



C-Auth: Exploring the Feasibility of Using Egocentric View of Face Contour for User Authentication on Glasses

Hyunchul Lim

Cornell University, USA

hl2365@cornell.edu

Hao Chen

Cornell University, USA

hc732@cornell.edu

Guilin Hu

Cornell University, USA

gh386@cornell.edu

Richard Jin

Cornell University, USA

rj284@cornell.edu

Ryan Mao

Cornell University, USA

rwm275@cornell.edu

Ruidong Zhang

Cornell University, USA

rz379@cornell.edu

Cheng Zhang

Cornell University, USA

chengzhang@cornell.edu

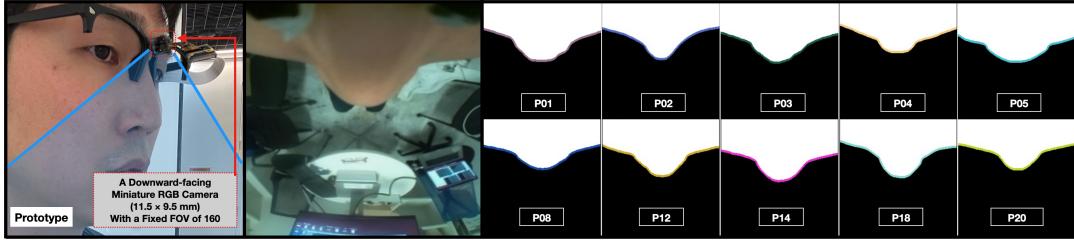


Figure 1: C-Auth: A novel authentication method for smart glasses utilizing egocentric view. By leveraging a downward-facing RGB camera to capture partial face information, C-Auth employs distinct contour lines of the nose and cheeks to establish secure user authentication.

ABSTRACT

C-Auth is a novel authentication method for smart glasses that explores the feasibility of authenticating users using the facial contour lines from the nose and cheeks captured by a down-facing camera in the middle of the glasses. To evaluate the system, we conducted a user study with 20 participants in three sessions on different days. Our system correctly authenticates the target participant versus the other 19 participants (attackers) with a true positive rate of 98.0% (SD: 2.96%) and a false positive rate of 4.97% (2.88 %) across all three days. We conclude by discussing current limitations, challenges, and potential future applications for C-Auth.

CCS CONCEPTS

- Human-centered computing → Ubiquitous and mobile computing systems and tools.

KEYWORDS

Smart glasses, Authentication, Ecocentric view

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ISWC '23, October 08–12, 2023, Cancun, Quintana Roo, Mexico

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-0199-3/23/10...\$15.00

<https://doi.org/10.1145/3594738.3611355>

ACM Reference Format:

Hyunchul Lim, Guilin Hu, Richard Jin, Hao Chen, Ryan Mao, Ruidong Zhang, and Cheng Zhang. 2023. C-Auth: Exploring the Feasibility of Using Egocentric View of Face Contour for User Authentication on Glasses. In *Proceedings of the 2023 International Symposium on Wearable Computers (ISWC '23), October 08–12, 2023, Cancun, Quintana Roo, Mexico*. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3594738.3611355>

1 INTRODUCTION

Smart glasses, or AR glasses, are rapidly advancing and hold significant potential in everyday use including accessing user's private information every day [14]. However, authenticating users on smart glasses is challenging and under-explored. Much of the prior work requires complex hardware customization to enable user authentication on glasses (e.g., Iris [12]) or requires the user to actively input gestures [3, 9, 26] or voice [16].

To address this issue, we present C-Auth, a novel and effortless authentication system on glasses that explores the feasibility of using the contour line of the nose and face as biometric information for user authentication. Our system offers seamless authentication experience by merely wearing the glasses, making it ideal for everyday applications like payments, banking, and accessing private information. It only required a miniature RGB camera placed at the nose bridge facing downward to capture the images of the face. These captured images are processed by a user-independent deep-learning algorithm that extracts the contour lines of the nose and cheeks for authentication. The reason we chose to use contour lines

instead of raw images for authentication is that images are known to be easily distorted by environmental noises such as lighting which can lead to unstable authentication performance. In contrast, the contour line of the nose and cheeks are relatively invariant to environmental factors and are relatively stable for each user. These extracted contour lines are sent to a customized and lightweight algorithm (dynamic time-warping (DTW)) for user authentication.

To evaluate this system, we conducted a non-consecutive three-day user study with 20 participants which is a comparable study size as prior similar works[5, 6, 13, 21, 22]. Using the data collected in this study, we showed that C-Auth can authenticate a target user against the attackers (the other 19 participants) with an average true positive rate (TPR) of 98.0% (SD: 2.96%) and an average false negative rate (FNR) of 4.97% (SD: 2.88%) across all days, indicating its effectiveness in user authentication. The contributions of this paper are as follows:

- An effort-less authentication system to explore the feasibility of using the contour line of the nose and cheek as the biometric information for a seamless authentication approach on smart glasses.
- Implementation of a user-independent segmentation model incorporating a tailored deep learning algorithm for extraction of users' facial contour lines from noisy backgrounds.
- Validation through a user study with 20 participants, affirming the effectiveness of the proposed authentication methods.

2 RELATED WORK

User authentication for smart glasses has been explored through various conventional methods such as voice recognition [16] and tapping/touch and head movement gestures [3, 9, 26]. However, these methods necessitate additional user input such as touch, gestures, or voice, which can be inconvenient or vulnerable to password exposure. Niche sensing methods, based on air and/or bone conduction [21, 22], have also been investigated for glasses authentication, achieving performance levels of 73.5% and 97.0% (true positive rate, TPR) respectively. However, these results may reliable in certain situations such as noisy environments. Another highly reliable approach involves iris recognition [12], attaining 100% TPR by employing infrared cameras to capture eye information. However, this method requires placing the camera in front of the eyes which is inconvenient. In contrast, we propose a novel authentication system that achieves a performance rate of 98.0% by utilizing egocentric images through the downward-facing camera located in the center of the smart glasses, providing a secure, affordable, and convenient authentication solution.

3 RESEARCH IDEA

Our primary research focus is to leverage partial facial images captured by a downward-facing camera positioned on the center of the glasses (See. Figure 1). This camera provides visibility of users' noses and upper cheeks. Despite being incomplete, we propose that the facial contours of the noses and upper cheeks reflect biometric distinctions between individuals and could potentially contain sufficient information to differentiate users. To validate our idea, we address two research questions in this paper: (1) Can C-Auth accurately extract facial contour lines from RGB images

captured by the downward-facing camera in smart glasses? and (2) Can partial facial contour lines serve as distinctive features for user authentication? To answer these questions, we developed a prototype and conducted a user study to assess the performance of our image segmentation and authentication models, which are detailed in the following sections.

4 SYSTEM DESIGN

4.1 Hardware Prototype

Our prototype utilizes a glass frame to accommodate a downward-facing RGB camera (b006605 RGB Arducam). The camera has dimensions of 11.5mm × 9.5mm and a fixed field of view (FOV) of 160°. We affixed the camera to the nose bridge of the frame and connected it to a Raspberry Pi using a flexible cable. The captured images are saved at a resolution of 200 by 200 pixels with a frame rate of 15 FPS. The lower half of the glass frame is removed to capture facial contours, as illustrated in Figure 1, which is similar to many of the everyday wearing smart glasses such as (Google Glass¹ and future smart glasses [15, 25].

4.2 Segmentation Model

In our preliminary experiments, we assessed the performance of existing deep learning segmentation algorithms, such as skin detection [17] and body segmentation [11]². However, these methods did not deliver satisfactory outcomes when applied to our dataset due to the unique camera angles employed in our images, specifically the downward-facing perspective. Thus, we decided to develop our own user-independent segmentation model, called EgoFaceNet, using the Unet architecture and synthetic data.



Figure 2: Synthetic image data

4.2.1 Synthetic Image Data. To develop our segmentation algorithm, a substantial dataset is required for model training. However, collecting the large dataset was challenging, so we employed image synthesis techniques to augment our dataset. Initially, we collected the dataset against a monochromatic background to generate ground-truth masks (see Section 5 for participant data details), which facilitated the segmentation of facial information, as depicted in Figure 2 (a). Subsequently, using these training data masks as a basis, we augmented our original dataset in the following ways: tuning

¹<https://www.google.com/glass/start/>

²<https://blog.tensorflow.org/2022/01/body-segmentation.html>

the skin to different colors (b), replacing the background with other background images³ (c), and varying the light conditions across different intensities (d). These augmentation techniques increased the diversity and size of our dataset, benefiting the segmentation process for images with unseen noise backgrounds.

4.2.2 Pipeline and Training. Our segmentation model uses a UNet [2, 8, 19, 24] architecture, taking input images from our prototype to generate a segmented mask with white pixels representing skin and black pixels representing the background. It takes input images from our prototype, producing a segmented mask with white pixels representing skin and black pixels representing the background. The Resnet18 encoder[7] extracts key features while mitigating gradient vanishing and information loss. The model outputs logits transformed into probabilities using a sigmoid layer, with a threshold of 0.5 for pixel classification to determine whether each pixel belongs to the background or the skin. Using the data collected from study (detailed in the next section), we employed the leave-one-participant-out evaluation, training individual models for each participant while excluding their own data and incorporating data from all other participants. During training, images underwent skin color changes (50%), background replacement with random noise (10%) or egocentric images (90%), and variations in light conditions. The models were trained with 40 epochs, Adam optimizer, 0.001 learning rate, and binary cross-entropy loss [20, 27].

4.3 Authentication Algorithm

We utilize a light-weighted Dynamic Time Warping (DTW) [10] as our authentication algorithm, which only needs a few samples from each user for authentication.

4.3.1 Adjusting and Extracting Facial Contour Lines. For each time a user mounts the glasses on, the glasses could slant at a slightly different angle, causing misalignment of the facial contour lines as shown in Figure 3 (a). For contour alignment, the segmented images would be rotated to adjust both sides of the nose contour to the same level (b). After adjustment (c), we extract the facial contour lines using a binary thresholding technique ⁴ using OpenCV (d).



Figure 3: Adjusting and extracting facial contour lines

4.3.2 Reference Set and Dynamic Time Warping (DTW) Algorithms for Authentication. To authenticate users, our system needs to first collect a few sample facial contour lines for each user. Then we calculate a threshold by calculating the average distances and standard deviation using the DTW algorithm [10] among the reference contour lines. The threshold was set as the mean plus 1.5 times the standard deviation. When an unknown user's contour line was used for authentication, this algorithm will compare the distance between this unknown contour line and the contour lines in the user's

³<https://ego4d-data.org/>

⁴https://docs.opencv.org/4.x/d7/d4d/tutorial_py_thresholding.html

reference set. If the minimum distance is below the aforementioned threshold, it indicates that this unknown user is authenticated as the target user. Otherwise, it will be rejected by the system as a different user.

5 USER STUDY

We conducted a user study involving 20 participants (mean age: 22, SD: 4.5), including 10 males, with an age range of 19 to 29, to evaluate our segmentation and authentication models. They visited the lab on three non-consecutive days while wearing different sets of clothing. The skin colors of them range from white to light tones.

5.1 Experiment Setup

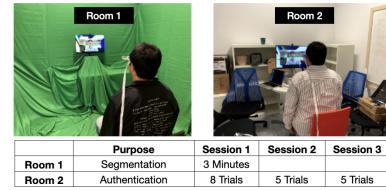


Figure 4: Study setting and procedure

The data collection process involved two separate rooms, each serving a specific purpose. Room 1 was used to collect training data for the segmentation model, utilizing a green chroma key background for clear facial area extraction to serve as the ground truth of the segmentation model training dataset. In contrast, Room 2 replicated a real-world environment with chairs, computers, and monitors, providing a noisy background to evaluate both segmentation and authentication models. Figure 4 illustrates these setups.

5.2 Procedure

To evaluate the long-term performance of our system, participants were involved in three separate sessions conducted on different days. Session 1 consisted of two parts conducted in different rooms. In the first part, participants wore the prototype and performed natural face movements for 3 minutes in Room 1. To account for potential shifts in the device position, participants remounted the glass frame by completely removing and wearing it back 10 times. Approximately 2,700 training images were recorded at a frame rate of 15 FPS per participant. These images served as the training dataset for the segmentation model. The second part, within Session 1, took place in Room 2, simulating a noisy office environment (Figure 4). Participants remounted the device every 10 seconds for a total of eight times, with three trials forming the reference set and the remaining five trials as the testing set. In Sessions 2 and 3, participants repeated the process of remounting the device five times in Room 2, serving as testing sets lasting for 10 seconds each. Each user contributed three reference samples and 15 testing samples (five per session) throughout the three-day study (Table 4).

6 EVALUATION

6.1 Segmentation Performance

The evaluation of our segmentation model was conducted using the images collected in Room 2, which featured a noisy background.

These images were not used in training the model. To assess the accuracy of model's predictions, we employed the Intersection over Union (IoU) metric, which is widely used in binary segmentation tasks [18, 23]. For each session, we randomly selected five images per participant, resulting in a total of 300 images. We manually segmented these images by hand-drawing the facial contour lines to create the ground-truth masks. Subsequently, we evaluated the performance of all our user-independent models. The results showed that our user-independent model achieved reliable segmentation of user facial parts, with an average IoU of 97.11% (0.47%) across all sessions, enabling accurate delineation of facial contour lines.

6.2 Authentication Performance

6.2.1 Evaluation Metrics. We evaluated the performance of the system to authenticate each user against the other 19 users using the contour line data collected in the studies. For evaluation, we assessed True Positive Rate (TPR) for login success and False Positive Rate (FPR) for attack success [4]. Each target participant underwent 15 login trials for TPR and faced 285 attack trials from the other 19 participants (attackers) for FPR evaluation.

6.2.2 Login Success: True Positive Rate. On average, our system achieved a TPR of 81.0% (SD: 16.93%) across all three sessions. In the first session, the TPR was high at 94.0% (SD: 11.42%), but some participants experienced a decline in TPR during the second and third sessions, with an average of 76.0% (SD: 29.45%) and 73.0% (SD: 29.2%), respectively (Figure 5). We attribute this decrease to the limited size of our reference set. In the second session, participants' data showed more variation due to changes in how they wore the device, resulting in slightly different captured facial contour lines compared to the first session (participants P01, P03, P12, and P14). To address this issue, we expanded the reference set by adding two facial contour lines from the first and second trials in session 2. After this adjustment, the system showed promising results, with an average TPR of 98.0% (SD: 2.96%) across all three sessions. The TPR performance improved to 97.0% (SD: 7.33%), 100.0% (SD: 0.0%), and 98.0% (SD: 6.16%) in the respective sessions. We discuss the form factor and remounting issues in the discussion section.

6.2.3 Attack Success: False Positive Rate. The FPR of our system achieved an average of 4.82% (SD: 4.37%) across three sessions, with rates of 4.47% (SD: 4.16%), 4.21% (SD: 4.11%), and 5.53% (SD: 4.39%) in each session, respectively. While the average FPR appears low (with all other participants achieving less than 10.0%), there were notable higher FPRs for participants P01 and P03, at 14.74% (SD: 3.16%) and 10.88% (SD: 0.66%), respectively. We observed that these participants had relatively similar facial contour lines, making it easier for them to attack each other. After adding two reference contour lines from Session 2, although the FPR slightly increased to 4.97% (2.88%) across the three sessions, we observed that the FPR for both participants decreased from 14.74% to 11.23% and from 10.88% to 9.47%, respectively. Adding more references helped set a tighter threshold for distinguishing between genuine users and attacks, thus improving the system's performance in detecting attacks. However, our model may still face challenges when dealing with individuals who have similar facial contour lines. We discuss potential strategies to handle this issue in the discussion section.

	# of Ref.	All	Session 1	Session 2	Session 3
TPR	3	81.0 (16.93)	94.0 (11.42)	76.0 (29.45)	73.0 (29.22)
	5	98.0 (2.96)	97.0 (7.33)	100.0 (0.00)	98.0 (6.16)
FPR	3	4.82 (4.20)	4.74 (4.61)	4.21 (4.11)	5.53 (4.39)
	5	4.97 (2.88)	4.68 (4.63)	4.32 (2.41)	5.90 (2.30)

Figure 5: Authentication Performance

7 DISCUSSION AND FUTURE WORK

Our paper aims to show the feasibility of this novel sensing approach, which is a starting point for future work. We evaluated the performance of 20 participants in a three-day experiment.

The study size is comparable (if not better) to similar prior work [5, 6, 13, 21, 22], which evaluated 10 to 23 participants. Furthermore, we evaluated C-Auth across three different days, which were not tested at all in many of these works [12, 13, 21, 22].

The performance is promising. However, there are many methods that can further improve the performance such as FPR. For instance, increasing the reference set through continuous and passive collection from users in various situations can help accommodate facial changes over time. Integrating our system with other authentication methods like gestures, touch, or voice and context information (e.g., Bluetooth pairing) can further help reduce the FPR. Furthermore, our system only used the contour lines for authentication to ensure the performance is robust against noises (e.g., lighting). However, using other biometric information, such as skin color and texture, can further enhance performance with advancements in CV algorithms and cameras.

Improving stability after remounting the device can help secure consistent facial contour lines. One solution is to add nose pads and design balanced glass legs, similar to commercial products like Epson's Moverio [1]. Our prototype can also be adapted to the current full-framed glasses form factor, which may partially cover some contour lines on the sides but still capture the middle of the facial contour lines, potentially contributing to unique patterns based on face size and nose height. Exploring different sizes and types of glasses is our future plan for system enhancement.

We believe our current study successfully shows the feasibility of our sensing method as the proof-of-concept. Our next step is to collect data from a larger and more diverse population (e.g., age) which will be used to evaluate and optimize our system for real-world deployment.

8 CONCLUSION

C-Auth allows smart glasses to authenticate users using contour lines of the nose and cheek captured by a downward-facing camera on glasses. Through a user study with 20 participants, it showed promising authentication performance with a TPR of 98%.

ACKNOWLEDGMENTS

We thank the participants for participating in the user study conducted on three different days. This material is based upon work supported by the National Science Foundation under Grant No.2239569.

REFERENCES

[1] 2023. Moverio Smart Glasses. <https://epson.com/moverio-smart-glasses-augmented-reality-devices-headsets>

[2] Hu Cao, Yueyue Wang, Joy Chen, Dongsheng Jiang, Xiaopeng Zhang, Qi Tian, and Manning Wang. 2021. Swin-unet: Unet-like pure transformer for medical image segmentation. *arXiv preprint arXiv:2105.05537* (2021).

[3] Jagmohan Chauhan, Hassan Jameel Asghar, Anirban Mahanti, and Mohamed Ali Kaafar. 2016. Gesture-Based Continuous Authentication for Wearable Devices: The Smart Glasses Use Case. In *Applied Cryptography and Network Security*, Mark Manulis, Ahmad-Reza Sadeghi, and Steve Schneider (Eds.). Springer International Publishing, Cham, 648–665.

[4] Alexander De Luca, Alina Hang, Frederik Brudy, Christian Lindner, and Heinrich Hussmann. 2012. Touch Me Once and i Know It's You! Implicit Authentication Based on Touch Screen Patterns. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (Austin, Texas, USA) (CHI '12). Association for Computing Machinery, New York, NY, USA, 987–996. <https://doi.org/10.1145/2207676.2208544>

[5] Yang Gao, Yincheng Jin, Jagmohan Chauhan, Seokmin Choi, Jiyang Li, and Zhanpeng Jin. 2021. Voice in ear: Spoofing-resistant and passphrase-independent body sound authentication. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 5, 1 (2021), 1–25.

[6] Yang Gao, Wei Wang, Vir V Phoha, Wei Sun, and Zhanpeng Jin. 2019. EarEcho: Using ear canal echo for wearable authentication. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 3, 3 (2019), 1–24.

[7] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 770–778.

[8] Huimin Huang, Lanfen Lin, Ruofeng Tong, Hongjie Hu, Qiaowei Zhang, Yutaro Iwamoto, Xianhua Han, Yen-Wei Chen, and Jian Wu. 2020. Unet 3+: A full-scale connected unet for medical image segmentation. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 1055–1059.

[9] MD. Rasel Islam, Doyoung Lee, Liza Suraiya Jahan, and Ian Oakley. 2018. Glass-Pass: Tapping Gestures to Unlock Smart Glasses. In *Proceedings of the 9th Augmented Human International Conference* (Seoul, Republic of Korea) (AH '18). Association for Computing Machinery, New York, NY, USA, Article 16, 8 pages. <https://doi.org/10.1145/3174910.3174936>

[10] Sandeep S Jambhale and Arti Khaparde. 2014. Gesture recognition using DTW & piecewise DTW. In *2014 International Conference on Electronics and Communication Systems (ICECS)*. IEEE, 1–5.

[11] Peike Li, Yunqiu Xu, Yunchao Wei, and Yi Yang. 2020. Self-correction for human parsing. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 44, 6 (2020), 3260–3271.

[12] Yung-Hui Li and Po-Jen Huang. 2017. An accurate and efficient user authentication mechanism on SMART glasses based on Iris recognition. *Mobile Information Systems* 2017 (2017), 1–14. <https://doi.org/10.1155/2017/1281020>

[13] Chen Liang, Chun Yu, Xiaoying Wei, Xuhai Xu, Yongquan Hu, Yuntao Wang, and Yuanchun Shi. 2021. Auth+ track: Enabling authentication free interaction on smartphone by continuous user tracking. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*. 1–16.

[14] Jins Meme. 2022. The glasses for self-care for your mind and body: Jins Meme. <https://jinsmeme.com/en/>

[15] Ryoma Ogawa, Kyosuke Futami, and Kazuya Murao. 2021. NasalBreathInput: A Hands-Free Input Method by Nasal Breath Gestures using a Glasses Type Device. In *The 23rd International Conference on Information Integration and Web Intelligence*. 620–624.

[16] Ge Peng, Gang Zhou, David T. Nguyen, Xin Qi, Qing Yang, and Shuangquan Wang. 2017. Continuous Authentication With Touch Behavioral Biometrics and Voice on Wearable Glasses. *IEEE Transactions on Human-Machine Systems* 47, 3 (2017), 404–416. <https://doi.org/10.1109/THMS.2016.2623562>

[17] Son Lam Phung, Abdesselam Bouzerdoum, and Douglas Chai. 2005. Skin segmentation using color pixel classification: analysis and comparison. *IEEE transactions on pattern analysis and machine intelligence* 27, 1 (2005), 148–154.

[18] Hamid Rezatofighi, Nathan Tsoi, JunYoung Gwak, Amir Sadeghian, Ian Reid, and Silvio Savarese. 2019. Generalized intersection over union: A metric and a loss for bounding box regression. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 658–666.

[19] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. 2015. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*. Springer, 234–241.

[20] Usha Ruby and Vamsidhar Yendapalli. 2020. Binary cross entropy with deep learning technique for image classification. *Int. J. Adv. Trends Comput. Sci. Eng* 9, 10 (2020).

[21] Stefan Schneegass, Youssef Oualil, and Andreas Bulling. 2016. SkullConduct: Biometric User Identification on Eyewear Computers Using Bone Conduction Through the Skull. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (San Jose, California, USA) (CHI '16). Association for Computing Machinery, New York, NY, USA, 1379–1384. <https://doi.org/10.1145/2858036.2858152>

[22] Shinichi Sugahara and Mikio Mori. 2021. Feasibility of using Air-conducted and Bone-conducted Sounds Transmitted through Eyeglasses Frames for User Authentication. In *2021 International Conference on Electrical, Communication, and Computer Engineering (ICECCE)*. 1–4. <https://doi.org/10.1109/ICECCE52056.2021.9514184>

[23] Floris van Beers. 2018. *Using intersection over union loss to improve binary image segmentation*. Ph.D. Dissertation.

[24] Yu Weng, Tianbao Zhou, Yujie Li, and Xiaoyu Qiu. 2019. Nas-unet: Neural architecture search for medical image segmentation. *IEEE Access* 7 (2019), 44247–44257.

[25] Zihan Yan, Yufei Wu, Yang Zhang, and Xiang'Anthony' Chen. 2022. EmoGlass: an End-to-End AI-Enabled Wearable Platform for Enhancing Self-Awareness of Emotional Health. In *CHI Conference on Human Factors in Computing Systems*. 1–19.

[26] Shanhe Yi, Zhengrui Qin, Ed Novak, Yafeng Yin, and Qun Li. 2016. GlassGesture: Exploring head gesture interface of smart glasses. In *IEEE INFOCOM 2016 - The 35th Annual IEEE International Conference on Computer Communications*. 1–9. <https://doi.org/10.1109/INFOCOM.2016.7524542>

[27] Zhilu Zhang and Mert Sabuncu. 2018. Generalized cross entropy loss for training deep neural networks with noisy labels. *Advances in neural information processing systems* 31 (2018).