Are Violations of Student Privacy “Quick and Easy”? Implications of K-12 Educational Institutions’ Posts on Facebook

Macy A. Burchfield¹, Joshua M. Rosenberg¹, Conrad Borchers², Tayla Thomas¹, Benjamin Gibbons³, Christian Fischer²

¹ University of Tennessee, Knoxville

² University of Tübingen

³ Emory University

ABSTRACT

As the use of social media increases in daily life, it has also increased for institutions in the field of education. While there may be benefits for schools to use this media outlet, the privacy of students within those schools may be at risk when their names and photos are shared on such a publicly accessible domain. In this study, we analyzed the extent to which students' privacy is protected by qualitatively coding a random sample of 100 Facebook posts made by U.S. school districts from a population of over 9.3 million photo posts that we collected. Using inferential techniques, we found that students are somewhat protected compared to teachers and community members, with only 2.67% of students' detected faces able to be identified by name. The same measure for staff and community members was 4.6% and 16%, respectively. These numbers at first appear small, but if applied to the entire population, this could potentially leave between 153,218 and 1,153,844 students identifiable to anyone on the internet. We discuss the severity and scale of these privacy threats and make recommendations for research on student privacy in social media and other informal education-related contexts.

Keywords

Privacy, Social Media, Facebook, Educational Institutions, Facial Recognition

1. INTRODUCTION & PRIOR RESEARCH

As the number of people using social media has increased, the risks to the privacy of social media users have also increased [23], and this is particularly true since social media use expands into areas of our lives that it did not previously occupy. Education is one such domain in which social media use is now widespread [2, 10, 11, 13, 21, 22]—and is one domain for which the privacy risks from social media use, in general, may be compounded because of the centrality of a particularly vulnerable population, minors at school.

For students in any given school district, the use of their name or face for social media may present notable privacy concerns. As many social media posts are made publicly available, they may be accessed by unexpected sets of individuals, even by those without an account on the corresponding social networking site. Such use may pose a legitimate threat that may be unknown to (or under-acknowledged by) teachers, administrators, and parents.

There is past research on the intersection of privacy and social media. For example, Fiesler and Proferes [6] examined what participants in social media studies thought of their data being used by others—particularly, by academic researchers. Only around one-quarter of participants in their survey study reported being comfortable with their data being used without being informed of such use.

Related lines of research explored the intersection of privacy and social media data for students. For instance, Ifenthaler and Schumacher [9] surveyed students about what they thought of their data being used in learning analytics systems. They found that while students expressed comfort with sharing some types of data (i.e., data on their course enrollments, for which less than 20% of students reported reluctance with sharing such data), for others, students were much less comfortable. Notably, highly personal data, such as medical records, data on one's personal income, and externally-produced data, including social media, were among those that students were the least willing to share. Less than 10% of students reported being willing for externally-produced data to be used within learning analytics systems. Other scholars have shown that pre-service teachers are highly uncomfortable with how social media companies use students' social media data, with more than two-thirds of teachers expressing discomfort with such uses [16].

While past research has explored the willingness of social media participants and students to share their data for research, a different—institutional rather than personal—context for social media use presents potentially notable privacy risks. Namely, past research has shown that both post-secondary [13] and K-12 educational institutions use social media extensively; particularly, Twitter and Facebook [10, 11]. However, to this point, no research has yet investigated privacy in the context of social media use by K-12 educational institutions.

This topic—K-12 institutions' use of social media from a privacy perspective—is relevant and timely for a number of reasons. Recent research has shown that institutions are very active on both Twitter and Facebook, being associated with more than 300,000 posts/month from the accounts of K-12 districts and schools [11]. As a consequence, there could be hundreds of thousands of students with their identities being posted in a highly public, searchable, persistent record, and in a way that could be misused in the future. In addition, these posts may contain information that would typically be thought of as information that should not be shared publicly and widely, but which may be shared because of limited understanding of how widely such posts (on public pages) can be viewed. The audiences of institutions are likely much greater than that of individual educators, meaning any potential privacy concerns may be much larger than that of independent users' accounts. Raising awareness of this issue may prompt some reflection on the part of those sharing this information.

This study involves an initial investigation into the extent to which students' privacy is protected through analysis of Facebook posts made by public schools and school districts. In doing so, we ask a question about the nature of Facebook. Facebook claims that its site

is “quick and easy,” [5] but the expediency and facility with which K-12 administrators and educators may use the platform may mean that it is also easy for school districts and schools to violate the privacy of students—with potentially difficult-to-anticipate negative ramifications at present and in the future. In particular, we aim to explore the degree to which the privacy of students might be compromised through public Facebook posts guided by the following research questions:

1. To what extent can students be identified by name and photo on public Facebook pages of schools and school districts?
2. How does the identifiability of students compare to that of staff and community members?

2. METHOD

2.1 Sample

We used a *public data mining* methodology, one that draws from *educational data mining* techniques [1, 7], but which is distinguished by the use of (largely unstructured) publicly available data, such as data from websites and social media platforms [12]. Specifically, to obtain our sample of 100 schools’ and school districts’ Facebook posts, we used CrowdTangle, Facebook’s platform for providing academics and journalists access to data about public content on Facebook, including the content of posts and links to associated media as well as their timestamps and number of comments and likes (and other interactions) [4]. This content includes historical data from public Facebook pages with more than 50,000 likes and verified profiles. In addition, individuals with access to CrowdTangle can access public pages—but not individual users’ pages.

We accessed all of the posts from K-12 institutions’ public Facebook pages in the United States, having obtained the URLs to 15,728 educational institutions’ Facebook pages. We did so by using the statistical software R [20] to programmatically access (or, to web scrape) their homepages using data provided by the Common Core of Data [19], and recording all links to Facebook pages from their home pages. When schools linked to the same page as the district, we considered the page as a district page. The total study population included roughly 18 million posts shared from 2005-2020, with about 9.3 million of these posts including at least one photo.

Carrying out a privacy-focused study ourselves, we took steps to protect the privacy of the individuals represented in our data. First, while we accessed and structured the data in a PostgreSQL database, we did not save the images themselves, instead using the Facebook posts and links therein to access the images through our web browser. More broadly, we determined early in our process that we were not prepared to analyze the photos algorithmically/automatically in a safe and ethical manner (e.g., using machine learning methods); we were concerned about uploading the images to a server, where they might be scanned and indexed. While we did not store the images in our database, we nevertheless took steps to protect this data, including permitting access only to authenticated members of the research team.

From the population of approximately 18 million posts, we randomly sampled 100 posts with photos for this analysis. Our random sample of posts and related coding data were stored in a private Google Sheets file stored within a University Google Account (in part because Google is less likely—based upon past

legislation, lawsuits, and company policies—to programmatically search the contents of educational accounts) to which only project contributors had access; this ensured that any data that could potentially be used to identify individuals was protected.

2.2 Measures

We analyzed the data qualitatively using a combination of two commonly-used qualitative analysis techniques [8], the use of priori codes that we developed based upon prior research and our research questions as well as an exploratory process that allowed us to elaborate on and substantiate those codes and to train as coders on the use of the coding frame. In particular, we analyzed the data in two ways, as we describe next.

First, to determine *whose privacy was at risk* using our sample of 100 posts, we accessed the images from each post through photo-specific URLs that are included along with information for each post in the data. Each image was accessed and analyzed individually. When there were more than ten images included in a post, we analyzed the first ten, reasoning that these first 10 were the most likely to be seen by viewers of the post. Each post of our sample was analyzed by two trained coders to evaluate the levels of identification for all names and faces included. Upon analysis of 15 posts, we drew three categories from similar research to distinguish individuals included in posts based on their role in the school or school district community [18]:

- *Students*: Any minor assumed to be enrolled in a school and/or participating in a school-hosted event or activity.
- *Staff*: Any known employee of the school or school district; including but not limited to teachers, administrators, paraprofessionals, and communications directors.
- *Community Members*: Any member of the school community who is not a verifiable student or staff member, including but not limited to parents, school board members, local business owners, and volunteers.

Second, to determine *how individuals’ privacy may be threatened*, we developed a coding frame that we used to assess whether individuals’ names and/or photos of individuals were shared in posts and whether it was possible to readily connect individuals’ names and photos of them. We will next describe our qualitative coding process for applying this coding frame.

2.3 Qualitative Coding

Coding proceeded by first determining the classification (student, staff, or community) of each individual detected by name or photo in a post, and then identifying the number of different first and last names included in the text of the post, as well as the number of individual faces shown in the posts’ images. In particular, the following four elements were recorded for each category of individuals:

- *Number of First and Last Names in Post*

First and last names were recorded separately within each category due to the fact that staff and community members are often mentioned using their professional prefix (Mr., Mrs., Dr., etc.) and only their last name.

- *Number of Faces in Images*

For identifying the presence of individuals’ faces, a detectable face was considered to be one for which three out of four of the following features were visible without enlarging the image: 1)

eyes, 2) nose, 3) ears, and 4) mouth. Any faces appearing in more than one photo within the entire post were only counted once.

- *How many Names and Faces Connected*

We looked in posts for specific indicators of an individual's location in an image, including the order in which individuals appear in an image or labels on images. In general, identifiability criteria appeared as any text that explicitly stated which name matched with which face in which image.

Our coding included an interrater reliability check for 15 posts. Two coders coded these 15 posts individually using the coding process outlined above. Agreement percentages for detecting names, detecting faces, and identifying faces were 100%, 77.77%, and 93.33% respectively. Total agreement between coders across all codes was 92.34%.

2.4 Illustration of the Coding Process

Image 1. Example Posts

To illustrate the coding process, we provide two example posts and how we coded them above (Image 1). In the image for the first example, two student names, both first and last, are included in the text of the post, and multiple student faces are included in the three images of the post. The “third and fourth graders playing soccer” are not named individually and cannot be distinguished from each other. The two listed student names, which have been covered along with their faces for their protection, are identified by their locations in the images, thus making their faces identifiable by name as well. In the second image, the post included the name of a staff member, as well as detectable faces of two staff members. Without clarification, neither of these faces could be identified with the name mentioned.

2.5 Inferential Analysis

To analyze data to answer our first question, on how students can be identified by name and photo on public posts by K-12 educational institutions, we evaluated the percentages of student faces that were able to be identified by name; for example, if, across the 100 posts, we detected 50 student faces in images, and one was identifiable by name, then the percentage of identifiable students would be 1% (rather than 2%, because we were interested in making inferences based on the number of identifiable students per post). We refer to this value in our results as the *percentage of identifiable faces per post*. Then, based on the observed frequencies (from which we calculated these percentages), we calculated binomial 95% confidence intervals for the ratio of identifiable faces and categories of faces. We did this to present an initial set of

estimates for how many faces in our population of 9.3 million photo posts may be identifiable.

To answer our second research question on relative differences in identifiability of individuals from different groups, we carried out the same analysis as above (for students) for teachers and community members. Then, to compare the percentages of photos with identifiable individuals across categories, we calculated a different percentage than for RQ #1, one based not upon the number of posts (i.e., one identifiable face across 100 posts; 1%), but, rather, one based upon the total number of faces detected for people in each category. For instance, if there were 50 faces of students detected, and one was identifiable, then the percentage would be 2%; we refer to this in our results as the *percentage of identifiable faces per category sum*. This number—and comparing the confidence intervals between groups—would allow us to speak to whether individuals were differentially identifiable when photos of them were detected, even if there were, for example, far more photos of students than community members detected.

3. RESULTS

Our coding resulted in the detection (but not identification) of 299 faces in the images from the 100 posts in our sample. Of these 299 faces, only 13 (4.35% of all detected faces [2.33%, 7.32%]) were able to be identified with the individuals' names from the text of the post.

RQ #1. These 13 identifiable faces were identified within 12 individual posts from schools or districts. Student faces comprised 5 of those 12 and thus, for every 100 posts, we estimated that there were 5 identifiable student faces, representing the rate of a single identifiable student face for every twenty posts. Put another way, we estimated that 5% ([1.64%, 11.28%]) of these posts contained identifiable student faces. While this rate is relatively low, if used to make an inference about the population of photo posts we collected, this would suggest that between 153,218 and 1,053,844 students could potentially be identified via their inclusion in school or school districts' posts.

RQ #2. For students, 187 faces were detected in photos and only 5 of those 187 faces were able to be identified by their names, meaning that 2.67% ([0.87%, 6.13%]) of student faces were identifiable by name. Similar percentages are given below for each of the other categories. These numbers indicate that students and staff had a much smaller percentage of identifiable faces than that of community members. The rest of our results are shown in the table below (Table 1).

Table 1. Identifiability Percentages by Category

Category	Total # of Faces	# of ID Faces	Percentage of Identifiable Faces per Post (RQ #1)	Percentage of Identifiable Faces Per Category Sum (RQ #2)
Student	187	5	5% [1.64%, 11.28%]	2.67% [0.87%, 6.13%]
Staff	87	4	4% [1.10%, 9.92%]	4.6% [1.27%, 11.36%]
Community	25	4	4% [1.10%, 9.92%]	16% [4.54%, 36.08%]

4. DISCUSSION

4.1 Key Findings

Upon the completion of coding our sample and numerical analysis for each category, we are able to make a few important claims about the protection of student privacy. First, students comprised a majority of the faces detected in images; however, compared to the large number of student faces, less than 3% of those faces were able to be identified by their names. While this low proportion may seem to indicate that students' privacy is well-protected, the massive scope of this data (more than nine million public posts by schools or school districts) nevertheless means that many students are at-risk to be identified by both face and name by anyone with internet access if expanded to the entire data set. In short, K-12 institutions' uses of social media could introduce very widespread threats to students' privacy.

How serious are these privacy threats? An identifiable photo presents a relatively low risk compared to, for example, one's address or grade-related information being shared. However, the risk of doing so is not zero: These posts could be used to identify information about individuals, and when accessed, could potentially be used to predict their personal characteristics, even those that require making strong inferences, such as those about individuals' political identities [14]—and, potentially, other identities. Adding to the problem, we note that each of these posts not only associates a name with a photo but also an identifiable photo to a particular location (a school or district) at a specific time. In summation, what seems like a low-risk form of identification, can reveal quite a bit of information on students, leaving their privacy vulnerable.

In addition to the number of photos of students that were able to be identified, the level of protection attached to the privacy of students was intriguing when compared to that of students and staff. More specifically, while students had the highest number of faces detected in images, their isolated level of identifiability was the lowest of all three categories. We can also note that students and staff members together have drastically lower isolated levels of identifiability compared to that of community members: Community members were generally easier to distinguish between than our other categories.

Taken together, these findings speak to concerns about privacy on social media, revealing that not only individuals' actions and posts (e.g. [6, 9, 23]), but also those of educational institutions may pose risks for the privacy of a vulnerable societal group: minors at school. They suggest that the wide use of Facebook and ease of accessing posts coupled with identifiable posts of students may make this particular use of social media a key avenue through which students' privacy is compromised. In this way, these findings add to prior research pointing out that young people may view privacy differently [17]. In addition, this research suggests to the educational data mining community that privacy risks to students may appear in unexpected contexts—and in contexts for which schools may, technically, not be violating the United States Family Educational Rights and Privacy Act (FERPA), but which may deserve greater scrutiny.

4.2 Limitations and Recommendations

This study represents an initial exploration of a topic that has been investigated extensively using other data sources and populations [6, 9]—and which could be investigated much further to better

understand the nature of how students' privacy may be threatened due to the increasingly widespread use of social media by K-12 educational institutions. Due to the small size of our sample (compared to that of the population of photo posts), while we made some inferences from our sample to the population, these were associated with very wide ranges of plausible values: for example, we estimated that the number of identifiable students ranges from between 150,000 and more than one million, a range that makes it difficult to inform other researchers as well as administrators, educators, parents, and students about the scale of the threat to students' privacy. In addition, there are certain statistical inferences that we are unable to make at this time: For instance, with a small number of posts from varying years, we must code a larger sample to be able to model change in privacy risk over time.

It is important to consider the issue of parental consent in the context of student photos via public pages of schools and districts. While our sample data does not include specific information on each educational institution's privacy policies, there has been past research performed regarding actions such as consent forms [3]. Students' parents or legal guardians typically act as their agents of consent, which may appear to legitimize the publicizing of student faces. However, those making these crucial decisions may not have all of the necessary information to make these choices on behalf of their students.

Future research may expand on the findings presented in this study by not only coding a greater number of posts but also coding for different features of them. For instance, we noted that because many images in the latter part of 2020 included students wearing masks, there may surprisingly be a decrease in the number of identifiable faces during the COVID-19 pandemic. Future research that aims to mitigate risks may also note some of the features of posts that protect the privacy of students, and posts by schools or districts that achieve some of the benefits of educational institutions' social media use. How accessible the posts we accessed via both the CrowdTangle [4] platform and other (authorized or unauthorized—e.g., through web-scraping) means is another topic future scholarship can explore in greater depth, as the extent to which others can reproduce our analysis has a bearing on how extensive the threats to students' privacy are. Limiting risks to students' privacy may serve as a model to inform or prompt reflection on the part of the administrators and educators using their school's or school districts' Facebook account. Finally, future research might investigate what key stakeholders—students, parents, and teachers—think of the potential privacy risks around social media use. While past research has reported that teachers are uncomfortable with how social media platforms use student data [16], our results suggest that key individuals in schools may not draw connections between this lack of comfort and how their school or district uses social media, and survey research methods may complement our public data mining approach.

5. REFERENCES

- [1] Baker, R.S., & Inventado, P.S. (2014). Educational data mining and learning analytics. In R.S. Baker & P.S. Inventado (Eds.), *Learning analytics* (pp. 61-75). New York: Springer
- [2] Carpenter, J., Tani, T., Morrison, S., & Keane, J. (2020). Exploring the landscape of educator professional activity on Twitter: an analysis of 16 education-related Twitter hashtags. *Professional Development in Education*, 1–22. <https://doi.org/10.1080/19415257.2020.1752287>.

- [3] Cino, D., & Vandini, C. D. (2020). "Why Does a Teacher Feel the Need to Post My Kid?": Parents and Teachers Constructing Morally Acceptable Boundaries of Children's Social Media Presence. *International journal of communication* [Online], 1153-1172. https://link.gale.com/apps/doc/A632440221/AONE?u=tel_a_utl&sid=AONE&xid=d42623382.
- [4] CrowdTangle Team (2021). *CrowdTangle*. Facebook, Menlo Park, California, United States. List ID: [all-k12-institutions]. Retrieved January 15, 2021.
- [5] Facebook. (n.d.). Retrieved March 12, 2021, from <https://www.facebook.com>
- [6] Fiesler, C., & Proferes, N. (2018). "Participant" perceptions of Twitter research ethics. *Social Media+Society*, 4(1), 2056305118763366.
- [7] Fischer, C., Pardos, Z. A., Baker, R. S., Williams, J. J., Smyth, P., Yu, R., ... & Warschauer, M. (2020). Mining big data in education: Affordances and challenges. *Review of Research in Education*, 44(1), 130-160.
- [8] Hatch, J. A. (2002). *Doing qualitative research in education settings*. Suny Press.
- [9] Ifenthaler, D., & Schumacher, C. (2016). Student perceptions of privacy principles for learning analytics. *Educational Technology Research and Development*, 64(5), 923-938.
- [10] Kimmons, R., Carpenter, J. P., Veletsianos, G., & Krutka, D. G. (2018). Mining social media divides: an analysis of K-12 US School uses of Twitter. *Learning, media and technology*, 43(3), 307-325.
- [11] Kimmons, R., Rosenberg, J.M., & Allman, B. (2021). Trends in educational technology: What Facebook, Twitter, and Scopus can tell us about current research and practice. *TechTrends*, 1-12. <https://link.springer.com/article/10.1007/s11528-021-00589-6>
- [12] Kimmons, R., & Veletsianos, G. (2018). Public internet data mining methods in instructional design, educational technology, and online learning research. *TechTrends*, 62(5), 492-500.
- [13] Kimmons, R., Veletsianos, G., & Woodward, S. (2017). Institutional uses of Twitter in US higher education. *Innovative Higher Education*, 42(2), 97-111.
- [14] Kosinski, M. (2021). Facial recognition technology can expose political orientation from naturalistic facial images. *Scientific Reports*, 11(1), 1-7.
- [15] Madden, M., Lenhart, A., Cortesi, S., Gasser, U., Duggan, M., Smith, A., & Beaton, M. (2013). Teens, social media, and privacy. *Pew Research Center*, 21(1055), 2-86.
- [16] Marin, V. I., Carpenter, J. P., & Tur, G. (2021). Pre-service teachers' perceptions of social media data privacy policies. *British Journal of Educational Technology*, 52(2), 519-535.
- [17] Marwick, A. E., & Boyd, D. (2014). Networked privacy: How teenagers negotiate context in social media. *New Media & Society*, 16(7), 1051-1067.
- [18] Michela, E., Rosenberg, J. M., Sultana, O., Burchfield, M.A., Thomas, T., & Kimmons, R. (2021, April). "Life will eventually get back to normal": School districts' Twitter use in response to COVID-19. Presentation at the American Educational Research Association Annual Meeting, Orlando, FL.
- [19] National Center for Education Statistics. (2021). *Common core of data*. <https://nces.ed.gov/ccd/>
- [20] R Core Team (2021). *R: A language and environment for statistical computing*. <https://www.r-project.org/s>
- [21] Romero-Hall, E., Kimmons, R., & Veletsianos, G. (2018). Social media use by instructional design departments. *Australasian Journal of Educational Technology*, 34(5).
- [22] 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.
- [23] Smith, M., Szongott, C., Henne, B., & Von Voigt, G. (2012, June). Big data privacy issues in public social media. In *2012 6th IEEE international conference on digital ecosystems and technologies (DEST)* (pp. 1-6). IEEE.