

Negative FaceBlurring: A Privacy-by-Design Approach to Visual Lifelogging with Google Glass

TengQi Ye, Brian Moynagh, Rami Albatal and Cathal Gurrin

Centre for
Data Analytics

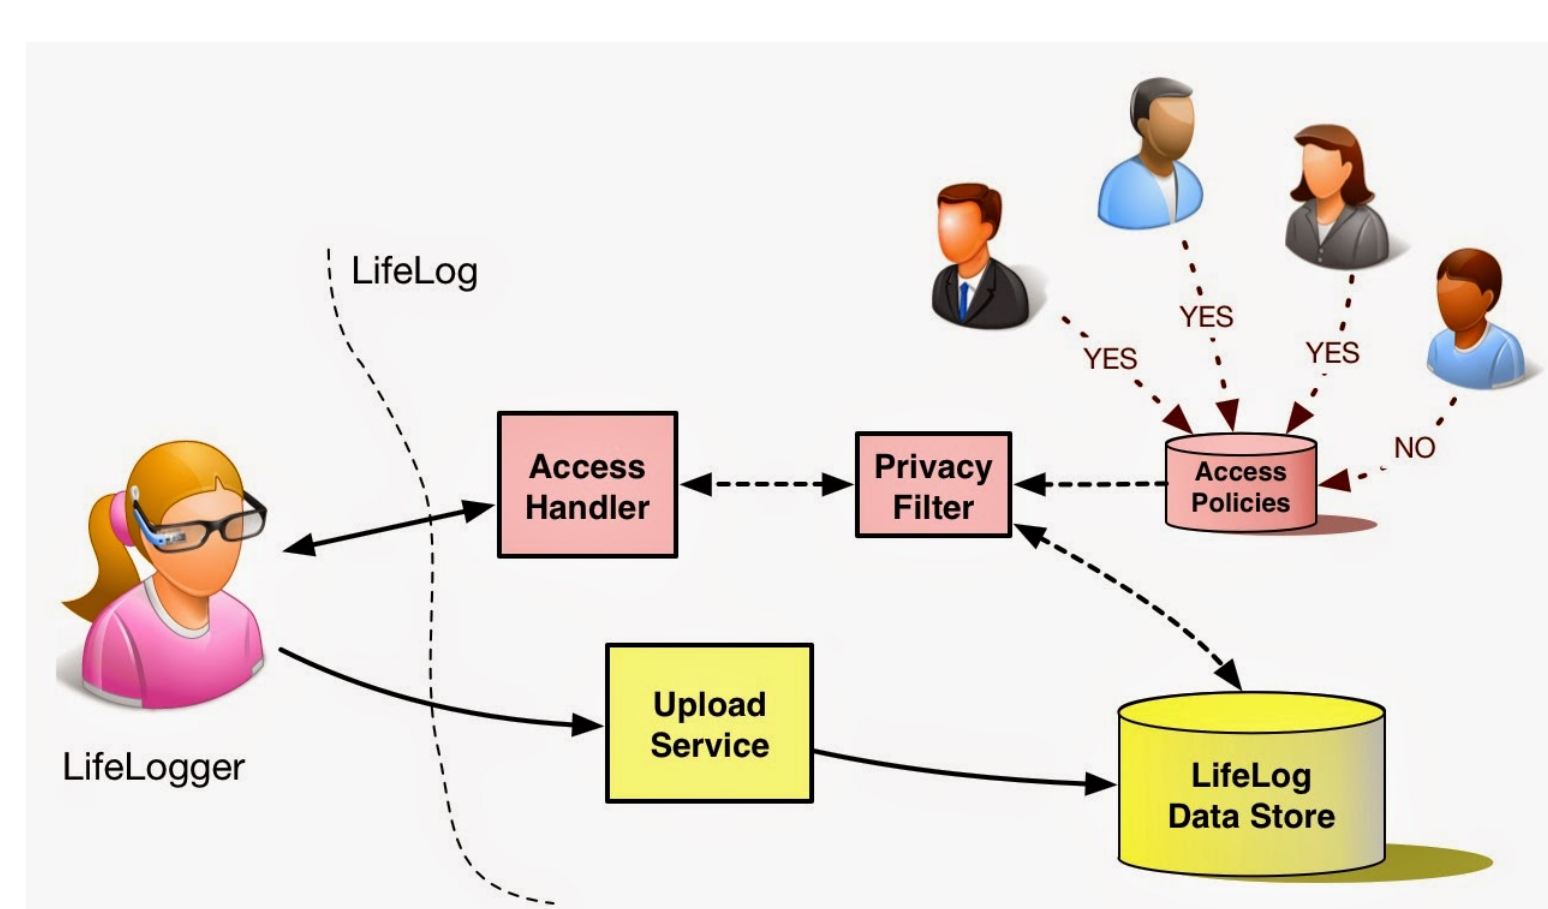
Insight



Visual Lifelogging

Wearable devices such as Google Glass are receiving increasing attention and look set to become part of our technical landscape over the next few years. At the same time, lifelogging is a topic that is growing in popularity with a host of new devices on the market that visually capture life experience in an automated manner. We describe a visual lifelogging solution for Google Glass that is designed to capture life experience in rich visual detail, yet maintain the privacy of unknown bystanders.

Our system is composed of two parts: a Google Glass glassware application and an independent online server (the lifelog). The glassware acts as the data gathering tool. The glassware uploads the images to the server post-capture where the images are stored (along with metadata) in their original form.



Conceptual Diagram of the System

We implement a Privacy-by-Design approach by treating privacy as the default configuration (i.e. we assume that all unknown bystanders and users wish to have their privacy protected unless they explicitly state otherwise).

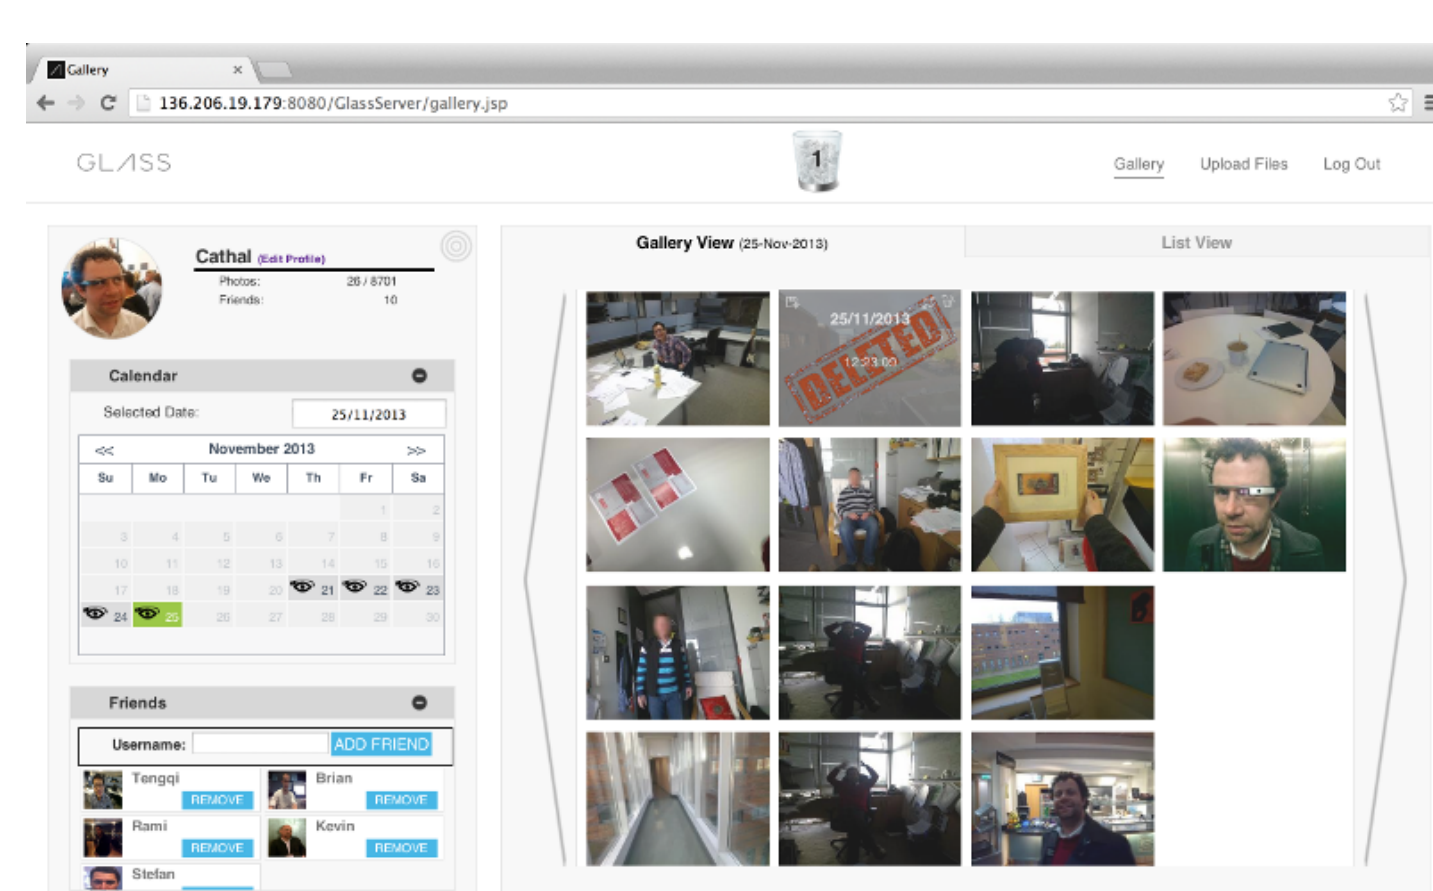
To evaluate our system 8,657 images were collected from three users over a period of 148 days.

Protecting the Privacy of Unknown Bystanders

We have taken the view that the visibility of an individual's face in an image is the main factor in preserving or violating that individual's privacy. Hence, successful implementation of a user's privacy policy is dependant upon two factors, accurate face detection and accurate face recognition.

The WWW-based Interface

A WWW-based interface allows users to view and maintain their own lifelog.



WWW-based Interface: Showing a day in the Lifelog

They can also choose the lifelogs in which their own image is permitted to appear (i.e. to define a privacy policy) by adding or removing friends. Separating the view of the lifelog data from the actual stored data, means user defined privacy policies can be updated in

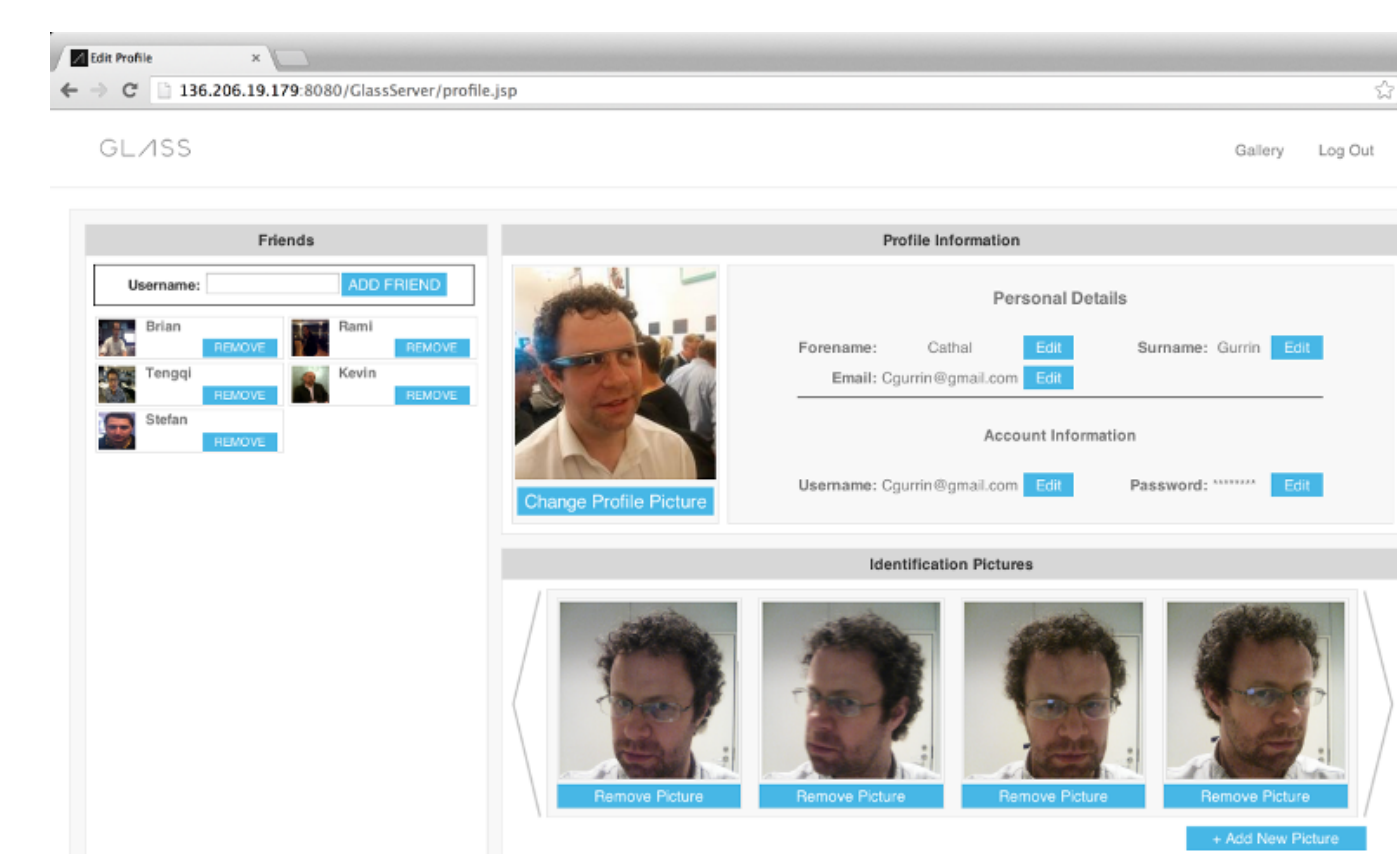
Face Detection and Recognition

Haar-like Feature-based Cascade Classifiers are used to detect any faces in images uploaded to the server. To obtain Haar-like features, an integral image is first generated to allow fast feature evaluation. However, a lot of Haar-like features are produced, most of which are useless for classification. Adaboost is employed to significantly reduce the space of Haar features

needed to classify a window on an image as containing or not containing a face.

Three different face recognition techniques are incorporated in this cascade:

1. Eigenfaces. Eigenfaces seek to find the principal components of the distribution of faces, or the eigenvectors of the covariance matrix of the set of face images. The face space is defined by training on a set of face images and calculating the Eigenfaces from user's profile face photos.
2. Fisherfaces The method is based on Fisher's Linear Discriminant and produces well separated classes in a low-dimensional subspace, even under severe condition like lighting and facial expressions. The idea is same classes should cluster tightly together, while different classes are as far away as possible from each other in the lower-dimensional representation.
3. Local Binary Patterns The method considers both shape and texture from face images. Local Binary Pattern histograms are extracted from small regions of face area and concatenated into a single, spatially enhanced feature histogram. The recognition is performed based on a nearest neighbour classifier in the computed feature space as a dissimilarity measure.



User Profile View: Showing user's profile photos

Thresholds for each method are obtained by analysing the Euclidean distances of each face from other faces in the training dataset. These thresholds are necessary to distinguish friends' faces from bystanders and are optimised by cross-validation on the training dataset. The higher the value of the threshold the more likely bystanders' faces are to be predicted as friends'; and vice versa. In our system, the three above recognition methods are all employed and assigned the same weight for the final decision.

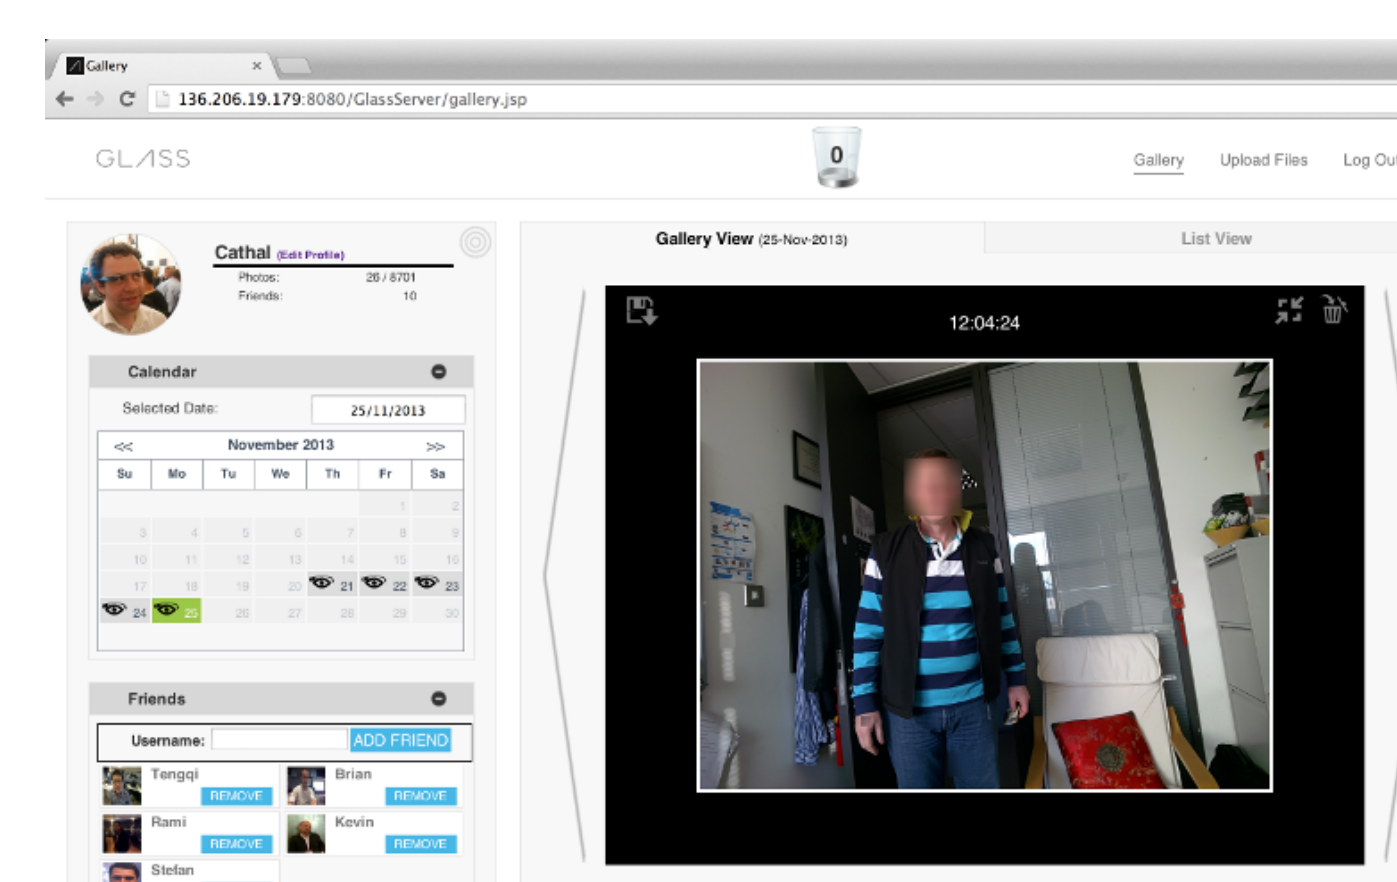
Results

To determine the effectiveness of our system, two evaluations were performed.

For face detection, Images containing faces in which any of the faces were not detected were counted as a fail. Of the 8,657 images, the system produced 6,985 passes and 1,672 fails generating a pass face detection accuracy of 80.68%.

For face recognition, we concerned ourselves with implementing the privacy policy and blurring unknown faces (the negative face blurring). 1,300 pictures with faces were randomly selected from the dataset. 1,310 faces were detected in this sample.

The false positive rate (i.e. bystanders classified as friends) was 0.76% and the false negative rate (i.e. friends classified as bystanders) was 29.01%. A total of 67.18% of the faces were correctly identified as bystanders.



Blurred Image of an unknown Bystander