hard? Empirical evaluation and new design for fighting evolving Twitter spammers. In R. Sommer, D. Balzarotti, & G. Maier (Eds.), *Proceedings of the 14th International Symposium on Recent Advances in Intrusion Detection (RAID)* (pp. 318-337). Menlo Park, CA: ACM. doi:10.1007/978-3-642-23644-0_17