

Unsupervised Representation Learning for Gaze Estimation

Yu Yu, Jean-Marc Odobez

Idiap Research Institute, CH-1920, Martigny, Switzerland

EPFL, CH-1015, Lausanne, Switzerland

{yyu, odobez}@idiap.ch

Abstract

Although automatic gaze estimation is very important to a large variety of application areas, it is difficult to train accurate and robust gaze models, in great part due to the difficulty in collecting large and diverse data (annotating 3D gaze is expensive and existing datasets use different setups). To address this issue, our main contribution in this paper is to propose an effective approach to learn a low dimensional gaze representation without gaze annotations, which to the best of our best knowledge, is the first work to do so. The main idea is to rely on a gaze redirection network and use the gaze representation difference of the input and target images (of the redirection network) as the redirection variable. A redirection loss in image domain allows the joint training of both the redirection network and the gaze representation network. In addition, we propose a warping field regularization which not only provides an explicit physical meaning to the gaze representations but also avoids redirection distortions. Promising results on few-shot gaze estimation (competitive results can be achieved with as few as ≤ 100 calibration samples), cross-dataset gaze estimation, gaze network pretraining, and another task (head pose estimation) demonstrate the validity of our framework.

1. Introduction

Gaze is a non-verbal cue with many functions. It can indicate attention, intentions, serve as communication cue in interactions, or even reveal higher level social constructs of people in relation with their personality. As such, it finds applications in many areas. For instance, it can be used in multi-party interaction analysis [2], Human-Robot-Interaction (HRI) for both floor control analysis and for robot behaviour synthesis to enable smooth interactions [1, 41]; in the virtual reality industry [3, 28], visual rendering can be improved by inferring the user gaze direction; in psychology, gaze behavior can contribute to mental health analysis and care [21, 49].

As with other computer vision tasks, the developments in deep neural networks have largely contributed to the progresses made on gaze estimation [4, 5, 37, 59, 60, 63]. For all these approaches, it is common sense that their performance depends to a large extent on the available amount of

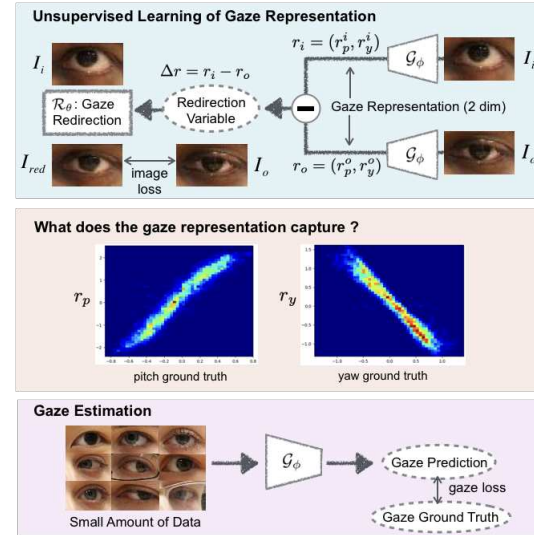


Figure 1. Proposed framework. *Top*: the networks \mathcal{G}_ϕ extracts gaze representations from two input eye images I_i and I_o . Their difference Δr is used as input to a gaze redirection network \mathcal{R}_θ along with the input image I_i to generate a redirected eye I_{red} which should be close to I_o . Both the \mathcal{G}_ϕ and \mathcal{R}_θ networks are trained jointly in an unsupervised fashion from unlabeled image pairs (I_i, I_o) . *Middle*. Thanks to our warping field regularization, the distribution of $(r_p$ vs pitch) and $(r_y$ vs yaw) exhibit high (almost linear) correlation. *Bottom*. The network \mathcal{G}_ϕ can further be used to train a gaze regressor.

data. Unfortunately, collecting and annotating 3D gaze data is complex and expensive, which introduces challenges and problems for gaze estimation, as summarized below:

- **Data amount.** The size of benchmark datasets [9, 11, 44, 45, 60, 61] including the number of people, is limited, making it difficult to train robust person-independent models. Synthetic data [53, 55] offers an alternative, but the domain gap is hard to eliminate.
- **Data annotation.** 3D gaze annotation can be noisy, due to (i) measurement errors: most datasets compute the 3D line of sight by visually estimating the 3D positions of eyes and gaze targets; (ii) participant distractions or blinks [7, 43], leading to wrong annotations.
- **Dataset bias.** Existing datasets rely on different cameras and setups, with important variations in visual

appearances (resolutions, lighting conditions). More importantly, they may only provide eye images obtained using different preprocessing techniques and gaze coordinate systems, making it almost impossible to merge datasets for training. It is thus hard to apply trained model to out-of-domain samples.

To address these challenges and lower the requirements for annotated gaze dataset, we propose an unsupervised approach which leverages large amounts of **unannotated eye images** for learning gaze representations, and only a **few calibration samples** to train a final gaze estimator. We show in experiments that with as low as 100 calibration samples we can already achieve competitive performances.

The main idea is illustrated in Fig. 1. The basis is a redirection network \mathcal{R}_θ which takes as input an eye image I_i as well as a gaze redirection variable Δr . It generates an output image I_{red} of the same eye but with the redirected gaze. In prior works [14, 20, 27, 59], Δr is explicitly set as a gaze offset, which means that gaze annotated images are required at training time (to set the gaze difference between I_i and the target output image I_o). In contrast, our method aims at using a network \mathcal{G}_ϕ to extract gaze representations from I_i and I_o and the simple representation difference provides the sufficient information required to do gaze retargeting. By imposing appropriate loss functions between the redirected output I_{red} and the target I_o , the framework can jointly train both the \mathcal{R}_θ and \mathcal{G}_ϕ networks from unlabelled images, implicitly enforcing the unsupervised learning of gaze representations. The middle part of Fig. 1 shows that this is achieved, as the 2-dimensional output of \mathcal{G}_ϕ is highly correlated (close to linear correlation) with groundtruth gaze angles. It is then possible to train a robust gaze estimator leveraging this representation. While investigating the above ideas, this paper makes the following contributions:

- **Unsupervised gaze representation learning.** We propose an approach to learn low dimensional gaze representations without gaze annotations, relying on a gaze redirection network and loss functions in image domain. To our best knowledge, this is the first work of unsupervised gaze representation learning.
- **Warping field regularization.** Similar to previous works, we rely on an inverse warping field w to perform gaze redirection. This paper proposed a warping field regularization which not only prevents possible overfitting or distortions, but also gives a physical meaning to the components of the learned unsupervised gaze representations.
- **Head pose extensions.** We also show that our unsupervised method is not limited to gaze estimation, but can also be used to process face images and learn a head pose related representation.

Experiments on three public datasets demonstrate the val-

idity of our approach, in particular when training with very few gaze calibrated datapoints and applying to cross-domain experiment (which shows that our method could successfully leverage large amount of Internet data to handle a much larger variety of eye shape, appearance, head poses, and illumination, ending in a more robust network for gaze representation extraction).

In the rest of the paper, we first summarize related works in Section 2. The method is detailed in Section 3. Section 4 explains our experiment protocol and reports our results. The conclusion is drawn in Section 5.

2. Related Work

Gaze estimation can be categorized into 3 classes, 2D Gaze Estimation, Gaze Following and 3D Gaze Estimation.

2D Gaze Estimation aims at predicting the 2D fixation point of gaze, e.g. on the screens of mobile devices [22, 29]. They usually rely on large datasets since annotating 2D gaze data is efficient. But it is hard to generalize a 2D gaze model to multiple devices or scenarios.

Gaze Following attempts to infer the object people are looking at. Recasens et al. proposed to use a saliency and a gaze pathways to predict the objects people look at in an image [39] or in a video [40]. Gaze following models tend to predict the head pose rather than the gaze, although recent works [6] attempted to jointly model gaze following and 3D gaze estimation, but without much improvement.

3D Gaze Estimation which retrieves the 3D line of sight of eyes is the main focus of this paper. Traditional approaches mainly include geometric based methods (GBM) and appearance based methods (ABM). GBM methods first extract features [12, 16, 17, 23, 47, 48, 50, 54, 56] from training images then estimate parameters of a geometric eye model which could be used to predict gaze. They usually require high resolution eye images and near frontal head poses, which limits its application scope. In contrast, ABM methods [13, 18, 22, 34–36, 45, 46] learn a direct mapping from eye appearance to the corresponding gaze.

ABM methods have attracted more attention in recent years with the development of deep learning. Zhang et al. [60] proposed a shallow network combining head pose along with extracted eye features, and later showed that a deeper network can further improve performance [62]. Moving beyond single eye gaze estimation, Cheng et al. [5] proposed to use two eyes while others relied on the full face, like in Zhang et al. [61] where a network process a full face but without using the head pose explicitly. Zhu et al. [63], however, proposed a geometric transformation layer to model the gaze and head pose jointly. Finally, as a recent trend, researchers start to work on building person-specific models from few reference samples to eliminate the person specific bias [4, 31–33, 37, 57, 59].

In general however, the performance of all the above

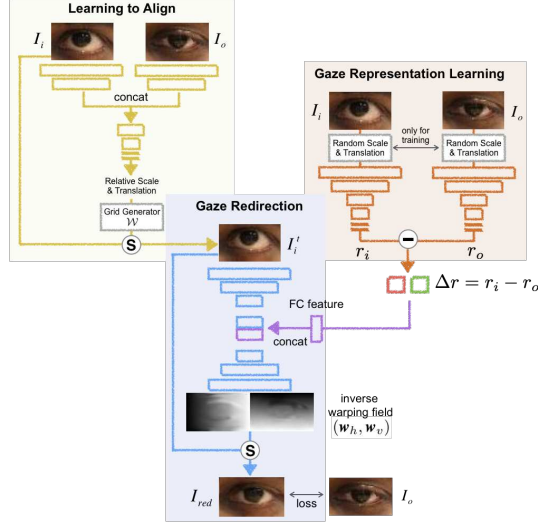


Figure 2. Unsupervised Learning of Gaze Representation.

models depends on the amount and diversity of training data. But as annotating 3D gaze is complex and expensive, it is difficult to collect data. Although synthetic data and domain adaptation [42, 51, 53, 55, 58] have been proposed, the domain gap between the synthetic data and real data is difficult to eliminate.

Representation Learning is also a topic related to our paper. Wiles et al. [26] proposed FAb-Net which learns a face embedding by retargetting the source face to a target face. The learned embedding encodes facial attributes like head pose and facial expression. Li et al. [30] later extended this work by disentangling the facial expression and the head motion through a TwinCycle Autoencoder. The training of the two approaches are conducted in a self-supervised way. Different from the above approaches which learn high dimensional embeddings with unclear physical meaning, our framework learns a low dimensional representations (2-Dim) with very clear meaning. Finally, following the face retargetting framework, an interesting gaze representation learning approach is proposed by Park et al. [37] where a face representation extracted from a bottleneck layer is disentangled as three components: appearance, gaze and head pose. The method however used a supervised approach for training, relying on head pose and gaze labels.

3. Method

3.1. Method Overview

The main idea behind our approach was introduced in Fig. 1: the aim is to jointly learn a representation network \mathcal{G}_ϕ and a redirection network \mathcal{R}_θ so that the difference $\Delta r = r_i - r_o = \mathcal{G}_\phi(I_i) - \mathcal{G}_\phi(I_o)$ between the extracted gaze representations indicates the gaze change to be used by the redirection network to generate a redirection image I_{red} which is as close as possible to I_o .

Our more detailed framework is shown in Fig. 2. Given a

training image pair (I_i, I_o) the network does the following. An alignment network \mathcal{A}_ψ aligns the input I_i to I_o using a global parametric motion model (translation and scale) according to: $I_i^t = \mathcal{A}_\psi(I_i, I_o)$. Then the redirection network takes this image as input, and produces a retargeted image as $I_{red} = \mathcal{R}_\theta(I_i^t, \Delta r)$, where as above Δr denotes the intended gaze change to be applied. In the following, we further motivate and detail our three networks. Then, we introduce the loss used for training the system, with a particular attention paid to the regularization of the warping field involved in the gaze redirection. Note that as a requirement for gaze redirection [59], the image pair (I_i, I_o) should be from the same person and share a similar head pose.

3.2. Gaze Representation Learning

The top right of Fig. 2 shows the architecture of the gaze representation learning module. It first extracts the gaze representations from the input images with \mathcal{G}_ϕ , a network based on ResNet blocks, and then computes the representation difference. In our approach, there are several elements which favor the learning of a gaze related representation rather than other information.

Gaze Representation r . We set r to be of dimension 2, which is motivated by two aspects. First, as the gaze direction is defined by the pitch and yaw angles, a 2D representation is enough. Secondly, a compact representation avoid the risk of capturing appearance information which should be extracted by the encoding part of the redirection network \mathcal{R}_θ from the input image. Otherwise, with a higher dimension, both \mathcal{R}_θ and \mathcal{G}_ϕ may encode eye appearance features, making the training of \mathcal{R}_θ and \mathcal{G}_ϕ less constrained.

Data Augmentation. To further enforce \mathcal{G}_ϕ to capture gaze-only information, we assume that the gaze representation should remain the same under small geometric perturbations. Thus, during training, we also apply random scaling and translation to the images before applying \mathcal{G}_ϕ . This data augmentation is a key step to achieve robust and accurate unsupervised gaze learning. Should this step be removed, \mathcal{G}_ϕ might learn to detect the pupil center position, which would be sufficient for the network \mathcal{R}_θ to generate a precise redirection output, but not be what we want. Thus, data augmentation enforces \mathcal{G}_ϕ to learn a scale and translation invariant representation, i.e. a gaze representation.

3.3. Global Alignment Network \mathcal{A}_ψ

As pointed in [59], training a gaze redirection network requires well aligned eye image pairs since global geometric transformation information can not be retrieved from the input image or the gaze difference. Hence, previous works used synthetic data [59] (domain adaptation required), landmark detection [60] or 3D head model [10] for eye alignment, which is not precise enough. Inspired by [24], we propose to learn to align an input image I_i with a target output I_o , as shown in the top left of Fig. 2. Concretely, an

alignment sub-network $f_{\mathcal{A}_\psi}$ takes I_i and I_o as input and predicts the motion parameters (translation and relative scale) between I_i and I_o . In the first few layers of $f_{\mathcal{A}_\psi}$, the two images are processed by separate network branches with shared weights. Then the extracted image features are concatenated and further processed to predict the geometric parameters. A grid generator \mathcal{W} [24] then converts these parameters into the inverse warping field transforming I_i into I_i^t (supposed to be aligned with I_o). The whole forward process can be formulated as:

$$I_i^t = \mathcal{A}_\psi(I_i, I_o) = I_i \circ \mathcal{W}(f_{\mathcal{A}_\psi}(I_i, I_o)) \quad (1)$$

where \circ denotes the grid sampling operator. Fig. 2 illustrate one alignment example, where I_i has been translated vertically to align with I_o .

3.4. Gaze Redirection Network \mathcal{R}_θ

The network \mathcal{R}_θ is shown in the bottom part of Fig. 2. The main part is an encoder-decoder network $f_{\mathcal{R}_\theta}$ trained to predict a warping field $\mathbf{w} = (\mathbf{w}_h, \mathbf{w}_v)$ which warps the (aligned) input I_i^t using a grid sampling operation [24] and synthesize a gaze redirection output I_{red} . In its bottleneck part, the network also receives feature maps generated from the retargeting gaze information Δr between I_i and I_o . As discussed in Sec. 3.2, the encoder of $f_{\mathcal{R}_\theta}$ ought to encode the eye structure (appearance) related information of I_i^t , while \mathcal{G}_ϕ (through Δr) should encode only the gaze change. The whole forward process can be summarized as:

$$I_{red} = \mathcal{R}_\theta(I_i^t, \Delta r) = I_i^t \circ f_{\mathcal{R}_\theta}(I_i^t, \mathcal{G}_\phi(I_i) - \mathcal{G}_\phi(I_o)) \quad (2)$$

3.5. Training Loss, Warping Field Regularization

The loss used to train the whole system is defined as a linear combination of several losses:

$$\mathcal{L} = \mathcal{L}_{img} + \lambda_w \mathcal{L}_w \text{ with } \mathcal{L}_{img} = \lambda_p \mathcal{L}_p + \lambda_f \mathcal{L}_f + \lambda_s \mathcal{L}_s \quad (3)$$

where \mathcal{L}_{img} is an image loss defined at the pixel (\mathcal{L}_p), feature (\mathcal{L}_f), and style (\mathcal{L}_s) levels, whereas \mathcal{L}_w is a regularization term on the warping field. In the following, we first introduce \mathcal{L}_{img} , and then emphasize the warping loss \mathcal{L}_w which plays an important role in our approach.

3.5.1 Image Loss \mathcal{L}_{img}

The main goal of the image loss is to measure the semantic difference between the generated image I_{red} and the target image I_o . It comprises three terms that we now describe.

Pixel Loss. It measures the discrepancy between I_{red} and I_o using a pixel level L1 loss (s_I denotes the image size).

$$\mathcal{L}_p = \frac{1}{s_I} \|I_{red} - I_o\|_1 \quad (4)$$

Perceptual Loss. \mathcal{L}_p is local and sensitive to illumination differences, and does not capture more structure and semantic information disparities. The latter and the robustness to illumination changes can be achieved using a perceptual

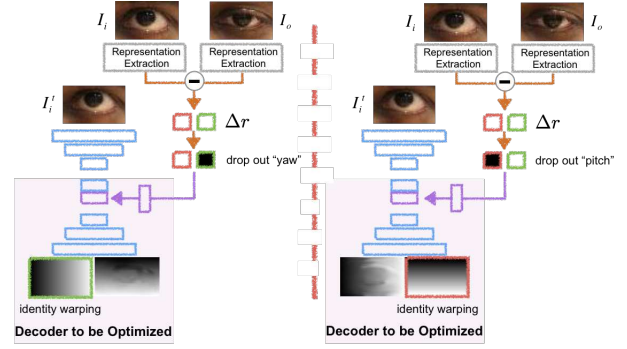


Figure 3. Representation drop out for warping field regularization.

loss comprising both feature and style reconstruction losses [25] which can be computed as follows. The I_{red} and I_o images are passed through a VGG16 network pretrained with ImageNet, from which we consider the features \mathbf{f}_j in the $j = 3, 8$, and 13^{th} layers. Accordingly, we can define the feature reconstruction loss \mathcal{L}_f as:

$$\mathcal{L}_f = \sum_j \frac{1}{c_j \cdot s_{f_j}} \|\mathbf{f}_j(I_{red}) - \mathbf{f}_j(I_o)\|_2 \quad (5)$$

in which s represents the spatial size of the feature maps and c the number of feature channels. To compute the style loss \mathcal{L}_s , the 3D feature maps \mathbf{f}_j are first reshaped into 2D matrices \mathbf{m}_j of size $c_j \times s_{f_j}$ from which we can compute the gram matrices \mathbf{g}_j (size $c_j \times c_j$), and then \mathcal{L}_s :

$$\mathbf{g}_j = \frac{1}{s_{f_j}} \mathbf{m}_j \cdot \mathbf{m}_j^T, \text{ and } \mathcal{L}_s = \sum_j \frac{1}{c_j^2} \|\mathbf{g}_j(I_{red}) - \mathbf{g}_j(I_o)\|_2. \quad (6)$$

3.5.2 Warping Field Regularization

Motivation. With the image loss \mathcal{L}_{img} alone, we are able to train the whole framework in an end-to-end fashion and achieve the unsupervised learning of gaze representation. However, the physical meaning of the gaze representation is not clear. In this section, we introduce a warping field regularization which not only gives physical meaning to the gaze representation but also regularizes the training of the whole framework (as shown in the experiments).

The main idea is to associate each gaze representation component with a specific warping field. Indeed, as shown in [58], the gaze yaw mainly involves an horizontal motion of the iris, and the pitch a vertical motion of the eyelid and iris. In other words, when there is only a yaw change (no pitch change), the vertical motion flow of eye region should be close to 0. Similarly, with only a pitch change, the horizontal flow should be close to 0. Note that no motion flow corresponds to an identity warping field.

Gaze Representation Dropout. To exploit the above assumption, we proposed the dropout mechanism illustrated in Fig. 3, in which we drop in turn each dimension of the gaze change Δr (setting it to 0), and enforce one of the warping field components to be an identity mapping \mathbf{w}_{id} , while keeping the other one unchanged.




Figure 4. Distribution of the $(pitch, r_p)$ and (yaw, r_y) data points (with $(yaw, pitch)$ being the ground truth gaze in head coordinate system (**HCS**)), for different training folds of a dataset (a) Eyediap. (b) Columbia Gaze (discrete and sparse gaze label). (c) UTMultiview.

More concretely, given a training image pair (I_i, I_o) , we first apply the forward pass and compute the representation difference as well as the warping field, according to:

$$[\Delta r_p, \Delta r_y] = \mathcal{G}(I_i) - \mathcal{G}(I_o), (\mathbf{w}_h, \mathbf{w}_v) = f_{\mathcal{R}}(I_i^t, [\Delta r_p, \Delta r_y]) \quad (7)$$

Then, we apply the dropout for each dimension, which results in the fields:

$$\begin{aligned} (\mathbf{w}_h^{\Delta r_y=0}, \mathbf{w}_v^{\Delta r_y=0}) &= f_{\mathcal{R}_{\theta^*}}(I_i^t, [\Delta r_p, 0]) \\ (\mathbf{w}_h^{\Delta r_p=0}, \mathbf{w}_v^{\Delta r_p=0}) &= f_{\mathcal{R}_{\theta^*}}(I_i^t, [0, \Delta r_y]) \end{aligned} \quad (8)$$

on which we apply the regularization loss:

$$\mathcal{L}_w = \frac{1}{S_I} (\|\mathbf{w}_h^{\Delta r_y=0} - \mathbf{w}_{id}\|_1 + \|\mathbf{w}_v^{\Delta r_y=0} - \mathbf{w}_v\|_1 + \|\mathbf{w}_h^{\Delta r_p=0} - \mathbf{w}_{id}\|_1 + \|\mathbf{w}_h^{\Delta r_p=0} - \mathbf{w}_h\|_1) \quad (9)$$

Note that for the dropout of each dimension, we not only enforce one field to be identity mapping, but also keep the other field unchanged since the other dimension of Δr is unchanged. In addition, note that for this regularization term, only the parameters θ^* of the decoder part of the redirection network are optimized, as shown in Fig. 3. This regularization is used along with the image loss when training the network (see Eq. 3). In essence, through this dropping and regularization mechanism, the network will be trained to associate the generation of one warping direction with one representation component, giving a physical meaning to the gaze representation. This is demonstrated in Fig. 4. In addition, as shown in Section 4.2, this regularization term also prevents potential overfitting or distortion of the warping field, leading to improved gaze redirection images and better gaze representations when used for gaze training.

3.6. Few-Shot Gaze Estimation

Linear Adaptation. To estimate the gaze (g_p, g_y) (in head coordinate system, **HCS**) from the unsupervised gaze representation (r_p, r_y) (also in **HCS**), we can first simply estimate two linear models (for pitch and yaw respectively):

$$g_p = k_p r_p + b_p, g_y = k_y r_y + b_y \quad (10)$$

using the very few calibration samples to rescale our representation, where k_p, b_p, k_y and b_y are model parameters.

Network Re-initialization and Finetuning. The second step is to fine-tune the network using the calibration samples. However, before doing this we re-initialized the

weight and bias ($\mathbf{k}^\phi = [k_p^\phi, k_y^\phi], \mathbf{b}^\phi = [b_p^\phi, b_y^\phi]$) of the last layer of gaze network \mathcal{G}_ϕ , using the above linear models, according to:

$$\begin{aligned} k_p(k_p^\phi \cdot \mathbf{x} + \mathbf{b}_p^\phi) + b_p &\rightarrow (k_p k_p^\phi) \cdot \mathbf{x} + (k_p b_p^\phi + b_p) \\ k_y(k_y^\phi \cdot \mathbf{x} + \mathbf{b}_y^\phi) + b_y &\rightarrow (k_y k_y^\phi) \cdot \mathbf{x} + (k_y b_y^\phi + b_y) \end{aligned} \quad (11)$$

where \mathbf{x} is feature forward to the last layer and $[k_p k_p^\phi, k_y k_y^\phi]$ and $[k_p b_p^\phi + b_p, k_y b_y^\phi + b_y]$ are the new weight and bias.

Gaze in World Coordinate System (WCS). To obtain a final gaze estimation in the **WCS**, the estimation in **HCS** is transformed using the head pose information.

3.7. Implementation Detail

Hyperparameters and Optimization. The framework is optimized by Adam with an initial learning rate of 10^{-4} and a small batch size of 16. 10 epochs are used to train the network and the learning rate is reduced by half every 3 epochs. The default values of loss weights $\lambda_p, \lambda_f, \lambda_s$ and λ_w are 1.0, 0.02, 0.1, 0.25 respectively. But for Eyediap samples which are blurry, we set λ_p to 0.2.

Activation Function. To bound the value of gaze representation when training begins, we used \tanh in the last layer of gaze network \mathcal{G}_ϕ . After 2 epochs, we removed the activation function, making the last layer a linear projection.

4. Experiment

4.1. Experiment Protocol

Dataset. We used three public datasets for experiment: Eyediap [11], Columbia Gaze [44] and UTMultiview [45]. Eyediap was collected with a RGBD sensor. It consists of sessions with different illumination conditions, gaze targets and head motion settings. We selected the session of HD video, condition B, floating target and static head pose for experiment, which results in 5 videos (5 subjects). Eye images were extracted and rectified to a frontal head pose [13]. Different from Eyediap, the gaze targets in Columbia Gaze and UTMultiview are discrete: only 7 horizontal and 3 vertical gaze directions in Columbia Gaze, 160 directions in UTMultiview where the gaze labels are further smoothed by a reconstruction-and-synthesis strategy. Both Columbia Gaze and UTMultiview use camera arrays to take multiview (head pose) samples. We use all available data for these two datasets (56 subjects and 50 subjects respectively).

Cross-Validation. For the 3 datasets, we perform $n =$

5, 5, 3-fold cross validation respectively (no subject overlap). In each fold, training data is used for unsupervised learning (without using gaze annotations) and then for few-shot gaze estimation by randomly selecting 10 to 100 samples with annotations. Test data is only used for evaluation. Please **note that this few shot setting is different from few shot personalization setting** as in [19, 37, 59], and all reported results in this paper are cross-subject.

Training Pair. For Columbia Gaze and UTMultiview, the image pairs (I_i, I_o) are randomly selected. However, for Eyediap which covers a larger gaze range, (I_i, I_o) are selected temporally by limiting their time gap within 10~20 frames. As already mentioned, I_i and I_o should be of the same person with a similar head pose. We used at most 200K pairs for training.

Network Models. Network details are given in the supplementary material. For the redirection network, it is based on ResNet blocks. Regarding the gaze network \mathcal{G}_ϕ , our default architecture *ResNet* is based on 4 ResNet blocks. For comparative experiments, we tested with *VGG16* pretrained with ImageNet (adopted in [62]) and with the *MnistNet* shallow architecture used in [60].

Tested Few-Shot Gaze Estimation Methods.

- ***U-LinFT***: our approach, consisting of unsupervised representation learning (U), linear adaptation (Lin) and network finetuning (FT), including re-initialization.
- ***U-Lin***: the same as above, but without network finetuning. A similar linear adaptation strategy was used in [32] (but this was for gaze personalization).
- ***U-SVR***: unsupervised representation learning followed by SVR adaptation. The SVR input features are the concatenation of the gaze representation and the output from the second last layer of \mathcal{G}_ϕ . A similar approach was used for gaze personalization in [29].
- ***DTrain***: a randomly initialized (or ImageNet pretrained for *VGG16*) baseline network directly trained with calibration samples.

Note that since to our best knowledge this is the first work to investigate unsupervised gaze representation learning and cross subject few-shot gaze estimation, it is difficult to find a state-of-the-art approach for comparison.

Performance Measure. We use the angle (in degree) between the estimated gaze and the ground truth gaze vectors as error measure. All reported results are the average of 10 runs (including random selection of calibration samples).

4.2. Qualitative results

Visualization of the Unsupervised Gaze Representation.

Fig. 4 shows the n distributions of unsupervised gaze representation w.r.t. ground truth. Each distribution corresponds to a gaze model obtained on the left-out folds of the datasets (see **Cross-Validation** above). As can be seen, these distributions are almost linear, validating the relevance of our ap-




Figure 5. Gaze redirection on (a) Eyediap. (b) Columbia Gaze. (c) UTMultiview. (d) Eyediap without the warping regularization loss. Each image triple represents: Left: input image, Middle: redirection output, Right: ground truth.

proach and warping regularization loss. An interesting point is that the gaze representation is inversely proportional to the ground truth sometimes, which might be due to the random factor during network initialization and training.

Gaze redirection. Fig. 5 illustrates the quality of our unsupervised gaze redirection results, where we remind that the gaze shift to be applied to the input image is provided by the representation difference obtained from the left and right images. As can be seen, our framework achieves accurate gaze redirection as well as eye alignment. Fig. 5(d) also demonstrates visually the overall benefit of our warping field regularization scheme.

4.3. Quantitative results

Few-Shot Gaze Estimation. The quantitative performances of few-shot gaze estimation approaches are reported in Fig. 6, where the results trained with all data and annotations (*DTrain (ResNet, full data)*) are plot as a lower bound. We can first notice that our approach (*U-LinFT*) achieves an acceptable accuracy ($7^\circ \sim 8^\circ$ error on all datasets) with only 100 calibration samples. In addition, all few-shot results based on our unsupervised learning (*U-LinFT*, *U-Lin*, *U-SVR*) are significantly better than the *DTrain* methods, including the *VGG* architecture pretrained on ImageNet. Furthermore, the performance of our approach with 10 calibration samples is still much better than the performance of *DTrain* methods with 100 samples.

U-LinFT performs very well on Eyediap and UTMultiview (only about 1° worse using 100 samples). In contrast, on Columbia Gaze, the performance gap is 1.9° . One possible reason regarding Eyediap is that its samples have been rectified to a frontal head pose and exhibit less variability, making our unsupervised gaze learning assumptions more valid. It is also reflected in Fig. 4 where the linear distributions from Eyediap are less dispersive. This implies that our approach could apply well to head mounted systems [15] where eyes have a fixed pose.

Amongst the unsupervised approaches, *U-LinFT* performs the best. *U-Lin* is in par using few calibration samples, but the performance gap increases with the number of samples since linear adaptation has fewer parameters. The *U-SVR* method is the worst when using few calibration




Figure 6. Few-shot gaze estimation results. (a) Eyediap. (b) Columbia Gaze. (c) UTMultiview.

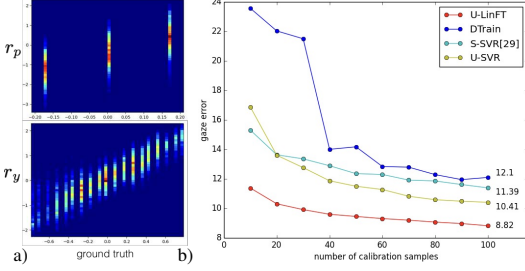


Figure 7. Cross dataset evaluation (unsupervised training on UTMultiview, test on Columbia Gaze) (a) Unsupervised gaze representation. (b) Few-shot gaze estimation.




Figure 8. Ablation study (Eyediap).




Figure 9. Few-shot gaze estimation with MnistNet.

samples because it has to train an SVR model from scratch, but it catches up as the number of samples increases.

Cross-Dataset Evaluation. We trained our unsupervised model on UTMultiview then tested it on Columbia Gaze (Columbia Gaze samples were converted to grayscale for consistency with UTMultiview). Fig. 7(a) visualizes the extracted representations vs ground truth distribution of the Columbia Gaze samples. They still follow a linear-like distribution. We then randomly select calibration samples from Columbia Gaze for few-shot gaze training. Results are reported in Fig. 7(b). Though we observe a small performance drop compared to results in Fig. 6(b), our unsupervised approaches are still much better than the *DTrain* method. More interestingly, we also trained a gaze estimator on UTMultiview in a supervised fashion (*DTrain (ResNet, full data)*) and adapted it on Columbia Gaze. This adaptation approach named *S-SVR* relies on an SVR model which uses features extracted from the last and the second last layer of the supervised model trained on UTMultiview. It was used in [29] (the original model only used features from the second last layer) for cross dataset experiment. Surprisingly, Fig. 7(b) shows that our unsupervised *U-SVR* (based on the same architecture but trained in an unsupervised fashion) is better than *S-SVR*, demonstrating that we achieved accurate unsupervised representation learning with good generalization capacity. These results show that our method can benefit from cross-domain data sources, and has the capacity to leverage large amount of Internet data to train robust models coping with diverse eye shapes, appearance, head poses, and illuminations.

Ablation Study. We study the impact of data augmenta-

tion and warping regularization by removing them. Note that as the physical meaning of the unsupervised representation is unclear when removing the warping regularization, we used a bilinear model to project the representation to gaze. Results are shown in Fig. 8. The performances without data augmentation or warping regularization are well below our proposed approach, but they remain better than *DTrain*. The gaze error increases by $\sim 3^\circ$ after removing data augmentation, showing that this is key to learn scale and translation invariant gaze representations. But removing the warping regularization leads to even more performance degradation. To explore the cause, some gaze redirection outputs without warping regularization are shown in Fig. 4(d). In the first three rows, skin pixels are re-projected to the sclera region because of wrong warping. In the last three rows, the outputs are totally distorted. This further demonstrates that the warping field regularization not only gives a physical meaning to the unsupervised gaze representations, but also prevents from possible distortions.

Shallow Architecture. As shallow networks can be of practical use in mobile devices, we tested our approach on the *MnistNet* gaze network \mathcal{G}_ϕ while keeping the same \mathcal{A}_ψ and \mathcal{R}_θ architectures. The performance is shown in Fig. 9. Compared with *ResNet*, the result of *MnistNet* is indeed worse. Nevertheless, we can notice that our approach *U-LinFT* works much better than the baseline *DTrain*, and that its performance is closer to the lower bound.

Unsupervised Learning for Pretraining. In this experiment, we use all the training data and their annotations to fine tune a model pretrained in an unsupervised fashion. Until now, tested architectures for \mathcal{G}_ϕ (ResNet based

Table 1. Gaze estimation with all the training data and annotations using *ResNet+HP* architecture. [†] $p < 0.01$

Error \ Dataset	Eyediap	Columbia Gaze	UTMultiview
Model			
U-Train	6.79	3.42	5.52
DTrain	7.09 [†]	3.63 [†]	5.72 [†]
Yu et al. [58]	8.5	-	5.7
Yu et al. [59] (cross subject)	-	3.54	-
Zhang et al. [60]	-	-	5.9
Park et al. [38]	-	3.59	-
Liu et al. [33] (cross subject)	-	-	5.95
Funes-Mora et al. [13]	11.6	-	-
Xiong et al. [57]	-	-	5.50 ± 1.03
Wang et al. [52]	-	-	5.4




Figure 10. Unsupervised head pose representation.

or *MnistNet*) were taking eye images as input and predicting gaze in *HCS* (final gaze obtained by transforming the estimate in *HCS* to *WCS* with the help of head pose). Such architectures are suitable for few-shot gaze estimation since the unsupervised gaze representation is also in *HCS*. However, when training with more data and annotations, a better strategy is to predict the gaze in *WCS* directly, by concatenating the convolution features with the head pose before the fully connected layers, as proposed in [62]. We denote this architecture as *ResNet+HP*. Due to this architecture difference in fully connected layers, we only use the convolutional layers of our pretrained *ResNet* to initialize *ResNet+HP*, and randomly initialized the fully connected layers which process the concatenated feature. Note that since in the Eyediap dataset eye samples are rectified to frontal head pose, we kept our *ResNet* architecture for Eyediap gaze prediction. Tab. 1 reports the results. As can be seen, using our unsupervised training leads to a performance gain of $0.2^\circ \sim 0.3^\circ$ compared to training from scratch. This is a small improvement, but given that results are based on 10 rounds of experiments, it is nevertheless stable and significant. Besides, we also compare our approach with SOTA results in Tab. 1. As can be seen, our approach is better or competitive. Please note that our Eyediap result can not be compared with SOTA directly since we used Eyediap session of HD video, condition B, floating target and static head pose while the listed SOTA used VGA video, condition A, floating target and static head pose. We can not find SOTA results with exactly the same session.

Application on Head Pose Estimation. We extended our approach to another task, head pose estimation. We used cropped faces from BIWI [8] for experiment and selected training pairs randomly within a temporal window. As head pose can be described by three rotation angles, pitch, yaw and roll, we used a 3 dimensional vector to represent it (instead of 2 dim for gaze). The second change we made con-

cerns the warping field regularization, where in order to relate the pitch with the vertical motion, we enforced the horizontal field to be the identity when dropping out the representation of yaw and roll; and similarly, when dropping out the pitch and roll, we defined a loss on the vertical field. The unsupervised head pose representation that was learned is illustrated in Fig. 10. The distribution of the pitch and yaw representations w.r.t the ground truth still exhibit a high correlation, but not so much for the roll representation. There might be two main reasons. First, none of our regularization terms involves the roll alone; second, the distribution of rolls in the data is uneven and concentrated, with 80% of them being within $-25^\circ \sim 5^\circ$. Although some future works could be done to improve the unsupervised learning of head pose representation, we demonstrated the potential of our framework for other tasks.

5. Conclusion and Discussion

We have proposed an unsupervised gaze learning framework which, to the best of our knowledge, is the first work on this topic. The two main contributing elements are the use of gaze redirection as auxiliary task for unsupervised learning, and the exploitation of a warping field regularization scheme which not only provides a physical meaning to the learned gaze representation dimensions, but also prevents overfitting or gaze redirection distortions. We demonstrate promising results on few-shot gaze estimation, network pretraining, and cross-dataset experiments in which the gaze representation (and network) learned in an unsupervised fashion proved to be better than a network trained supervisedly with gaze data. In this view, we believe that our method could successfully leverage internet data to train an unsupervised network robust to a large variety of eye shapes, appearance, head poses, and illumination.

Our work can be expanded along two interesting directions. First, our work can be used for few shot person-specific gaze estimation. We believe large performance improvement can be achieved since there is less bias or noise among person specific samples. It is especially beneficial for few shot gaze estimation (less bias in few data). Second, given the unsupervised performance on head pose estimation, our work can also be used for full face gaze estimation. The framework can be implemented by using two network branches, as done in [63].

Notwithstanding the above advantages, a limitation of our method is that it requires image pairs with the same or at least close head poses for training, setting some constraints on our approach. We leave the evaluation of this requirement as a future work.

Acknowledgement. This work was funded by the European Unions Horizon 2020 research and innovation programme -grant agreement no. 688147 (MuMMER, mummer-project.eu).

[†] indicates an error significantly higher than our method ($p < 0.01$).

References

- [1] Sean Andrist, Xiang Zhi Tan, Michael Gleicher, and Bilge Mutlu. Conversational gaze aversion for humanlike robots. In *Proceedings of the 2014 ACM/IEEE International Conference on Human-robot Interaction, HRI '14*, pages 25–32, New York, NY, USA, 2014. ACM. 1
- [2] S. Ba and J-M. Odobez. Multi-party focus of attention recognition in meetings from head pose and multimodal contextual cues. In *IEEE Int. Conf. on Acoustics, Speech, and Signal Processing (ICASSP)*, Las-Vegas, march 2008. 1
- [3] Meixu Chen, Yize Jin, Todd Goodall, Xiangxu Yu, and Alan C. Bovik. Study of 3d virtual reality picture quality, 2019. 1
- [4] Zhaokang Chen and Bertram E. Shi. Appearance-based gaze estimation via gaze decomposition and single gaze point calibration, 2019. 1, 2
- [5] Yihua Cheng, Feng Lu, and Xucong Zhang. Appearance-based gaze estimation via evaluation-guided asymmetric regression. In *The European Conference on Computer Vision (ECCV)*, September 2018. 1, 2
- [6] Eunji Chong, Nataniel Ruiz, Yongxin Wang, Yun Zhang, Agata Rozga, and James M. Rehg. Connecting gaze, scene, and attention: Generalized attention estimation via joint modeling of gaze and scene saliency. In *The European Conference on Computer Vision (ECCV)*, September 2018. 2
- [7] Kevin Cortacero, Tobias Fischer, and Yiannis Demiris. Rt-bene: A dataset and baselines for real-time blink estimation in natural environments. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, 2019. 1
- [8] Gabriele Fanelli, Matthias Dantone, Juergen Gall, Andrea Fossati, and Luc Van Gool. Random forests for real time 3d face analysis. *Int. J. Comput. Vision*, 101(3):437–458, February 2013. 8
- [9] Tobias Fischer, Hyung Jin Chang, and Yiannis Demiris. RT-GENE: Real-Time Eye Gaze Estimation in Natural Environments. *European Conference on Computer Vision (ECCV)*, 2018. 1
- [10] K. Funes and J.-M. Odobez. Gaze estimation from multimodal kinect data. In *CVPR Workshop on Face and Gesture and Kinect demonstration competition (Best Student Paper Award)*, Providence, june 2012. 3
- [11] Kenneth Alberto Funes Mora, Florent Monay, and Jean-Marc Odobez. Eyediap: A database for the development and evaluation of gaze estimation algorithms from rgb and rgb-d cameras. In *Proceedings of the Symposium on Eye Tracking Research and Applications (ETRA)*, pages 255–258, 2014. 1, 5
- [12] Kenneth Alberto Funes Mora and Jean-Marc Odobez. Geometric Generative Gaze Estimation (G3E) for Remote RGB-D Cameras. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1773–1780, jun 2014. 2
- [13] Kenneth A. Funes-Mora and Jean-Marc Odobez. Gaze estimation in the 3d space using rgb-d sensors, towards head-pose and user invariance. *International Journal of Computer Vision (IJCV)*, 118(2):194–216, 2016. 2, 5, 8
- [14] Yaroslav Ganin, Daniil Kononenko, Diana Sungatullina, and Victor Lempitsky. DeepWarp: Photorealistic image resynthesis for gaze manipulation. *European Conference on Computer Vision (ECCV)*, pages 311–326, 2016. 2
- [15] Stephan J. Garbin, Yiru Shen, Immo Schuetz, Robert Cavin, Gregory Hughes, and Sachin S. Talathi. Openeds: Open eye dataset. *CoRR*, abs/1905.03702, 2019. 6
- [16] Chao Gou, Yue Wu, Kang Wang, Fei-Yue Wang, and Qiang Ji. Learning-by-synthesis for accurate eye detection. *ICPR*, 2016. 2
- [17] Chao Gou, Yue Wu, Kang Wang, Kunfeng Wang, Fei Yue Wang, and Qiang Ji. A joint cascaded framework for simultaneous eye detection and eye state estimation. *Pattern Recognition*, 67:23–31, 2017. 2
- [18] Dan Witzner Hansen and Qiang Ji. In the eye of the beholder: A survey of models for eyes and gaze. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 32(3):478–500, 2010. 2
- [19] Junfeng He, Khoi Pham, Nachiappan Valliappan, Pingmei Xu, Chase Roberts, Dmitry Lagun, and Vidhya Navalpakkam. On-device few-shot personalization for real-time gaze estimation. In *The IEEE International Conference on Computer Vision (ICCV) Workshops*, Oct 2019. 6
- [20] Zhe He, Adrian Spurr, Xucong Zhang, and Otmar Hilliges. Photo-realistic monocular gaze redirection using generative adversarial networks. In *The IEEE International Conference on Computer Vision (ICCV)*, October 2019. 2
- [21] Michael Xuelin Huang, Jiajia Li, Grace Ngai, and Hong Va Leong. Stressclick: Sensing stress from gaze-click patterns. In *Proceedings of the ACM on Multimedia Conference (ACMMM)*, pages 1395–1404, 2016. 1
- [22] Qiong Huang, Ashok Veeraraghavan, and Ashutosh Sabharwal. Tablet gaze: unconstrained appearance-based gaze estimation in mobile tablets. *Machine Vision and Applications (MVAP)*, 2017. 2
- [23] Takahiro Ishikawa. Passive driver gaze tracking with active appearance models. 2004. 2
- [24] Max Jaderberg, Karen Simonyan, Andrew Zisserman, and koray kavukcuoglu. Spatial transformer networks. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 2017–2025. 2015. 3, 4
- [25] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *European Conference on Computer Vision*, 2016. 4
- [26] A. Sophia Koepke, Olivia Wiles, and Andrew Zisserman. Self-supervised learning of a facial attribute embedding from video. In *British Machine Vision Conference 2018, BMVC 2018, Northumbria University, Newcastle, UK, September 3-6, 2018*, page 302, 2018. 3
- [27] Daniil Kononenko, Yaroslav Ganin, Diana Sungatullina, and Victor S. Lempitsky. Photorealistic Monocular Gaze Redirection Using Machine Learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, pages 1–15, 2017. 2
- [28] Robert Konrad, Anastasios Angelopoulos, and Gordon Wetzstein. Gaze-contingent ocular parallax rendering for virtual reality, 2019. 1

- [29] Kyle Krafka, Aditya Khosla, Petr Kellnhofer, and Harini Kannan. Eye Tracking for Everyone. *IEEE Conference on Computer Vision and Pattern Recognition*, pages 2176–2184, 2016. 2, 6, 7
- [30] Yong Li, Jiabei Zeng, Shiguang Shan, and Xilin Chen. Self-supervised representation learning from videos for facial action unit detection. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. 3
- [31] Erik Lindén, Jonas Sjöstrand, and Alexandre Proutiere. Learning to personalize in appearance-based gaze tracking, 2018. 2
- [32] Gang Liu, Yu Yu, Kenneth Alberto Funes-Mora, and Jean-Marc Odobez. A Differential Approach for Gaze Estimation with Calibration. *British Machine Vision Conference (BMVC)*, 2018. 2, 6
- [33] Gang Liu, Yu Yu, Kenneth Alberto Funes Mora, and Jean-Marc Odobez. A differential approach for gaze estimation. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 2020. 2, 8
- [34] Feng Lu, Yusuke Sugano, Takahiro Okabe, and Yoichi Sato. Inferring human gaze from appearance via adaptive linear regression. *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pages 153–160, 2011. 2
- [35] Francis Martinez, Andrea Carbone, and Edwige Pissaloux. Gaze estimation using local features and non-linear regression. pages 1961–1964, 2012. 2
- [36] Basilio Noris, Jean-Baptiste Keller, and Aude Billard. A wearable gaze tracking system for children in unconstrained environments. *Computer Vision and Image Understanding*, 115(4):476–486, 2011. 2
- [37] Seonwook Park, Shalini De Mello, Pavlo Molchanov, Umar Iqbal, Otmar Hilliges, and Jan Kautz. Few-shot adaptive gaze estimation, 2019. 1, 2, 3, 6
- [38] Seonwook Park, Adrian Spurr, and Otmar Hilliges. Deep Pictorial Gaze Estimation. In *European Conference on Computer Vision (ECCV)*, pages 741–757. 2018. 8
- [39] Adria Recasens, Aditya Khosla, Carl Vondrick, and Antonio Torralba. Where are they looking? *Advances in Neural Information Processing Systems*, pages 199–207, 2015. 2
- [40] A. Recasens, C. Vondrick, A. Khosla, and A. Torralba. Following gaze in video, Oct 2017. 2
- [41] S. Sheikhi and J.M. Odobez. Combining dynamic head pose and gaze mapping with the robot conversational state for attention recognition in human-robot interactions. *Pattern Recognition Letters*, 66:81–90, Nov. 2015. 1
- [42] Ashish Shrivastava, Tomas Pfister, Oncel Tuzel, Josh Susskind, Wenda Wang, and Russ Webb. Learning from simulated and unsupervised images through adversarial training. *CoRR*, abs/1612.07828, 2016. 3
- [43] Rémy Siegfried, Yu Yu, and Jean-Marc Odobez. Towards the use of social interaction conventions as prior for gaze model adaptation. In *Proceedings of the 19th ACM International Conference on Multimodal Interaction, ICMI '17*, pages 154–162, New York, NY, USA, 2017. ACM. 1
- [44] Brian A Smith, Qi Yin, Steven K Feiner, and Shree K Nayar. Gaze locking: passive eye contact detection for human-object interaction. In *Proceedings of the 26th annual ACM symposium on User interface software and technology (UIST)*, pages 271–280, 2013. 1, 5
- [45] Yusuke Sugano, Yasuyuki Matsushita, and Yoichi Sato. Learning-by-synthesis for appearance-based 3d gaze estimation. In *Computer Vision and Pattern Recognition (CVPR), 2014 IEEE Conference on*, pages 1821–1828. IEEE, 2014. 1, 2, 5
- [46] Kar-Han Tan, David J Kriegman, and Narendra Ahuja. Appearance-based eye gaze estimation. pages 191–195, 2002. 2
- [47] Fabian Timm and Erhardt Barth. Accurate eye centre localisation by means of gradients. In *Proceedings of the Int. Conference on Computer Theory and Applications (VISAPP)*, volume 1, pages 125–130, Algarve, Portugal, 2011. INSTICC. 2
- [48] Ronda Venkateswarlu et al. Eye gaze estimation from a single image of one eye. pages 136–143, 2003. 2
- [49] Mélodie Vidal, Jayson Turner, Andreas Bulling, and Hans Gellersen. Wearable eye tracking for mental health monitoring. *Computer Communications*, 35(11):1306–1311, 2012. 1
- [50] Arantxa Villanueva, Victoria Ponz, Laura Sesma-Sanchez, Mikel Ariz, Sonia Porta, and Rafael Cabeza. Hybrid method based on topography for robust detection of iris center and eye corners. *ACM Trans. Multimedia Comput. Commun. Appl.*, 9(4), 2013. 2
- [51] Kang Wang. A Hierarchical Generative Model for Eye Image Synthesis and Eye Gaze Estimation. 2018. 3
- [52] Kang Wang, Rui Zhao, Hui Su, and Qiang Ji. Generalizing eye tracking with bayesian adversarial learning. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. 8
- [53] E Wood, T Baltruaitis, X Zhang, Y Sugano, P Robinson, and A Bulling. Rendering of Eyes for Eye-Shape Registration and Gaze Estimation. *IEEE International Conference on Computer Vision (ICCV)*, pages 3756–3764, 2015. 1, 3
- [54] Erroll Wood, Tadas Baltrušaitis, Louis Philippe Morency, Peter Robinson, and Andreas Bulling. A 3D morphable eye region model for gaze estimation. *European Conference on Computer Vision (ECCV)*, pages 297–313, 2016. 2
- [55] Erroll Wood, Tadas Baltrušaitis, Louis-Philippe Morency, Peter Robinson, and Andreas Bulling. Learning an appearance-based gaze estimator from one million synthesised images. In *Proceedings of the Ninth Biennial ACM Symposium on Eye Tracking Research & Applications (ETRA)*, pages 131–138, 2016. 1, 3
- [56] Erroll Wood and A Bulling. Eytetab: Model-based gaze estimation on unmodified tablet computers. *ACM Symposium on Eye Tracking Research & Applications (ETRA)*, pages 3–6, 2014. 2
- [57] Yunyang Xiong, Hyunwoo J. Kim, and Vikas Singh. Mixed effects neural networks (menets) with applications to gaze estimation. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. 2, 8
- [58] Yu Yu, Gang Liu, and Jean-Marc Odobez. Deep multitask gaze estimation with a constrained landmark-gaze model. *European Conference on Computer Vision Workshop (ECCVW)*, 2018. 3, 4, 8

- [59] Yu Yu, Gang Liu, and Jean-Marc Odobez. Improving few-shot user-specific gaze adaptation via gaze redirection synthesis. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019. 1, 2, 3, 6, 8
- [60] Xucong Zhang, Yusuke Sugano, Mario Fritz, and Andreas Bulling. Appearance-based gaze estimation in the wild. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4511–4520, jun 2015. 1, 2, 3, 6, 8
- [61] Xucong Zhang, Yusuke Sugano, Mario Fritz, and Andreas Bulling. It’s Written All Over Your Face: Full-Face Appearance-Based Gaze Estimation. *IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2016. 1, 2
- [62] Xucong Zhang, Yusuke Sugano, Mario Fritz, and Andreas Bulling. MPIIGaze: Real-World Dataset and Deep Appearance-Based Gaze Estimation. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, pages 1–14, 2017. 2, 6, 8
- [63] Wangjiang Zhu and Haoping Deng. Monocular free-head 3d gaze tracking with deep learning and geometry constraints. In *The IEEE International Conference on Computer Vision (ICCV)*, Oct 2017. 1, 2, 8