Reactive Security: Responding to Visual Stimuli from Wearable Cameras

Robert Templeman
School of Informatics and Computing
Indiana University
Bloomington, IN, USA
templeman@indiana.edu

Roberto Hoyle
School of Informatics and Computing
Indiana University
Bloomington, IN, USA
rjhoyle@indiana.edu

David Crandall
School of Informatics and Computing
Indiana University
Bloomington, IN, USA
djcran@indiana.edu

Apu Kapadia
School of Informatics and Computing
Indiana University
Bloomington, IN, USA
kapadia@indiana.edu

UbiComp '14, September 13 - 17 2014, Seattle, WA, USA
Copyright 2014 ACM 978-1-4503-3047-3/14/09 ... $15.00.
http://dx.doi.org/10.1145/2638728.2641708

Abstract

Consumer electronic devices like smartphones increasingly feature arrays of sensors that can 'see', 'hear', and 'feel' the environment around them. While these devices began with primitive capabilities, newer generations of electronics offer sophisticated sensing arrays that collect high-fidelity representations of the physical world. For example, wearable cameras are becoming more prevalent with new consumer lifelogging products including the Narrative Clip, Autographer, and Google Glass. These wearable cameras give computing devices a persistent sense of sight, raising important concerns about protecting people's privacy. At the same time, these devices also provide opportunities for enhancing security, by allowing trusted devices to observe and react to the physical environment surrounding the user and the device. We propose Attribute Based Access Control (ABAC) to mediate access to sensors and their data using attributes of the context and content of sensor information. Attributes extracted from sensor data could be used to trigger policy actions ranging from sharing or not sharing images, to invoking system changes in reaction to outside visual stimuli such as automatically shutting down network interfaces when in the presence of unknown people. While prior work has addressed some specific actions, like preventing potentially private images from being shared based on their location, in this paper we present and advocate for a more general working definition of ABAC that applies to sensors and sensor data. We also present use cases for how this reactive security approach may help protect the privacy and security of users.

Author Keywords

Lifelogging; wearable cameras; privacy

ACM Classification Keywords

K.4.2. [Social Issues]; K.4.1. [Public Policy Issues]: Privacy

Introduction

> *"One can now picture a future investigator in his laboratory. His hands are free, and he is not anchored. As he moves about and observes, he photographs and comments. Time is automatically recorded to tie the two records together. If he goes into the field, he may be connected by radio to his recorder. As he ponders over his notes in the evening, he again talks his comments into the record. His typed record, as well as his photographs, may both be in miniature, so that he projects them for examination."* (V. Bush, 1945 [7])

In 1945, Vannevar Bush, then Director of the U.S. Office of Scientific Research and Development, penned a prescient essay charging scientists to turn to peaceful pursuits after World War II. The advancement and accessibility of technology today reaches far beyond Dr. Bush's then-visionary ideas. Cameras and other sensors are commonly with us wherever we go and collect data about our daily lives, and their persistent access to the global network allows them to share this data instantly with the world. These devices witness the minutiae of our routines, and we often give them access to contexts and environments that we would not give to other people.

Smartphones have become the *de facto* standard sensor platform, and a recent Pew study shows that more than half of Americans now own and use one [33]. Current-generation smartphones can persistently estimate our geographical locations, 'see' the space and objects around us, 'listen' to ambient noise and conversations, and 'feel' our activity. This sensing capability makes many popular smartphone apps possible: for example, Instagram lets us capture and share snapshots of our lives, while Foursquare permits us to broadcast where we are. Meanwhile, new wearable cameras such as the Narrative Clip [28], Autographer [4], and Google Glass [16] that can serve as 'lifelogging' devices are on the horizon that will give such sensing functionality center stage. For example, the Narrative Clip records a geotagged picture every 30 seconds in perpetuity.

Lifelogging devices can record wholesale narratives of our activities, introducing a notion of 'perfect memory' [3] where more data is archived than we as humans are capable of remembering or processing. This transcription of our lives may contain data that we want to share with others, such as images that are shared on Facebook, exercise data shared on Strava, and our geographical locations shared on Foursquare. These devices may also offer therapeutic benefits: archived logs can help people with memory impairments (including Alzheimer's disease) remember important details [23], and may also help people cope with social anxiety disorder [30].

But while ubiquitous sensors enable many exciting applications, they also introduce obvious security and privacy challenges. Modern devices make it arguably too easy for users to collect and share images and other data. For example, when a user takes a picture with Google Glass, the photo is automatically uploaded to the cloud by default, and then Glass presents the user with a menu of sharing options. The user must then manually decide whether and how to share the image with others. This manual process typical of modern devices can be thought of as a form of *Discretionary Access Control* (DAC): the user's sharing decisions define an implicit access policy, with the responsibility of defining and enforcing it lying solely with the user (the data owner). But requiring users to manually review each image becomes intractable when images and other data are being constantly collected, so today's mobile applications are often granted permission to collect and transmit sensor data with impunity and without the user's explicit knowledge.

We claim that current approaches for managing personal sensor data are inadequate as we move into an era of first-person sensing and lifelogging. We envision using semantic content within an image to assist a user in enforcing privacy policies: attributes of the content and context surrounding the person can be estimated from the visual information seen by a wearable camera, and then these attributes can be used to apply user-defined policies that could react to changing environments. This 'reactive security' could range from applying an appropriate sharing policy to an image without requiring user intervention (e.g. *"Do not share any pictures of my children with the public"*), to modifying device behavior based on visual surroundings (e.g. locking a smartphone's screen whenever a stranger is seen in the vicinity), to configuring automatic actions whenever specific visual content is detected (e.g. subscribing to an Amber Alert feed to automatically notify authorities if a missing child is observed).

In the following, we first focus on the sharing of sensor data (with our primary interested being images), and then broaden our discussion to other applications of reactive security.

Sensing and Sharing Today

Most mobile devices today contain suites of sensors including cameras, microphones, accelerometers, gyroscopes, magnetometers, and GPS receivers, with the accuracy, sophistication, and quantity of these sensors steadily increasing with each new generation of device. These sensors are *user-owned resources* and reside at the boundary of the digital and physical worlds. For mobile devices, the digitization of physical world data is generally controlled through *permissions*. While details of permission systems vary across different platforms, the general concept is that users 'grant' applications permission to access sensor data, with most platforms offering only coarse-grained permissions and limited control to users [9, 21]. For example, in Android, the user is presented with a list of requested permissions when a new application is about to be installed. After the user grants permission, the newly-installed application has unfettered access to the approved resources. Revocation of permissions is only possible by uninstalling the application (although future versions of Android may offer finer-grained control [36]).

This and other well-documented shortcomings in existing permission systems have allowed powerful sensor-based attacks that actually give remote hackers 'virtual access' to the physical environment surrounding a device [32,35]. One solution is to use fine-grained permissions consistent with the *principle of least privilege* [9]. Unfortunately, this requires verbose policies or a high degree of user workload to manage access to user-owned resources.

Instead of focusing on whether and when certain sensors can be accessed, we consider how to control access to sensor data *after* it is collected, and how this sensor data can trigger broader actions. Here we are agnostic about *how* sensor data was collected. For example, consider a set of images that has already been captured. Our problem is now cast as determining how to provide access to the images (i.e., share them, or allow other system actions to be taken based on their contents). Errors in judgment or even simple mistakes can result in damaging consequences (e.g., by sending an incriminating or embarrassing image to a colleague, or by not realizing that one is in an unsafe environment and leaving a device in a promiscuous state of accepting connections).

Control of sensors and sensor data is challenging enough with smartphones but becomes even more difficult with wearable devices. These devices have the ability to collect large volumes of information, often in an opportunistic manner. Current solutions for managing sensors and sensor data are not sufficient, and the magnitude of this problem will only grow as these devices become more popular and more powerful.

Attribute-Based Access Control for Sensors and Sensor Data

There is a fundamental tension between protecting the confidentiality of sensor data while allowing data to be collected and shared by applications. In fact, one could argue that the popularity of mobile devices is largely driven by their ability to share personal information; YouTube, Vine, Instagram, and Facebook are among the most popular apps on both iOS and Android devices and are all based on collecting and sharing data. Existing access control frameworks and implementations of mobile file systems are sufficient when data is collected and used solely by the owner of the device. Virtualization, sandboxing, and discretionary access control preserve confidentiality across applications. Since mobile devices typically do not support multiple users, external user access to data only occurs when it is transferred out of the device. At that point, users implement informal discretionary access control via sharing actions (e.g., by emailing a photo or sharing a geospatial location).

Consider Bob who knowingly has compromising and sensitive photos of himself in the Gallery on his phone where they reside alongside pictures that he takes as part of his job. Bob understands the semantic differences amongst these photos and only shares them with qualified individuals or services. For example, Bob would never knowingly send a compromising photo to a customer (although like anyone, he may make mistakes or irrational decisions that violate his own internal access control ‘policy’).

Attribute-Based Access Control (ABAC) mechanisms that act on the ‘sensitivity’ of sensor content could balance the desire to share data with the need to control access to private personal data. Given an image, an ideal system would evaluate its semantic meaning and permit sharing only with the subset of contacts or applications for which access is appropriate. A key technical challenge, however, is how to infer the content of an image automatically, especially since subtle differences between images can result in dramatic differences in privacy implications (e.g., different cultural norms between topless men and women). While complete semantic understanding of images and other sensor data using computational methods may be a very long-term goal, we believe that even partial and noisy semantic understanding may be sufficient to assist access control systems and users in managing their data. We posit that systems based on ABAC may satisfy users' desires to use social networking, lifelogging, and other applications that depend on collecting and sharing data, while maintaining control over their own privacy. Further, we claim that ABAC systems can offer functions to improve security and privacy postures of the platforms that they reside on.

Interpretation of ABAC

NIST defines ABAC as follows: "Attribute Based Access Control (ABAC): *An access control method where subject requests to perform operations on objects are granted or denied based on assigned attributes of the subject, assigned attributes of the object, environmental conditions, and a set of policies that are specified in terms of those attributes and conditions*" [20]. We now offer our own definition of ABAC in the context of sensors and sensor data. We begin with the deconstruction and interpretation of the NIST definition of ABAC: **Attributes** are characteristics of the object sensor content or its metadata. We also include, but do not limit to, the attributes related to subjects and environmental conditions. **Subjects** are users and a superset of principals that include other people, applications, and other potential recipients of object sensor data (e.g., network-reachable services). (We assume that the user is granted full access to their objects; access control decisions pertain to the remaining 'external' subjects.) **Objects** are sensor resources, discrete elements of sensor data (e.g., a geospatial location, an image, or an accelerometer reading), interfaces (e.g., Wifi, USB, Ethernet, etc), or other system resources. We can also aggregate atomic objects (e.g., a stream of photos, the geospatial location where they were taken, and an audio clip) to create more complex objects. **Operations** are functions to be performed on objects at the request of a subject. Without loss of generality, we restrict our model to *read* operations of object sensor data and *collect* operations of sensors. **Policies** are rules that use attributes to allow or disallow operations. An obvious application is the mediation of access requests to objects by subjects given the attributes of the subject, object, and environment. We consider policies that are both specific to individual users and those that are generalized to all users. Likewise, we include both policies that are manually defined and those that are generated using machine learning, statistical, or heuristic approaches. **Environmental conditions** include the object and situational contexts of the access request.

Features and Requirements

An access control framework for mobile sensor devices should provide five key features: **Usability.** Any proposed implementation of a system must be usable, requiring a modest amount of user effort to use and maintain the framework by inexperienced users, while running efficiently on mobile hardware. **Attribute extraction.** A system needs to define a set of attributes that it can recognize with reasonable accuracy, from the infinite number of possibilities (e.g. *presence of computer monitor in image, image taken outdoors, daughter present in image, am I shaven in image, etc.*). **Sensor access control.** Access control decisions may mediate the *collection* of data (e.g., because it is in a locker room the phone is not permitted to take a picture). **Sensor data access control.** Access control decisions may mediate operations performed on objects *after* collection (e.g., an existing photo is not tweeted publicly because it contains confidential business information). **Platform resource access control.** Access control decisions may control access to any number of system resources (e.g., external interfaces) and can include policy rules that adapt the system configuration to the environment (context).

Attribute extraction

A key challenge in implementing ABAC for wearable cameras is how to extract semantic attributes from images. Despite over 50 years of study, computer vision remains a very difficult problem, with accuracy of state-of-the-art recognition techniques paling in comparison to the human visual system. Practical ABAC systems will thus require identifying attributes that can be extracted quickly and accurately, while still providing sufficient discrimination for access control decisions. Our PlaceAvoider system [34] is one initial attempt at this trade-off: that approach is ABAC-based, but relies on a single attribute (image location) that could be recognized accurately through a combination of GPS and automatic image recognition techniques.

We envision two main types of attributes. Low-level attributes could be based on simple processing of raw sensor data, including time of day, day of the week, physical GPS location, amount of ambient noise, degree of device motion, ambient light level, etc. These attributes can be easily inferred from modern sensors, and in some cases might contribute to useful policies (e.g. restricting sharing of images taken at night inside a home). Higher-level attributes reflecting content and context of an image are arguably more important because of their greater expressive power, but of course require much more advanced automatic analysis techniques. Here we envision two levels of this analysis: mid-level attributes that focus on syntactic properties of images, and high-level attributes that relate to more complex, semantic properties. For example, mid-level attributes might include the number of faces in an image, the genders and estimated ages of the people, whether the people are sitting or standing up, whether the photo is taken inside or outside, the types of objects in the scene, etc. Higher-level attributes could include more complex semantic properties of the scene, like whether the photo is taken in a private or public space, what people in the scene are doing, whether the environment appears secure or not, etc.

Although computer vision techniques are far from perfect, we believe some higher-level attributes could be detected with sufficient accuracy to enable useful ABAC policies with today's technology, and image-based ABAC may help define new research challenges for computer vision research. For example, face detection and face recognition [1] technology have become reliable enough to find widespread use in consumer devices like digital cameras, with accurate and easy-to-use APIs commercially available [11]. The accuracy of other tasks like recognizing places [27], scene types [37], events [26], activities [2], objects [15], and human poses [10] and interactions [24] are less predictable, but all are very active research areas that are making rapid progress. Attributes involving these problems could be chosen judiciously based on the requirements of the application and available technology. Much of this work was designed for consumer images; studying these problems under the unique conditions of egocentric images and video is also an active research area [12, 13, 25, 29].

ABAC Applications & Research Directions

We now explore the use of ABAC with wearable cameras and reactive security, and highlight promising research directions. An immediate need for this technology is in the area of managing sensor data (e.g., sharing images), from either lifelogging applications or deliberately-captured data which a user wants to share automatically. We also consider generalizing this technology to a broader field by allowing system events to be triggered by visual stimuli from a wearable camera.

Controlling the Sharing of Images

A policy exchange mechanism has been proposed by Roesner et. al [31] that would create *passports*, which are certificates that bind policies so that they can be exchanged between users. Our interest is in the attributes that make these policies possible, as determining which attributes are necessary is an open problem. Our prior PlaceAvoider work successfully uses the *scene location* attribute of an image to inform sharing decisions [34], so that (for example) users can choose to share photos taken in the kitchen but not in the bathroom. We used a computer vision approach that recognizes the visual appearance of different rooms, in order to mitigate limitations of location services (since GPS does not work well indoors) and allow users to construct policies of blacklisted spaces in usable ways. However, this work is limited in that it supports a single attribute, and the sensitivity of a photo depends on many more semantic dimensions than location alone. But besides location, which attributes would be most useful?

We conducted a user study in the context of lifelogging to investigate which scene attributes are most needed in determining the sensitivity of an image [19]. We asked participants to wear a custom lifelogging system for Android smartphones (around the neck in a lanyard) that recorded images and other sensor data throughout the day. Afterwards, we showed them the images and asked if and how they would share them, and the reasons behind their choices. Our study tracked 36 participants over the course of a week. We found that sharing behavior could not be determined by any single image factor, but instead depended on a combination of features including time, image content, and location. Some features were more likely to trigger privacy constraints by participants. Images that featured computers, for example, were more likely to be kept private, while some participants did not share images that contained specific people out of concern for the privacy of those people. Participants were also less likely to share photos taken in their home versus those taken elsewhere. More detailed results are presented in [19].

A future challenge is how to detect these features automatically via computer vision. An initial step could be an algorithm that is able to detect the presence of a computer screen. Facial detection algorithms could be used to trigger policies to automatically share or not share an image. A child-detection algorithm could

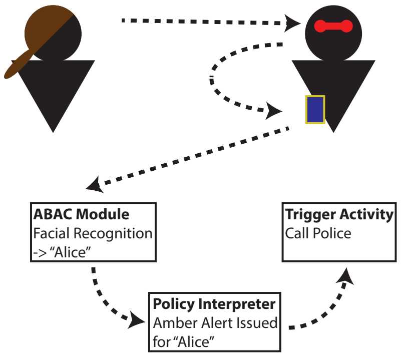

Figure 1: ABAC architecture overview showing how seeing Alice via Bob's wearable camera can trigger Bob's smartphone to take an action.

detect whether a person in a picture was an adult or a child, enabling a parent to protect any images that were captured containing one of their children.

Reactive Security using Sensors

More generally, image analysis of wearable camera data could be used not only to protect images taken by the camera, but also to augment the security of other devices. As illustrated in Figure 1, we envision a scenario in which a first-person wearable camera is paired with one or more systems having resources that a user wants to be protected. The system runs an ABAC module that reads in data from the camera and other sensors, and then triggers actions in the system via messages or hooks. Applications that own protected resources would establish hooks that would perform operations on objects when indicated by the ABAC system, based on policies. We envision actions that are triggered both within the operating system as well as at the application layer. For the latter case, applications may subscribe to events based on ABAC policies. As with the more basic question of whether an app should be allowed a certain set of permissions, a major research challenge is when to grant apps access to ABAC based triggers. For example, an Amber Alert app may not be granted access to the raw video feed, and instead be triggered when a certain attribute (e.g., a particular child's face) is detected. Research is needed to make such policies understandable to users, who can then make informed decisions about whether these triggers are reasonable. While researchers have begun to address this challenge in the context of permissions requested by an app [5, 6, 14], extending such research to application triggers is needed.

We anticipate several promising applications of reactive security. Researchers have proposed approaches for 'continuous authentication,' where anomalous activity of the user can signal an intruder who is impersonating the owner of the account [8, 17, 18, 22]. We envision a camera-based approach that can detect anomalous users based on a visual inspection of their surroundings. For example, a user's camera could enable a simpler phone unlocking system (like a short passcode or a simple swipe) when it detects a user's home or office, but require a more complex passcode when in unfamiliar locations. The camera could also be used to detect unfamiliar faces to signal anomalous situations. Even when the authorized user is in possession of their device, a paired camera could customize the device's behavior based on the situation. For example, a smartphone could silence its ringer when the user is in certain social contexts, like a business meeting, but could enable additional interaction (like automatically reading incoming text messages outloud) when the user is alone. It could also 'look' for situations that pose a threat to the device's security or the user's privacy, and react accordingly. For example, detecting that an employee is visiting a competitor's facility could trigger more restrictive firewall policies and more aggressive encryption of stored data and communications. A crowd detection mechanism could trigger blocking access to devices on a user's body-area network when the user is moving quickly past a large number of people, while relaxing the restrictions when a user is in an environment with fewer people where sharing of files and information may be more common.

Finally, aggregating data from many first-person devices could create new opportunities for crowdsourcing, that could in turn enhance society's security and safety. Law enforcement could crowdsource efforts to locate missing children to users who opt into an Amber Alert service, by pushing images of missing children to the devices, which could analyze surrounding imagery for missing children. Of course, a key challenge for crowdsourced applications is how this detection can be performed while maintaining the privacy of camera owners, as imagery will be shared with law enforcement and users' tolerance for false positives will be low. A countervailing crowdsourcing application could be to monitor authorities to detect and report abuse: people could opt into an effort to monitor law enforcement (when legally allowable) and share imagery where uniformed officers are detected.

Conclusion

We expect wearable cameras to become commonplace in the near future, to the extent that they may soon log our *entire lives* with photos and other sensor data. While this will create exciting applications and opportunities to record and share aspects of our lives, it will also introduce substantial privacy concerns. Visual imagery is extremely rich in context, such that leaking image data could be particularly damaging. At the same time, the possibility of extracting this rich semantic context automatically opens up exciting opportunities for trusted security applications that respond to visual stimuli. We propose the use of attribute-based access control implementations for mobile sensors and sensor data, and discuss various applications and challenges related to extracting suitable attributes from visual imagery. Systems built on this architecture open up the possibility of reactive security actions based on visual stimuli from a wearable camera, freeing users from having to actively manage their privacy and instead triggering security actions as needed. As we enter an era of pervasive cameras, we hope to spur further research in the area of security and privacy that leverages visual sensing.

Acknowledgments

This material is based upon work supported by the National Science Foundation under grants CNS-1016603, CNS-1252697, and IIS-1253549. This work was also partially funded by the Office of the Vice Provost for Research at Indiana University Bloomington through the Faculty Research Support Program.

References

[1] Abate, A., Nappi, M., Riccio, D., and Sabatino, G. 2D and 3D face recognition: A survey. *Pattern Recognition Letters 28*, 14 (Oct. 2007).

[2] Aggarwal, J., and Ryoo, M. Human activity analysis: A review. *ACM CSUR 43*, 16 (2011).

[3] Allen, A. Dredging up the past: Lifelogging, memory, and surveillance. *Univ. of Chicago Law Review* (2008).

[4] *Autographer*. http://autographer.com.

[5] Bauer, L., Cranor, L., Reeder, R. W., Reiter, M. K., and Vaniea, K. A User Study of Policy Creation in a Flexible Access-control System. In *CHI* (2008).

[6] Benton, K., Camp, L. J., and Garg, V. Studying the effectiveness of android application permissions requests. In *PERCOM* (2013).

[7] Bush, V. As we may think. *Atlantic Monthly 176* (1945).

[8] Clarke, N., and Furnell, S. Authenticating mobile phone users using keystroke analysis. *IJIS 6*, 1 (2007), 1–14.

[9] Conti, M., Nguyen, V., and Crispo, B. CRePE: Context-related policy enforcement for Android. In *ICISC* (2011), 331–345.

[10] Duan, K., Batra, D., and Crandall, D. A multi-layer composite model for human pose estimation. In *British Machine Vision Conference* (2012).

[11] *Face++*. http://www.faceplusplus.com/.

[12] Fathi, A., Farhadi, A., and Rehg, J. Understanding egocentric activities. In *CVPR* (2011).

[13] Fathi, A., Ren, X., and Rehg, J. Learning to recognize objects in egocentric activities. In *CVPR* (2011).

[14] Felt, A. P., Chin, E., Hanna, S., Song, D., and Wagner, D. Android permissions demystified. In *CCS* (2011).

[15] Felzenszwalb, P., Girshick, R., McAllester, D., and Ramanan, D. Object detection with discriminatively trained part-based models. *IEEE T. PAMI 32*, 9 (2010).

[16] *Google Glass*. http://glass.google.com.

[17] Gunetti, D., and Picardi, C. Keystroke analysis of free text. *ACM TISSEC 8*, 3 (2005), 312–347.

[18] Gupta, A., Miettinen, M., Asokan, N., and Nagy, M. Intuitive Security Policy Configuration in Mobile Devices Using Context Profiling. In *SocialCom/PASSAT* (2012).

[19] Hoyle, R., Templeman, R., Armes, S., Anthony, D., Crandall, D., and Kapadia, A. Privacy behaviors of lifeloggers using wearable cameras. In *UbiComp* (2014).

[20] Hu, V., Ferraiolo, D., Kuhn, R., Schnitzer, A., Sandlin, K., Miller, R., and Scarfone, K. *Guide to Attribute Based Access Control Definition and Considerations*. NIST, 2013.

[21] Jeon, J., Micinski, K., Vaughan, J., Fogel, A., Reddy, N., Foster, J., and Millstein, T. Dr. Android and Mr. Hide: Fine-grained permissions in Android applications. In *ACM Workshop on Security & Privacy in Smartphones & Mobile Devices* (2012).

[22] Killourhy, K. S., and Maxion, R. A. Comparing anomaly-detection algorithms for keystroke dynamics. In *DSN* (2009), 125–134.

[23] Lee, M., and Dey, A. Lifelogging memory appliance for people with episodic memory impairment. In *UbiComp* (2008).

[24] Lee, S., Bambach, S., Crandall, D., Franchak, J., and Yu, C. This hand is my hand: A probabilistic approach to hand disambiguation in egocentric video. In *CVPR Egovision* (2014).

[25] Lee, Y. J., Ghosh, J., and Grauman, K. Discovering important people and objects for egocentric video summarization. In *CVPR* (2012).

[26] Li, L.-J., and Fei-Fei, L. What, where and who? classifying events by scene and object recognition. In *ICCV* (2007).

[27] Li, Y., Crandall, D., and Huttenlocher, D. Landmark classification in large-scale image collections. In *ICCV* (2009).

[28] *Narrative Clip*. http://getnarrative.com.

[29] Pirsiavash, H., and Ramanan, D. Detecting activities of daily living in first-person camera views. In *CVPR* (2012).

[30] Rennert, K., and Karapanos, E. Faceit: Supporting reflection upon social anxiety events with lifelogging. In *CHI* (2013).

[31] Roesner, F., Molnar, D., Moshchuk, A., Kohno, T., and Wang, H. World-driven access control for continuous sensing. Tech. rep., Microsoft Research, 2014.

[32] Schlegel, R., Zhang, K., Zhou, X., Mehool, I., Kapadia, A., and Wang, X. Soundcomber: A stealthy and context aware sound trojan for smartphones. In *NDSS* (2011).

[33] Smith, A. Nearly half of American adults are smartphone owners. Tech. rep., Pew Research, 2012.

[34] Templeman, R., Korayem, M., Crandall, D., and Kapadia, A. PlaceAvoider: Steering first-person cameras away from sensitive spaces. In *NDSS* (2014).

[35] Templeman, R., Rahman, Z., Crandall, D., and Kapadia, A. PlaceRaider: Virtual theft in physical spaces with smartphones. In *NDSS* (2013).

[36] Tung, L. Google removes ‘awesome’ but unintended privacy controls in Android 4.4.2. *ZDNet* (Dec 16, 2013).

[37] Xiao, J., Hays, J., Ehinger, K., Oliva, A., and Torralba, A. Sun database: Large-scale scene recognition from abbey to zoo. In *CVPR* (2010).