

Recurrent CNN for 3D Gaze Estimation using Appearance and Shape Cues

Cristina Palmero^{1,2}

crpalmec7@alumnes.ub.edu

Javier Selva¹

javier.selva.castello@est.fib.upc.edu

Mohammad Ali Bagheri^{3,4}

mohammadali.bagheri@ucalgary.ca

Sergio Escalera^{1,2}

sergio@maia.ub.es

¹ Dept. Mathematics and Informatics
Universitat de Barcelona, Spain

² Computer Vision Center
Campus UAB, Bellaterra, Spain

³ Dept. Electrical and Computer Eng.
University of Calgary, Canada

⁴ Dept. Engineering
University of Larestan, Iran

Abstract

Gaze behavior is an important non-verbal cue in social signal processing and human-computer interaction. In this paper, we tackle the problem of person- and head pose-independent 3D gaze estimation from remote cameras, using a multi-modal recurrent convolutional neural network (CNN). We propose to combine face, eyes region, and face landmarks as individual streams in a CNN to estimate gaze in still images. Then, we exploit the dynamic nature of gaze by feeding the learned features of all the frames in a sequence to a many-to-one recurrent module that predicts the 3D gaze vector of the last frame. Our multi-modal static solution is evaluated on a wide range of head poses and gaze directions, achieving a significant improvement of 14.6% over the state of the art on EYEDIAP dataset, further improved by 4% when the temporal modality is included.

1 Introduction

Eyes and their movements are considered an important cue in non-verbal behavior analysis, being involved in many cognitive processes and reflecting our internal state [1]. More specifically, eye gaze behavior, as an indicator of human visual attention, has been widely studied to assess communication skills [2] and to identify possible behavioral disorders [3]. Therefore, gaze estimation has become an established line of research in computer vision, being a key feature in human-computer interaction (HCI) and usability research [4, 5].

Recent gaze estimation research has focused on facilitating its use in general everyday applications under real-world conditions, using off-the-shelf remote RGB cameras and removing the need of personal calibration [6]. In this setting, appearance-based methods, which learn a mapping from images to gaze directions, are the preferred choice [7]. However, they need large amounts of training data to be able to generalize well to in-the-wild situations, which are characterized by significant variability in head poses, face appearances and lighting conditions. In recent years, CNNs have been reported to outperform classical methods. However, most existing approaches have only been tested in restricted HCI tasks,

Method	3D gaze direction	Unrestricted gaze target	Full face	Eye region	Facial landmarks	Sequential information
Zhang <i>et al.</i> (1) [42]	✓	✗	✗	✓	✗	✗
Krafka <i>et al.</i> [46]	✗	✗	✓	✓	✗	✗
Zhang <i>et al.</i> (2) [43]	✓	✗	✓	✗	✗	✗
Deng and Zhu [9]	✓	✓	✓	✓	✗	✗
Ours	✓	✓	✓	✓	✓	✓

Table 1: Characteristics of recent related work on person- and head pose-independent appearance-based gaze estimation methods using CNNs.

where users look at the screen or mobile phone, showing a low head pose variability. It is yet unclear how these methods would perform in a wider range of head poses.

On a different note, until very recently, the majority of methods only used static eye region appearance as input. State-of-the-art approaches have demonstrated that using the face along with a higher resolution image of the eyes [46], or even just the face itself [43], increases performance. Indeed, the whole-face image encodes more information than eyes alone, such as illumination and head pose. Nevertheless, gaze behavior is not static. Eye and head movements allow us to direct our gaze to target locations of interest. It has been demonstrated that humans can better predict gaze when being shown image sequences of other people moving their eyes [9]. However, it is still an open question whether this sequential information can increase the performance of automatic methods.

In this work, we show that the combination of multiple cues benefits the gaze estimation task. In particular, we use face, eye region and facial landmarks from still images. Facial landmarks model the global shape of the face and come at no cost, since face alignment is a common pre-processing step in many facial image analysis approaches. Furthermore, we present a subject-independent, free-head recurrent 3D gaze regression network to leverage the temporal information of image sequences. The static streams of each frame are combined in a late-fusion fashion using a multi-stream CNN. Then, all feature vectors are input to a many-to-one recurrent module that predicts the gaze vector of the last sequence frame.

In summary, our contributions are two-fold. First, we present a Recurrent-CNN network architecture that combines appearance, shape and temporal information for 3D gaze estimation. Second, we test static and temporal versions of our solution on the EYEDIAP dataset [9] in a wide range of head poses and gaze directions, showing consistent performance improvements compared to related appearance-based methods. To the best of our knowledge, this is the first third-person, remote camera-based approach that uses temporal information for this task. Table 1 outlines our main method characteristics compared to related work. Models and code are publicly available at <https://github.com/crisie/RecurrentGaze>.

2 Related work

Gaze estimation methods are typically categorized as model-based or appearance-based [6, 10, 15]. **Model-based approaches** use a geometric model of the eye, usually requiring either high resolution images or a person-specific calibration stage to estimate personal eye parameters [22, 53, 54, 57, 40]. In contrast, **appearance-based methods** learn a direct mapping from intensity images or extracted eye features to gaze directions, thus being

potentially applicable to relatively low resolution images and mid-distance scenarios. Different mapping functions have been explored, such as neural networks [4], adaptive linear regression (ALR) [19], local interpolation [52], gaussian processes [60, 65], random forests [11, 61], or k-nearest neighbors [40]. Main challenges of appearance-based methods for 3D gaze estimation are head pose, illumination and subject invariance without user-specific calibration. To handle these issues, some works proposed compensation methods [18] and warping strategies that synthesize a canonical, frontal looking view of the face [6, 13, 21]. Hybrid approaches based on analysis-by-synthesis have also been evaluated [69].

Currently, data-driven methods are considered the state of the art for person- and head pose-independent appearance-based gaze estimation. Consequently, a number of gaze estimation datasets have been introduced in recent years, either in controlled [29] or semi-controlled settings [8], in the wild [16, 42], or consisting of synthetic data [51, 58, 40]. Zhang *et al.* [42] showed that CNNs can outperform other mapping methods, using a multi-modal CNN to learn the mapping from 3D head poses and eye images to 3D gaze directions. Krafka *et al.* [16] proposed a multi-stream CNN for 2D gaze estimation, using individual eye, whole-face image and the face grid as input. As this method was limited to 2D screen mapping, Zhang *et al.* [42] later explored the potential of just using whole-face images as input to estimate 3D gaze directions. Using a spatial weights CNN, they demonstrated their method to be more robust to facial appearance variation caused by head pose and illumination than eye-only methods. While the method was evaluated in the wild, the subjects were only interacting with a mobile device, thus restricting the head pose range. Deng and Zhu [9] presented a two-stream CNN to disjointly model head pose from face images and eyeball movement from eye region images. Both were then aggregated into 3D gaze direction using a gaze transform layer. The decomposition was aimed to avoid head-correlation overfitting of previous data-driven approaches. They evaluated their approach in the wild with a wider range of head poses, obtaining better performance than previous eye-based methods. However, they did not test it on public annotated benchmark datasets.

In this paper, we propose a multi-stream recurrent CNN network for person- and head pose-independent 3D gaze estimation for a mid-distance scenario. We evaluate it on a wider range of head poses and gaze directions than screen-targeted approaches. As opposed to previous methods, we also rely on temporal information inherent in sequential data.

3 Methodology

In this section, we present our approach for 3D gaze regression based on appearance and shape cues for still images and image sequences. First, we introduce the data modalities and formulate the problem. Then, we detail the normalization procedure prior to the regression stage. Finally, we explain the global network topology as well as the implementation details. An overview of the system architecture is depicted in Figure 1.

3.1 Multi-modal gaze regression

Let us represent gaze direction as a 3D unit vector $\mathbf{g} = [g_x, g_y, g_z]^T \in \mathbb{R}^3$ in the Camera Coordinate System (CCS), whose origin is the central point between eyeball centers. Assuming a calibrated camera, and a known head position and orientation, our goal is to estimate \mathbf{g} from a sequence of images $\{\mathbf{I}^{(i)} \mid \mathbf{I} \in \mathbb{R}^{W \times H \times 3}\}$ as a regression problem.




Figure 1: Overview of the proposed network. A multi-stream CNN jointly models full-face, eye region appearance and face landmarks from still images. The combined extracted features from each frame are fed into a recurrent module to predict last frame’s gaze direction.

Gazing to a specific target is achieved by a combination of eye and head movements, which are highly coordinated. Consequently, the apparent direction of gaze is influenced not only by the location of the irises within the eyelid aperture, but also by the position and orientation of the face with respect to the camera. Known as the Wollaston effect [66], the exact same set of eyes may appear to be looking in different directions due to the surrounding facial cues. It is therefore reasonable to state that eye images are not sufficient to estimate gaze direction. Instead, whole-face images can encode head pose or illumination-specific information across larger areas than those available just in the eyes region [16, 43].

The drawback of appearance-only methods is that global structure information is not explicitly considered. In that sense, facial landmarks can be used as global shape cues to encode spatial relationships and geometric constraints. Current state-of-the-art face alignment approaches are robust enough to handle large appearance variability, extreme head poses and occlusions, being especially useful when the dataset used for gaze estimation does not contain such variability. Facial landmarks are mainly correlated with head orientation, eye position, eyelid openness, and eyebrow movement, which are valuable features for our task.

Therefore, in our approach we jointly model appearance and shape cues (see Figure 1). The former is represented by a whole-face image \mathbf{I}_F , along with a higher resolution image of the eyes \mathbf{I}_E to identify subtle changes. Due to dealing with wide head pose ranges, some eye images may not depict the whole eye, containing mostly background or other surrounding facial parts instead. For that reason, and contrary to previous approaches that only use one eye image [31, 42], we use a single image composed of two patches of centered left and right eyes. Finally, the shape cue is represented by 3D face landmarks obtained from a 68-landmark model, denoted by $\mathbf{L} = \{(l_x, l_y, l_z)_c \mid \forall c \in [1, \dots, 68]\}$.

In this work we also consider the dynamic component of gaze. We leverage the sequential information of eye and head movements such that, given appearance and shape features of consecutive frames, it is possible to better predict the gaze direction of the current frame. Therefore, the 3D gaze estimation task for a 1-frame sequence is formulated

as $\mathbf{g}^{(i)} = f\left(\{\mathbf{I}_F^{(i)}\}, \{\mathbf{I}_E^{(i)}\}, \{\mathbf{L}^{(i)}\}\right)$, where i denotes the i -th frame, and f is the regression function.

3.2 Data normalization

Prior to gaze regression, a normalization step in the 3D space and the 2D image, similar to [80], is carried out. This is performed to reduce the appearance variability and to allow the gaze estimation model to be applied regardless of the original camera configuration.

Let $\mathbf{H} \in \mathbb{R}^{3 \times 3}$ be the head rotation matrix, and $\mathbf{p} = [p_x, p_y, p_z]^T \in \mathbb{R}^3$ the reference face location with respect to the original CCS. The goal is to find the conversion matrix $\mathbf{M} = \mathbf{S}\mathbf{R}$ such that (a) the X -axes of the virtual camera and the head become parallel using the rotation matrix \mathbf{R} , and (b) the virtual camera looks at the reference location from a fixed distance d_n using the Z -direction scaling matrix $\mathbf{S} = \text{diag}(1, 1, d_n/\|\mathbf{p}\|)$. \mathbf{R} is computed as $\mathbf{a} = \hat{\mathbf{p}} \times \mathbf{H}^T \mathbf{e}_1$, $\mathbf{b} = \hat{\mathbf{a}} \times \hat{\mathbf{p}}$, $\mathbf{R} = [\hat{\mathbf{a}}, \hat{\mathbf{b}}, \hat{\mathbf{p}}]^T$, where \mathbf{e}_1 denotes the first orthonormal basis and $\langle \hat{\cdot} \rangle$ is the unit vector.

This normalization translates into the image space as a cropped image patch of size $W_n \times H_n$ centered at \mathbf{p} where head roll rotation has been removed. This is done by applying a perspective warping to the input image \mathbf{I} using the transformation matrix $\mathbf{W} = \mathbf{C}_o \mathbf{M} \mathbf{C}_n^{-1}$, where \mathbf{C}_o and \mathbf{C}_n are the original and virtual camera matrices, respectively.

The 3D gaze vector is also normalized as $\mathbf{g}_n = \mathbf{R}\mathbf{g}$. After image normalization, the line of sight can be represented in a 2D space. Therefore, \mathbf{g}_n is further transformed to spherical coordinates (θ, ϕ) assuming unit length, where θ and ϕ denote the horizontal and vertical direction angles, respectively. This 2D angle representation, delimited in the range $[-\pi/2, \pi/2]$, is computed as $\theta = \arctan(g_x/g_z)$ and $\phi = \arcsin(-g_y)$, such that $(0, 0)$ represents looking straight ahead to the CCS origin.

3.3 Recurrent Convolutional Neural Network

We propose a Recurrent CNN Regression Network for 3D gaze estimation. The network is divided in 3 modules: (1) *Individual*, (2) *Fusion*, and (3) *Temporal*.

First, the *Individual* module learns features from each appearance cue separately. It consists of a two-stream CNN, one devoted to the normalized face image stream and the other to the joint normalized eyes image. Next, the *Fusion* module combines the extracted features of each appearance stream in a single vector along with the normalized landmark coordinates. Then, it learns a joint representation between modalities in a late-fusion fashion. Both *Individual* and *Fusion* modules, further referred to as *Static* model, are applied to each frame of the sequence. Finally, the resulting feature vectors of each frame are input to the *Temporal* module based on a many-to-one recurrent network. This module leverages sequential information to predict the normalized 2D gaze angles of the last frame of the sequence using a linear regression layer added on top of it.

3.4 Implementation details

3.4.1 Network details

Each stream of the *Individual* module is based on the VGG-16 deep network [27], consisting of 13 convolutional layers, 5 max pooling layers, and 1 fully connected (FC) layer with Rectified Linear Unit (ReLU) activations. The full-face stream follows the same configuration

as the base network, having an input of 224×224 pixels and a 4096D FC layer. In contrast, the input joint eye image is smaller, with a final size of 120×48 pixels, so the number of parameters is decreased proportionally. In this case, its last FC layer produces a 1536D vector. A 204D landmark coordinates vector is concatenated to the output of the FC layer of each stream, resulting in a 5836D feature vector. Consequently, the *Fusion* module consists of 2 5836D FC layers with ReLU activations and 2 dropout layers between FCs as regularization. Finally, to model the temporal dependencies, we use a single GRU layer with 128 units.

The network is trained in a stage-wise fashion. First, we train the *Static* model and the final regression layer end-to-end on each individual frame of the training data. The convolutional blocks are pre-trained with the VGG-Face dataset [27], whereas the FCs are trained from scratch. Second, the training data is re-arranged by means of a sliding window with stride 1 to build input sequences. Each sequence is composed of $s = 4$ consecutive frames, whose gaze direction target is the gaze direction of the last frame of the sequence $(\{\mathbf{I}^{(i-s+1)}, \dots, \mathbf{I}^{(i)}\}, \mathbf{g}^{(i)})$. Using this re-arranged training data, we extract features of each frame of the sequence from a frozen *Individual* module, fine-tune the *Fusion* layers, and train both, the *Temporal* module and a new final regression layer from scratch. This way, the network can exploit the temporal information to further refine the fusion weights.

We trained the model using ADAM optimizer with an initial learning rate of 0.0001, dropout of 0.3, and batch size of 64 frames. The number of epochs was experimentally set to 21 for the first training stage and 10 for the second. We use the average Euclidean distance between the predicted and ground-truth 3D gaze vectors as loss function.

3.4.2 Input pre-processing

For this work we use head pose and eye locations in the 3D scene provided by the dataset. The 3D landmarks are extracted using the state-of-the-art method of Bulat and Tzimiropoulos [9], which is based on stacked hourglass networks [24].

During training, the original image is pre-processed to get the two normalized input images. The normalized whole-face patch is centered 0.1 meters ahead of the head center in the head coordinate system, and \mathbf{C}_n is defined such that the image has size of 250×250 pixels. The difference between this size and the final input size allows us to perform random cropping and zooming to augment the data (explained in Section 4.1). Similarly, each normalized eye patch is centered in their respective eye center locations. In this case, the virtual camera matrix is defined so that the image is cropped to 70×58 , while in practice the final patches have size of 60×48 . Landmarks are normalized using the same procedure and further pre-processed with mean subtraction and min-max normalization per axis. Finally, we divide them by a scaling factor w such that all coordinates are in the range $[0, w]$. This way, all concatenated feature values are in a similar range. After inference, the predicted normalized 2D angles are de-normalized back to the original 3D space.

4 Experiments

In this section, we evaluate the cross-subject 3D gaze estimation task on a wide range of head poses and gaze directions. Furthermore, we validate the effectiveness of the proposed architecture comparing both static and temporal approaches. We report the error in terms of mean angular error between predicted and ground-truth 3D gaze vectors. Note that due to the requirements of the temporal model not all the frames obtain a prediction. Therefore, for a




Figure 2: Ground-truth eye gaze \mathbf{g} and head orientation \mathbf{h} distribution on the filtered EYE-DIAP dataset for CS and FT settings, in terms of x- and y- angles.

fair comparison, the reported results for static models disregard such frames when temporal models are included in the comparison.

4.1 Training data

There are few publicly available datasets devoted to 3D gaze estimation and most of them focus on HCI with a limited range of head pose and gaze directions. Therefore, we use VGA videos from the publicly-available EYEDIAP dataset [2] to perform the experimental evaluation, as it is currently the only one containing video sequences with a wide range of head poses and showing the full face. This dataset consists of 3-minute videos of 16 subjects looking at two types of targets: continuous *screen* targets on a fixed monitor (CS), and *floating* physical targets (FT). The videos are further divided into *static* (S) and *moving* (M) head pose for each of the subjects. Subjects 12-16 were recorded with 2 different lighting conditions.

For evaluation, we filtered out those frames that fulfilled at least one of the following conditions: (1) face or landmarks not detected; (2) subject not looking at the target; (3) 3D head pose, eyes or target location not properly recovered; and (4) eyeball rotations violating physical constraints ($|\theta| \leq 40^\circ$, $|\phi| \leq 30^\circ$) [23]. Note that we purposely do not filter eye blinking moments to learn their dynamics with the temporal model, which may produce some outliers with a higher prediction error due to a less accurate ground truth. Figure 2 shows the distribution of gaze directions and head poses for both filtered CS and FT cases.

We applied data augmentation to the training set with the following random transformations: horizontal flip, shifts of up to 5 pixels, zoom of up to 2%, brightness changes by a factor in the range $[0.4, 1.75]$, and additive Gaussian noise with $\sigma^2 = 0.03$.

4.2 Evaluation of static modalities

First, we evaluate the contribution of each static modality on the FT scenario. We divided the 16 participants into 4 groups, such that appearance variability was maximized while maintaining a similar number of training samples per group. Each static model was trained end-to-end performing 4-fold cross-validation using different combinations of input modalities. Since the number of fusion units depends on the number of input modalities, we also compare different fusion layer sizes. The effect of data normalization is also evaluated by training a not-normalized face model where the input image is the face bounding box with square size the maximum distance between 2D landmarks.




Figure 3: Performance evaluation of the *Static* network using different input modalities (*O* - Not normalized, *N* - Normalized, *F* - Face, *E* - Eyes, *L* - 3D Landmarks) and size of fusion layers on the *FT* scenario.

As shown in Figure 3, all models that take normalized full-face information as input achieve better performance than the eyes-only model. More specifically, the combination of face, eyes and landmarks outperforms all the other combinations by a small but significant margin (paired Wilcoxon test, $p < 0.0001$). The standard deviation of the best-performing model is reduced compared to the face and eyes model, suggesting a regularizing effect due to the addition of landmarks. The not-normalized face-only model shows the largest error, proving the impact of normalization to reduce the appearance variability. Furthermore, our results indicate that the increase of fusion units is not correlated with a better performance.

4.3 Static gaze regression: comparison with existing methods

We compare our best-performing static model with three baselines. **Head:** Treating the head pose directly as gaze direction. **PR-ALR:** Method that relies on RGB-D data to rectify the eye images viewpoint into a canonical head pose using a 3DMM. It then learns an RGB gaze appearance model using ALR [14]. Predicted 3D vectors for *FT-S* scenario are provided by EYEDIAP dataset. **MPIIGaze:** State-of-the-art full-face 3D gaze estimation method [15]. They use an Alexnet-based CNN model with spatial weights to enhance information in different facial regions. We fine-tuned it with the filtered EYEDIAP subsets using our training parameters and normalization procedure.

In addition to the aforementioned *FT*-based evaluation setup, we also evaluate our method on the *CS* scenario. In this case there are only 14 participants available, so we divided them in 5 groups and performed 5-fold cross-validation. In Figure 4 we compare our method to MPIIGaze, achieving a statistically significant improvement of 14.6% and 19.5% on *FT* and *CS* scenarios, respectively (paired Wilcoxon test, $p < 0.0001$). We can observe that a restricted gaze target benefits the performance of all methods, compared to a more challenging unrestricted setting with a wider range of head poses and gaze directions.

Table 2 provides a detailed comparison on every participant, performing leave-one-out cross-validation on the *FT* scenario for *static* and *moving* head separately. Results show that, as expected, facial appearance and head pose have a noticeable impact on gaze accuracy, with average error differences of up to 7.7° among participants.




Figure 4: Performance comparison among MPIIGaze method [15] and our *Static* and *Temporal* versions of the proposed network for *FT* and *CS* scenarios.

Method	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	Avg.
Head	23.5	22.1	20.3	23.6	23.2	23.2	23.6	21.2	26.7	23.6	23.1	24.4	23.3	24.0	24.5	22.8	23.3
PR-ALR	12.3	12.0	12.4	11.3	15.5	12.9	17.9	11.8	17.3	13.4	13.4	14.3	15.2	13.6	14.4	14.6	13.9
MPIIGaze	5.3	5.1	5.7	4.7	7.3	15.1	10.8	5.7	9.9	7.1	5.0	5.7	7.4	3.8	4.8	5.5	6.8
Static	3.9	4.1	4.2	3.9	6.0	6.4	7.2	3.6	7.1	5.0	5.7	6.7	3.9	4.7	5.1	4.2	5.1
Temporal	4.0	4.9	4.3	4.1	6.1	6.5	6.6	3.9	7.8	6.1	4.7	5.6	4.7	3.5	5.9	4.6	5.2
Head	19.3	14.2	16.4	19.9	16.8	21.9	16.1	24.2	20.3	19.9	18.8	22.3	18.1	14.9	16.2	19.3	18.7
MPIIGaze	7.6	6.2	5.7	8.7	10.1	12.0	12.2	6.1	8.3	5.9	6.1	6.2	7.4	4.7	4.4	6.0	7.3
Static	5.8	5.7	4.4	7.5	6.7	8.8	11.6	5.5	8.3	5.5	5.2	6.3	5.3	3.9	4.3	5.6	6.3
Temporal	6.1	5.6	4.5	7.5	6.4	8.2	12.0	5.0	7.5	5.4	5.0	5.8	6.6	4.0	4.5	5.8	6.2

Table 2: Gaze angular error comparison for *static* (top half) and *moving* (bottom half) head pose for each subject in the *FT* scenario. Best results in bold.




Figure 5: Angular error distribution across gaze (a) and head orientation (b) spaces in the *FT* setting, in terms of x- and y- angles. For each space, we depict the *Static* model performance (left) and the contribution of the *Temporal* model versus *Static* (right). In the latter, positive difference means higher improvement of the *Temporal* model.

4.4 Evaluation of the temporal network

In this section, we evaluate the contribution of adding the temporal module to the static model. To do so, we trained a lower-dimensional version of the static network with comparable performance to the original, reducing the number of units of the second fusion layer to 2918. Results are reported in Figure 4 and Table 2. One can observe that using sequential information is helpful on the *FT* scenario, outperforming the static model by a statistically significant 4.4% (paired Wilcoxon test, $p < 0.0001$). This contribution is more noticeable in the *moving* head setting, proving that the temporal model can benefit from head motion information. In contrast, such information seems to be less meaningful in the *CS* scenario, where the obtained error is already very low for a cross-subject setting and the amount of head movement declines.

Figure 5 further explores the error distribution of the static network and the impact of sequential information. We can observe that the accuracy of the static model drops with extreme head poses and gaze directions, which can also be correlated to having less data in those areas. Compared to the static model, the temporal model particularly benefits gaze targets from mid-range upwards. Its contribution is less clear for extreme targets, probably again due to data imbalance.

Finally, we evaluated the effect of different recurrent architectures for the temporal model. In particular, we tested 1 (128 units) and 2 (256-128 units) LSTM and GRU layers, with 1 GRU layer obtaining slightly superior results (up to 0.12°). We also assessed the effect of sequence length fixing s in the range $\{4, 7, 10\}$, with $s = 7$ performing worse than the other two (up to 0.14°).

5 Conclusions

In this work, we studied the combination of full-face and eye images along with facial landmarks for person- and head pose-independent 3D gaze estimation. Consequently, we proposed a multi-stream recurrent CNN network that leverages the sequential information of eye and head movements. Both static and temporal versions of our approach significantly outperform current state-of-the-art 3D gaze estimation methods on a wide range of head poses and gaze directions. We showed that adding geometry features to appearance-based methods has a regularizing effect on the accuracy. Adding sequential information further benefits the final performance compared to static-only input, especially from mid-range upwards and in those cases where head motion is present. The effect in very extreme head poses is not clear due to data imbalance, suggesting the importance of learning from a continuous, balanced dataset including all head poses and gaze directions of interest. To the best of our knowledge, this is the first attempt to exploit the temporal modality in the context of gaze estimation from remote cameras. As future work, we will further explore extracting meaningful temporal representations of gaze dynamics, considering 3DCNNs as well as the encoding of deep features around particular tracked face landmarks [14].

Acknowledgements

This work has been partially supported by the Spanish project TIN2016-74946-P (MINECO/FEDER, UE), CERCA Programme / Generalitat de Catalunya, and the FP7 people program (Marie Curie Actions), REA grant agreement no FP7-607139 (iCARE - Improving Children Auditory REhabilitation). We gratefully acknowledge the support of NVIDIA Corporation with the donation of the GPU used for this research. Portions of the research in this paper used the EYEDIAP dataset made available by the Idiap Research Institute, Martigny, Switzerland.

References

- [1] Nicola C Anderson, Evan F Risko, and Alan Kingstone. Motion influences gaze direction discrimination and disambiguates contradictory luminance cues. *Psychonomic bulletin & review*, 23(3):817–823, 2016.
- [2] Shumeet Baluja and Dean Pomerleau. Non-intrusive gaze tracking using artificial neural networks. In *Advances in Neural Information Processing Systems*, pages 753–760, 1994.
- [3] Adrian Bulat and Georgios Tzimiropoulos. How far are we from solving the 2d & 3d face alignment problem? (and a dataset of 230,000 3d facial landmarks). In *International Conference on Computer Vision*, 2017.
- [4] Haoping Deng and Wangjiang Zhu. Monocular free-head 3d gaze tracking with deep learning and geometry constraints. In *Computer Vision (ICCV), 2017 IEEE International Conference on*, pages 3162–3171. IEEE, 2017.
- [5] Onur Ferhat and Fernando Vilariño. Low cost eye tracking. *Computational intelligence and neuroscience*, 2016:17, 2016.

- [6] Kenneth A Funes-Mora and Jean-Marc Odobez. Gaze estimation in the 3D space using RGB-D sensors. *International Journal of Computer Vision*, 118(2):194–216, 2016.
- [7] 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 ACM Symposium on Eye Tracking Research and Applications*. ACM, March 2014. doi: 10.1145/2578153.2578190.
- [8] 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*, pages 255–258. ACM, 2014.
- [9] Quentin Guillon, Nouchine Hadjikhani, Sophie Baduel, and Bernadette Rogé. Visual social attention in autism spectrum disorder: Insights from eye tracking studies. *Neuroscience & Biobehavioral Reviews*, 42:279–297, 2014.
- [10] 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*, 32(3): 478–500, 2010.
- [11] Qiong Huang, Ashok Veeraraghavan, and Ashutosh Sabharwal. Tabletgaze: dataset and analysis for unconstrained appearance-based gaze estimation in mobile tablets. *Machine Vision and Applications*, 28(5-6):445–461, 2017.
- [12] Robert JK Jacob and Keith S Karn. Eye tracking in human-computer interaction and usability research: Ready to deliver the promises. In *The mind’s eye*, pages 573–605. Elsevier, 2003.
- [13] László A Jeni and Jeffrey F Cohn. Person-independent 3d gaze estimation using face frontalization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 87–95, 2016.
- [14] Heechul Jung, Sihaeng Lee, Junho Yim, Sunjeong Park, and Junmo Kim. Joint fine-tuning in deep neural networks for facial expression recognition. In *Computer Vision (ICCV), 2015 IEEE International Conference on*, pages 2983–2991. IEEE, 2015.
- [15] Anuradha Kar and Peter Corcoran. A review and analysis of eye-gaze estimation systems, algorithms and performance evaluation methods in consumer platforms. *IEEE Access*, 5:16495–16519, 2017.
- [16] Kyle Krafka, Aditya Khosla, Petr Kellnhofer, Harini Kannan, Suchendra Bhandarkar, Wojciech Matusik, and Antonio Torralba. Eye tracking for everyone. In *Computer Vision and Pattern Recognition (CVPR), 2016 IEEE Conference on*, pages 2176–2184. IEEE, 2016.
- [17] Simon P Liversedge and John M Findlay. Saccadic eye movements and cognition. *Trends in cognitive sciences*, 4(1):6–14, 2000.
- [18] Feng Lu, Takahiro Okabe, Yusuke Sugano, and Yoichi Sato. A head pose-free approach for appearance-based gaze estimation. In *BMVC*, pages 1–11, 2011.

- [19] Feng Lu, Yusuke Sugano, Takahiro Okabe, and Yoichi Sato. Inferring human gaze from appearance via adaptive linear regression. In *Computer Vision (ICCV), 2011 IEEE International Conference on*, pages 153–160. IEEE, 2011.
- [20] Päivi Majaranta and Andreas Bulling. Eye tracking and eye-based human–computer interaction. In *Advances in physiological computing*, pages 39–65. Springer, 2014.
- [21] Kenneth Alberto Funes Mora and Jean-Marc Odobez. Gaze estimation from multi-modal kinect data. In *Computer Vision and Pattern Recognition Workshops (CVPRW), 2012 IEEE Computer Society Conference on*, pages 25–30. IEEE, 2012.
- [22] Carlos Hitoshi Morimoto, Arnon Amir, and Myron Flickner. Detecting eye position and gaze from a single camera and 2 light sources. In *Pattern Recognition, 2002. Proceedings. 16th International Conference on*, volume 4, pages 314–317. IEEE, 2002.
- [23] IMO MSC. Circ. 982 (2000) guidelines on ergonomic criteria for bridge equipment and layout.
- [24] Alejandro Newell, Kaiyu Yang, and Jia Deng. Stacked hourglass networks for human pose estimation. In *European Conference on Computer Vision*, pages 483–499. Springer, 2016.
- [25] Yasuhiro Ono, Takahiro Okabe, and Yoichi Sato. Gaze estimation from low resolution images. In *Pacific-Rim Symposium on Image and Video Technology*, pages 178–188. Springer, 2006.
- [26] Cristina Palmero, Elisabeth A. van Dam, Sergio Escalera, Mike Kelia, Guido F. Lichtert, Lucas P.J.J Noldus, Andrew J. Spink, and Astrid van Wieringen. Automatic mutual gaze detection in face-to-face dyadic interaction videos. In *Proceedings of Measuring Behavior*, pages 158–163, 2018.
- [27] Omkar M. Parkhi, Andrea Vedaldi, and Andrew Zisserman. Deep face recognition. In *British Machine Vision Conference*, 2015.
- [28] Derek R Rutter and Kevin Durkin. Turn-taking in mother–infant interaction: An examination of vocalizations and gaze. *Developmental psychology*, 23(1):54, 1987.
- [29] 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*, pages 271–280. ACM, 2013.
- [30] Yusuke Sugano, Yasuyuki Matsushita, and Yoichi Sato. Appearance-based gaze estimation using visual saliency. *IEEE transactions on pattern analysis and machine intelligence*, 35(2):329–341, 2013.
- [31] 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.
- [32] Kar-Han Tan, David J Kriegman, and Narendra Ahuja. Appearance-based eye gaze estimation. In *Applications of Computer Vision, 2002.(WACV 2002). Proceedings. Sixth IEEE Workshop on*, pages 191–195. IEEE, 2002.

- [33] Ronda Venkateswarlu et al. Eye gaze estimation from a single image of one eye. In *Computer Vision, 2003. Proceedings. Ninth IEEE International Conference on*, pages 136–143. IEEE, 2003.
- [34] Kang Wang and Qiang Ji. Real time eye gaze tracking with 3d deformable eye-face model. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1003–1011, 2017.
- [35] Oliver Williams, Andrew Blake, and Roberto Cipolla. Sparse and semi-supervised visual mapping with the $s^{\wedge}3gp$. In *Computer Vision and Pattern Recognition, 2006 IEEE Computer Society Conference on*, volume 1, pages 230–237. IEEE, 2006.
- [36] William Hyde Wollaston et al. Xiii. on the apparent direction of eyes in a portrait. *Philosophical Transactions of the Royal Society of London*, 114:247–256, 1824.
- [37] Erroll Wood and Andreas Bulling. Eytatb: Model-based gaze estimation on unmodified tablet computers. In *Proceedings of the Symposium on Eye Tracking Research and Applications*, pages 207–210. ACM, 2014.
- [38] Erroll Wood, Tadas Baltrušaitis, Xucong Zhang, Yusuke Sugano, Peter Robinson, and Andreas Bulling. Rendering of eyes for eye-shape registration and gaze estimation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3756–3764, 2015.
- [39] Erroll Wood, Tadas Baltrušaitis, Louis-Philippe Morency, Peter Robinson, and Andreas Bulling. A 3d morphable eye region model for gaze estimation. In *European Conference on Computer Vision*, pages 297–313. Springer, 2016.
- [40] 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*, pages 131–138. ACM, 2016.
- [41] Dong Hyun Yoo and Myung Jin Chung. A novel non-intrusive eye gaze estimation using cross-ratio under large head motion. *Computer Vision and Image Understanding*, 98(1):25–51, 2005.
- [42] Xucong Zhang, Yusuke Sugano, Mario Fritz, and Andreas Bulling. Appearance-based gaze estimation in the wild. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4511–4520, 2015.
- [43] Xucong Zhang, Yusuke Sugano, Mario Fritz, and Andreas Bulling. It’s written all over your face: Full-face appearance-based gaze estimation. In *Proc. IEEE International Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2017.