Thanks to AI, the future of 'fake news' may be easily-faked video

Julian Sanchez

February 8, 2018

From the printing press and home VCRs to Snapchat and virtual reality, the pervasive desire to look at attractive naked people has been a great unsung driver of technological progress. If you want to know where technology is going, in other words, a good rule of thumb is: Look to porn.

When it comes to the future of news, however, that advice may leave you feeling unsettled — and for reasons having nothing to do with prudery.

As the technology news site Motherboard reported late last month, the latest merger of high tech and low urges is a phenomenon dubbed "deepfakes." Using free, readily available software, the everyday horndog can now swap the faces of celebrities — or anyone else — into pornographic videos. While once such fakery would have required advanced video editing skills, the FakeApp, designed for the convenience of deepfake aficionados, makes use of machine learning algorithms to produce what is, in effect, a video editing Artificial Intelligence.

The upshot is that shoehorning an onscreen — or real life — crush into an ersatz but highly convincing porn no longer requires a serious technical background.

That ought to be disturbing enough: Most of us would rather not contemplate the prospect of discovering we've been unwillingly cast in an obscene video that's gone viral online, even if it's known to be a fake. (Some major porn sites are now barring the phony videos, though plenty remain in circulation.)

But perhaps even more unsettling should be the inevitable application of this free-to-download tech to politics and journalism. Combined with software like Adobe Voco, which can create a pitch-perfect virtual simulation of anyone's voice based on a short audio sample, you've got a recipe for realistic viral "fake news" fodder that the average prankster can manufacture in an afternoon.

Just imagine the October Surprise potential: The candidate caught cavorting with prostitutes, spewing racial epithets, outlining a plan to round up Lutherans for secret medical experiments! Even the most brazen political campaign might fear the damage of such a forgery being traced back to its own doorstep — but when the software to pull it off is available to anyone with a broadband connection, they likely won't have to.

In an ecosystem flooded with forged amateur videos, of course, many viewers will naturally become more skeptical about the idea that "seeing is believing." But that, too, has a cost: Recall Donald Trump's strange, belated efforts to raise doubts among his associates about the veracity of the infamous Access Hollywood "grab 'em by the pussy" tape.

In a world of fake video, such a denial might well have seemed plausible, at least to those who wished to believe. A sufficiently shameless politician might deny even actions caught on tape, with supporters given license to trust their preconceptions over their eyes.

Democratized digital fakery is nothing new, of course: Photoshopped images of political figures have long been a staple of those chain e-mails your uncle forwards along periodically. But they've typically remain confined to the fringes of political discourse for a few important reasons. One is that amateur Photoshop jobs are usually not too hard for even untrained eyes to detect: Zoom in close enough, and the pixelated hallmarks of a sloppy edit are apparent.

Just as importantly, however, is the fact that it's harder than you might initially think to construct a still image that's unambiguously scandalous without being so comically heavy-handed as to raise instinctive suspicion in the minimally savvy viewer. (For instance, most of us understand that, when politicians take bribes, they rarely come in the form of giant sacks of cash emblazoned with dollar signs.)

Audio alone, by contrast, offers more opportunity for creating plausibly scandalous content — it's much easier to concoct damaging things that a politician might unwisely say out loud in an unguarded moment — but we're all accustomed enough to hearing uncanny impersonations of famous people that an audio recording alone lacks persuasive power without a relatively ironclad provenance.

All of that, taken together, make mainstream media outlets less likely to be taken in by and amplify such forgeries. Thus, what harm they do stays confined to chain e-mails.

Video combined with audio, however, is another matter, especially as algorithmic assistants get better at concealing the more obvious digital artifacts of editing. Even in an era of sophisticated CGI, we are all still inclined to believe what we can both see and hear.

Maybe more importantly, something that's caught on video makes for good television. And the technology is arriving precisely as the incentives that media outlets face make them less able to resist paying attention to something that's gone viral.

Recall the path taken by the now-infamous "Steele Dossier," the research compiled by a former British Intelligence officer purporting to document collusion between the Trump campaign and the Russian government. Many media outlets had obtained copies of the dossier, but because —

however much of it ultimately proves accurate — they could not verify its claims, it remained unpublished.

Until, that is, the online news site Buzzfeed decided that the dossier was sufficiently newsworthy to publish with the caveat that its allegations could not be confirmed. Instantly, the fact that one news organization had run with the story was itself a newsworthy development that others could justify covering.

In the Internet Era, there are no more regional media oligopolies. Every news outlet is, essentially, in competition with hundreds, if not thousands, of others. That makes the traditional benign paternalism exercised by news organizations much harder to sustain economically: If you don't run with that explosive video, your competitors may — and when there are thousands of competitors, it becomes a near certainty that at least one will.

A site with dubious journalistic standards deciding that a fake clip is "newsworthy" merely on the grounds that it has gone viral on social media can plausibly kick off a chain reaction, as more credible outlets rush to cover the coverage, lest they be left last in the increasingly pitiless competition for eyeballs.

Technology has made it easier to fake; the economics of the internet make it increasingly likely that the fakes become news. And the inevitable blunders will confirm diminishing public trust in professional news media — the effect of which to date, ironically, has been to drive many viewers and readers into the arms of outlets with even fewer journalistic scruples.

Eventually, of course, both news producers and news consumers will adapt to the new reality, with some combination of professional protocols and personal skepticism. But the chaotic period of fumbling toward a new equilibrium promises to be a wild ride.

Julian Sanchez is a senior fellow at the Cato Institute studying issues at the busy intersection of technology, privacy, and civil liberties. He is a contributing editor for Reason magazine and a founding editor of the policy blog Just Security.