



Since January 2020 Elsevier has created a COVID-19 resource centre with free information in English and Mandarin on the novel coronavirus COVID-19. The COVID-19 resource centre is hosted on Elsevier Connect, the company's public news and information website.

Elsevier hereby grants permission to make all its COVID-19-related research that is available on the COVID-19 resource centre - including this research content - immediately available in PubMed Central and other publicly funded repositories, such as the WHO COVID database with rights for unrestricted research re-use and analyses in any form or by any means with acknowledgement of the original source. These permissions are granted for free by Elsevier for as long as the COVID-19 resource centre remains active.

The Potential Dangers of Artificial Intelligence for Radiology and Radiologists

Linda C. Chu, MD, Anima Anandkumar, PhD, Hoo Chang Shin, PhD, Elliot K. Fishman, MD

With the advent of artificial intelligence (AI) across many fields and subspecialties, there are considerable expectations for transformative impact. However, there are also concerns regarding the potential abuse of AI. Many scientists have been worried about the dangers of AI leading to “biased” conclusions, in part because of the enthusiasm of the inventor or overenthusiasm among the general public. Here, though, we consider some scenarios in which people may intend to cause potential errors within data sets of analyzed information, resulting in incorrect conclusions and leading to potential problems with patient care and outcomes.

A generative adversarial network (GAN) is a recently developed deep-learning model aimed at creating new images. It simultaneously trains a generator and a discriminator network, which serves to generate artificial images and to discriminate real from artificial images, respectively. We have recently described how GANs can produce artificial images of people and audio content that fool the recipient into believing that they are authentic. As applied to medical imaging, GANs can generate synthetic images that can alter lesion size, location, and transpose abnormalities onto normal examinations (Fig. 1) [1]. GANs have the potential to improve image quality, reduce radiation dose, augment data for training algorithms,

and perform automated image segmentation [2]. However, there is also the potential for harm if these artificial images infiltrate our health care system by hackers with malicious intent. As proof of principle, Mirsky et al [3] showed that they were able to tamper with CT scans and artificially inject or remove lung cancers on the images. When the radiologists were blinded to the attack, this hack had a 99.2% success rate for cancer injection and a 95.8% success rate for cancer removal. Even when the radiologists were warned about the attack, the success of cancer injection decreased to 70%, but the cancer removal success rate remained high at 90%. This illustrates the sophistication and realistic appearance of such artificial images. These hacks can be targeted against specific patients or can be used as a more general attack on our radiologic data. It is already challenging enough to keep up with the daily clinical volume when the radiology system is running smoothly. Our clinical workflow would be paralyzed if we could not trust the authenticity of the images and must spend extra effort searching for evidence of image tampering in every case.

There are multiple access points within the chain of image acquisition and delivery that can be corrupted by attackers, including the scanner,

PACS, server, and workstations [3]. Unfortunately, data security is poorly developed and poorly standardized in radiology. In 2016, Stites and Pinykh [4] performed a scan through the World Wide Web of networked computers and devices and showed that there were 2,774 unprotected radiology or DICOM servers worldwide, most of them located in the United States. To date, there has been no known hack into the radiology system, aside from the research study demonstrating its feasibility [3]. However, the vulnerability is clearly present and may be exploited by hackers.

Such threats could affect not only radiology departments but also entire health systems. We have all read articles about security breaches of medical records. There have been almost 3,000 breaches (involving more than 500 medical records) in the United States within the past 10 years. This includes high-profile cases such as the 2015 breach of the Anthem medical insurance company that potentially exposed the medical records of 78 million Americans and led to a \$115 million settlement [5]. Hospitals and clinics have been held hostage when their data were corrupted by a third party that demanded payment (ransom) to release the data [5]. In 2017, ransomware WannaCry and NotPetya spread through thousands of institutions worldwide, including




Fig 1. Examples of images artificially generated using generative adversarial network of brain tumor MR images. First column: T1-weighted images; second column: T1-weighted images with contrast; third column: T2-weighted images; fourth column: fluid-attenuated inversion recovery images. First row: original images with tumor in the right frontal lobe (arrows). Second row: tumor is made 16% larger. Third row: tumor is made 16% smaller. Fourth row: tumor is artificially placed on an otherwise tumor-free brain.

many hospitals, and caused \$18 billion in damages [5]. Hospitals and clinics have not been the only targets. The city of Baltimore was essentially out of business for a month this past year because of such a ransomware attack. At first glance, all of these situations seem more likely in a movie made for Netflix or HBO. However, the truth is that we must be prepared to deal with such scenarios in the near future. As electronic health records and hospital data become more centralized and more computerized, the dangers only multiply.

However, there are several ways to mitigate potential AI-based hacks and attacks. These include clear security guidelines and protocols that are uniform across the globe. As deep-fake technology gets more sophisticated, there is emerging research on AI-driven defense strategies. One example features the training of an AI to detect artificial images by image artifacts induced by GAN [6]. However, AI-driven defense mechanisms have a long way to catch up, as seen in the related problem of defense against adversarial attacks. Recognizing these challenges, the Defense Advanced Research

Projects Agency has launched the Media Forensics program to research against deep fakes [7]. Hence, for now, the best defense against deep fakes is based on traditional cybersecurity best practices: secure all stages in the pipeline and enable strong encryption and monitoring tools.

In the current coronavirus disease 2019 pandemic, many clinicians and radiologists have turned to working remotely in attempts to “flatten the curve” and slow the spread of disease. In the body imaging division at our institution, currently, approximately

half of the radiologists are working remotely. Many of our clinicians are transitioning to telemedicine visits, which adds tremendous stress on our networks. Our IT department has been proactive in setting up a dedicated virtual private network for radiology to ensure that there is sufficient bandwidth for our clinical work. On our few on-site rotations, we practice “social distancing,” and we have suspended our all side-by-side readouts and in-person lectures. We have turned to Zoom (San Jose, California) and other mobile platforms for managing our rapidly changing clinical operations, educating trainees, or simply staying in touch during these uncertain times. The daily meeting participants rose from 10 million daily users in December 2019 to 200 million daily users in March 2020 [8]. Our reliance on Zoom and other mobile platforms has exposed a new vulnerability. There has been a proliferation of “Zoombombing,” in which intruders hijack video calls and paste hate speech and offensive images. Furthermore, additional vulnerabilities in Zoom can allow hackers to gain control of users’ microphones and webcams and steal login credentials. Zoom video meetings did not provide end-to-end encryption as promised, and a large number of Zoom video meeting recordings, many of which contain private information, are left

unprotected and viewable on the web. The Federal Bureau of Investigation has issued security warnings about Zoom, and a number of organizations, including SpaceX, Google, New York’s Department of Education, and the US Senate, have banned or discouraged the use of Zoom [8]. The meteoric rise and fall of Zoom is a cautionary tale about the importance of data security.

With the development of AI and all its potential wonders in terms of increasing the accuracy of our diagnostic capabilities and potentially improving patient care, we must also be concerned about the potential dark side by bad actors. The sooner organized radiology and organized medicine address these issues with clarity, the more stable and protected the health care system and our patients will be from those intent on creating harm and havoc by abusing AI. The acceleration of data sharing during the current pandemic exposes critical vulnerabilities in data security. It reminds us of the pervasive threat that bad actors can and will exploit any technology for their selfish gains. Doing nothing is not a viable strategy, but acting in a concerted effort will lead us to the protection we need and is important as we push AI development over the next several years.

ACKNOWLEDGMENTS

The authors thank senior science editor Edmund Weisberg, MS, MBE, for his editorial assistance.

REFERENCES

1. Shin H-C, Tenenholtz NA, Rogers JK, et al. Medical image synthesis for data augmentation and anonymization using generative adversarial networks. In: Gooya A, Goksel O, Oguz I, Burgos N, eds. *Simulation and synthesis in medical imaging*. Cham, Switzerland: Springer.
2. Sorin V, Barash Y, Konen E, Klang E. Creating artificial images for radiology applications using generative adversarial networks (GANs)—a systematic review. *Acad Radiol*. In press.
3. Mirsky Y, Mahler T, Shelef I, Elovici Y. GAN CT-. malicious tampering of 3D medical imagery using deep learning. Available at: https://www.usenix.org/system/files/sec19-mirsky_0.pdf. Accessed April 19, 2020.
4. Stites M, Pinykh OS. How secure is your radiology department? Mapping Digital radiology adoption and security worldwide. *AJR Am J Roentgenol* 2016;206:797-804.
5. Desjardins B, Mirsky Y, Ortiz MP, et al. DICOM images have been hacked! Now what? *AJR Am J Roentgenol* 2020;214:727-35.
6. Zhang X, Karaman S, Chang SF. Detecting and simulating artifacts in GAN fake images. In: 2019 IEEE International Workshop on Information Forensics and Security (WIFS). Piscataway, NJ: Institute of Electrical and Electronics Engineers.
7. Turek M. Media Forensics (MediFor). Available at: <https://www.darpa.mil/program/media-forensics>. Accessed April 9, 2020.
8. Hodge R. Zoom: every security issue uncovered in the video chat app. Available at: <https://www.cnet.com/news/zoom-every-security-issue-uncovered-in-the-video-chat-app/>. Accessed April 9, 2020.

Linda C. Chu, MD, Associate Professor, The Russell H. Morgan Department of Radiology and Radiological Science, Johns Hopkins University School of Medicine, Baltimore, Maryland. Anima Anandkumar, PhD, Department of Computing and Mathematical Science, California Institute of Technology, Pasadena, California; Director of Machine Learning Research, NVIDIA Corporation, Santa Clara, California. Hoo Chang Shin, PhD, NVIDIA Corporation, Santa Clara, California. Elliot K. Fishman, MD, Professor, Johns Hopkins Medicine, Department of Radiology and Radiological Science, Department of Oncology and Department of Surgery. He serves as Director of Diagnostic Imaging and Body CT at Johns Hopkins University School of Medicine, Baltimore, Maryland.

The authors state that they have no conflict of interest related to the material discussed in this article.

Linda C. Chu, MD: The Russell H. Morgan Department of Radiology and Radiological Science, Hal B168, Johns Hopkins Hospital, 600 North Wolfe Street, Baltimore, MD 21287; e-mail: lindachu@jhmi.edu