

REPORT DOCUMENTATION PAGE			Form Approved OMB NO. 0704-0188		
The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA, 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.					
1. REPORT DATE (DD-MM-YYYY) 17-07-2017		2. REPORT TYPE Final Report		3. DATES COVERED (From - To) 21-Jul-2014 - 20-Jul-2015	
4. TITLE AND SUBTITLE Final Report: Danger; Understanding Privacy Risks of Ubiquitous Personal Augmented Reality Head-mounted Displays			5a. CONTRACT NUMBER W911NF-14-1-0438		
			5b. GRANT NUMBER		
			5c. PROGRAM ELEMENT NUMBER 611103		
6. AUTHORS Jan-Michael Frahm			5d. PROJECT NUMBER		
			5e. TASK NUMBER		
			5f. WORK UNIT NUMBER		
7. PERFORMING ORGANIZATION NAMES AND ADDRESSES University of North Carolina - Chapel Hill 104 Airport Drive, CB 1350 Suite 2200 Chapel Hill, NC 27599 -1350			8. PERFORMING ORGANIZATION REPORT NUMBER		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS (ES) U.S. Army Research Office P.O. Box 12211 Research Triangle Park, NC 27709-2211			10. SPONSOR/MONITOR'S ACRONYM(S) ARO		
			11. SPONSOR/MONITOR'S REPORT NUMBER(S) 65102-CS-RIP.1		
12. DISTRIBUTION AVAILABILITY STATEMENT Approved for Public Release; Distribution Unlimited					
13. SUPPLEMENTARY NOTES The views, opinions and/or findings contained in this report are those of the author(s) and should not be construed as an official Department of the Army position, policy or decision, unless so designated by other documentation.					
14. ABSTRACT Under this effort, we investigated the space of realistic information leakage by personal mobile platforms such as mobile phones, AR devices, and their purposefully and accidentally obtained imagery. The hardware supported efforts in two research areas, namely in leveraging compromising observations of those devices to infer information about the user's environment and her position. Second, we investigated the leakage potential of the emanations of the devices when in use.					
15. SUBJECT TERMS mobile phones, AR devices					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT UU	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON Jan-Michael Frahm
a. REPORT UU	b. ABSTRACT UU	c. THIS PAGE UU			19b. TELEPHONE NUMBER 919-590-6003

Report Title

Final Report: Danger; Understanding Privacy Risks of Ubiquitous Personal Augmented Reality Head-mounted Displays

ABSTRACT

Under this effort, we investigated the space of realistic information leakage by personal mobile platforms such as mobile phones, AR devices, and their purposefully and accidentally obtained imagery. The hardware supported efforts in two research areas, namely in leveraging compromising observations of those devices to infer information about the user's environment and her position. Second, we investigated the leakage potential of the emanations of the devices when in use.

Enter List of papers submitted or published that acknowledge ARO support from the start of the project to the date of this printing. List the papers, including journal references, in the following categories:

(a) Papers published in peer-reviewed journals (N/A for none)

<u>Received</u>	<u>Paper</u>
-----------------	--------------

TOTAL:

Number of Papers published in peer-reviewed journals:

(b) Papers published in non-peer-reviewed journals (N/A for none)

<u>Received</u>	<u>Paper</u>
-----------------	--------------

TOTAL:

Number of Papers published in non peer-reviewed journals:

(c) Presentations

Number of Presentations: 0.00

Non Peer-Reviewed Conference Proceeding publications (other than abstracts):

Received Paper

TOTAL:

Number of Non Peer-Reviewed Conference Proceeding publications (other than abstracts):

Peer-Reviewed Conference Proceeding publications (other than abstracts):

Received Paper

TOTAL:

Number of Peer-Reviewed Conference Proceeding publications (other than abstracts):

(d) Manuscripts

Received Paper

TOTAL:

Number of Manuscripts:

Books

Received Book

TOTAL:

TOTAL:

Patents Submitted

Patents Awarded

Awards

Graduate Students

<u>NAME</u>	<u>PERCENT SUPPORTED</u>
FTE Equivalent:	
Total Number:	

Names of Post Doctorates

<u>NAME</u>	<u>PERCENT SUPPORTED</u>
FTE Equivalent:	
Total Number:	

Names of Faculty Supported

<u>NAME</u>	<u>PERCENT SUPPORTED</u>
FTE Equivalent:	
Total Number:	

Names of Under Graduate students supported

<u>NAME</u>	<u>PERCENT SUPPORTED</u>
FTE Equivalent:	
Total Number:	

Student Metrics

This section only applies to graduating undergraduates supported by this agreement in this reporting period

The number of undergraduates funded by this agreement who graduated during this period:

The number of undergraduates funded by this agreement who graduated during this period with a degree in science, mathematics, engineering, or technology fields:.....

The number of undergraduates funded by your agreement who graduated during this period and will continue to pursue a graduate or Ph.D. degree in science, mathematics, engineering, or technology fields:.....

Number of graduating undergraduates who achieved a 3.5 GPA to 4.0 (4.0 max scale):.....

Number of graduating undergraduates funded by a DoD funded Center of Excellence grant for Education, Research and Engineering:.....

The number of undergraduates funded by your agreement who graduated during this period and intend to work for the Department of Defense

The number of undergraduates funded by your agreement who graduated during this period and will receive scholarships or fellowships for further studies in science, mathematics, engineering or technology fields:

Names of Personnel receiving masters degrees

NAME

Total Number:

Names of personnel receiving PHDs

NAME

Total Number:

Names of other research staff

NAME

PERCENT SUPPORTED

FTE Equivalent:

Total Number:

Sub Contractors (DD882)

Inventions (DD882)

Scientific Progress

Hardware was purchased.

Technology Transfer

Final Report for “Danger; Understanding Privacy Risks of Ubiquitous Personal Augmented Reality Head-mounted Displays”

Under this effort, we investigated the space of realistic information leakage by personal mobile platforms such as mobile phones, AR devices, and their purposefully and accidentally obtained imagery. The hardware supported efforts in two research areas, namely in leveraging compromising observations of those devices to infer information about the user’s environment and her position. Second, we investigated the leakage potential of the emanations of the devices when in use.

Compromising Observations

The hardware obtained under this grant provided the enabling compute and experimental platforms for several research efforts in the area of exploitation of the data captured on mobile AR platforms [1, 5, 6, 7, 9].

Mobile AR platforms and mobile phones have developed and will further developed into our daily capture platforms and will capture more and more of our lives in the near future. This is especially true when used in ubiquitous AR platforms. In Heinly et al. [1] we proposed a novel, large-scale, structure-from-motion framework that advances the state of the art in data scalability from city-scale modeling (millions of images) to world-scale modeling (several tens of millions of images) using just a single computer. This platform shows that realistically an attacker can join the information across the streams of large numbers of different users to obtain a model of their environment and the user’s position at the time of the capture. The main enabling technology in this effort is the use of a streaming-based framework for connected component discovery. Moreover, the system employs an adaptive, online, iconic image clustering approach based on an augmented bag-of-words representation, in order to balance the goals of registration, comprehensiveness, and data compactness. We demonstrated our method on a world-scale dataset the publicly available 100 million image crowd- sourced photo collection containing images geographically distributed throughout the entire world. Results illustrate that our streaming-based approach does not compromise model completeness, but achieves unprecedented levels of efficiency and scalability. In summary, this shows that even large-scale datasets as soon will be created by personal augmented reality devices can realistically be efficiently analyzed to infer information about the user and his peers capturing the data.

To drive efficiency of large-scale reconstruction systems such as in the system of Heinly et al. [1] to even larger scales of user data we investigated methods into even faster image overlap detection (geometric verification), which is the main bit in correlating data across user’s and across different images of the same user. In Schönberger et al. [5,7] we targeted a more efficient method for geometric verification. Typically, the geometric verification recovers the epipolar geometry of two views for a moving camera by estimating a fundamental or essential matrix. The essential matrix describes the relative geometry for two views up to an unknown scale. Then two-view triangulation or multi-model estimation approaches can reveal the relative geometric

configuration of two views, e.g., small or large baseline and forward or sideward motion. Information about the relative configuration is essential for many problems such as recovering the environment model and the user's position in the environment. However, essential matrix estimation and assessment of the relative geometric configuration are computationally expensive and hence a bottleneck in any large-scale exploit of user photos. To overcome this bottleneck, we proposed a learning-based approach for efficient two-view geometry classification, leveraging the by-products of feature matching. Our approach can predict whether two views have scene overlap and for overlapping views it can assess the relative geometric configuration. Our experiments show that through combining the above methods [1,5] efficient user localization and environment extraction is possible.

We also investigated the amount of information that can be extracted from the imagery of these devices. Specifically, in Schönberger et al. [6] we researched the level of detail information that is retrievable from the images. High resolution detail is often of interest in privacy attacks, as it for example allows the attacker to read text in the scene and notice even small objects of interest. While the above methods [1, 5, 7] boost scalability, they traditionally also limit the amount of detail that the large-scale reconstruction systems are able to produce. In Schönberger et al. [6] we introduced a joint reconstruction and retrieval system that maintains the scalability of large-scale Structure-from-Motion systems [1, 5, 7], while also recovering the often lost ability of reconstructing fine details of the scene through efficiently combining image recognition and reconstruction strategies. We demonstrated the method on a large-scale dataset of 7.4 million images downloaded from the Internet. This effort proves that the obtained reconstructions allow recovery even of small details of the captured environment.

AR glasses as an Ad Hoc Surveillance Network

Another existing limitation in the reconstruction of the environment is the fact that dynamic moving objects such as people and cars could not be reconstructed. We analyzed if this automatically provides privacy for the captured users but found this not to be the case since we were able to research a method of correctly localizing imaged users or other moving objects. Specifically, in Price et al. [9] we proposed a method to bring 3D reconstructions to life by populating them with transient objects as observed in the real world. We focus on detecting people in individual images and accurately placing them into an existing 3D model of the environment. As part of this placement, our method also estimates the metric scale of the scene from object semantics, namely the size distribution of the population. Alternatively, if the observed users are known we could use their specific body measures. Moreover, our method models a smooth approximation of the ground surface, which is typically difficult to reconstruct in crowd-sourced image sets. We have tested our approach on a large number of unordered Internet photo collections and demonstrate realistic visualizations of such previously unobtainable scene elements. This effort demonstrates that AR glasses and their observations can be effectively used as a distributed surveillance system by exploiting their observations.

Privacy Threats through Compromising Emanations

The purchased hardware was critical in enabling our research in leveraging information of the user and content on the device available to an attacker [2,3,8]. Privacy on mobile devices has multiple facets on one hand the privacy of the user using the device and on the other hand the privacy of the people in the user's environment. We researched both aspect.

First, we analyzed the private information of the AR device's user that could be leaked. One interesting source of information is the recovery of the user's gaze direction and hence the ability to recover her interest objects in the environment. In Perra et al. [2] we introduced a continuous, locally optimal calibration scheme for use with head-worn devices. It removes the need of existing calibration schemes to solve for a globally optimal model of the eye-device transformation by performing calibration on a per-user or once-per-use basis. Removing, this requirement is critical in enabling the powerful threat of tracking the users gaze from exposed or intercepted data not intended for tracking. Our calibration scheme allows an attacker to calculate a locally optimal eye-device transformation on demand by computing an optimal model from a local window of previous frames. By leveraging naturally occurring interest regions within the user's environment, our system can calibrate itself without the user's active participation. Experimental results demonstrate that our proposed calibration scheme outperforms the existing state of the art systems while being significantly less restrictive to the user and the environment.

Another aspect of leaked information from head-sets, mobile devices, and even static screens is the recovery of information displayed on the user's screen. In Xu et al. [3] we proved that video content played on a device can efficiently be recovered even when not directly seen. In Price et al. [8], we extended this method to show more robustness against a number of transformations on the video and in the observed video screen, i.e. perspective foreshortening of the video.

Another aspect of mobile device privacy is the more and more popular biometrics based user authentication. We evaluated the safety of protecting user accounts and information through facial identification. In Xu et al. [4] we showed that, by leveraging a handful of pictures of the target user taken from social media, we are able to create realistic, textured, 3D facial models that undermine the security of widely used face authentication solutions. Our framework makes use of virtual reality (VR) systems, incorporating along the way the ability to perform animations (e.g., raising an eyebrow or smiling) of the facial model, in order to trick liveness detectors into believing that the 3D model is a real human face. The synthetic face of the user is displayed on the screen of the VR device, and as the device rotates and translates in the real world, the 3D face moves accordingly. To an observing face authentication system, the depth and motion cues of the display match what would be expected for a human face. Our analysis lets us argue that such VR-based spoofing attacks constitute a fundamentally new class of attacks that point to a serious weakness in camera-based authentication systems: Unless they incorporate other sources of verifiable data, systems relying on color image data and camera motion are prone to attacks via virtual realism. To demonstrate the practical nature of this threat, we conducted thorough experiments using an end-to-end implementation of our approach and show how it

undermines the security of several face authentication solutions that include both motion-based and liveness detectors. Our innovative research showing this fundamental flaw in the security of biometrics for user authentication on mobile devices was enabled by the hardware purchased under this grant [4].

Bibliography

[1] Jared Heinly, Johannes L. Schönberger, Enrique Dunn, Jan-Michael Frahm, “Reconstructing the World* in Six Days *(As Captured by the Yahoo 100 Million Image Dataset)”, IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2015

[2] David Perra, Rohit Kumar Gupta, Jan-Michael Frahm, “Adaptive Eye-Camera Calibration for Head-Worn Devices”, IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2015

[3] Yi Xu, Fabian Monroe, Jan-Michael Frahm, “Watching the Watchers: Inferring TV Content From Outdoor Light Effusions”, 21st ACM Conference on Computer and Communications Security, 2015

[4] Yi Xu, True Price, Jan-Michael Frahm, Fabian Monroe, “Virtual U: Defeating Face Liveness Detection by Building Virtual Models From Your Public Photos”, USENIX Security, 2015

[5] Johannes L. Schönberger, Alexander C. Berg, Jan-Michael Frahm, “Efficient Two-View Geometry Classification”, German Conference on Pattern Recognition (GCPR), 2015, (best paper honorable mention)

[6] Johannes L Schönberger, Filip Radenovi, Ondrej Chum, Jan-Michael Frahm, “From Single Image Query to Detailed 3D Reconstruction”, IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2015

[7] Johannes L Schönberger, Alexander C Berg, Jan-Michael Frahm, “PAIGE: PAirwise Image Geometry Encoding for Improved Efficiency in Structure-from-Motion”, IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2015

[8] Yi Xu, True Price, Fabian Monroe, Jan-Michael Frahm. “Caught Red-Handed: Toward Practical Video-based Subsequences Matching in the Presence of Real-World Transformations”, CV-COPS workshop CVPR 2017

[9] True Price, Johannes L Schönberger, Marc Pollefeys, Jan-Michael Frahm, “Crowding Out: Augmenting 3D Reconstructions using Human Detections”, under submission, 2017