

# Resolving Camera Position for a Practical Application of Gaze Estimation on Edge Devices

Linh Van Ma, Tin Trung Tran, Moongu Jeon  
School of Electrical Engineering and Computer Science  
Gwangju Institute of Science and Technology  
Gwangju, South Korea  
{linh.mavan, ttrungtin, mgjeon}@gist.ac.kr

**Abstract**—Most Gaze estimation research only works on a setup condition that a camera perfectly captures eyes gaze. They have not literarily specified how to set up a camera correctly for a given position of a person. In this paper, we carry out a study on gaze estimation with a logical camera setup position. We further bring our research in a practical application by using inexpensive edge devices with a realistic scenario. That is, we first set up a shopping environment where we want to grasp customers gazing behaviors. This setup needs an optimal camera position in order to maintain estimation accuracy from existing gaze estimation research. We then apply the state-of-the-art of few-shot learning gaze estimation to reduce training sampling in the inference phase. In the experiment, we perform our implemented research on NVIDIA Jetson TX2 and achieve a reasonable speed, 12 FPS which is faster compared with our reference work, without much degradation of gaze estimation accuracy. The source code is released at <https://github.com/linh-gist/GazeEstimationTX2>.

**Index Terms**—Gaze estimation, Few shot learning, Edge devices, Customers' Attention, Triangulation.

## I. INTRODUCTION

Businesses are now benefiting from computer vision applications. One of the well-known businesses in the USA, Amazon Go [1] has successfully applied deep learning, sensor fusion, and computer vision for checkout, purchase, and proceed for payment automatically without any human interactions. Hence, understanding customers' attention/behavior/preferences during shopping is crucial to increase avenues. For example, a customer who buys milk usually looks for bread. We can place milk and bread next to each other. Gaze is an individual's perception and awareness of human visual attention. We can track the gaze to know which products customers prefer the most and their shopping behaviors.

Many eye commercial trackers perfectly measure the motion of an eye relative to the head. For example, Tobii [2] can estimate and track eye gaze without requiring recalibration. However, it is expensive and hard to upgrade once a business decides to buy those physical devices. More importantly, up-

This work was partly supported by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) (No. 2019R1A2C2087489), and Ministry of Culture, Sports and Tourism(MCST) and Korea CreativeContent Agency(KOCCA) in the Culture Technology(CT) Research & Development (R2020070004) Program 2021.

gradeable, inexpensive, and easy-to-deploy are the key points to leverage business avenues.

In this paper, we introduce research on nearly real-time eye gaze estimation. We use low-priced RGB cameras to capture eyes gaze frames. Those frames are processed by economical edge devices to find out where a person stares at. More specifically, we first reflect a shopping environment where a customer looks at a shelf to buy goods, groceries. We then find an optimal position inside the shelf to set up an RGB camera. This position supports the camera appropriately grasping the customer's eyes gaze. Afterward, we build a deep eye gaze estimation network for edge devices. In this network, images captured from an RBG camera are first fed into a face detection module. Subsequently, facial landmark detection is employed to find key points on facial images. Finally, gaze estimation is accurately estimated from those obtained key points. We use few-shot learning to reduce the number of samples require to fine-tune a deep gaze estimation network. This few-shot learning also makes our work more easily to be deployed in a real-world application because we can fine-tune our deep network within a few samples. In the experiment, we perform our implemented system on NVIDIA Jetson TX2. We achieve a reasonable speed, 12 FPS which is faster compared with our reference work, without much degradation of gaze estimation accuracy. We also prove that the camera should be set up in an optimal position to increase gaze estimation accuracy.

This paper is organized as follows. Section II presents our based research. In Section III, we logically find an optimal camera position that supports the camera to accurately estimate a person's eyes gaze. Subsequently, we demonstrate our method with several experiments in Section IV. Finally, we conclude this paper with a future research direction.

## II. RELATED WORKS

In gaze estimation, we first need to detect faces using face detectors [3]–[6]. MTCNN proposed in [4] is a fast and efficient facial detection model. MTCNN composes of deep cascaded networks, Proposal Network (P-Net), Refinement Network (R-Net), and Output Network (O-Net). These three networks exploit the properties correlation between face alignment (in R-Net) and face detection (in P-Net) to increase the performance of face detection (in O-Net). Furthermore, the authors propose a new online hard sample mining strategy

leading to an improvement during the training process with fewer manual factors.

Facial landmarks [7], [8] to find key points (68 landmarks) on detected faces, is the next crucial step for gaze estimation. Facial landmark detectors essentially try to label and localize the seven facial regions as follows: (1) Right eyebrow, (2) Left eyebrow, (3) Right eye, (4) Left eye, (5) Nose, (6) Mouth, and (7) Jaw. More specifically, the authors [7] propose a general framework based on gradient boosting for learning an ensemble of regression trees that optimizes the sum of square error loss and naturally handles missing or partially labeled data. In short, an ensemble regression trees is trained to estimate the face’s landmark positions directly from a sparse subset of pixel intensities. Notably, their result can achieve face alignment in milliseconds for a single image. This result brings us a chance to implement facial landmark detection on edge devices where we do not have many computational resources. Fortunately, this module was well implemented in Dlib [9]. In contrast, in [8] the authors propose a HRNetV2, a modification on HRNet [10], to work high-resolution representation learning. This learning leads to stronger representations and higher landmark localization accuracy. Hence, we use this HRNetV2 to detect facial landmarks while fine-tuning our gaze estimation network. It has high accuracy but not fast compared to [7].

Gaze estimation [11], [12] is a process to predict where a person is gazing at given an image with the person’s full face. Similar to our approach [13] argued that each person has a distinct gaze, a person-specific gaze. It leads to limit the accuracy of person-independent gaze estimation networks. Hence, they propose personalizing gaze networks. This network encodes face appearance, gaze direction, and head rotation into latent space by using a disentangling encoder-decoder architecture. Their method allows the network to learn person-specific gaze within a few samples (less than nine samples).

Thanks to a real-time detecting eye blink algorithm proposed in [14], we can determine whether eyes are widely opened or slightly closed. They first use landmark positions to calculate the eye aspect ratio (EAR). A Support Vector Machine [15] classifier is subsequently employed to determine whether eyes are blinking or non-blinking pattern based on EAR value. In our approach, this EAR allows us to check whether our camera position set up on the shelf is optimal or not because eyes are slightly narrower when a person’s eyes look downward. Given a camera position, we can check this EAR value to determine the position can avoid the above effect or not.

In [16], [17], the authors argue that the limitation of accuracy from facial landmark detection comes from the training process, which is lack of quality and quantity of annotated face databases. The databases mostly were manually annotated by a trained expert and the fatigue factor is hard to avoid which lead to the error during the works. Hence, Sagonas et al. [16], [17] propose a unified annotation pipeline with a semi-automatic annotation system. They use Active Orientation Models (AOMs) generative models [18] to train their network

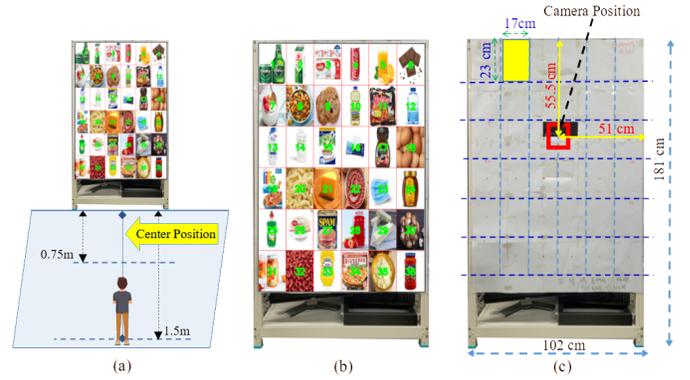


Fig. 1. (a) We ideally model a shelf in a store, (b) each item is mapped to a labeled rectangle number from the top-left to the bottom-right, (c) a simplified version of a shelf in a store.

with an image from mixed expression and viewing angle. The resulted train model can generate accurate annotation in different conditions and can be generalized to unseen images. Along with the problem of the database, the authors [19] further introduce a data normalization method to combine the images and gaze direction to a normalized space, which can cancel out the varieties of head and eye pose positions.

### III. SYSTEM OVERVIEW

We ideally model a physical store in our research environment with a shelf as shown in Fig. 1. It has goods and groceries placed separately in different positions on the shelf. In Fig. 1 (a), a user stands closely 0.75 meters and far at most 1.5 meters from the shelf. In Fig. 1 (b), we divide the shelf into 36 (6x6) rectangles labeled from 1 at the top-left to 36 at the bottom right. Each rectangle has a size of 17 centimeters in width and 23 centimeters in height. Each item is mapped/encoded to one or two rectangles. To simplify, in Fig. 1 (c), we remove groceries and put a large paper with printed 36 rectangles onto the shelf. The camera is located somewhere inside the shelf behind the large white paper. If a unit measurement is not specified, we use centimeters throughout our paper.

Fig. 2 illustrates a side view projection of Fig. 1 (a) (projection from left to right). We name  $B$  is at the top,  $C$  is at the bottom (of the large white paper, not the floor),  $D$  is at the camera location on the shelf.  $A$  locates at the person’s eyes (or middle of them). From the eyes to  $B$ ,  $D$ , and  $C$ , we have two angles  $\alpha_1 = \widehat{BAD}$ ,  $\alpha_2 = \widehat{DAC}$ . The camera should be in the optimal position so that when a person’s head rotates vertically (up and down, or pitch in Euler angle),  $AD$  should equally divide the angle  $\widehat{BAC}$  created when the person looks at the top and bottom of the shelf as shown in Fig. 2. In other words, the camera should make this equality  $\alpha_1 = \alpha_2$  happens. A person looks at the shelf while standing in the range 750 to 1500 centimeters far from the shelf.

Fig. 3 shows an example of nonoptimal camera position (24.5 cm from the top of the shelf) where  $\alpha_1 \neq \alpha_2$ . A person who looks downward with his open eyes. However, the camera determines his eyes are mostly closed. We (humans) tend to

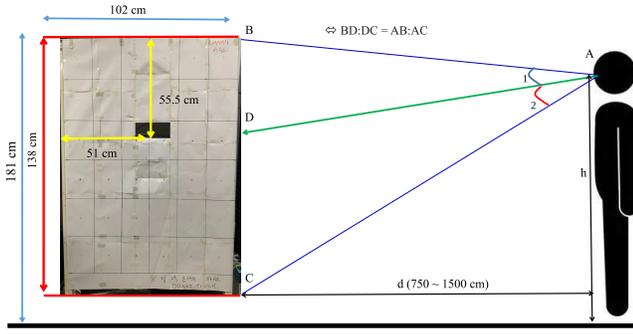


Fig. 2. A shelf with 181 (cm) height and width of 102 (cm). A person stands away from the shelf with a distance ranging from 750 to 1500 (cm). The camera places optimally at the center width and 55.5 (cm) from the top. (A) represents the person’s eyes (or middle of them), (B) top of the shelf, (C) bottom of the large white paper, not floor, (D) is the optimal camera position on the self.



Fig. 3. A person looks downward to the bottom of the shelf. His eye is widely opened, but the camera position is set up inappropriately (about 24.5 cm from the top of the shelf) resulting in the camera looks at him with mostly closed eyes. We experimentally detect facial landmarks [16], [17], [20] to calculate Eye Aspect Ratio (EAR) proposed in [14]. This EAR characterizes how eyes are largely or small opened. In [14], the authors determine that eyes are widely opened with EAR above 0.2. If we inappropriately set up the camera, we can obtain EAR 0.0667 with open eyes as shown on the left side.

open our eyes wider while looking upward and oppositely narrower while looking downward. This unwanted effect leads to our eye gaze algorithm (deep learning algorithm) misunderstands that his eyes are closed. The reason is that we train our deep learning model with a dataset with eyes straightly look at a camera. It has no idea to determine whether eyes look downward and are widely opened. We experimentally detect facial landmarks [16], [17], [20] to calculate Eye Aspect Ratio (EAR) proposed in [14]. This EAR characterizes how eyes are largely or small opened. In [14], the authors determine that eyes are widely opened with EAR above 0.2. If we inappropriately set up the camera, we can obtain EAR 0.0667 with open eyes as shown on the left side.

In Fig. 2,  $BC = 138$  cm is the height of the white large paper,  $750 \leq d \leq 1500$  cm is the distance between the shelf and a person,  $h$  is the height of a person minus 4.8 cm (approximation distance from eyes to the top of head).  $b = 181$  cm is the height of the shelf. We apply one property of bisection, if  $\alpha_1 = \alpha_2$  then  $BD : DC = AB : AC$ . In this equality, we have already known  $BC$ , while  $AB$  and  $AC$  can be calculated by using Pythagorean theorem. Given  $h$  is the

TABLE I

CAMERA IS PLACED 55.5 CM FROM THE TOP OF THE SHELF. A PERSON WITH 1800 CM HEIGHT IS RECOMMENDED TO STAY AWAY FROM THE SHELF AT 1212.1 CM.

Item	Person Height	Distance from the shelf
1	1500	851.453
2	1550	928.174
3	1600	996.515
4	1650	1058.101
5	1700	1114.053
6	1750	1165.182
7	1800	1212.100

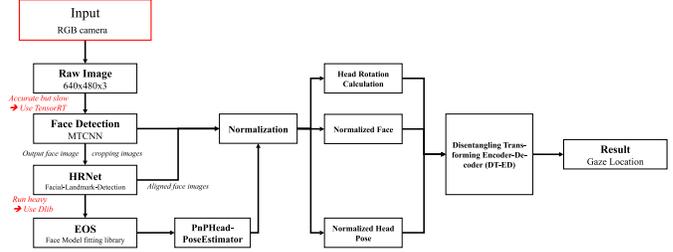


Fig. 4. Workflow of our gaze estimation system. It is largely inspired by FAZE proposed in [13]. We use MTCNN [4] for face detection and HRNetv2 [8] for facial landmark detection in fine-tuning. In inference, these two tasks are computationally expensive for edge devices. Hence, we replace the MTCNN module with a similar one but implemented in TensorRT [22], TensorRT MTCNN. Though HRNetv2 has high accuracy, we replace it with a landmark detection module in Dlib [7], [9].

height of a person, we look to find a position  $d$ . The height follows Gaussian distribution with the global mean height for 159 cm for women and men 171 cm [21]. From  $\frac{DC}{BD} = \frac{AC}{AB}$ , we have  $\frac{AC}{AB} + 1 = \frac{BC}{DB}$ . Using this fact, we put (1), (2) altogether and have equation (3). It has three variables  $d, h$  and  $DB$ , ( $b - BC = 43$ ). We randomly generate  $h$  with its Gaussian distribution (mean is 165, standard deviation is 6), and  $d$  is uniformly distributed in the range  $[750, 1500]$  and obtain  $DB$  is optimal at 55.5 cm from the top of the shelf.

$$AB = \sqrt{d^2 + (b - h)^2}. \quad (1)$$

$$AC = \sqrt{d^2 + (h - (b - 138))^2}. \quad (2)$$

$$DB = \frac{138\sqrt{d^2 + (h - 43)^2}}{\sqrt{d^2 + (b - h)^2} + \sqrt{d^2 + (h - 43)^2}}. \quad (3)$$

If the camera optimally places at 55.5 cm from the top of the shelf, Table I depicts that a person with a specific height must stay far from the shelf to increase gaze estimation accuracy. For example, a person with 170 cm height (we ignore a space from eyes to the top of the human head for simplifying explanation) is recommended to stay away from the shelf at 1114.053 cm.

We use few shot learning techniques to reduce time to fine-tune our deep gaze estimation network. These techniques make our gaze estimation module applicable in real-world

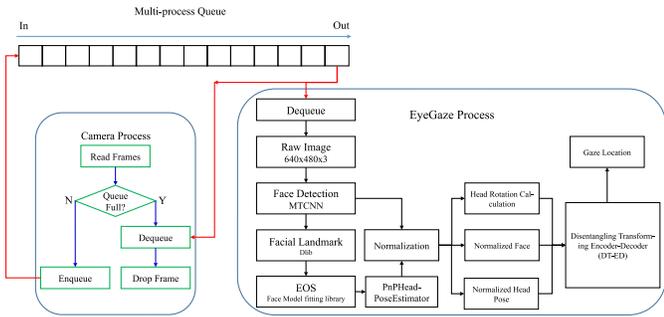


Fig. 5. Working with multiprocessing. Our algorithm cannot process to find gaze location in real-time (30 FPS). We use a queue with two processes. One process named “Camera Process” puts images captured from a camera into the queue. Another process named “EyeGaze Process” obtains the latest recent images in that queue and starts processing to estimate gaze location

applications where we do not have much data to train. It also supports us utilize existing eye gaze research in our configured environment. We intensively inherit FAZE proposed in [13] to build our gaze estimation network. Fig. 4 shows the overall workflow of FAZE. They use MTCNN [4] for face detection and HRNet [8] for facial landmark detection. These two tasks are computationally expensive. Hence, we replace the MTCNN module with a similar one but implemented in TensorRT [22], TensorRT MTCNN Face Detector. Though HRNetv2 has high accuracy, we replace it with a landmark detection module in Dlib [7], [9]. Facial landmark detection in Dlib is fast, which is fairly real-time without degrading accuracy compared with HRNetv2. To maintain high accuracy of inference, we still use the original HRNetv2 and MTCNN to finetune FAZE model. We only use these two replacement modules in inference on edge devices. We use data normalization proposed in [19] to cancel out the significant variability in head pose such as head rotation with respect to the camera.

Despite our great effort, we cannot achieve real-time processing (30 FPS) in edge devices such as Jetson TX2. It can only run 10 to 15 FPS. Hence, we use multiprocessing and queue to ignore frames from the camera. Suppose our gaze estimation algorithm works on frame  $f_i$ , it moves to work on the next latest frame capture from the camera  $f_j$  ( $i > 0, j > i + 1, i, j \in \mathbb{N}$ ). All frames  $\{f_{i+1}, \dots, f_{j-1}\}$  captured during processing frame  $f_i$  are discarded. This procedure ensures that we always capture changes in a person’s head such as rotation. Fig. 5 illustrates our idea of using a queue. We use a queue with two processes. One process named “Camera Process” puts images captured from a camera into the queue. Another process named “EyeGaze Process” obtains the latest recent images in that queue and starts processing to estimate gaze location.

#### IV. EXPERIMENTS

We put a large white paper into the shelf to simulation our real-world shopping environment as shown in Fig. 1 and Fig. 2. It has a width of 102 cm and a height of 138 cm. We divide it into 36 equal rectangles size 17 cm in width and 23 cm in

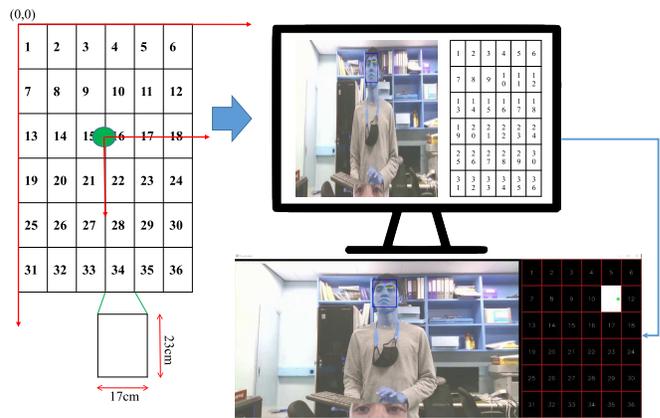


Fig. 6. Mapping from the physical world to screen monitor for visualization. We use the coordinates axis (0,0) at the top-left corner of the shelf to label location of each point center at each rectangle. However, our algorithm only works with the coordinates axis (0,0) at the camera pinhole location (large green dot).

height. The camera is always at the center of width and down from the top at 55.5 cm (calculated optimal position in Section III) as shown in Fig. 6. We use the coordinates axis (0,0) at the top-left corner of the shelf to label the location of each point center at each rectangle. However, our algorithm works with the coordinates axis (0,0) at the camera pinhole location. We obtain our custom dataset to fine-tune our gaze estimation module by asking a user constantly looking at each point centered at each rectangle. In this experiment, we only have one participant. The paper (size 102x138 cm) is mapped to a screen monitor and still maintains the width height ratio. This mapping supports us to visualize our process on the monitor while collecting dataset.

More specifically, we make a Python script to output one video file and one Pickle file [23] contains ground truth (with the origin of the coordinates axis (0,0) at the top-left corner of the shelf). A person needs to look at 36 points (centered at each rectangle) to collect a calibration dataset. We select ten frames at each point. These ten frames contain images that specify a person constantly looks at one point (out of 36). Hence, we have a total of 360 frames in one calibration video. However, to fine-tune our deep learning model, we only use three random images for training and one image for validation from those ten collected images of a given point. In Fig. 6, there have thirty-six labeled squares ranging from 1 to 36. We experimentally find that fine-tuning and validation locations should be equally distributed in the shelf in order to avoid bias in any region (left, right, top, down) respect to camera location. Hence, We use squares with number  $\{8, 11, 26, 29\}$  for validation. Those points location are equally located with respected to camera location. The rest 32 points are used in different sets of fine-tuning as shown in Table II.

Our implementation is intensively based on a few shot learning [13]. We ignore the training phrase since their model was fairly trained with the following number of calibration points  $\{1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16,$

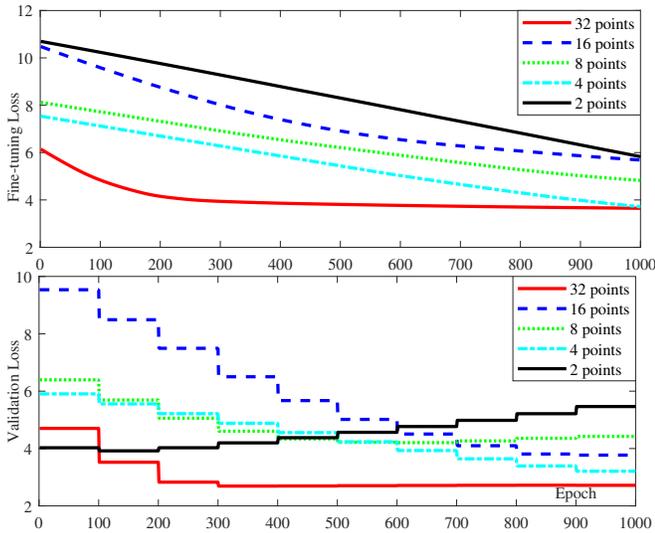


Fig. 7. We fine-tune FAZE network using five different sets of points (located at each rectangle) shown in Table II. If we fine-tune with two points, validation loss increases compared to other sets. Hence, each person needs at least four samples to fine-tune FAZE network. As we increase the number of samples (points) to fine-tune the gaze estimation network, validation loss reduces accordingly.

TABLE II

DIFFERENCE TRAINING (FINE-TUNING) SETS OF POINTS. EACH POINT CENTERED AT EACH RECTANGLE.

Item	Tuning points	Point indexes
1	2	6, 31
2	4	3, 13, 18, 33
3	8	1, 3, 6, 13, 18, 31, 33, 36
4	16	1, 3, 4, 6, 13, 15, 16, 18, 19, 21, 22, 24, 31, 33, 34, 36
5	32	1, 2, 3, 4, 5, 6, 7, 9, 10, 12, 13, 14, 15, 16, 17, 18, 19, 20, 21, 22, 23, 24, 25, 27, 28, 30, 31, 32, 33, 34, 35, 36

17, 18, 32, 64, 128, 256}. We do not re-train their network and only use their pre-trained weight on inference. Logically, gaze estimation accuracy can be improved if we increase the number of calibration points. However, we do not have enough time and labor resources to label our custom dataset. Hence, we use different sets of points to find the reasonable number of calibration points shown in Table II. As mentioned earlier, we fix four points to validate the fine-tuning process.

Fig. 7 illustrates fine-tuning loss (training and validation) with respect to an epoch. If we fine-tune with two points, validation loss increases compared to other sets. Hence, each person needs at least four samples to fine-tune FAZE network. As we increase the number of samples (points), to fine-tune the gaze estimation network, validation loss reduces accordingly.

Fig. 8 visually shows the output while fine-tuning our gaze network. RGB images are feed into our system. Our system outputs (1) face bounding box [4], (2) facial landmark [8], (3) head pose, (4) eyes patch after normalization [19]. We also output (5) the location of where the user gazes at. For

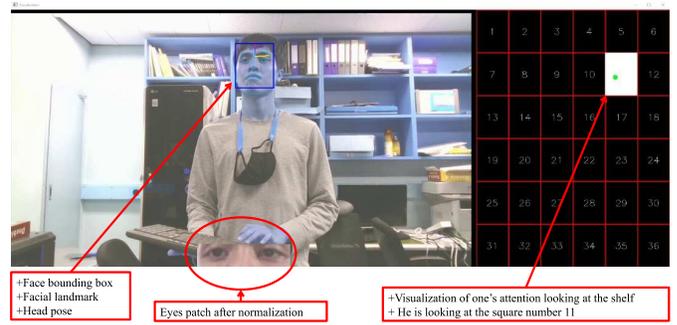


Fig. 8. We display the appearance of a person with (1) face bounding box [4], (2) facial landmark [8], (3) head pose, (4) eyes patch after normalization [19]. We also output (5) the location of where the user gazes at. For example, a user attentively looks at the square number eleven.

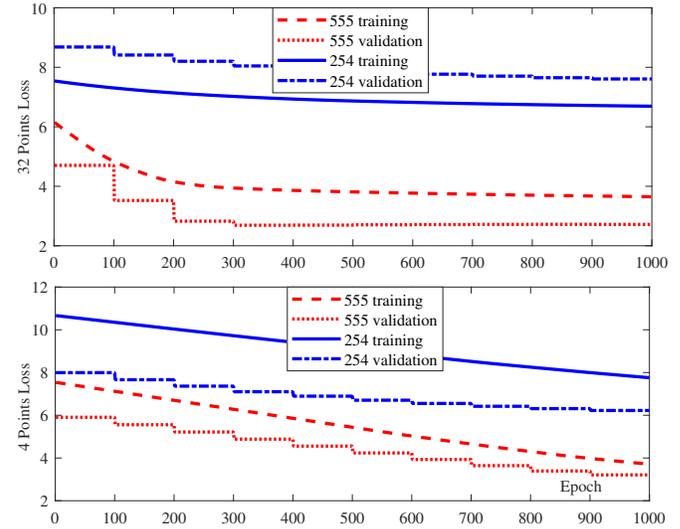


Fig. 9. We set up an RGB camera to capture frames from two different locations centered along the width. (1) optimal position 55.5 cm, and non-optimal position 24.5 cm from the top of the shelf. Fine-tuning results with two sets of 32 and 4 points from Table II. It shows that the camera at position 25.4 cm cannot reduce training and validation loss even though we increase the number of samples from 4 to 32.

example, a user attentively looks at the square number eleven.

If we set up a camera to capture frames inappropriately, we cannot archive high-accuracy gaze estimation. As shown in Fig. 9, we set up an RGB camera to capture frames from two different locations. We compare training and validation loss of two camera positions which are 55.5 and 25.4 cm from the top of the shelf. It shows that the camera at position 25.4 cm cannot reduce training and validation loss even though we increase the number of samples from 4 to 32. Especially, fine-tuning with 32 points even increases validation loss and maintains at 8. Oppositely, the camera at position 55.5 cm logically reduces training and validation loss to 2. Despite our hard efforts in finding a camera position, we still face the eye-looking effects shown in Fig. 10. Though, we mathematically find the optimal camera position 55.5 cm from the top. The eye-looking effect still prevents us to obtain equally opened



Fig. 10. We (humans) tend to open our eyes wider while looking upward and oppositely narrower while looking downward. We have found the optimal position of the camera (55.5 cm from the top) but this effect prevents us to equally capture opened eyes when looking upward and downward. However, we have overcome the effect and estimate eyes gaze more accurately compared to what we can obtain in Fig. 3.

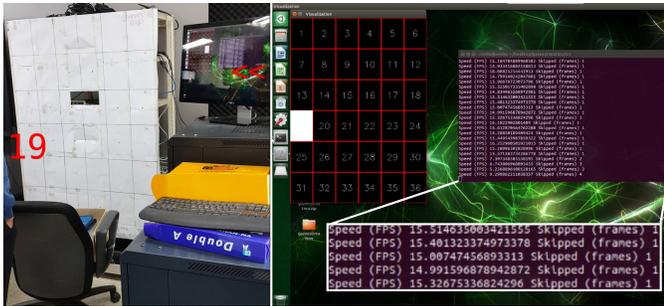


Fig. 11. Our experiment with the optimal camera position at 55.5 cm from the top of the shelf. A user gazes at rectangle number 19 (leftmost side). Our gaze estimation module correctly identifies the user's gazing (shown in a white rectangle). It skips one to five frames in order to process real-time images captured from the camera. The processing speed is about 10 to 15 FPS.

eyes when looking upward and downward but it is much better compared to what we can obtain in Fig. 3.

Fig. 11 shows the final output of our experiment with the optimal camera position at 55.5 cm from the top of the shelf. A user gazes at rectangle number 19 (left side). Our gaze estimation module correctly identifies the user's gazing (shown in a white rectangle, right side). It skips one to five frames to process the latest images captured from the camera. A few captured images are discarded to ensure that our gaze estimation algorithm is not stuck or keeps processing past frames. Every change of personal appearance such as head rotation, eyes blink is constantly processed. The processing speed is about 10 to 15 FPS.

## V. CONCLUSION

In this paper, we bring research in gaze estimation to a real-world application. Gaze estimation has been done intensively in the literature, but it has many limitations and far beyond real-world constraints. We bridge the gap to find the optimal position for the camera. This effort maintains high accuracy grasped from our reference research. Then, we implement to run state-of-the-art gaze estimation on edge devices. Our experiment proves that setting the camera at an appropriate position supports us to obtain eyes gaze correctly. It results in maintaining high accuracy compared to the results that have been done in the research environment. However, we only consider a scenario in which only one person statically

standing in front of the camera and has no movement such as walking. Furthermore, the 3D shape of the face around the eyes strongly affects the gaze estimation accuracy so that the optimal camera position may be different among different people. In the future, we intend to experimentally and quantitatively confirm our optimal camera position with data from multiple participants. Additionally, we plan multiple cameras to avoid constraints from our current research environment making it more realistic. We can also extend it to a scenario where many people move around while looking at the camera.

## REFERENCES

- [1] K. Wankhede, B. Wukkadada, and V. Nadar, "Just walk-out technology and its challenges: A case of amazon go," in *2018 International Conference on Inventive Research in Computing Applications (ICIRCA)*. IEEE, 2018, pp. 254–257.
- [2] A. Gibaldi, M. Vanegas, P. J. Bex, and G. Maiello, "Evaluation of the tobii eyeX eye tracking controller and matlab toolkit for research," *Behavior research methods*, vol. 49, no. 3, pp. 923–946, 2017.
- [3] P. Viola and M. J. Jones, "Robust real-time face detection," *International journal of computer vision*, vol. 57, no. 2, pp. 137–154, 2004.
- [4] K. Zhang, Z. Zhang, Z. Li, and Y. Qiao, "Joint face detection and alignment using multitask cascaded convolutional networks," *IEEE Signal Processing Letters*, vol. 23, no. 10, pp. 1499–1503, 2016.
- [5] S. S. Farfadi, M. J. Saberian, and L.-J. Li, "Multi-view face detection using deep convolutional neural networks," in *Proceedings of the 5th ACM International Conference on Multimedia Retrieval*, 2015, pp. 643–650.
- [6] F. Schroff, D. Kalenichenko, and J. Philbin, "Facenet: A unified embedding for face recognition and clustering," in *CVPR*, 2015, pp. 815–823.
- [7] V. Kazemi and J. Sullivan, "One millisecond face alignment with an ensemble of regression trees," in *CVPR*, 2014, pp. 1867–1874.
- [8] K. Sun, Y. Zhao, B. Jiang, T. Cheng, B. Xiao, D. Liu, Y. Mu, X. Wang, W. Liu, and J. Wang, "High-resolution representations for labeling pixels and regions," *arXiv preprint arXiv:1904.04514*, 2019.
- [9] D. E. King, "Dlib-ml: A machine learning toolkit," *The Journal of Machine Learning Research*, vol. 10, pp. 1755–1758, 2009.
- [10] K. Sun, B. Xiao, D. Liu, and J. Wang, "Deep high-resolution representation learning for human pose estimation," in *CVPR*, 2019, pp. 5693–5703.
- [11] T. Fischer, H. J. Chang, and Y. Demiris, "Rt-gene: Real-time eye gaze estimation in natural environments," in *ECCV*, 2018, pp. 334–352.
- [12] S. Park, A. Spurr, and O. Hilliges, "Deep pictorial gaze estimation," in *ECCV*, 2018, pp. 721–738.
- [13] S. Park, S. D. Mello, P. Molchanov, U. Iqbal, O. Hilliges, and J. Kautz, "Few-shot adaptive gaze estimation," in *ICCV*, 2019, pp. 9368–9377.
- [14] T. Soukupova and J. Cech, "Real-time eye blink detection using facial landmarks," *21st Computer Vision Winter Workshop*, pp. 1–8, 2016.
- [15] W. S. Noble, "What is a support vector machine?" *Nature biotechnology*, vol. 24, no. 12, pp. 1565–1567, 2006.
- [16] C. Sagonas, G. Tzimiropoulos, S. Zafeiriou, and M. Pantic, "300 faces in-the-wild challenge: The first facial landmark localization challenge," in *ICCV Workshops*, 2013, pp. 397–403.
- [17] —, "A semi-automatic methodology for facial landmark annotation," in *CVPR workshops*, 2013, pp. 896–903.
- [18] G. Tzimiropoulos, J. Alabort-i Medina, S. P. Zafeiriou, and M. Pantic, "Active orientation models for face alignment in-the-wild," *IEEE transactions on information forensics and security*, vol. 9, no. 12, pp. 2024–2034, 2014.
- [19] X. Zhang, Y. Sugano, and A. Bulling, "Revisiting data normalization for appearance-based gaze estimation," in *Proceedings of the 2018 ACM Symposium on Eye Tracking Research & Applications*, 2018, pp. 1–9.
- [20] C. Sagonas, E. Antonakos, G. Tzimiropoulos, S. Zafeiriou, and M. Pantic, "300 faces in-the-wild challenge: Database and results," *Image and Vision Computing*, vol. 47, pp. 3–18, 2016.
- [21] M. Roser, C. Appel, and H. Ritchie, "Human height," *Our world in data*, 2013.
- [22] H. Vanholder, "Efficient inference with tensorsrt," 2016.
- [23] M. Pilgrim and S. Willison, *Dive Into Python 3*. Springer, 2009, vol. 2.