Landon Cox *Duke University, Durham, N.C.*

Editor: Geoffrey Challen

HIDING SECRETS IN PLAIN SIGHT

Photo, istockphoto.com

The FIFA World Cup and NFL Super Bowl were two of the highest-profile sporting events of 2014, each drawing hundreds of millions of viewers from around the world. Such large audiences create tempting targets for all manner of ne'er-do-well, and event organizers went to great lengths to secure their games. Brazil was reported to have deployed more than 150,000 police and military officers across the World Cup's 12 sites[1]. The NFL, working with the state and federal government, used a swarm of helicopters to patrol the stadium airspace, and kept a team of F16 aircraft poised nearby, ready to jump into action if the no-fly zone surrounding the stadium was breached[2]. Likely as a deterrent, and to assuage potentially skittish attendees, FIFA and the NFL were eager to publicize how extensive their security efforts were. However, in doing so, both organizations committed identical facepalm-inducing security mistakes.

The first breach occurred during a pre-game telecast for the Super Bowl, during which CBS News highlighted the video-surveillance system NFL security was using to monitor activity in and around the stadium[3]. CBS broadcast pictures of security officers carefully watching video feeds from the stadium playing field, parking lot and spectator seats on wall of screens. Unfortunately, the NFL failed to notice a screen on the wall clearly containing a username ("marko") and password ("w3lc0m3!HERE") for the league's Wi-Fi network.

FIFA demonstrated that it had not learned from the NFL's mistake when it invited the Brazilian newspaper *Correio Braziliense* into its security command center several months later. Just like the NFL, FIFA failed to notice that the lower right-hand corner of its wall of surveillance videos contained the password for its Wi-Fi network ("b5a2112014")[4]. Although it is unclear whether either leak caused any damage, both sets of credentials were widely disseminated on social media.

Lest one conclude that only highly trained professionals in charge of multi-million dollar security apparatuses are prone to accidentally disclosing sensitive data in the presence of cameras, consider two other recent leaks. In the spring of 2015 the French television station TV5Monde accidentally broadcast the login details for its Twitter, Instagram, and YouTube accounts when a reporter forgot that the credentials were printed on a sheet of paper on his cubicle[5] – the station's YouTube password was "lemotdepassedeyoutube", "youtubepassword" in French. And passwords are not the only kind of sensitive data that can be inadvertently leaked: our esteemed colleague and friend Victor Bahl accidentally left a rough roadmap for the future of Windows Phone on his whiteboard during a publicity shot[6].

What lesson can we learn from these leaks? One is that as inept we are at protecting digital assets, such as sensitive documents and photos, we are even less prepared to protect the sensitive information on our desks, screens and whiteboards from cameras in the environment. Most people believe that their digital data resides within a tangled web of computers and networks and that only this data can be accessed

through these systems. We assume that Scandinavian script kiddies can't hack into the margins of our spiral-bound notebooks or spy the Post-it notes affixed to the frames of our monitors.

The incidents above highlight the limitations of this mental model, and how rapidly it is becoming obsolete. Even if one dismisses the hype surrounding wearable systems like Google Glass, it is impossible to deny the growing number of computer systems in our environment that sense and record the physical world. Nearly all modern consumer computers have cameras and microphones. Laptops, smartphones and gaming consoles all have these sensors, and we mainly use them for social purposes – to chat with friends, document our lives or play games. But sensors, such as cameras and microphones are poised to reach well beyond their current domains into nearly every corner of our lives.

Other networked devices in the home, such as security cameras and smart thermostats, are emerging, and robotic systems like self-driving cars, autonomous security guards and domestic health-care assistants will use cameras to navigate our homes, workplaces and roads. Wearable devices, worn on the wrist or face or even embedded into our clothes and bodies, likely represent the next wave. These pervasive sensing systems are only the beginning of a mass digitization of the physical world. If reliably protecting secrets from pervasive cameras and microphones seems daunting, trying to protect information revealed through one's body will be nightmarish.

So, what is to be done? Short of developing Harry Potter's invisibility cloak, it is probably impossible to prevent a determined adversary from recording sensitive information. Such an attacker will almost always be capable of sneaking a recording device into environments with sensitive information. I believe a more realistic goal is to prevent inadvertent leaks by trusted recording devices. This goal not only seems feasible, but may capture the vast majority settings in which cameras will be introduced. In all of the examples above – whether FIFA, the NFL, or TV5Monde – information was leaked by trusted cameras that had been invited into the workplace.

Under the assumption that recording devices are trusted, the simplest solution is to turn sensors off when sensitive information is present. This solution may work in some situations, but in general it is too coarse-grained. Sensitive and non-sensitive information often co-mingle in the same view, and as we increasingly rely on applications that require sensor data, completely shutting off the camera will become less and less appealing. For example, consider a video chat between a worker and an external collaborator. The chat feed will capture the speaker and everything in the background. To regulate access to sensitive information in a more fine-grained manner, researchers have generally taken two approaches.

First, applications can be forced to state which classes of objects they wish to access, with the system revealing only instances of those objects to the application. For example, a parking-garage robot might only need access to license plates through its camera. Alternatively, users can specify sensitive objects and decide which applications are allowed to access those objects. For example, a user could tell a recorder to remove her face from its images. The primary problem with both approaches is that the world is likely too rich for system designers to anticipate all of the objects that an application might need to access or a user might want to protect. How would one create a classifier to reliably recognize something as abstract as a Wi-Fi password or a list of upcoming mobile-operating system features?

My collaborators and I have been exploring another approach that allows security labels identifying sensitive objects to be captured in tandem with the information of interest. For images, one basic idea is to allow users to mark two-dimensional regions in their environment with a special shape, such as a rectangle with a dotted border, that is easy for someone to draw on a whiteboard and easy for a recording device to identity and block when recognized. The biggest drawback of this approach is that security can only be as robust as the computer-vision algorithms used to detect marked regions.

Our experience thus far has not been encouraging. Long-standing computer-vision challenges like the effects of lighting, occlusion and motion blur conspire to make computer vision an unstable foundation for security. Even if a vision algorithm could achieve 99 percent accuracy and recall for an object, it would still regularly miss the object over any extended period of time. For settings in which security is critical, misidentifying a single frame in a video feed is just as bad as misidentifying all of the frames. We are unaware of any vision algorithm that can achieve 99 percent accuracy and recall across a wide range of realistic settings, including the best face-detection algorithms.

Furthermore, we are unsure how to extend the idea of a marker to other kinds of perceptual data. How would one specify which sounds or movements are sensitive and which are not? And what hope would we have that the system could protect sensitive information without significantly undermining application utility? Perhaps the right approach is to ask the system to assume that all sensor data is sensitive and allow it to reveal only information that is marked for sharing. This may prove more secure, but will likely limit usability. For example, a video chat application that only broadcasts the participants' faces would be a strange experience, and would inhibit natural forms of collaboration.

In the end, the kinds of security breaches that we are beginning to see at the interface of the physical and digital worlds are not going away and seem likely to get worse. That there are no obvious solutions may seem dispiriting, but this also gives mobile and security researchers the opportunity to work on hard problems and the potential to have tremendous positive impact. ■

[1] http://www.bbc.com/news/world-latin-america-26283966
[2] http://www.foxnews.com/us/2014/02/02/law-enforcement-agencies-ready-security-plans-in-preparation-for-super-bowl/
[3] http://www.zdnet.com/article/super-bowl-wi-fi-password-credentials-broadcast-in-pre-game-security-gaffe/
[4] http://www.theregister.co.uk/2014/06/25/brace_yourselves_brazil_dill_in_world_cup_wifi_spill/
[5] http://www.independent.co.uk/life-style/gadgets-and-tech/news/tv5monde-hack-staff-accidentally-show-passwords-in-report-about-huge-cyberattack-10168475.html
[6] http://techcrunch.com/2011/06/28/whiteboard-in-microsoft-video-hints-at-wp7s-future/