Characterizing Image Sharing Behaviors in US Politically Engaged, Random, and Demographic Audience Segments

Keng-Chi Chang,[1] **Cody Buntain**[2]

[1] Department of Political Science, University of California, San Diego, CA, USA
[2] College of Information Studies, University of Maryland, College Park, MD, USA
kechang@ucsd.edu, cbuntain@umd.edu

Abstract

This work advances understandings of image-sharing behavior on Twitter, across race, gender, age, and political engagement. We infer account-level demographic measures via profile pictures of US Twitter accounts and characterize 20 types of images. Several of these types predict one's demographics using account-level logistic regression models. Around half of the learned clusters (e.g., infographics, natural scenery, sports) are predictive of the user's age, race, or gender, while several other clusters appear to be popular among politically engaged accounts (e.g., images of groups and images of single individuals, which often contain politicians). Our findings suggest it is possible to characterize certain audiences via different types of visual imagery, which has implications for information quality, online engagement, and communications.

Introduction

As visual media–i.e., images and video–become increasingly popular in the online information space, insights into and methods for measuring how such media is used and the audiences most engaged with such media are increasingly important for understanding *and improving* both online and offline behaviors and information spaces. Many studies support these online/offline implications for media in online spaces, as we have good evidence that including visual media in textual posts increases engagement (Li and Xie 2020), mobilizes individuals to protest (Casas and Williams 2019), exposes individuals to anti-social QAnon content (Buntain et al. 2022), and often provides a vector for hate speech (Kiela et al. 2020). Recent advances in image generation and large multi-modal models like GPT4,[1] only amplify these needs, as they reduce cost and effort necessary to create visual media.

Visual media plays a substantial role in contemporary political discourse, and while individuals on Twitter rarely post political content (Bestvater et al. 2022), posting political imagery has the potential to increase exposure to political content. At the same time, ideologically cross-cutting exposure on social media can drive polarization (Bail et al. 2018), and studies suggest that the political right enjoys additional amplification in online spaces (Huszár et al. 2022), especially with respect to visual content (Munger and Phillips 2022). Barberá further explores how better capturing audience demographics is needed to improve our understanding of these online dynamics (2016). As a community, we must understand these dynamics thoroughly to better mitigate such inequities and improve information spaces. It is in this context that this paper is situated, where we contribute to this space by characterizing how different types of images are used by politically engaged Twitter audiences versus a general Twitter audience and how different demographic segments engage with images. To this end, we answer two main research questions:

RQ1: Do politically engaged Twitter accounts share different types of imagery on Twitter than the general US Twitter audience?

RQ2: What types of image sharing behavior are predictive of the account's demographic backgrounds?

To answer these questions, this paper combines two large-scale samples of US Twitter audiences from Alizadeh et al. (2020) with an automated method for demographic inference from profile pictures, called FairFace (Karkkainen and Joo 2021). This analysis uses behavior from 10,000 US Twitter accounts, covering more than 66 million tweets, and 10 million images. Uniquely, we apply FairFace to public profile images from these accounts at scale, inferring age, race, and presented gender–a departure from prior work, that has leveraged surveys (Barberá 2016) or matching Twitter accounts with "voter files" (Hughes et al. 2021; Barberá 2016). Then applying a clustering scheme to these 10 million images to construct a set of image *types*, we assess the predictive power of these image types as they relate to demographics and political engagement.

Results show that demographics exhibit little variation between politically engaged and randomly selected US Twitter audiences. For image types, we likewise see several types of images appear common across ages, gender presentations, and political engagement. That said, using logistic regression models to predict demographics, we see about half of the image types (i.e., around ten of the image clusters) correlate significantly with race, gender, age, and political interest; which clusters correlate with these attributes vary across the attributes, however. These logistic regressions capture only a limited amount of variation in these attributes though,

[1]https://openai.com/research/gpt-4

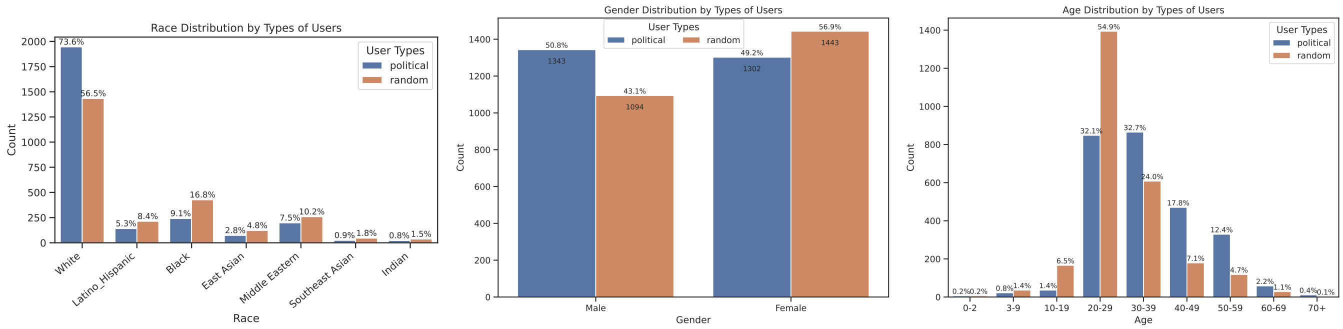

Figure 1: Distributions of Predicted Account-Level Demographics–Race, Gender, and Age–Across Random and Politically Engaged Audiences. Demographic attributes come from the FairFace library (Karkkainen and Joo 2021).

Sample	Random	Political
# users	5,000	5,000
# tweets & retweets	31,038,705	35,932,231
% tweets	40%	37%
% retweets	60%	63%
% has image	18%	14%

Table 1: Summary statistics of Twitter sample

ranging from McFadden pseudo $R^2 = 0.04$ for race and $R^2 \in [0.10, 0.12]$ for gender, age, and political engagement.

Data and Methodology

Our sample consists of two sets of Twitter users, gathered from Alizadeh et al. (2020): one based on timelines from a random set of 5,000 accounts geolocated to the United States (“random”), and another based on timelines from a collection of 5,000 US accounts who are politically engaged (“political”). Politically engaged accounts are defined as those accounts following at least 5 political Twitter accounts –that is, they follow Twitter accounts of US congresspeople in the Senate, House, Governors, and in the executive branch. These samples are also restricted to users who posted at least 100 times in 2015–2017. Table 1 summarizes statistics for these two audience samples.

Demographic Inference from Profile Pictures

From these audience samples, we assess differences in how accounts of various demographic attributes present themselves via the images they share. To that end, beyond dividing between politically engaged versus random accounts, we also investigate account-level demographics using FairFace (Karkkainen and Joo 2021) applied to an account’s profile picture. FairFace uses a face attribute model that is balanced on race, gender, and age to mitigate the potential bias toward Caucasian faces. Figure 1 shows distributions of the predicted demographics. Compared with the Twitter samples linked to voter files (Hughes et al. 2021), the marginal distributions of race and gender are similar (e.g., 70% White, 11% Black; 52% Female), while predicted age from profile pictures appears biased toward younger age.

Collecting and Characterizing Images in Tweets

After assessing demographics, we then collect images in the timeline of these account samples, so that we can assess differences in visual presentation. We focus on the original tweets–i.e., we exclude retweets–on the user timelines containing images. To featurize images, we use a pre-trained ResNet50 deep learning model (He et al. 2016) to generate 2048-dimensional embedding for each image.

To characterize disparate types of visuals, as well as the general content of these image types, we apply k-means clustering on the image embeddings to group these images into clusters. Using cluster-quality metrics of within-cluster sum of squared distances and silhouette scores to determine the number of clusters (see Figure 2), we set $k = 20$. See Figure 5 for random images for some clusters.

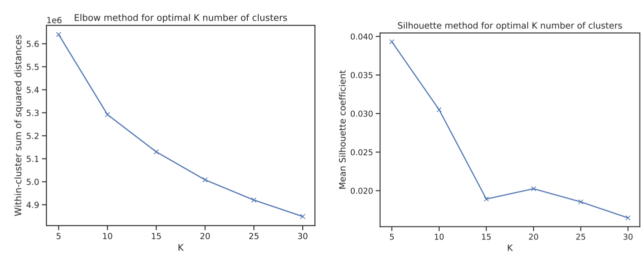

(a) Within-Cluster Distances (b) Silhouette Score

Figure 2: Cluster Quality Metrics. Cluster counts between 15-20 seem reasonable based on elbows in these curves.

Results

To address **RQ1**, on whether politically engaged accounts share different types of images than general US Twitter audiences, we first present the distribution of clusters by different types of audiences–politically engaged versus random accounts–as shown in Figure 3. Results show overlap between image-types shared by political and random users: Most clusters are not discriminant in separating political and random accounts. That said, a few clusters appear over-represented among politically engaged accounts, specifically clusters 3, 4, and 7, whereas clusters 0 and 16 are over-represented among general audiences.

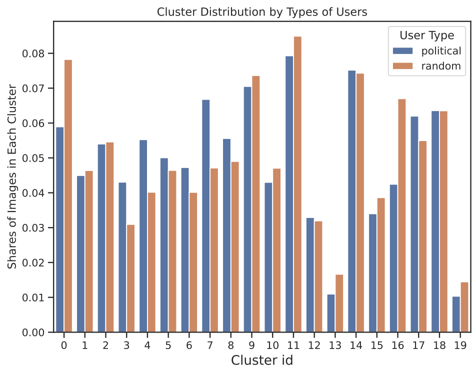

Figure 3: Cluster Distribution by Audience. Most clusters have similar distributions across politically engaged and random accounts, but a few clusters appear over-represented in one of the other sets (e.g., clusters 0, 3, 4, 7, and 16).

To understand predictive power of image clusters for demographics (**RQ2**), we run an account-level logistic regression of demographic variables (race, gender, age, political engagement) using an account's percent of images in each cluster; that is, for user i with demographics y_i, we estimate:

$$\sum_{k=1}^{20} \beta_k \frac{\text{\# images in cluster } k \text{ shared by user } i}{\text{\# images shared by user } i} + \varepsilon_i$$

In this equation, β_k show the correlation between sharing cluster k and demographic variables. To binarize demographic attributes returned by FairFace, we collapse race to white or non-white, gender to female or non-female, age to less than 30 years old or older, and politically engaged or not. Figure 4 shows the results. Our findings are in general interpretable, providing insight into how increases in sharing a particular type of image change the probability .

Each logistic regression model has an associated McFadden Pseudo-R^2, in the range $[0.04 - 0.12]$, with race appearing the most difficult to predict given an account's distribution of images over image clusters. In contrast, gender appears to perform best with a pseudo-$R^2 = 0.12$, while identifying age and political engagement have the same $R^2 = 0.1$. Though these values are low, guidance on interpreting McFadden's Pseudo-R^2 state it is never 1 (Hu, Shao, and Palta 2006) and suggest values in the range $[0.2 - 0.4]$ "represent an excellent fit" (McFadden 2021). As such, we interpret at least the gender, age, and politically engaged models to have some predictive power.

Discussion and Conclusions

Our finding that the distributions of clusters over politically engaged versus general audiences is relatively stable, as shown in Figure 3, is interesting in that it suggests, in general, information sharing behaviors on Twitter are not massively driven by political interest. This result may be a reflection of our choice to exclude retweeted images, as retweets are a significant indicator of political affiliation (Conover

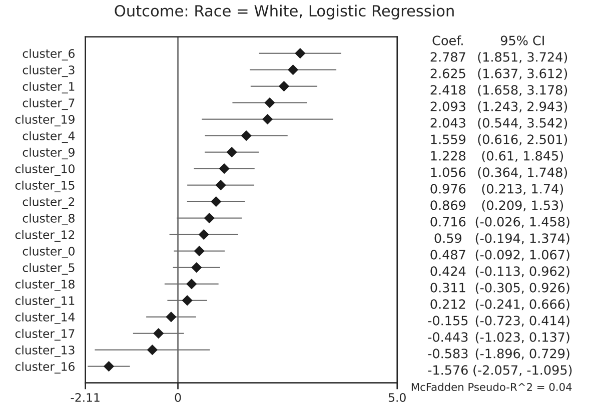

(a) Race = White (Pseudo $R^2 = 0.04$)

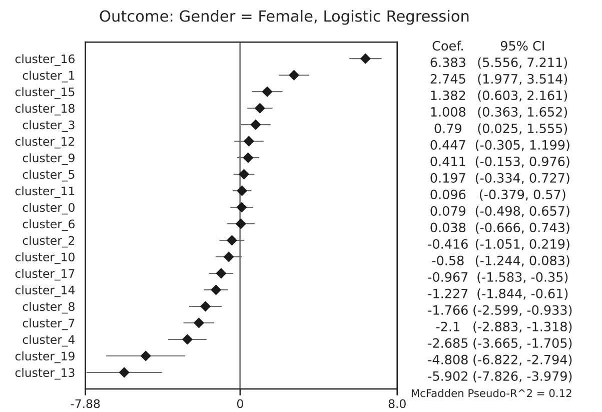

(b) Gender = Female (Pseudo $R^2 = 0.12$)

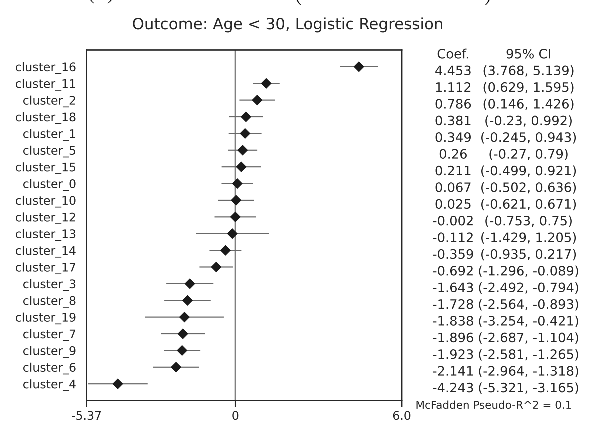

(c) Age < 30 (Pseudo $R^2 = 0.1$)

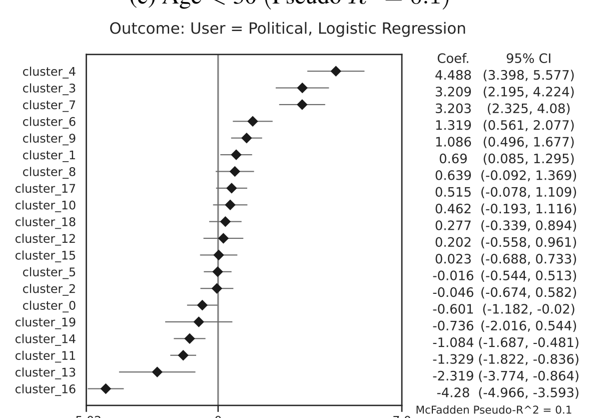

(d) Politically Engaged (Pseudo $R^2 = 0.1$)

Figure 4: Regression Coefficients of Account-Level Cluster Distributions on Demographics. Cluster distributions appear to have some predictive power for gender, age, and political engagement, whereas white-versus-non-white appears difficult to capture with this data.

et al. 2012), coupled with the rarity of posting political content organically (Bestvater et al. 2022). Additional research is needed to assess whether these results on image-sharing behaviors are consistent at the account level–that is, whether politically engaged and general audiences post similar distributions of political imagery as they do political text.

If we inspect the clusters in Figure 5 that are over-represented in politically engaged audiences, these clusters of images do appear particularly politically relevant. For example, in cluster 3, we see many images of groups of people; this finding is consistent with Joshi and Buntain (Joshi and Buntain 2022), where politicians share images of constituents. Likewise, cluster 4 appears comprised of images containing one or two faces, mainly politicians (e.g., Donald Trump, Justin Trudeau, etc.), with clear political relevance. Cluster 7 is less clearly political relevant though, suggesting more research is needed to assess how these images are being used in a potentially political context.

Regarding other demographic attributes, sharing infographics (cluster 9), faces of politicians (cluster 4), natural scenery (cluster 6), or street views (cluster 7) is predictive of users being older than 30; sharing images of natural scenery (cluster 6) or social gatherings of White people (cluster 3) is highly predictive of users being White; sharing sports (clusters 13, 19) is highly predictive of users being male.

Decades of social science research suggest sociodemographic traits are major drivers for behaviors online and offline. Our study illustrates a way to proxy such information from profile pictures and infers how image-sharing behavior varies with demographic segments. Using visual features extracted from deep learning, our initial finding suggests that around half of the image clusters contain predictive information about the account's race, gender, age, and political engagement. The proposed method is interpretable and scalable, allowing for more images, more fine-tuned feature extractors, and more fine-grained demographic variables.

Our study has clear implications for studies of digital literacy and misinformation. If certain users tend to share certain types of images, information actors can utilize this information to design visuals if they would like to target some particular populations. It is also possible that this "content-based targeting" is harder to achieve in text than images.

Limitations include the fact that profile pictures are not always representing the users themselves. However, manual inspections and the customs on Twitter as a platform (e.g., unlike Reddit, where most users do not use profile pictures) convinced us that it is a reasonable measure. Second, we have only studied a narrow aspect of online image-sharing behavior–tweeting. Retweeting or reacting to visual-based content is also of interest to the study of information space.

References

Alizadeh, M.; Shapiro, J. N.; Buntain, C.; and Tucker, J. A. 2020. Content-Based Features Predict Social Media Influence Operations. *Science Advances*, 6(30): eabb5824.

Bail, C. A.; Argyle, L. P.; Brown, T. W.; Bumpus, J. P.; Chen, H.; Hunzaker, M. B. F.; Lee, J.; Mann, M.; Merhout, F.; and Volfovsky, A. 2018. Exposure to opposing views on social media can increase political polarization. *Proceedings of the National Academy of Sciences*, 115(37): 9216–9221.

Barberá, P. 2016. Less is more? How demographic sample weights can improve public opinion estimates based on Twitter data.

Bestvater, S.; Shah, S.; Rivero, G.; and Smth, A. 2022. Politics On Twitter: One-third Of Tweets From U.s. Adults Are Political. Technical report, Pew Research Center.

Buntain, C.; Barlow, M. D.; Bloom, M.; and Johns, M. A. 2022. Paved with Bad Intentions: QAnon's Save the Children Campaign. *Journal of Online Trust and Safety*, 1(2).

Casas, A.; and Williams, N. W. 2019. Images that Matter: Online Protests and the Mobilizing Role of Pictures. *Political Research Quarterly*, 72(2): 360–375.

Conover, M. D.; Gonçalves, B.; Flammini, A.; and Menczer, F. 2012. Partisan asymmetries in online political activity. *EPJ Data Science*, 1(1): 6.

He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.

Hu, B.; Shao, J.; and Palta, M. 2006. Pseudo-R 2 in logistic regression model. *Statistica Sinica*, 847–860.

Hughes, A. G.; McCabe, S. D.; Hobbs, W. R.; Remy, E.; Shah, S.; and Lazer, D. M. J. 2021. Using Administrative Records and Survey Data to Construct Samples of Tweeters and Tweets. *Public Opinion Quarterly*, 85(S1): 323–346.

Huszár, F.; Ktena, S. I.; O'Brien, C.; Belli, L.; Schlaikjer, A.; and Hardt, M. 2022. Algorithmic amplification of politics on Twitter. *Proceedings of the National Academy of Sciences*, 119(1): e2025334119.

Joshi, A.; and Buntain, C. 2022. Examining Similar and Ideologically Correlated Imagery in Online Political Communication. arxiv:arXiv:2110.01183.

Karkkainen, K.; and Joo, J. 2021. FairFace: Face attribute dataset for balanced race, gender, and age for bias measurement and mitigation. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, 1548–1558.

Kiela, D.; Firooz, H.; Mohan, A.; Goswami, V.; Singh, A.; Ringshia, P.; and Testuggine, D. 2020. The Hateful Memes Challenge: Detecting Hate Speech in Multimodal Memes. In Larochelle, H.; Ranzato, M.; Hadsell, R.; Balcan, M.; and Lin, H., eds., *Advances in Neural Information Processing Systems*, volume 33, 2611–2624. Curran Associates, Inc.

Li, Y.; and Xie, Y. 2020. Is a Picture Worth a Thousand Words? An Empirical Study of Image Content and Social Media Engagement. *Journal of Marketing Research*, 57(1): 1–19.

McFadden, D. 2021. Quantitative methods for analysing travel behaviour of individuals: some recent developments. In *Behavioural travel modelling*, 279–318. Routledge.

Munger, K.; and Phillips, J. 2022. Right-Wing YouTube: A Supply and Demand Perspective. *The International Journal of Press/Politics*, 27(1): 186–219.

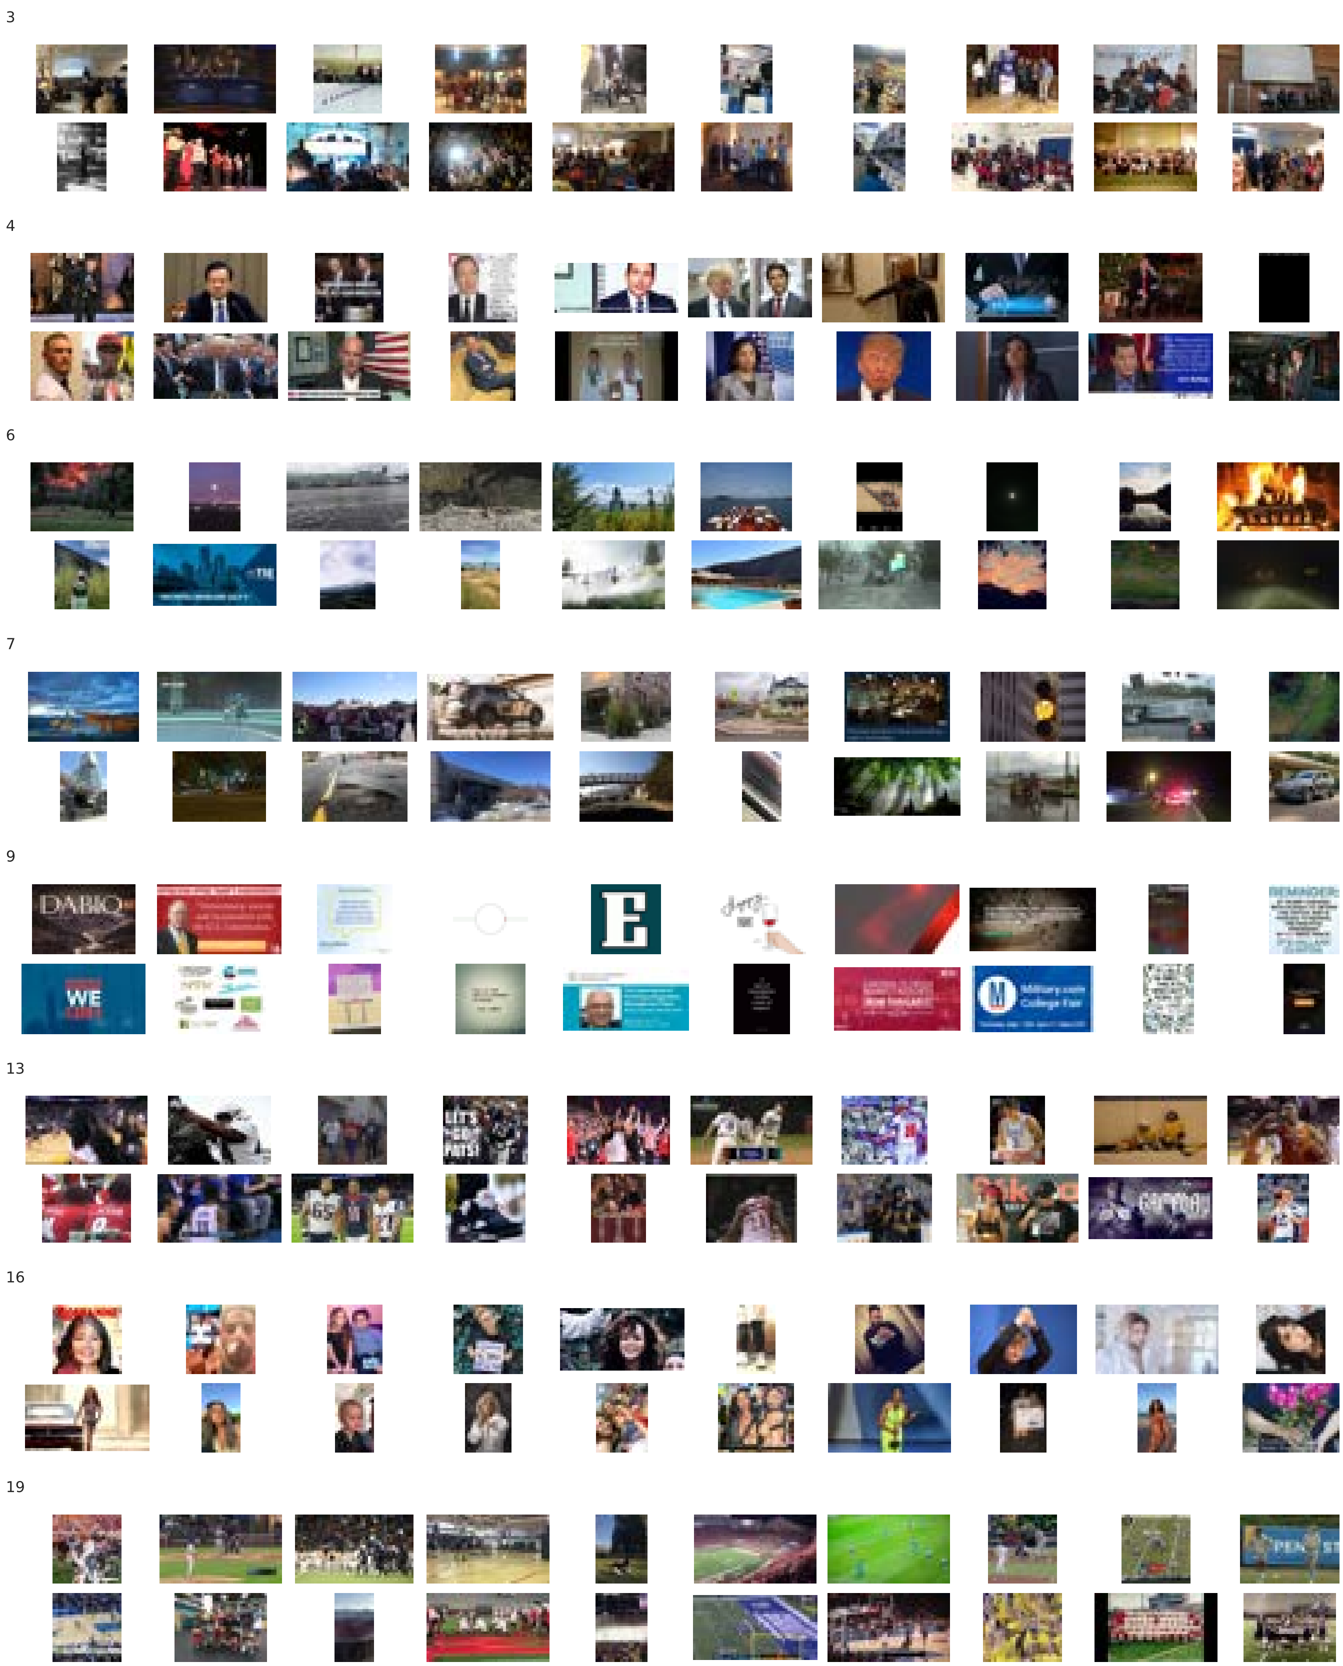

Figure 5: Random images from predictive clusters (cluster ids on top left of each panel)