

This ICCV workshop paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

# MAAD: A Model and Dataset for "Attended Awareness" in Driving

Deepak Gopinath\*

Guy Rosman<sup>†</sup> Si

Brenna Argall\*

Simon Stent<sup>†</sup> K

Katsuya Terahata<sup>#</sup>

John Leonard<sup>†</sup>

Luke Fletcher<sup>‡</sup>

\*Department of Mechanical Engineering, Northwestern University, Evanston, Illinois, † Toyota Research Institute, Cambridge, MA, #Woven Planet Holdings, Tokyo, Japan, ‡ Motional, Cambridge, MA.

Corresponding Author: deepakgopinath@u.northwestern.edu

# Abstract

We propose a computational model to estimate a person's attended awareness of their environment. We define "attended awareness" to be those parts of a potentially dynamic scene which a person has attended to in recent history and which they are still likely to be physically aware of. Our model takes as input scene information in the form of a video and noisy gaze estimates, and outputs visual saliency, a refined gaze estimate and an estimate of the person's attended awareness. In order to test our model, we capture a new dataset with a high-precision gaze tracker including 24.5 hours of gaze sequences from 23 subjects attending to videos of driving scenes. The dataset also contains third-party annotations of the subjects' attended awareness based on observations of their scan path. Our results show that our model is able to reasonably estimate attended awareness in a controlled setting, and in the future could potentially be extended to real egocentric driving data to help enable more effective ahead-of-time warnings in safety systems and thereby augment driver performance. We also demonstrate our model's effectiveness on the tasks of saliency, gaze calibration and denoising, using both our dataset and an existing saliency dataset. We make our model and dataset available at https://github.com/ToyotaResearchInstitute/att-aware/.

# 1. Introduction

We define "attended awareness" to be those parts of a potentially dynamic scene which a person has attended to in recent history and which they are still likely to be physically aware of.<sup>1</sup> While this measure of a person's aware-



Figure 1: Model for Attended Awareness in Driving (MAAD) overview. Our model takes as input a video of a scene, as seen by a person performing a task in the scene, along with noisily registered ego-centric gaze sequences from that person. The model estimates (i) a saliency heat map, (ii) a refined gaze estimate, and (iii) an estimate of the subject's attended awareness of the scene. We evaluate our model using a unique annotated dataset of third-person estimates of a driver's attended awareness. By explicitly estimating a person's attended awareness from noisy measurements of their gaze, MAAD can improve human-machine interactions. In the driving example, such interactions might include safety warnings in situations where attended awareness is deemed insufficient.

ness is difficult to objectively measure, in certain situations such as driving, it is possible to infer, at least to some useful degree, the attended awareness of a driver. Driving instructors routinely assess a driver's behavior based on their estimated attended awareness of a given driving scene, and provide real-time feedback to ensure safety. In the context of human-machine interaction, inferring human attended awareness in a given scenario is valuable for machines to facilitate seamless interaction and effective interventions. Human gaze information can be used by machines for this purpose and many models have been developed to relate scene understanding and overt attention as saliency estimates [7, 10, 15, 22, 41, 42] and objectness measures [1, 6]. However, further exploration of the link between visual at-

<sup>&</sup>lt;sup>1</sup>Attended awareness will also be referred to as awareness in the rest of the paper for the sake of brevity.

tention and processes of scene understanding and decision making has been constrained by a limited ability to reason about such cognitive processes, which is difficult in general contexts [63, 53, 49].

In this work, we propose a model to estimate attended awareness as a spatial heatmap. The input to our model is a video of the scene under observation, and a noisy estimate of the person's gaze. Our model allows us to (i) compute visual saliency for the scene, (ii) leverage this to refine the noisy gaze estimate, and (iii) combine this information over time to infer the person's attended awareness. The construction of our model is driven by axiomatic considerations that define an image translation network from the scene image to the saliency and awareness heat maps. Within a larger human-machine interaction scenario, the model becomes useful as an inference engine for human awareness enabling seamless interaction (for e.g., a person operating a semi-autonomous car or a smart wheelchair). We adopt a data-driven approach (as opposed to more limited-scope analytical models [30]) to allow for scalable and more comprehensive modeling of attended awareness that can potentially rely on supervision from multiple sources (such as resulting actions/self-reporting).

**Contributions** 1) We propose a learned model that affords estimation of attended awareness based on noisy gaze estimates and scene video over time. 2) We further demonstrate how the model affords saliency estimation as well as the refinement of a noisy gaze signal. 3) We present a new dataset that explores the gaze and perceived attended awareness of subjects as they observe a variety of driving and cognitive task conditions. While the dataset is captured via a proxy hazard awareness task rather than through real or simulated driving, it serves as a useful starting point to study visual saliency and awareness in driving scenes.

# 2. Related Work

Our work builds on prior work in visual saliency estimation, situational awareness, and driving-specific exploration. We briefly summarize related work from these areas.

**Visual saliency.** Vision scientists have long sought to model and understand the mechanisms behind our allocation of attention in a visual scene [63]. Visual salience is a function of many factors, including the spatio-temporal nature of the stimulus itself as well as its relationship to neighboring stimuli and the nature of the visual system perceiving it [23]. The evolution of computational approaches to estimating visual image salience has mirrored that of other popular computer vision problems, with largely hand-engineered models—designed to mirror certain bottom-up and top-down attention processes in the human visual system (e.g. [8, 23])—giving way to supervised, data-driven approaches (e.g. [22, 24]). Unlike image saliency, video saliency models consider spatial as well as temporal infor-

mation to detect objects of interest in a dynamic way.

Spatio-temporal saliency estimation opens inquiry into how processes of visual attention, situational awareness and task-related decision making are connected. Previous attempts have been made to computationally model situational awareness (see e.g. [4, 38, 59]). Our approach to modeling situational awareness is unique in that we try to explicitly estimate the parts of the visual scene to which a person has attended using a spatio-temporal model for gaze and scene understanding. The three stages of forming situational awareness consists of perception, comprehension and projection [19] and in our work we focused primarily on perception. We aim to model, from noisy observations of a person's scan path, the set of objects and scene structures which that person is likely to have attended to, and therefore might be better able to incorporate into their future decision-making. While we note that peripheral vision alone can achieve visual awareness in many settings [58], we focus on objects of fixation, since we are concerned primarily with estimating when drivers fail to notice potential driving hazards, which are known to strongly induce fixations [16].

**Data-Driven Saliency Modeling and Datasets.** Datadriven approaches to image and video saliency rely on stateof-the-art deep learning architectures: CNNs [10, 22, 41], GANs [42] and LSTMs [3, 15, 56], and single image inputs through to multi-stream inputs incorporating video, along with optical flow, depth or semantic segmentation estimates [33, 40] and even additional modalities such as audio [55]. In our work, similar to [39], we adopt a 3D CNN–based approach, due to its simplicity and success on other video understanding tasks such as action recognition.

While the majority of saliency datasets explore image and video saliency under controlled viewing conditions (e.g. [32, 37, 56]), in recent years, numerous ego-centric video gaze datasets have been developed in which subjects perform tasks as varied as cooking and meal preparation [34], playing video games [8] and driving [40], in parallel to developments in applications of saliency (see, *e.g.* [51, 45, 11]. On the other hand in the driving domain, in-lab data collection procedures have been extensively adopted as they have the advantages of high accuracy, repeatability and the ability to focus on rare scenarios such as critical situations and accidents [61, 21, 17, 2], where understanding human perception can inform safety systems approaches.

Our dataset uses an in-lab data collection paradigm, however it differs from prior work for several reasons. Firstly and most notably, we capture multiple subjects observing the same visual stimuli under *different cognitive task modifiers*. Our dataset therefore allows for reasoning about the effect of different cognitive task modifiers on the visual gaze patterns, given identical visual input. Secondly, we provide annotations for third party estimates of a subject's attended awareness, based on observations of their scan path. For this purpose, we devise a novel reproducible annotation scheme. Finally, as our dataset is gathered using a high precision gaze tracker with a chin rest, the precision of the scan paths is extremely high when compared to that of eye-glass gaze tracking ego-centric datasets such as [40].

Driving specific applications. Driving is a predominantly visual task. The first studies into driver attention and eye scanning patterns date back over half a century [50]. Since then, certain driving behaviors have been well established and modelled, such as the "perceptual narrowing" effect in which drivers increasingly fixate on the road ahead as task demands increase (e.g. through greater speed, increased traffic or lack of familiarity with a route) [20], or the benefits of understanding driver attention when predicting a driver's future intent [18, 60]. However, to the best of our knowledge, no models exist with the purpose of quantitatively estimating a driver's spatial awareness of a scene. In recent years, works such as [40, 61] have used highprecision gaze trackers to create video gaze datasets both in the car and in the lab, allowing for data-driven approaches to modelling driver attention. While we make use of the roadfacing data from [40] in our experiments, our model differs in one key respect. Rather than estimating visual saliency from video alone, we demonstrate how, given access to a noisy gaze estimate of a driver it is possible to simultaneously estimate scene saliency, a denoised gaze signal and an estimate of the driver's overall awareness of the scene.

# 3. Method

The computational model we propose is guided by several assumptions related to human attended awareness and its relation to gaze patterns (Section 3.1). These assumptions are implemented through a mixture of explicit objectives and behavioral regularizations (Section 3.3).

# 3.1. Assumptions and Supervisory Cues

We use several assumptions about gaze patterns and attended awareness to define the priors in training our model:

- A1 Saliency: Gaze tends to focus on specific regions [7], both object and stuff [31].
- A2 Attended awareness: People tend to become aware of the objects they look at [52, 12]. Their attention is, however, limited in its capacity.
- A3 Awareness decay: Awareness of an object can decrease (due to forgetting), or increase (when looking at an object) at different rates [46].
- A4 Regularity of gaze, awareness: Gaze and awareness maps should be regular and smooth unless otherwise warranted [9, 25].

• A5 Awareness and motion: As an observer moves through a dynamic scene, their awareness moves along with objects and regions in the scene and exhibits temporal persistence. [35, 5].

# 3.2. Model

We define the *Model for Attended Awareness in Driving* (*MAAD*), shown in Figure 2, as a fully convolutional image encoding-decoding network with shortcut links such as U-Net [47]. Sensor images are encoded and decoded into a latent feature map, M(x,t), from which two convolutional modules emit the estimated gaze distribution  $p_G(x,t)$ , and the awareness image  $M_A(x,t)$ .

The gaze distribution is normalized as a probability density function (via a softmax operator). We note that  $p_G(x, t)$ is a unified notation for gaze probability maps with and without a noisy gaze input from an external gaze tracker. In the absence of a noisy gaze input,  $p_G(x, t)$  is a probabilistic form of saliency. In the rest of the paper we use  $p_G(x, t)$  to denote both forms to simplify the notation. The individual modules in the decoder are fed additional information: the (noisy) driver gaze measurement over time in the form of Voronoi maps, and optical flow maps [40] encoded as two feature layers for horizontal and vertical motion.

### **3.3.** Loss Function Design

At training time, the gaze and awareness maps are used to compute several supervisory and regularization terms, whose design is guided by the assumptions in Section 3.1.

#### 3.3.1 Supervisory Terms

**Gaze Prediction.** We want  $p_G$  to predict a subject's gaze as accurately as possible. This is encouraged via the primary data term:

$$\mathcal{L}_{G} = -\sum_{t} \sum_{x \in X_{g}(t)} \log p_{G}(x, t),$$
(1)

where  $X_q(t)$  are the 2D ground truth gaze points at time t.

**Perceived Awareness.** We include a supervisory source for awareness estimation. This term surrogates awareness estimation training by a perceived awareness estimate. One approach to obtain a perceived awareness estimate is to provide a person with the gaze estimate of the driver overlaid on the road scene and query how aware the driver is of specific locations in the image at particular points in time. This is further described in Section 4. The cost term reads:

$$\mathcal{L}_{\text{ATT}} = \sum_{(x,t) \in \text{labeled}} \left( M_A(x,t) - L_A(x,t) \right)^2, \quad (2)$$

where the summation is over all annotated samples in location x at time t, and  $L_A$  denotes the annotated measure of awareness in the range [0, 1] as described in Section 4.2.



Figure 2: **Model architecture.** During training, our model takes as input a video sequence with associated scan paths from a gaze tracker, along with optical flow estimates. The video sequence is processed using a 3D convolutional encoder/decoder, and information from gaze, and optic flow is transformed and injected during the decoding stage. The model outputs two heat maps: a gaze probability density heat map,  $p_G$ , which provides a clean estimate of the noisy gaze input, and an attended awareness heat map,  $M_A$ , which provides an estimate of the driver awareness, both at the final frame of the input video sequence. For further details, see Section 3.2.

Awareness of Attended Objects. Based on (A2), we add a term that encourages awareness to be high when a person is gazing at a scene location:

$$\mathcal{L}_{AA} = \sum_{t} \sum_{x \in X_g(t)} (M_A(x, t) - 1)^2.$$
(3)

#### 3.3.2 Regularization Terms

**Spatial Smoothness.** Based on (A4-5), we added regularity terms to both the awareness and gaze maps:

$$\mathcal{L}_{\mathbf{S},\cdot} = \int \frac{|\nabla \phi|^2}{\sqrt{|\nabla I|^2 + \epsilon}} dx, \tag{4}$$

where  $\phi$  is  $M_A$  and  $p_G$  for  $\mathcal{L}_{S,A}$  and  $\mathcal{L}_{S,G}$  respectively and the integral is computed over all pixels. This regularization is a variant of anisotropic diffusion [44, 43] with crossdiffusivity based on the scene image I.

**Temporal Smoothness.** Based on (A3-5), we apply also temporal regularization for awareness:

In order to make the map temporally consistent with respect to the locations and not just object boundaries, we use a smoothness / decay term based on the image optical flow:

$$\mathcal{L}_{\rm T} = \sum_{x,t} f_{w_{\rm OF}} \left( M_A(x + v_{\rm OF}(x), t+1), M_A(x,t) \right), \quad (5)$$
$$f_{w_{\rm OF}}(a,b) = c_1 \left( (a - w_{\rm OF}b)_+ \right) + c_2 \left( (a - w_{\rm OF}b)_- \right)^2,$$

where  $v_{OF}(x)$  is the optical flow computed on the input images and  $w_{OF} < 1$  is a weight factor that is close to 1. ()<sub>+</sub> and ()<sub>-</sub> denote positive and negative change respectively in awareness values when going from t to t + 1. Particularly,  $f_{w_{OF}}(a, b)$  is set to be an *asymmetric* loss function that penalizes awareness that increases instantaneously when attending to an object less compared to awareness that decreases rapidly via forgetting. This is accomplished by having the forgetting term, ()<sub>-</sub>, to be quadratic.

Awareness Decay. Based on (A3), we expect the level of awareness to decay over time, which we model via:

$$\mathcal{L}_{\text{DEC}} = \sum_{x,t} \begin{pmatrix} (1 - \epsilon_{\text{DEC}}) M_A(x,t) \\ -M_A(x,t+1) \end{pmatrix}^2, \quad (6)$$

where  $\epsilon_{\text{DEC}}$  is a decay factor.

Attention Capacity. Based on (A2), we expect that the cognitive resources available to the driver do not change over time. This assumption captures the fact that the overall awareness should be similar between frames on average.

$$\mathcal{L}_{\text{CAP}} = \sum_{t} \left( \sum_{x} M_A(x,t) - \sum_{x} M_A(x,t+1) \right)^2.$$
(7)

**Block-level consistency.** We denote by  $M_A(x, t; t_1)$  the awareness estimate at (x, t) that is emitted based on the training snippet started at  $t_1$ , similarly for  $p_G$ . We define a consistency term [48] between consecutive estimates via

the loss:

$$\mathcal{L}_{\text{CON},\cdot} = \sum_{t_1,t} \sum_{x} \left( \phi(x,t;t_1) - \phi(x,t;t_1+1) \right)^2, \quad (8)$$

where  $\phi$  is  $M_A, p_G$  for  $\mathcal{L}_{\text{CON},A}, \mathcal{L}_{\text{CON},G}$  respectively.  $\mathcal{L}_{\text{CON},\cdot}$  helps to minimize the difference between the predictions of multiple passes of samples at the same timestamps through the network. The overall training loss is a linear combination of the previously described loss terms. See supplementary material for additional details.

#### 3.4. Training Procedure

We trained our model using PyTorch 1.7 using NV100 GPUs. The training was carried out using video input subsequences of length 4 seconds, sampled at 5Hz with frames resized to  $240 \times 135$ . Our model was trained using Adam optimizer [29] with a learning rate of  $5 \times 10^{-3}$  and a batch size of 4. The model weights were updated after the gradients were aggregated for a fixed number of steps in order to reduce the variance in the loss. The batch aggregation size was set at 8. The first two layers of the encoder (kept frozen during training) are taken from a pretrained Resnet18 implementation in Torchvision [36]. Later spatio-temporal encoding is done by three layers of separable 3D convolution modules denoted as S3D in Figure 2 [62]. To provide optical flow estimates, we used RAFT pre-trained on FlyingChairs and FlyingThings datasets [54].

The decoder consists of stacked decoder units each of which receives input from three sources a) side-channel information b) skip connections from the encoder layer and c) the output of the previous decoder unit, when available. Each decoder unit consists of two submodules: 1) The skip connections are first processed via an S3D module whose output is then concatenated (channel-wise) with the sidechannel information and the output of the previous decoder unit. This concatenated input is processed by another S3D module followed by bilinear upsampling that brings the output to the proper resolution for the next decoder unit.

The decoder emits a latent map M which is subsequently processed by two separate convolutional models to emit a gaze probability map,  $p_G$  and an awareness heatmap denoted as a  $M_A$ . The softmax in the gaze convolutional module ensures that the gaze probability map is a valid probability distribution. More details regarding network architecture and training to be found in supplementary material.

# 4. Dataset Description

Our complete dataset comprises approximately 24.5 hours of gaze tracking data captured via multiple exposures from different subjects to 6.2 hours of road-facing video drawn from the DR(eye)VE dataset [40]. We concentrate our gaze capture on repeated exposures of downtown (as

opposed to highway and rural) driving scenes (77%) and daylight scenes (90%), since these contain the most diverse visual scenarios. While the original DR(eye)VE dataset captured and registered gaze to the road-facing video using a driver head-mounted eye tracker and feature-based matching for homography estimation, accumulating significant measurement errors, we opted for an in-lab experiment. In-lab experiment offers several advantages such as higher accuracy and repeatability across subjects and cognitive task conditions. Furthermore, models trained on in-lab data has already been shown to be effective when tested on in-the-wild data [61]. We measure gaze to an extremely high precision  $(0.15^{\circ} - 0.50^{\circ})$  typical accuracy,  $0.01^{\circ}$  RMS resolution, 0.05° microsaccade resolution) using a towermounted SR Research EyeLink 1000 Plus tracker. A main novelty of the dataset is that in addition to high precision gaze, we supplement the dataset with third party annotations of perceived awareness. The annotation procedure is outlined in Section 4.2. We recruited 23 subjects (aged 20-55), who each watched a subset of video clips with their heads mounted in a chin-rest after a 9-point calibration procedure. The subjects all carried US driving licenses and had at least two years of driving experience. Their primary task was to monitor the driving scene as a safety driver might monitor an autonomous vehicle. While not a perfect substitute for in-car driving data collection, this primary task allowed for the capture of many of the characteristics of attentive driving behavior. In order to explore the effect of the cognitive task difference (vs. in-car data) on the gaze and awareness estimates, subjects viewed the video under different cognitive task modifiers, as detailed in Section 4.1 (data collected with non-null cognitive task modifiers comprise 30% of total captured gaze data). Around 45% of video stimuli were watched more than once, of which 11% (40 minutes) was observed by 16 or more subjects.

#### 4.1. Cognitive Task Modifiers

Although our dataset is collected in an in-lab experiment, we are still interested in subjects' behavior under a different task (in-car driving, engaged as well as distracted). We therefore included in our experiments several secondary cognitive task modifiers to artificially alter participant behavior in a way which might mimic certain variations in the behavior of drivers in real conditions. We aimed at modifiers that affected visual search patterns, but did not explicitly instruct the participants toward specific search targets. These modifiers affect the visual patterns by either changing the goal of visual search or by changing the scene appearance. The cognitive task modifiers were as follows:

**1. Null condition:** The subjects were given the task of supervising the driving of the car, looking for and flagging possible risky events or obstacles.

2. Blurred condition: Same as 1, but stimulus videos



Figure 3: What is the driver aware of? Our dataset is unique in containing repeated views of the same driving video by 23 subjects performing the same proxy supervisory task but under different task modifiers: (i) null (green), i.e. with no secondary condition, (ii) with blurred video (yellow), (iii) with vertically flipped video (red), (iv) fixating on the road only (dark blue), and (v) reading occasional on-screen text (light blue) (see Section 4.1 for more details). Here we overlay a sub-sampled one second gaze history per subject for a given frame in a video sequence. From observation of the scan paths of subjects, it is evident that some are aware and fixate strongly on the risk in the scene (e.g. crossing vehicles or pedestrians), while others fixate fleetingly or, due to the presence of cognitive task modifiers, are unaware of the risk.

were blurred with a Gaussian kernel corresponding to  $N \deg$  of visual field, making scene understanding and therefore the supervisory task harder, and affecting the visual search pattern.

3. Vertical-flip condition: Same as 1, but stimulus videos were flipped upside down, making scene understanding and therefore the supervisory task counter-intuitive, and affecting the visual search pattern.

4. Road-only condition: Same as 1, but subjects were asked to only fixate on road structure and not on dynamic obstacles such as cars and people.

5. Reading-text condition: Same as 1, but stimulus videos were overlaid with snippets of text of approximately even length at random image locations for P seconds at an average interval of Q seconds. Subjects were asked to read each text snippet while supervising the driving.

4.2. Annotations

We annotated 53,983 sequences of approximately 10 seconds sampled randomly from within the data for attended awareness. While inferring awareness is difficult and subjective, and probing the subject's working memory directly is impossible, we devised a reproducible annotation scheme to explore a third person's estimate of a person's attended awareness. Our annotation protocol leverages the fact that humans are able to develop a theory of mind of other peoples' mental states from cues inherent in eve gaze [28]. While the labels provided (as in any manually annotated dataset) are imperfect and have certain limitations, they are an important first step towards data-driven modeling of awareness. Annotators watched a video snippet where the subject's gaze was marked by two circles centered at the gaze point. One circle (green) size was set to the diameter of a person's central foveal vision area (2 degrees) at the viewing distance. Another circle (red) was set to a diameter four times the foveal vision circle. At the end of the video snippet, a random location was chosen and the annotators were asked whether they believe the subject has attended to that location on a scale between 1 and 5 (1-no, definitely not aware, 5-yes, definitely aware). Three different sampling types (object, edges, and non-objects) were used for sampling the final location. The annotations were linearly transformed to [0,1] in  $L_A$  and provided a supervisory signal that the network (awareness predictor) tried to match. Annotators were asked whether the cursor corresponded to a well-defined object, whether they would expect to be aware of the location if they were driving and how surprised they by the appearance/behavior of the highlighted region. The annotations had good coverage across awareness levels and contained sufficient number of examples of both highly aware as well as unaware examples. Figure 4 shows frames from an example annotation video. More annotation details and statistics in supp. material.

# 5. Results

We now demonstrate the results from our model on several tasks of interest such as saliency, gaze refinement and attended awareness estimation. Our model, while applied here to a dataset which isn't strictly egocentric, could be straightforwardly extended to a true egocentric setting.

#### 5.1. Saliency

In order to confirm that our model is capable of estimating visual saliency, we trained it on the DR(eye)VE dataset images and splits [40]. As our approach assumes individual gaze samples per frame as the primary supervisory cue, we sampled gaze points from the fixation maps provided in the original dataset for every frame. We generated  $p_G$ , and compared it against the ground truth gaze map. For the



Figure 4: **Annotation example.** Left, center: Frames from the sequence with gaze estimates, along with the estimate of the central foveal vision area, and a larger red circle to mark what is probably outside the central vision area. Annotators were guided to treat these as approximations of the subject's attended gaze. Right: Annotators were asked to say whether the subject has attended to the red cross region and is aware of it. See the supp. mat. for additional details on the annotation.

standard measures of KL, cross-correlation and information gain (see [40] for details), we obtained 1.734, 0.565, and -0.0002 respectively, comparing favorably to other algorithms tested on that dataset, such as [14, 57].

# 5.2. Gaze refinement

In the following experiments, we show how our model is able to correct imperfect gaze estimates. In driving, driver gaze estimates can be obtained using camera-based driver monitoring systems (DMS). However, these estimates can be highly noisy due to the inherent process noise introduced during gaze tracking, or biased due to calibration errors in the DMS [26]. Our hypothesis is that knowledge of where the person might look in the scene can help the model refine noisy and biased (miscalibrated) coarse estimates. We describe two experiments that were conducted to address these typical modes of failure. For all experiments herein, we trained our model with a training dataset encompassing 8 urban day-time videos from our dataset with the highest subject and task coverage and adopted a fixed 80%/20% non-overlapping training/test split.

#### 5.2.1 Gaze Denoising

For this experiment the gaze input to the gaze transform module is corrupted by noise, mimicking imperfect knowledge of a person's gaze. We use our model to try to identify the correct gaze point taking scene saliency into consideration. We use a spatially-varying additive white Gaussian noise model, where the input to the network  $x_{noisy}$  is computed according to [27]:

$$\begin{aligned} x_{\text{noisy}}^{(i)}(t) = & x_{\text{true}}^{(i)}(t) + \eta^{(i)}, \ \eta^{(i)} \sim \mathcal{N}(0, (\sigma^{(i)})^2), \\ \sigma^{(i)} = & \max(\sigma_n^{(i)}, w * |x^{(i)} - x_0^{(i)}|) \end{aligned}$$
(9)

where  $\eta$  is the additive noise and the standard deviation  $\sigma^{(i)}$ increases as we get further from the center of the image along each coordinate  $x^{(i)} \in \{x^{(1)}, x^{(2)}\}$ . Our network denoises the gaze input by leveraging scene saliency information as encoded in the network. Fusing the noisy gaze location allows us to surpass the capability of a pure-saliency based model. The latter merely finds saliency image peaks that are close to the noisy gaze location and would be the straightforward way of incorporating saliency information. We also compare to an approach that relates gaze to the nearest object. In all cases we use a meanshift approach [13] to find nearby objects or peaks, with a standard deviation given by  $\sigma_n \sqrt{(H^2 + W^2)}$ , where H and W are the dimensions of the map. The results are summarized in Table 1 and demonstrate significant improvement with MAAD.

Scenario	Raw	OBJ	SAL	MAAD
$\sigma_n = 0.05$	15.7	62.0	25.7	18.1
$\sigma_n = 0.10$	28.1	55.6	25.4	19.8
$\sigma_n = 0.15$	40.9	53.3	24.9	21.0
$\sigma_n = 0.20$	51.7	53.1	24.9	22.1

Table 1: Mean absolute error (in pixels) of noisy gaze recovery based on object attention: meanshift into nearby objects (OBJ), meanshift according to pure saliency map (SAL), and meanshift correction based on the MAAD gaze map, for different noise levels. Our approach improved upon other alternatives over a wide variety of input noise levels, far beyond the noise level at train time ( $\sigma_n = 0.03$ ).

#### 5.2.2 Gaze Recalibration

We model imperfect calibration as an affine transformation. For this experiment, the DMS gaze input,  $x_{noisy}$ , to the gaze transform module is given by:

$$x_{\text{noisy}}(t) = T_{\text{correct}}\left(T_{\text{corrupt}}\left(x\right)\right) \tag{10}$$

where  $T_{\text{correct}}$ ,  $T_{\text{corrupt}}$  are both 2D affine transforms. We model the correcting transform,  $T_{\text{correct}}$ , as an MLP with one hidden layer. The corruption transform is created by adding element-wise zero-mean Gaussian noise with standard deviation  $\sigma_n^2$  to the transformation matrix and vector of an identity transform. We show the reduction in average error after calibration in Table 2. By leveraging saliency information, we are able to naturally compensate for the calibration errors using the model.

#### 5.3. Driver Attended Awareness

In this experiment, we measure the model's ability to infer attended awareness. We do so by measuring the model's

Scenario	Before	After
$\sigma_n = 0.1$	0.23	0.13
$\sigma_n = 0.2$	0.58	0.16
$\sigma_n = 0.3$	1.12	0.33
$\sigma_n = 0.5$	1.23	0.59

Table 2: The error was computed as the sum of the meansquared error of the elements of the  $T_{\text{correct}}$  and  $T_{\text{corrupt}}^{-1}$  before and after recovery from miscalibration. For each noise level the errors were computed by averaging over eight optimization runs. In Figure 4 (supp. material) we show the gaze map emitted by the model before and after recalibration with miscalibrated gaze input.

Noise level	MSE, FG	MSE, MAAD
$\sigma_n = 0.01$	0.357	0.138
$\sigma_n = 0.05$	0.359	0.135
$\sigma_n = 0.1$	0.425	0.138
$\sigma_n = 0.15$	0.464	0.140

Table 3: Mean squared error awareness estimates with spatio-temporal Gaussian with optic flow (FG) and MAAD, as a function of input gaze noise level. Our approach significantly outperforms the baseline and is robust to the noise present in the gaze input.

agreement with annotated third person estimate of awareness. We compare our approach to the following alternative of filtered gaze (FG) using a spatio-temporal Gaussian filter. We convolve each of the past gaze samples with a spatial Gaussian and utilize optic flow to temporally propagate the gaze information to the subsequent frames and aggregate them to form an awareness estimate. The optic flow mimics the subject's notion of object permanence under motion, and the spatial Gaussian account for subjective uncertainty accumulated over time as well as track limitations with optic flow. The results are given in Table 3.

#### 5.4. Ablation Experiments

We performed a series of leave-one out ablations to investigate the impact of the various cost terms on the attended awareness estimation task. Both  $\mathcal{L}_{ATT}$  and  $\mathcal{L}_{AA}$  are crucial for more accurate awareness estimation (Table 4).

# 6. Discussion and Conclusion

We have introduced a new model which uses imperfect gaze information and visual saliency to reason about perceived attended awareness within a single model. To train and validate our model, we generated a dataset which includes both high-accuracy gaze tracks as well as third person annotations for estimated attended awareness. MAAD

Ablation	Awareness Estimate
$\mathcal{L}_{ATT}$	0.262
$\mathcal{L}_{AA}$	0.199
$\mathcal{L}_T$	0.159
$\mathcal{L}_{CON}, \cdot$	0.164
Full model	0.138

Table 4: Attended awareness estimation (mean squared error) on the test set using different ablations of MAAD. The testing noise level was set to be  $\sigma_n = 0.1$ . More ablation results are presented in the supplementary material.



Figure 5: Time evolution of gaze (left) and awareness (right) under the reading-text modifier condition. Top: The subject gazes straight ahead and is aware of the car in front. Middle: The subject attends to text block and the gaze map shifts accordingly. The awareness map has two regions: one for the newly attended and the other for the previously attended location. Bottom: As the subject continues to read the text, the gaze map is localized in the bottom right, and the awareness map correctly captures how the subject is unaware of the scene in front.

can easily be extended to work with multi-image streams or scene graphs. Although our subjects viewed pre-collected video stimuli as opposed to being part of a true egocentric vision-action loop, one advantage is that we could acquire multiple observations of the same video, enabling the study of *distributions* rather than single scanpaths of attention for any given video. Our dataset can be compared in many ways to the related egocentric driving dataset from [40]. Extending our work and model to study subjects who are in control of real vehicles is a topic for future work.

# References

- Bogdan Alexe, Thomas Deselaers, and Vittorio Ferrari. Measuring the objectness of image windows. *IEEE Transactions* on Pattern Analysis and Machine Intelligence, 34(11):2189– 2202, 2012.
- [2] Sonia Baee, Erfan Pakdamanian, Inki Kim, Lu Feng, Vicente Ordonez, and Laura Barnes. Medirl: Predicting the visual attention of drivers via maximum entropy deep inverse reinforcement learning. arXiv preprint arXiv:1912.07773, 2019.
- [3] Loris Bazzani, Hugo Larochelle, and Lorenzo Torresani. Recurrent mixture density network for spatiotemporal visual attention. In *International Conference on Learning Representations*, 2017.
- [4] Thierry Bellet, Pierre Mayenobe, Jean-Charles Bornard, Dominique Gruyer, and Bernard Claverie. A computational model of the car driver interfaced with a simulation platform for future virtual human centred design applications: Cosmo-sivic. *Engineering Applications of Artificial Intelligence*, 25(7):1488–1504, 2012.
- [5] Paul J Boon, Jan Theeuwes, and Artem V Belopolsky. Updating visual–spatial working memory during object movement. *Vision Research*, 94:51–57, 2014.
- [6] Ali Borji, Ming-Ming Cheng, Huaizu Jiang, and Jia Li. Salient object detection: A benchmark. *IEEE Transactions* on *Image Processing*, 24(12):5706–5722, 2015.
- [7] Ali Borji and Laurent Itti. State-of-the-art in visual attention modeling. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(1):185–207, 2012.
- [8] Ali Borji, Dicky N Sihite, and Laurent Itti. Probabilistic learning of task-specific visual attention. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 470–477. IEEE, 2012.
- [9] Timothy F Brady and Joshua B Tenenbaum. A probabilistic model of visual working memory: Incorporating higher order regularities into working memory capacity estimates. *Psychological Review*, 120(1):85, 2013.
- [10] Neil DB Bruce, Christopher Catton, and Sasa Janjic. A deeper look at saliency: Feature contrast, semantics, and beyond. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 516–524, 2016.
- [11] Dario Cazzato, Marco Leo, Cosimo Distante, and Holger Voos. When i look into your eyes: A survey on computer vision contributions for human gaze estimation and tracking. *Sensors*, 20(13):3739, 2020.
- [12] Marvin M Chun. Visual working memory as visual attention sustained internally over time. *Neuropsychologia*, 49(6):1407–1409, 2011.
- [13] Dorin Comaniciu and Peter Meer. Mean shift: A robust approach toward feature space analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24(5):603–619, 2002.
- [14] Marcella Cornia, Lorenzo Baraldi, Giuseppe Serra, and Rita Cucchiara. A deep multi-level network for saliency prediction. In *International Conference on Pattern Recognition*, pages 3488–3493. IEEE, 2016.
- [15] Marcella Cornia, Lorenzo Baraldi, Giuseppe Serra, and Rita Cucchiara. Predicting human eye fixations via an lstm-based

saliency attentive model. *IEEE Transactions on Image Processing*, 27(10):5142–5154, 2018.

- [16] David Crundall, Peter Chapman, Steven Trawley, Lyn Collins, Editha Van Loon, Ben Andrews, and Geoffrey Underwood. Some hazards are more attractive than others: Drivers of varying experience respond differently to different types of hazard. *Accident Analysis & Prevention*, 45:600– 609, 2012.
- [17] Tao Deng, Hongmei Yan, Long Qin, Thuyen Ngo, and BS Manjunath. How do drivers allocate their potential attention? driving fixation prediction via convolutional neural networks. *IEEE Transactions on Intelligent Transportation Systems*, 21(5):2146–2154, 2019.
- [18] A. Doshi and M. M. Trivedi. On the roles of eye gaze and head dynamics in predicting driver's intent to change lanes. *IEEE Transactions on Intelligent Transportation Systems*, 10(3):453–462, 2009.
- [19] Mica R Endsley. Toward a theory of situation awareness in dynamic systems. *Human Factors*, 37(1):32–64, 1995.
- [20] Johan Engström, Emma Johansson, and Joakim Östlund. Effects of visual and cognitive load in real and simulated motorway driving. *Transportation Research Part F: Traffic Psychology and Behaviour*, 8(2):97–120, 2005.
- [21] Jianwu Fang, Dingxin Yan, Jiahuan Qiao, and Jianru Xue. Dada: A large-scale benchmark and model for driver attention prediction in accidental scenarios. arXiv preprint arXiv:1912.12148, 2019.
- [22] Xun Huang, Chengyao Shen, Xavier Boix, and Qi Zhao. Salicon: Reducing the semantic gap in saliency prediction by adapting deep neural networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 262– 270, 2015.
- [23] Laurent Itti, Christof Koch, and Ernst Niebur. A model of saliency-based visual attention for rapid scene analysis. *IEEE Transactions on Pattern Analysis and Machine Intelli*gence, 20(11):1254–1259, 1998.
- [24] Tilke Judd, Krista Ehinger, Frédo Durand, and Antonio Torralba. Learning to predict where humans look. In 2009 IEEE 12th International Conference on Computer Vision,, pages 2106–2113. IEEE, 2009.
- [25] Daniel Kaiser, Timo Stein, and Marius V Peelen. Real-world spatial regularities affect visual working memory for objects. *Psychonomic Bulletin & Review*, 22(6):1784–1790, 2015.
- [26] Anuradha Kar and Peter Corcoran. Performance evaluation strategies for eye gaze estimation systems with quantitative metrics and visualizations. *Sensors*, 18(9):3151, 2018.
- [27] Petr Kellnhofer, Adria Recasens, Simon Stent, Wojciech Matusik, and Antonio Torralba. Gaze360: Physically unconstrained gaze estimation in the wild. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6912–6921, 2019.
- [28] Saara Khalid, Jason C Deska, and Kurt Hugenberg. The eyes are the windows to the mind: Direct eye gaze triggers the ascription of others minds. *Personality and Social Psychology Bulletin*, 42(12):1666–1677, 2016.
- [29] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*, 2015.

- [30] Katja Kircher and Christer Ahlström. The driver distraction detection algorithm AttenD. In Trent W. Victor Michael A. Regan, John D. Lee, editor, *Driver Distraction and Inattention*, pages 327–348. Ashgate, 2013.
- [31] Alexander Kirillov, Kaiming He, Ross Girshick, Carsten Rother, and Piotr Dollár. Panoptic segmentation. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 9404–9413, 2019.
- [32] Matthias Kümmerer, Zoya Bylinskii, Tilke Judd, Ali Borji, Laurent Itti, Frédo Durand, Aude Oliva, and Antonio Torralba. MIT/Tübingen Saliency Benchmark. https://saliency.tuebingen.ai/.
- [33] Congyan Lang, Tam V Nguyen, Harish Katti, Karthik Yadati, Mohan Kankanhalli, and Shuicheng Yan. Depth matters: Influence of depth cues on visual saliency. In *European Conference on Computer Vision*, pages 101–115. Springer, 2012.
- [34] Yin Li, Zhefan Ye, and James M Rehg. Delving into egocentric actions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 287–295, 2015.
- [35] Tal Makovski and Yuhong V Jiang. The role of visual working memory in attentive tracking of unique objects. *Journal* of Experimental Psychology: Human Perception and Performance, 35(6):1687, 2009.
- [36] Sébastien Marcel and Yann Rodriguez. Torchvision the machine-vision package of torch. In *Proceedings of the* 18th ACM International Conference on Multimedia, MM '10, page 14851488, New York, NY, USA, 2010. Association for Computing Machinery.
- [37] Stefan Mathe and Cristian Sminchisescu. Actions in the eye: Dynamic gaze datasets and learnt saliency models for visual recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 37(7):1408–1424, 2014.
- [38] Jason McCarley, Christopher Wickens, Juliana Goh, and William Horrey. A computational model of attention/situation awareness. *Proceedings of the Human Factors* and Ergonomics Society Annual Meeting, 46, 09 2002.
- [39] Kyle Min and Jason J Corso. Tased-net: Temporallyaggregating spatial encoder-decoder network for video saliency detection. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2394–2403, 2019.
- [40] Andrea Palazzi, Davide Abati, Francesco Solera, Rita Cucchiara, et al. Predicting the driver's focus of attention: the dr (eye) ve project. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(7):1720–1733, 2018.
- [41] Junting Pan, Elisa Sayrol, Xavier Giro-i Nieto, Kevin McGuinness, and Noel E O'Connor. Shallow and deep convolutional networks for saliency prediction. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 598–606, 2016.
- [42] Junting Pan, Elisa Sayrol, Xavier Giro-i Nieto, Cristian Canton Ferrer, Jordi Torres, Kevin McGuinness, and Noel E OConnor. Salgan: Visual saliency prediction with adversarial networks. In CVPR Scene Understanding Workshop (SUNw), 2017.

- [43] Sylvain Paris and Frédo Durand. A fast approximation of the bilateral filter using a signal processing approach. *International Journal of Computer Vision*, 81(1):24–52, 2009.
- [44] Pietro Perona and Jitendra Malik. Scale-space and edge detection using anisotropic diffusion. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(7):629–639, 1990.
- [45] Shafin Rahman, Sejuti Rahman, Omar Shahid, Md Tahmeed Abdullah, and Jubair Ahmed Sourov. Classifying eyetracking data using saliency maps. *arXiv*, 2020.
- [46] Timothy J Ricker and Nelson Cowan. Loss of visual working memory within seconds: The combined use of refreshable and non-refreshable features. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 36(6):1355, 2010.
- [47] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. Unet: Convolutional networks for biomedical image segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 234–241. Springer, 2015.
- [48] Mehdi Sajjadi, Mehran Javanmardi, and Tolga Tasdizen. Regularization with stochastic transformations and perturbations for deep semi-supervised learning. In *NeurIPS*, 2016.
- [49] Dario D Salvucci. An integrated model of eye movements and visual encoding. *Cognitive Systems Research*, 1(4):201– 220, 2001.
- [50] JW Senders, AB Kristofferson, WH Levison, CW Dietrich, and JL Ward. The attentional demand of automobile driving. *Highway Research Record*, 1526(195), 1967.
- [51] Yusuke Sugano, Yasuyuki Matsushita, and Yoichi Sato. Calibration-free gaze sensing using saliency maps. In Proceedings of the Twenty-Third IEEE Conference on Computer Vision and Pattern Recognition, pages 2667–2674, 2010.
- [52] Benjamin W Tatler. Characterising the visual buffer: Realworld evidence for overwriting early in each fixation. *Perception*, 30(8):993–1006, 2001.
- [53] Benjamin W Tatler, Nicholas J Wade, Hoi Kwan, John M Findlay, and Boris M Velichkovsky. Yarbus, eye movements, and vision. *i-Perception*, 1(1):7–27, 2010.
- [54] Zachary Teed and Jia Deng. Raft: Recurrent allpairs field transforms for optical flow. arXiv preprint arXiv:2003.12039, 2020.
- [55] Antigoni Tsiami, Petros Koutras, and Petros Maragos. Stavis: Spatio-temporal audiovisual saliency network. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- [56] Wenguan Wang, Jianbing Shen, and Haibin Ling. A deep network solution for attention and aesthetics aware photo cropping. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(7):1531–1544, 2018.
- [57] Wenguan Wang, Jianbing Shen, and Ling Shao. Consistent video saliency using local gradient flow optimization and global refinement. *IEEE Transactions on Image Processing*, 24(11):4185–4196, 2015.
- [58] Benjamin Wolfe, Jonathan Dobres, Ruth Rosenholtz, and Bryan Reimer. More than the useful field: Considering peripheral vision in driving. *Applied Ergonomics*, 65:316–325, 2017.

- [59] Bertram Wortelen, Martin Baumann, and Andreas Lüdtke. Dynamic simulation and prediction of drivers attention distribution. *Transportation Research part F: Traffic Psychol*ogy and Behaviour, 21:278–294, 2013.
- [60] 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 *International Conference* on *Intelligent Robots and Systems*, 2019.
- [61] Ye Xia, Danqing Zhang, Jinkyu Kim, Ken Nakayama, Karl Zipser, and David Whitney. Predicting driver attention in critical situations. pages 658–674. Springer, 2018.
- [62] Saining Xie, Chen Sun, Jonathan Huang, Zhuowen Tu, and Kevin Murphy. Rethinking spatiotemporal feature learning: Speed-accuracy trade-offs in video classification. In *Proceedings of the European Conference on Computer Vision* (ECCV), pages 305–321, 2018.
- [63] Alfred L Yarbus. Eye Movements and Vision. Plenum Press, 1967.