VIEWPOINT

Why concern about facial recognition from brain scans is overblown

BY **KEVIN PELPHREY**

18 FEBRUARY 2020

How secure are brains scans of research participants? Can they, in the wrong hands, compromise the participants' identity?

Perhaps. But I firmly believe that the risk is minimal, and there are things scientists can do — such as masking the participants' faces — to safeguard this information.

The privacy concerns arose when a study published in October reported that a facial-recognition software program had correctly matched photos of research volunteers with magnetic resonance imaging (MRI) scans of their head and brain[1]. This feat was a technical tour de force, representing important advances in biomedical image analysis.

However, most of the media reports surrounding the paper struck an alarmist tone. For example, an article in The New York Times warned that MRI scans could now be used by to **mine medical records** and sell targeted ads.

Now that the dust has settled, I would like to revisit this study's findings and the concerns they fostered.

Despite the care the researchers took to temper and constrain their conclusions, media coverage of their work raised two potentially disturbing implications: First, that given the right software and access to the internet's bounty of selfies, anyone might use a person's brain scan to determine that person's identity. Second, these stories suggested there would be no easy way to overcome this privacy risk without jeopardizing scientific progress.

Both conclusions are overblown.

The validity of the first inference made in popular reports rests upon several prerequisites. To begin with, one would need the brain scan — but that is not so easily obtained. Researchers can download brain scans of research participants who provided informed consent for data sharing from various databases, but they must first promise that they will not attempt to identify the participants and swear to protect the downloaded images on secure servers.

Thwarting 'bad actors':

In the autism community, for instance, the two most widely used databases are the **NIMH Data Archive** and the **Autism Brain Imaging Data Exchange**, whose acronym, ABIDE, emphasizes the fact that researchers must abide by an agreement to safeguard the identity of their research participants.

Any scientist who uses an MRI scan from either of these (or any other) databases to identify a participant would be violating ethics rules, institutional review board policy and U.S. federal law — specifically, the Health Insurance Portability and Accountability Act's (HIPAA) **privacy rule**. In essence, this would be an offense punishable by disciplinary action at work, personal and institutional civil or financial liability and perhaps even criminal prosecution.

Assuming there are always ‘bad actors’ willing to risk these severe punishments, how easy is it to obtain the other ingredients for this brainy form of cyber-stalking?

Admittedly, it is not difficult to obtain photos of people’s faces, given the ubiquity of selfies posted on the web and the power of image search engines. But it is far more difficult to access and employ the necessary reconstruction and image-matching algorithms.

There are key differences between the algorithmic and computational demands needed to identify the participants in the original study and the more general challenge of identifying individuals within the general public, which was played up in media reports.

Orwellian possibilities:

The researchers addressed a constrained problem, accurately matching an MRI reconstruction of someone’s brain to a panel of a few dozen photos, well-matched in terms of size, angle, luminance and intensity. Comparatively speaking, they were shooting fish in a barrel. Identifying members of the public would call for taking many MRI reconstructions from a group of people and finding the best match among millions of available photographs.

The researchers discussed a number of ways in which bad actors could narrow their search. For instance, they listed several methods by which these bad actors might determine whether an individual exists within a particular dataset. This is where some truly Orwellian possibilities arise.

For example, in the **Supplementary Methods** portion of the study, the researchers noted that "access to de-identified medical images could also be obtained via re-purposed clinical data shared with data brokers and analytics enterprises, or by having or gaining access to computer systems being used to store, transfer, or analyze de-identified data (e.g. cloud- based storage/analytics)." Because these sources are not covered under HIPAA, they are not beholden to the privacy rule requiring that researchers remove identifying details, such as "full face photographic images and other comparable images."

Numerous media outlets incorrectly reported that removing or masking an individual’s face renders it difficult or impossible to use that person’s data for research. But this is incorrect. It is not difficult to remove faces from MRI scans stored in databases. In fact, there are a number of readily available software tools that do this accurately and automatically. Several are bundled with the most common software packages used to process MRI data.

Using these tools to remove faces from MRI databases would not negatively affect our ability to conduct science. It would be a relatively easy and highly effective fix for any privacy concerns. And it is a fix that many research labs already implement.

Kevin Pelphrey is Harrison-Wood Professor of Neurology at the University of Virginia in

Charlottesville.

REFERENCES:

1. Schwarz C.G. *et al. N. Engl. J. Med.* **381**, 1684-1686 (2019) **PubMed**