

Utilizing Eye Gaze to Enhance the Generalization of Imitation Networks to Unseen Environments

Congcong Liu^{*1} Yuying Chen^{*1} Lei Tai¹ Ming Liu¹ Bertram Shi¹

Abstract

Vision-based autonomous driving through imitation learning mimics the behaviors of human drivers by training on pairs of data of raw driver-view images and actions. However, there are other cues, e.g. gaze behavior, available from human drivers that have yet to be exploited. Previous research has shown that novice human learners can benefit from observing experts' gaze patterns. We show here that deep neural networks can also benefit from this. We demonstrate different approaches to integrating gaze information into imitation networks. Our results show that the integration of gaze information improves the generalization performance of networks to unseen environments.

1. Introduction

End-to-end deep learning has captured much attention and has been widely applied in many autonomous control systems. Successful attempts have been made in end-to-end driving through imitation learning (Bojarski et al., 2016; Codevilla et al., 2018).

Behavioral cloning has been successfully used for many tasks including off-road driving (Muller et al., 2006) and lane following (Bojarski et al., 2016). It is both simple and efficient. It follows a teacher-student paradigm, where students learn to mimic the teacher's demonstration. Previous behavioral cloning work for autonomous driving mainly focused on learning only the explicit mapping from the sensory input to the control output, paying no consideration to other potential cues from the teacher that might be beneficial.

^{*}Equal contribution ¹Department of Electronic and Computer Engineering, The Hong Kong University of Science and Technology, Hong Kong SAR, China. Correspondence to: Yuying Chen <ychen@connect.ust.hk>, Congcong Liu <cliubh@connect.ust.hk>.

While executing tasks, humans attend to behaviorally relevant information using saccades, which direct gaze towards important areas. In the context of driving, a driver's gaze contains rich information related to his/her intent and decision making. Researches have shown that novice human learners can benefit from observing experts' gaze (Vine et al., 2012). For example, Yamini et al. (2017) showed that viewing the expert gaze videos can improve the hazard anticipation ability of novice drivers. Therefore, it is very promising to investigate whether deep driving networks trained by behavioral cloning might also benefit from exposure to expert gaze patterns.

Whether and how human gaze can help autonomous driving has been under-explored (Alletto et al., 2016). Palazzi et al. (2017) analyzed gaze data in different driving conditions. They trained a network to predict eye gaze and demonstrated a strong relationship between gaze patterns and driving conditions. However, they did not apply their results to autonomous driving.

We trained a conditional generative adversarial network (GAN) to predict human gaze maps accurately while driving in both familiar and unseen environments. We incorporated the predicted gaze maps into end-to-end networks through two different methods.

First, we added the gaze map as an additional input to the network. This is a fairly straightforward approach to incorporating additional information. Treating the gaze map as an additional image input has the disadvantage that it increases the complexity of the network. Since most values of the gaze map are close to zero, this additional complexity is inefficiently utilized. Second, we used the gaze map to modulate the dropout probability. Since the human saccadic eye movements directs high resolution processing and attention towards different area of the visual scene, we hypothesized that gaze behavior may be better treated as a modulating effect than as an additional input.

Both approaches improve model generalization to unseen environments. We demonstrate that modelling auxiliary cues not directly related to the control commands improves the performance of imitation learning. This work takes imitation learning to the next step, by showing how deep networks can benefit from a more complete understanding

of human expert behaviors.

2. Methodology

In this section, we first introduce the network architecture used to estimate the gaze map. Then we describe the two ways we used to incorporate the estimated gaze map into imitation network.

2.1. Gaze Network




Figure 1. The estimated gaze map and ground truth gaze trajectories visualized as heatmaps and green lines superimposed on the driver-view images. The first row corresponds to the two scenes from environments seen during training. The second row show scenes from two environments unseen during training. On the heatmaps, red areas indicate areas with more fixations.

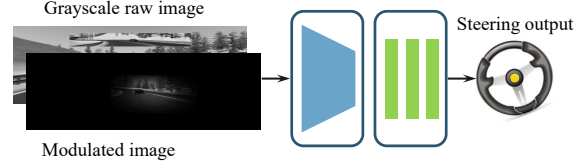
We generate estimated gaze maps by a conditional GAN following the *Pix2Pix* architecture (Isola et al., 2017). The gaze network was trained on pairs of driver-view images and ground truth gaze maps in a manner similar to the way deep networks have been trained to generate saliency maps (Cornia et al., 2016).

Based on our observation of the ground truth gaze trajectories, we find the subject mostly looks at the center area of the image. Therefore, we used a static gaze map consisting of a single Gaussian at the center of the image as a baseline for gaze network evaluation.

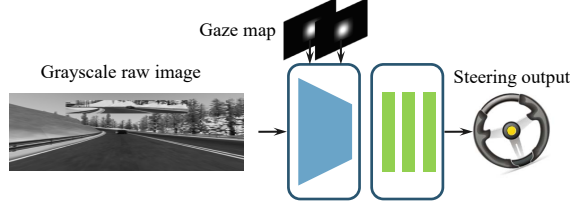
2.2. Imitation Learning with Gaze

We implemented two different methods to incorporate gaze information into the driving network as shown in Fig.2. Both gaze incorporation methods were implemented and evaluated with the networks based on the PilotNet (Bojarski et al., 2016), which has five convolutional layers and four fully connected layers. The activation function for the networks was the ReLU.

We studied two methods for integrating gaze into the imitation learning network. In the first method, we used the estimated gaze map generated by the gaze network to create an additional input to the network. As shown in Fig.2 (a), we pixel-wise multiplied the greyscale original driver-



(a) Architecture with gaze map as input



(b) Architecture with gaze map modulated dropout

Figure 2. (a) The gaze map directly modulates the input image by pixel-wise multiplying it. The grayscale image and the modulated image are stacked as the input to the network. (b) The gaze map modulates the dropout probability spatially. The gaze modulated dropout selectively keep the features after first and second convolutional layers.

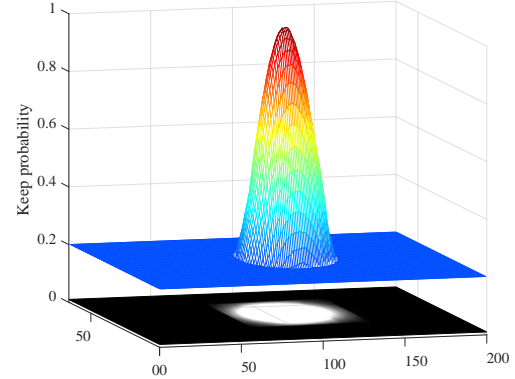


Figure 3. Keep probability settings for gaze-modulated dropout. The bottom image is the corresponding gaze map. At the training stage, keep probabilities are utilized to control whether units are kept or dropped. At the testing state, similar to dropout, the gaze modulated dropout multiplies the features with the keep probability mask generated by the gaze map.

view image with the estimated gaze map. The modulated and original images were stacked as input to the imitation network. For more details, please refer to (Liu et al., 2019).

In the second method, we used the generated gaze map as the mask to modulate dropout in the network. Based on the idea that gaze serves as a filter to focus on important areas and de-emphasize task-irrelevant areas, we utilized the gaze to modulate the dropout probability, so that it was higher in uninteresting regions. The probability setting for the gaze modulated dropout is shown in Fig.3. As shown in Fig.2 (b), we applied the gaze modulated dropout after the first

two convolutional layers. For more details, please refer to (Chen et al., 2019).

		Seen tracks	Unseen tracks
KL	Estimated gaze	0.69	0.88
	Central blob	2.83	2.32
CC	Estimated gaze	0.86	0.83
	Central blob	0.70	0.76

Table 1. The similarity of gaze map estimates to the real gaze map for seen and unseen tracks. KL denotes Kullback-Leibler divergence, and CC denotes Correlation Coefficient.

3. Results

3.1. Gaze Network Evaluation

The examples of estimated gaze maps superimposed with ground truth gaze trajectories shown in Fig. 1 illustrate the good performance of the gaze map prediction for both seen and unseen environments.

To evaluate the gaze network quantitatively, we compute the two standard metrics for similarity evaluation: the Kullback-Leibler divergence (KL) and the Correlation Coefficient (CC). Smaller KL and larger CC denote better similarity.

As shown in Table.1, the estimated gaze map closely matched the ground truth gaze map. Compared with the baseline (central gaussian blob), the KL divergence between the estimated gaze map and the real gaze map is significantly smaller on average (75.6% smaller for seen tracks and 62.1% smaller for unseen tracks). The CC between estimated gaze map and real gaze map is significantly larger on average (22.9% larger for seen tracks and 9.2% larger for unseen tracks).

3.2. Imitation Network Evaluation

We trained the imitation network with uniform dropout as the baseline, which we refer to as **No gaze**. We considered the cases where the real gaze map, the estimated gaze map and the central Gaussian blob were used as the input to or as dropout modulator for the network. We refer to the networks with the addition of the input image modulated by the real and estimated gaze maps as **real/estimated gaze as input** respectively. We refer to the architecture that modulates dropout with real and estimated gaze map as **real/estimated gaze dropout** respectively. The network with the central Gaussian blob (**Central blob dropout**) tested the effect of simply emphasizing the center region.

3.2.1. TEST ON DATASET

We evaluate the average prediction errors between the commands generated by the various models and the human

	Seen tracks (deg)	Unseen tracks (deg)
No gaze	2.90	5.58
Real gaze as input	2.86	4.29
Estimated gaze as input	2.85	4.63
Real gaze dropout	2.82	4.00
Estimated gaze dropout	2.84	4.27
Central blob dropout	2.84	4.67

Table 2. Average prediction errors between the commands generated by the various models and the human driver. Errors for the seen tracks are for tracks used in training, but trials not in the training set. Errors for unssen tracks were for tracks and trials not in the training set.

		W/cars	W/o cars
Success rate of cars overtaking (%)	No gaze	67.6	N/A
	Gaze as input	85.2	N/A
	Gaze dropout	88.9	N/A
Ave. dist. traveled between two infractions (km)	No gaze	0.40	0.48
	Gaze as input	0.54	0.67
	Gaze dropout	0.61	0.79

Table 3. Quantitative performance running on an unseen track in the simulator. Gaze maps are obtained from the gaze network running in real time. We compare the network with estimated gaze-modulated dropout (Gaze dropout), the network with estimated gaze map as input (Gaze as input) and the network with uniform dropout (No gaze). We measure the percentage of successful cars overtaking and the average distance driven between infractions (km). Higher numbers are better for both measurements.

driver. The testing results shown in Table.2 demonstrate that both approaches decrease the action estimation error in unseen environments. The network with gaze as input outperforms the baseline (No gaze) by 20.1% and outperforms the central blob dropout network by 4.5% for unseen tracks on average. The network with gaze dropout outperforms the baseline (No gaze) by 25.9% and outperforms the central blob dropout network by 11.5% for unseen tracks on average.

The gaze modulated dropout shows better performance than gaze as input. The network with real gaze dropout outperforms the network with real gaze as input by 6.8% for unseen tracks and 1.4% for seen tracks. The network with estimated gaze dropout outperforms the network with estimated gaze as input by 7.8% for unseen tracks and 0.35% for seen tracks.

3.2.2. CLOSE LOOP PERFORMANCE

We further test the performance of the networks running either "on an unseen track" or "on X unseen tracks" where X is equal to the number of unssen tracks in the simulator.

As the driving networks are intended for autonomous driving applications, the estimated gaze maps from the gaze network are applied in. For each episode, the agent started from a random location and drove along the path, overtaking cars if needed. Human intervention was given when infractions (collisions or drive outside the lane) occur until the car is back on the road. We tested in two cases: track with cars (W/cars) and without cars (W/o cars). In the case without cars, the agent simply follows the road. We evaluated the performance by the success rate in overtaking cars and the average distance traveled between infractions.

The results are summarized in Table 3. Both approaches improve the close loop performance. The network with estimated gaze dropout outperforms the network with estimated gaze as input in both cases. For success rate of overtaking cars, the network with estimated gaze dropout is 31.5% better than baseline and 4.3% better than the network with estimated gaze as input. For the average distance traveled between two infractions averaged over the two test case, the network with estimated gaze dropout is 58.5% better than the baseline and 15.4% better than the network with estimated gaze as input.

4. Conclusions

In this paper, we proposed the use of gaze information contained in expert demonstrations to improve the generalization of imitation learning networks for autonomous driving to unseen environments. We show that a conditional GAN can generate an accurate estimation of human gaze maps during driving. We studied two ways to incorporate gaze information. Both significantly improve human action estimation accuracy. Better performance is obtained with the gaze modulated dropout. This work demonstrates for the first time that it is possible to incorporate human information about gaze behavior into deep driving networks so that they receive similar benefits as obtained when novice human drivers view expert gaze patterns. This work is an effort to take imitation learning to the next level by implicitly exploiting expert behavior.

Acknowledgements

This work was supported by the National Natural Science Foundation of China (Grant No. U1713211); partially supported by the HKUST Project IGN16EG12 and Shenzhen Science, Technology and Innovation Commission (SZSTI) JCYJ20160428154842603, awarded to Prof. Ming Liu; partially supported by the Hong Kong Research Grants Council under grant 16213617.

References

- Alletto, S., Palazzi, A., Solera, F., Calderara, S., and 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.
- Bojarski, M., Del Testa, D., Dworakowski, D., Firner, B., Flepp, B., Goyal, P., Jackel, L. D., Monfort, M., Muller, U., Zhang, J., et al. End to end learning for self-driving cars. *arXiv preprint arXiv:1604.07316*, 2016.
- Chen, Y., Liu, C., Tai, L., Liu, M., and Shi, B. E. Gaze training by modulated dropout improves imitation learning. *arXiv preprint arXiv:1904.08377*, 2019.
- Codevilla, F., Müller, M., López, A., Koltun, V., and Dosovitskiy, A. End-to-end driving via conditional imitation learning. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 1–9. IEEE, 2018.
- Cornia, M., Baraldi, L., Serra, G., and Cucchiara, R. A deep multi-level network for saliency prediction. In *2016 23rd International Conference on Pattern Recognition (ICPR)*, pp. 3488–3493. IEEE, 2016.
- Isola, P., Zhu, J.-Y., Zhou, T., and Efros, A. A. Image-to-image translation with conditional adversarial networks. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 5967–5976. IEEE, 2017.
- Liu, C., Chen, Y., Tai, L., Ye, H., Liu, M., and Shi, B. E. A gaze model improves autonomous driving. In *Proceedings of the 2019 ACM Symposium on Eye Tracking Research & Applications*. ACM, 2019. to appear.
- Muller, U., Ben, J., Cosatto, E., Flepp, B., and Cun, Y. L. Off-road obstacle avoidance through end-to-end learning. In *Advances in neural information processing systems*, pp. 739–746, 2006.
- Palazzi, A., Abati, D., Calderara, S., Solera, F., and Cucchiara, R. Predicting the driver’s focus of attention: the dr (eye) ve project. *arXiv preprint arXiv:1705.03854*, 2017.
- Vine, S. J., Masters, R. S., McGrath, J. S., Bright, E., and Wilson, M. R. Cheating experience: Guiding novices to adopt the gaze strategies of experts expedites the learning of technical laparoscopic skills. *Surgery*, 152(1):32–40, 2012.
- Yamani, Y., Bıçaksız, P., Palmer, D. B., Cronauer, J. M., and Samuel, S. Following expert’s eyes: Evaluation of the effectiveness of a gaze-based training intervention on young drivers’ latent hazard anticipation skills. In *9th International Driving Symposium on Human Factors in Driver Assessment, Training, and Vehicle Design*, 2017.