Eyemotion: Classifying facial expressions in VR using eye-tracking cameras

Steven Hickson Georgia Institute of Technology shickson@gatech.edu

> Vivek Kwatra Google kwatra@google.com

Nick Dufour Google ndufour@google.com Avneesh Sud Google avneesh@google.com

Irfan Essa Georgia Institute of Technology irfan@gatech.edu

Abstract

One of the main challenges of social interaction in virtual reality settings is that head-mounted displays occlude a large portion of the face, blocking facial expressions and thereby restricting social engagement cues among users. We present an algorithm to automatically infer expressions by analyzing only a partially occluded face while the user is engaged in a virtual reality experience. Specifically, we show that images of the user's eyes captured from an IR gaze-tracking camera within a VR headset are sufficient to infer a subset of facial expressions without the use of any fixed external camera. Using these inferences, we can generate dynamic avatars in real-time which function as an expressive surrogate for the user. We propose a novel data collection pipeline as well as a novel approach for increasing CNN accuracy via personalization. Our results show a mean accuracy of 74% (F1 of 0.73) among 5 'emotive' expressions and a mean accuracy of 70% (F1 of 0.68) among 10 distinct facial action units, outperforming human raters.

1. Introduction

Facial expressions are essential for interpersonal communication and social interaction. They provide a means for conveying thought and emotion through visual cues that may not be easy to articulate verbally. However, virtual reality (VR) equipment using head-mounted displays (HMD) makes natural expressions difficult to recognize as half the face is occluded. Thus for VR systems to provide rich social interaction, faithfully representing these expressions in some manner is absolutely critical. We propose to recognize and convey facial expressions from inside a VR HMD.

Visual classification of expressions has been a well studied topic in computer vision. Most of this relies on faces [14, 23, 31]. We focus on a no extern, where the user is wearing a head-mounted display (HMD) in a VR setting



Figure 1: Eyemotion visual schematic. A: A user wearing the VR HMD used for expression tracking (Note that no external camera is used in our method; this is just for visualization). B: Interior of the HMD, with IR LEDs visible around the radius of the eyepieces, highlighted with red circles. C: Captured eye data. D: Model inference with dynamic avatar representation.

as shown in (Fig. 1A).

One could attach a user facing external camera that captures the lower face but that may not always be feasible such as in a mobile setting. An external camera is also not able to capture upper-facial expressions as they are occluded. We propose a new approach aimed at classification of facial action units (AUs) [12] and 'emotive' expressions using only internally mounted infrared cameras within the HMD. We are motivated by the recent availability of commercial HMDs with eye-tracking cameras [1]. This uses infrared cameras (Fig. 1B). These are used for tracking [29], but in our work we use the same input images for expression classification. A key aspect of our work is a *labeling*-based approach – as opposed to visual tracking of facial features*eye images*.

Our model classifies user expressions using only limited periocular eye image data, as shown in Fig. 1C, which is further limited by the large amount of intra-class variation among users. To account for limited data, a new type of data, and large variations, we turn to deep learning techniques. Recently convolutional neural networks (CNNs) [20, 15, 35] have performed very well on image classification tasks and are pervasive in machine learning and computer vision. Additionally, deep learning methods have the benefit of learning important invariant features and embeddings without requiring any hand-crafted feature representations. Deep learning has also been shown to give state of the art results on faces [22, 34]. Our approach, based on deep learning, outperforms normal human accuracy and even advanced (trained users) human accuracy for categorizing select facial expressions from our dataset of only IR eye images. Human ratings form the primary baseline for our work. We use these ratings for comparison and evaluation, but not as labels during training.

We also demonstrate an application of our classification framework that animates facial models and avatars in realtime, which could be used in social VR apps to convey and interpret users' facial expressions. Note that our *classification* based approach has applications beyond synthesis since it also provides semantics of the expression (potentially correlated with the emotive response to the VR environment), which may then be used as feedback to the system.

Our primary contributions are: (1) Demonstrating that the information required to classify a variety of facial expressions is reliably present in IR eye images captured by a commercial HMD sensor, and that this information can be decoded using a CNN-based method. (2) A novel personalization technique to improve CNN accuracy on new users. Across experiments, personalization resulted in a 4% accuracy improvement on average, and was statistically significant for a set of basic 'emotive' expressions (p = 0.018)and AUs (p = 0.001) (Section 4.2). (3) The collection of a unique dataset (Section 3) of eye images paired with expression labels, collected with two separate commercial HMDs each with 23 different users. (4) We show our method can be used to generate expressive avatars in real-time, which can function as an expressive surrogate for users engaged in VR environments (Section 5.2).

2. Related work

The problem of expression inference from limited facial data in a VR setting exists at the intersection of well studied problems in a variety of disciplines, discussed below.

Expression classification from visual data: While much work has been done on automatically inferring human expressions from images, nearly all of it focuses on unoccluded, frontal faces (see [27, 14, 5, 32] for recent surveys). Tian *et al.* [39] employ facial action units for finegrained facial expression recognition, using feature tracking and neural networks. Saatci *et al.* [31] use active appearance models to construct cascaded SVM classifiers for gender and 4 expressions/emotional states. The popularity of deep learning produced renewed interest in the field, with new challenges like [11] requiring emotion recognition on video data, and [13] requiring classification of 11 AUs and recognition of basic and compound emotions. Recent work on expression inference includes Liu et al. [23], who propose a boosted deep belief network, and Barsoum et al. [4], who use the VGG network [35] to learn emotions from noisy crowd-sourced labels. Kahou et al. [17] and Bargal et al. [3] use deep convolutional networks combined with SVM classifiers for expressions from the EmotiW dataset challenge [11]. Benitez-Quiroz et al. [13] use Kernel Subclass Discriminant Analysis (KSDA) to classify facial action units, intensities and emotion categories on a large image dataset in the wild. All of these works require full, unoccluded face images, unavailable in our scenario. There has been some work with partial occlusion of the face such as [16, 18, 7].

Expression classification with alternate sensors: There has been some recent research on expression classification using wearable sensors. Scheirer et al. [33] use face mounted piezoelectric sensors to discriminate between confused and interested expressions as well as discriminating between expressive states in general and neutral states. More recently Masai et al. [24] used optical sensors mounted on glasses to determine 8 expressions similar to [11] with a small set of users, and performed hardware modifications for facial expression mapping inside a HMD [36]. This work is the closest to our proposal of expression classification in virtual reality headsets. However, we propose a method for expression classification using gaze tracking cameras rather than embedded optical or piezoelectric sensors, which is more robust to personalized fit, and does not require customized sensors beyond a commercial eye-tracking system. Much work using alternate sensors has been motivated by facial re-enactment (e.g. [22]) and is covered below.

Gaze tracking in VR: Gaze tracking is the subject of intense and sustained research, and used for many interactive applications [25]. Gaze tracking has recently seen applications in virtual and mixed reality [21, 38] and on mobile devices using convolutional neural networks [19]. However, these methods do not classify expressions or AUs.

Facial re-enactment in VR: Unoccluded face synthesis, facial reenactment and avatar re-targeting in VR has been an active area of recent research. These approaches use a combination of visual tracking of the lower (unoccluded) face, along with custom sensors mounted inside the HMD. Burgos *et al.* [8] composite the occluded part of the face by aligning and blending the unoccluded parts with a pre-

captured facial expression database. Li et al. [22] measure strain signals with electronic sensors to estimate facial expressions of the occluded face, and combine with a user facing RGB-D camera to train a linear 3D parametric model of blendshapes. Olszewski et al. [26] propose an approach for 3D avatar control with eye-tracking cameras and a user facing RGB camera. A CNN model is trained to regress from these streams to blendshape coefficients. Thies et al. [38] also perform real-time gaze-aware facial reenactment in VR using a RGB-D camera to capture the unoccluded regions, and two internal infra-red (IR) cameras to track the eye gaze. A multi-linear blendshape model is fit to these streams by optimizing for photometric and geometric alignment. Zhao et al. [43] synthesize unoccluded face images using an HMD case fitted with wide-angle near IR cameras, along with a user facing RGB camera. A 3D bilinear blendshape model is fit to the input, the occluded part is synthesized from warped, and colorized IR eye images with evaluation on unoccluded face images.

In contrast, our work aims to classify a set of facial expressions from cameras present in eye-tracking enabled VR HMDs, with all optical and field-of-view constraints. As such our work cannot be directly compared against these; instead we show comparisons among various CNN-based models trained on eye images and against human ratings to benchmark our results. Most attempts at personalization use per-subject samples and quick retraining [9, 42]. However there has also been some work at personalizing expression classification without retraining with new samples [10] based on unsupervised generalization approach which requires no retraining with a deep learning framework.

3. Our Dataset

We perform supervised training using a CNN to classify face expressions from eye images recorded inside the HMD. We collect a large amount of data which is cumbersome or impossible to label by hand. This section describes our data collection method, which obtains ground truth expression data from participants using infra-red eye images and without the need for manual image annotations such as that required by spontaneous expressions.

3.1. Data Collection

There are some public datasets of near infra-red eye images [6] that target iris detection and biometric authentication, but contain no expression labels. In addition, these datasets do not directly translate to the novel sensors used in VR. Therefore, we designed a system to collect such data from participants in a controlled setting.

We collected a subset of facial action units that influence the upper face, and could be reliably performed by multiple subjects. We also distinguish between left and



Figure 2: Inferring expressions from eye images alone is significantly different from doing so using the unoccluded face. Here we have four participants making four different expressions: without wearing the HMD (column 1, for reference only), within the HMD (column 2), and within the HMD during session 2 (column 3). The expressions are (A) happiness, (B) anger, (C) surprise, (D) and squint. The difficulty of this task, even for humans, is clear as the amount of expressive information conveyed is significantly reduced while variability is readily visible.

right AUs, where applicable. These are Neutral (AU0), Left Brow Raise (AU1+2L), Right Brow Raise (AU1+2R), Brow Lower (AU4), Upper Lid Raise (AU5), Squint (AU44), Both Eyes Closed (AU43), Left Wink (AU46L), Right Wink (AU46R), and Cheek Raise (AU6). We also collect 'emotive' expressions for basic emotions as defined by [12], which are Neutral, Anger, Surprise, and Happiness. See Fig. 2 for an illustration of the variability of the data. We experiment with training classifiers on facial action units, a subset of useful and non-overlapping facial action units, expressions, and a subset of expressions useful for VR social environments. The mapping from AUs to emotional expressions has been well characterized [12].

3.2. Collection Setup

We collected data with two separate HMDs, with near IR (880nm) cameras mounted between the eye lens and display screen using a beam splitter. HMD1 and HMD2 capture 200x200-pixel and 320x240-pixel eye image pairs respectively, both at 10Hz. The different headset cameras and camera requirements of each HMD allow characterization of the generalization of our technique (they are different



Figure 3: The Eyemotion pipeline, including personalization. A: Raw eye images from the HMD. B: Rectified eye images. C: The average neutral image for this user session, used for personalization. D: The difference between the rectified headset image and the mean neutral image is the input to a deep neural network. In the non personalization case, the mean neutral image is not subtracted from the rectified image. E: Output takes the form of a distribution over expressions. F: This distribution is used to generate an expressive avatar.

manufacturered headsets from SMI). 23 different participants were collected with each HMD with different genders (for a total of 46), ethnicities, and hair color.

We collected these data by asking users to form an expression, giving them an example from an exemplar video. While this may not result in spontaneous expressions [41], it provides explicit labels for each expression. Removing the need for expert ratings allows larger-scale data collection than would otherwise be possible. To provide a realistic exemplar, we first recorded videos of trained actors performing each expression for the participant to use as a reference. During the collection process, for each expression, we provide to the participant the name of the expression, a looped clip of an actor performing the expression, and a live video of the participant in order for them to practice the expression. If the participant can't perform the expression, they are able to skip it. Otherwise they hold the expression and follow a randomly moving target on the screen with their gaze (to encourage gaze diversity) or head pose (to encourage variations in HMD pose). This continues for all expressions and AUs (these are the images in column 1 of Fig. 2). We then have them put on the HMD and repeat the process twice more, taking the headset off and putting it back on to account for slippage and variation in fit. Each of these headset repetitions constitutes a 'session.'

3.3. Data Characteristics

To ensure diversity, we collected each participant's personal information through a data collection form. Of the 46 participants: 16 were female; 16 were aged 35 or over from an age range of 18 to 64 with a median age of 30; 11 participants had non-brown eyes and 4 had non-brown or black hair. 25 of our participants were nonwhite, with 9 Asian, 7 east Indian or south Asian, 4 two or more races, 3 Hispanic or Latino, 2 African American, and 2 preferring not to say.

Approximately 50,000 eye image pairs were collected per HMD (about 2,000 per participant). Each expression was collected for the same amount of time; however, as some participants were unable to perform all of the expressions, there exists slight variability in the number of images per expression.

3.4. Data Cleanup

Since participants blink during the data collection process, and we want to remove this source of variance, we removed closed-eye images with a custom classifier, similar to the approach in Section 4.1. The classifier was trained to recognize eyeblinks using the neutral and the eyes closed images from our participants. Neutral and closed-eye images were validated and cleaned manually by ensuring that neutral images did not include any blinks and that eyes were actually closed during closed-eye periods. Approximately 400 images were removed by hand and using the classifier.

3.5. Data Augmentation

The tightly controlled acquisition environment in the headset means lighting and eye-camera viewpoint are largely fixed, with nearly all of the variation coming from differences in participants and the position of the headset on the head during acquisition. Thus data augmentation was performed carefully and monitored for consistency. Further, the semantic meaning of many of the labels (*e.g.* 'Left Eye Wink') precludes random flips. Random rotations are limited so that they do not exceed human variation in eye orientation. We found a 2% variation (rather than 10% used frequently) is appropriate when performing random augmentation of rotation, scale, and brightness.

4. Our Approach

Initially, we experimented with features obtained from facial iris and eye landmark positions, similar to those produced in [14] and [44]. However, testing with a proprietary eye landmarker produced frequent failures when the partic-



(a) A correct example of the (b) Most eye landmark tracker eye landmark tracker. results contained errors.

Figure 4: Eye tracker results with blue dots as the iris and pupil landmark locations and green dots as the palpebral fissure landmark locations.

ipants made expressions that distorted the shape of the eyes (see Fig. 4). Using those kinds of existing hand-crafted features on our new domain does not yield good features and ignores the periocular data which is found to contain useful information [28]. Recent work on facial animation using VR and an external camera [22] has also shown that deep models work well to estimate facial expressions, which is similar to our problem. The entire pipeline, including personalization, is represented schematically in Fig. 3.

4.1. CNN Architecture

Our proposed method leverages a CNN to learn an embedding describing expressions and emotions using infrared eye images. Specifically, we train a variant of the widespread Inception architecture [37] using the Tensor-Flow library [2] motivated by experimentation on various CNN architectures, described in Section 5. The model used was pre-trained for 150,000 iterations on the Imagenet data. Data are registered and augmented as described in Section 3. The HMD eye cameras are calibrated, and both eye images are rectified, concatenated and scaled to 299x299 pixels. The network was trained using a learning rate of 0.045 which decays stepwise by 0.94 every epoch. To prevent overfitting, we used L2 weight decay (0.0004). The softmax cross-entropy function was chosen as the loss along with L2 regularization.

Optimization was performed with a 'RMSProp' optimizer [40], with momentum 0.9, decay factor 0.9, and ϵ of 1.0.

4.2. Personalization

One of the dominant sources of variance in our data is individual variation in appearance. We attempt to partially remove this variance since it does not vary with, and may not be predictive of, affective state. Since this variance removal occurs within-subject, it is effectively 'personalization.' Our approach is inspired by the standard practice of mean image subtraction. To remove appearance based variation, we construct a mean neutral image, one for each person/session pair by averaging together first 5 seconds of



Figure 5: Personalization is performed by subtracting each user's mean neutral data from the current image to reduce the unimportant sources of variance and highlight important ones. Columns are 1: the original image, 2: the mean neutral image, and 3: the difference between the two (which are contrast normalized for demonstrative purposes). Row are different expressions, with A: happiness, B: anger, and C: brow raise.

their neutral images, as neutral images are stable over time. This image is subtracted from all other images derived from the originating person/session pair to effectively normalize it for a user (see Fig. 5) per Equation 1 where P(I) is the CNN test input, I is the original image, and N_u is the set of neutral images for that user. An 'in practice' realization of this technique would be requesting a user maintain a minimally emotive expression for a short period of time and using the accumulated data to construct a similar normalizing image effectively giving each person a different mean subtraction. This generates P(I) for new/test users.

$$P(I) = I - \frac{1}{|N_u|} \sum_{I_u \in N_u} I_u$$
 (1)

Repeated experiments demonstrated that 'personalizing' images by subtracting a separate mean image per user is an effective means of increasing the accuracy of the classifier.

To avoid introducing a bias towards statistical significance when testing the effectiveness of personalization, we first collapsed across results within a single subject-session-condition triplet (*e.g.*, subject 28, session 2, condition 'brow raise') into a single average accuracy and F1 score, then conducted a paired 1-tailed *t*-test to test the significance of the difference in average accuracy across the 23 users between results obtained with and without 'personalization.' Action unit expression classification was significantly increased by the introduction of personalization (p < 0.001) as was 'emotive' expression classification (p = 0.018).

5. Experiments & Results

We experiment with benchmarks of different convolutional neural networks in a small, single validation set study on the 10 action units listed in Section 3 with 4 holdout participants and 19 participants to train on from HMD1. This yields a comparison metric for our proposed architecture. The different architectures we tried as benchmarks are listed in Table 1. For all our uses of InceptionV3, we fine-tune a model trained off of the ImageNet dataset [30].

Methods 1-4 were results from human raters as described in Section 5.1. Method 5 is the result from professions FACs coders validating our ground truth data. Method 6 is the same ignoring left/right brow raise. Method 7 is obtained from a two tower (one-per-eye) gaze tracking network adapted from [19]. Method 8 is InceptionV3 with two towers (one for each eye) and tied weights. Method 9 is a single-tower InceptionV3, which receives concatenated eye images, but only the last fully connected softmax layers are allowed to vary. Method 10 is the same approach with the alexnet architecture. Methods 11 and 12 are fine-tuned InceptionV3 with data augmentation as described in Section 3.5. Method 13/14 is our Eyemotion approach detailed in Section 4.

Initial tests we performed on different participants showed that results across individuals (without personalization) vary drastically and range between 31% to 90% between users/headsets (see Supplementary materials). Results depend on factors such as ergonomic fit, eyebrow color, eyebrow position, and expressiveness. Because of this, one hold out test set of participants is acceptable in a small study but not an effective enough performance measurement for comparison of our actual model accuracy. Therefore, in the experiments below, we use 5-fold cross

Method	Accuracy
1: Average of human raters	47%
2: Average of human raters w/neutral	50%
3: Advanced human rater	58%
4: Advanced human raters w/neutral	62%
5: FACs Coder	73%
6: FACs Coder without individual brow raises	81%
7: Variants based off of [19]	26%
8: InceptionV3 one tower per eye	48%
9: InceptionV3 w/ frozen weights	55%
10: Alexnet HMD1	57%
11: InceptionV3 HMD1	65%
12: InceptionV3 HMD2	69%
13: Eyemotion (our approach) HMD1	70%
14: Eyemotion (our approach) HMD2	73 %

Table 1: Preliminary tests using 4 left out participants.

validation by holding out participants on our dataset to get a more robust comparison between our approach and the baseline model.



Figure 6: Visualizing the correct and incorrect classifications of the model for a subset of expressions. Columns 1 and 2 are correct classifications (green border), while column 3 are false negatives (red border). Rows are different expressions. A: Surprise, B: Neutral, C: Happiness, D: Eyes closed, E: Anger. In some cases, it is clear why the incorrect classification was made.

We found results on HMD1 and HMD2 were roughly the same as shown in Table 1, so for brevity we present results for HMD2 only and include more results in our supplementary material. Our model achieves a mean accuracy on 'emotive' expressions of 66.6% and 73.7% without and with personalization, respectively. Facial action unit classification accuracy was 63.7% without personalization and 70.2% with personalization (results in Fig 6, per class details shown in Tables 2, 3, 4, and 5). Both results have higher than the average advanced human rater accuracy of 60.8% discussed in Section 5.1.

5.1. Human Accuracy

We took a subset of 350 images of the data with different users and expressions in order to form a human benchmark of the data given that humans excel at pattern matching. We had 2 beginner classifiers (who had not seen any eye images previously), 1 intermediate classifier (who had seen some eye images previously), and 2 advanced, welltrained classifiers look at the concatenated eye images and guess the label from the 10 facial action units on each of those images. We also do this with a sample image of each person as a personalization method can be useful for the

Action Unit	Precision	Recall	F_1	Support
Brow lower	0.59	0.66	0.63	2897
Upper lid raise	0.74	0.76	0.75	4912
Cheek raise	0.56	0.64	0.60	2659
Eyes closed	0.54	0.51	0.53	793
Left brow raise	0.52	0.28	0.36	1598
Left wink	0.95	0.86	0.90	2638
Neutral	0.48	0.59	0.53	2309
Right brow raise	0.29	0.18	0.22	1855
Right wink	0.83	0.87	0.85	2220
Squint	0.62	0.71	0.66	2135
Avg / total	0.64	0.65	0.64	24016

Table 2: Facial action units finetuned with InceptionV3. The overall mean accuracy 63.7%.

Action Unit	Precision	Recall	F_1	Support
Brow lower	0.67	0.74	0.70	2576
Upper lid raise	0.75	0.70	0.73	4956
Cheek raise	0.65	0.64	0.64	2903
Eyes closed	0.92	0.54	0.68	872
Left brow raise	0.61	0.33	0.43	1616
Left wink	0.93	0.90	0.91	2577
Neutral	0.53	0.96	0.68	2372
Right brow raise	0.50	0.24	0.32	1701
Right wink	0.82	0.87	0.84	2235
Squint	0.65	0.67	0.66	2208
Avg / total	0.70	0.69	0.68	24016

Table 3: Facial action units finetuned with InceptionV3 using our personalization method. Mean accuracy 70.2%.

classifier. The inter-rater kappa from our raters without and with a sample neutral image were 0.61 and 0.64 respectively showing good agreement among raters. The F1 score of the best rater was 0.63 without and 0.65 with neutral images (worse than our method's results). This suggests that some amount of one-shot learned personalization should be able to improve results. We also computed the inter-rater kappa from 4advanced human raters on 350 random images without andwith a sample neutral image per sample. The scores were0.61and0.64respectively showing good agreement amongraters. We also computed the f1 score for the advancedhuman rater: 0.63 without and 0.65 with a neutral image;worse/lower than our methods (0.64 without, 0.68 with).

5.2. Applications

Given our model produces a probability distribution of AUs or expressions, we can use that output to infer what a user is feeling and/or drive an avatar showing their ex-

Expression	Precision	Recall	F_1	Support
Anger	0.72	0.72	0.72	2695
Closed Eyes	0.77	0.78	0.77	899
Happiness	0.66	0.67	0.66	2610
Neutral	0.60	0.78	0.68	2274
Surprise	0.71	0.51	0.59	2425
Avg / total	0.68	0.68	0.67	10903

Table 4: Expressions finetuned with InceptionV3. Overall mean accuracy 66.6%.

Expression	Precision	Recall	F_1	Support
Anger	0.71	0.83	0.77	2455
Closed Eyes	0.90	0.67	0.77	863
Happiness	0.81	0.60	0.69	2644
Neutral	0.64	0.95	0.77	2425
Surprise	0.81	0.60	0.69	2516
Avg / total	0.76	0.74	0.73	10903

Table 5: Expressions finetuned with InceptionV3 using our personalization method. Overall mean accuracy 73.7%.

pressions in real-time. Smoothing with exponential decay is applied on the inferred expression probabilities to get temporally stable labels. The model also yields a more intuitive expression medium for VR as opposed to gestures or keyboard inputs. We can also adjust a user's VR environment based on their expression and hence their inferred emotional

state This can be demonstrated with a simple example where we change the ambient light and color based on the user's expressions while they watch a film clip as shown in the supplementary material. An example of this is included in the supplementary video, and shown in Fig. 8a and Fig. 8b. Clearly these are very simple applications of our framework, and far more elaborate adjustments of the ambient environment are possible.

6. Discussion

Our initial studies rapidly demonstrated the superiority of the concatenated-eye architecture, in which eye images are concatenated together before being input (*see* Table 1). This architecture outperformed all others by ten percentage points or more when evaluated on the hold-out participants. We hypothesize that the improvement is due to concatenating the eye images at the outset, giving the CNN the ability to combine information from each eye. When the testing was extended to four other folds, such that each participant is excluded from training in exactly one fold, this accuracy improvement remained. Introducing a novel eye personalization technique, motivated by mean image subtraction, further improved the accuracy by nearly five per-



Figure 7: Confusion matrices of advanced human raters showing the accuracy of a)Advanced human, b) Advanced human with a given neutral image, c) Our network tested on holdout users never seen before. Labels are *Neutral (AU0), Left Brow Raise (AU1+2L), Right Brow Raise (AU1+2R), Brow Lower (AU4), Upper Lid Raise (AU5), Squint (AU44), Both Eyes Closed (AU43), Left Wink (AU46L), Right Wink (AU46R), and Cheek Raise (AU6).*



(a) Animation still with our (b) Animation still with our model inferring happiness. model inferring anger.
Figure 8: An animation that can be driven by our model's inference.

centage points on average, which *t*-tests showed was significant (all p < 0.05). Typical errors occured when a visible component (*e.g.* eyebrows) was occluded (*see* Fig. 6).

We compare our personalization method to nonpersonalization methods and an implementation inspired by [19]. Compared to other attempts at personalization, we instead explicitly create a difference from the mean neutral image and show this yields better results on average across all classes and especially improving neutral classification. In addition, some comparison could be made to [24] where they classify 8 expressions on 8 users achieving good accuracy on holdout sessions. However, on hold-out users, they achieve an accuracy of at most 48% and determined that a new user requires individual training for classification. Our deep learning personalization method scales well to new users and requires no individual training.

We also found that our implementation is robust to different hair colors and makeup. However, incorrect labels during data collection harm our results, which is why left brow raise and right brow raise have such low precision and recall as many people cannot do both and, critically, participants during data collection would attempt to do them even if they could not. This is a downside of autonomous data collection in which participants are asked to perform, as we have no means other than the participant's self report that the performances are accurate. It is, however, non-trivial to collect in-the-wild or unoccluded data for this task.

7. Conclusions

Future work would consider Using naturalistic stimuli (such as video) and treating the task as a clustering or an alignment would spare participants the requirement of having to 'fake' expressions and remove the requirement of imposing a fixed set of expressions.

A large amount of information about facial expressions is encoded in the eyes alone. Perhaps more surprising, this work shows that such information is interpretable by a CNN and can automate expression inference, even when the images are relatively low fidelity and derived from consumergrade equipment. Even with 50 users, our model handily exceeds human classification accuracy and can power an avatar that serves as an expressive surrogate for a user wearing an HMD. We have demonstrated using consumergrade eye-tracking cameras, which are already being included in VR headsets, a means to preserve and transmit social information among users engaged in VR. *We also propose a novel personalization method with no retraining that achieves 74% accuracy on classifying expressions and 70% accuracy on classifying facial action units*.

References

- [1] Smi eye tracking hmd upgrade for the oculus rift dk2. http://www.smivision.com/en/ gaze-and-eye-tracking-systems/products/ eye-tracking-oculus-rift.html. 1
- M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, et al. Tensorflow: Large-scale machine learning on heterogeneous distributed systems. *arXiv preprint arXiv:1603.04467*, 2016.
- [3] S. A. Bargal, E. Barsoum, C. C. Ferrer, and C. Zhang. Emotion recognition in the wild from videos using images. In *Proceedings of the 18th ACM International Conference on Multimodal Interaction*, ICMI 2016, pages 433–436, New York, NY, USA, 2016. ACM. 2
- [4] E. Barsoum, C. Zhang, C. C. Ferrer, and Z. Zhang. Training deep networks for facial expression recognition with crowdsourced label distribution. In *Proceedings of the 18th ACM International Conference on Multimodal Interaction*, ICMI 2016, pages 279–283, New York, NY, USA, 2016. ACM. 2
- [5] V. Bettadapura. Face expression recognition and analysis: The state of the art. *CoRR*, abs/1203.6722, 2012. 2
- [6] K. W. Bowyer and P. J. Flynn. The ND-IRIS-0405 iris image dataset. CoRR, abs/1606.04853, 2016. 3
- [7] I. Buciu, I. Kotsia, and I. Pitas. Facial expression analysis under partial occlusion. In *Acoustics, Speech, and Signal Processing, 2005. Proceedings.(ICASSP'05). IEEE International Conference on*, volume 5, pages v–453. IEEE, 2005.
 2
- [8] X. P. Burgos-Artizzu, J. Fleureau, O. Dumas, T. Tapie, F. LeClerc, and N. Mollet. Real-time expression-sensitive hmd face reconstruction. In *SIGGRAPH Asia 2015 Technical Briefs*, page 9. ACM, 2015. 2
- [9] D. Castro, S. Hickson, V. Bettadapura, E. Thomaz, G. Abowd, H. Christensen, and I. Essa. Predicting daily activities from egocentric images using deep learning. In *Proceedings of the 2015 ACM International symposium on Wearable Computers*, pages 75–82. ACM, 2015. 3
- [10] W.-S. Chu, F. De la Torre, and J. F. Cohn. Selective transfer machine for personalized facial action unit detection. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2013. 3
- [11] A. Dhall, R. Goecke, J. Joshi, J. Hoey, and T. Gedeon. Emotiw 2016: Video and group-level emotion recognition challenges. In *Proceedings of the 18th ACM International Conference on Multimodal Interaction*, pages 427– 432. ACM, 2016. 2
- [12] P. Ekman and W. V. Friesen. Facial action coding system. Consulting Psychologists Press, Stanford University, Palo Alto, 1978. 1, 3
- [13] C. Fabian Benitez-Quiroz, R. Srinivasan, and A. M. Martinez. Emotionet: An accurate, real-time algorithm for the automatic annotation of a million facial expressions in the wild. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016. 2
- [14] B. Fasel and J. Luettin. Automatic facial expression analysis: a survey. *Pattern recognition*, 36(1):259–275, 2003. 1, 2, 4

- [15] R. Girshick, J. Donahue, T. Darrell, and J. Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 580–587, 2014. 2
- [16] X. Huang, G. Zhao, W. Zheng, and M. Pietikäinen. Towards a dynamic expression recognition system under facial occlusion. *Pattern Recognition Letters*, 33(16):2181–2191, 2012.
 2
- [17] S. E. Kahou, C. Pal, X. Bouthillier, P. Froumenty, Ç. Gülçehre, R. Memisevic, P. Vincent, A. Courville, Y. Bengio, R. C. Ferrari, et al. Combining modality specific deep neural networks for emotion recognition in video. In *Proceedings of the 15th ACM on International conference on multimodal interaction*, pages 543–550. ACM, 2013. 2
- [18] I. Kotsia, I. Buciu, and I. Pitas. An analysis of facial expression recognition under partial facial image occlusion. *Image* and Vision Computing, 26(7):1052–1067, 2008. 2
- [19] K. Krafka, A. Khosla, P. Kellnhofer, H. Kannan, S. Bhandarkar, W. Matusik, and A. Torralba. Eye tracking for everyone. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 2, 6, 8
- [20] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012. 2
- [21] E. C. Lee, K. R. Park, M. C. Whang, and J. Park. Robust gaze tracking method for stereoscopic virtual reality systems. In *International Conference on Human-Computer Interaction*, pages 700–709. Springer, 2007. 2
- [22] H. Li, L. Trutoiu, K. Olszewski, L. Wei, T. Trutna, P.-L. Hsieh, A. Nicholls, and C. Ma. Facial performance sensing head-mounted display. ACM Transactions on Graphics (Proceedings SIGGRAPH 2015), 34(4), July 2015. 2, 3, 5
- [23] P. Liu, S. Han, Z. Meng, and Y. Tong. Facial expression recognition via a boosted deep belief network. In *Proceed*ings of the IEEE Conference on Computer Vision and Pattern Recognition, pages 1805–1812, 2014. 1, 2
- [24] K. Masai, Y. Sugiura, M. Ogata, K. Kunze, M. Inami, and M. Sugimoto. Facial expression recognition in daily life by embedded photo reflective sensors on smart eyewear. In *Proceedings of the 21st International Conference on Intelligent User Interfaces*, pages 317–326. ACM, 2016. 2, 8
- [25] C. H. Morimoto and M. R. Mimica. Eye gaze tracking techniques for interactive applications. *Computer Vision and Image Understanding*, 98(1):4–24, 2005. 2
- [26] K. Olszewski, J. J. Lim, S. Saito, and H. Li. High-fidelity facial and speech animation for vr hmds. ACM Transactions on Graphics (Proceedings SIGGRAPH Asia 2016), 35(6), December 2016. 3
- [27] M. Pantic and L. J. M. Rothkrantz. Automatic analysis of facial expressions: The state of the art. *IEEE Trans. Pattern Anal. Mach. Intell.*, 22(12):1424–1445, Dec. 2000. 2
- [28] U. Park, R. R. Jillela, A. Ross, and A. K. Jain. Periocular biometrics in the visible spectrum. *IEEE Transactions on Information Forensics and Security*, 6(1):96–106, 2011. 5
- [29] A. Patney, J. Kim, M. Salvi, A. Kaplanyan, C. Wyman, N. Benty, A. Lefohn, and D. Luebke. Perceptually-based

foveated virtual reality. In *ACM SIGGRAPH 2016 Emerging Technologies*, SIGGRAPH '16, pages 17:1–17:2, New York, NY, USA, 2016. ACM. 1

- [30] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)*, 115(3):211–252, 2015. 6
- [31] Y. Saatci and C. Town. Cascaded classification of gender and facial expression using active appearance models. In 7th International Conference on Automatic Face and Gesture Recognition (FGR06), pages 393–398. IEEE, 2006. 1, 2
- [32] E. Sariyanidi, H. Gunes, and A. Cavallaro. Automatic analysis of facial affect: A survey of registration, representation, and recognition. *IEEE transactions on pattern analysis and machine intelligence*, 37(6):1113–1133, 2015. 2
- [33] J. Scheirer, R. Fernandez, and R. W. Picard. Expression glasses: a wearable device for facial expression recognition. In CHI'99 Extended Abstracts on Human Factors in Computing Systems, pages 262–263. ACM, 1999. 2
- [34] F. Schroff, D. Kalenichenko, and J. Philbin. Facenet: A unified embedding for face recognition and clustering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 815–823, 2015. 2
- [35] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556, 2014. 2
- [36] K. Suzuki, F. Nakamura, J. Otsuka, K. Masai, Y. Itoh, Y. Sugiura, and M. Sugimoto. Facial expression mapping inside head mounted display by embedded optical sensors. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology*, pages 91–92. ACM, 2016. 2
- [37] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna. Rethinking the inception architecture for computer vision. *arXiv preprint arXiv:1512.00567*, 2015. 5
- [38] J. Thies, M. Zollhöfer, M. Stamminger, C. Theobalt, and M. Nießner. Facevr: Real-time facial reenactment and eye gaze control in virtual reality. *arXiv preprint arXiv:1610.03151*, 2016. 2, 3
- [39] Y.-I. Tian, T. Kanade, and J. F. Cohn. Recognizing action units for facial expression analysis. *IEEE Trans. Pattern Anal. Mach. Intell.*, 23(2):97–115, Feb. 2001. 2
- [40] T. Tieleman and G. Hinton. Lecture 6.5-rmsprop: Divide the gradient by a running average of its recent magnitude. *COURSERA: Neural Networks for Machine Learning*, 4(2), 2012. 5
- [41] X. Zhang, L. Yin, J. F. Cohn, S. Canavan, M. Reale, A. Horowitz, P. Liu, and J. M. Girard. Bp4d-spontaneous: a high-resolution spontaneous 3d dynamic facial expression database. *Image and Vision Computing*, 32(10):692–706, 2014. 4
- [42] X. Zhao, X. Liang, L. Liu, T. Li, Y. Han, N. Vasconcelos, and S. Yan. Peak-piloted deep network for facial expression recognition. In *European Conference on Computer Vision*, pages 425–442. Springer, 2016. 3
- [43] Y. Zhao, Q. Xu, X. Huang, and R. Yang. Mask-off: Synthesizing face images in the presence of head-mounted displays. *arXiv preprint arXiv:1610.08481*, 2016. 3

[44] Z. Zhu and Q. Ji. Robust real-time eye detection and tracking under variable lighting conditions and various face orientations. *Computer Vision and Image Understanding*, 98(1):124–154, 2005. 4