GazeDirector: Fully Articulated Eye Gaze Redirection in Video

Erroll Wood[†1,4], Tadas Baltrušaitis[†2,4], Louis-Philippe Morency[2], Peter Robinson[1], and Andreas Bulling[3]

[1]University of Cambridge, UK
[2]Carnegie Mellon University, USA
[3]Max Planck Institute for Informatics, Germany
[4]Microsoft

Abstract

We present GazeDirector, a new approach for eye gaze redirection that uses model-fitting. Our method first tracks the eyes by fitting a multi-part eye region model to video frames using analysis-by-synthesis, thereby recovering eye region shape, texture, pose, and gaze simultaneously. It then redirects gaze by 1) warping the eyelids from the original image using a model-derived flow field, and 2) rendering and compositing synthesized 3D eyeballs onto the output image in a photorealistic manner. GazeDirector allows us to change where people are looking without person-specific training data, and with full articulation, i.e. we can precisely specify new gaze directions in 3D. Quantitatively, we evaluate both model-fitting and gaze synthesis, with experiments for gaze estimation and redirection on the Columbia gaze dataset. Qualitatively, we compare GazeDirector against recent work on gaze redirection, showing better results especially for large redirection angles. Finally, we demonstrate gaze redirection on YouTube videos by introducing new 3D gaze targets and by manipulating visual behavior.

1. Introduction

Gaze redirection is an upcoming research topic where the goal is to alter an image to change where someone is looking (see Figure 1) [GKSL16, TZS*16b]. This is an important generalization of the classic gaze correction task [ZGT99, CSBT03], where someone's gaze is adjusted to a single new direction to simulate eye contact. With gaze redirection, gaze can be adjusted to any direction.

The ability to freely change where someone is looking paves the way for a variety of compelling new applications (see Figure 2). For example, taking a group picture with everyone is looking at the camera at the same time can be difficult [SYFN13]. Imagine a gaze-correcting camera that could always enforce eye contact, no matter where people are actually looking. Also, one challenge for actors nowadays is performing alone before other computer-generated characters are composited in. Where are they supposed to look? With gaze redirection their apparent point-of-regard could be controlled in post-production, ensuring they look at virtual characters. Gaze direction is also an important social signal [Eme00] – the ability to redirect gaze or even impose specific visual behaviours on video content in real-time could serve as a useful experimental tool, e.g. to study gaze following or joint attention in autism research [JC04].

A reliable and robust gaze redirection algorithm should work with
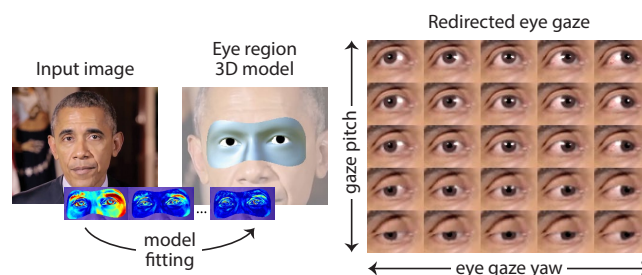
Figure 1: *GazeDirector is a new 3D model based approach for gaze redirection. We first recover the shape and appearance of the eyes by fitting a 3D eye region model. We then redirect gaze by warping the eyelids and rendering new redirected eyeballs. Examples of redirected gaze can be seen on the right.*

previously unseen people and handle desired gaze directions which differ significantly from the original gaze. Thies et al. [TZS*16b] recently proposed an approach which requires per-user calibration, a tedious process that is unsuitable for many scenarios. More relevant to our goal of user-independent gaze redirection is Deep-Warp [GKSL16], an approach that uses a deep neural network to directly predict an image-warping flow field between two eye images with a known gaze "correction" angular offset between them. This flow field is applied to the original image to redirect gaze. In this way, DeepWarp can only redirect gaze by shifting it by an angular

---

Figure 2: *GazeDirector enables new applications that were previously impossible. (a) Taking group pictures with everyone looking at the camera can be tricky. (b) A gaze correcting camera can ensure this is always the case. (c,d) A challenge for actors is knowing where to look before visual effects are added to a scene. This can be modified in post-production, so if a CGI character is changed, the actor's gaze can be adjusted accordingly. The highlighted faces in (a,b) and the face in (c,d) have been modified by GazeDirector.*

offset; it cannot specify new gaze directions explicitly. Furthermore, this approach is prone to producing unsightly artefacts when redirecting gaze over large angles. This problem is fundamental in any purely warping-based approach since it is impossible to warp parts of the eye that were occluded in the original image.

In this work we present *GazeDirector*, a new approach for person-independent gaze redirection. The main idea of our approach is to model the eye region in 3D instead of trying to predict a flow field directly from an input image [GKSL16]. Since we recover the shape and pose of the eyes in 3D, our approach can redirect gaze with *full articulation*: GazeDirector can precisely specify new desired gaze targets or directions in 3D instead of using gaze angle correction offsets [GKSL16]. To model the eye in 3D, we extend a recently proposed method [WBM*16a] to fit a 3D morphable model of the eye region to both eyes in an input image using analysis-by-synthesis. Once we have recovered the shape, pose, and appearance of the eyes, we redirect gaze in two steps. First, we compute a dense model-derived flow field corresponding to eyelid motion between the original and desired gaze directions. This dense flow field is efficiently extrapolated from sparse per-vertex flow values using GPU rasterization. We apply this flow field to the input image to warp the eyelids. Second, we render and composite our redirected eyeball models onto the output image in a photorealistic manner.

Contributions   1) Our primary contribution is GazeDirector: a new method that demonstrates how eye-region model fitting using analysis-by-synthesis enables superior gaze redirection compared to previous approaches (§3). In addition, we present the following secondary contributions: 2) A practical approach for rapid synthesis of dense model-derived optical flow fields using GPU rasteriza-

tion (§5.1). 3) Improvements over the state-of-the-art in gaze estimation using our dataset-independent model fitting approach (§6.1).

2. Related Work

Eye gaze manipulation   The lack of eye contact during video-conferencing is a well-known problem in computer graphics. There are three main approaches to tackle it: 1) novel-view synthesis, 2) eye-replacement, and 3) eye-warping.

Novel-view synthesis methods re-render the subject's face so they appear to be looking at the camera. The first step is recovering a dense depthmap of the face – this can been done with stereo vision [YZ02, CSBT03], RGB-D (color with depth) cameras [KPB*12], and monocular RGB cameras [GBK*14]. This facial depthmap is then rotated and re-rendered from a new viewpoint along a frontal gaze path. However, as these methods distort the face as a whole, they are not suitable for more general forms of of gaze manipulation.

Eye-replacement methods replace eyes in the original image with new eye images representing different gaze. The most realistic approaches collect a set of person-specific images of eyes looking at a camera, and composite them into the original face [WFA10, QLTH15, SSSH16]. These methods require person-specific eye images to pick from, and encounter issues when compositing eyes across different head poses or illumination conditions. Other eye-replacement approaches synthesize new eyeballs with graphics [GTZ*00, WK03]. However, these methods do not move the eyelids – an important cue for vertical gaze, and only use rudimentary 2D graphics techniques that ignore iris color, head pose, and scene illumination. Our method instead synthesizes new eyeballs taking eyelid motion, iris color, and illumination into account.

Warping-based methods can redirect gaze without requiring person-specific training data. These methods learn to generate a flow field from one eye image to another using training pairs of eye images with known gaze offsets between them. This flow field is used to warp pixels in the original image, thus modifying gaze [KL15, GKSL16]. However, purely warping-based methods suffer three major limitations: First, they can only offset the original unknown gaze direction, so cannot specify a new gaze direction explicitly. Second, the range of possible redirection is limited by the gaze directions in the training set. Third, warping artefacts appear for large redirection angles as parts of the eye that were originally occluded cannot be synthesized correctly. Using 3D models, GazeDirector can explicitly specify new gaze directions in 3D, without training data, and without introducing artefacts.

Like us, Banf and Blanz [BB09] used morphable models to redirect gaze. They fit a single-part face model to an image, and redirect gaze by deforming the eyelids using an example-based approach, and sliding the iris across the model surface using texture-coordinate interpolation. Since they use a mesh where the face and eyes are joined, their method only works when people look straight ahead. GazeDirector instead models the face and eyeballs as separate parts, allowing for non-frontal input gaze.

Facial performance capture

Since GazeDirector recovers the shape, texture, pose, and gaze of

Figure 3: *Given observed image I_{obs}, we first initialize our model using 25 facial landmarks from a face tracker [BRM16]. We then find optimal model parameters Φ^* using analysis-by-synthesis, minimizing a reconstruction energy $E(\Phi)$. We then modify Φ^* with the desired gaze and eyelid behaviour, resulting in a new I_{syn} which we blend onto I_{obs}, giving a redirected gaze image.*

the facial eye region, it is also related to work on monocular facial performance capture – a well established research topic [KRP*15]. The goal is to recover dynamic facial geometry and appearance using commodity cameras alone.

Monocular facial performance capture is a highly under-constrained problem, so a parametric face model [BV99] is often used as a prior to help recover shape and albedo. Such models can then be fit to either RGB-D data [WBLP11, TZN*15] or RGB data [CHZ14, SSKS15, TZS*16a]. However, these approaches generally avoid the eyes, cutting them out of the mesh [CBZB15, TZS*16a]. This is because the parametric face model they use only represents the surface of the skin, and has reduced fidelity around the eye due to poor correspondences in the source head scan data. For GazeDirector, we extended a previous model that was built using high quality scans [WBM*16a], with care taken to maintain correspondences around the eyelids and eye corners. Critically, this model treats the eyeballs as separate parts that move independently from the face.

Some previous work tracked the eyes as a part of the face. Garrido et al. [GZC*16] include eyeball geometry in a "detail" layer of their facial mesh. Though this can lead to acceptable re-rendering, it does not allow gaze redirection as the eyeballs and face are joined in a single mesh. Suwajanakorn et al. [SSKS15] model eyeball movement by interpolating between facial textures. This does not allow smooth arbitrary eyeball motion, and requires a large training set of person-specific images with eye movement. Recent work has combined a facial skin surface capture system with a separate gaze tracker [TZS*16a, WSXC16, CWW*16]. Our approach instead captures the facial eye region and eyeball simultaneously. This lets us reliably recover eyeball shape and texture parameters – important for realistic gaze redirection.

There have been recent breakthroughs in capturing the eyeballs and eyelids in extreme detail using special equipment [BBN*14, BBK*15, BBGB16]. Our work does not come close to this level of detail. Instead, we focus on capturing the eye for gaze redirection in commodity monocular images and video.

3. Overview

As shown in Figure 3, our approach consists of two main stages: eye region tracking and eye gaze redirection.

Tracking Given a monocular RGB image frame, we first capture the eyes by fitting our eye region model. This model consists of two parts: a generative facial part and an articulated eyeball part. It is defined by a set of parameters Φ that describe shape, texture, pose, and scene illumination. We fit our model to the image using analysis-by-synthesis, searching for optimal parameters Φ^* by minimizing a photometric reconstruction energy.

Redirection We redirect gaze in two steps: 1) We warp the eyelids in the original image using a flow field derived from our 3D model. We efficiently calculate this flow field by re-posing our eye region model to change gaze, and rendering the image-space flow between tracked and re-posed eye regions. 2) We then render the redirected eyeballs and composite them back into the image. We blur the boundary between the skin and eyeball to soften the transition so the eyes they "fit in" better.

4. Eye region tracking

For our gaze redirection to look plausible, we must first recover the original shape and texture of the eye region. Given an image frame I_{obs}, we therefore wish to recover a set of optimal parameters Φ^* that best explains it in terms of our eye region model. We search for Φ^* using analysis-by-synthesis: iteratively rendering a synthetic eye region image I_{syn}, comparing it to I_{obs} using our reconstruction energy E (defined in Equation 4), and updating Φ accordingly.

4.1. Eye region model

At the heart of our method lies a multi-part eye region model based on that by Wood et al. [WBM*16a]. For GazeDirector, we extended it to model two eyes rather than one, simplified the iris color model to improve robustness, and added aesthetic improvements (subdivision surfaces, ambient occlusion, reflections) to improve realism. Our model contains four main parts: the left and right facial eye regions, and the left and right eyeballs. It is parameterized by Φ:

$$\Phi = \{\beta, \tau, \theta, \iota\}, \tag{1}$$

where β are the set of shape parameters, τ the texture parameters, θ the pose parameters, and ι the illumination parameters. We now describe each parameter below.

Shape model $\mathcal{M}_{geo}$ Texture model $\mathcal{M}_{tex}$

Figure 4: *The average facial shape μ_{geo} and texture μ_{tex}, along with the top modes of variation. The first mode of shape variation moves between hooded and protruding eyes, and the first mode of texture variation moves between dark and light skin.*

Shape β The geometric shape of each eye region is described by a linear Principal Component Analysis (PCA) model $\mathcal{M}_{geo} \in \mathbb{R}^{3n}$ in the style of previous work [BV99]. This comprises $n = 229$ vertices and was built from a collection of 22 high resolution scans acquired online [WBM*16b]. We assume faces are symmetrical, so the shapes of both eye regions are controlled with a single set of coefficients $\beta_{face} \in \mathbb{R}^{16}$,

$$\mathcal{M}_{geo}(\beta_{face}) = \mu_{geo} + U\text{diag}(\sigma_{geo})\beta_{face} \qquad (2)$$

where μ_{geo} is the average face shape, U the modes of shape variation, and σ_{geo} the standard deviations of these modes (see Figure 4). For simplicity, each $\beta_i \in \beta_{face}$ is scaled so that $\beta_i = 1$ represents one standard deviation's worth of variation in that dimension. For the eyeball we use a standard two-sphere model based off physiological averages [RAB*14]. We also include a parameter β_{iris} that controls iris size by scaling vertices on the iris boundary about the pupil.

Texture τ We use a linear PCA texture model $\mathcal{M}_{tex} \in \mathbb{R}^{3m}$ of the facial eye region, built from the same set of scans. Rather than model the color of each vertex [BV99], $\mathcal{M}_{tex}$ generates RGB texture maps sized $m = 512 \times 512$px that we apply to both eye regions. This linear texture model is controlled with texture coefficients $\tau_{face} \in \mathbb{R}^8$,

$$\mathcal{M}_{tex}(\tau_{face}) = \mu_{tex} + V\text{diag}(\sigma_{tex})\tau_{face} \qquad (3)$$

where μ_{tex} is the average face texture, V the modes of texture variation, and σ_{tex} the respective standard deviations. Each coefficient is scaled in a similar way to $\mathcal{M}_{geo}$, so it represents one standard deviation in its dimension. As shown in Figure 5, we vary the iris by multiplying the iris region of the base eyeball texture with an RGB color τ_{iris}. Since the "white of the eye" is rarely purely white, we also tint it with another color τ_{tint}

Pose θ Our pose parameters describe both global and local pose. Globally, the eye regions are positioned with rotation θ_R and translation θ_T. The interocular distance is controlled via θ_{iod} The eyeball positions are fixed in relation to the eye regions. Our local pose parameters allow the eyeballs to rotate independently from the face, controlling gaze. The general gaze direction is given by pitch and yaw angles θ_p and θ_y, and vergence is controlled with θ_v. When the eyeball looks up or down, the eyelids follow it. We use procedural animation to pose the eyelids in the facial mesh by rotational ammount θ_{lid} [WBM*16b].

Figure 5: *Our eyeball model captures iris color variation with an RGB color τ_{iris} (yellow arrow). Environmental reflections are added with spherical environment maps (red arrow).*

Figure 6: *Our eyelid posed using procedural animation for eyelid gaze pitch angles θ_{lid} between $-20°$ and $+20°$.*

Illumination ι We assume a simple illumination model of ambient light coupled with a single directional light. The ambient light has intensity $\iota_{amb} \in \mathbb{R}^3$, and the directional light has intensity $\iota_{dir} \in \mathbb{R}^3$ and direction defined by rotation $\iota_R \in \mathbb{R}^2$ (pitch and yaw angles). We assume all surfaces are Lambertian. Though ι cannot describe complex scene illumination, we found it was sufficient in many cases considering the small facial region that we consider.

In total we have $17 + 14 + 11 + 9 = 51$ parameters of Φ to optimize over.

Rendering the model Once our model has been configured with parameters Φ, we render synthetic images $I_{syn}(\Phi)$ using a DirectX-based rasterizer. We fix our virtual camera location at the world origin, and assume knoweldge (or estimate) of camera intrinsic parameters.

Realistically rendering eyes is a challenge [RAB*14]. We implement three additional effects to improve the realism of our output. First, as our model is low-resolution, it appears blocky when rendered. We therefore smooth the skin's surface using a single step of Loop subdivision [Loo87] with precomputed stencils for efficiency. Second, we use physically correct corneal refraction techniques in the eyeball shader to better model its layered transparent structure [JDvdP12]. Third, we approximate ambient occlusion shadowing on the eyeball using a single-pass analytic techniqe: we project the positions of eyelid vertices into eyeball *uv* space, fit a 2D cubic polynomial to them, and apply per-pixel ambient occlusion as a function of distance to each eyelid polynomial. Refraction, reflection, and ambient occlusion are applied by the fragment shader.

4.2. Energy formulation

A good energy function is critical to the success of any analysis-by-synthesis method. Our proposed energy $E(\Phi)$ is a weighted sum of several terms, each encoding a different requirement of our model
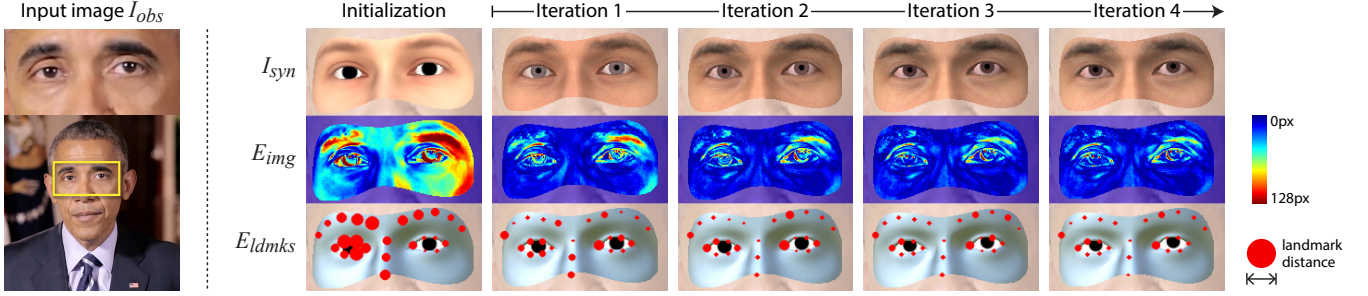
Figure 7: *We fit our 3D eye region model to an image my minimizing a reconstruction energy $E(\Phi)$. Our two main energy terms are a dense photometric error term E_{img} and a sparse landmark similarity term E_{ldmks}. This figure shows the energies decreasing over four iterations of the Gauss Newton algorithm.*

fit. Each term is expressable as a sum-of-squares, allowing us to minimize $E(\Phi)$ using the Gauss-Newton algorithm.

$$E(\Phi) = \underbrace{E_{img}(\Phi) + E_{ldmks}(\Phi)}_{\text{Data terms}} + \underbrace{E_{stats}(\Phi) + E_{pose}(\Phi)}_{\text{Prior terms}} \quad (4)$$

Our data terms (see Figure 7) guide our model fit using image pixels and facial landmarks, while our prior terms penalize unlikely facial shape and texture, and eyeball orientations. We now describe each term in detail.

Image similarity E_{img}   Our primary goal is to minimize the photometric reconstruction error between I_{syn} and I_{obs}. The data term E_{img} expresses how well the fitted model explains I_{obs} by densely measuring pixel-wise differences across the images using a robust mean squared error. We promote image similarity with the term

$$E_{img}(\Phi) = \frac{1}{|\mathcal{P}|} \sum_{p \in \mathcal{P}} \rho\big(|I_{syn}(p) - I_{obs}(p)|\big)^2 \quad (5)$$

where $\mathcal{P} \subset I_{syn}$ represents the set of rendered foreground pixels belonging to our 3D model. The background pixels are ignored. The robust function $\rho(e) = \min(\sqrt{T}, e)$, for threshold T, alleviates the effects of outliers; this is important for recovering iris color in the presence of strong specular highlights on the eye.

Landmark similarity E_{ldmks}   The face contains several landmark feature points that can be tracked reliably. We therefore regularize our dense data term (E_{img}) using a sparse set of landmarks $\mathcal{L}$ provided by a face tracker [BRM16]. $\mathcal{L}$ consists of 25 points that describe the eyebrows, nose and eyelids. For each 2D tracked landmark $l \in \mathcal{L}$, we also compute a corresponding synthesized 2D landmark l' as a linear combination of projected vertices in our shape model. Facial landmark similarities are incorporated into our energy using

$$E_{ldmks}(\Phi) = \lambda_{ldmks} \cdot \frac{1}{|\mathcal{P}|} \sum_{i=0}^{|\mathcal{L}|} |l_i - l_i'|^2 \quad (6)$$

As landmark distances $|l_i - l_i'|$ are measured in image-space, we normalize the energy by dividing through by foreground area $|\mathcal{P}|$ to avoid bias from eye region size in the image. The importance of E_{ldmks} is controlled with weight λ_{ldmks}.

Statistical prior E_{stats}   We penalize unlikely facial shape and texture using a statistical prior [BV99]. As we assume a normally distributed population, our PCA model parameters should be close to the mean 0:

$$E_{stats}(\Phi) = \lambda_{geo} \cdot \sum_{i=0}^{|\beta|} \beta_i^2 + \lambda_{tex} \cdot \sum_{i=0}^{|\tau|} \tau_i^2 \quad (7)$$

Recall that $\beta_i \in \beta$ and $\tau_i \in \tau$ are scaled by their respective standard deviations in our model. This energy helps our fit avoid degenerate facial shapes and texture, and guides its recovery from poor local minima found in previous frames. The penalties for unlikely shape and texture are weighted separately with λ_{geo} and λ_{tex}.

Pose prior E_{pose}   Our final energy penalizes mismatched parameters for eyeball gaze direction and eyelid position. The eyelids follow eye gaze, so if the eyeball is looking upwards, the eyelids should be rotated upwards, and visa versa. We enforce eyelid pose consistency with

$$E_{pose}(\Phi) = \lambda_{pose} \cdot |\theta_{lid} - \theta_p|^2 \quad (8)$$

where θ_{lid} is the eyelid pitch angle of our model's face parts, and θ_p is the gaze pitch angle of our eyeball parts. Its relative importance is controlled by weight λ_{pose}.

4.3. Optimization procedure

Minimizing our proposed objective $E(\Phi)$ is a challenging high-dimensional non-convex optimization problem. We use a GPU-assisted, annealed form of the Gauss-Newton algorithm, where the parameter update for Φ is as follows:

$$\Phi^{i+1} = \Phi^i - \eta^i \, (J_r{}^T J_r)^{-1} \cdot J_r{}^T r \quad (9)$$

where r is the vector of energy function residuals, J_r the Jacobian matrix of residuals r evaluated at Φ^i, $J_r{}^T J_r$ the approxmation to the Hessian matrix, and η the annealing rate. We perform a variable number of Gauss-Newton iterations, terminating early if our energy stops decreasing. Figure 7 shows four iterations of our model fit.

To compute the Jacobian we use numerical central derivatives. This is an expensive operation, requiring two images to be rendered for every parameter. We keep our system performant by calculating J_r and $J_r{}^T J_r$ entirely on the GPU, avoiding expensive pipeline stalls
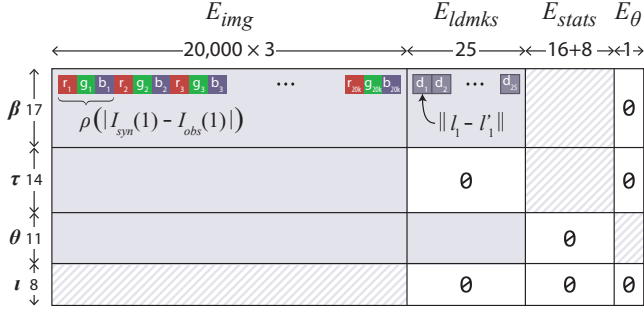
Figure 8: *The non-zero structure of our Jacobian J_r for a $200 \times 100px$ eye region. J_r is calculated entirely on the GPU. Dashed regions represent sparse blocks.*

Figure 9: *We efficiently convert sparse per-vertex image-space flows (left) to a dense per-pixel flow field (right) using GPU rasterization. We use this dense flow-field to warp the eyelids.*

from cross-system data transfer. Additionally, since image rendering is a key operation for our system, we use a tailored DirectX rasterizer that can render I_{syn} over 5000 times per second. To further lighten the computational load of our numerical derivatives, we mask out a subset of Φ when tracking in a video, so optimize over a smaller set of parameters frame-to-frame. As a result, GazeDirector can run at interactive rates.

Initialization The energy landscape of $E(\Phi)$ is riddled with local minima, so we must start from a good initializion. Our face tracker provides 3D estimates for the facial landmark positions. We initialize global translation to the mean landmark position and set global rotation parameters using the the Kabsch [Kab76] algorithm, minimizing the root mean squared deviation between the detected and initial set of landmark points in 3D. Other parameters are initialized to 0 by default, except for interocular distance and iris size, for which we use anthropomorphic averages, and illumination, for which we experimentally chose a basic setup. When tracking in video, we exploit temporal similarities by initializing Φ_{init} with Φ^* from the previous frame.

5. Eye gaze redirection

Once we have obtained a set of fitted model parameters Φ^* for an image I_{obs}, our next step is to redirect gaze to point at a new 3D target g'. We first modify Φ^* to obtain Φ' that represents the redirected gaze. We then calculate the optical flow between eye region models with Φ^* and Φ', and use this to warp the eyelids in the source image. Finally, we render the redirected eyeballs and seamlessly composite them into the output image.

Re-posing our model The first step of gaze re-direction is straightforward: given a new gaze target g', we calculate new values for gaze pitch θ'_p, yaw θ'_y, and vergence θ'_v so each eyeball points towards g'. Furthermore, we calculate θ'_{lid} to match the new gaze direction. Altogether, these new gaze parameters are encoded in Φ'.

5.1. Warping the eyelids

When the eyeball rotates, the eyelids move with it. To simulate this, we warp the eyelids from the original image using a model-derived

optical flow field O. To calculate O, we first calculate the sparse screen-space flow $o_i \in \mathbb{R}^2$ for each vertex $v_i \in \mathbb{R}^3$ in both facial parts of the eye region:

$$o_i = \Pi\left(\Theta'(v_i)\right) - \Pi\left(\Theta^*(v_i)\right) \quad i \in [0, 458] \quad (10)$$

where Π is the projection defined by our camera parameters, and $\Theta^{*|'}$ are the transforms that combine eyelid motion (θ_{lid}) with model-to-world transforms θ_R and θ_T. It is common for analysis-by-synthesis methods to use GPU rasterization to evaluate an objective function [SKR*15, TZS*16a]. We propose a simple and efficient approach for computing dense flow-fields using the same framework. To efficiently distribute sparse flow values across image space, we load per-vertex flows o_i into our renderer as vertex attributes and let the rasterization stage interpolate between them and handle occlusions between different model parts (see Figure 9). This takes ~5ms. The result is a dense flow field O that we use to remap source image pixels to simulate eyelid motion.

5.2. Compositing redirected eyeballs

Once the eyelids have been warped, we render the portion of the eyeballs between the eyelids and composite them onto the output image. Following rasterization, the eyelid edges will be perfectly sharp and unlikely to match the observed image. We therefore follow the approach adopted by the real-time rendering community [JDvdP12, KACM16], and blur the seam where the eyeballs meet the eyelids with a small Gaussian.

A shortcoming of our underlying scene model is the lack of specular reflections on the eyeball surface. Real world eye images often exhibit strong highlights or glints. We decided not to explicitly model multiple light sources in Φ because of the additional computational cost with numerical derivatives. We instead pre-rendered a set of five spherical reflection maps that model common environmental lighting scenarios (see Figure 5), and use them to apply specular reflections on the eyeball at runtime. This choice is made by seeking the reflection map that minimizes image error. While this cannot model complex environmental reflections, it improves the perceived quality of the eyeball re-rendering.

6. Evaluations

In this section we evaluate GazeDirector. Quantitatively, we evaluate our model fitting stage with a gaze estimation experiment, and our

Figure 10: *Eye region model fits on the Columbia gaze dataset [SYFN13] showing true gaze (red) and estimated gaze (cyan).*
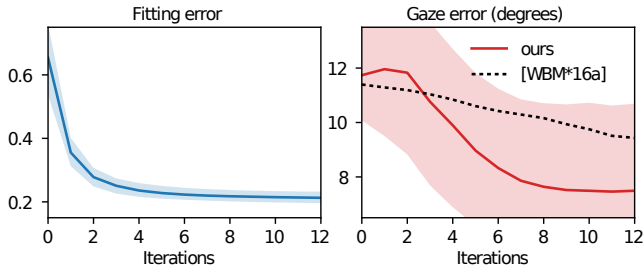

Figure 11: *Fitting error and gaze error for the Columbia dataset [SYFN13]. Line is median, filled region is interquartile range. Our second-order optimization strategy converges much faster than previous first-order methods [WBM*16a].*

Figure 12: *DeepWarp (top rows) [GKSL16] and GazeDirector (bottom rows) showing horizontal gaze redirection up to $45°$. Our model based approach avoids the smudging artefacts (red) encountered from large redirection angles with DeepWarp.*

Figure 13: *Redirection error decreases as we enable more parts of our redirection pipeline. The x-axis represents image error, and the y-axis represents the proportion of data under that error.*

gaze synthesis stage with a gaze redirection experiment. Qualitatively, we compare our method against recent work [GKSL16], and demonstrate gaze redirection and visual behaviour manipulation on YouTube videos.

6.1. Model fitting performance

We performed an experiment to assess our fitting strategy. We measured two factors: 1) photometric error to determine how well we reconstructed the image, and 2) gaze estimation error to see if we can correctly recover eyeball pose. We used the Columbia gaze dataset [SYFN13], which contains images of 56 people looking at a target grid on the wall. The participants were constrained by a head-clamp, and images were taken from five different head orientations. In our experiments we used a subset of 34 people (excluding those with eyeglasses) with 20 images per person.

Results of our experiment can be seen in Figure 11, and example model fits can be seen in Figure 10. Photometric error and gaze estimation error decrease with the number of model fitting iterations. This confirms the effectiveness of our fitting strategy. If we examine the pitch and yaw components of gaze separately, we outperform recent work [JC16] in terms of gaze yaw ($3.13°$ vs $3.51°$), though perform worse in terms of gaze pitch ($6.92°$ vs $4.27°$). This result is promising since GazeDirector operates in a dataset agnostic manner, while previous work [JC16] was trained on the Columbia dataset specifically. Furthermore, our second-order optimization strategy leads to faster convergence than first-order methods used in previous work [WBM*16a], despite performing a similar amount of computation per iteration.

6.2. Gaze redirection

We performed an experiment to evaluate our gaze redirection stages. We prepared another subset of the Columbia gaze dataset [SYFN13] with neutral head pose. We aligned images of each participant using facial landmarks [BRM16], and used the aligned images with different gaze as ground truth for "redirected gaze". Following model fitting on the frontal gaze image, we produced three output images for each different gaze image: a) with no gaze redirection, b) with gaze redirection with the eyeballs only, and c) with gaze redirection with eyeballs and eyelids. We measured the per-pixel image difference between GazeDirector images and the ground truth redirected gaze images (see Figure 13). The benefits of both eyeball redirection and eyelid redirection are clear.

Comparison to DeepWarp [GKSL16] Previous work produces unsightly smudging artefacts when starting from non-central gaze, and redirecting gaze over large angles. This is because their method fails to correctly hallucinate parts of the eyeball that were originally occluded. As can be seen in Figure 12, these issues do not arise with GazeDirector as we use a 3D model. Furthermore, since DeepWarp can only apply an angular gaze offset to an input gaze direction, it cannot be used to produce results like those in Figure 14 where someone has been made to look at 3D gaze targets. Please see our supplementary video for additional comparisons.

Input frame $\pm15°$ pitch, $\pm20°$ yaw   Redirected eye gaze in YouTube videos
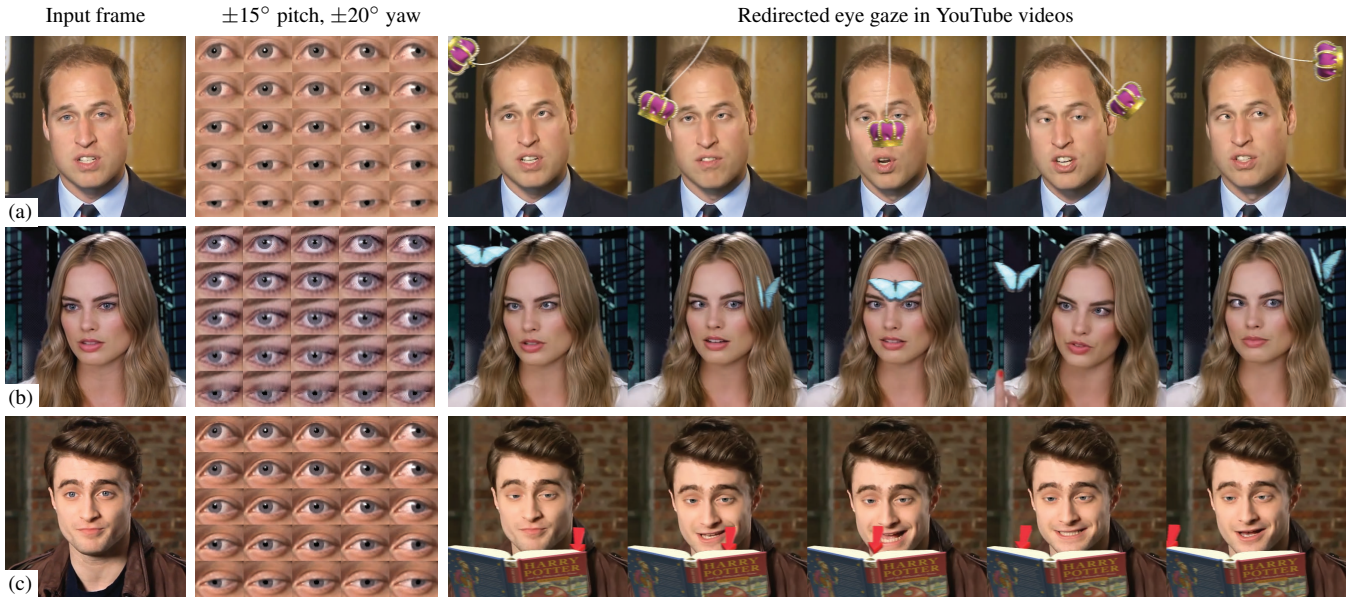
Figure 14: *Example input frame, redirected eye gaze grid, and example output frames for three separate YouTube videos. (a,b): gaze has been redirected to new 3D gaze targets. (c): we have modified visual behaviour, making the video subject appear to read a book.*

6.3. Redirecting gaze in YouTube videos

We demonstrate GazeDirector on videos with a variety of eye appearances, head pose, and illumination conditions by redirecting gaze in YouTube videos. We downloaded videos from YouTube and resized them to a resolution of 640×480px. New 3D gaze targets were specified through physics simulations and procedural programming using the Unity engine [Uni]. Figure 14 shows some examples. Please refer to our supplementary video for the full results.

Runtime   GazeDirector runs on a commodity desktop machine (3.3Ghz CPU, Nvidia GTX 1080). Runtime is split between fitting and redirection. We first process the entire video to recover Φ^* for each frame. This model fitting stage ran at 11.6fps, 12.5fps, and 12.1fps for the three YouTube videos in Figure 14. We then redirect gaze for each frame in the video. Gaze redirection is less computationally demanding, and ran at 80fps for each video.

7. Discussion

In this work we described GazeDirector, a novel method for gaze redirection that uses model-fitting. Unlike previous work, GazeDirector does not require person-specific training data, and can redirect eye gaze to new 3D targets explicitly. We fit a parametric eye region model to images using analysis-by-synthesis, minimizing a reconstruction energy to recover shape, texture, pose, gaze, and illumination simultaneously. Gaze redirection is then performed by warping eyelids, and compositing eyeballs onto the output in a photorealistic manner.

Limitations remain. We do not explicitly model a full range of facial expressions such as blinking or squinting. While these expressions can be modelled to a certain extent with the linear shape

model, specific expression blendshapes would assist greatly assist tracking in these cases. Furthermore, we do not handle occlusions or distortion effects from eyeglasses [KRK*16]. Our model does not include the eyelashes; these are hard to model realistically, but can provide an important cue for downwards looking eye gaze. We also do not consider cast shadows from hooded eyes or eyelashes, reducing the quality of our model fit in their presence.

One failure mode is redirecting gaze downwards when little of eyelid is visible, i.e. eyes open wide. When the source eyelid warp region is too small, smudge artefacts are present in the result. This could be addressed with a temporal model of subject eye region appearance. Despite these limitations, we believe our work will enable a range of interesting and novel applications.

Acknowledgements

References

[BB09] BANF M., BLANZ V.: Example-based rendering of eye movements. In *Computer Graphics Forum* (2009). 2

[BBGB16] BÉRARD P., BRADLEY D., GROSS M., BEELER T.: Lightweight eye capture using a parametric model. *ACM Transactions on Graphics (TOG) 35*, 4 (2016), 117. 3

[BBK*15] BERMANO A., BEELER T., KOZLOV Y., BRADLEY D., BICKEL B., GROSS M.: Detailed spatio-temporal reconstruction of eyelids. *ACM Trans. Graph. 34*, 4 (July 2015). 3

[BBN*14] BÉRARD P., BRADLEY D., NITTI M., BEELER T., GROSS M. H.: High-quality capture of eyes. *ACM Trans. Graph. 33*, 6 (2014), 223–1. 3

[BRM16] BALTRUŠAITIS T., ROBINSON P., MORENCY L.-P.: OpenFace: an open source facial behavior analysis toolkit. In *IEEE WACV* (2016). 3, 5, 7

[BV99] BLANZ V., VETTER T.: A morphable model for the synthesis of 3d faces. In *Proc. 26th conf. on Computer graphics and interactive techniques* (1999). 3, 4, 5

[CBZB15] CAO C., BRADLEY D., ZHOU K., BEELER T.: Real-time high-fidelity facial performance capture. *ACM Transactions on Graphics (TOG) 34*, 4 (2015), 46. 3

[CHZ14] CAO C., HOU Q., ZHOU K.: Displaced dynamic expression regression for real-time facial tracking and animation. *ACM Transactions on Graphics (TOG)* (2014). 3

[CSBT03] CRIMINISI A., SHOTTON J., BLAKE A., TORR P. H.: Gaze manipulation for one-to-one teleconferencing. In *Computer Vision, 2003. Proceedings. Ninth IEEE International Conference on* (2003), IEEE, pp. 191–198. 1, 2

[CWW*16] CAO C., WU H., WENG Y., SHAO T., ZHOU K.: Real-time facial animation with image-based dynamic avatars. *ACM Transactions on Graphics (TOG)* (2016). 3

[Eme00] EMERY N. J.: The eyes have it: the neuroethology, function and evolution of social gaze. *Neuroscience & Biobehavioral Reviews 24*, 6 (2000), 581–604. 1

[GBK*14] GIGER D., BAZIN J.-C., KUSTER C., POPA T., GROSS M.: Gaze correction with a single webcam. In *Intl. Conf. on Multimedia and Expo (ICME)* (2014), IEEE. 2

[GKSL16] GANIN Y., KONONENKO D., SUNGATULLINA D., LEMPITSKY V.: Deepwarp: Photorealistic image resynthesis for gaze manipulation. In *European Conference on Computer Vision* (2016), Springer, pp. 311–326. 1, 2, 7

[GTZ*00] GEMMELL J., TOYAMA K., ZITNICK C. L., KANG T., SEITZ S.: Gaze awareness for video-conferencing: A software approach. *IEEE Multimedia* (2000). 2

[GZC*16] GARRIDO P., ZOLLHOEFER M., CASAS D., VALGAERTS L., VARANASI K., PEREZ P., THEOBALT C.: Reconstruction of personalized 3d face rigs from monocular video. 3

[JC04] JONES E. A., CARR E. G.: Joint attention in children with autism theory and intervention. *Focus on autism and other developmental disabilities 19*, 1 (2004), 13–26. 1

[JC16] JENI L. A., COHN J. F.: Person-independent 3d gaze estimation using face frontalization. In *Proc. CVPR Workshops (CVPRW)* (2016), IEEE. 7

[JDvdP12] JIMENEZ J., DANVOYE E., VON DER PAHLEN J.: Photo-realistic eyes rendering. In *SIGGRAPH Talks, Advances in Real-Time Rendering* (2012), ACM. 4, 6

[Kab76] KABSCH W.: A solution for the best rotation to relate two sets of vectors. *Acta Crystallographica Section A: Crystal Physics, Diffraction, Theoretical and General Crystallography 32*, 5 (1976), 922–923. 6

[KACM16] KARIS B., ANTONIADES T., CAULKIN S., MASTILOVIC V.: Digital humans: Crossing the uncanny valley in unreal engine 4. In *GDC Talks* (2016), EPIC. 6

[KL15] KONONENKO D., LEMPITSKY V.: Learning to look up: realtime monocular gaze correction using machine learning. In *Proc. Computer Vision and Pattern Recognition (CVPR)* (2015), IEEE. 2

[KPB*12] KUSTER C., POPA T., BAZIN J.-C., GOTSMAN C., GROSS M.: Gaze correction for home video conferencing. *ACM Transactions on Graphics (TOG)* (2012). 2

[KRK*16] KÜBLER T. C., RITTIG T., KASNECI E., UNGEWISS J., KRAUSS C.: Rendering refraction and reflection of eyeglasses for synthetic eye tracker images. In *Proc. ETRA* (2016), ACM. 8

[KRP*15] KLEHM O., ROUSSELLE F., PAPAS M., BRADLEY D., HERY C., BICKEL B., JAROSZ W., BEELER T.: Recent advances in facial appearance capture. In *Computer Graphics Forum* (2015), Wiley Online Library. 3

[Loo87] LOOP C.: Smooth subdivision surfaces based on triangles. 4

[QLTH15] QIN Y., LIEN K.-C., TURK M., HÖLLERER T.: Eye gaze correction with a single webcam based on eye-replacement. In *Advances in Visual Computing*. Springer, 2015. 2

[RAB*14] RUHLAND K., ANDRIST S., BADLER J., PETERS C., BADLER N., GLEICHER M., MUTLU B., MCDONNELL R.: Look me in the eyes: A survey of eye and gaze animation for virtual agents and artificial systems. In *Eurographics State-of-the-Art Report* (2014), pp. 69–91. 4

[SKR*15] SHARP T., KESKIN C., ROBERTSON D., TAYLOR J., SHOTTON J., KIM D., RHEMANN C., LEICHTER I., VINNIKOV A., WEI Y., ET AL.: Accurate, robust, and flexible real-time hand tracking. In *Proc. 33rd Annual ACM Conference on Human Factors in Computing Systems* (2015), ACM. 6

[SSKS15] SUWAJANAKORN S., SEITZ S. M., KEMELMACHER-SHLIZERMAN I.: What makes tom hanks look like tom hanks. In *Proc. International Conference on Computer Vision (ICCV)* (2015), IEEE. 3

[SSSH16] SHU Z., SHECHTMAN E., SAMARAS D., HADAP S.: Eye-opener: Editing eyes in the wild. *ACM Transactions on Graphics (TOG) 36*, 1 (2016), 1. 2

[SYFN13] SMITH B., YIN Q., FEINER S., NAYAR S.: Gaze Locking: Passive Eye Contact Detection for Human–Object Interaction. In *ACM User Interface Software and Technology (UIST)* (2013). 1, 7

[TZN*15] THIES J., ZOLLHÖFER M., NIESSNER M., VALGAERTS L., STAMMINGER M., THEOBALT C.: Real-time expression transfer for facial reenactment. *ACM Transactions on Graphics (TOG) 34*, 6 (2015), 183. 3

[TZS*16a] THIES J., ZOLLHÖFER M., STAMMINGER M., THEOBALT C., NIESSNER M.: Face2Face: Real-time Face Capture and Reenactment of RGB Videos. In *Proc. Computer Vision and Pattern Recognition (CVPR)* (2016), IEEE. 3, 6

[TZS*16b] THIES J., ZOLLHÖFER M., STAMMINGER M., THEOBALT C., NIESSNER M.: Facevr: Real-time facial reenactment and eye gaze control in virtual reality. *arXiv preprint arXiv:1610.03151* (2016). 1

[Uni] UNITY: Game engine. *Online – http://unity3d.com*. 8

[WBLP11] WEISE T., BOUAZIZ S., LI H., PAULY M.: Realtime performance-based facial animation. In *ACM Transactions on Graphics (TOG)* (2011), vol. 30, ACM, p. 77. 3

[WBM*16a] WOOD E., BALTRUŠAITIS T., MORENCY L.-P., ROBINSON P., BULLING A.: A 3d morphable eye region model for gaze estimation. In *European Conference on Computer Vision* (2016), Springer. 2, 3, 7

[WBM*16b] WOOD E., BALTRUŠAITIS T., MORENCY L.-P., ROBINSON P., BULLING A.: Learning an appearance-based gaze estimator from one million synthesised images. In *Proc. ETRA* (2016), ACM. 4

[WFA10] WOLF L., FREUND Z., AVIDAN S.: An eye for an eye: A single camera gaze-replacement method. In *Proc. Computer Vision and Pattern Recognition (CVPR)* (2010), IEEE. 2

[WK03] WEINER D., KIRYATI N.: Virtual gaze redirection in face images. In *Proc. Conf. Image Analysis and Processing* (2003), IEEE. 2

[WSXC16] WANG C., SHI F., XIA S., CHAI J.: Realtime 3d eye gaze animation using a single rgb camera. *ACM Transactions on Graphics (TOG) 35*, 4 (2016), 118. 3

[YZ02] YANG R., ZHANG Z.: Eye gaze correction with stereovision for video-teleconferencing. In *European Conference on Computer Vision* (2002), Springer, pp. 479–494. 2

[ZGT99] ZITNICK C. L., GEMMELL J., TOYAMA K.: Manipulation of video eye gaze and head orientation for video teleconferencing. *Microsoft Research MSR-TR-99-46* (1999). 1