Privacy Preserving Optics for Miniature Vision Sensors

Francesco Pittaluga[1], Sanjeev J. Koppal[2]
[1,2]Department of Electrical and Computer Engineering, University of Florida.

We present a novel framework, which enables "pre-capture" privacy, for miniature vision sensors. Most privacy preserving systems for computer vision, process images after capture. There exists a moment of vulnerability in such systems, *after* capture, when privacy has not yet been enforced. Our privacy sensors filter the incident light-field *before* image capture, while light passes through the sensor optics, so sensitive information is never measured by the sensor. Within this framework, we introduce, to our knowledge, the first ever sensor that enables pre-capture k-anonymity and multiple sensors that achieve pre-capture privacy through optical defocus. We also show theory for miniaturizing the proposed designs, including a novel "optical knapsack" solution for finding a field-of-view-optimal arrangement of optical elements. Our privacy preserving sensors enable applications such as accurate depth sensing, full-body motion tracking, multiple people tracking and low-power blob detection.

Related Work. *Ad hoc* privacy preserving algorithms for video data, such as naive blurring, trade off data utility for privacy protection, as they rely on heavy distortion of the data, to thwart restoration attacks [3]. In contrast, formal methods, such as those based on k-anonymity [2], maintain a high degree of data utility and provide provable privacy guarantees. One-dimensional sensors also maintain some data utility while preserving privacy [1]. In this paper, we show applications, for time-of-flight and thermal sensors, where heavy distortion, via defocus blur, maintains both data utility and privacy, and an optical implementation of k-anonymity. Our designs offer significant improvement over one-dimensional privacy sensors, in terms of data utility, as they capture appropriately modulated two-dimensional sensor readings.

1 Optical Knapsack

Our optical knapsack algorithm takes into account angular coverage by first discretizing the field-of-view into β angular regions, each with a solid angle of $\frac{\pi}{\beta}$. We define an array $K(n,\beta)$, where $K(i,b) = 1$ if that optical element covers the angular regions b in its field-of-view, and is zero everywhere else. We also define the array M to be three-dimensional of size $n \times A \times \beta$. As before, each entry of $M(i,a,0)$ contains the maximum field of view that can be obtained with the first i elements with a sensor of size a and $M(n,A,0)$ contains the solution to the knapsack problem. Entries $M(i,a,1)$ through $M(i,a,\beta)$ are binary, and contain a 1 if that angular region is covered by the elements corresponding to the maximum field-of-view $M(i,a,0)$ and a zero otherwise. The array M is initialized as, $M[i,a,b] = 0,\;$ if $\;0 \leq a \leq A,\;\; 0 \leq i \leq n\;$ and $\;0 \leq b \leq \beta$ and is recursively updated as

$$\begin{array}{ll}
\textbf{If } a < 0 & M[i,a,0] = -\infty \\
\textbf{For any other } a, \textbf{ for any } i & \\
\quad \textbf{If} & \\
\quad M[i-1,a,0] < & \left\{\begin{array}{l} M[i,a,0] = \\ f_i + M[i-1,a-a_i,0] \end{array}\right. \\
\quad f_i + M[i-1,a-a_i,0] & \\
\quad \textbf{and} & \\
\quad \sum\limits_{1 \leq b \leq \beta} M[i-1,a,b] < & \left\{\begin{array}{l} M[i,a,b] = \\ M[i-1,a-a_i,b] \vee \\ K[i,b],\; b \in (1,\beta) \end{array}\right. \\
\quad \sum\limits_{1 \leq b \leq \beta} M[i-1,a-a_i,b] \vee K[i,b] & \\
\quad \textbf{Otherwise } \forall b & M[i,a,b] = M[i-1,a,b]
\end{array}$$

where $\vee$ represents the logical OR function. This optical knapsack packing algorithm adds a β multiplications and $\beta + 2$ additions to the computational cost of the algorithm. This results in a $O(nA\beta)$ algorithm, which is still pseudo-polynomial. As with the original knapsack problem, if the discretization of A and the angular regions β are reasonable, the implementation is tractable.

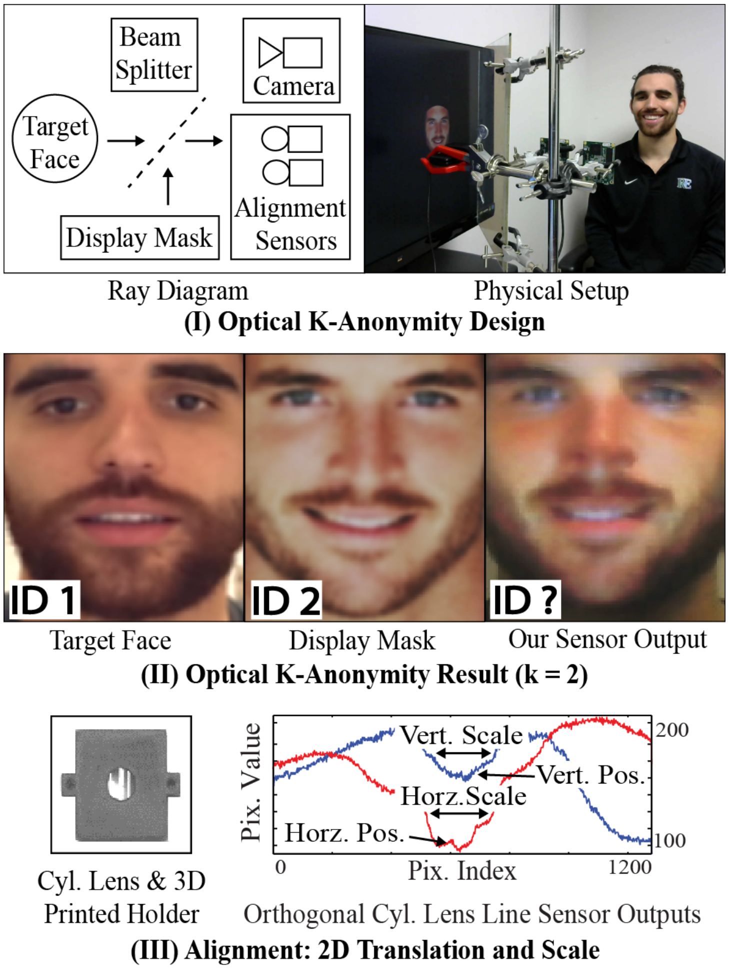

Ray Diagram Physical Setup

(I) Optical K-Anonymity Design

Target Face Display Mask Our Sensor Output

(II) Optical K-Anonymity Result (k = 2)

Cyl. Lens & 3D Printed Holder Orthogonal Cyl. Lens Line Sensor Outputs

(III) Alignment: 2D Translation and Scale

Figure 1: Optical K-Anonymity

2 Optical K-Anonymity

In Fig. 1 we show our design and results for our optics-based implementation of k-anonymity [2]. In Fig. 1(I) we show the ray diagram and physical setup for our design, whose primary input is k, the number of faces to anonymize a target face with. Light from a real target face is merged via a beamsplitter with illumination from a display showing the $k-1$ neighbor-faces and captured by a conventional sensor. The output is a k-anonymized face, directly captured by our sensor, at 30 FPS, as shown in Fig. 1(II). Finding the $k-1$ neighbors and 2D translation/scaling alignment, between the target and the neighbors, is achieved using two line sensors with orthogonally-oriented cylindrical lenses. The scale and position of the target face is found by identifying local extrema of the intensity profiles, Fig. 1(III).

[1] Shota Nakashima, Yuhki Kitazono, Lifeng Zhang, and Seiichi Serikawa. Development of privacy-preserving sensor for person detection. *Procedia-Social and Behavioral Sciences*, 2(1):213–217, 2010.

[2] Elaine M Newton, Latanya Sweeney, and Bradley Malin. Preserving privacy by de-identifying face images. *Knowledge and Data Engineering, IEEE Transactions on*, 17(2):232–243, 2005.

[3] Masashi Nishiyama, Hidenori Takeshima, Jamie Shotton, Tatsuo Kozakaya, and Osamu Yamaguchi. Facial deblur inference to improve recognition of blurred faces. In *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*, pages 1115–1122. IEEE, 2009.

This is an extended abstract. The full paper is available at the Computer Vision Foundation webpage.