Fatigue Detection on Face Image Using FaceNet Algorithm and K-Nearest Neighbor Classifier

Faisal Dharma Adhinata[1] * ID **, Diovianto Putra Rakhmadani[2]** ID **, Danur Wijayanto[3]** ID

[1][2] *Department of Software Engineering, Faculty of Informatics, Institut Teknologi Telkom Purwokerto, Indonesia*
Jalan D.I Panjaitan 128, 53147, Purwokerto
[1]faisal@ittelkom-pwt.ac.id, [2]diovianto@ittelkom-pwt.ac.id

[3]*Information Technology, Faculty of Science and Technology, Universitas 'Aisyiyah Yogyakarta, Indonesia*
Jalan Siliwangi 63, 55133, Yogyakarta
[3]danurwijayanto@unisayogya.ac.id

Abstract

Background: The COVID-19 pandemic has made people spend more time on online meetings more than ever. The prolonged time looking at the monitor may cause fatigue, which can subsequently impact the mental and physical health. A fatigue detection system is needed to monitor the Internet users well-being. Previous research related to the fatigue detection system used a fuzzy system, but the accuracy was below 85%. In this research, machine learning is used to improve accuracy.
Objective: This research examines the combination of the FaceNet algorithm with either k-nearest neighbor (K-NN) or multiclass support vector machine (SVM) to improve the accuracy.
Methods: In this study, we used the UTA-RLDD dataset. The features used for fatigue detection come from the face, so the dataset is segmented using the Haar Cascades method, which is then resized. The feature extraction process uses FaceNet's pre-trained algorithm. The extracted features are classified into three classes—focused, unfocused, and fatigue—using the K-NN or multiclass SVM method.
Results: The combination between the FaceNet algorithm and K-NN, with a value of $K = 1$ resulted in a better accuracy than the FaceNet algorithm with multiclass SVM with the polynomial kernel (at 94.68% and 89.87% respectively). The processing speed of both combinations of methods has allowed for real-time data processing.
Conclusion: This research provides an overview of methods for early fatigue detection while working at the computer so that we can limit staring at the computer screen too long and switch places to maintain the health of our eyes.

Keywords: Artificial Intelligence, FaceNet Algorithm, Fatigue Detection, K-Nearest Neighbor, Multiclass Support Vector Machine

Article history: Received 28 December 2020, first decision 2 February 2021, accepted 22 February 2021, available online 28 April 2021

I. INTRODUCTION

The COVID-19 pandemic requires office workers to work from home (WFH), which means more exposure to computer or laptop screens. Teachers or lecturers also have to conduct teaching and learning activities virtually. Looking at the computer or laptop screen in a prolonged period may strain the eyes and make people feel tired. The average adult maximum screen time is two hours [1]. After that, it is advisable to move to distant viewing or to look at green objects. Staring at a computer screen for a long period of time can affect brain development [2] and may cause visual impairment [3]. Some signs of fatigue due to spending too much time in front of a computer are face and eye fatigue or the computer vision syndrome (CVS) [4], which may start to appear after two hours. Based on a survey [4], 75% of computer users who work at the computer for too long often complain of CVS symptoms, such as eyestrain, blurred vision, and dry eyes. Therefore, a face-based fatigue detection system is needed as an early warning of fatigue in viewing the screen [5][6][7]. This system can also help doctors in early detection analysis of CVS symptoms to take action on patients.

Fatigue can be seen from the facial expressions [8] and so the features can be used to detect fatigue when people are sitting in front of a screen. Recently, researchers have developed a fatigue detection system by using facial data and the findings are used, for example, to avoid accidents among drivers [9][10][11][12]. Kong and Li [13] developed a fatigue detection system based on facial data by using a fuzzy system. The research examined the face data from

---

* Corresponding author

online learning classes. The performance accuracy of the fatigue detection system developed was still below 85%. This indicates that the system is not effective enough. For example, the face or eyes may look tired, but the system detects that the person is still focused. Therefore, we need a method that can improve the accuracy in the system. The novelty of this research is to replace the feature extraction methods with machine learning techniques for data classification in order to increase the accuracy of the system.

The key is on the feature extraction and the classification stages. A common feature extraction algorithm is the FaceNet pre-trained models, which is commonly used, for example, to detect a face presence by searching for the facial features [14][15]. The use of the FaceNet algorithm for this such a system results in an accuracy of above 90%. The FaceNet algorithm's ability to detect facial data has been trained by Google researchers, and this has been tested for verification and recognition of people's identities with an accuracy of above 95% [16].

Classification is also an important stage to groop the features according to their class. Artificial intelligence (AI) applies supervised machine learning techniques to labeled data. K-NN algorithm [17] and SVM [18][19] can classify identities based on the facial data. These systems result in an accuracy of above 90%. In this research, we propose a combination between the FaceNet algorithm for facial feature extraction and the K-NN or multiclass SVM classification methods for the fatigue detection system. This study uses data that consist of three classes, namely focused, unfocused, and fatigue.

This paper is organized as follows. Section 2 describes the fatigue detection system's architecture, namely data acquisition, face detection, data pre-processing, feature extraction, and classification. Section 3 presents the results of fatigue detection by examining the results of K-NN or multiclass SVM methods. Section 4 presents a discussion of this system regarding the improved results compared to the previous research. Section 5 concludes the discussion and provides suggestions for future work.

II. Methods

The fatigue detection system starts with the dataset acquisition consisting of three classes of person's expression, namely focused, unfocused, and fatigue. Haar Cascades method [20] is used because this is suitable to detect facial features. Next, the face image is pre-processed by resizing it into 160 x 160 before being extracted by using the FaceNet algorithm. The features obtained are classified accordingly by using the K-NN and multiclass SVM method. The best results will be the conclusion of our research. Fig. 1 shows the fatigue detection architecture.
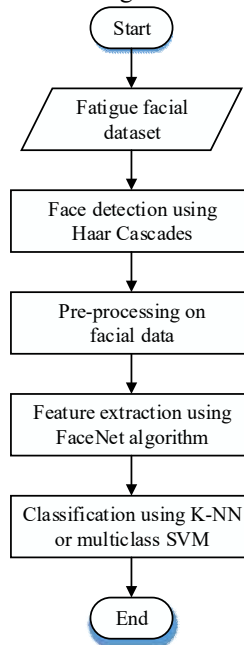

Fig. 1 The architecture of fatigue detection system

A. Fatigue Facial Dataset

This research uses the UTA-RLDD dataset [21]. Actors in the dataset were volunteers, consisting of undergraduate or graduate students, and staff members. Participants were aged over 18 years and consisted of 51 men

and nine women from various ethnicities. The dataset consists of three classes, namely focused, unfocused, and fatigue. The image from the dataset is Full HD (1080x1920) resolution. All image datasets contain photos of people sitting in front of the computer. The training data were 1000 training data and the validation data were 100 for each class. Fig. 2 shows an example of the UTA-RLDD dataset.
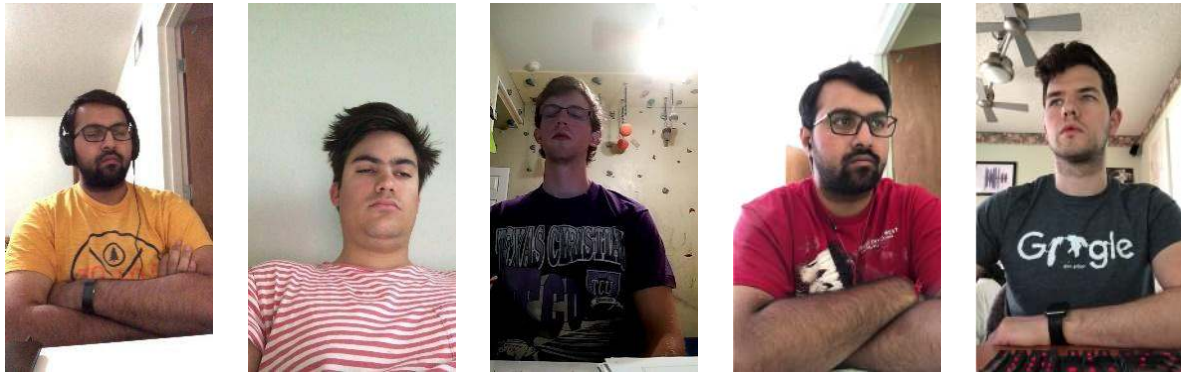

Fig. 2 The example of UTA-RLDD dataset

B. Face Detection using Haar Cascades

Since fatigue detection systems use facial features, the Haar Cascades is used. This method uses a machine learning approach with a cascade function that has been trained with many positive and negative datasets [22]. Fig. 3 shows an example of face detection results.
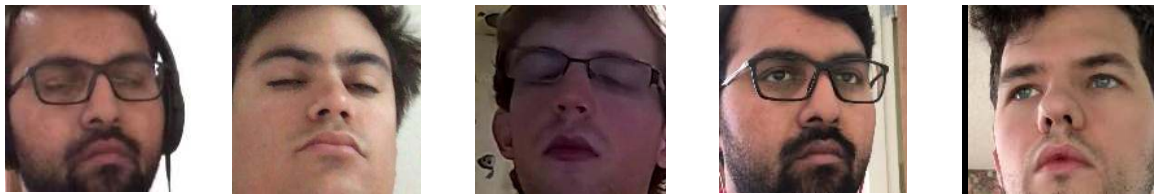

Fig. 3 The example of face detection

C. Pre-processing Facial Data

The pre-processing stage aims to prepare data before extraction. In this study, facial data was resized into a resolution of 160 x 160. We made a function to convert all images into a uniform size format. The results of face detection using Haar Cascades vary in resolution sizes (more than 160x 160), as shown in Fig. 3, so we needed to make them uniform. The square size of 160 x 160 was chosen because the feature-extraction process using the FaceNet algorithm requires conversion to that size. Fig. 4 shows an example of the results of pre-processing images into uniform size.

Fig. 4 The example result of pre-processing on facial data

D. Feature Extraction using FaceNet Algorithm

In the FaceNet model, there is a triplet loss function, a set of three images [16] which consists of an anchor image, a positive image (an identical image to the anchor image), and a negative image (image that is different from the anchor image). Fig. 5 illustrates an example of a triplet.
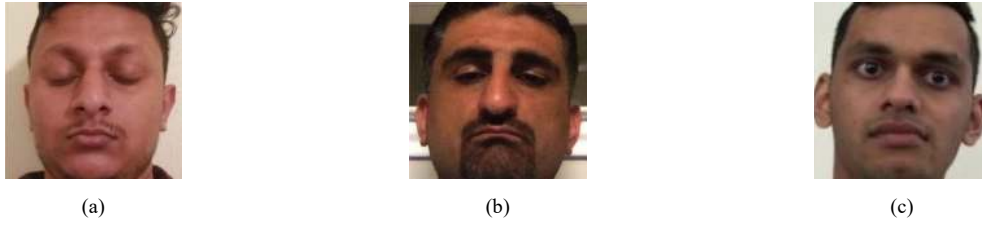
| (a) | (b) | (c) |

Fig. 5 The example of triplet in FaceNet algorithm (a) Fatigue face (anchor); (b) Fatigue face (positive); (c) Focus face (negative)

A triplet set is said to be valid if it has the anchor, positive (same class), and negative (different class) images. The triplet's objective function is the distance between the embedding anchors with a positive smaller than the negative anchor distance. An illustration of learning triplet loss is shown in Fig. 6.
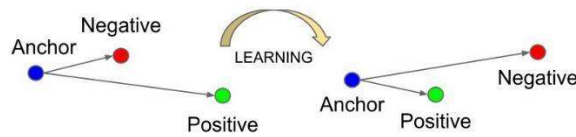

Fig. 6 The triplet loss [16]

The model must be trained to produce a distance of an anchor that is closer to the positive image than the negative image. This research uses the FaceNet algorithm that uses RGB images with the size of 160 x160 with three channel colors. The result is a face embedding vector with 128 dimensions.

E. Classification using K-NN or Multiclass SVM

In this research, two supervised learning methods are studied: the K-NN and multiclass SVM.

1) K-Nearest Neighbor

The K-Nearest Neighbor (K-NN) method classifies objects based on the learning data that are the closest to the object. K-NN is a supervised classification method where input data are labeled before training. This method is often used in the fields of pattern recognition [23] and image processing [24]. Learning data will be projected into multiple dimensional spaces, with each dimension representing a feature of the data. This space is divided into sections based on the learning data classification, namely focused, unfocused, and fatigue. Fig. 7 shows an illustration of the K-NN method.
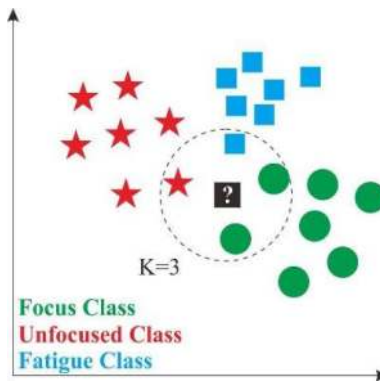

Fig. 7 K-Nearest Neighbor illustration

The dots this space are marked as a focused class if the focused class is the most common classification for the K-Nearest Neighbor at that point. Near or far, neighbors are usually calculated based on the Euclidean distance with the (1).

$$dist(X,Y) = \sqrt{\sum_{i=1}^{n} (X_i - Y_i)^2} \qquad (1)$$

where,

X_i : The training data
Y_i : The testing data
n : The number of training data

2) *Multiclass SVM*

Multiclass SVM is a method for dividing classes of more than two. The technique used in this multiclass SVM is one-vs-all, or often called one-vs-rest, which is a method of training one classifier in each class, with one class sample as a positive class and all other samples as a negative sample [25]. There are three classes used in this system, namely focused, unfocused, and fatigue. Fig. 8 represents the implementation of multiclass SVM in this research. Focused class is a positive class, and this is compared with unfocused class and fatigue class. Then the unfocused class data is compared with focused class and fatigue class. Likewise, the fatigue class is compared with the focused class and unfocused class.
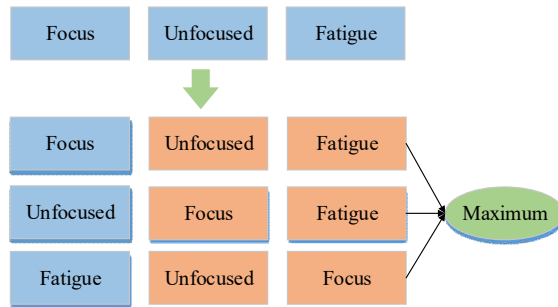

Fig. 8 The implementation of multiclass SVM

SVM is a linear classifier so it can be separated linearly. However, SVM can be developed to work on non-linear problems by incorporating the kernel concept in a high-dimensional workspace. A hyperplane will be sought in high-dimensional space that can maximize the distance (margin) between data classes. Multiclass SVM in this study uses (2). In the prediction of multiclass SVM, a maximum value is calculated from each class comparison [26]. Fig. 9 is an illustration of multiclass SVM in this research.

$$[W_{11}\ W_{12}\ W_{13}\ ..\ W_{1D}\ W_{21}\ W_{22}\ W_{23}\ ..\ W_{2D}\ ....\ W_{K1}\ W_{K2}\ W_{K3}\ ..\ W_{KD}\ ][X_1\ X_2\ X_3\ ..\ X_D\ ] = [s_1\ s_2\ s_3\ ..\ s_K\ ]\quad (2)$$

where:
W_{KD} : Feature vector of the generated model
X_D : Feature vector of testing image
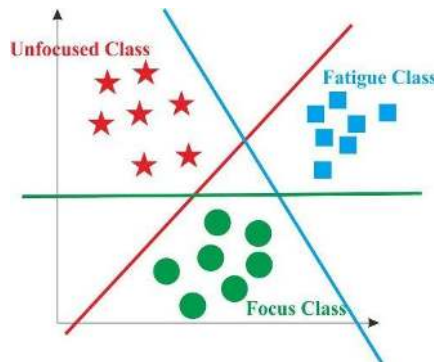s_K : Class prediction results

Fig. 9 Multiclass Support Vector Machine illustration

III. RESULTS

The fatigue detection system is run on a computer with an Intel Core i3-9100F CPU @ 3.60 GHz, 8192 MB of RAM, and a 64-bit Windows 10 Pro operating system. This research tested the accuracy of the model building after the model is formed. Then, new data testing will be carried out to test the speed and accuracy.

A. Learning using K-Nearest Neighbor

The feature classification experiment to detect fatigue uses the Euclidean distance to calculate the closest distance. Calculation of the closest distance is both for testing or validation data against all training data. This experiment's K value uses odd values 1, 3, 5, 7, and 9. The odd K value is used to avoid the appearance of the same number of distances in the classification process. Fig. 10 shows the experimental results of variations in the K value on K-NN.
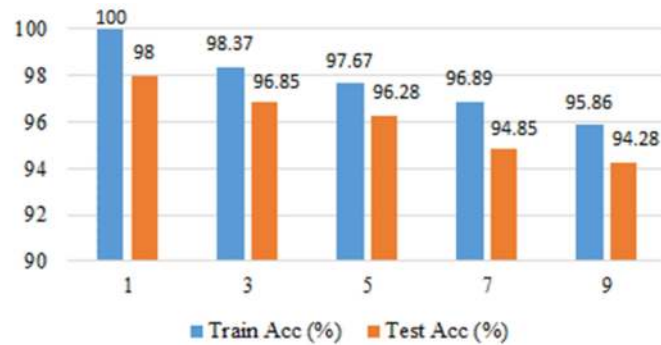

Fig. 10 The result of the experiment on various K value on K-NN toward accuracy

Based on Fig. 10, the higher the K value has the effect of a decreasing accuracy. The decrease in accuracy is because the features in each class are close from each other. The proximity of this feature makes it difficult for the system to differentiate between classes and detect fatigue. In this experiment, the best accuracy results were obtained at $K = 1$ because there is only one closest neighbor to determine the class. From the results of the accuracy of the variation in the K value, it can be seen that the accuracy value tends to decrease with the increasing the value of K. At $K = 1$, the training accuracy value is 100%, then at $K = 9$ the training accuracy value is 95.86%. Similarly, the testing accuracy value tends to decrease. The decrease in this accuracy value is due to the distribution of image data features that are closely together between classes. A higher K value can reduce noise because there are several features in the same class (not just one feature) to cover the noise feature. However, the increased K value causes the boundaries between classes to become increasingly blurred, especially in image data with very complex features. Therefore, the K-NN method with $K = 1$ is used for experiments on new testing data.

B. Learning using Multiclass SVM

The three-class classification uses the one-vs-all SVM method. The SVM kernel is used to separate data between classes. This research studies three kernels, namely radial basis function (RBF), polynomial, and linear. The role of the kernel in SVM greatly affects the accuracy value. Fig. 11 shows the results of the modelling accuracy of with SVM kernel variations.
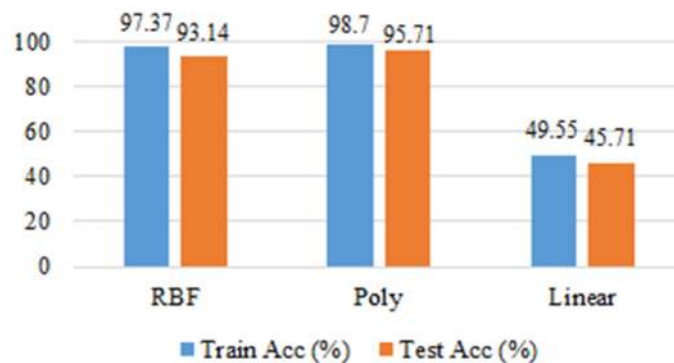
Fig. 11 The result of the experiment on various SVM kernel toward accuracy

Fig. 11 shows that a multiclass SVM with a linear kernel produces the worst accuracy compared to the RBF and polynomial kernels. The linear kernel precisely divides the three classes of fatigue detection systems with a straight line. There are many image features and in the dimensional space, not all of them are clustered in their respective classes. Therefore, straight lines cannot separate them. Meanwhile, the RBF kernel and polynomial separation do not use a straight line. In this experiment, the best result is obtained using the polynomial kernel. Therefore, multiclass SVM using the polynomial kernel is used for experiments on new testing data.

IV. DISCUSSION

AI helps humans solve everyday life problems. One of the COVID-19 pandemic's effects is that the teaching and learning activities are required to be online. Staring at the screen for too long in front of the screen may cause fatigue. People who are too long in front of the computer can trigger Computer Vision Syndrome (CVS), which is caused by visual fatigue. People can prevent CVS symptoms by creating an early fatigue detection system to limit staring at computer screens. We propose fatigue detection systems based on facial expressions. One application development uses a fuzzy system approach, but the accuracy results are less than 85%.

Testing using new data aims to test the accuracy and the processing time in fatigue detection of one image. The latest data used in this test amounted to 100 facial images in each class. Table 1 shows the accuracy and processing time of fatigue detection testing in the multiclass SVM and K-NN classification methods.

TABLE 1
THE EXPERIMENT RESULT OF ACCURACY AND PROCESSING TIME ON MATCHING IMAGE

Methods	Accuracy (%)		Processing Time (s)
	RLDD	**RLDD + New data**	
FaceNet + Multiclass SVM	89.87	89.88	0.058
FaceNet + K-NN	94.66	94.68	0.061

The FaceNet algorithm depends on the similarity of facial features in each class [16]. Therefore, two faces with focused expressions have more similarity than two faces with focused and fatigue expressions. The Euclidean distance at the anchor between the two faces with focused expression is smaller (positive) than the face with a fatigue expression (negative). Thus, the embeddings of all faces from the same class are grouped. The facial embeddings for fatigue detection are calculated to produce three clusters based on facial expressions (focused, unfocused, and fatigue). Fig. 11 shows an example of the results of matching the fatigue detection on the image.

Table 1 shows that the combination of FaceNet and K-NN algorithms results in better accuracy than using multiclass SVM as a classifier. The speed of using the K-NN method as a classifier is not much different from the multiclass SVM. The difference in speed is due to the K-NN method comparing the new data features with all the training data features. However, the combination of FaceNet and K-NN with a value of $K = 1$ has an accuracy value of 94.68%, which is much better than the multiclass SVM. This accuracy result is obtained from the combination of new data from video recordings with data on the UTA-RLDD reference [21]. The addition of new data consists of 10 facial images in each focused, unfocused, and fatigue class. The addition of this data does not change the accuracy value significantly either in the combination of FaceNet and multiclass SVM or FaceNet and K-NN. These results are consistent with the accuracy of training and testing as shown in Fig. 9 and 10, where the K-NN method with a value of $K = 1$ has a better accuracy than the SVM method with a polynomial kernel. The accuracy of the fatigue detection system using the K-NN method is better than the multiclass SVM because the classification on the K-NN uses the closest neighbors that do not need to separate the data into three classes. Whereas in the multiclass SVM method, features need to be separated according to the kernel and this can cause certain class features to be positioned incorrectly. This subsequently causes errors in the classification results. This result is better than the previous study [13], which resulted in an accuracy of below 85%. Moreover, in the combination of our proposed methods, the processing speed of one image for fatigue detection is only 0.061 seconds. Therefore, the use of a combination of the FaceNet algorithm with K-NN can be used for real-time video data processing. In this study, there are still weaknesses in face detection. The detected face still has a background that is not the face area, as shown in Fig. 12. Future research can use dynamic face detection so that only the face area is detected.
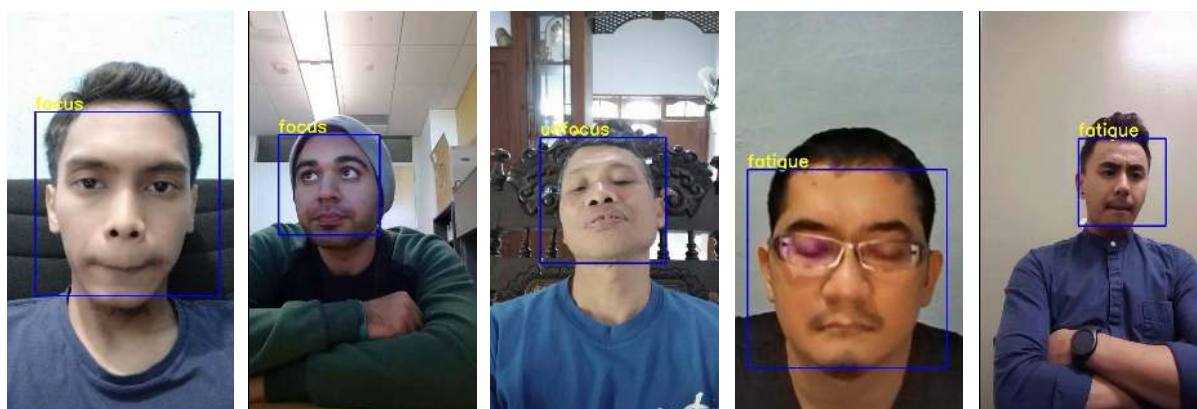
Fig. 12 Example of fatigue detection results

V. CONCLUSIONS

In this paper, the fatigue detection system uses the FaceNet algorithm for feature extraction and examines two classifiers (K-NN and multiclass SVM) to get the best accuracy results. The combination of FaceNet and K-NN algorithms with a value of $K = 1$ produces the best accuracy of 94.68%, and the speed in processing one image of 0.061 seconds. These results allow for real-time data processing using video data. Meanwhile, the method we propose still has weaknesses in face detection, that is the captured background area. Future work can make modifications in the face detection so that the features obtained are only from the face area.

Author Contributions: *Faisal Dharma Adhinata*: Conceptualization, Methodology, Software, Writing - Original Draft, Supervision. *Diovianto Putra Rakhmadani*: Investigation, Data Curation, Resources, Funding acquisition, Project administration. *Danur Wijayanto*: Writing - Review & Editing, Validation, Visualization, Formal Analysis.

Conflicts of Interest: The authors declare no conflict of interest.

REFERENCES

[1] T. L. Wang and D. A. Vella-Brodrick, "Examining Screen Time, Screen Use Experiences, and Well-Being in Adults," *Soc. Netw.*, vol. 07, no. 01, pp. 32–44, 2018, doi: 10.4236/sn.2018.71003.

[2] E. Neophytou, L. A. Manwell, and R. Eikelboom, "Effects of Excessive Screen Time on Neurodevelopment, Learning, Memory, Mental Health, and Neurodegeneration: a Scoping Review," *Int. J. Ment. Health Addict.*, no. December, 2019, doi: 10.1007/s11469-019-00182-2.f

[3] S. Poudel, "A research report about effect of display gadgets on eyesight quality (computer vision syndrome) Of M.Sc.(CSIT) students in Tribhuvan University," *Int. J. Sci. Eng. Res.*, vol. 9, no. 8, 2018, doi: 10.14299/ijser.2018.99.

[4] Z. Yan, L. Hu, H. Chen, and F. Lu, "Computer Vision Syndrome: A widely spreading but largely unknown epidemic among computer users," *Comput. Human Behav.*, vol. 24, no. 5, pp. 2026–2042, 2008, doi: 10.1016/j.chb.2007.09.004.

[5] S. H. Meng, S. B. Hu, A. C. Huang, T. J. Huang, Z. Xie, and C. Jian, "Research on Eye Detection and Fatigue Early Warning Technologies," in *Advances in Intelligent Systems and Computing*, 2018, vol. 682, pp. 3–9, doi: 10.1007/978-3-319-68527-4_1.

[6] M. Toniolo-Barrios and L. Pitt, "Mindfulness and the challenges of working from home in times of crisis," *Bus. Horiz.*, no. January, pp. 1–19, 2020.

[7] W. Shi, J. Li, and Y. Yang, "Face Fatigue Detection Method Based on MTCNN and Machine Vision," in *International Conference on Applications and Techniques in Cyber Intelligence ATCI 2019*, 2020, pp. 233–240.

[8] S. Palaniswamy and S. Tripathi, "Emotion recognition from facial expressions using images with pose, illumination and age variation for human-computer/robot interaction," *J. ICT Res. Appl.*, vol. 12, no. 1, pp. 14–34, 2018, doi: 10.5614/itbj.ict.res.appl.2018.12.1.2.

[9] N. Zhang, H. Zhang, and J. Huang, "Driver fatigue state detection based on facial key points," *2019 6th Int. Conf. Syst. Informatics, ICSAI 2019*, no. Icsai, pp. 144–149, 2019, doi: 10.1109/ICSAI48974.2019.9010478.

[10] H. Tao, G. Zhang, Y. Zhao, and Y. Zhou, "Real-time driver fatigue detection based on face alignment," *Ninth Int. Conf. Digit. Image Process. (ICDIP 2017)*, vol. 10420, no. Icdip, p. 1042003, 2017, doi: 10.1117/12.2282043.

[11] T. Qunzhu, Z. Rui, Y. Yufei, Z. Chengyao, and L. Zhijun, "Improvement of random forest cascade regression algorithm and its application in fatigue detection," *2019 2nd Int. Conf. Electron. Technol. ICET 2019*, pp. 499–503, 2019, doi: 10.1109/ELTECH.2019.8839317.

[12] Z. Li and L. Nianqiang, "Fatigue driving detection system based on face feature," *2019 2nd Int. Conf. Electron. Technol. ICET 2019*, pp. 525–529, 2019, doi: 10.1109/ELTECH.2019.8839479.

[13] Y. Kong and W. Li, "Research on recognition method of learning concentration based on face feature," *2017 IEEE Int. Conf. Cybern. Intell. Syst. CIS 2017 IEEE Conf. Robot. Autom. Mechatronics, RAM 2017 - Proc.*, vol. 2018-Janua, pp. 334–338, 2017, doi: 10.1109/ICCIS.2017.8274797.

[14] F. Cahyono, W. Wirawan, and R. Fuad Rachmadi, "Face Recognition System using Facenet Algorithm for Employee Presence," pp. 57–62, 2020, doi: 10.1109/icovet50258.2020.9229888.

[15] T. Nyein and A. N. Oo, "University Classroom Attendance System Using FaceNet and Support Vector Machine," in *2019 International Conference on Advanced Information Technologies (ICAIT)*, Nov. 2019, pp. 171–176, doi: 10.1109/AITC.2019.8921316.

[16] F. Schroff, D. Kalenichenko, and J. Philbin, "FaceNet: A unified embedding for face recognition and clustering," in *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2015, pp. 815–823, doi: 10.1109/CVPR.2015.7298682.

[17] A. Swaminathan, M. Chaba, D. K. Sharma, and Y. Chaba, "Gender Classification using Facial Embeddings: A Novel Approach," *Procedia Comput. Sci.*, vol. 167, no. 2019, pp. 2634–2642, 2020, doi: 10.1016/j.procs.2020.03.342.

[18] T. T. N. Thi and K. N. Trong, "An Efficient Face Detection and Recognition," *Int. J. Innov. Technol. Explor. Eng.*, vol. 7, no. 5, pp. 35–39, 2018.

[19] C. Guo and Y. Yang, "Implementation of a specified face recognition system based on video," *Proc. 2019 IEEE 4th Adv. Inf. Technol. Electron. Autom. Control Conf. IAEAC 2019*, no. Iaeac, pp. 79–84, 2019, doi: 10.1109/IAEAC47372.2019.8997627.

[20] P. Viola and M. J. Jones, "Robust Real-Time Face Detection," *Int. J. Comput. Vis.*, vol. 57, no. 2, pp. 137–154, 2004, doi: 10.1023/B:VISI.0000013087.49260.fb.

[21] R. Ghoddoosian, M. Galib, and V. Athitsos, "A realistic dataset and baseline temporal model for early drowsiness detection," *IEEE Comput. Soc. Conf. Comput. Vis. Pattern Recognit. Work.*, vol. 2019-June, pp. 178–187, 2019, doi: 10.1109/CVPRW.2019.00027.

[22] P. Viola and M. Jones, "Rapid Object Detection using a Boosted Cascade of Simple Features," in *Proceedings of the 2001 IEEE Computer Society Conference on Computer Vision and Pattern Recognition. CVPR 2001*, 2001, pp. 511–518.

[23] A. Susanto, D. Sinaga, C. A. Sari, E. H. Rachmawanto, and D. R. I. M. Setiadi, "A High Performance of Local Binary Pattern on Classify Javanese Character Classification," *Sci. J. Informatics*, vol. 5, no. 1, p. 8, 2018, doi: 10.15294/sji.v5i1.14017.

[24] V. Ong and D. Suhartono, "Using K-Nearest Neighbor in Optical Character Recognition," *ComTech Comput. Math. Eng. Appl.*, vol. 7, no. 1, p. 53, 2016, doi: 10.21512/comtech.v7i1.2223.

[25] K. B. Duan, J. C. Rajapakse, and M. N. Nguyen, "One-versus-one and one-versus-all multiclass SVM-RFE for gene selection in cancer classification," *Lect. Notes Comput. Sci. (including Subser. Lect. Notes Artif. Intell. Lect. Notes Bioinformatics)*, vol. 4447 LNCS, pp. 47–56, 2007, doi: 10.1007/978-3-540-71783-6_5.

[26] F. D. Adhinata, A. Harjoko, and Wahyono, "Object Searching on Video Using ORB Descriptor and Support Vector Machine," in *Advances in Computational Collective Intelligence*, 2020, pp. 239–251.