AI & SOCIETY (2024) 39:1969–1973
https://doi.org/10.1007/s00146-023-01655-8

OPEN FORUM

Privacy and surveillance concerns in machine learning fall prediction models: implications for geriatric care and the internet of medical things

Russell Yang[1]

Received: 15 February 2022 / Accepted: 29 March 2023 / Published online: 18 April 2023
© The Author(s), under exclusive licence to Springer-Verlag London Ltd., part of Springer Nature 2023

Abstract
Fall prediction using machine learning has become one of the most fruitful and socially relevant applications of computer vision in gerontological research. Since its inception in the early 2000s, this subfield has proliferated into a robust body of research underpinned by various machine learning algorithms (including neural networks, support vector machines, and decision trees) as well as statistical modeling approaches (Markov chains, Gaussian mixture models, and hidden Markov models). Furthermore, some advancements have been translated into commercial and clinical practice, with companies in various stages of development capitalizing on the aging population to develop new commercially available products. Yet despite the marvel of modern machine learning-enabled fall prediction, little research has been conducted to shed light on the security and privacy concerns that such systems pose for older adults. The present study employs an interdisciplinary lens in examining privacy issues associated with machine learning fall prediction and exploring the implications of these models in elderly care and the Internet of Medical Things (IoMT). Ultimately, a justice-informed set of best practices rooted in social geroscience is suggested to help fall prediction researchers and companies continue to advance the field while preserving elderly privacy and autonomy.

Keywords Machine learning · Artificial intelligence · Computer vision · Cybersecurity · Privacy · Security · Gerontology · Social gerontology · Internet of medical things · Best practices

1 Introduction

The 2004 paper (Sixsmith and Johnson 2004) represented one of the first end-to-end deep learning-based systems for prediction/detection of falls in older adults. SIMBAD, as coined by the researchers from the University of Liverpool and InfraRed Integrated Systems (IRISYS), was a neural network-enabled system that analyzed thermal data from an IRISYS-produced sensor to provide detection of falls in a specific set of curated actor scenarios. Although the model performance had some notable flaws and shortcomings (i.e., detecting a fall just 30% of the time and being focused on detection as opposed to prediction), the study itself had a profound impact on machine learning fall-related technology and also sheds light on how the cooperation between academia and industry within the fall prediction space is a topic worthy of critical examination.

It is precisely this lens of private–public cooperation and interfacing, paired with the fact that many fall prediction models employ real-time video and physiological data, that distinguishes the fall prediction problem from other prediction tasks in the field of geroscience. For instance, machine learning-based drug discovery for diseases/conditions that predominantly affect older adults represents another significant intersection of machine learning and geroscience (Zhavoronkov et al. 2021). Drug discovery methods may present some privacy issues related to big data collection and storage, but once developed into clinically applicable medications, said privacy and security issues become moot. By contrast, machine learning models that rely on real-time video and physiological sensor feeds continuously read and analyze patient data during their operation. Thus, the inherently longitudinal nature of fall prediction as an instrument of preventative medicine suggests that it is coupled with unique privacy concerns. This becomes especially clear

✉ Russell Yang
russell.yang@yale.edu

[1] Yale University, New Haven, CT, USA

when returning to the topic of private–public cooperation, as companies are held to different ethical standards than IRB-approved human subject research endeavors. To add yet another complication, data policies can vary by country (McIntyre 2008; Rojszczak 2021; Fabbrini 2015) and are continually evolving to match the pace of workplace adoption of data-driven methodologies.

The present study examines fall prediction machine learning approaches from perspectives of research frontiers, commercial products, and social gerontology, ultimately finding that if not properly addressed, fall prediction models can pose a serious threat to the privacy and autonomy of older adults. More broadly, the ways in which older adults interact with the Internet of Medical Things (IoMT) are also discussed in both practical and conceptual terms.

1.1 Research frontiers in fall prediction

Integration of machine learning into fall prediction required a convergence of several research movements: modern machine learning, evidence-based medicine, gerontology, and the proliferation of computing resources. In contemporary society, all four of these are widely available and active topics of discovery, pushing the subfield of machine learning fall prediction into the spotlight. As one review explains, fall prevention and detection methods have employed myriad machine learning algorithms (including neural networks, support vector machines, and decision trees) as well as probabilistic/statistical modeling approaches (Markov chains, Gaussian mixture models, and hidden Markov models) (Pannurat et al. 2014). Current frontiers being explored include a Timed Up and Go (TUG) based approach (Roshdibenam et al. 2021), motion-pose geometric descriptor (MPGD) based approach (Alazrai et al. 2015), as well as a thermal imaging approach (Song et al. 2017) not too dissimilar from Sixsmith and Johnson (2004). However, perhaps the most important discriminating factor between modern approaches to machine learning fall prediction is the data being fed to the model. Using this paradigm, we principally distinguish between three different classes of models and summarize potential privacy issues in Table 1 below.

As evinced in the 'Connected Sensing' column of Table 1, any device in the Internet of Things (IoT) is subject to potential compromise. (Patton et al. 2014) estimated that the vulnerability rate can run anywhere from less than half a percent to 40% for certain types of IoT devices. Moreover, this challenge becomes especially important for multimodal fall prediction systems – systems that utilize more than one data source in their predictions. As the number of connected devices increases, the chance of any one device being compromised increases, as does the possibility that multiple types of data (e.g., video and sensor data) might be exposed.

1.2 Commercial fall prediction systems

There are various solutions for fall prediction that are being developed or are already available for purchase, although they vary in terms of business development stage and target market. One company, Ocuvera, produces a portable system that includes a depth camera and can notify on-site professionals in case a patient might be about to "exit a bed or chair" (Ocuvera 2022). However, not all companies are using cameras for prediction. Qventus, for instance, computes risk scores based on call-light usage, medications, and data from an electronic health record (EHR) info (Qventus. 2022; Scott 2017). Dele utilizes sensors in conjunction with EHRs as inputs for its modeling (Dele 2022). Although all three of these services achieve a similar goal—notifying a professional if a patient is likely to experience a fall–they vary in terms of data sources. The latter two approaches involve computations of risk scores based on clinical indicators, which is less intrusive than the former, with its camera-enabled system. However, a sweeping generalization of all camera-based machine learning fall prediction systems as insecure belies the import and practical significance of such systems. While risk scores might give medical professionals

Table 1 Classes of different machine learning fall prediction models, characteristics, and associated potential privacy issues

Input data type	Characteristics	Video of patient	Audio of patient	Connected sensing
Video or Camera	May employ convolutional neural networks (CNNs) or other computer vision techniques	Potential concern, especially if continuously employed	Not a significant concern in most cases	Cameras can have security issues, could be compromised by bad actors
Physiological Sensor	Generally wearable Could be implemented using smartphone telemetry data	Not a significant concern in most cases	Not a significant concern in most cases	Physiological indicators are personal, but not as significant as video/camera
Multimodal/Combination	Combination of data types May employ data fusion to reconcile different information	Varies depending on data sources	Varies depending on data sources	Potential concern, especially because multiple IoT devices compounds likelihood of security issues

broad insight into patient likelihood of falling, they don't give the same granular, precise, and actionable information that camera-based systems do. In the end, the current commercial landscape of machine learning fall prediction solutions is characterized by a tradeoff between privacy and utility. Purely sensor-based systems might be able to solve this tradeoff by eliminating the need for cameras, but said approaches also bring unique challenges like patient compliance (which might be especially important if patients are affected by medication, memory issues, or other health conditions) and a need to charge devices.

1.3 Technology awareness in older populations

Computer vision-based surveillance and analysis techniques are unique in historical context: although machine learning traces back to the 1940s or 1950s, the bulk of the field was expostulated in the past twenty years. Thus, in the design of research studies or acceptance of commercial fall systems, older adults might not have the technical understanding of machine learning fall based prediction systems to give informed consent to their utilization. Furthermore, when older adults are affected by neurodegenerative diseases like dementia, ethical concerns arise as to whether patients can truly give informed consent (Whitehouse 2000). Beyond the specific methodologies employed for prediction (i.e., machine learning), it is also imperative that older adults understand how their data is being stored and used. It is clear that in the development of machine learning fall prediction technology, especially that which relies on comparatively intrusive data collection, researchers and companies must carefully consider the needs of older adults and their education and awareness related to technical topics. Enforcement of intrusive machine learning fall prediction systems onto unconsented elderly individuals presents a significant concern that could be addressed by national legislation or guidance from trusted medical institutions. Furthermore, the sparse, uneven nature of data protection policy in the United States contrasts with legislation in other high-income regions (i.e., the European Union) (Steinke 2002). We call for new data protection legislation that is sensible, consistent, specific, and adequately addresses the needs of older adults as well as the usage of data in data-driven processes like machine learning systems.

1.4 Security implications for the internet of medical things

The Internet of Medical Things (IoMT), describes the subset of IoT as applied in healthcare and medicine. (Yaacoub et al. 2020) delineated key security issues related to IoMT devices. Specifically, improperly secured networks could allow bad actors to "eavesdrop and intercept incoming and outgoing data" and IoMT devices could be subject to attacks like DDoSing (distributed denial of service) (Yaacoub et al. 2020). While data compromise is an extremely important issue, the latter possibility arguably poses an even greater threat to patients, because it could result in the inactivation or hampering of medical devices. The paper (Yaacoub et al. 2020) provides a salient case in point–VP Dick Cheney asked medical professionals to remove wireless communication abilities from his pacemaker so that he couldn't be targeted by bad actors (Yaacoub et al. 2020; Peterson 2013). In the case of machine learning fall prediction, malicious cyberattacks might result in leakage of personal health information or disabling of fall prediction systems, but they most likely would not immediately lead to serious injury or death as a compromised pacemaker might. Nonetheless, examining IoMT devices through a lens of cybersecurity is critically important, even more so than it is for non-medical IoT devices like smart speakers and appliances because of the potential to impact human health.

1.5 Solutions for facial de-identification

Solutions to the privacy issues raised by fall prediction systems must address all stages of the data lifecycle—from collection to destruction and/or archival. Next, we review state-of-the-art strategies for the "de-identification" of faces in image/video data and comment on how these strategies might help solve privacy issues for fall prediction systems.

A broad segment of the literature focuses on the so-called "de-identification" of faces in surveillance videos using novel computer vision algorithms (Newton et al. 2005; Korshunov et al. 2013; Nakashima et al. 2015). A 2005 study by Newton et al. (Newton et al. 2005) described an approach called k-Same aimed to combat the legacy principal components analysis (PCA)-based recognition model that was introduced in the 1990s (Turk and Pentland 1991). The k-Same approach looks for "close" faces in the input data and swaps out those with an averaged face. When evaluated against the U.S. Army Face Recognition Technology (FERET) dataset, the researchers demonstrated that the proposed k-Same algorithm prevented the identification of individual faces by the PCA procedure (Newton et al. 2005). Notably, Newton et al. also demonstrated that ad-hoc ("common-sense") transformations such as pixelation and thresholding did little to impede the recognition algorithm (Newton et al. 2005). Later work in 2013 (Korshunov et al. 2013) studied geometrical warping of images as a defense against face recognition and tested against the Viola-Jones model, a boosted classifier paradigm from the early 2000s (Viola et al. 2001). Ultimately, the researchers concluded that sufficient warping could render a face unidentifiable by the Viola-Jones model while also quantifying the tradeoff

between the severity of warping and probability of identification (Korshunov et al. 2013). Finally, Nakashima et al. developed a melding method that successfully conserved facial expression as measured by a survey (Nakashima et al. 2015) while also circumventing recognition by a 2006 facial recognition procedure involving local binary patterns (Ahonen et al. 2006). One persistent challenge shared by these "de-identification" methodologies is that nearly every approach performs testing on legacy recognition systems. Yet these studies undoubtedly demonstrate that robust facial recognition de-identification can be achieved through a myriad of approaches, suggesting that the identification of particular individuals among a group can be avoided. This might find clinical applications in a senior living environment, where many older adults share a common space and fall prediction is outsourced to an external vendor or images/video are shared with non-medical personnel for quality control or other purposes.

2 Edge computing: a paradigm for privacy-conscious data storage/analysis

Comprehensively enforcing privacy and security in automated fall prediction systems also demands an examination of the ways in which data is stored and analyzed. Edge computing, a new paradigm for IoT devices, promises to solve the dual problems of storage and analysis, especially in healthcare settings (Hartmann et al. 2022). In edge computing systems, data can be processed and predictions can be made on the local device, which decreases latency and improves privacy by minimizing the exchange of data between local devices and a cloud server (Hartmann et al. 2022). Edge computing has successfully been applied in wearable devices, where local algorithms provide cognitive cues to the users in a variety of simple spatiotemporal tasks (Chen et al. 2017). In (Chen et al. 2017), the researchers found that a multi-algorithm edge computing approach helped cut latency by more than 60%, without any notable change in accuracy. Both of these advantages are crucially important in the problem of fall prediction: reducing latency would help professionals more quickly respond to emergencies, and keeping processing local ensures that fewer bad actors have the opportunity to intercept or interfere with data transmission. One of the principal disadvantages of an edge computing paradigm is that more hardware and resources are required on the local devices, which might introduce deployment challenges for administrators of fall prediction systems (Liu et al. 2019).

3 Best practices and a call to action

In synthesizing the current research and practice related to machine learning fall prediction systems, we propose the following set of consistent best practices:

Machine learning fall prediction research studies and commercial systems must seek to promote comprehensive patient education (especially for older populations) and informed consent.
Internet-connected systems should seek to minimize cybersecurity risks by complying with best practices and working with cybersecurity and audit professionals.
Future research in the field of machine learning fall prediction should include discussion of sociocultural implications of predictive systems and potential mental health implications of said technology.
Best data privacy policies must be followed, and new legislation and guidance related to data privacy for older adults would be instrumental in ensuring that machine learning fall prediction preserves privacy for the elderly.
Computer vision researchers must evaluate and compare methods for de-identification of faces using state-of-the-art models
Industry and academia should convene broad, interdisciplinary teams to study edge computing as a paradigm for low latency, privacy-conscious automated fall prediction systems

Ultimately, machine learning fall prediction systems comprise an interesting solution to an age-old problem and have the potential to prevent untold injuries and deaths. However, these systems also demand a justice-informed, privacy-preserving approach to ensure their ethical development and real-world deployment. It is only through an interdisciplinary lens combining gerontology, machine learning science, and ethics that the subfield can thrive.

Declarations

Conflict of interest The author declares that there is no conflict of interest.

References

Ahonen T, Hadid A, Pietikainen M (2006) Face description with local binary patterns: application to face recognition. IEEE Trans Pattern Anal Mach Intell 28(12):2037–2041
Alazrai R, Mowafi Y, Hamad E, 2015 editors. A fall prediction methodology for elderly based on a depth camera. 2015 37th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC). Pp. 25–29

Chen Z, Hu W, Wang J, Zhao S, Amos B, Wu G, et al. An empirical study of latency in an emerging class of edge computing applications for wearable cognitive assistance. Proceedings of the Second ACM/IEEE Symposium on Edge Computing. 2017.

Dele. Fall Prevention [cited 2022 Jan 18, 2022]. Available from: https://delehealth.com/solutions/prevention/.

Fabbrini F (2015) Human rights in the digital age: The European Court of Justice ruling in the data retention case and its lessons for privacy and surveillance in the United States. Harv Hum Rights J 28:65–95

Hartmann M, Hashmi US, Imran A (2022) Edge computing in smart health care systems: review, challenges, and research directions. Trans Emerg Telecommun Technol 33(3):e3710

Korshunov P, Ebrahimi T, editors. Using warping for privacy protection in video surveillance. 2013 18th International Conference on Digital Signal Processing (DSP); 2013 1–3 July 2013.

Liu F, Tang G, Li Y, Cai Z, Zhang X, Zhou T (2019) A Survey on edge computing systems and tools. Proc IEEE 107(8):1537–1562

McIntyre TJ (2008) Data retention in Ireland: privacy, policy and proportionality. Comput Law Secur Rev 24(4):326–334

Nakashima Y, Koyama T, Yokoya N, Babaguchi N, editors. Facial expression preserving privacy protection using image melding. 2015 IEEE International Conference on Multimedia and Expo (ICME); 2015 29 June-3 July 2015.

Newton EM, Sweeney L, Malin B (2005) Preserving privacy by de-identifying face images. IEEE Trans Knowl Data Eng 17(2):232–243

Ocuvera. AI technology that empowers nurses to prevent patient falls. [cited 2022 Jan 18, 2022]. Available from: https://ocuvera.com/.

Pannurat N, Thiemjarus S, Nantajeewarawat E (2014) Automatic fall monitoring: a review. Sensors 14(7):12900

Patton M, Gross E, Chinn R, Forbis S, Walker L, Chen H, 2014 editors. Uninvited Connections: A Study of Vulnerable Devices on the Internet of Things (IoT). 2014 IEEE Joint Intelligence and Security Informatics Conference; 2014 24–26

Peterson A. Yes, terrorists could have hacked Dick Cheney's heart: The Washington Post; 2013 [cited 2022 Jan 18, 2022]. Available from: https://www.washingtonpost.com/news/the-switch/wp/2013/10/21/yes-terrorists-could-have-hacked-dick-cheneys-heart/.

Qventus. The Leader in AI-Powered Healthcare Operations [cited 2022 Jan 18, 2022]. Available from: https://qventus.com/.

Rojszczak M (2021) The uncertain future of data retention laws in the EU: Is a legislative reset possible? Comput Law Secur Rev 41:105572

Roshdibenam V, Jogerst GJ, Butler NR, Baek S (2021) Machine learning prediction of fall risk in older adults using timed up and go test kinematics. Sensors 21(10):3481

Scott J. Preventing Seniors From Falling is Going to Be a Huge Market: KQED; 2017 [cited 2022 Jan 18, 2022]. Available from: https://www.kqed.org/futureofyou/435417/dying-from-a-fall-is-top-danger-for-seniors-tech-devices-may-help.

Sixsmith A, Johnson N (2004) A smart sensor to detect the falls of the elderly. IEEE Pervasive Comput 3(2):42–47

Song KS, Nho YH, Kwon DS, 2017 editors. Histogram based fall prediction of patients using a thermal imagery camera. 2017 14th International Conference on Ubiquitous Robots and Ambient Intelligence (URAI); 2017 28 June-1.

Steinke G (2002) Data privacy approaches from US and EU perspectives. Telemat Inform 19(2):193–200

Turk M, Pentland A (1991) Eigenfaces for recognition. J Cogn Neurosci 3(1):71–86

Viola P, Jones M, editors. Rapid object detection using a boosted cascade of simple features. Proceedings of the 2001 IEEE Computer Society Conference on Computer Vision and Pattern Recognition CVPR 2001; 2001 8–14 Dec. 2001.

Whitehouse PJ (2000) Ethical issues in dementia. Dialogues Clin Neurosci 2(2):162–167

Yaacoub JPA, Noura M, Noura HN, Salman O, Yaacoub E, Couturier R et al (2020) Securing internet of medical things systems: Limitations, issues and recommendations. Future Gener Comput Syst 105:581–606

Zhavoronkov A, Bischof E, Lee K-F (2021) Artificial intelligence in longevity medicine. Nature Aging 1(1):5–7

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.