

10k Twitter Profiles Use AI-Generated Faces Daily!

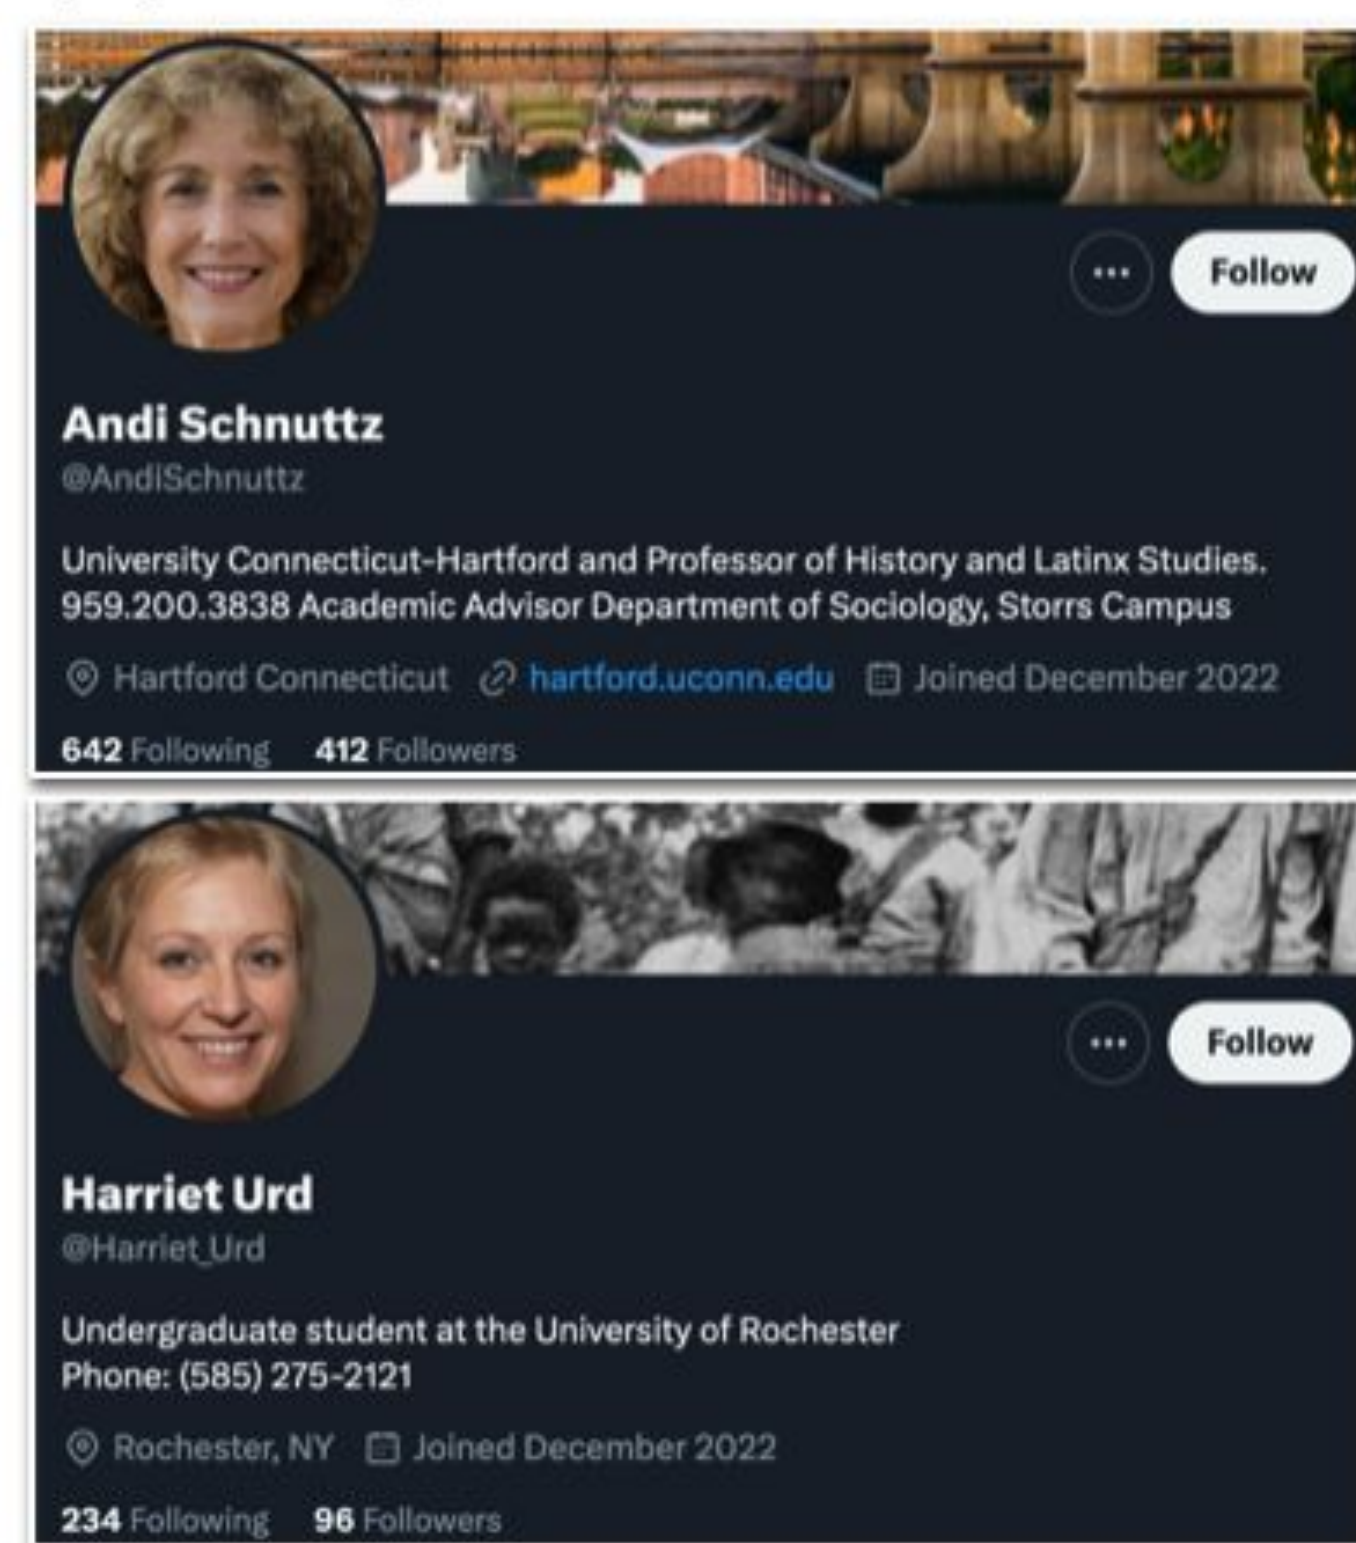
Characteristics and prevalence of fake social media profiles with AI-generated faces

Kai-Cheng Yang, Danishjeet Singh, and Filippo Menczer

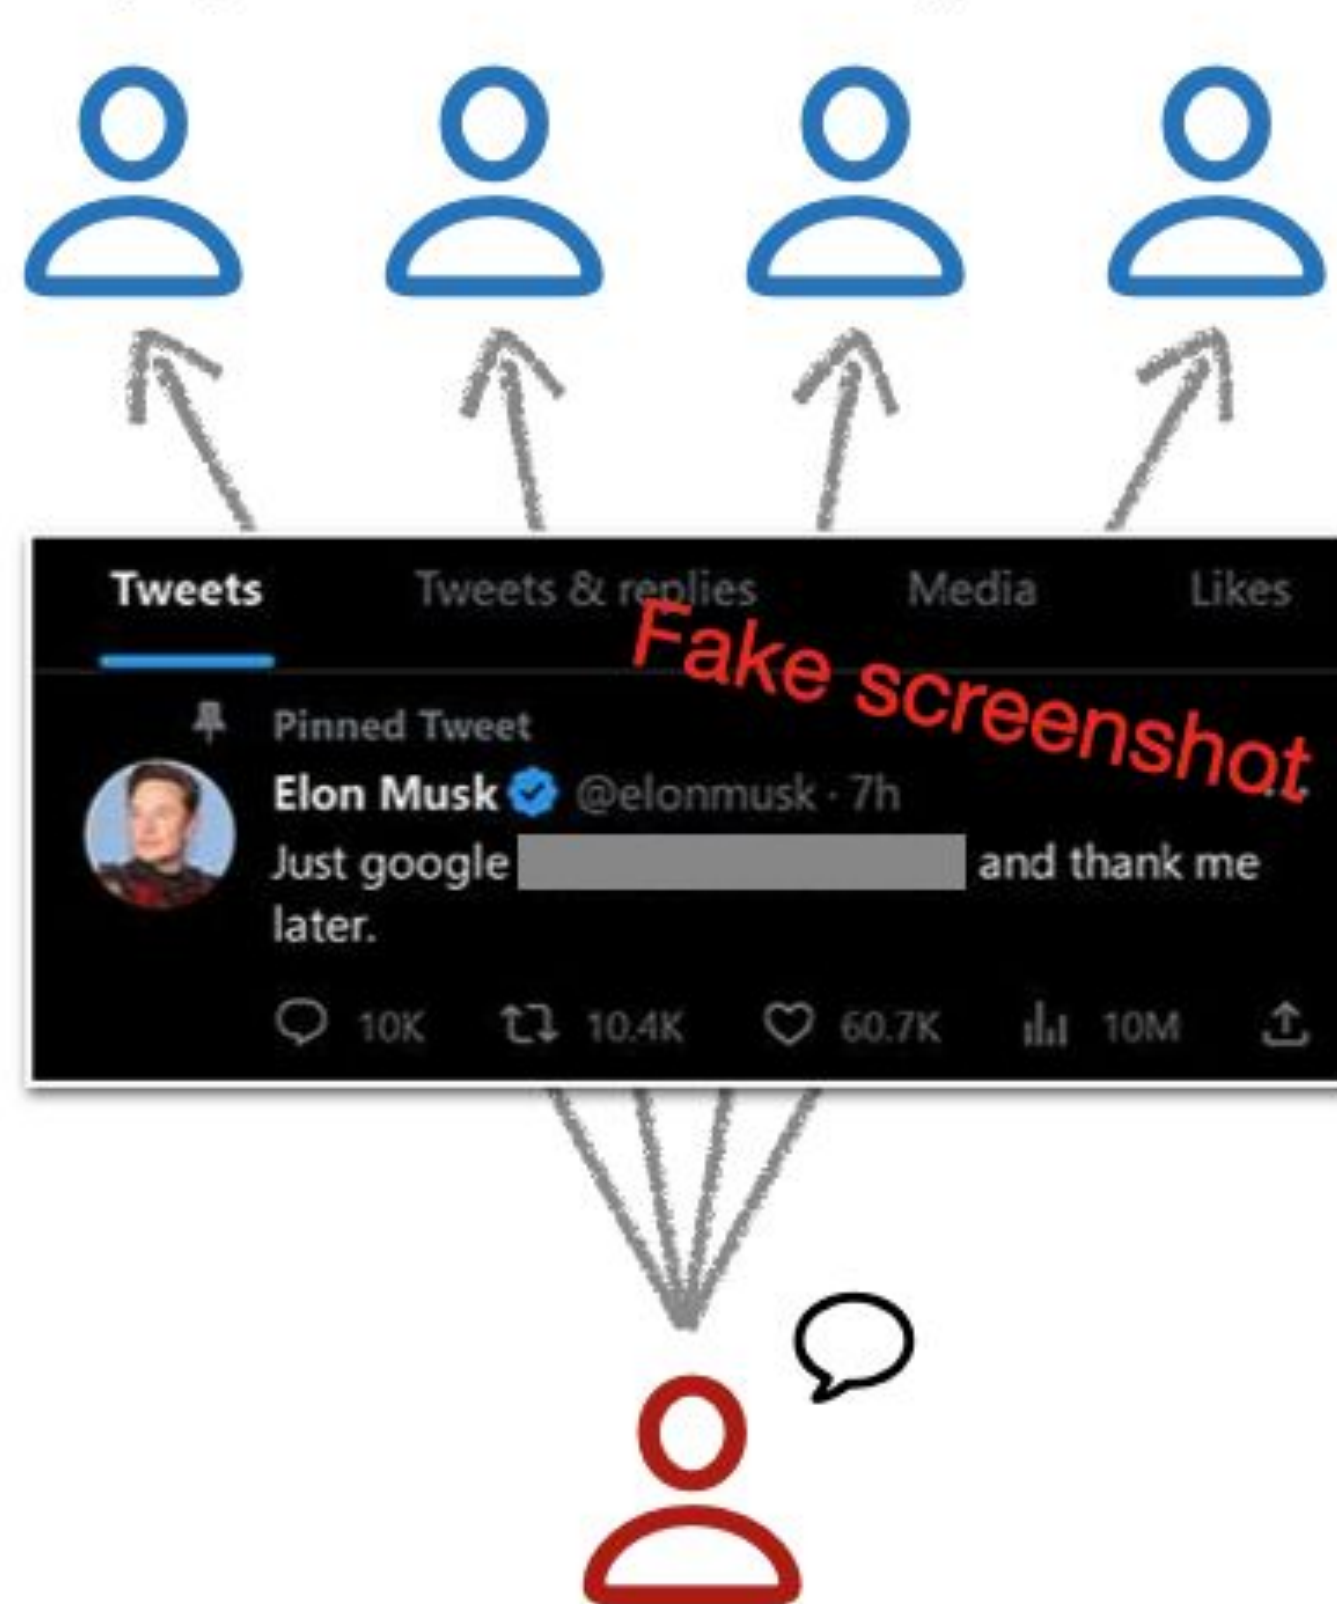


We found 1400+ Twitter accounts using human faces generated by Generative Adversarial Networks (GANs) as their profile pictures. They demonstrate various behaviors and tactics.

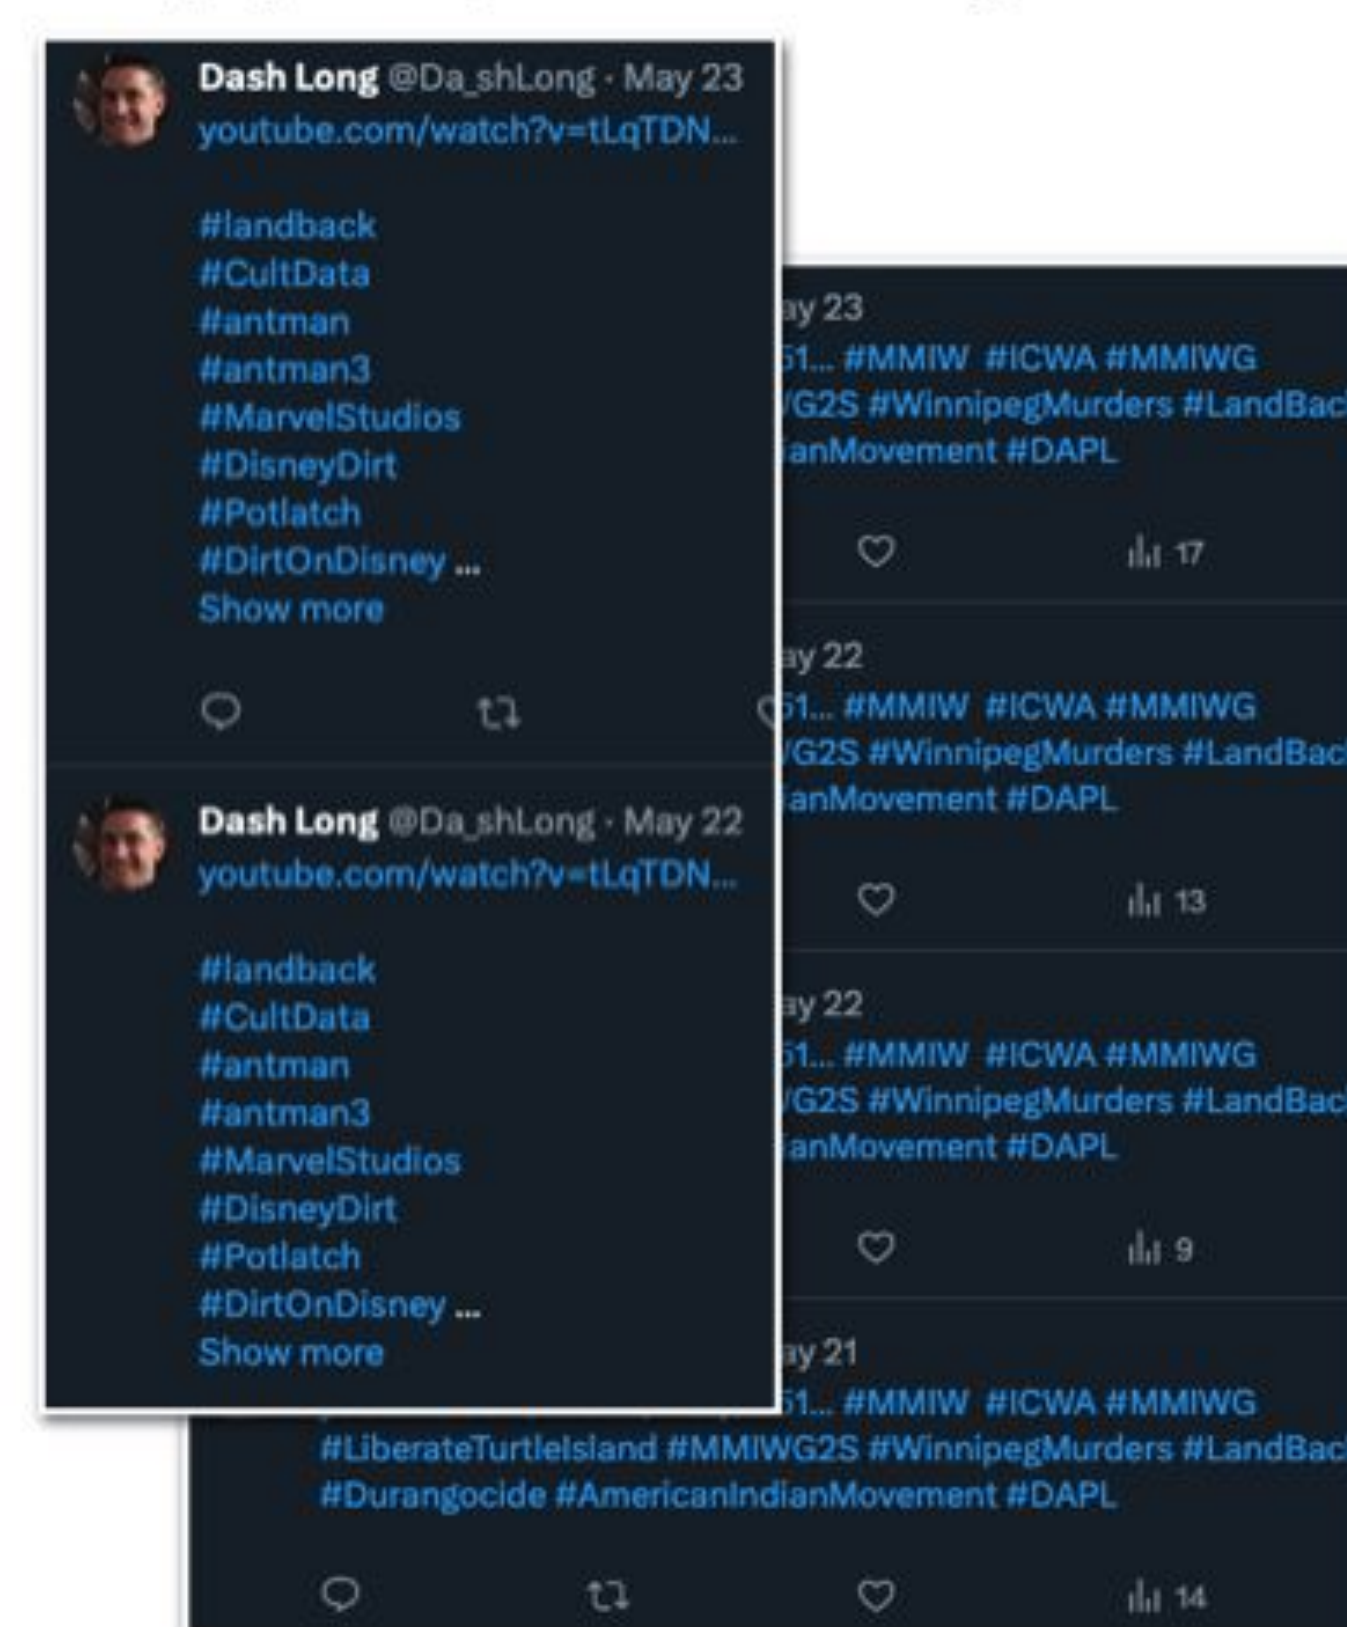
(a) Impersonation



(b) Scamming



(c) Spamming



(d) Amplification



(e) ChatGPT for content generation



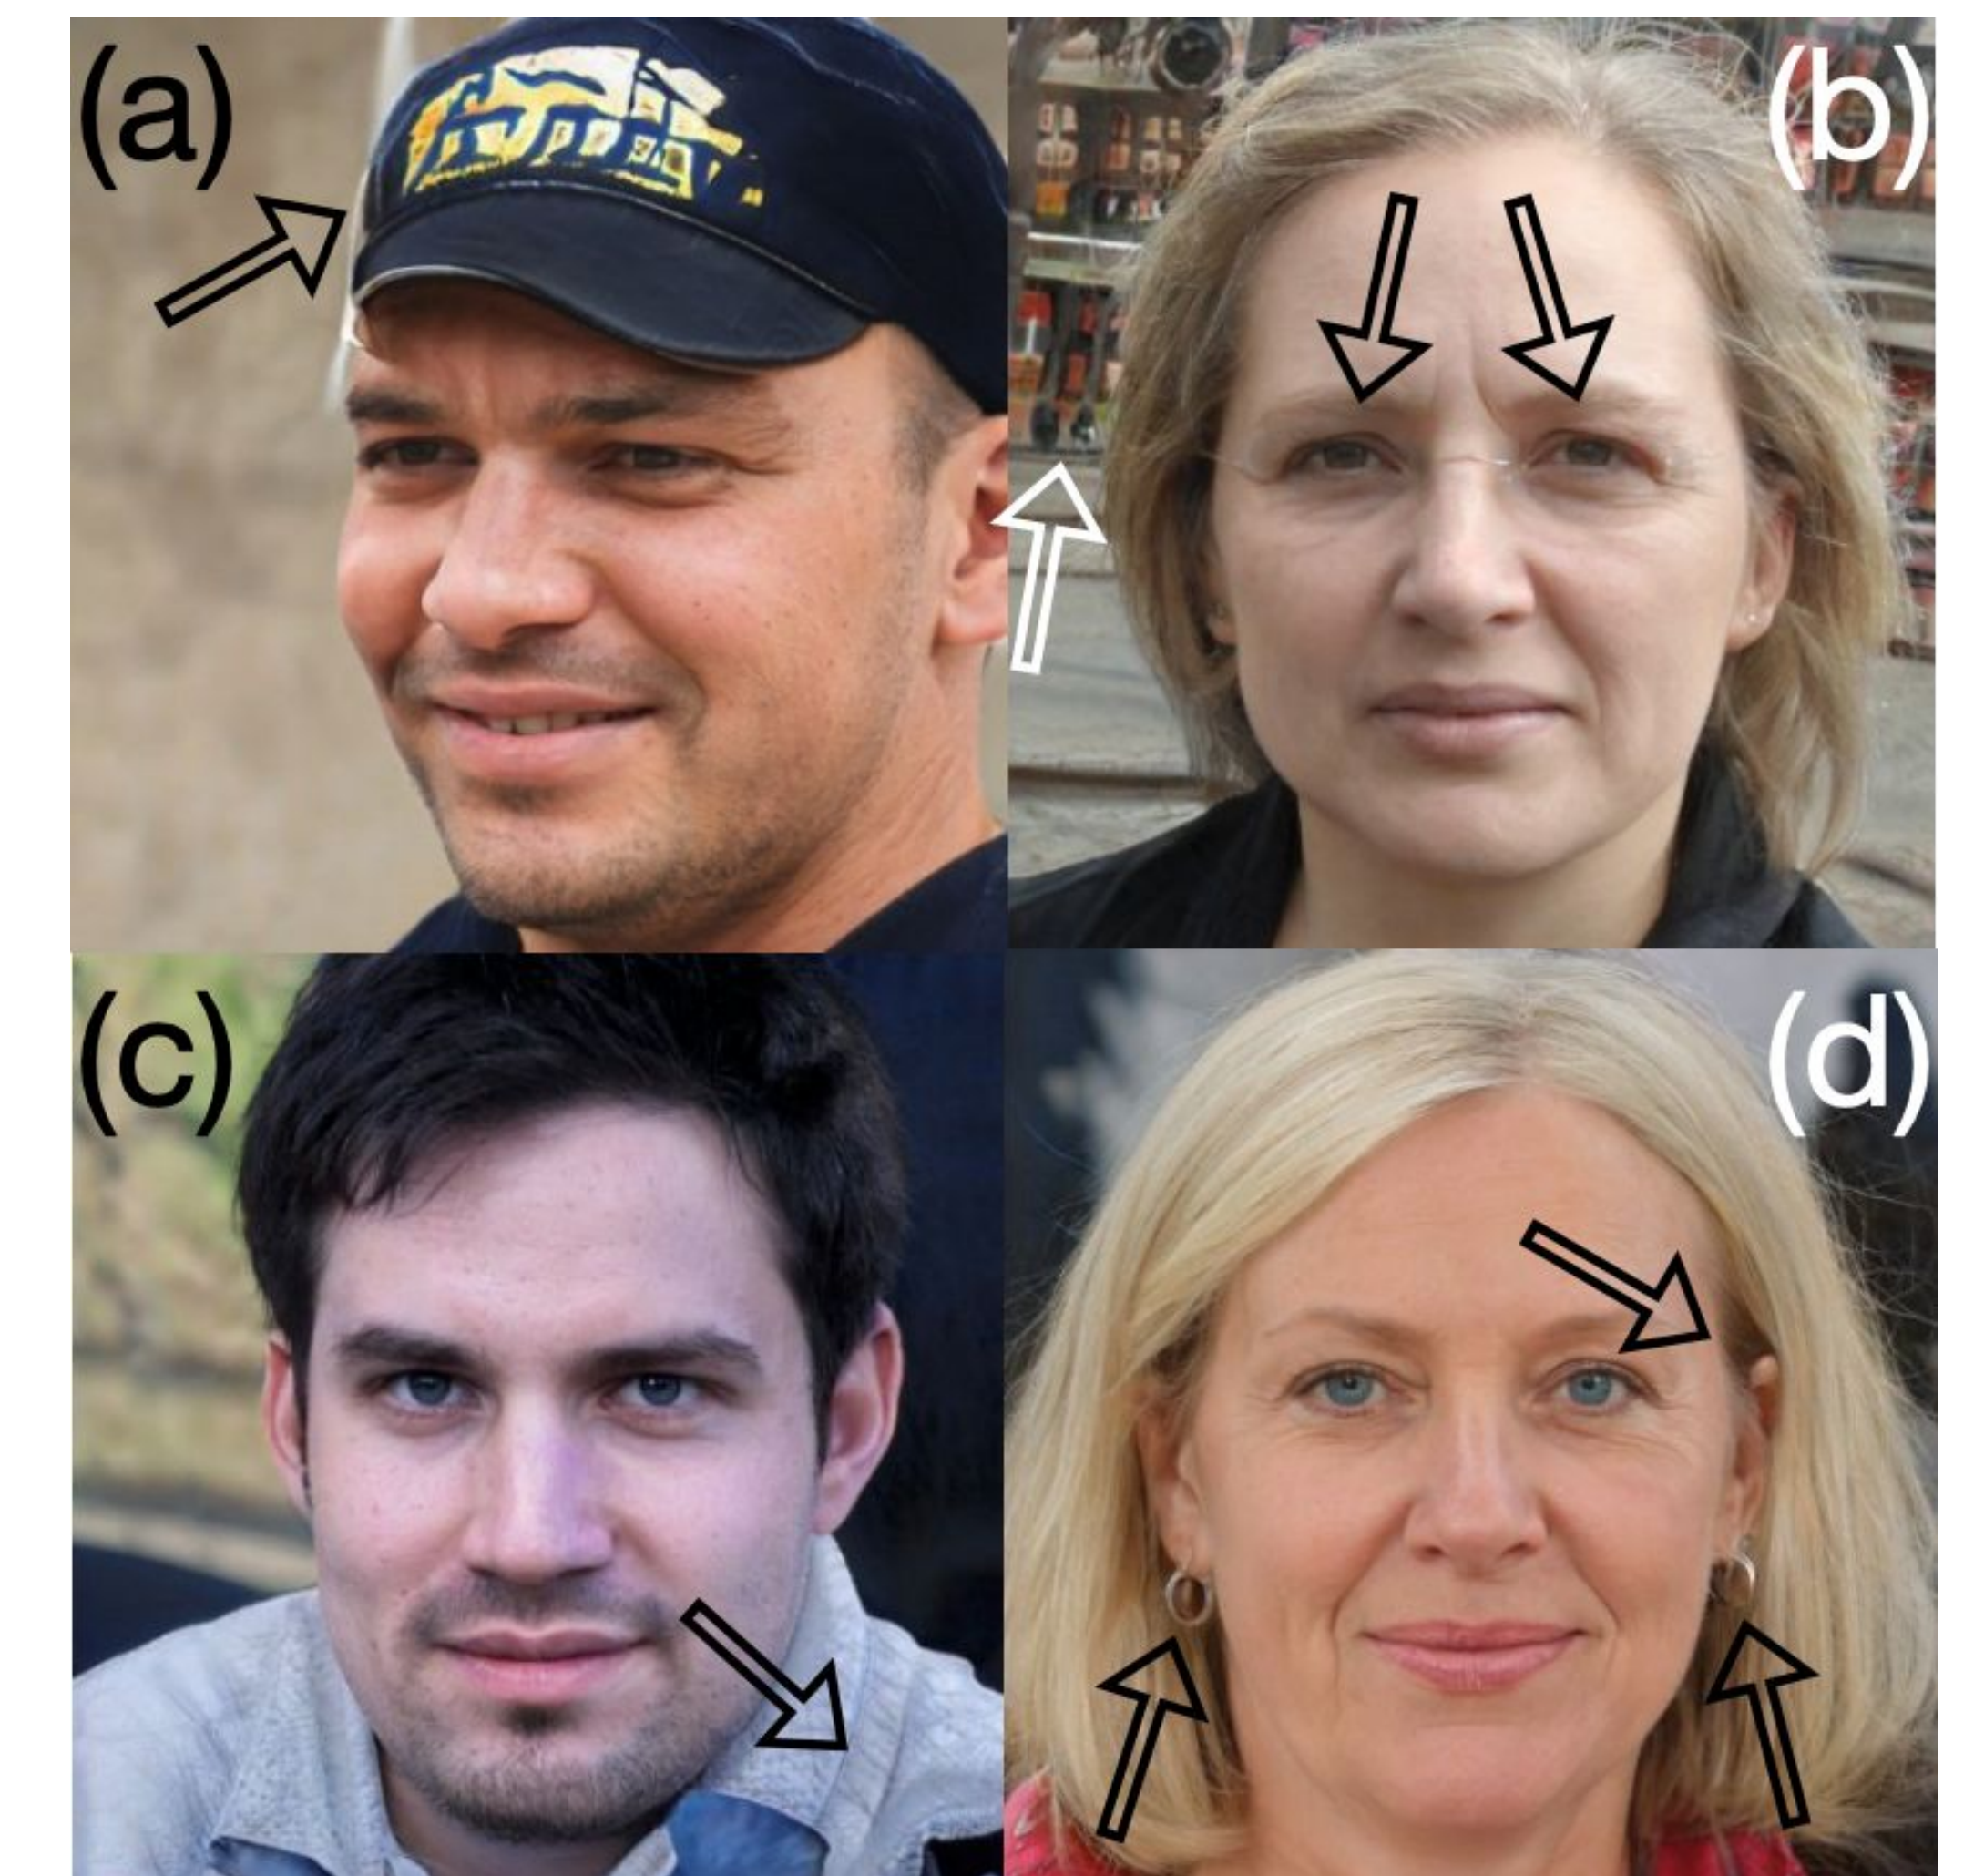
How to spot a GAN generated face?

EYES at same locations



(a) Unreal hat

(b) Glasses frame blended



(c) Unreal collar

(d) Earrings blended

We built a detection algorithm and estimated their prevalence on Twitter

0.021% - 0.044% Twitter users have a GAN generated profile picture
i.e **8,537 to 17,846** daily active accounts