

GazeMoE: Perception of Gaze Target with Mixture-of-Experts

Zhuangzhuang Dai^{*1}, Zhongxi Lu², Vincent G. Zakka¹, Luis J. Manso¹, Jose M Alcaraz Calero¹, Chen Li³

¹Dept. of Applied AI and Robotics, Aston University, Birmingham, United Kingdom

²Computing Science, University of Leicester, Leicester, United Kingdom

³Dept. of Materials and Production, Aalborg University, Aalborg, Denmark

*Corresponding address: z.dai1@aston.ac.uk

Abstract—Estimating human gaze target from visible images is a critical task for robots to understand human attention, yet the development of generalizable neural architectures and training paradigms remains challenging. While recent advances in pre-trained vision foundation models offer promising avenues for locating gaze targets, the integration of multi-modal cues—including eyes, head poses, gestures, and contextual features—demands adaptive and efficient decoding mechanisms. Inspired by Mixture-of-Experts (MoE) for adaptive domain expertise in large vision-language models, we propose GazeMoE, a novel end-to-end framework that selectively leverages gaze-target-related cues from a frozen foundation model through MoE modules. To address class imbalance in gaze target classification (in-frame vs. out-of-frame) and enhance robustness, GazeMoE incorporates a class-balancing auxiliary loss alongside strategic data augmentations, including region-specific cropping and photometric transformations. Extensive experiments on benchmark datasets demonstrate that our GazeMoE achieves state-of-the-art performance, outperforming existing methods on challenging gaze estimation tasks. The code and pre-trained models are released at: <https://huggingface.co/zdai257/GazeMoE>.

I. INTRODUCTION

Gaze serves as a fundamental indicator of human intention, interest, and cognitive state, making its accurate estimation critical for a wide range of applications. Gaze plays a forefront role for robots and autonomous systems to understand user engagement [1], driver fatigue [2], consumer interest [3], and autism through identifying abnormal visual attention patterns [4]. Estimation of gaze targets has gained increasing traction in the past decade [5], [6]. Unlike gaze direction analysis, gaze target aims to derive where a person is looking through absorbing the semantics of a scene.

Directly estimating subjects’ attended targets from visible data unlocks more advanced applications such as joint attention detection [7], human-robot interaction [8], or intention prediction [9], and provides the additional benefit of being less invasive than other alternatives such as headsets. Gaze-Follow [10] marks the first sizeable dataset in the community. By leveraging standard cameras and Deep Neural Networks, data-driven methods have grown rapidly since then. However, successes are limited by specific gaze evaluation setting, such as movies [11] or early education [4]. It is yet difficult to use such models to predict human gaze targets reliably in real-world scenarios. In the literature, gaze target has been found highly dependent on four visual cues: (a) visible eye landmarks [12], [10], (b) head pose [13], (c) gestures

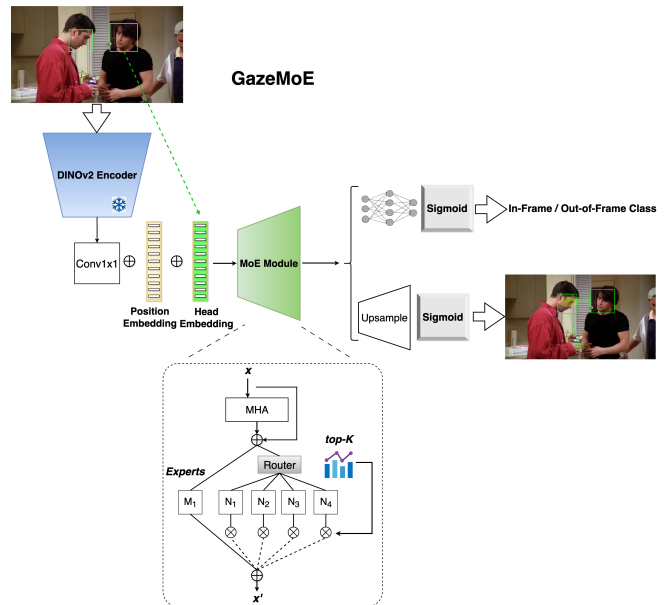


Fig. 1. GazeMoE data flow diagram. A frozen DINOv2 is used to extract fine-grained scene representations alongside the Mixture-of-Experts (MoE) module specialized in selectively routing gaze-related cues.

and geometry of the scene [14], [15], and (d) contextual saliency [10], [11]. In reality, gaze targets can be more diverse than those annotated in datasets, which tend to curate data from lifelogging, Youtube, and other public imaging databases. Moreover, out-of-frame gaze targets, in which a subject fixates at things outside camera’s view, will be more prevalent and unpredictable than in-frame perceivable objects in real-world scenarios [16].

Recently, vision foundation models demonstrated outstanding performance [17], [18] proving that representations of eyes, head poses, gestures, and environmental saliency have been superbly learned in a DINOv2 model [19]. Unfortunately, prior works have not thoroughly investigated an optimized neural architecture of decoder to selectively consume scene-specific cues for effective gaze target evaluation, nor the optimal augmentation and training strategies. How to allow models to adapt to missing features, e.g., due to occlusions or image distortion, remains largely unexplored.

The Mixture-of-Experts (MoE) [20], [21] is a neural network block designed to improve model capacity and efficiency by dynamically routing input data to specialized

sub-networks, i.e., experts [22]. Instead of processing all inputs through the entire network, MoE selectively activates only a subset of experts for each input, enabling scalable and computationally efficient learning. We note that gaze target estimation requires adaptively integrating eyes, head, gesture, and contextual information [11], some of which may or may not be available in different scenes. In this work, we propose an end-to-end model, GazeMoE, leveraging a shared expert and four routed experts that selectively gate input features via ranking independent feedforward networks in transformer decoders as shown in Fig. 1. Each expert learns distinct features which improves representation diversity and model performance on different gaze data distributions. GazeMoE leverages an encoder-decoder architecture [17], [18] with a frozen DINOv2 of ViT_{large} [19] as encoder to extract low-level, gaze-related spatial representation. Two output heads are designed to predict gaze target location as a heatmap of probability and to classify in-frame or out-of-frame cases, respectively. The modular structure of GazeMoE helps avoid overfitting by distributing learning across experts. Since only a few experts are used per input, GazeMoE preserves computational cost while achieving new state-of-the-art performance on various datasets.

Previous methods predominantly use a Binary Cross-Entropy (BCE) loss for in-frame or out-of-frame classification, overlooking class imbalance in common gaze datasets. In GazeMoE, a focal loss is adopted to address class imbalance by focusing more on hard-to-classify minority samples. To maximize the performance as well as generalization, GazeMoE encompasses a full stack of augmentations, including random cropping, horizontal flipping, head prompt jittering, and photometric changes (colour jittering, random grayscaling, sharpness and contrast). It has been found that such augmentation techniques are essential to optimize training. GazeMoE achieves state-of-the-art performance on all common benchmarks. Importantly, it shows excellent robustness and generalization in out-of-distribution visual scenes, such as fisheye lens imaging [23] and children gaze estimation [4].

The main contributions of this work are as follows. (1) We propose GazeMoE, an end-to-end gaze target estimator leveraging Mixture-of-Experts modules that produces state-of-the-art results on a variety of benchmarks; (2) We show optimal training strategies including the introduction of a custom class-balancing loss and a suite of augmentations; (3) Extensive experiments demonstrate the superiority of our proposed GazeMoE in accuracy, robustness, and generalization. Particularly, GazeMoE is the first-ever attempt that demonstrates reliable gaze target evaluation across both fisheye lens imaging and standard lens imaging domains.

II. RELATED WORK

Data-driven Gaze Target Evaluation. Ever since Re-casens *et al.* [10] curated the GazeFollow challenge – a large scale dataset aimed at locating subjects’ gaze target in 2D images, researchers have been pushing the boundaries of gaze target estimation in terms of precision [7], [11], [24], [18],

[17], generalization [3], [4], efficiency [11], [17], as well as 3D scene semantics [25], [26], [24]. Traditional methods typically focus on fusing human facial and gestural features with scene context through multiple encoder pathways [11], [25] or with the aid of multi-channel sensing modalities [27]. They realize successes in common benchmarks constructed for specific scenarios, such as TV shows [7] or kids behaviour monitoring [4], whereas, engendering trade-offs in generalization, complexity of multi-objective training, and runtime efficiency.

Ryan [17] and Song [18] recently showed that a large pre-trained vision foundation model, like DINOv2 [19], is effective at encoding gaze target-related information. A frozen pre-trained encoder combined with a decoder deliver state-of-the-art performance across many benchmarks and generalized well for out-of-distribution data. In GazeLLE [17], head positions are incorporated as positional embeddings without the need for a separate learnable pathway. A compact, transformer-based decoder can jointly locate gaze targets and classify whether a target is in or outside a frame [11]. GazeTarget360 [16] is recently proposed – a system leveraging multi-scale fusion in decoder – for efficient gaze target evaluation incorporating eye contact detection. However, the gaze target localization precision is suboptimal.

Training Paradigm. Different learning targets have been explored to train a gaze target estimator. Using L2 loss to supervise predicted heatmaps is adopted by most works [10], [7], [25], [24], [18] in the community. Some recent works [23], [17] found pixel-wise Binary Cross-Entropy loss more effective in deriving fine-grained gaze target. Cosine distance of predicted gaze vectors or head poses are utilized as auxiliary losses which prove effective to guide the learning process for many [7], [4], [24], [18].

What remains to be answered is what the optimal decoder network is and its training paradigm. Recently, Large-Language Models benefit tremendously from MoE architectures which allow dynamic input routing [22]. Would a MoE-based decoder bring more robustness to gaze target estimation which depends on selective scene features? What is the best loss function? Data augmentation plays a crucial role in computer vision but none of the prior work has concurred in the best augmentations for gaze target tasks. In this work, we are motivated to address these questions.

III. METHOD

Given an input image $I \in \mathbb{R}^{h \times w \times 3}$, gaze target estimation aims to localize the gaze attended position in pixel coordinates, and discern whether the target is inside or outside the camera’s field of view. Suppose there are multiple subjects in the scene, their head bounding boxes, $bbox_i = [x_{min}, y_{min}, x_{max}, y_{max}]_i$, form the other input pathway of head prompts. The gaze target task can be expressed as

$$y, hm = F_{\theta}(I, bbox) \quad (1)$$

where y is the probability of an in-frame target; hm is a heatmap containing target location probabilities within the frame I ; and $\theta(\cdot)$ represents all parameters of a neural network.

A. Mixture-of-Expert Module

We notice visual scenes of gaze target tasks vary enormously necessitating versatile representation learning of the scene depending on availability of different visual cues. Previous methods predominately exploit handcrafted cues including eye appearance [18], head poses [13], or a conical field-of-view from a head adhering to straight-line propagation of light [25]. However, eyes and head poses may be occluded when subjects look away from camera [10] imposing high uncertainties. The rule-of-thumb that eyes and gaze target should be line-of-sight connected no longer holds true under imaging distortion, especially in panorama images where sphere-to-plane projection brings severe distortion [23]. Moreover, traditional methods often find performance degradation in different subject groups like children or autism patients [4]. Recent work has shown that a separate encoder pathway for head features is unnecessary, in which head positional embeddings suffice [17]. A convolutional layer with 1×1 kernel can be used to pull features from a frozen DINOv2 encoder which gives robust representations of human gaze [18]. Subsequently, these latent features can be decoded with standard transformer blocks.

To this end, we propose an MoE module to enhance feature feed-forward propagation in transformer decoder blocks, as illustrated in Figure 2. The MoE module can adaptively integrate shared and scenario-specific (selectively routed) knowledge based on the input image [22], [28]. This allows for segregation of routed information, i.e., the four pivotal gaze target cues – eye appearance, head poses, gestures, and contextual saliency, when some are available and others are not across different scenes. MoE acts as a strong and efficient gated decoder to select only relevant cues as contender for gaze target evaluation whilst pruning visual noises. Mixing shared experts and routed experts enable pathways for both common knowledge and conditional features which stabilize model learning.

Given the self-attention scores acquired by Multi-Head Attention (MHA) in each transformer block, input to the feed-forward network (FFN), $x \in \mathbb{R}^{B \times Seq \times d_{model}}$, will be flattened before the MoE module, where B is batch size and Seq is sequence length. Each shared expert, E_j^{shared} , will contribute to shared features,

$$x_{shared} = \frac{1}{M} \sum_{j=1}^M E_j^{shared}(x) \quad (2)$$

$$E_j^{shared}(x) = GeLU(W_{j,1} \cdot x)W_{j,2} \quad (3)$$

where E_j^{shared} are MLP with weights $W_{j,1} \in \mathbb{R}^{d_{model} \times d_h}$ and $W_{j,2} \in \mathbb{R}^{d_h \times d_{model}}$ with a GeLU activation in between; M is the number of shared experts; d_h represents hidden layer dimension of the shared expert.

The N routed experts, E_i^{routed} , take an identical shape as each shared expert. Routed experts will go through a gating mechanism where the gating network $W_g \in \mathbb{R}^{d_{model} \times N}$ will select top- K routed experts to use for each input. The gating weights can be expressed as,

$$g(x) = \sigma(W_g \cdot x) \in \mathbb{R}^N \quad (4)$$

where σ is Softmax function. The top- K selection reads;

$$\text{Top-}K \text{ indices: } \{i_1, \dots, i_K\} = \text{argsort}(g(x))[: K] \quad (5)$$

$$\text{Normalized weights: } w_k = \frac{g(x)_{i_k}}{\sum_{k'=1}^K g(x)_{i_{k'}}} \quad (6)$$

The routed feature is comprised of weighted sum of routed experts' outputs;

$$x_{routed} = \sum_{k=1}^K w_k \cdot E_{i_k}^{routed}(x) \quad (7)$$

Final output of the MoE module can be written as;

$$x' = x_{shared} + x_{routed} \quad (8)$$

In this work, we set $N = 4$ routed experts to selectively capture the four gaze target cues identified in previous work; $M = 1$ shared expert for passage of common scene features; each expert with $d_{model} = 256$ hidden size in MLP layers; and $K = 2$ meaning only the top-2 weighted experts will be used in the calculations for optimized performance and efficiency. Detailed validation of these parameters is discussed later in the Ablation Study section.

B. Heatmap Loss and Auxiliary Focal Loss

Previous works mostly adopted heatmap L2, i.e., Mean-Squared Error (MSE) loss [29], [7], [26], sometimes combined with gaze angular loss [25] for training. Other research reveal pixel-wise BCE loss produces satisfactory performance [17], [23] raising concerns about an optimal loss setting. The BCE loss treats each pixel as an independent binary classification. This aligns well with the heatmap's role of modelling probability distribution across the gaze target region. Our experiments show strong evidence that pixel-wise BCE loss is less sensitive to the magnitude of probabilistic error than MSE which may overly penalize peak (0.9 prediction vs. 1.0 ground-truth gaze) or valleys (0.1 prediction vs. 0 ground-truth gaze). Given gaze target is taken as the max probability coordinate, the absolute probability values do not matter as much. Thus, BCE loss generally handles noisy or imperfect target heatmaps better.

As for in/out classification, the vanilla BCE loss can be expressed as;

$$\text{BCE}(p, y) = -y \log(p) - (1 - y) \log(1 - p) \quad (9)$$

where p denotes probability of an in-frame target and y is the ground-truth label. We introduce a focal loss component to allow better robustness for in/out-of-frame

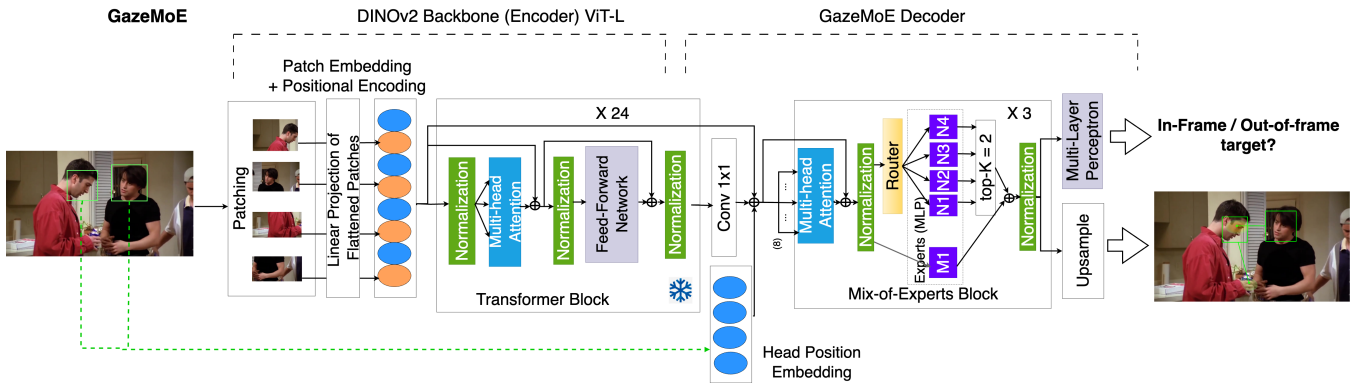


Fig. 2. GazeMoE architecture. Given an input image, the model predicts whether a person’s gaze target is in-frame or out-of-frame and where target is. A frozen DINOv2 ViT-L backbone is used to extract fine-grained scene representations. The GazeMoE decoder with three Mix-of-Experts blocks is specialized in selectively routing gaze-target-related features.

classification tasks of imbalanced classes, which is typical in related datasets. For instance, VideoAttentionTarget [11] contains approximately 40% out-of-frame samples, while ChildPlay [4] has fewer than 10% out-of-frame samples. To address class imbalance, let

$$p_t = \begin{cases} p & \text{if } y = 1 \\ 1 - p & \text{if } y = 0 \end{cases} \quad (10)$$

then, the focal loss can be written as

$$\mathcal{L}_{focal} = \alpha \cdot (1 - p_t)^\gamma \cdot \text{BCE}(p, y) \quad (11)$$

where α is the class weighting factor to handle class imbalance and γ is the focusing parameter. We set α as the ratio between minor class and major class in our experiments and γ is fixed at 2.0. This allows enhanced learning of minority classes and difficult samples.

The final loss function is expressed as,

$$\mathcal{L} = \mathcal{L}_{heatmap} + \lambda \cdot \mathcal{L}_{focal} \quad (12)$$

where $\mathcal{L}_{heatmap}$ is the per-pixel BCE loss of the gaze target heatmap.

Through extensive testing, we demonstrate focal loss acts as a powerful auxiliary loss to enable smoother convergence of the in/out classification. The focal loss regulates the latent representation of the decoder’s feature map which complements the pixel-independent heatmap prediction to improve overall performance.

C. Augmentation

Prior approaches [11], [17], [18] indicate three commonly used input augmentations, including 1) random image cropping which randomly zooms in an image whilst preserving head bounding boxes and gaze target within the frame, 2) horizontal flipping, and 3) random jittering of head bounding boxes. Note that data for gaze target tasks are often acquired from well-lit and high-fidelity visible data such as COCO dataset and movies. We hypothesize more imaging augmentation techniques are needed to enhance model robustness

and generalization. We introduce an additional suite of photometric augmentations, namely, 4) colour jittering, 5) random grayscaling, 6) auto-contrasting, and 7) random sharpness adjustment. Experiment results validate the significance of these augmentation steps.

IV. EVALUATION

We conduct experiments on GazeFollow, VideoAttention-Target, ChildPlay, GazeFollow360, EYEDIAP datasets, and ablation studies to unveil GazeMoE’s optimum setting and runtime performance. We evaluate GazeMoE and comparative methods using the following metrics: area under the curve (AUC) of the probability heatmap, mean error distance (Mean L2) of peak gaze location, average precision (AP) of in-frame and out-of-frame classification, and a proprietary spherical distance for 360-degree images of GazeFollow360 [23]. All experimental results are derived from an average of five repeated training campaigns of different seeds.

The input image resolution is set to 448×448 . The output heatmap is of 64×64 in consistency with previous works. We use three transformer blocks in the decoder with 8 self-attention heads. The models are first pre-trained on GazeFollow for 15 epochs with batch size 60 and a $1e - 3$ learning rate. Then, the models are fine-tuned on respective datasets with batch size 36, $1e - 3$ as the learning rate for the in/out classification head and $1e - 5$ as the learning rate for other layers. We utilize *Adam* optimizer, a 0.1 dropout, and a cosine learning rate scheduler.

A. Comparative Study

GazeFollow. We pre-train the GazeMoE on GazeFollow [10] first, a large-scale dataset containing over 110K training samples from diverse visual scenes. The recommended full stack of augmentations are applied in conjunction with per pixel BCE loss. As is seen from Table I, the GazeMoE demonstrates competitive performance with state-of-the-art AUC. Note the advances of GazeMoE is not significant compared to Gaze-LLE models. This is primarily because GazeFollow is known to collect data samples from well-lit,

TABLE I
METHODS COMPARISON ON GAZE FOLLOW AND VIDEO ATTENTION TARGET (VAT) DATASETS.

Method	Learnable Params	GazeFollow		VideoAttentionTarget		
		AUC \uparrow	Mean L2 \downarrow	AUC \uparrow	Mean L2 \downarrow	AP $_{in/out}$ \uparrow
Human Expert	-	0.924	0.096	0.921	0.051	0.925
Recasens et al. [10]	50M	0.878	0.19	-	-	-
Chong et al. [7]	51M	0.896	0.187	0.833	0.171	0.712
Chong et al. [11]	61M	0.921	0.137	0.86	0.134	0.853
Fang et al. [25]	68M	0.922	0.124	0.905	0.108	0.896
Hu et al. [27]	61M	0.923	0.128	0.88	0.118	0.881
Gupta et al. [30]	35M	0.943	0.114	0.914	0.11	0.879
Tafasca et al. [4]	25M	0.939	0.122	0.914	0.109	0.834
ESCNet [26]	29M	0.928	0.122	0.885	0.12	0.869
Gaze-LLE $_{ViT-B}$ [17]	2.8M	0.956	0.104	0.933	0.107	0.897
Gaze-LLE $_{ViT-L}$ [17]	2.9M	0.958	0.099	0.937	0.103	0.903
GazeMoE	3.4M	0.959	0.101	0.939	0.097	0.917

TABLE II
EVALUATION RESULTS ON CHILDPLAY DATASET. \dagger WE TEST RELEASED GAZE-LLE MODELS ON THE TEST SET THAT THE AUTHORS OF CHILDPLAY DATASET SUGGESTED.

Method	ChildPlay		
	AUC \uparrow	Mean L2 \downarrow	AP $_{in/out}$ \uparrow
Gupta et al. [30]	0.919	0.113	0.983
Tafasca et al. [4]	0.935	0.107	0.986
Gaze-LLE $_{ViT-B}^{\dagger}$ [17]	0.939	0.123	0.992
Gaze-LLE $_{ViT-L}^{\dagger}$ [17]	0.942	0.111	0.993
GazeMoE	0.945	0.106	0.994

high-fidelity imaging datasets like COCO and Flickr. Selective feature routing and sharing is therefore not remarkable on this benchmark.

VideoAttentionTarget (VAT). Chong et al. [11] expanded the challenge by curating a new gaze target benchmark - VideoAttentionTarget (VAT) - to simultaneously classify in-frame and out-of-frame targets. We conduct supervised fine-tuning on VAT for 10 epochs combining heatmap BCE loss and the auxiliary focal loss in Eqn. 12. We set $\lambda = 1$. As shown in Table I, GazeMoE outperforms prior methods demonstrating state-of-the-art performance with a significant margin over prior methods.

ChildPlay. Tafasca et al. [4] recently created the ChildPlay dataset containing over 250K instances of child-to-child and child-to-adult gaze interactions from diverse Youtube sources. Previous datasets mainly focus on adult gazes, whereas, children’s gaze can be tricky to understand due to the particular stage of their cognitive development. Besides in-frame and out-of-frame, we note that ChildPlay also annotates other gaze classes, such as *gaze-shift*, *eyes-closed*, and *uncertain*, accounting for 10% of total data. We exclude other classes before testing GazeMoE in comparison to existing methods in Table II. We set $\lambda = 0.1$ and train for 3 epochs. Again, our method represents state-of-the-art solution to ChildPlay.

GazeFollow360. GazeFollow360 contains 360-degree images with omnidirectional field-of-views. These images un-

dergo heavy distortion brought by sphere-to-plane projection raising difficulties for gaze target evaluation. The original work [23] used sophisticated multiple pathways to diminish the assumption of light-of-sight in the frame. We follow the same split of 8,225 images for training, about 900 for validation, and 900 for testing. We conduct fine-tuning on GazeFollow360 with a batch size of 36 and a $1e - 4$ learning rate. The models converge after 10 epochs. As shown in Table III, the evaluation results give a strong indication that GazeMoE adapts well to out-of-distribution data, surpassing previous methods on AUC and approaching human-level performance. This verifies that a pre-trained DINOv2 excels in gaze target estimation regardless of line-of-sight hypothesis.

EYEDIAP zero-shot inference. The EYEDIAP [35] dataset curates videos of human participants continuously gazing at a floating ping-pong ball in a laboratory. We evenly sample 50 frames from these videos and construct a dataset of 1,750 samples including 38.6% out-of-frame images where the ping-pong ball is beyond camera’s view. Because participants always fixated at the ball, it is unsuitable for supervised fine-tuning where models quickly overfit to attend to the ball. We conduct zero-shot testing as shown in Table IV. Again, GazeMoE delivers state-of-the-art performance. The MoE architecture allows the model to adaptively integrate gaze-related cues in challenging unseen scenarios.

TABLE III
EVALUATION RESULTS ON GAZE FOLLOW 360 DATASET.

Method	GazeFollow360	
	Spherical Dist↓	AUC↑
Human Expert	0.2531	0.935
Random	1.5357	0.5056
Lian et al. [31]	1.254	0.6057
SALICON [32]	1.094	0.8002
Chong et al. [7]	0.9183	0.7765
Ruiz et al. [33]	0.7612	0.7188
Zhang et al. [34]	0.7366	0.7213
Li et al. [23]	0.6067	0.8104
Gaze-LLE $_{ViT-L}$ [17]	0.759	0.9197
GazeMoE	0.771	0.9232

TABLE IV
EVALUATION RESULTS ON EYEDIAP DATASET.

Method	EYEDIAP		
	AUC↑	Mean L2↓	AP $_{in/out}$ ↑
Gaze-LLE $_{ViT-B}$	0.617	0.411	0.725
Gaze-LLE $_{ViT-L}$	0.596	0.421	0.614
Gaze-LLE $_{ViT-B}$ + head prompt	0.597	0.423	0.73
Gaze-LLE $_{ViT-L}$ + head prompt	0.593	0.431	0.662
GazeMoE	0.618	0.312	0.702

TABLE V
ABLATION STUDY ON DIFFERENT LOSS FUNCTIONS.

Loss	Type	GazeFollow		VideoAttentionTarget		
		AUC↑	Mean L2↓	AUC↑	Mean L2↓	AP $_{in/out}$ ↑
Heatmap Head	BCELoss (GazeMoE)	0.959	0.101	0.939	0.097	0.917
	MSELoss + KL-D	0.955	0.108	0.938	0.106	0.909
In/out Head	Focal Loss (GazeMoE)	-	-	0.939	0.097	0.917
	BCELoss	-	-	0.938	0.098	0.912



Fig. 3. Qualitative results. We use *dlib* face detector to automate head bounding boxes prompting as green frames.

B. Ablation Study

Loss function. We conduct extensive experiments to validate GazeMoE settings. We prove that BCE loss for *hm* with an auxiliary focal loss for the in/out classification is

optimal (Table V), though MSE loss with a 0.05 weighted KL-divergence (KL-D) loss can still lead to convergence.

MoE configuration. We investigated the effects of MoE configuration including vanilla decoder without MoE (*FFN*

TABLE VI
EVALUATION ON MOE SIZE.

Model	GazeFollow	
	AUC \uparrow	Mean L2 \downarrow
GazeMoE _{FFN only}	0.955	0.105
GazeMoE _{1shared, 2mlp}	0.957	0.101
GazeMoE _{1shared, 4mlp}	0.956	0.103
GazeMoE _{2shared, 1mlp}	0.958	0.101
GazeMoE _{1shared, 1mlp}	0.959	0.101

TABLE VII
EVALUATION OF AUGMENTATIONS.

Augmentations	GazeFollow	
	AUC \uparrow	Mean L2 \downarrow
No Augmentation	0.954	0.111
Cropping	0.957	0.101
Photometric	0.956	0.103
Cropping + Photometric	0.959	0.101

TABLE VIII
RUNTIME PERFORMANCE.

Method	Latency (ms)	Memory (MB)
GT360 [16]	70.3	689
Gaze-LLE _{ViT-L}	72.5	930
GazeMoE	74.2	984

only), ratios between expert hidden size and transformer’s d_{model} (k *mlp*), and the number of shared experts (m *shared*) in Table VI. Performance boost brought by MoE module is apparent in comparison to vanilla feedforward decoders, GazeMoE_{FFN only}.

Augmentation. We verify the efficacy of the proposed suite of augmentations as shown in Table VII and at the same time verify the robustness of the proposed GazeMoE. Random cropping in junction with photometric augmentations are essential to optimize the model though the MoE decoder already pushes performance to a high level without augmentations.

C. Qualitative Analysis

We showcase the efficacy of GazeMoE across various datasets as can be seen from Figure 3. GazeMoE works robustly in multi-person, diverse scenes including outdoor sports, TV advertisements, children learning environments, different public and private locations. It also performs well identifying in-frame vs. out-frame gazing (with a red colour bounding box over the face). The example where the algorithm has not performed well is EYEDIAP where there are extreme gaze angles plus it is zero-shot.

D. Inference Evaluation and Implementation

We evaluate GazeMoE’s runtime performance with an Intel Core i7 processor and an NVIDIA GeForce RTX 3080 GPU as shown in Table VIII. Per-sample latency of GazeMoE equals 74.2 ms, comparable to that of Gaze-LLE_{ViT-L} [17] and GT360 [16] which claimed to be an

efficient solution. This allows GazeMoE performing at 13 frames per second, which is sufficient for a number of applications such as human-robot interaction. GazeMoE’s peak memory usage is 984 MB. Although it is greater than Gaze-LLE_{ViT-L} and much more than GT360 at peak, GazeMoE realizes state-of-the-art performance making it an appealing solution in terms of precision and robustness.

GazeMoE is implemented using Ubuntu 22.04 vanilla with PyTorch 2.5.1 with NVIDIA CUDA support enabled. CUDA library required is v11.8 or higher.

V. CONCLUSION

We propose a novel Mixture-of-Experts decoder for human gaze target estimation in imaging data. Based on this approach, we introduce GazeMoE, which leverages a unique combination of shared and routed experts for adaptive gaze feature learning under different scenarios. Our proposed model is the first to utilize a specialized Mixture-of-Experts structure in gaze target tasks. In all the datasets being evaluated - GazeFollow, VideoAttentionTarget, ChildPlay, GazeFollow360, EYEDIAP - our model outperforms existing methods to the best of our knowledge, which shows the potential to become a new benchmark for future gaze target research. The runtime performance of GazeMoE clearly sheds light on capabilities of deploying GazeMoE in robotic systems for real-time processing.

In future work, our research will focus on investigating and developing methods that exploit temporal information and contextual cues to handle extreme head poses and partially visible faces in real-world scenarios.

ACKNOWLEDGEMENT

Experiments were run on Aston Engineering and Physical Science Machine Learning Server, funded by the EPSRC Core Equipment Fund, Grant EP/V036106/1. The authors would like to thank to UKRI EPSRC EP/Y028813/1 “National Edge AI Hub for Real Data: Edge Intelligence for Cyberdisturbances and Data Quality.” The authors would like to acknowledge support by Villum Experiment project (00058627) and the Aston Research Flourish Together Support Fund.

REFERENCES

- [1] Thorsten Schodde, Laura Hoffmann, Sonja Stange, and Stefan Kopp. Adapt, explain, engage—a study on how social robots can scaffold second-language learning of children. *J. Hum.-Robot Interact.*, 9(1), dec 2019.
- [2] Fan Liu, Delong Chen, Jun Zhou, and Feng Xu. A review of driver fatigue detection and its advances on the use of rgb-d camera and deep learning. *Engineering Applications of Artificial Intelligence*, 2022.
- [3] Henri Tomas, Marcus Reyes, Raimarc Dionido, Mark Ty, Joel Casimiro, Rowel Atienza, and Richard Guinto. Goo: A dataset for gaze object prediction in retail environments. In *CVPR Workshops (CVPRW)*, 2021.
- [4] Samy Tafasca, Anshul Gupta, and Jean-Marc Odobez. Childplay: A new benchmark for understanding children’s gaze behaviour. In *IEEE International Conference on Computer Vision (ICCV)*, October 2023.
- [5] Zhengxi Hu, Kunxu Zhao, Bohan Zhou, Han Guo, Shichao Wu, Yuxue Yang, and Jingtai Liu. Gaze target estimation inspired by interactive attention. *IEEE Transactions on Circuits and Systems for Video Technology*, 32:8524–8536, 2022.

- [6] Gang Liu, Yu Yu, Kenneth Alberto Funes Mora, and J. Odobez. A differential approach for gaze estimation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43:1092–1099, 2019.
- [7] Eunji Chong, Nataniel Ruiz, Yongxin Wang, Yun Zhang, Agata Rozga, and James M. Rehg. Connecting gaze, scene, and attention: Generalized attention estimation via joint modeling of gaze and scene saliency. In *Computer Vision – ECCV 2018: 15th European Conference, Munich, Germany, September 8–14, 2018, Proceedings, Part V*, page 397–412, Berlin, Heidelberg, 2018. Springer-Verlag.
- [8] Zhuangzhuang Dai, Jinhua Park, Aleksandra Kaszowska, and Chen Li. Detecting worker attention lapses in human-robot interaction: An eye tracking and multimodal sensing study. In *IEEE 28th International Conference on Automation and Computing (ICAC)*, 2023.
- [9] Min Wu, Tyron Louw, Morteza Lahijanian, Wenjie Ruan, Xiaowei Huang, Natasha Merat, and Marta Kwiatkowska. Gaze-based intention anticipation over driving manoeuvres in semi-autonomous vehicles. In *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 6210–6216, 2019.
- [10] Adria Recasens*, Aditya Khosla*, Carl Vondrick, and Antonio Torralba. Where are they looking? In *Advances in Neural Information Processing Systems (NIPS)*, 2015. * indicates equal contribution.
- [11] Eunji Chong, Yongxin Wang, Nataniel Ruiz, and James M. Rehg. Detecting attended visual targets in video. *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5395–5405, 2020.
- [12] Eunji Chong, Elysha Clark-Whitney, Audrey Southerland, Elizabeth Stubbs, Chanel Miller, Eliana L. Ajodan, Melanie R. Silverman, Catherine Lord, Agata Rozga, Rebecca Merrill Jones, and James M. Rehg. Detection of eye contact with deep neural networks is as accurate as human experts. *Nature Communications*, 11, 2020.
- [13] Shivam Gupta and Jean-Marc Odobez. A modular multimodal architecture for gaze target prediction: Application to privacy-sensitive settings. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 1650–1659. IEEE, 2022.
- [14] Haofei Wang, Jimin Pi, Tong Qin, Shaojie Shen, and Bertram E. Shi. Slam-based localization of 3d gaze using a mobile eye tracker. In *Proceedings of the 2018 ACM Symposium on Eye Tracking Research & Applications*, ETRA '18, New York, NY, USA, 2018. Association for Computing Machinery.
- [15] Nora Horanyi, Linfang Zheng, Eunji Chong, Aleš Leonardis, and Hyung Jin Chang. Where are they looking in the 3d space? In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, pages 2678–2687, June 2023.
- [16] Zhuangzhuang Dai, Vincent Gbouna Zakka, Luis J. Manso, and Chen Li. Gazetarget360: Towards gaze target estimation in 360-degree for robot perception. *arXiv preprint arXiv:2507.00253*, 2025. Cross-listed in Human-Computer Interaction (cs.HC).
- [17] Fiona Ryan, Ajay Bati, Sangmin Lee, Daniel Bolya, Judy Hoffman, and James M Rehg. Gaze-llc: Gaze target estimation via large-scale learned encoders. *arXiv preprint arXiv:2412.09586*, 2024.
- [18] Yuehao Song, Xinggang Wang, Jingfeng Yao, Wenyu Liu, Jinglin Zhang, and Xiangmin Xu. Vitgaze: Gaze following with interaction features in vision transformers. *Visual Intelligence*, 2(31), 2024.
- [19] Maxime Oquab, Timothée Darcet, Theo Moutakanni, Huy V. Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, Russell Howes, Po-Yao Huang, Hu Xu, Vasu Sharma, Shang-Wen Li, Wojciech Galuba, Mike Rabbat, Mido Assran, Nicolas Ballas, Gabriel Synnaeve, Ishan Misra, Herve Jegou, Julien Mairal, Patrick Labatut, Armand Joulin, and Piotr Bojanowski. Dinov2: Learning robust visual features without supervision, 2023.
- [20] Saeed Masoudnia and R. Ebrahimpour. Mixture of experts: a literature survey. *Artificial Intelligence Review*, 42:275 – 293, 2012.
- [21] Hien Nguyen and Faicel Chamroukhi. Practical and theoretical aspects of mixture-of-experts modeling: An overview. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 8:e1246, 02 2018.
- [22] Weilin Cai, Juyong Jiang, Fan Wang, Jing Tang, Sunghun Kim, and Jiayu Huang. A survey on mixture of experts in large language models. *IEEE Transactions on Knowledge and Data Engineering*, PP:1–20, 01 2025.
- [23] Yunhao Li, Wei Shen, Zhongpai Gao, Yucheng Zhu, Guangtao Zhai, and Guodong Guo. Looking here or there? gaze following in 360-degree images. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 3742–3751, October 2021.
- [24] Francesco Tonini, Nicola Dall’Asen, Cigdem Beyan, and Elisa Ricci. Object-aware gaze target detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 21860–21869, 2023.
- [25] Yi Fang, Jiapeng Tang, Wang Shen, Wei Shen, Xiao Gu, Li Song, and Guangtao Zhai. Dual attention guided gaze target detection in the wild. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11390–11399, June 2021.
- [26] Jun Bao, Buyu Liu, and Jun Yu. Escnet: Gaze target detection with the understanding of 3d scenes. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 14106–14115, 2022.
- [27] Zhengxi Hu, Dingye Yang, Shilei Cheng, Lei Zhou, Shichao Wu, and Jingtai Liu. We know where they are looking at from the rgb-d camera: Gaze following in 3d. *IEEE Transactions on Instrumentation and Measurement*, 71:1–14, 2022.
- [28] Anke Tang, Li Shen, Yong Luo, Nan Yin, Lefei Zhang, and Dacheng Tao. Merging multi-task models via weight-ensembling mixture of experts. In *Proceedings of the 41st International Conference on Machine Learning, ICML’24*. JMLR.org, 2024.
- [29] Shubham Dokania, Vasudev Singh, and Shuaib Ahmed. Gazehell: Gaze estimation with hybrid encoders and localised losses with weighing. In *35th British Machine Vision Conference 2024, BMVC 2024, Glasgow, UK, November 25-28, 2024*. BMVA, 2024.
- [30] Anshul Gupta, Samy Tafasca, and Jean-Marc Odobez. A modular multimodal architecture for gaze target prediction: Application to privacy-sensitive settings. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, pages 5041–5050, June 2022.
- [31] Dongze Lian, Zehao Yu, and Shenghua Gao. Believe it or not, we know what you are looking at! In *Asian Conference on Computer Vision*, 2018.
- [32] Xun Huang, Chengyao Shen, Xavier Boix, and Qi Zhao. Salicon: Reducing the semantic gap in saliency prediction by adapting deep neural networks. In *2015 IEEE International Conference on Computer Vision (ICCV)*, pages 262–270, 2015.
- [33] Nataniel Ruiz, Eunji Chong, and James M. Rehg. Fine-grained head pose estimation without keypoints. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 2155–215509, 2018.
- [34] Xucong Zhang, Yusuke Sugano, Mario Fritz, and Andreas Bulling. It’s written all over your face: Full-face appearance-based gaze estimation. In *2017 IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 2299–2308, 2017.
- [35] Kenneth Alberto Funes Mora, Florent Monay, and Jean-Marc Odobez. Eyediap: a database for the development and evaluation of gaze estimation algorithms from rgb and rgb-d cameras. In *Proceedings of the Symposium on Eye Tracking Research and Applications*, ETRA '14, page 255–258, New York, NY, USA, 2014. Association for Computing Machinery.