Pictures as a Form of Protest: A Survey and Analysis of Images Posted During the Stop Asian Hate Movement on Twitter

1st Oliver Melbourne Allen
Observatory on Social Media
Indiana University
Bloomington, USA
alleno@iu.edu

2nd Emily Chen
Information Sciences Institute
University of Southern California
Los Angeles, USA
echen920@usc.edu

3rd Emilio Ferrara
Information Sciences Institute
University of Southern California
Los Angeles, USA
ferrarae@isi.edu

***Abstract*—Modern protests are not limited to on-the-ground operations, and the ease and speed at which users can upload images to social media platforms has enabled protests to manifest online. Previous analysis of protest imagery from social media sites categorized these images into groups including texts, screenshots, memes, and artwork. However, large-scale manual annotation to identify different types of images is not feasible. By applying machine learning to a large Twitter dataset focused on the Stop Asian Hate movement, we found the type of image an account posted during protests on Twitter is tied to the credibility and political leaning of posted content, type of witnessing (remote or connective), and community formation.**

***Index Terms*—images, protest, social media, social movements, Twitter, visual analysis, machine learning**

I. Introduction

Social media platforms such as Twitter have enabled users to participate in protests by simply uploading images. Photographs have become a form of visual protest through this "connective witnessing" [1], where individuals contribute to the flow of information by sharing eyewitness photographs of an event. Researchers have also observed the phenomenon of "remote witnessing" [2] where accounts appropriate secondary sources, e.g. posting screenshots of news articles.

Previous work focuses on small samples of popular or commonly re-posted images understood through manual annotation [2], [3], but there is little research on how these different types of images exist across an entire movement due to the sheer volume of images involved. In order to bridge this gap, we ask the following questions: 1) how can we apply machine learning techniques to identify and classify types of protest imagery on Twitter and 2) how do accounts leverage images to participate in protest discourse on social media?

II. Method

We extended previous work by developing a method to classify images by type in a large Twitter dataset using the convolutional neural network ResNet50 [4]. We then applied this method to the recent Stop Asian Hate movement to better understand how different types of images contributed to political- and misinformation-related discourse.

We used a dataset focused on the Stop Asian Hate movement, which appeared in response to the rise of anti-asian hate crimes in the wake of the Covid-19 pandemic. We started with a dataset of 4.2 million tweets and selected a subset of accounts who posted at least one domain and at least one image. Images posted by these accounts were classified into one of 7 categories (infographics, photographs with graphics or text, photographs, illustrations with text, illustrations, texts, and screenshots). We also collected the number of right-leaning, left-leaning, low-credibility and high-credibility domains[1] that each account posted and built a retweet network where each account was labelled with the type of image that it retweeted the most.

III. Results

A. Network Structure

In the over 1.3 million images that were analyzed approximately 50% of the images shared were photographs, which supports Jenzen et. al.'s findings [3]. Infographics made up approximately 25% of the corpus, and screenshots and texts each made up about 12% of the corpus.

As seen in Fig. 1, the core of the network (E) is split into two halves: on the left, accounts shared mostly photographs while on the right, accounts shared mostly infographics. Smaller communities formed outside of the cluster E. For example, cluster A is a tightly knit community of accounts that mostly shared screenshots, cluster B shared mostly photographs, while clusters C and D shared primarily infographics. These clusters indicate that accounts who posted protest-related images on social media interacted with accounts who participated in a similar way.

B. Images and Protest Discourse

Overall, we found that the types of images that an account shared were tied to the political lean and credibility of

The authors gratefully acknowledge support from the National Science Foundation Research Experiences for Undergraduates grant #2051101.

[1]Ratings for credibility and political leaning of domains were taken from Media Bias/Fact Check (https://mediabiasfactcheck.com/)