3.2 billion images and 720,000 hours of video are shared online daily. Sorting real from fake is an increasing challenge.

T.J. Thomson, *Queensland University of Technology*; Daniel Angus, *Queensland University of Technology*, and Paula Dootson, *Queensland University of Technology*

Twitter screenshots/Unsplash, Author provided

Twitter over the weekend "tagged" as manipulated a video showing US Democratic presidential candidate Joe Biden supposedly forgetting which state he's in while addressing a crowd.

Biden's "hello Minnesota" greeting contrasted with prominent signage reading "Tampa, Florida" and "Text FL to 30330".

The Associated Press's fact check confirmed the signs were added digitally and the original footage was indeed from a Minnesota rally. But by the time the misleading video was removed it already had more than one million views, The Guardian reports.

If you use social media, the chances are you see (and forward) some of the more than 3.2 billion images and 720,000 hours of video shared daily. When faced with such a glut of content, how can we know what's real and what's not?

While one part of the solution is an increased use of content verification tools, it's equally important we all boost our digital media literacy. Ultimately, one of the best lines of defence — and the only one you can control — is you.

Seeing shouldn't always be believing

Misinformation (when you accidentally share false content) and disinformation (when you intentionally share it) in any medium can erode trust in civil institutions such as news organisations, coalitions and social movements. However, fake photos and videos are often the most potent.

For those with a vested political interest, creating, sharing and/or editing false images can distract,

confuse and manipulate viewers to sow discord and uncertainty (especially in already polarised environments). Posters and platforms can also make money from the sharing of fake, sensationalist content.

Only 11-25% of journalists globally use social media content verification tools, according to the International Centre for Journalists.

Could you spot a doctored image?

Consider this photo of Martin Luther King Jr.

This altered image clones part of the background over King Jr's finger, so it looks like he's flipping off the camera. It has been shared as genuine on Twitter, Reddit and white supremacist websites.

In the original 1964 photo, King flashed the "V for victory" sign after learning the US Senate had passed the civil rights bill.

Beyond adding or removing elements, there's a whole category of photo manipulation in which images are fused together.

Earlier this year, a photo of an armed man was photoshopped by Fox News, which overlaid the man onto other scenes without disclosing the edits, the Seattle Times reported.

Similarly, the image below was shared thousands of times on social media in January, during Australia's Black Summer bushfires. The AFP's fact check confirmed it is not authentic and is actually a combination of several separate photos.
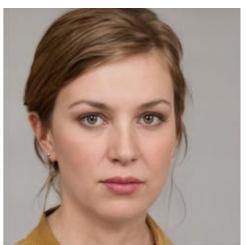
Fully and partially synthetic content

Online, you'll also find sophisticated "deepfake" videos showing (usually famous) people saying or doing things they never did. Less advanced versions can be created using apps such as Zao and Reface.

A team from the Massachusetts Institute of Technology created this fake video showing US President Richard Nixon reading lines from a speech crafted in case the 1969 moon landing failed. (Youtube)

Or, if you don't want to use your photo for a profile picture, you can default to one of several websites offering hundreds of thousands of AI-generated, photorealistic images of people.

These people don't exist, they're just images generated by artificial intelligence.

Editing pixel values and the (not so) simple crop

Cropping can greatly alter the context of a photo, too.

We saw this in 2017, when a US government employee edited official pictures of Donald Trump's inauguration to make the crowd appear bigger, according to The Guardian. The staffer cropped out the empty space "where the crowd ended" for a set of pictures for Trump.

Views of the crowds at the inaugurations of former US President Barack Obama in 2009 (left) and President Donald Trump in 2017 (right).
AP

But what about edits that only alter pixel values such as colour, saturation or contrast?

One historical example illustrates the consequences of this. In 1994, Time magazine's cover of OJ Simpson considerably "darkened" Simpson in his police mugshot. This added fuel to a case already plagued by racial tension, to which the magazine responded:

> No racial implication was intended, by Time or by the artist.

Tools for debunking digital fakery

For those of us who don't want to be duped by visual mis/disinformation, there are tools available — although each comes with its own limitations (something we discuss in our recent paper).

Invisible digital watermarking has been proposed as a solution. However, it isn't widespread and requires buy-in from both content publishers and distributors.

Reverse image search (such as Google's) is often free and can be helpful for identifying earlier, potentially more authentic copies of images online. That said, it's not foolproof because it:

- relies on unedited copies of the media already being online
- doesn't search the *entire* web
- doesn't always allow filtering by publication time. Some reverse image search services such as TinEye support this function, but Google's doesn't.
- returns only exact matches or near-matches, so it's not thorough. For instance, editing an image and then flipping its orientation can fool Google into thinking it's an entirely different one.Most reliable tools are sophisticated

Meanwhile, manual forensic detection methods for visual mis/disinformation focus mostly on edits visible to the naked eye, or rely on examining features that aren't included in every image (such as shadows). They're also time-consuming, expensive and need specialised expertise.

Still, you can access work in this field by visiting sites such as Snopes.com — which has a growing repository of "fauxtography".

Computer vision and machine learning also offer relatively advanced detection capabilities for images and videos. But they too require technical expertise to operate and understand.

Moreover, improving them involves using large volumes of "training data", but the image repositories used for this usually don't contain the real-world images seen in the news.

If you use an image verification tool such as the REVEAL project's image verification assistant, you might need an expert to help interpret the results.

The good news, however, is that before turning to any of the above tools, there are some simple questions you can ask yourself to potentially figure out whether a photo or video on social media is fake. Think:

- was it originally made for social media?
- how widely and for how long was it circulated?
- what responses did it receive?
- who were the intended audiences?

Quite often, the logical conclusions drawn from the answers will be enough to weed out inauthentic visuals. You can access the full list of questions, put together by Manchester Metropolitan University experts, here.

T.J. Thomson, Senior Lecturer in Visual Communication & Media, *Queensland University of Technology*; Daniel Angus, Associate Professor in Digital Communication, *Queensland University of Technology*, and Paula Dootson, Senior Lecturer, *Queensland University of Technology*