

Learning Where to Attend Like a Human Driver

Andrea Palazzi*, Francesco Solera*, Simone Calderara, Stefano Alletto, Rita Cucchiara

Abstract—Despite the advent of autonomous cars, it’s likely - at least in the near future - that human attention will still maintain a central role as a guarantee in terms of legal responsibility during the driving task. In this paper we study the dynamics of the driver’s gaze and use it as a proxy to understand related attentional mechanisms. First, we build our analysis upon two questions: *where* and *what* the driver is looking at? Second, we model the driver’s gaze by training a coarse-to-fine convolutional network on short sequences extracted from the DR(eye)VE dataset. Experimental comparison against different baselines reveal that the driver’s gaze can indeed be learnt to some extent, despite i) being highly subjective and ii) having only one driver’s gaze available for each sequence due to the irreproducibility of the scene. Eventually, we advocate for a new assisted driving paradigm which suggests to the driver, with no intervention, where she should focus her attention.

I. INTRODUCTION

While autonomous driving is quickly reaching maturity, it’s not clear how far in time society will overlook the legal responsibility of the human driver [1]. Conversely, Advanced Driver Assistance Systems (ADAS) are human-centric and already established both in literature and in the market.

Some of the most ambitious examples of assisted driving are related to driver monitoring systems [2], [3], [4], [5], where the attentional behavior of the driver is parsed together with the road scene to predict potentially unsafe manoeuvres and act on the car in order to avoid them (either by signaling the driver or braking). However, all these approaches are limited by their ability to capture the true attentional and intentional behavior of the driver, which is still a complex and largely unsolved task today. Conversely, we advocate for a new assisted driving paradigm which suggests to the driver, with no hard intervention, where she should focus her attention. The problem is thus shifted from a personal level (*what the driver is looking at*) to a task-driven level (*what the driver should be looking at*). Following the notion that gaze is a primary cue to human visual attention, the contributions of this paper are twofold:

- we investigate the attentional dynamics during the driving task by employing data collected from eye tracking glasses and semantic cues, Fig. 1(c) and (d) respectively (Sec. III);
- upon these findings, we build a deep network architecture to model human attention while driving, Fig. 1(b), and evaluate the ability of the proposed approach to replicate what we observed on humans (Sec. IV-VI).

* Equal contribution

Authors are with the Department of Computer Engineering, University of Modena and Reggio Emilia, Italy `name.surname@unimore.it`

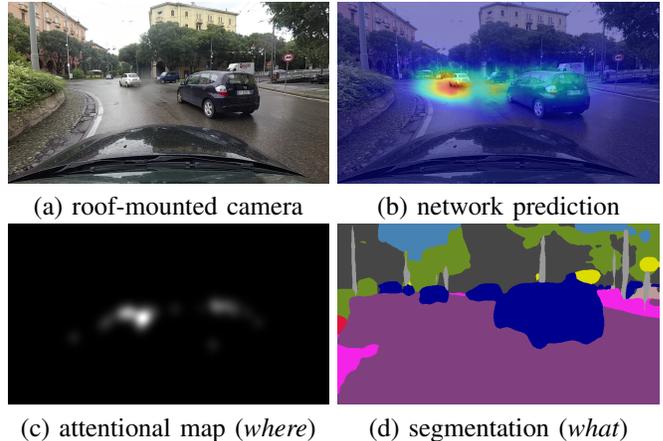


Fig. 1. Starting from raw frames and attentional maps provided with the DR(eye)VE [6] dataset, we investigate *where* the driver focuses (c) and *what* does she look at (d). We then replicate the observed attentional behavior through a deep network model (b).

We train our model on the DR(eye)VE data [6], under the hypothesis that i) individuals were thoughtfully driving and ii) distractions have no predefined pattern and can thus be considered outliers. Given that the provided gaze annotation is reliable, we aim to generalize from *what the driver was looking at* in annotated sequences to *what the driver should be looking at* in unseen scenarios.

Code and pre-trained model are available at:

<https://github.com/francescosolera/dreyeving>.

II. RELATED WORKS

Our proposal relates to the following topics.

Head pose and gaze estimation. [7], [8], [9], [10] estimate the driver’s gaze direction using both head and eye location cues. They aim to infer eyes off-the-road and other dangerous driving habits, as [11], [12], [13]. Others exploit driver’s attention to predict next maneuvers, and use either head-pose estimation and GPS/maps [2] or collect precise gaze data through eye tracking devices [14]. Conversely, [15] doesn’t employ driver’s attention information (neither measured nor estimated) to predict actions. They show a strict correlation between the visual appearance of the scene and driver’s action such as steering, turning or braking. This finding also motivates our work, where we try to model the driver’s attention from the external scene only.

Video saliency. In this work we model the driver’s focus through attentional maps, which resemble video saliency maps. Video saliency is typically tackled by extracting low level (or bottom-up) features from each frame an unsupervised [16], [17], [18] or supervised [19], [20] manner.

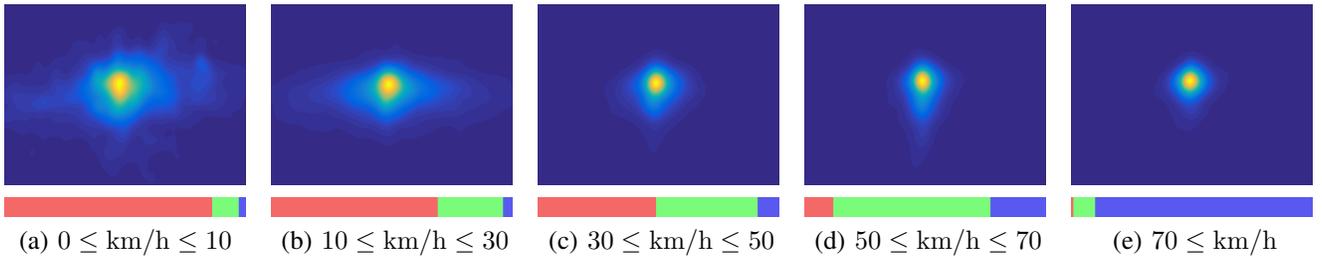


Fig. 2. The driver’s attention converges towards the vanishing point as the speed gradually increases. (a) When the car is almost stationary, the driver is distracted by various objects in the scene. (b) At slightly higher speeds, the focus extends towards the sides of the frame: indeed this is the typical speed at which we take on curves. (c-e) As the speed increases further, the driver’s gaze deviates less and less from the vanishing point of the road. See Sec. III-A for details about the red, green and blue bars. Best viewed on screen.

Temporal dependencies are enforced by conditioning the prediction on previous frames [21] or by means of optical flow [19], [20]. To the best of our knowledge, this work is the first one to exploit deep learning to tackle video saliency.

III. EXPERIMENTAL ANALYSIS OF DRIVER’S ATTENTION

In this section we investigate the attentional mechanisms involved in driving. To this end, we rely on the recently proposed DR (*eye*) VE [6] dataset, which is a collection of 74 sequences 5 minutes long, featuring different landscapes, weather scenarios, lighting conditions and 8 drivers. Every sequence is composed of two videos, capturing both the driver’s and the car point of view. While the driver’s gaze is synchronized with the former video, ground truth attentional maps are provided on the latter through a homographic projection followed by spatial smoothing and time integration. These latter post-processing steps attenuate subjectivity of the drivers’ scan-path and reduce measurement inaccuracies in the ground truth. More formally, we can define an attentional map as a 2D grid where each cell represents the probability that the corresponding pixel of the roof-mounted camera image is within the driver’s focus of attention.

A. Where do we attend while driving?

The first observation that emerges from the analysis of the dataset is that drivers tend to focus their gaze towards the vanishing point of the road – often uncaring of close-by road signals, pedestrians walking on a sidewalk or cars coming from the opposite direction. This phenomenon can be qualitatively appreciated in Fig. 3. This finding can be explained through the notion of *peripheral vision*, introduced in physiological studies to indicate the area around the fixation

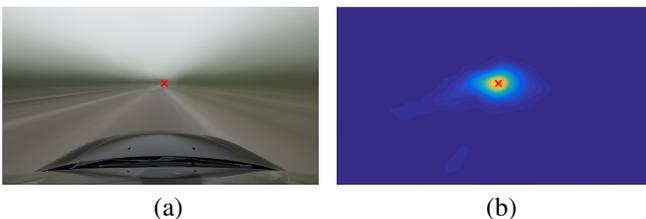


Fig. 3. Frames (a) and attentional maps (b) averaged across the whole sequence 02. The red cross marks the mode of all fixations for this sequence, qualitatively highlighting the link between driver’s focus and the vanishing point of the road.

point where objects and events can still be perceived and interpreted [22]. Accordingly, the largest area of attention is covered by focusing on the vanishing point of the road because, in the task of driving, many of the objects worth of attention have already been perceived when distant.

Additionally, we measured that the average location of the gaze deviates from the road vanishing point as a function of speed and landscape. In fact, due to limited resources our brain compensates the increase of information density - either *spatial* or *temporal* - with a reduction of the visual field size, so that everything that is perceived can also be elaborated [23]. We have an increase in information *spatial density* in complex scenes, where both many interesting elements and distractors are present. Similarly, when the driver travels at higher speed, we have an increase in information *temporal density*, e.g. the amount of information we need to elaborate per unit of time. Contextually to a decreased visual field, the driver needs to move the gaze around the scene more often to capture all the informative elements. Fig. 2 illustrates how the average gaze position changes at different speeds. When the speed increases - from less than 10km/h (a) to more than 70km/h (e) - the average gaze position converges towards the vanishing point of the road, suggesting an increased visual field size able to capture enough of the scene without shifting the gaze. The increased field size is the result of a linear increase of *temporal density* counterbalanced by a stronger decrease of information *spatial density*. In fact, low speed frames occur frequently in downtown sequences where high information *spatial density* is expected, while high speed frames mainly belong to highway driving events, where the amount of information is scarcer. The bar plots in Fig. 2 measure the amount of downtown (red), countryside (green) and highway (blue) frames that concurred to generate the average gaze position for a specific speed range.

B. What do we attend while driving?

To further investigate the attentional mechanism involved in the driving task, in this section we consider *what* the driver is looking at. To conduct this analysis we leverage on semantic segmentation - the task of assigning a label to each pixel according to what object it belongs to. We rely on the network proposed by Yu and Koltun [24], trained on the *cityscapes* [25] dataset, to recognize the following 10 categories: Road, sidewalk, buildings, traffic lights / signs,

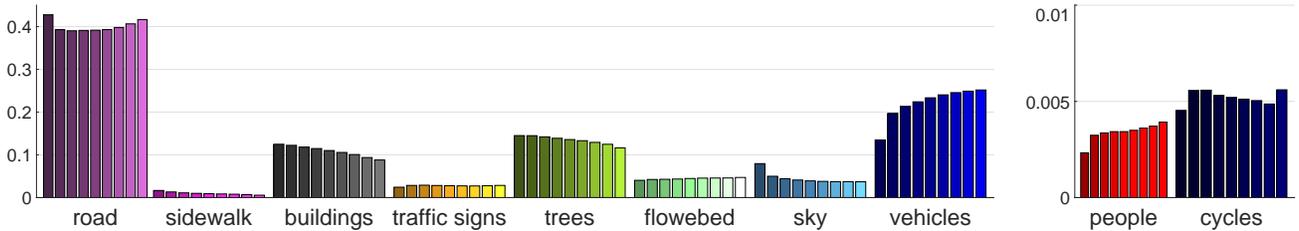


Fig. 4. Proportion of semantic categories hit by the driver gaze map when thresholded at increasing values (from left to right). Upward trends suggest the element was the focus of the gaze, while downward trends mean the presence of the gaze on the object was only circumstantial. See Sec. III-B for details.

trees, road limits, sky, people, vehicles and cycles. Once all frames are labeled we can answer the question of *what is the driver looking at* by measuring the segmentation labels density inside the attentional map¹, see Fig. 1(d). Since the attentional maps have continuous values in $[0, 1]$, we first build binary maps through linearly spaced thresholds. As the threshold approaches 1, we shrink the observed portion of the scene close to the fixation point. Fig. 4 shows how different object categories react when such threshold is changed. In particular, an upward trend reveals that a specific category was indeed the focus of the gaze since, as the observed area decreases, the proportion of pixels labeled in that way increases. Interestingly, this is true for traffic lights, road limits, vehicles and people. Similarly, a downward trend reveals that focus on many categories is often only circumstantial, *i.e.* the categories contain objects that happened to be close to what the driver was really looking at. Examples include sidewalks, buildings, trees and sky. By fixing a threshold, the plot elucidates on the proportion of observed categories across all the sequences. Our focus is dominated by road and vehicles, but we often observe buildings and trees even if they contain little information useful to drive.

IV. NETWORK MODELS

In the previous section we studied the driver’s focus under a variety of conditions, aiming at providing insights on the mechanisms that govern human attention while driving. Below, we learn a deep attentional model that, given a driving sequence, is able to focus where the human driver would.

As various recent works show [26], [27] when dealing with videos, taking explicitly into account the temporal dimension of the input in the network architecture can lead to results that easily outclass the single-frame-input baselines in various high-level video analysis tasks such as video classification and action recognition among others. Here, we make the assumption that a short video sequence (*e.g.* half a second) contains enough information to successfully predict where the driver should look in that situation. Indeed, it can be argued that humans take even less time to react to a stimulus. For this reason, we build our approach on 3D convolutions which take as input a fixed-size sequence of 16 consecutive frames from a video (called from now on *videoclip*) and outputs the gaze map for the last frame of the input clip.

¹This analysis comes with the caveat of being dependant on the quality of the semantic segmentation algorithm (which we cannot evaluate on our dataset, if not qualitatively) and the precision of the eye tracking glasses.

A. Coarse gaze prediction module

The core of our deep network model is a fully convolutional network whose architecture is represented in Fig. 5. The first half of the network acts as an encoder and maps the input videoclip in feature space. Conversely, the second block decodes the feature representation in an attentional map which has the same width and height of the input videoclip, but singleton temporal dimension. In order to perform the encoding, we employ the C3D network by [26] with pre-trained weights and few minor modifications, such as dropping the last convolutional and the fully connected layers in order to maintain spatial information that would be otherwise discarded. See Fig. 5 for further details.

During training we resize the training images to 128×128 and then randomly crop them to 112×112 , following the training process described in [26]. Due to the strong bias towards the vanishing point, this standard cropping policy turned out to be insufficient for creating a real variety in the location of the attentional maps; thus the network prediction resulted strongly attracted towards the center of the image.

B. Complete architecture

In order to solve the aforementioned problem, we moved to an extremely more aggressive crop policy. In this second attempt the training sequences were resized to 256×256 before cropping them to 112×112 . This way the crop constituted less than a quarter of the original image, creating enough variety in the location of the ground truth. This posed however a new challenge: As the network was trained on small crops, it learned to predict better localized but significantly wider attentional maps, reflecting the proportion between salient and non-salient areas seen at training time.

We thus enhanced the network architecture towards a more sophisticated multi-input multi-output approach, shown in Fig. 6. The network still takes a 16 frames videoclip as input, although now there are three different data streams. The first stream provides the model with a randomly cropped videoclip, as explained above. This videoclip passes through the encoding-decoding pipeline and produces a coarse prediction: A first loss on this output is employed to force the model to learn variety in the prediction position and avoid the trivial *hit-the-center* solution. The second stream feeds the model with the same uncropped videoclip, but resized to match the input shape. This videoclip also goes through the encoding-decoding stack, producing a new coarse uncropped saliency prediction. The prediction is

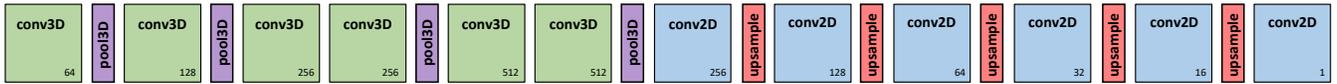


Fig. 5. COARSE prediction architecture. The first part of the network performs the feature encoding. The input videoclip is a tensor of size $3 \times 16 \times 112 \times 112$ that undergoes a sequence of conv3D and pool3D layers that gradually squeeze it to size $512 \times 1 \times 7 \times 7$. All conv3D have kernel size (3,3,3) and ReLU activation units; all pool3D have pool size (2,2,2) except the first one that has pool size (1,2,2). In order to obtain a saliency map with the same spatial size of the input frame, the feature representation is decoded through a series of intertwined layers of conv2D and $\times 2$ upsampling on the spatial dimensions. All conv2D have kernel size (3,3) and are followed by leaky ReLU activations with $\alpha = .001$. As a result, the output of the network is a tensor of size $1 \times 112 \times 112$, *i.e.* the predicted attentional map.

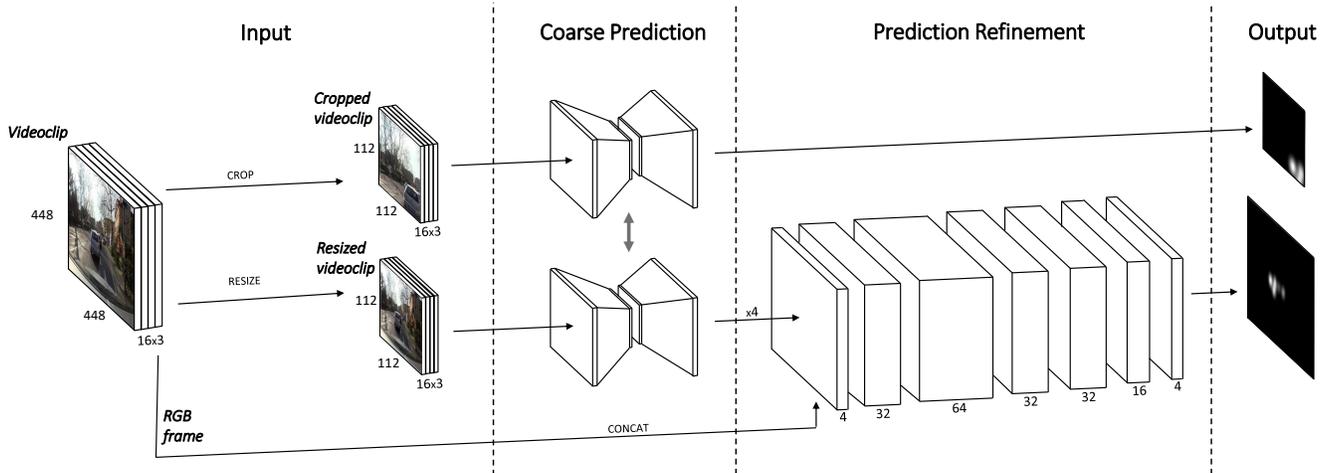


Fig. 6. COARSE+FINE prediction architecture. The COARSE module (see Figure 5) is applied to both a cropped and a resized version of the input tensor, which is a videoclip of 16 consecutive frames. The cropped input is used during training to augment the data and the variety of ground truth attentional maps. The prediction of the resized input is stacked with the last RGB frame of the videoclip and fed to a series of convolutional layers (FINE prediction module) with the aim of refining the prediction. All convolutions have kernel size (3,3). Training is performed end-to-end and weights between COARSE modules are shared. At test time, only the refined predictions are used.

then upsampled and stacked on the RGB image of the last frame of the videoclip, which is provided as third input. Eventually, the concatenated tensor is passed through a last block of convolutions, with the purpose of refining the prediction. This last step also exploits the appearance information of the original videoclip. On the output of this block - our final attentional map prediction - we compute the second loss.

In the following experimental sections we evaluate both quantitatively and qualitatively the proposed model. In Sec. V we rely on saliency metrics to measure the network’s performance against several baselines and state-of-the-art video saliency methods; while in Sec. VI we analyze how well the model mimics the driver’s focus dynamics.

V. EXPERIMENTS

In this section we quantitatively measure the performance of the proposed COARSE and COARSE+FINE models against different baselines. Following the guidelines in [28], for the evaluation phase we rely on Pearson’s Correlation Coefficient (CC) and Kullback–Leibler divergence (KL) measures.

Training details. The encoding half of the COARSE network is initialized with pre-trained weights [26]. Training sequences are randomly mirrored to augment the data. End-to-end training minimizes the Mean Squared Error for both losses of the COARSE+FINE model; we employ Adam optimizer with parameters suggested in the original paper [29]. Train and test are split according to [6] and 500 central frames from each training sequence are used for validation.

A. Keeping it simple: Baselines from saliency

It is widely known that a centered Gaussian, stretched to the aspect ratio of the image, makes for an incredibly effective baseline for the visual saliency task. This static baseline scores better than many methods benchmarked on the MIT300 [30] dataset. Section III-A revealed that a similar bias affects the DR (eye) VE dataset (see Fig. 3). Thus, to validate the proposed model, we compare it against both the aforementioned baseline and a more task-driven version of it built as the average of all training set attentional maps, Fig. 7(c) and (d). Results are reported in Tab. I: The second baseline performs better than both the Gaussian baseline and the COARSE model on the test set, but significantly worse than COARSE+FINE model. Indeed the latter has a relative gain of about +25% over the COARSE model and about +15% over the average training ground truth.

TABLE I

TRAINING AND TEST RESULTS OBTAINED BY BOTH THE BASELINES AND THE PROPOSED NETWORKS. SEE TEXT FOR DETAILS.

	Train sequences		Test sequences		CC(GT, MEAN) < 0.3	
	CC	KL	CC	KL	CC	KL
Baseline (gaussian)	0.33 ± 0.13	2.50 ± 0.63	0.33 ± 0.13	2.50 ± 0.63	0.22 ± 0.15	2.70 ± 0.74
Baseline (mean train GT)	0.47 ± 0.28	1.65 ± 1.13	0.48 ± 0.27	1.65 ± 1.17	0.17 ± 0.20	2.85 ± 1.31
COARSE	0.47 ± 0.22	1.65 ± 0.96	0.44 ± 0.23	1.73 ± 1.00	0.19 ± 0.18	2.74 ± 1.07
COARSE+FINE	0.62 ± 0.26	1.21 ± 0.99	0.55 ± 0.28	1.42 ± 1.07	0.30 ± 0.24	2.24 ± 1.10

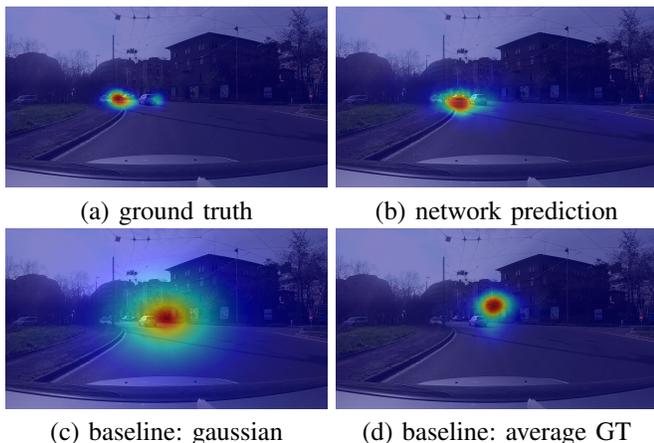


Fig. 7. Comparison between the ground truth, the prediction and the baselines for an example frame.

TABLE II

COMPARISON WITH VIDEO SALIENCY STATE-OF-THE-ART METHODS.

	CC	KL
Wang <i>et al.</i> [18]	0.08 ± 0.11	3.77 ± 0.77
Wang <i>et al.</i> [17]	0.03 ± 0.09	4.24 ± 1.13
Mathe <i>et al.</i> [20]	0.04 ± 0.08	3.92 ± 0.53

B. Comparison with state-of-the-art

In Tab. II we report results from two recent unsupervised video saliency methods [17], [18] and a supervised one [20] on the test set. Both unsupervised methods rely on appearance and motion discontinuities and are easily fooled by the motion of the roof-mounted camera. Unfortunately, [20] is trained on the *Action in The Eye* dataset and thus performs poorly on this domain. As far as we know, no supervised video saliency method releases the source code allowing us to re-train it, nor reports performance on DR(eye)VE. Results shown in Tab. II call for supervised methods aware of both the semantic of the scene and the peculiarities of the task. To our knowledge, our proposal is the first deep model for driving attention, and video saliency to a greater extent.

C. New annotations to escape the bias

Despite showing good results, the baselines introduced in Sec. V-A are of no interest for the driving task as they are not able to generalize when required. There is a strong unbalance between lots of trivial-to-predict scenarios of little interest and few but important hard-to-predict events. To enable the evaluation of our model under such circumstances, we select

from the DR(eye)VE dataset those sub-sequences whose ground truth poorly correlates with the average ground truth of the whole sequence ($CC < 0.3$), under the assumption that in these situations something worth noticing is drifting the driver's focus from the vanishing point of the road. The last column of Tab. I reports the results computed on such subset. When tested on these sequences, the Gaussian baseline outperforms both the COARSE model and the average training ground truth baseline. To interpret this result, consider that when measuring a distance between distributions, a high probability but wrong prediction is severely penalized over a somehow uncertain prediction (*i.e.* a Gaussian with high variance). Nonetheless, the COARSE+FINE model scores higher than all other methods with a relative gain of about +35% over the second best.

VI. DO WE CAPTURE THE ATTENTIONAL DYNAMICS?

In Sec. V we quantitatively evaluated the proposed network. Here, we qualitatively investigate the ability of the model to learn both *where* and *what* a human driver would focus while driving. The results are then compared against the analysis previously introduced in Sec. III.

Figure 8 shows the average attentional maps predicted by our model, arranged by speed range. On each plot we also overlay precision errors (green) and recall errors (red), *i.e.* pixels whose value differs by more than 10% from the analogous ground truth plots reported in Fig 2. We observe that i) the model generally succeeds in capturing the location of the driver gaze at different speed, ii) errors are mostly due to precision (prediction is wider than ground truth) and iii) errors decrease as the speed increases, as the lower variance of the gaze at high speed makes the modeling task easier.

We repeat the analysis of Sec. III-B employing our predictions to determine how often our model focuses on different object categories. Such comparison is shown in Fig. 9 where values are reported on a log scale in order to highlight differences for low populated categories. Gray bars represent the ground truth reference values averaged across different thresholds (see Sec. III-B for details), while prediction results bars are colored based on the semantic segmentation color map introduced in Fig. 1(d). According to Fig. 9, some level of absolute error is found for all categories. Yet, our model was able to coarsely recognize the relative importance of different categories, *e.g.* road and vehicles still dominate the list, while people and cycles are at the bottom. This correlation is confirmed by Kendall Rank Coefficient, which scored 51.11% when computed on the two bar series.

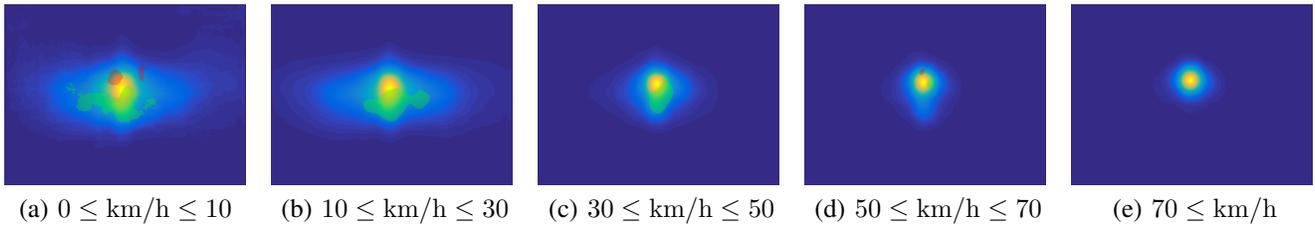


Fig. 8. Model prediction averaged across all test sequences and grouped by driving speed. We highlight areas in which the mean prediction deviates by more than 10% from the mean ground truth. Precision errors are overlaid in green, while recall error in red.

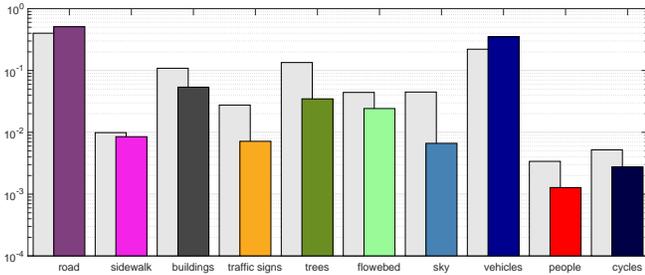


Fig. 9. Comparison between ground truth (gray bars) and predicted attentional maps (colored bars) when used to mask semantic segmentation of the scene. Absolute errors exist, but the two bar series agree on the relative importance of different categories.

VII. CONCLUSIONS

In this work we investigated the spatial and semantic attentional dynamics of the human driver and designed a deep network able to replicate such behavior while driving. These results eventually pave the way for a new assisted driving module, where real-time attentional maps support the driver by both decreasing fatigue and helping in keeping focus. We argue that such attentional maps can be less invasive than other Advanced Driver Assistance Systems that directly act on the car (e.g. by activating breaks), whereas attentional maps leave full control to the driver.

REFERENCES

- [1] U. D. of Transportation, “Federal automated vehicles policy,” <https://www.transportation.gov/sites/dot.gov/files/docs/AV%20policy%20guidance%20PDF.pdf>, September 2016.
- [2] A. Jain, H. S. Koppula, B. Raghavan, S. Soh, and A. Saxena, “Car that knows before you do: Anticipating maneuvers via learning temporal driving models,” in *ICCV*, 2015, pp. 3182–3190.
- [3] B. Fröhlich, M.ENZweiler, and U. Franke, “Will this car change the lane? turn signal recognition in the frequency domain,” in *IEEE IVS*, 2014, pp. 37–42.
- [4] P. Kumar, M. Perrollaz, S. Lefevre, and C. Laugier, “Learning-based approach for online lane change intention prediction,” in *IEEE IVS*, 2013, pp. 797–802.
- [5] B. Morris, A. Doshi, and M. Trivedi, “Lane change intent prediction for driver assistance: On-road design and evaluation,” in *IEEE IVS*, 2011, pp. 895–901.
- [6] S. Alletto, A. Palazzi, F. Solera, S. Calderara, and R. Cucchiara, “DR(EYE)VE: A dataset for attention-based tasks with applications to autonomous and assisted driving,” in *CVPR Workshops*, 2016.
- [7] A. Tawari, K. H. Chen, and M. M. Trivedi, “Where is the driver looking: Analysis of head, eye and iris for robust gaze zone estimation,” in *ITSC*, 2014, pp. 988–994.
- [8] R. Valenti, N. Sebe, and T. Gevers, “Combining head pose and eye location information for gaze estimation,” *Trans. on Image Processing*, vol. 21, no. 2, pp. 802–815, 2012.
- [9] T. Bär, J. F. Reuter, and J. M. Zöllner, “Driver head pose and gaze estimation based on multi-template icp 3-d point cloud alignment,” in *Trans. on Intelligent Transportation Systems*, 2012, pp. 1797–1802.
- [10] R. O. Mbouna, S. G. Kong, and M.-G. Chun, “Visual analysis of eye state and head pose for driver alertness monitoring,” *Trans. on Intelligent Transportation Systems*, 2013.
- [11] A. Tawari and M. M. Trivedi, “Robust and continuous estimation of driver gaze zone by dynamic analysis of multiple face videos,” in *IEEE IVS*, 2014, pp. 344–349.
- [12] F. Vicente, Z. Huang, X. Xiong, F. De la Torre, W. Zhang, and D. Levi, “Driver gaze tracking and eyes off the road detection system,” *Trans. on Vehicular Tech.*, vol. 16, no. 4, pp. 2014–2027, 2015.
- [13] L. Fridman, P. Langhans, J. Lee, and B. Reimer, “Driver gaze estimation without using eye movement,” *arXiv:1507.04760*, 2015.
- [14] F. Lethaus, M. R. Baumann, F. Köster, and K. Lemmer, “Using pattern recognition to predict driver intent,” in *Int’l Conf. on Adaptive and Natural Computing Algorithms*. Springer, 2011, pp. 140–149.
- [15] N. Pugeault and R. Bowden, “How much of driving is pre-attentive?” *Trans. on Vehicular Tech.*, 2015.
- [16] T. Mauthner, H. Possegger, G. Waltner, and H. Bischof, “Encoding based saliency detection for videos and images,” in *CVPR*, 2015.
- [17] W. Wang, J. Shen, and F. Porikli, “Saliency-aware geodesic video object segmentation,” in *CVPR*, 2015, pp. 3395–3402.
- [18] W. Wang, J. Shen, and L. Shao, “Consistent video saliency using local gradient flow optimization and global refinement,” *Trans. on Image Processing*, vol. 24, no. 11, pp. 4185–4196, 2015.
- [19] S.-h. Zhong, Y. Liu, F. Ren, J. Zhang, and T. Ren, “Video saliency detection via dynamic consistent spatio-temporal attention modelling,” in *AAAI*, 2013.
- [20] S. Mathe and C. Sminchisescu, “Actions in the eye: Dynamic gaze datasets and learnt saliency models for visual recognition,” *Trans. on PAMI*, vol. 37, no. 7, pp. 1408–1424, July 2015.
- [21] D. Rudoy, D. B. Goldman, E. Shechtman, and L. Zelnik-Manor, “Learning video saliency from human gaze using candidate selection,” in *CVPR*, 2013, pp. 1147–1154.
- [22] J. Sardegna, S. Shelly, and S. Steidl, *The encyclopedia of blindness and vision impairment*. Infobase Publishing, 2002.
- [23] J. Rogé, T. Pébayle, E. Lambilliotte, F. Spitzenstetter, D. Giselbrecht, and A. Muzet, “Influence of age, speed and duration of monotonous driving task in traffic on the drivers useful visual field,” *Vision research*, vol. 44, no. 23, pp. 2737–2744, 2004.
- [24] F. Yu and V. Koltun, “Multi-scale context aggregation by dilated convolutions,” *arXiv:1511.07122*, 2015.
- [25] M. Cordts, M. Omran, S. Ramos, T. Rehfeld, M.ENZweiler, R. Benenson, U. Franke, S. Roth, and B. Schiele, “The cityscapes dataset for semantic urban scene understanding,” in *CVPR*, 2016.
- [26] D. Tran, L. Bourdev, R. Fergus, L. Torresani, and M. Paluri, “Learning spatiotemporal features with 3d convolutional networks,” *arXiv:1412.0767*, 2014.
- [27] A. Karpathy, G. Toderici, S. Shetty, T. Leung, R. Sukthankar, and L. Fei-Fei, “Large-scale video classification with convolutional neural networks,” in *CVPR*, 2014, pp. 1725–1732.
- [28] Z. Bylinskii, T. Judd, A. Oliva, A. Torralba, and F. Durand, “What do different evaluation metrics tell us about saliency models?” *arXiv:1604.03605*, 2016.
- [29] D. Kingma and J. Ba, “Adam: A method for stochastic optimization,” *arXiv:1412.6980*, 2014.
- [30] Z. Bylinskii, T. Judd, A. Borji, L. Itti, F. Durand, A. Oliva, and A. Torralba, “Mit saliency benchmark,” 2015.