

UVAGaze: Unsupervised 1-to-2 Views Adaptation for Gaze Estimation

Ruicong Liu, Feng Lu*

State Key Laboratory of VR Technology and Systems, School of CSE, Beihang University, Beijing, China
{liuruicong, lufeng}@buaa.edu.cn

Abstract

Gaze estimation has become a subject of growing interest in recent research. Most of the current methods rely on single-view facial images as input. Yet, it is hard for these approaches to handle large head angles, leading to potential inaccuracies in the estimation. To address this issue, adding a second-view camera can help better capture eye appearance. However, existing multi-view methods have two limitations. 1) They require multi-view annotations for training, which are expensive. 2) More importantly, during testing, the exact positions of the multiple cameras must be known and match those used in training, which limits the application scenario. To address these challenges, we propose a novel 1-view-to-2-views (1-to-2 views) adaptation solution in this paper, the Unsupervised 1-to-2 Views Adaptation framework for Gaze estimation (UVAGaze). Our method adapts a traditional single-view gaze estimator for flexibly placed dual cameras. Here, the “flexibly” means we place the dual cameras in arbitrary places regardless of the training data, without knowing their extrinsic parameters. Specifically, the UVAGaze builds a dual-view mutual supervision adaptation strategy, which takes advantage of the intrinsic consistency of gaze directions between both views. In this way, our method can not only benefit from common single-view pre-training, but also achieve more advanced dual-view gaze estimation. The experimental results show that a single-view estimator, when adapted for dual views, can achieve much higher accuracy, especially in cross-dataset settings, with a substantial improvement of 47.0%. Project page: <https://github.com/MickeyLLG/UVAGaze>.

Introduction

Gaze estimation is a crucial indicator of user attention and has been widely adopted in various applications, including human-robot interaction (Admoni and Scassellati 2017; Terzioğlu, Mutlu, and Şahin 2020; Wang et al. 2015), semi-autonomous driving (Demiris 2007; Majaranta and Bulling 2014; Park, Jain, and Sheikh 2013), medical diagnostics (Castner et al. 2020), and augmented/virtual reality games (Burova et al. 2020; Konrad, Angelopoulos, and Wetzstein 2020; Wang et al. 2020). With the advancements in deep learning, numerous gaze estimation networks (Krafka et al.

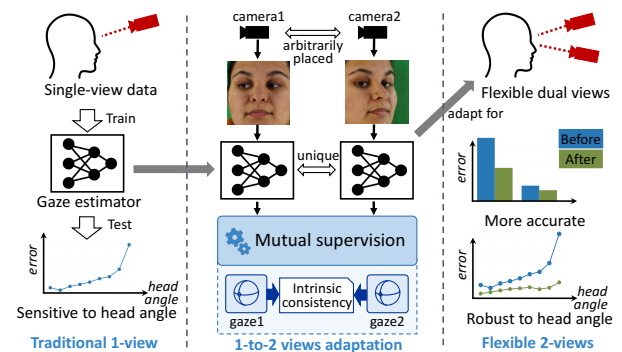


Figure 1: Overview of the proposed unsupervised 1-view-to-2-views adaptation method for gaze estimation. A single-view estimator can be adapted for flexible dual views.

2016; Park, Spurr, and Hilliges 2018; Zhang, Liu, and Lu 2022; Zhang et al. 2017a; Wang and Li 2023) and adaptation/optimization methods (Bao et al. 2022; Cheng, Bao, and Lu 2022; Liu et al. 2021; Park et al. 2019; Yu, Liu, and Odobez 2019) have been proposed in the last decade.

However, the majority of existing gaze estimation methods still use single-view facial images as input. As shown in Tab. 1, although using single-view facial images is simple and effective for learning-based approaches, it poses several challenges. A drawback of traditional single-view gaze estimation is its limited reliable output range for head poses. Empirically and experimentally, when the camera is facing the side of the face, the testing accuracy is significantly reduced due to occlusion and deformation problems.

To address this challenge, a natural solution is to add a second-view camera to better capture the eye appearance. In this paper, we also demonstrate the benefits of using dual cameras. Concurrently, several existing studies have paid attention to gaze estimation in multi-view settings (Arar and Thiran 2017; Kim and Jeong 2020; Lian et al. 2018; Cheng and Lu 2023). However, all these methods require training with at least two views. Their models consistently utilize input images from multiple cameras during both training and testing, predominantly employing feature concatenation or fusion to leverage multi-view information. Such methodologies exhibit two main limitations, as outlined in Tab. 1: 1)

*Corresponding Author

Method	Training	Test	Output
Traditional	1-view (common data)	1-view (arbitrary camera pose)	1-view (limited head pose range)
Two-views training	2-views (multi-view data)	2-views (same camera poses with training)	Select/Average/etc. (extended head pose range)
Ours (Adaptation)	1-view (common data)	2-views (arbitrary camera poses)	Select/Average/etc. (extended head pose range)

Table 1: Comparison of our method with traditional single-view gaze estimation and existing multi-view methods that require at least two-views training (*e.g.*, feature concatenation). “Select/Average/etc.” indicates the strategy to produce the gaze output. In particular, “Select” means selecting one view for the output, “Average” means averaging the predictions from the two views.

They require costly multi-view annotations for training, especially for the feature concatenation or fusion module. 2) Crucially, during testing, the positions and poses of the multiple cameras must be known and match those employed in training, thus limiting the application scenarios.

Unlike those existing methods, this paper proposes a novel unsupervised 1-view-to-2-views adaptation solution. Our approach requires only a pre-trained estimator and unlabeled input images from flexible dual views. By “flexible”, we mean that the dual cameras can be placed arbitrarily, regardless of the pre-training data, without the needing to know their specific extrinsic parameters. As shown in Tab. 1, our method can not only benefit from common single-view pre-training, but also adapt a single-view estimator for more advanced dual-view inference, resulting in reliable outputs in an extended head pose range.

In this paper, an Unsupervised 1-to-2 Views Adaptation framework for Gaze estimation (UVAGaze) is proposed. To adapt a single-view estimator for dual-view inference, we leverage the intrinsic consistency of the gaze directions between the two views. Specifically, our method builds a mutual supervision strategy that dynamically judges the reliability of each view, allowing the predictions from one view to supervise the other. The primary contributions of this paper are summarized as follows.

- For the first time, we propose a novel unsupervised 1-view-to-2-views adaptation (1-to-2 views UVA) framework for gaze estimation. A traditional single-view estimator can be adapted for flexible dual views without gaze annotations or camera extrinsic parameters.
- We design a dual-view mutual supervision strategy for the 1-to-2 views UVA. It utilizes the intrinsic consistency of gaze directions and can dynamically judge the reliability of each view.
- Empirical evaluation highlights the advantage of UVAGaze in adapting a single-view estimator, particularly in cross-dataset scenarios, with a notable 47.0% improvement over traditional single-view methods.

Related Work

Appearance-Based Gaze Estimation

Appearance-based gaze estimation predicts 3D gaze directions from facial images. Various networks have been designed for better accuracy, such as adversarial-learning (Wang et al. 2019), two-eye-asymmetry (Cheng, Lu, and Zhang 2018), and coarse-to-fine (Cheng et al. 2020) structures. Recently, studies (Bao et al. 2022; Cheng, Bao, and Lu




Figure 2: Number of samples and gaze errors with respect to head angles. Left: the model is trained on Gaze360 (Kellnhofer et al. 2019) dataset and tested on ETH-XGaze (Zhang et al. 2020). Right: both the training and testing datasets are ETH-XGaze.

2022; Liu et al. 2021; Kellnhofer et al. 2019) have utilized ResNet (He et al. 2016), inspired by previous work (Zhang et al. 2020) that proves its effectiveness for gaze estimation.

To achieve better accuracy, increasing research is focusing on aspects like cross-dataset performance (Liu et al. 2021; Bao et al. 2022) and generalization (Cheng, Bao, and Lu 2022). However, these studies predominantly approach gaze estimation as a single-view task, ignoring the potential challenges posed by large head angles. In this paper, we will design a method to extend gaze estimation to dual views.

Multi-View Gaze Estimation

Multi-view gaze estimation remains a relatively unexplored research area. Prior work (Gideon, Su, and Stent 2022) learns effective gaze representations from multiple views but can be only used for single-view inference. Some studies (Arar and Thiran 2017; Kim and Jeong 2020; Lian et al. 2018) aim to train more accurate gaze estimators using data from multiple views. However, the aforementioned studies share two drawbacks. 1) They require multi-view annotations to train the models. It is expensive and challenging to collect reliable gaze data under multiple views. 2) During testing, the poses of the multiple cameras are assumed to be known and covered by the training data, hereby limiting their applicability. In contrast, our unsupervised view adaptation (UVA) method eliminates the need for multi-view annotations and is adaptable to arbitrarily placed dual views.

Motivation

Gaze Estimation Using Two Views

Upon analyzing the error of traditional single-view gaze estimation, we have observed the error increases as the head

	Gaze360 → ETH	ETH → ETH
Baseline	18.73	4.75
Select-front	16.03 ▼ 14.4%	4.08 ▼ 14.1%
Average	17.69 ▼ 5.6%	4.05 ▼ 14.7%

Table 2: The gaze estimation results of the two approaches to produce the output based on two views. Angular error is used as the metric.

angle increases (capturing the side of the face), as illustrated in Fig. 2. This phenomenon can be attributed to the fact that large head angles can cause occlusion and deformation. Therefore, introducing a second-view camera could enhance the capture of eye appearance.

As stated in Tab. 1, existing multi-view methods are not practical, since they have limitations in the training data and camera poses. Therefore, we plan to propose a novel unsupervised view adaptation (UVA) method, which is the first to adapt a gaze estimator from 1-view to 2-views.

Output of Two-Views Based Gaze Estimation

As shown in Tab. 1, to calculate the gaze estimation results based on dual views, there are two typical approaches: select-front and average. Given predictions from both views, we can either 1) select the prediction with a smaller head angle as the input for estimation or 2) transform the gaze predictions from both views to the same coordinate system and average them. In Tab. 2, we evaluate both approaches by applying them to a single-view baseline estimator. Specifically, we use the dual-view data ETH-XGaze dataset (Zhang et al. 2020) for testing, and the head pose labels are used for selecting or averaging.

We conduct experiments under two settings: Gaze360 and ETH→ETH, with the only difference being the pre-training dataset. As indicated in Tab. 2, both methods enhance accuracy, confirming the benefit of dual-view information. However, this improvement stems only from merging at the final prediction stage. If we fully harness dual-view data for model optimization and combine it with the two methods, we anticipate even higher accuracy.

One-to-two Views UVA for Gaze Estimation

Based on the discussed ideas, we propose an unsupervised 1-view-to-2-views adaptation task (1-to-2 views UVA). The aim is to adapt a single-view gaze estimator for dual-view inference in an unsupervised manner. Compared with existing multi-view methods, this approach offers two main benefits: 1) It avoids the need for multi-view annotations, reducing data collection costs. 2) Dual cameras can be arbitrarily placed regardless of training data, without knowing their extrinsic parameters, which enhances its practicality. In this way, our method not only benefits from common single-view training, but also achieves advanced dual-view gaze estimation. To achieve this, we design a mutual supervision strategy that uses the dual-view information to optimize the model, and the details are explained in the following section.

UVAGaze: 1-to-2 Views Gaze Estimation

We propose the Unsupervised 1-to-2 Views Adaptation framework for Gaze estimation (UVAGaze), which adapts a single-view estimator for flexibly placed dual cameras. An overview is shown in Fig. 3. The central idea is to exploit the intrinsic consistency of the gaze directions under both views. To ensure the consistency, we also design a head pose stabilization module, supervised by the fixed rotation transformation between the two cameras.

Mutual Supervision Strategy

Forward propagation. The UVAGaze merges the models under both views to be a unique $G(\cdot|\Theta)$, which receives input from the dual-camera. The parameters of the model Θ are pre-trained before adaptation and can predict gaze direction and head pose from an input image \mathbf{x} . The forward function is:

$$\mathbf{g}, (\alpha, \beta, \gamma) = G(\mathbf{x}|\Theta), \quad (1)$$

where \mathbf{g} represents the 3D gaze direction, and (α, β, γ) denote the yaw, pitch, and roll angles of the head pose, respectively. Note that both gaze and head pose are defined in the camera coordinate system. Building on this, the head angle θ is defined as the angle between the z-axes of both the head and camera coordinate systems, *i.e.*:

$$\theta = \langle R(\alpha, \beta, \gamma) \times [0, 0, 1]^T, [0, 0, 1]^T \rangle, \quad (2)$$

where R indicates the rotation matrix, and $\langle \mathbf{a}, \mathbf{b} \rangle$ indicates the angle between \mathbf{a} and \mathbf{b} .

Mutual supervision. Since gaze estimation under dual-view is a stereo problem, the two gaze directions should be identical in the head coordinate system, *i.e.*, theoretically $R_1^T \mathbf{g}_1 = R_2^T \mathbf{g}_2$ (where the subscript indicates the input camera ID). However, directly imposing $R_1^T \mathbf{g}_1 = R_2^T \mathbf{g}_2$ during adaptation may not be optimal. As previously discussed, the reliability of predictions varies with head angles. Therefore, our method dynamically judges the reliability of gaze predictions based on their head pose predictions. The more reliable gaze predictions serve as pseudo-labels to supervise consistency. This process is defined as follows:

$$(p, \hat{\mathbf{y}}) = \begin{cases} (0, R_1^T \mathbf{g}_1), & \theta_1 \leq \theta_2, \\ (1, R_2^T \mathbf{g}_2), & \text{else,} \end{cases} \quad (3)$$

where $\hat{\mathbf{y}}$ is the pseudo-label, and p serves as a flag indicating the more reliable view. For example, $p = 0$ means the view from camera 1 can generate more reliable predictions. Note that the pseudo-label is transformed into head coordinate system by left multiplying with R^T .

Then, the less reliable gaze predictions are supervised by the pseudo-label. The mutual supervision loss function is:

$$\mathcal{L}_{mut} = p \langle R_1^T \mathbf{g}_1, stop(\hat{\mathbf{y}}) \rangle + (1-p) \langle R_2^T \mathbf{g}_2, stop(\hat{\mathbf{y}}) \rangle, \quad (4)$$

where $stop(\cdot)$ cuts off the back propagation of gradient. In this manner, UVAGaze dynamically judges the reliability of both views and enables mutual supervision by leveraging the consistency of gaze directions between both views.

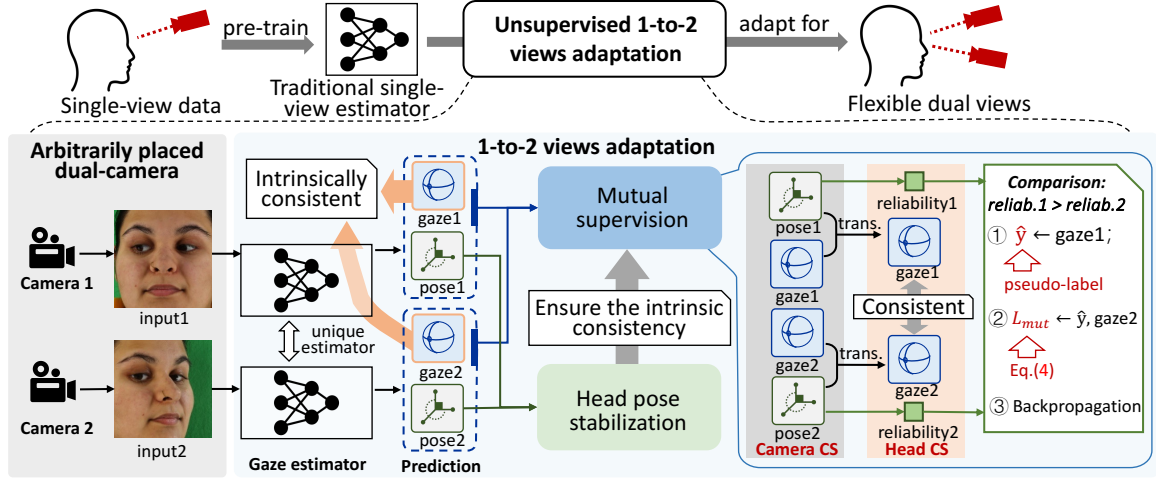


Figure 3: Overview of the proposed UVAGaze, images captured from an arbitrarily placed dual-camera pair are used for UVA. Their gaze and head pose predictions are then used for mutual supervision. In the mutual supervision, gaze predictions are transformed from the camera coordinate system (Camera CS) into the head coordinate system (Head CS) using head pose predictions. Subsequently, with the reliability generated from the head pose predictions, the loss for mutual supervision is calculated. The core idea is to leverage the intrinsic consistency of gaze directions. And to ensure the consistency, a head pose stabilization module is proposed.




Figure 4: The idea of the head pose stabilization module.

Momentum Head Pose Stabilization

Unfortunately, during mutual supervision, the consistency of gaze may fail, and error amplification can occur easily since the model trains itself by using its own predictions. Previous work (Liu et al. 2021) has also described this phenomenon. To ensure consistency, we design a module to stabilize head pose predictions. This is because it is the head pose predictions that are used by mutual supervision to transform the gaze predictions to the head coordinate system.

In this module, as shown in Fig. 4, we use the fixed positional relationship between the dual-camera to stabilize the head pose predictions. Since the positions of the two cameras are fixed, the rotation matrix between the two camera coordinate systems is also constant, *i.e.*, $R_1 R_2^T = \text{constant}$. From this formula, we can draw the following equation:

$$C = f(\alpha_1, \alpha_2, \beta_1, \beta_2) = \sin \beta_1 \sin \beta_2 + \sin \alpha_1 \sin \alpha_2 \cos \beta_1 \cos \beta_2 + \cos \alpha_1 \cos \alpha_2 \cos \beta_1 \cos \beta_2. \quad (5)$$

Eq. (5) computes the value of the (3, 3) element of the rotation matrix $R_1 R_2^T$. We use it to represent the entire matrix and check its stability. Note that the roll angle γ is not introduced, because the effect of γ is eliminated in the normalization step of the input image. Also, its predicted values

are often in the $[-0.01, 0.01]$ interval, having no effect on the task. The normalization step is described in the ‘‘Data Preparation’’ section.

In practice, the input facial image is normalized using an affine transformation matrix W , derived from the head pose and camera focal length. Thus, our method initially performs an inverse transformation on R_1 and R_2 , leading to a revised equation: $(W_1^{-1} R_1)(W_2^{-1} R_2)^T = C$.

To stabilize the head pose predictions using the constant C , we construct a momentum variable during adaptation by temporal averaging. Specifically, for the current iteration T , the variable is denoted as $C^{(T)}$ and can be updated as:

$$C^{(T+1)} = \eta C^{(T)} + (1 - \eta) f^{(T)}(\alpha_1, \alpha_2, \beta_1, \beta_2), \quad (6)$$

where $C^{(T+1)}$ represents the momentum variable for the next iteration. η represents the momentum, which is usually set to 0.99. The momentum variable is employed to stabilize the head pose predictions in the following manner:

$$\mathcal{L}_{stb} = |f^{(T)}(\alpha_1, \alpha_2, \beta_1, \beta_2) - C^{(T)}|. \quad (7)$$

Constraints from the pre-training dataset. Besides stabilizing the head pose predictions, leveraging information from the *pre-training dataset* is also found beneficial. Its information is learned by minimizing the angular error between the gaze predictions \mathbf{g}_{pre} and the labels \mathbf{y}_{pre} .

$$\mathcal{L}_{pre}(\mathbf{x}_{pre}; \Theta) = \langle \mathbf{g}_{pre}, \mathbf{y}_{pre} \rangle. \quad (8)$$

Our method is still unsupervised since the actual adaptation data (dual-view data) is unlabeled.

Total loss. The total loss is computed by summing the aforementioned loss functions, as described below:

$$\mathcal{L} = \mathcal{L}_{mut} + \lambda_1 \mathcal{L}_{stb} + \lambda_2 \mathcal{L}_{pre}, \quad (9)$$

empirically, we set $\lambda_1 = 50$ and $\lambda_2 = 10$.

Algorithm 1: Unsupervised 1-to-2 views adaptation.

Input: Dual-camera input \mathcal{D} , pre-training dataset \mathcal{D}_{pre} and G pre-trained on \mathcal{D}_{pre}

Output: $G(\cdot|\Theta)$

- 1: Initialize: $C^{(1)} \leftarrow G(\mathbf{x}_1, \mathbf{x}_2|\Theta^{(1)})$
- 2: **for** $T \leftarrow 1$ **to** N **do**
- 3: $(\mathbf{x}_{pre}, \mathbf{y}_{pre}), (\mathbf{x}_1, \mathbf{x}_2) \leftarrow \mathcal{D}_{pre}, \mathcal{D}$
- 4: $p, \hat{\mathbf{y}} \leftarrow G(\mathbf{x}_1, \mathbf{x}_2|\Theta^{(T)})$ with Eq. (3).
- 5: $\mathcal{L}_{mut} \leftarrow p, \hat{\mathbf{y}}, G(\mathbf{x}_1, \mathbf{x}_2|\Theta^{(T)})$ with Eq. (4).
- 6: $f^{(T)} \leftarrow G(\mathbf{x}_1, \mathbf{x}_2|\Theta^{(T)})$ with Eq. (5).
- 7: $\mathcal{L}_{stb} \leftarrow f^{(T)}, C^{(T)}$ with Eq. (7).
- 8: $\mathcal{L}_{pre} \leftarrow \mathbf{y}_{pre}, G(\mathbf{x}_{pre}|\Theta^{(T)})$ with Eq. (8).
- 9: Train $G(\cdot|\Theta^{(T+1)})$ with Eq. (9).
- 10: Update $C^{(T+1)}$ with Eq. (6)
- 11: **end for**




Figure 5: The visualization of the postures of the 18 cameras in ETH-XGaze and their captured image samples.

One-to-Two Views Adaptation Procedure

The unsupervised 1-to-2 views adaptation is outlined in Algorithm 1. We start with a pre-trained single-view model $G(\cdot|\Theta^{(1)})$. Labeled images (\mathcal{D}_{pre}) from the pre-training dataset and unlabeled images from the dual cameras (\mathcal{D}) are input. During UVA, model G is adapted by minimizing Eq. (9). The momentum variable C is updated using Eq. (6).

Training details. We use PyTorch on an NVIDIA 3090 GPU. Pre-training employs the Adam optimizer at a learning rate of 10^{-4} , while UVA uses 10^{-5} .

Data Preparation

Dual-camera division. We pre-train using two datasets: ETH-XGaze (Zhang et al. 2020) and Gaze360 (Kellnhofer et al. 2019). But only the ETH-XGaze dataset is for testing, since only it provides multi-view data. We split its 18 cameras into 9 dual-camera pairs (camera 0-9; 1-10; 2-11; ...). This simple division strategy is chosen given the vast number of potential pair combinations. The positions of the 18 cameras are visualized in Fig. 5.

Fine-tuning the ETH-XGaze. The original ETH-XGaze dataset doesn't fully adhere to the intrinsic consistency of gaze directions. That is, after transforming to the head coordinate system, the gaze directions aren't strictly identical.

We find it is due to the inaccurate head pose labels. To rectify this, we refine the head pose labels, starting with the fine-tuning of the head position, denoted as \mathbf{t} :

$$\mathbf{t}_{new} = \frac{1}{18} \sum_{i=0}^{17} \mathbf{t}_i, \quad (10)$$

next, we fine-tune the rotation angles, denoted as \mathbf{h} , by adding a correction term $\Delta\mathbf{h}$, which is determined as:

$$\arg \min_{\Delta\mathbf{h}} \text{var}_{0 \leq i \leq 17} [R^T(\mathbf{h}_i + \Delta\mathbf{h}_i) \times \mathbf{g}_i] + \delta \sum_{i=0}^{17} |\Delta\mathbf{h}_i|, \quad (11)$$

where the var indicates the variance, and the gaze direction \mathbf{g} is calculated as the vector from head position \mathbf{t} to the gaze target point. We have named the fine-tuned dataset as ETH-MV (Multi-View), which is the foundation for all our experiments. The ETH-MV is available at the project page.

Image normalization. We adopt the commonly-used approach (Zhang et al. 2017b) to normalize the data. The normalization aims to ensure consistency in the form of all input images by eliminating head rolls and translating all facial images to the same distance from the camera. To achieve this, the images are wrapped and cropped using an affine transformation matrix W , derived from the head pose and camera focal length.

Experiments

Metric Definition

In this paper, we study the 1-to-2 views UVA task, which adapts a single-view gaze estimator for flexible dual cameras. Therefore, traditional monocular angular error (Mono) metric is insufficient. Inspired by the two approaches presented previously, we propose two new dual-view metrics, select-front error (Dual-S) and average error (Dual-A).

It is important to note that the average error can be significantly affected by the head pose error. To account for this, we also measure the head pose error (HPose) using mean-angle-error for reference.

In summary, we use four metrics to evaluate performance on the dual-view task in this paper, namely:

- Mono: the traditional monocular angular error, which collects all the single-view errors from both views and calculates their average.
- Dual-S: the error under dual-view, calculated by selecting the input image with a smaller head angle.
- Dual-A: the error under dual-view, calculated by averaging the predictions from both views.
- HPose: the mean error across the three angles of head pose (α, β, γ) . Note that this is only a reference metric and not a measure of performance.

One-to-Two Views UVA Results for Flexible Dual Views

In this section, we use the UVAGaze to adapt a single-view estimator for each of the 9 dual-camera pairs from ETH-MV. The estimator is pre-trained on either the Gaze360 (Kellnhofer et al. 2019) or ETH-MV dataset, leading to

Camera pair	Method	Cross-dataset (Gaze360→ETH)				In-dataset (ETH→ETH)			
		Mono	Dual-S	Dual-A	HPosE (ref)	Mono	Dual-S	Dual-A	HPosE (ref)
cam 0,9	Baseline	13.57	11.70	13.90	5.44	3.55	3.37	3.47	1.40
	UVAGaze	8.75 ▼ 35.5%	6.69 ▼ 42.8%	9.45 ▼ 32.0%	8.13 ▲ 49.6%	3.03 ▼ 14.5%	2.68 ▼ 20.6%	2.89 ▼ 16.7%	1.35 ▼ 3.3%
cam 1,10	Baseline	15.06	12.54	15.03	5.49	4.27	3.48	3.87	1.76
	UVAGaze	10.43 ▼ 30.7%	7.11 ▼ 43.3%	10.10 ▼ 32.8%	7.80 ▲ 42.2%	4.05 ▼ 5.2%	2.87 ▼ 17.5%	3.81 ▼ 1.5%	1.69 ▼ 4.4%
cam 2,11	Baseline	18.03	16.61	19.53	7.68	5.66	4.42	5.25	1.60
	UVAGaze	13.13 ▼ 27.2%	10.26 ▼ 38.2%	15.39 ▼ 21.2%	9.76 ▲ 27.1%	4.68 ▼ 17.4%	3.43 ▼ 22.4%	4.20 ▼ 19.9%	1.48 ▼ 7.5%
cam 3,12	Baseline	22.85	18.34	23.20	7.72	4.66	4.54	4.55	1.76
	UVAGaze	14.93 ▼ 34.7%	9.99 ▼ 45.4%	15.48 ▼ 33.3%	8.22 ▲ 6.5%	4.30 ▼ 7.7%	3.79 ▼ 16.5%	4.41 ▼ 3.1%	1.96 ▲ 11.0%
cam 4,13	Baseline	24.27	21.70	24.24	6.06	4.66	4.01	4.10	1.40
	UVAGaze	16.87 ▼ 30.4%	13.39 ▼ 38.3%	17.74 ▼ 26.8%	6.75 ▲ 11.4%	4.09 ▼ 12.1%	3.50 ▼ 12.7%	3.63 ▼ 11.4%	1.36 ▼ 2.7%
cam 5,14	Baseline	24.03	17.55	27.19	9.57	5.42	3.93	5.02	1.56
	UVAGaze	13.76 ▼ 42.7%	8.64 ▼ 50.7%	20.99 ▼ 22.8%	10.26 ▲ 7.3%	4.62 ▼ 14.9%	3.18 ▼ 19.2%	4.39 ▼ 12.6%	1.74 ▲ 11.8%
cam 6,15	Baseline	18.57	17.31	20.61	7.29	5.19	3.81	4.44	1.51
	UVAGaze	14.00 ▼ 24.6%	11.20 ▼ 35.3%	16.62 ▼ 19.4%	8.22 ▲ 12.7%	4.78 ▼ 7.9%	3.24 ▼ 15.0%	4.20 ▼ 5.5%	1.48 ▼ 2.4%
cam 7,16	Baseline	17.09	14.95	17.61	6.77	4.99	3.93	4.37	1.65
	UVAGaze	12.48 ▼ 27.0%	9.98 ▼ 33.2%	12.32 ▼ 30.0%	7.66 ▲ 13.1%	4.26 ▼ 14.7%	3.30 ▼ 16.0%	3.88 ▼ 11.2%	1.81 ▲ 9.3%
cam 8,17	Baseline	15.10	13.82	14.95	5.46	4.36	3.96	4.06	1.63
	UVAGaze	13.05 ▼ 13.6%	12.10 ▼ 12.4%	12.40 ▼ 17.0%	5.93 ▲ 8.8%	4.06 ▼ 6.8%	3.47 ▼ 12.4%	3.83 ▼ 5.6%	1.69 ▲ 4.0%
Overall	Baseline	18.73	16.06	19.58	6.83	4.75	3.94	4.35	1.59
	UVAGaze	13.05 ▼ 30.3%	9.93 ▼ 38.2%	14.50 ▼ 26.0%	8.08 ▲ 18.3%	4.21 ▼ 11.4%	3.27 ▼ 16.9%	3.92 ▼ 9.9%	1.62 ▲ 2.0%

Table 3: Unsupervised 1-to-2 views results for 9 dual-camera pairs. The pairs are from the ETH-MV dataset. In “In-dataset” and “Cross-dataset” settings, the baseline model (ResNet18 (He et al. 2016)) is pre-trained on Gaze360 or ETH-MV, respectively. The adaptation yields a uniquely adapted model for each pair, and “Overall” averages the results across all nine pairs. Note that the head pose error (HPosE) is only for reference.

two distinct tasks: Gaze360-to-ETH and ETH-to-ETH. The same baseline model is adapted to each dual-camera pairs with unlabeled data independently, yielding 9 adapted models. Each adapted model is then tested independently.

The results are presented in Tab. 3. Compared to the pre-trained model (Baseline), our method shows significant improvements in both tasks for all dual-camera pairs, as measured by the three gaze accuracy metrics (Mono, Dual-S, and Dual-A). This indicates that our method is effective for flexibly placed dual camera pairs. There’s a slight increase in the head pose error (HPosE), primarily because our method does not impose strong constraints on it.

Due to the inclusion of the head pose error (HPosE) in the average of gaze predictions, Dual-A is larger than the Dual-S in both tasks. This suggests that using the select-front strategy is optimal when accurate head poses are unobtainable.

One-to-Two Views UVA Results under Different Head Angles

To verify the performance of the UVAGaze across various head angles, we conduct experiments using the overall error across all dual-camera pairs, illustrating accuracy variations with respect to the head angles. As depicted in Fig. 6, the monocular gaze error (Adapt-Mono) of the adapted model increases with increasing head angle. However, there’s a notable improvement in accuracy after adaptation.

Remember, our method is grounded in dual-view inference, allowing us to combine predictions from both views using strategies such as the select-front approach. After applying this strategy, the select-front error (Dual-S, green lines in Fig. 6) is reduced to the level of 0-10° of head angle before adaptation. This indicates that our method can largely




Figure 6: The blue and red lines show monocular gaze errors with respect to the head angle on two tasks before and after adaptation, respectively. The green line shows the errors after selecting the predictions with smaller head angles.

eliminate the negative effects caused by large head angles.

Ablation Study

An ablation study is conducted to evaluate the effectiveness of each component in UVAGaze. The components are:

- Baseline: a standard CNN-based gaze estimation network using architecture of ResNet18 (He et al. 2016).
- mut: the mutual supervision module, which uses the intrinsic consistency of gaze directions under both views to supervise the adaptation.
- stb: the momentum head pose stabilization module, which uses the fixed rotation transformation between the two camera coordinate systems for stabilization.
- pre: the gaze constraint that leverages the information from the pre-training dataset.

We base our experiments on the Gaze360-to-ETH task for clearer observation of differences. Results are shown in

Method	Mono	Dual-S	Dual-A	HPOse
Baseline (w/o mut)	18.73	16.06	19.58	6.83
CNN+mut [†]	57.73	57.09	56.98	13.67
CNN+mut+stb	17.43	14.81	18.62	7.60
CNN+mut+pre	13.33	10.23	15.21	9.95
CNN+mut+stb+pre	13.05	9.93	14.50	8.08

Table 4: Ablation study results of our dual-view adaptation method on the Gaze360-to-ETH task. [†] indicates error amplification occurs.




Figure 7: Comparison of the consistency between our method with/without the head pose stabilization module, lower is better.




Figure 8: Adaptation results of UVAGaze for pre-trained models with different initial head pose accuracy.

Tab. 4. Notably, the results for “CNN+mut” reveal that error amplification occurs if the model uses its own predictions for mutual supervision without any other constraint.

In contrast, with the momentum head pose stabilization module (CNN+mut+stb), the mutual supervision module achieves a significant improvement over the baseline. Additionally, adding the gaze constraint from the pre-training dataset can also stabilize the adaptation process (CNN+mut+pre) and provide additional accuracy improvement (CNN+mut+stb+pre). Ultimately, we select the “CNN+mut+stb+pre” combination as our final version.

Effect of the Head Pose Stabilization Module

The ablation study indicates that the head pose stabilization module can prevent error amplification. Here, we delve deeper into the module’s impact during adaptation.

Our experiments compare the gaze prediction consistency between the two views, both with and without the stabilization module. To measure the consistency, first, we transform the two gaze predictions to the same head coordinate system using the head pose predictions. We then define the angular difference between them as the consistency metric, *i.e.*, $\langle R_1^T \mathbf{g}_1, R_2^T \mathbf{g}_2 \rangle$. The comparison results are shown in Fig. 7. The consistency metric between the two views increases sharply without the stabilization module, highlighting a failure in dual-view consistency. In contrast, the stabilization module ensures consistency even with a slight improvement, indicating why this module prevents error amplification.

System Characteristics Analysis

In our experiments, we observe that the initial head pose error of the pre-trained model has a significant impact on the

	Gaze360 → ETH			ETH → ETH		
	Mono	Dual-S	Dual-A	Mono	Dual-S	Dual-A
Baseline	18.73	16.06	19.58	4.75	4.08	4.35
PnP-GA	26.36	18.60	26.23	4.63	4.10	4.31
RUDA	14.01	10.29	15.40	4.60	3.62	5.31
DAGEN	16.26	13.48	16.58	5.14	4.48	4.75
Gaze360	17.10	12.44	17.53	5.78	4.60	8.07
GazeAdv	15.00	11.12	15.93	9.09	7.89	11.00
UVAGaze (Ours)	13.05	9.93	14.50	4.21	3.40	3.92

Table 5: Comparison with other approaches that have state-of-the-art performance under cross-dataset settings.

final adaptation result. In this section, we compare the adaptation results of UVAGaze for pre-trained models with different initial head pose errors.

We choose several models with different head pose errors from their pre-training phase and then apply the UVAGaze for UVA. The adapted models are then tested on the Gaze360-to-ETH task, as shown in Fig. 8. Clearly, our method enhances accuracy across the board, irrespective of the head pose accuracy variations in the pre-trained model (Baseline: 18.73). Moreover, higher accuracy in head pose estimation directly correlates with higher adaptation results. Hence, in practical, we recommend opting for a pre-trained model with the highest accuracy in head pose estimation.

Comparison under Cross-Dataset Settings

In this section, we compare our method with leading techniques that demonstrate state-of-the-art results in cross-dataset scenarios. Given the practical significance of cross-dataset capability, this experiment is designed to evaluate the capability of our UVAGaze in comparison to leading single-view methods. Specifically, the methods under comparison include PnP-GA (Liu et al. 2021), RUDA (Bao et al. 2022), DAGEN (Guo et al. 2020), Gaze360 (Kellnhofer et al. 2019), and GazeAdv (Wang et al. 2019).

The comparison is shown in Tab. 5. For the Gaze360-to-ETH task, our method has the best performance. Conversely, the PnP-GA method experiences error amplification when applied to the task, and the performance of other methods pales in comparison to our UVAGaze.

In addition, we present results for the ETH-to-ETH task, even though this is an in-dataset task. It is shown that only our method has a consistent improvement across all three metrics. Evidently, UVAGaze continues to surpass the performance of all the existing methods.

Conclusion

In this paper, we introduce the UVAGaze, a groundbreaking 1-view-to-2-views adaptation framework for gaze estimation. This method adapts a traditional single-view gaze estimator for flexible dual views. Central to UVAGaze is a mutual supervision strategy, which takes advantage of the intrinsic consistency of the gaze directions. Notably, our method demonstrates superior performance when adapting a single-view estimator for dual-view scenarios.

Acknowledgments

This work was supported by National Natural Science Foundation of China (NSFC) under Grant 62372019.

References

- Admoni, H.; and Scassellati, B. 2017. Social eye gaze in human-robot interaction: a review. *Journal of Human-Robot Interaction*, 6(1): 25–63.
- Arar, N. M.; and Thiran, J.-P. 2017. Robust real-time multi-view eye tracking. *arXiv preprint arXiv:1711.05444*.
- Bao, Y.; Liu, Y.; Wang, H.; and Lu, F. 2022. Generalizing gaze estimation with rotation consistency. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4207–4216.
- Burova, A.; Mäkelä, J.; Hakulinen, J.; Keskinen, T.; Heinonen, H.; Siltanen, S.; and Turunen, M. 2020. Utilizing VR and gaze tracking to develop AR solutions for industrial maintenance. In *CHI*.
- Castner, N.; Kuebler, T. C.; Scheiter, K.; Richter, J.; Eder, T.; Hüttig, F.; Keutel, C.; and Kasneci, E. 2020. Deep semantic gaze embedding and scanpath comparison for expertise classification during OPT viewing. In *ACM symposium on eye tracking research and applications*, 1–10.
- Cheng, Y.; Bao, Y.; and Lu, F. 2022. Puregaze: Purifying gaze feature for generalizable gaze estimation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, 436–443.
- Cheng, Y.; Huang, S.; Wang, F.; Qian, C.; and Lu, F. 2020. A coarse-to-fine adaptive network for appearance-based gaze estimation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, 10623–10630.
- Cheng, Y.; and Lu, F. 2023. DVGaze: Dual-View Gaze Estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 20632–20641.
- Cheng, Y.; Lu, F.; and Zhang, X. 2018. Appearance-based gaze estimation via evaluation-guided asymmetric regression. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 100–115.
- Demiris, Y. 2007. Prediction of intent in robotics and multi-agent systems. *Cognitive processing*, 8(3): 151–158.
- Gideon, J.; Su, S.; and Stent, S. 2022. Unsupervised multi-view gaze representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5001–5009.
- Guo, Z.; Yuan, Z.; Zhang, C.; Chi, W.; Ling, Y.; and Zhang, S. 2020. Domain adaptation gaze estimation by embedding with prediction consistency. In *Proceedings of the Asian Conference on Computer Vision*.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.
- Kellnhofer, P.; Recasens, A.; Stent, S.; Matusik, W.; and Torralba, A. 2019. Gaze360: Physically unconstrained gaze estimation in the wild. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 6912–6921.
- Kim, J.-H.; and Jeong, J.-W. 2020. A Preliminary Study on Performance Evaluation of Multi-View Multi-Modal Gaze Estimation under Challenging Conditions. In *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems*, 1–7.
- Konrad, R.; Angelopoulos, A.; and Wetzstein, G. 2020. Gaze-contingent ocular parallax rendering for virtual reality. *ACM Transactions on Graphics (TOG)*.
- Krafka, K.; Khosla, A.; Kellnhofer, P.; Kannan, H.; Bhandarkar, S.; Matusik, W.; and Torralba, A. 2016. Eye tracking for everyone. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2176–2184.
- Lian, D.; Hu, L.; Luo, W.; Xu, Y.; Duan, L.; Yu, J.; and Gao, S. 2018. Multiview multitask gaze estimation with deep convolutional neural networks. *IEEE transactions on neural networks and learning systems*, 30(10): 3010–3023.
- Liu, Y.; Liu, R.; Wang, H.; and Lu, F. 2021. Generalizing Gaze Estimation with Outlier-guided Collaborative Adaptation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 3835–3844.
- Majaranta, P.; and Bulling, A. 2014. Eye tracking and eye-based human-computer interaction. In *Advances in physiological computing*, 39–65. Springer.
- Park, H. S.; Jain, E.; and Sheikh, Y. 2013. Predicting primary gaze behavior using social saliency fields. In *Proceedings of the IEEE International Conference on Computer Vision*, 3503–3510.
- Park, S.; Mello, S. D.; Molchanov, P.; Iqbal, U.; Hilliges, O.; and Kautz, J. 2019. Few-shot adaptive gaze estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 9368–9377.
- Park, S.; Spurr, A.; and Hilliges, O. 2018. Deep pictorial gaze estimation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 721–738.
- Terzioğlu, Y.; Mutlu, B.; and Şahin, E. 2020. Designing social cues for collaborative robots: the role of gaze and breathing in human-robot collaboration. In *Proceedings of the 2020 ACM/IEEE International Conference on Human-Robot Interaction*, 343–357.
- Wang, H.; Dong, X.; Chen, Z.; and Shi, B. E. 2015. Hybrid gaze/EEG brain computer interface for robot arm control on a pick and place task. In *2015 37th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, 1476–1479. IEEE.
- Wang, K.; Zhao, R.; Su, H.; and Ji, Q. 2019. Generalizing eye tracking with bayesian adversarial learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11907–11916.
- Wang, L.; and Li, S. 2023. Wheelchair-Centered Omnidirectional Gaze-Point Estimation in the Wild. *IEEE Transactions on Human-Machine Systems*.
- Wang, Z.; Yu, H.; Wang, H.; Wang, Z.; and Lu, F. 2020. Comparing Single-modal and Multimodal Interaction in an Augmented Reality System. In *ISMAR*. IEEE.

Yu, Y.; Liu, G.; and Odobez, J.-M. 2019. Improving few-shot user-specific gaze adaptation via gaze redirection synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11937–11946.

Zhang, M.; Liu, Y.; and Lu, F. 2022. Gazeonce: Real-time multi-person gaze estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4197–4206.

Zhang, X.; Park, S.; Beeler, T.; Bradley, D.; Tang, S.; and Hilliges, O. 2020. ETH-XGaze: A large scale dataset for gaze estimation under extreme head pose and gaze variation. In *European Conference on Computer Vision*, 365–381. Springer.

Zhang, X.; Sugano, Y.; Fritz, M.; and Bulling, A. 2017a. It's written all over your face: Full-face appearance-based gaze estimation. In *Computer Vision and Pattern Recognition Workshops (CVPRW), 2017 IEEE Conference on*, 2299–2308. IEEE.

Zhang, X.; Sugano, Y.; Fritz, M.; and Bulling, A. 2017b. Mpiigaze: Real-world dataset and deep appearance-based gaze estimation. *IEEE transactions on pattern analysis and machine intelligence*, 41(1): 162–175.