 An audio interview with Dr. Wilensky is available at NEJM.org

representative medical groups and, if necessary, to their representatives in Congress as well. In the past, practicing clinicians have been woefully bad at making their voices heard. Now is a good time for that to change.

Disclosure forms provided by the author are available at NEJM.org.

From Project HOPE, Bethesda, MD.

This article was published on February 28, 2018, at NEJM.org.

1. Wilensky GR. Reforming Medicare's physician payment system. *N Engl J Med* 2009;360:653-5.

2. Wilensky GR. Improving value in Medicare with an SGR fix. *N Engl J Med* 2014;370:1-3.

3. Medicare Payment Advisory Commission. Medicare and the health care delivery system. Report to the Congress. June 2017 (http://www.medpac.gov/docs/default-source/reports/jun17_reporttocongress_sec.pdf).

DOI: 10.1056/NEJMp1801673

Copyright © 2018 Massachusetts Medical Society.

Bedside Computer Vision — Moving Artificial Intelligence from Driver Assistance to Patient Safety

Serena Yeung, M.S., N. Lance Downing, M.D., Li Fei-Fei, Ph.D., and Arnold Milstein, M.D., M.P.H.

The Institute of Medicine's 1999 report on preventable patient harm in hospitals provoked requirements for public reporting of errors and financial penalties for preventable hospital-acquired conditions. Checklists, protocols, root-cause analyses, programs for creating a culture of safety, and early forms of technological assistance followed. Yet the goal of zero patient harm continues to elude hospitals, particularly when it comes to deviations from intended bedside practices ranging from reliable hand hygiene to central line insertions.

We may be approaching the limits of what is achievable through improvements in clinical processes, culture, and narrowly focused technological assistance. Expectations that fatigued clinicians will reliably execute each behavioral step of complex hospital treatments ignore evidence from cognitive science that humans usually operate in error-prone “fast thinking” mode.¹ Even remotely located hospital staff watching intensive care beds by video feed cannot immediately detect and correct bedside behavioral errors such as failing to re-


set bedrails, restraints, or inflatable calf boots.

A source of clinician assistance may lie in a rapidly progressing domain of artificial intelligence (AI) known as computer vision. Broadly defined as the development of intelligent machines, the field of AI focuses on both capabilities, such as understanding spoken language, and development methods, such as machine learning. Computer vision allows machines to see and understand the visual world. Machine learning entails building knowledge from patterns in data rather than being specified by human programmers. When applied to computer-vision tasks such as discernment of people, objects, and their motion, cameras and imaging sensors supply data for learning. For example, exposed to data of thousands of digitized dog photographs labeled according to breed, computers can deploy machine-learning methods to digest the data during a “training” phase and devise an algorithm that accurately distinguishes among dog breeds.

No longer science fiction, computer vision is improving rapidly,

in part thanks to “deep learning,” a type of machine learning that uses multilayered neural networks whose hierarchical computational design is partly inspired by a biologic neuron's structure. A reference point for the speed of improvement in computer vision is Google's computer-vision system for supporting self-driving vehicles. Over a recent 12-month span, its performance advanced from requiring human intervention every 700 miles to fully autonomous driving for more than 5000 miles at a time (see graph). If computer vision can detect when drivers initiate dangerous lane changes and safely control vehicular steering, can it similarly analyze motion to detect unintended deviations in important clinician behaviors or patient activities?

There are reasons to be optimistic that computer-vision applications will prove clinically useful. Computer vision is poised to gain a foothold in screening medical images for clinician analysis. A recent study found that computer vision performed on par with 21 board-certified dermatologists in classifying digital images of benign and malignant



Progress of Computer Vision for Self-Driving Vehicles, 2014–2015.

Data are the average numbers of miles driven autonomously by the Google self-driving car before human intervention was required. Data are from the Autonomous Vehicle Engagement Reports of the California Department of Motor Vehicles.

skin lesions.² Small studies show similar early progress in the interpretation of radiologic and pathological images. Beyond static medical images, research is expanding into machine interpretation of video data of clinician and patient behaviors. University of Strasbourg researchers equipped an operating room with sensors and demonstrated accurate computer recognition of surgical workflow. Johns Hopkins researchers applied computer vision in an intensive care unit to quantify progress in patients' mobility.

Now, researchers from Stanford's engineering and medical schools, Lucile Packard Children's Hospital (LPCH), and Intermountain LDS Hospital are collaborating on a hospital unit-wide deployment of computer vision to discern real-time clinician behaviors. Because of concerns about staff and patient privacy, depth and thermal sensors rather than video cameras are used to gather data for machine interpretation (see video, available with the full text of this article at NEJM.org). Depth sensors gather

rebounding infrared signals to create silhouette-like images based on the sensor's distance from the surface features of people and objects. However, the images lack some surface details that would be captured by color video. Thermal sensors, by detecting small differences in temperature on the surfaces of people and objects, enable the creation of heat-map images that reveal human shapes in motion as well as physiological events such as shallow breathing and episodes of urinary incontinence in both lighted and dark environments. The researchers are investigating whether combinations of image-sensing methods will allow accurate identification of clinically important bedside behaviors in hospital rooms, while protecting privacy.

Researchers at Stanford and Intermountain selected hand-hygiene compliance as their first target for computer-vision discernment because of its clinical importance and the evidence of its limited responsiveness to management efforts.³ Though approaches using other data types,

such as those produced by radio-frequency identification systems, have also targeted discernment of hand-hygiene behavior, the researchers hypothesized that because depth sensors capture richer, continuous image data, they would provide greater accuracy and finer discernment without interrupting clinical workflow. To evaluate the effectiveness of a computer vision-based approach, the researchers used deep learning to train a neural network to detect hand-hygiene events. Since training involves providing the neural network with labeled images from which it can learn, research staff annotated depth images of hand-hygiene events and non-hand-hygiene events at patient-room doorways. The resulting machine algorithm continuously detects hand hygiene at LPCH at a 95.5% level of accuracy using depth data alone. When applied to images from LDS Hospital, the algorithm developed at LPCH attained 84.6% accuracy without additional training with locally collected images, despite interhospital differences in the location of wall-mounted depth sensors, types of hand-sanitizer dispensers, and doorway features.

This type of computer vision, using data from ambient sensors, offers structural advantages over current systems for assessing bedside behavior such as monthly "secret shopper" observation of hand-hygiene compliance or nurse observation of protocol-based central line insertion by physicians. Ambient computer vision is ceaseless and fatigue-free, operates at very low variable cost, and is unaffected by an imperfect safety culture. Since computer vision-based discernment systems can be trained to identify diverse bedside activities, if it were integrat-



A video is available at NEJM.org

ed with electronic health records, it might also free clinicians to shift from dispiriting documentation and data-entry tasks to patient-focused activity.

Clinical uses of AI have aroused skepticism, as early applications have struggled in some settings.⁴ Threats to success include poor data quality, the difficulty of explaining the complex computational steps leading to a machine-generated clinical determination, and failure to dovetail with customary clinical workflow. By collecting data not subject to human documentation error, computer vision may mitigate one of the threats. Given health care's mixed experience with information technology, AI applications will need to overcome these challenges to move quickly from the "hype peak" to steady gains in health care value. If successfully developed and deployed, ambient computer vision carries the potential to discern diverse bedside clinician and patient behaviors at superhuman performance levels⁵ and send user-designed prompts in real time. Such systems could remind a doctor or nurse to per-

form hand hygiene if they begin to enter a patient room without doing so, alert a surgeon that an important step has been missed during a complex procedure, or notify a nurse that an agitated patient is dangerously close to pulling out an endotracheal tube. The use of computer vision to continuously monitor bedside behavior could offload low-value work better suited to machines, augmenting rather than replacing clinicians.

Much remains to be learned before such technology can be adopted widely. An apt analogy may be self-driving vehicles: they will not dominate roads immediately, yet their intermediate-term feasibility is highly plausible. Although safe hospital care presents unique challenges, if productivity gains seen in other industries are an indication, computer vision may contribute significantly to clinical quality and efficiency while freeing clinicians to focus on nuanced decision making, engaging with patients, and delivering empathic care. Given its rapid pace of improvement in accuracy and affordability in other indus-

tries, computer vision may soon bring us closer to resolving a seemingly intractable mismatch between the growing complexity of intended clinician behaviors and human vulnerability to error.

Disclosure forms provided by the authors are available at NEJM.org.

From the Department of Computer Science (S.Y., L.F.-F.), the Center for Biomedical Informatics Research (N.L.D.), the Department of Medicine (A.M.), and the Clinical Excellence Research Center (S.Y., N.L.D., L.F.-F., A.M.), Stanford University, Stanford, CA.

1. Kahneman D. Thinking, fast and slow. New York: Farrar, Straus and Giroux, 2011.
2. Esteva A, Kuprel B, Novoa RA, et al. Dermatologist-level classification of skin cancer with deep neural networks. *Nature* 2017;542:115-8.
3. Haque A, Guo M, Alahi A, et al. Towards vision-based smart hospitals: a system for tracking and monitoring hand hygiene compliance. *Proc Mach Learn Res* 2017;68:75-87 (<https://arxiv.org/abs/1708.00163>).
4. Chen JH, Asch SM. Machine learning and prediction in medicine — beyond the peak of inflated expectations. *N Engl J Med* 2017;376:2507-9.
5. He K, Zhang X, Ren S, Sun J. Delving deep into rectifiers: surpassing human-level performance on ImageNet classification. In: *Proceedings of the 2015 IEEE International Conference on Computer Vision*, Santiago, Chile, December 11–18, 2015:1026-34.

DOI: 10.1056/NEJMp1716891

Copyright © 2018 Massachusetts Medical Society.

Cousin Pam

Ranjana Srivastava, F.R.A.C.P.

"Cousin Pam! Here, Cousin Pam!"

"Good morning. I'm your doctor. How do you feel today?"

"Cousin Pam, you're my doctor?" He looks puzzled.

"You are in the hospital, and I'm your doctor. May I have a look at your leg?"

"Not today, Cousin Pam. But you call Uncle Andy and we will have a meeting, and he will tell

you it's all good and where we will go, and he generally knows this stuff, you just have to call him, arrange a phone call, or tell him to be here"

Mr. J. is a new patient with a diagnosis of sepsis and probable delirium on a background of schizophrenia initially diagnosed 40 years ago. Now 60, he lives alone in supported accommodation, although the fact that he'd

absconded a week before being found makes me wonder how supported he really is.

"What happened to you?" I ask.

His calf is ripe with infection — red, angry, and disfigured. A sorry-looking crepe bandage swirls around his ankle as a hapless trainee nurse talks him into letting her clean the gaping wound.

"Is something wrong with my leg?" he asks.