



# Predicting Driver Attention in Critical Situations

Ye Xia<sup>(✉)</sup>, Danqing Zhang, Jinkyu Kim, Ken Nakayama, Karl Zipser,  
and David Whitney

University of California, Berkeley, CA 94720, USA  
yexia@berkeley.edu

**Abstract.** Robust driver attention prediction for critical situations is a challenging computer vision problem, yet essential for autonomous driving. Because critical driving moments are so rare, collecting enough data for these situations is difficult with the conventional in-car data collection protocol—tracking eye movements during driving. Here, we first propose a new in-lab driver attention collection protocol and introduce a new driver attention dataset, Berkeley DeepDrive Attention (BDD-A) dataset, which is built upon braking event videos selected from a large-scale, crowd-sourced driving video dataset. We further propose Human Weighted Sampling (HWS) method, which uses human gaze behavior to identify crucial frames of a driving dataset and weights them heavily during model training. With our dataset and HWS, we built a driver attention prediction model that outperforms the state-of-the-art and demonstrates sophisticated behaviors, like attending to crossing pedestrians but not giving false alarms to pedestrians safely walking on the sidewalk. Its prediction results are nearly indistinguishable from ground-truth to humans. Although only being trained with our in-lab attention data, the model also predicts in-car driver attention data of routine driving with state-of-the-art accuracy. This result not only demonstrates the performance of our model but also proves the validity and usefulness of our dataset and data collection protocol.

**Keywords:** Driver attention prediction · BDD-A dataset · Berkeley DeepDrive

## 1 Introduction

Human visual attention enables drivers to quickly identify and locate potential risks or important visual cues across the visual field, such as a darting-out pedestrian, an incursion of a nearby cyclist or a changing traffic light. Drivers' gaze behavior has been studied as a proxy for their attention. Recently, a large driver

---

**Electronic supplementary material** The online version of this chapter ([https://doi.org/10.1007/978-3-030-20873-8\\_42](https://doi.org/10.1007/978-3-030-20873-8_42)) contains supplementary material, which is available to authorized users.

attention dataset of routine driving [1] has been introduced and neural networks [21, 25] have been trained end-to-end to estimate driver attention, mostly in lane-following and car-following situations. Nonetheless, datasets and prediction models for driver attention in rare and critical situations are still needed.

However, it is nearly impossible to collect enough driver attention data for crucial events with the conventional in-car data collection protocol, *i.e.*, collecting eye movements from drivers during driving. This is because the vast majority of routine driving situations consist of simple lane-following and car-following. In addition, collecting driver attention in-car has two other major drawbacks. (i) Single focus: at each moment the eye-tracker can only record one location that the driver is looking at, while the driver may be attending to multiple important objects in the scene with their covert attention, *i.e.*, the ability to fixate one's eyes on one object while attending to another object [6]. (ii) False positive gazes: human drivers also show eye movements to driving-irrelevant regions, such as sky, trees, and buildings [21]. It is challenging to separate these false positives from gazes that are dedicated to driving.

An alternative that could potentially address these concerns is showing selected driving videos to drivers in the lab and collecting their eye movements with repeated measurements while they perform a proper simulated driving task. Although this third-person driver attention collected in the lab is inevitably different from the first-person driver attention in the car, it can still potentially reveal the regions a driver should look at in that particular driving situation from a third-person perspective. These data are greatly valuable for identifying risks and driving-relevant visual cues from driving scenes. To date, a proper data collection protocol of this kind is still missing and needs to be formally introduced and tested.

Another challenge for driver attention prediction, as well as for other driving-related machine learning problems, is that the actual cost of making a particular prediction error is unknown. Attentional lapses while driving on an empty road does not cost the same as attentional lapses when a pedestrian darts out. Since current machine learning algorithms commonly rely on minimizing average prediction error, the critical moments, where the cost of making an error is high, need to be properly identified and weighted.

Here, our paper offers the following novel contributions. First, in order to overcome the drawbacks of the conventional in-car driver attention collection protocol, we introduce a new protocol that uses crowd-sourced driving videos containing interesting events and makes multi-focus driver attention maps by averaging gazes collected from multiple human observers in lab with great accuracy (Fig. 1). We will refer to this protocol as the in-lab driver attention collection protocol. We show that data collected with our protocol reliably reveal where a experienced driver should look and can serve as a substitute for data collected with the in-car protocol. We use our protocol to collect a large driver attention dataset of braking events, which is, to the best of our knowledge, the richest to-date in terms of the number of interactions with other road agents. We call this dataset Berkeley DeepDrive Attention (BDD-A) dataset and will make it



**Fig. 1.** An example of input raw images (*left*), ground-truth human attention maps collected by us (*middle*), and the attention maps predicted by our model (*right*). The driver had to sharply stop at the green light to avoid hitting two pedestrians running the red light. The collected human attention map accurately shows the multiple regions that simultaneously demand the driver’s attention. Our model correctly attends to the crossing pedestrians and does not give false alarms to other irrelevant pedestrians (Color figure online)

publicly available. Second, we introduce Human Weighted Sampling (HWS), which uses human driver eye movements to identify which frames in the dataset are more crucial driving moments and weights the frames according to their importance levels during model training. We show that HWS improve model performance on both the entire testing set and the subset of crucial frames. Third, we propose a new driver attention prediction model trained on our dataset with HWS. The model shows sophisticated behaviors such as picking out pedestrians suddenly crossing the road without being distracted by the pedestrians safely walking in the same direction as the car (Fig. 1). The model prediction is nearly indistinguishable from ground-truth based on human judges, and it also matches the state-of-the-art performance level when tested on an existing in-car driver attention dataset collected during driving.

## 2 Related Works

**Image/Video Saliency Prediction.** A large variety of the previous saliency studies explored different bottom-up feature-based models [3, 4, 9, 20, 28, 32] combining low-level features like contrast, rarity, symmetry, color, intensity and orientation, or topological structure from a scene [12, 29, 32]. Recent advances in deep learning have achieved a considerable improvement for both image saliency prediction [13, 15–17] and video saliency prediction [2, 8, 18]. These models have achieved start-of-the-art performance on visual saliency benchmarks collected mainly when human subjects were doing a free-viewing task, but models that are specifically trained for predicting the attention of drivers are still needed.

**Driver Attention Datasets.** DR(eye)VE [1] is the largest and richest existing driver attention dataset. It contains 6 h of driving data, but the data was collected from only 74 rides, which limits the diversity of the dataset. In addition,

the dataset was collected in-car and has the drawbacks we introduced earlier, including missing covert attention, false positive gaze, and limited diversity. The driver’s eye movements were aggregated over a small temporal window to generate an attention map for a frame, so that multiple important regions of one scene might be annotated. But there was a trade-off between aggregation window length and gaze location accuracy, since the same object may appear in different locations in different frames. Reference [10] is another large driver attention dataset, but only six coarse gaze regions were annotated and the exterior scene was not recorded. References [24] and [27] contain accurate driver attention maps made by averaging eye movements collected from human observers in-lab with simulated driving tasks. But the stimuli were static driving scene images and the sizes of their datasets are small (40 frames and 120 frames, respectively).

**Driver Attention Prediction.** Self-driving vehicle control has made notable progress in the last several years. One of major approaches is a mediated perception-based approach – a controller depends on recognizing human-designated features, such as lane markings, pedestrians, or vehicles. Human driver’s attention provides important visual cues for driving, and thus efforts to mimic human driver’s attention have increasingly been introduced. Recently, several deep neural models have been utilized to predict where human drivers should pay attention [21, 25]. Most of existing models were trained and tested on the DR(eye)VE dataset [1]. While this dataset is an important contribution, it contains sparse driving activities and limited interactions with other road users. Thus it is restricted in its ability to capture diverse human attention behaviors. Models trained with this dataset tend to become vanishing point detectors, which is undesirable for modeling human attention in urban driving environment, where drivers encounter traffic lights, pedestrians, and a variety of other potential cues and obstacles. In this paper, we provide our human attention dataset as a contribution collected from a publicly available large-scale crowd-sourced driving video dataset [30], which contains diverse driving activities and environments, including lane following, turning, switching lanes, and braking in cluttered scenes.

### 3 Berkeley DeepDrive Attention (BDD-A) Dataset

**Dataset Statistics.** The statistics of our dataset are summarized and compared with the largest existing dataset (DR(eye)VE) [1] in Table 1. Our dataset was collected using videos selected from a publicly available, large-scale, crowd-sourced driving video dataset, BDD100k [30, 31]. BDD100K contains human-demonstrated dashboard videos and time-stamped sensor measurements collected during urban driving in various weather and lighting conditions. To efficiently collect attention data for critical driving situations, we specifically selected video clips that both included braking events and took place in busy areas (see supplementary materials for technical details). We then trimmed videos to include 6.5s prior to and 3.5s after each braking event. It turned

out that other driving actions, *e.g.*, turning, lane switching and accelerating, were also included. 1,232 videos (=3.5 h) in total were collected following these procedures. Some example images from our dataset are shown in Fig. 6. Our selected videos contain a large number of different road users. We detected the objects in our videos using YOLO [22]. On average, each video frame contained 4.4 cars and 0.3 pedestrians, multiple times more than the DR(eye)VE dataset (Table 1).

**Table 1.** Comparison between driver attention datasets

Dataset	# Rides	Durations (hours)	# Drivers	# Gaze providers	# Cars (per frame)	# Pedestrians (per frame)	# Braking events
DR(eye)VE [1]	74	6	8	8	1.0	0.04	464
BDD-A	1,232	3.5	1,232	45	4.4	0.25	1,427

**Data Collection Procedure.** For our eye-tracking experiment, we recruited 45 participants who each had more than one year of driving experience. The participants watched the selected driving videos in the lab while performing a driving instructor task: participants were asked to imagine that they were driving instructors sitting in the copilot seat and needed to press the space key whenever they felt it necessary to correct or warn the student driver of potential dangers. Their eye movements during the task were recorded at 1000 Hz with an EyeLink 1000 desktop-mounted infrared eye tracker, used in conjunction with the EyeLink Toolbox scripts [7] for MATLAB. Each participant completed the task for 200 driving videos. Each driving video was viewed by at least 4 participants. The gaze patterns made by these independent participants were aggregated and smoothed to make an attention map for each frame of the stimulus video (see Fig. 6 and supplementary materials for technical details).

Psychological studies [11, 19] have shown that when humans look through multiple visual cues that simultaneously demand attention, the order in which humans look at those cues is highly subjective. Therefore, by aggregating gazes of independent observers, we could record multiple important visual cues in one frame. In addition, it has been shown that human drivers look at buildings, trees, flowerbeds, and other unimportant objects non-negligibly frequently [1]. Presumably, these eye movements should be regarded as noise for driving-related machine learning purposes. By averaging the eye movements of independent observers, we were able to effectively wash out those sources of noise (see Fig. 2B).

**Comparison with In-car Attention Data.** We collected in-lab driver attention data using videos from the DR(eye)VE dataset. This allowed us to compare in-lab and in-car attention maps of each video. The DR(eye)VE videos we used were 200 randomly selected 10-second video clips, half of them containing braking events and half without braking events.

We tested how well in-car and in-lab attention maps highlighted driving-relevant objects. We used YOLO [22] to detect the objects in the videos of our



**Fig. 2.** Comparison between in-car and in-lab driver attention maps. (A) Proportions of attended objects of different categories for in-car and in-lab driver attention maps. In-car attention maps tend to highlight significantly fewer driving-relevant objects than in-lab attention maps. (B) An example of in-car driver attention maps showing irrelevant regions. The in-lab attention map highlights the car in front and a car that suddenly backed up, while the in-car attention map highlights some regions of the building

dataset. We identified three object categories that are important for driving and that had sufficient instances in the videos (car, pedestrian and cyclist). We calculated the proportion of attended objects out of total detected instances for each category for both in-lab and in-car attention maps (see supplementary materials for technical details). The results showed that in-car attention maps highlighted significantly less driving-relevant objects than in-lab attention maps (see Fig. 2A).

The difference in the number of attended objects between the in-car and in-lab attention maps can be due to the fact that eye movements collected from a single driver do not completely indicate all the objects that demand attention in the particular driving situation. One individual’s eye movements are only an approximation of their attention [23], and humans can also track objects with covert attention without looking at them [6]. The difference in the number of attended objects may also reflect the difference between first-person driver attention and third-person driver attention. It may be that the human observers in our in-lab eye-tracking experiment also looked at objects that were not relevant for driving. We ran a human evaluation experiment to address this concern.

**Human Evaluation.** To verify that our in-lab driver attention maps highlight regions that should indeed demand drivers’ attention, we conducted an online study to let humans compare in-lab and in-car driver attention maps. In each trial of the online study, participants watched one driving video clip three times: the first time with no edit, and then two more times in random order with overlaid in-lab and in-car attention maps, respectively. The participant was then asked to choose which heatmap-coded video was more similar to where a good driver would look. In total, we collected 736 trials from 32 online participants. We found that our in-lab attention maps were more often preferred by the participants than the in-car attention maps (71% versus 29% of all trials, statistically significant as  $p = 1 \times 10^{-29}$ , see Table 2). Although this result cannot suggest that in-lab driver attention maps are superior to in-car attention maps in general, it

**Table 2.** Two human evaluation studies were conducted to compare in-lab human driver attention maps with in-car human driver attention maps and attention maps predicted by our HWS model, respectively. In-car human driver attention maps were preferred in significantly less trials than the in-lab human driver attention maps. The attention maps predicted by our HWS model were not preferred in as many trials as the in-lab human driver attention maps, but they achieved significantly higher preference rate than the in-car human driver attention maps

	# trials	Attention maps	Preference rate
Study 1	736	In-car human driver	29%
		In-lab human driver	71%
Study 2	462	HWS model predicted	41%
		In-lab human driver	59%

does show that the driver attention maps collected with our protocol represent where a good driver should look from a third-person perspective.

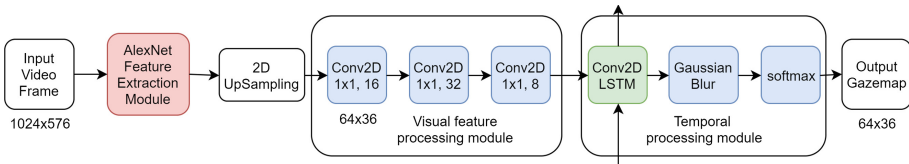
In addition, we will show in the Experiments section that in-lab attention data collected using our protocol can be used to train a model to effectively predict actual, in-car driver attention. This result proves that our dataset can also serve as a substitute for in-car driver attention data, especially in crucial situations where in-car data collection is not practical.

To summarize, compared with driver attention data collected in-car, our dataset has three clear advantages: multi-focus, little driving-irrelevant noise, and efficiently tailored to crucial driving situations.

## 4 Attention Prediction Model

### 4.1 Network Configuration

Our goal is to predict the driver attention map for a video frame given the current and previous video frames. Our model structure can be divided into a visual feature extraction module, a visual feature processing module, and a temporal processing module (Fig. 3).

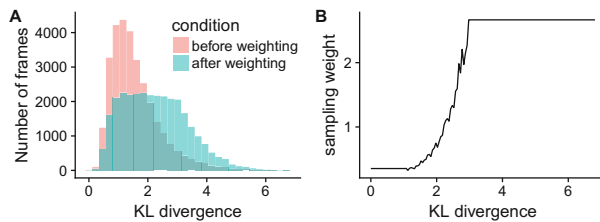


**Fig. 3.** An overview of our proposed model that predicts human driver’s attention from input video frame. We use AlexNet pre-trained on ImageNet as a visual feature extractor. We also use three fully convolutional layers (Conv2D) followed by a convolutional LSTM network (Conv2D LSTM)

The visual feature extraction module is a pre-trained dilated fully convolutional neural network, and its weights are fixed during training. We used ImageNet pre-trained AlexNet [14] as our visual feature extraction module. We chose to use the features from the conv5 layer. In our experiment, the size of the input was set to  $1024 \times 576$  pixels, and the feature map by AlexNet was upsampled to  $64 \times 36$  pixels and then fed to the following visual feature processing module.

The visual feature processing module is a fully convolutional neural network. It consists of three convolutional layers with  $1 \times 1$  kernels and a dropout layer after each convolutional layer. It further processes the visual features from the previous extraction module and reduces the dimensionality of the visual features from 256 to 8. In our experiments, we observed that without the dropout layers, the model easily got stuck in a suboptimal solution which simply predicted a central bias map, i.e. an attention map concentrated in a small area around the center of the frame.

The temporal processor is a convolutional LSTM network with a kernel size of  $3 \times 3$  followed by a Gaussian smooth layer ( $\sigma$  set to 1.5) and a softmax layer. It receives the visual features of successive video frames in sequence from the visual feature processing module and predicts an attention map for every new time step. Dropout is used for both the linear transformation of the inputs and the linear transformation of the recurrent states. We had also experimented with using an LSTM network for this module and observed that the model tended to incorrectly attend to only the central region of the video frames. The final output of this model is a probability distribution over  $64 \times 36$  grids predicting how likely each region of the video frame is to be looked at by human drivers. Cross-entropy is chosen as the loss function to match the predicted probability distribution to the ground-truth.



**Fig. 4.** Human Weighted Sampling: (A) For each video frame, we measure the KL divergence between the collected driver attention maps and the mean attention map for that entire video clip ( $\approx 10$  s). We use this computed KL divergence as a weight value to sample image frames during training phase, *i.e.*, training a model more often with uncommon attention maps. Histograms show that more uncommon attention maps were selected for training the model, *e.g.*, seeing pedestrians or traffic lights is weighted more than just seeing the vanishing point of roads. (B) Normalized sampling weights as a function of KL divergence values. A normalized sampling weight value of 1 indicates that the video frame is sampled once on average during a single epoch



## 4.2 Human Weighted Sampling (HWS)

Human driver attention datasets, as well as many other driving related datasets, share a common bias: the vast majority of the datasets consist of simple driving situations such as lane-following or car-following. The remaining small proportion of driving situations, such as pedestrians darting out, traffic lights changing, etc., are usually more crucial, in the sense that making errors in these moments would lead to greater cost. Therefore, ignoring this bias and simply using mean prediction error to train and test models can be misleading. In order to tackle this problem, we developed a new method that uses human gaze data to determine the importance of different frames of a driving dataset and samples the frames with higher importance more frequently during training.

In simple driving situations human drivers only need to look at the center of the road or the car in front, which can be shown by averaging the attention maps of all the frames of one driving video. When the attention map of one frame deviates greatly from the average default attention map, it is usually an important driving situation where the driver has to make eye movements to important visual cues. Therefore, the more an attention map varies from the average attention map of the video, the more important the corresponding training frame is. We used the KL divergence to measure the difference between the attention map of a particular frame and the average attention map of the video. The KL divergence determined the sampling weight of this video frame during training.

The histogram of the KL divergence of all the training video frames of our dataset is shown in Fig. 4. As we expected, the histogram was strongly skewed to the left side. Our goal was to boost up the proportion of the frames of high KL divergence values by weighted sampling. The sampling weight was determined as a function of KL divergence ( $D_{KL}$ ) illustrated in Fig. 4B. The middle part of this function ( $D_{KL} \in [1,3]$ ) was set to be proportional to the inverse of the histogram so that after weighted sampling the histogram of KL divergence would become flat on this range. The left part of the function ( $D_{KL} < 1$ ) was set to a low constant value so that those frames would be sampled occasionally but not completely excluded. The right part of the function was set to a saturated constant value instead of monotonically increasing values in order to avoid overfitting the model to this small proportion of data. Besides, the attention maps collected in the beginning and the end of each video clip can deviate from the average default attention map merely because the participants were distracted by the breaks between video clips. We therefore restricted the sampling weights of the first second and the last 0.5s of each video to be less or equal to once per epoch. The histogram of KL divergence after weighted sampling is shown in Fig. 4A. In our experiment, we needed to sample the training frames in continuous sequences of 6 frames. For a particular sequence, its sampling weight was equal to the sum of the sampling weights of its member frames. These sequences were sampled at probabilities proportional to the sequence sampling weights.

## 5 Results and Discussion

Here, we first provide our training and evaluation details, then we summarize the quantitative and qualitative performance comparison with existing gaze prediction models and variants of our model. To test how natural and reasonable our model prediction look to humans, we conduct a human evaluation study and summarize the results. We further test whether our model trained on in-lab driver attention data can also predict driver attention maps collected in-car.

### 5.1 Training and Evaluation Details

We made two variants of our model. One was trained with a regular regime, *i.e.*, equal sampling during training, and the other was trained with Human Weighted Sampling (HWS). Except for the sampling method during training, our default model and HWS model shared the same following training settings. We used 926 videos from our BDD-A dataset as the training set and 306 videos as the testing set. We downsampled the videos to  $1024 \times 576$  pixels and 3 Hz. After this preprocessing, we had about 30k frames in our training set and 10k frames in our testing set. We used cross-entropy between predicted attention maps and human attention maps as the training loss, along with Adam optimizer (learning rate = 0.001,  $\beta_1 = 0.9$ ,  $\beta_2 = 0.999$ ,  $\epsilon = 1 \times 10^{-8}$ ). Each training batch contained 10 sequences and each sequence had 6 frames. The training was done for 10,000 iterations. The two models showed stabilized testing errors by iteration 10,000.

To our knowledge, [21] and [25] are the two deep neural models that use dash camera videos alone to predict human driver’s gaze. They demonstrated similar results and were shown to surpass other deep learning models or traditional models that predict human gaze in non-driving-specific contexts. We chose to replicate [21] to compare with our work because their prediction code is public. The model designed by [21] was trained on the DR(eye)VE dataset [1]. We will refer to [21]’s model as DR(eye)VE model in the following. The training code of [21] is not available. We implemented code to fine-tune their model on our dataset, but the fine-tuning did not converge to any reasonable solution, potentially due to some training parameter choices that were not reported. We then tested their pre-trained model directly on our testing dataset without any training on our training dataset. Since the goal of the comparison was to test the effectiveness of the combination of model structure, training data and training paradigm as a whole, we think it is reasonable to test how well DR(eye)VE model performs on our dataset without further training. For further comparison, we fine-tuned a publicly available state-of-the-art image gaze prediction model, SALICON [13] on our dataset. We used the open source implementation [26]. We also tested our models against a baseline model that always predicts the averaged human attention map of training videos.

Kullback-Leibler divergence (KL divergence,  $D_{KL}$ ), Pearson’s Correlation Coefficient (CC), Normalized Scanpath Saliency (NSS) and Area under ROC Curve (AUC) are four commonly used metrics for attention map prediction [5, 21, 25]. We calculated the mean prediction errors in these four metrics on the

**Table 3.** Performance comparison of human attention prediction. Mean and 95% bootstrapped confidence interval are reported

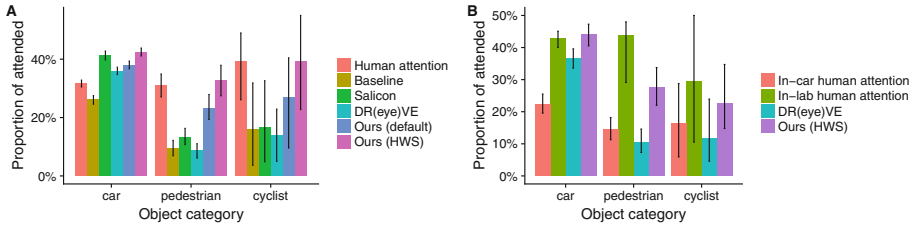
	Entire testing set				Testing subset where $D_{KL}(GT, \text{Mean}) > 2$			
	KL divergence		Correlation coefficient		KL divergence		Correlation coefficient	
	Mean	95% CI	Mean	95% CI	Mean	95% CI	Mean	95% CI
Baseline	1.50	(1.45, 1.54)	0.46	(0.44, 0.48)	1.87	(1.80, 1.94)	0.36	(0.34, 0.37)
SALICON [13]	1.41	(1.39, 1.44)	0.53	(0.51, 0.54)	1.76	(1.72, 1.80)	0.39	(0.37, 0.41)
DR(eye)VE [21]	1.95	(1.87, 2.04)	0.50	(0.48, 0.52)	2.63	(2.51, 2.77)	0.35	(0.33, 0.37)
Ours (default)	1.24	(1.21, 1.28)	0.58	(0.56, 0.59)	1.71	(1.65, 1.79)	0.41	(0.40, 0.43)
Ours (HWS)	<b>1.24</b>	<b>(1.21, 1.27)</b>	<b>0.59</b>	<b>(0.57, 0.60)</b>	<b>1.67</b>	<b>(1.61, 1.73)</b>	<b>0.44</b>	<b>(0.42, 0.45)</b>

testing set to compare the different models. In order to test how well the models perform at important moments where drivers need to watch out, we further calculated the mean prediction errors on the subset of testing frames where the attention maps deviate significantly from the average attention maps of the corresponding videos (defined as KL divergence greater than 2.0). We will refer to these frames as non-trivial frames. Our models output predicted attention maps in the size of  $64 \times 36$  pixels, but the DR(eye)VE model and the SALICON outputs in bigger sizes. For a fair comparison, we scaled the DR(eye)VE model and the SALICON model’s predicted attention maps into  $64 \times 36$  pixels before calculating the prediction errors.

Another important evaluation criterion of driver attention models is how successfully they can attend to the objects that demand human driver’s attention, e.g. the cars in front, the pedestrians that may enter the roadway, etc. Therefore, we applied the same attended object analysis described in the Berkeley DeepDrive Attention Dataset section. We used YOLO [22] to detect the objects in the videos of our dataset. We selected object categories that are important for driving and that have enough instances in both our dataset and the DR(eye)VE dataset for comparison (car, pedestrian and cyclist). We calculated the proportions of all the detected instances of those categories that were actually attended to by humans versus the models. The technical criterion of determining attended objects was the same as described in the Berkeley DeepDrive Attention Dataset section.

## 5.2 Evaluating Attention Predictor

**Quantitative Analysis of Attention Prediction.** The mean prediction errors of different models are summarized in Table 3 (measured in  $D_{KL}$  and CC) and Table S1 (measured in NSS and AUC) in supplementary materials. Both of our models significantly outperformed the DR(eye)VE model, the SALICON model and the baseline model in all metrics on both the entire testing set and the subset of non-trivial frames. Our model trained with HWS was essentially trained on a dataset whose distribution was altered from the distribution of the testing set. However, our HWS model showed better results than our default model even when being tested on the whole testing set. When being tested on



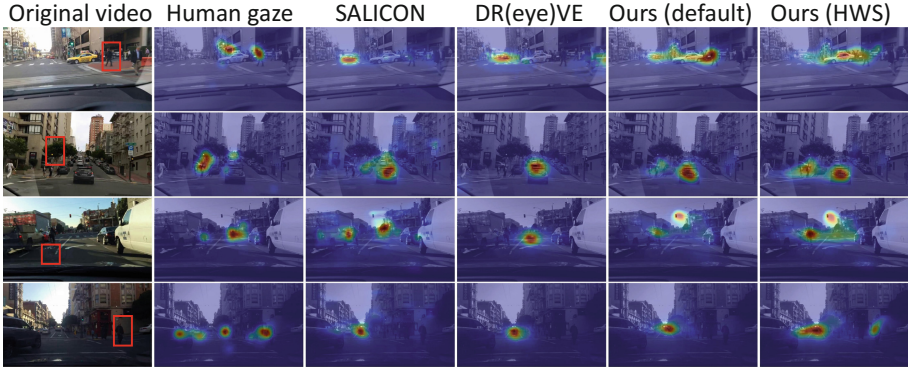
**Fig. 5.** Analysis of attended objects for human attention and different models tested on our dataset (A) and the DR(eye)VE dataset (B). Error bars show 95% bootstrapped confidence intervals

the subset of non-trivial frames, our HWS model outperformed our default model even more significantly. These results suggest that HWS has the power to overcoming the dataset bias and better leveraging the knowledge hidden in crucial driving moments.

The results of the attended object analysis are summarized in Fig. 5A. Cars turned out to be easy to identify for all models. This is consistent with the fact that a central bias of human attention is easy to learn and cars are very likely to appear in the center of the road. However, for pedestrians and cyclists, the DR(eye)VE model, SALICON model and baseline model all missed a large proportion of them compared with human attention ground-truth. Both of our models performed significantly better than all the other competing models in the categories of pedestrians and cyclists, and our HWS model matched the human attention performances the best.

Importantly, our HWS model did not simply select objects according to their categories like an object detection algorithm. Considering the category that has the highest safety priority, pedestrian, our models selectively attended to the pedestrians that were also attended to by humans. Let us refer to the pedestrians that were actually attended to by humans as the important pedestrians and the rest of them as non-important pedestrians. Among all the pedestrians detected by the object detection algorithm, the proportion of important pedestrians was 33%. If our HWS model were simply detecting pedestrians at a certain level and could not distinguish between important pedestrians and non-important pedestrians, the proportion of important pedestrians among the pedestrians attended to by our model should also be 33%. However, the actual proportion of important pedestrians that our HWS model attended to was 48% with a bootstrapped 95% confidence interval of [42%, 55%]. Thus, our HWS model predicts which of the pedestrians are the ones most relevant to human drivers.

**Qualitative Analysis of Attention Prediction.** Some concrete examples are shown in Fig. 6 (see supplementary information for videos). These examples demonstrates some important driving scenarios: pedestrian crossing, cyclist getting very close to the vehicle and turning at a busy crossing. It can be seen from these examples that the SALICON model and the DR(eye)VE model mostly



**Fig. 6.** Examples of the videos in our dataset, ground-truth human attention maps and the prediction of different models. The red rectangles in the original video column highlight the pedestrians that pose a potential hazard. Row 1: the driver had the green light, but a pedestrian was about to cross the road while speaking on a phone without looking at the driver. Another pedestrian was present in the scene, but not relevant to the driving decision. Row 2: the driver had a yellow light and some pedestrians were about to enter the roadway. Another pedestrian was walking in the same direction as the car and therefore not relevant to the driving decision. Row 3: a cyclist was very close to the car. Row 4: the driver was making a right turn and needed to yield to the crossing pedestrian. Other pedestrians were also present in the scene but not relevant to the driving decision (Color figure online)

only predicted to look at the center of the road and ignored the crucial pedestrians or cyclists. In the examples of row 1, 2 and 3, both our default model and HWS model successfully attended to the important pedestrian/cyclist, and did not give false alarm for other pedestrians who were not important for the driving decision. In the challenging example shown in row 4, the driver was making a right turn and needed to yield to the crossing pedestrian. Only our HWS model successfully overcame the central bias and attended to the pedestrian appearing in a quite peripheral area in the video frame.

**Human Evaluation.** To further test how natural and reasonable our HWS model’s predicted attention maps look to humans, we conducted an online Turing Test. In each trial, a participant watched one driving video clip three times: the first time with no edit, and then two times in random order with the ground-truth human driver attention map and our HWS model’s predicted attention map overlaid on top, respectively. The participant was then asked to choose whether the first or the second attention map video was more similar to where a good driver would look.

Note that the experiment settings and instructions were the same as the online study described in the dataset section, except that one compares model prediction against the in-lab driver attention maps, and the other compares the in-car driver attention maps against the in-lab driver attention maps. Therefore,

the result of this Turing Test can be compared with the result of the previous online study. In total, we collected 462 trials from 20 participants. If our HWS model's predicted attention maps were perfect and indistinguishable from the ground-truth human driver attention maps, the participants would have had to make random choices, and therefore we would expect them to choose our model prediction in about 50% of the trials. If our HWS model's prediction was always wrong and unreasonable, we would expect a nearly zero chosen rate for our model prediction. Our results showed that in 41% of all trials the participants chose our HWS model's predicted attention maps as even better than the in-lab human attention maps (see Table 2). In the previous online study, the in-car attention maps of DR(eye)VE only achieved a chosen rate of 29%. This result suggests that our HWS model's predicted attention maps were even more similar to where a good driver should look than the human driver attention maps collected in-car (permutation test  $p = 4 \times 10^{-5}$ ).

### 5.3 Predicting In-car Driver Attention Data

To further demonstrate that our model has good generalizability and that our driver attention data collected in-lab is realistic, we conducted a challenging test: we trained our model using only our in-lab driver attention data, but tested it on the DR(eye)VE dataset, an in-car driver attention dataset. Note that the DR(eye)VE dataset covers freeway driving, which is not included in our dataset due to the small density of road user interactions on freeway. The high driving speed on freeway introduces strong motion blur which is not present in our dataset videos. Furthermore, drivers need to look further ahead in high speed situations, so the main focus of driver gaze pattern shifts up as the driving speed increases. In order to adapt our model to these changes, we selected 200 ten-second-long video clips from the training set of the DR(eye)VE dataset and collected in-lab driver attention maps for those video clips (already described in the Berkeley DeepDrive Attention Dataset section). We fine-tuned our HWS model with these video clips (30 min in total only) and the corresponding in-lab driver attention maps, and then tested the model on the testing set of the DR(eye)VE dataset (with in-car attention maps). The mean testing errors were calculated in  $D_{KL}$  and CC because the calculation of NSS and AUC requires the original fixation pixels instead of smoothed gaze maps and the original fixation pixels of the DR(EYE)VE dataset were not released. Our fine-tuned model showed a better mean value in KL Divergence and a worse mean value in CC than the DR(eye)VE model (see Table 4). But the 95% bootstrapped confidence intervals for the two models in both metrics overlapped with each other. So overall we concluded that our fine-tuned model matched the performance of the DR(eye)VE model. Note that the DR(eye)VE model was trained using the DR(eye)VE dataset and represents the state-of-the-art performance on this dataset.

We also calculated proportions of attended objects of important categories for our fine-tuned model and the DR(eye)VE model (Fig. 5B). Our fine-tuned model showed significantly higher proportions of attended objects in the car, pedestrian and cyclist categories and was more similar to the in-lab driver attention than

**Table 4.** Test results obtained on the DR(eye)VE dataset by the state-of-the-art model (DR(eye)VE) and our finetuned model. Mean and 95% bootstrapped confidence interval are reported

	KL divergence		Correlation coefficient	
	Mean	95% CI	Mean	95% CI
DR(eye)VE	1.76	(1.65, 1.87)	0.54	(0.51, 0.56)
Ours (finetuned)	1.72	(1.66, 1.81)	0.51	(0.48, 0.53)

the DR(eye)VE model. Note that we have shown in the Berkeley DeepDrive Attention Dataset section that humans rated the in-lab attention maps as more similar to where a good driver should look from a third-person perspective than the in-car attention maps.

## 6 Conclusions

In this paper, we introduce a new in-lab driver attention data collection protocol that overcomes drawbacks of in-car collection protocol. We contribute a human driver attention dataset which is to-date the richest and will be made public. We propose Human Weighted Sampling which can overcome common driving dataset bias and improve model performance in both the entire dataset and the subset of crucial moments. With our dataset and sampling method we contribute a novel human driver attention prediction model that can predict both in-lab and in-car driver attention data. The model demonstrates sophisticated behaviors and show prediction results that are nearly indistinguishable from ground-truth to humans.

## References

1. Alletto, S., Palazzi, A., Solera, F., Calderara, S., Cucchiara, R.: DR(eye)VE: a dataset for attention-based tasks with applications to autonomous and assisted driving. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops, pp. 54–60 (2016)
2. Bazzani, L., Larochelle, H., Torresani, L.: Recurrent mixture density network for spatiotemporal visual attention. arXiv preprint [arXiv:1603.08199](https://arxiv.org/abs/1603.08199) (2016)
3. Bruce, N., Tsotsos, J.: Saliency based on information maximization. In: Advances in Neural Information Processing Systems, pp. 155–162 (2006)
4. Bruce, N.D., Tsotsos, J.K.: Saliency, attention, and visual search: an information theoretic approach. *J. Vis.* **9**(3), 5–5 (2009)
5. Bylinskii, Z., Judd, T., Oliva, A., Torralba, A., Durand, F.: What do different evaluation metrics tell us about saliency models? *IEEE Trans. Pattern Anal. Mach. Intell.* **41**, 740–757 (2018)
6. Cavanagh, P., Alvarez, G.A.: Tracking multiple targets with multifocal attention. *Trends Cogn. Sci.* **9**(7), 349–354 (2005)



7. Cornelissen, F.W., Peters, E.M., Palmer, J.: The eyelink toolbox: eye tracking with matlab and the psychophysics toolbox. *Behav. Res. Methods Instrum. Comput.* **34**(4), 613–617 (2002)
8. Cornia, M., Baraldi, L., Serra, G., Cucchiara, R.: Predicting human eye fixations via an LSTM-based saliency attentive model. arXiv preprint [arXiv:1611.09571](https://arxiv.org/abs/1611.09571) (2016)
9. Erdem, E., Erdem, A.: Visual saliency estimation by nonlinearly integrating features using region covariances. *J. Vis.* **13**(4), 11–11 (2013)
10. Fridman, L., Langhans, P., Lee, J., Reimer, B.: Driver gaze region estimation without use of eye movement. *IEEE Intell. Syst.* **31**(3), 49–56 (2016)
11. Groner, R., Walder, F., Groner, M.: Looking at faces: local and global aspects of scanpaths. In: *Advances in Psychology*, vol. 22, pp. 523–533. Elsevier (1984)
12. Harel, J., Koch, C., Perona, P.: Graph-based visual saliency. In: *Advances in Neural Information Processing Systems*, pp. 545–552 (2007)
13. Huang, X., Shen, C., Boix, X., Zhao, Q.: SALICON: reducing the semantic gap in saliency prediction by adapting deep neural networks. In: *Proceedings of the IEEE International Conference on Computer Vision*, pp. 262–270 (2015)
14. Krizhevsky, A., Sutskever, I., Hinton, G.E.: ImageNet classification with deep convolutional neural networks. In: *Advances in Neural Information Processing Systems*, pp. 1097–1105 (2012)
15. Kümmerer, M., Theis, L., Bethge, M.: Deep Gaze I: boosting saliency prediction with feature maps trained on ImageNet. In: *International Conference on Learning Representations (ICLR 2015)* (2015)
16. Kümmerer, M., Wallis, T.S., Bethge, M.: DeepGaze II: reading fixations from deep features trained on object recognition. arXiv preprint [arXiv:1610.01563](https://arxiv.org/abs/1610.01563) (2016)
17. Liu, N., Han, J., Zhang, D., Wen, S., Liu, T.: Predicting eye fixations using convolutional neural networks. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 362–370 (2015)
18. Liu, Y., Zhang, S., Xu, M., He, X.: Predicting salient face in multiple-face videos. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 4420–4428 (2017)
19. Mannan, S., Ruddock, K., Wooding, D.: Fixation sequences made during visual examination of briefly presented 2D images. *Spat. Vis.* **11**(2), 157–178 (1997)
20. Murray, N., Vanrell, M., Otazu, X., Parraga, C.A.: Saliency estimation using a non-parametric low-level vision model. In: *2011 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 433–440. IEEE (2011)
21. Palazzi, A., Solera, F., Calderara, S., Alletto, S., Cucchiara, R.: Learning where to attend like a human driver. In: *2017 IEEE Intelligent Vehicles Symposium (IV)*, pp. 920–925. IEEE (2017)
22. Redmon, J., Farhadi, A.: YOLO9000: better, faster, stronger. In: *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 6517–6525. IEEE (2017)
23. Rizzolatti, G., Riggio, L., Dascola, I., Umiltá, C.: Reorienting attention across the horizontal and vertical meridians: evidence in favor of a premotor theory of attention. *Neuropsychologia* **25**(1), 31–40 (1987)
24. Simon, L., Tarel, J.P., Brémond, R.: Alerting the drivers about road signs with poor visual saliency. In: *2009 IEEE Intelligent Vehicles Symposium*, pp. 48–53. IEEE (2009)
25. Tawari, A., Kang, B.: A computational framework for driver’s visual attention using a fully convolutional architecture. In: *2017 IEEE Intelligent Vehicles Symposium (IV)*, pp. 887–894. IEEE (2017)



26. Thomas, C.L.: OpenSalicon: an open source implementation of the salicon saliency model. Technical report. TR-2016-02, University of Pittsburgh (2016)
27. Underwood, G., Humphrey, K., Van Loon, E.: Decisions about objects in real-world scenes are influenced by visual saliency before and during their inspection. *Vis. Res.* **51**(18), 2031–2038 (2011)
28. Valenti, R., Sebe, N., Gevers, T.: Image saliency by isocentric curvedness and color. In: 2009 IEEE 12th International Conference on Computer Vision, pp. 2185–2192. IEEE (2009)
29. Wei, Y., Wen, F., Zhu, W., Sun, J.: Geodesic saliency using background priors. In: Fitzgibbon, A., Lazebnik, S., Perona, P., Sato, Y., Schmid, C. (eds.) ECCV 2012. LNCS, vol. 7574, pp. 29–42. Springer, Heidelberg (2012). [https://doi.org/10.1007/978-3-642-33712-3\\_3](https://doi.org/10.1007/978-3-642-33712-3_3)
30. Xu, H., Gao, Y., Yu, F., Darrell, T.: End-to-end learning of driving models from large-scale video datasets. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (2017)
31. Yu, F., et al.: BDD100K: a diverse driving video database with scalable annotation tooling. arXiv preprint [arXiv:1805.04687](https://arxiv.org/abs/1805.04687) (2018)
32. Zhang, J., Sclaroff, S.: Saliency detection: a boolean map approach. In: 2013 IEEE International Conference on Computer Vision (ICCV), pp. 153–160. IEEE (2013)