Image-based Face Verification for Student Identity Management — the TRUSTID Case Study

José N. Faria
jose.faria@isr.uc.pt
Institute of Systems and Robotics, University of Coimbra
Coimbra, Portugal

David Portugal
davidbsp@isr.uc.pt
Institute of Systems and Robotics, University of Coimbra
Coimbra, Portugal

Pedro Martins
perdromartins@isr.uc.pt
Institute of Systems and Robotics, University of Coimbra
Coimbra, Portugal

Marios Belk
belk@cognitiveux.de
Cognitive UX GmbH
Germany

Argyris Constantinides
constantinides.argyris@ucy.ac.cy
University of Cyprus, Cyprus and Cognitive UX GmbH
Germany

Andreas Pitsillides
andreas.pitsillides@ucy.ac.cy
University of Cyprus and University of Johannesburg (Visiting Professor)
South Africa

Christos A. Fidas
fidas@upatras.gr
University of Patras
Greece

ABSTRACT

Managing attendance and confirming student identity in online lessons is a contemporary problem of Higher Education Institutions (HEIs) that comes with its challenges, especially in classes with large numbers of students. The development and deployment of fully automatic and continuous image-based face verification techniques may provide a natural solution for this problem and assist instructors in handling the complexity of managing student identity in a remote classroom.

In this paper, we present an automatic image-based student identification framework. Our approach proposes a biometric authentication platform using a Residual Network (ResNet) Learning method for face verification. The system described is intended for use with consumer grade web cameras, providing a tradeoff between reliability and computational performance, as no assumption can be done regarding the target student hardware (i.e. CPU or GPU). This follows from the case study of TRUSTID, an European R&D initiative for intelligent student identity management in distance learning scenarios.

CCS CONCEPTS

• **Security and privacy → Biometrics**; • **Computing methodologies → Biometrics**; **Object identification**; **Neural networks**.

KEYWORDS

Face Verification, Image Processing, Student Identity Management

UMAP '23 Adjunct, June 26–29, 2023, Limassol, Cyprus

ACM ISBN 978-1-4503-9891-6/23/06.
https://doi.org/10.1145/3563359.3597397

ACM Reference Format:
José N. Faria, David Portugal, Pedro Martins, Marios Belk, Argyris Constantinides, Andreas Pitsillides, and Christos A. Fidas. 2023. Image-based Face Verification for Student Identity Management — the TRUSTID Case Study. In *UMAP '23 Adjunct: Adjunct Proceedings of the 31st ACM Conference on User Modeling, Adaptation and Personalization (UMAP '23 Adjunct), June 26–29, 2023, Limassol, Cyprus.* ACM, New York, NY, USA, 6 pages. https://doi.org/10.1145/3563359.3597397

1 INTRODUCTION

The COVID-19 pandemic has forced worldwide HEIs to adopt remote learning as a means to continue providing education to their students. This transition has brought challenges such as ensuring that the students are who they claim to be during online assessments. Automatic face verification strategies have the potential to revolutionize student identity management in remote learning. By using appropriate technology, HEIs can correctly authenticate the identity of their students, thereby reducing the possibility of academic dishonesty.

The potential benefits of automatic face verification for student identity management in remote learning are numerous. It can provide a secure and reliable method of verifying student identities, reduce the possibility of impersonation and cheating, and increase the credibility of online assessments [6]. Additionally, it can save time and resources for both educators and students. By using automatic face verification, educational institutions can streamline their identity management processes, minimize the risk of error, and ensure the integrity of their assessments. With the advancement of artificial intelligence and machine learning, automatic face verification technology is becoming more accessible and accurate, making it an increasingly viable solution for identity management in remote learning environments.

Recent years have seen a surge in research focused on the application of face verification to address the challenges of identity management in online academic assessments. Numerous studies

have demonstrated the potential of these technologies in ensuring academic integrity in remote learning environments. For instance, in [8], the authors have proposed a deep learning-based framework for face recognition-based identity verification in online exams, achieving high accuracy rates. Moreover, other studies such as [20] use image processing techniques for eye-gaze tracking and mouth-opening detection to maintain academic integrity in online learning by identifying and preventing suspicious behavior during online learning activities.

In fact, AI-based proctoring systems have gained attention recently as a means to ensure academic integrity in remote learning environments. A study by Labayen et al. [13] investigates the effectiveness of an AI-based proctor that combines facial and voice recognition, as well as keystroke typing analysis to monitor student behavior during online exams. The results showed that the system could accurately detect students identity as well as instances of cheating, such as the student not being alone or using an electronic device or book during online activities, and deter students from engaging in such behavior. In a similar vein, a study by Fenu et al. [4] proposes a multi-biometric system for the continuous and transparent authentication of students in e-learning activities, which uses score-level fusion of face, voice, touch, mouse and keystroke data, showing preliminary yet promising results for online student identification.

These studies highlight the potential of face verification and AI-based proctoring systems in remote learning environments, while also underscoring the need for ethical considerations and transparency in the use of these technologies to prevent privacy violations and ensure fairness in assessment. For a more detailed survey on continuous user identification in distance learning, the interested reader is referred to [17].

In TRUSTID, we seek to address the challenges of academic integrity in the digital age by designing, developing and evaluating a multi-tier continuous user identification system for higher education institutions. The system uses a combination of technologies, including biometrics, face verification, and behavioral analysis, to ensure the identification and authentication of students and detect any instances of cheating or misconduct during online assessments. In this paper, we report work in progress focused on our image-based face verification framework for remote learning. In Section 2, the requirements and challenges of face verification are discussed. Section 3 presents the proposed face verification system to ensure the identification and authentication of students during online activities. Section 4 preliminarily evaluates the effectiveness and reliability of the proposed face verification system. Finally, Section 5 describes concluding remarks and identifies areas for future research and development.

2 REQUIREMENTS AND CHALLENGES

It is important to understand the difference between face verification and face recognition when developing a system for continuous student identification. Face verification involves verifying whether a particular face matches a single pre-registered identity, while face recognition involves identifying a face from a pool of registered identities. For our case study, face verification is more suitable since it is generally faster and more accurate than face recognition.

A **face verification system for remote learning activities has become a necessity** due to the rise of online learning in recent times. This is critical for HEIs to provide credible, trustworthy and accurate degrees to their students, and therefore sustain their credibility in society. Moreover, face verification should be easily integrated with HEIs' Learning Management Systems (LMSs) to provide a seamless experience for both students and educators. This requirement is crucial because it ensures that the system can be easily incorporated, making it accessible to a larger number of users. Current state-of-the-art online education LMS of HEIs currently compromise students' continuous identification with traditional user authentication methods (e.g., passwords). It is urgent to adopt more secure and usable strategies for continuous student identification management, going beyond the single entry-point of authentication, and removing the tedious and low-value-added task by instructors of manually confirming individual student's identification.

Another important requirement for the system is that it **should work with any Commercial off-the-shelf (COTS) webcam, any Central Process Unit (CPU) without requiring a dedicated GPU and run on any Operating System (OS)**. This ensures that it can be easily deployable and accessible to students and educators regardless of their hardware and software configurations. Additionally, it reduces the need for costly hardware upgrades, making the system more affordable for institutions. Furthermore, the client-server deployment of the solution is a major challenge, as the system needs to be capable of handling a large number of simultaneous connections from different clients while maintaining fast response times. To overcome this, the system should be **designed with a scalable architecture** that can handle increasing traffic as the number of users increases without compromising usability.

The abovementioned requirement raises an important concern, which is the **computational trade-off between local-based and online-based identification systems**. Popular methods based on machine learning, and specifically deep learning, are improving the accuracy of user detection every day. However, these methods are also data hungry and computationally expensive. Local solutions require the availability of computational resources on the end-users workstation, while improving privacy and keeping infrastructure costs low; whereas online verification solutions, such as cloud-based architectures, allow continuous identification of users on virtually any terminal, e.g., smartphones, yielding other concerns such as transmission of sensitive data over the network, costs and even scalability issues.

In fact, **privacy-preserving challenges** are one of the major concerns for face verification systems. The system should be designed to ensure that students' privacy is protected while still providing accurate identification. This can be achieved by implementing techniques such as secure communication protocols, encryption, and anonymization. Core threats and challenges for designing secure and privacy-preserving biometric technologies relate to security, privacy and revocability of biometric data, as discussed in [5]. Solutions must assure that biometric data are processed and stored in a way to achieve high levels of security and sustain privacy-preservation aspects.

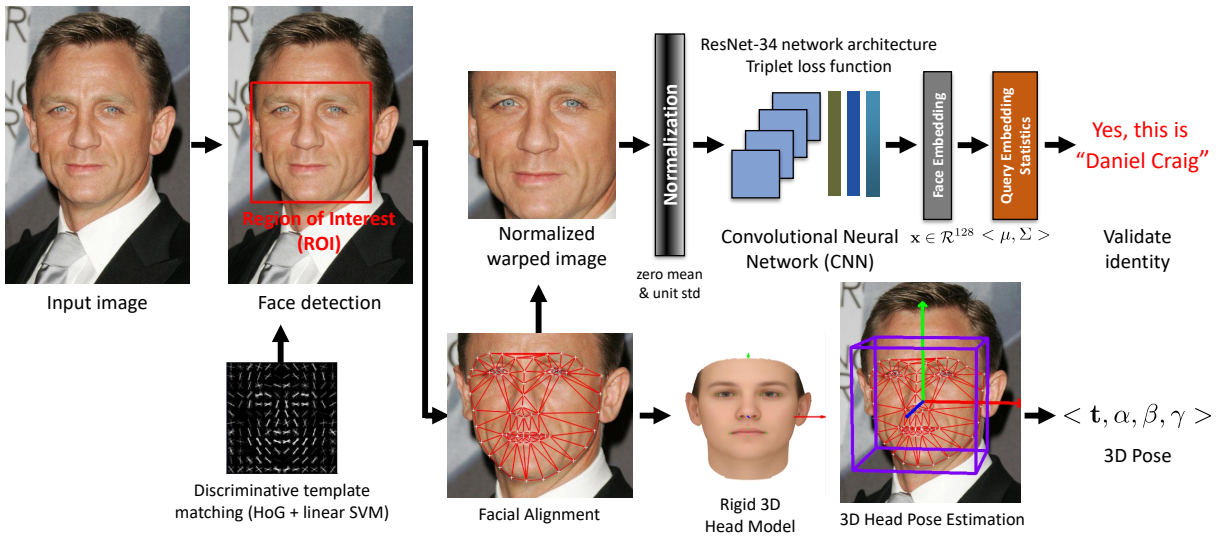

Figure 1: Proposed Face Verification Architecture of TRUSTID.

There are also **data protection barriers** and concerns that should be considered when collecting biometrics data from students. For instance, in the European Union, HEIs need to comply with the General Data Protection Regulation (GDPR) for continuous user identification systems. Thus, ethical considerations and transparency are fundamental in the use of these systems to avoid privacy violations and ensure fairness in assessment.

Finally, other important aspects that should be considered when developing continuous student identification systems include **ease of use, accessibility, and reliability**. The system should be easy to use and accessible to all students regardless of their technical proficiency. Additionally, it should be reliable and accurate, with minimal false positives, false negatives and manual intervention from instructors.

3 TRUSTID FACE VERIFICATION SYSTEM

With the previously identified requirements and challenges in mind, we have been developing and testing an online student identification framework using image-based face verification. Our system uses a desktop client application that supports Windows and Mac OS (see [19] for more details on the overarching TRUSTID architecture), which is a design choice that provides superior monitoring capabilities when compared to competing browser-based solutions. The system is open source[1], enables privacy preserving continuous user authentication using a client-server architecture, and only requires from the student a personal computer with internet and webcam access.

Since the system runs on any consumer grade webcam, it must consider low quality images, i.e. with low resolution, noise and blur, poor illumination and unexpected occlusions resulting from completely unconstrained environments. Moreover, it is designed to run with computational performance limitations, as the target student Personal Computers (PCs) are unknown and resource constrained. In spite of this, the system should be as accurate and reliable as possible.

Significant recent advances in the field of face recognition have been noticed, particularly with the widespread adoption of Deep Convolutional Neural Networks (DCNN) [12]. Although several methods can be found in the literature, DCNN have continuously shown superior performance and is thus considered the main line of research in this topic. In general, we find that there are two approaches when using CNNs to develop facial recognition algorithms: *i)* Training a CNN as **multi-class classifier** that separates different identities in the training set, such as using a softmax classifier; *ii)* Training a CNN for **face verification**, by mapping face images to a compact Euclidean space (face embeddings), where smaller distances between these embeddings correspond to similar faces.

The former (**face recognition**) is usually more adequate in cases when we work with a pre-defined number of identities, as the learned features are separable for the closed-set classification problem but not discriminative enough for the open-set face recognition problem, while the learned features in the latter (**face verification**) can even be used in open-set classification problems, but are harder to train due to the usage of complex loss functions.

In a preliminary stage of the work, we considered the face recognition approach. However, it needs to fully retrain the network to add new individuals, forcing the enrollment of new users to be a synchronous process and adding a significant delay in the collection of data of every individual to be recognized, which hinders the scalability of the approach. As such, we have moved towards the face verification approach, by collecting data and creating models for each individual, instead of training the network for all the users in the system.

In general, an image-based face verification framework follows a pipeline that requires three main modules: *i)* face detection; *ii)* image normalization and; *iii)* classification. For improved performance, our system has been developed in C++ using `dlib` [11], an

[1] https://github.com/cognitiveux/trustid

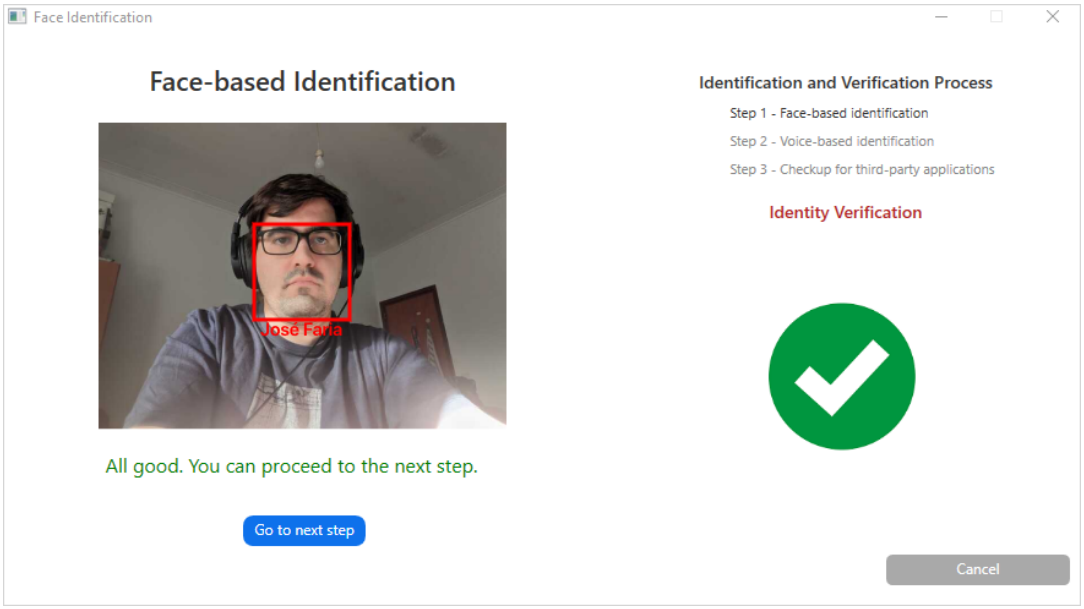

Figure 2: Face verification output as perceived by the student while running the TRUSTID client.

open-source toolkit that contains several machine learning algorithms.

For detecting faces in images, a linear detector combined with a Histogram of Oriented Gradients (HOG)-based features [2], and pyramidal image search is used. Afterwards, we run a nonrigid facial alignment method [10] to obtain the spatial face geometry, and then apply a similarity image warp (scaled crop) within a Region of Interest (RoI) to a normalized image of 150×150 pixels. Additionally, a pre-processing image normalization step with zero mean and unit variance is applied to change the range of pixel intensity of input images to feed the machine learning model.

Face verification is then pursued using a CNN. In our case, we make use of a ResNet-34 model [9] with $150 \times 150 \times 3$ sized inputs, and 128 dimensional vector output, pre-trained on a dataset of about 3 million faces derived from the Visual Geometry Group Face (VGG-Face) [16] and FaceScrub [15] datasets. To fit the model to the training data, we use Triplet loss [18] as a loss function. This network extracts a 128-D embedding, a mapping from facial images to a compact Euclidean space. We then use the Euclidean distance to estimate the similarity between the embeddings of the acquired images and the embeddings of the users' trained image data to ultimately verify the user identity, coupled with a k-Nearest Neighbors [7] voting classifier.

Figure 1 presents an overview of the architecture described for the face verification process of TRUSTID. It is noteworthy that we also extract a 3D head pose estimate [14] from the input image. At this stage of the work, this is used to check whether the data collected during student enrollment for authentication with the system has sufficient variability, e.g. looking straight at the camera, head titled, head cocked to the side, etc.

The TRUSTID student identity management system follows a client-server architecture [19], which incorporates the face verification framework gracefully as a modular component. The server consists of a web application that exposes endpoints through which the TRUSTID client applications can interact with and exchange data.

As mentioned before, special attention has been taken to data privacy and computational concerns. The server only maintains the student account management information, while the client-side application (see Figure 2) is responsible to run the full face identity verification pipeline, following the student's secure login. This means that all personal data, including enrollment images (user-specific training data), acquired images, face embeddings, etc. do not leave the student's computer, preserving his/her privacy.

Since all data processing occurs on the client, this avoids overloading the server, which in turn allows the system to be inherently scalable, as intended. Moreover, the ResNet model adopted is chosen due to its reduced computational cost in training and deployment. ResNets have a unique architecture that addresses the well-known vanishing gradient problem [3] in DCNNs, by using skip connections to allow for the flow of information across multiple layers. Therefore, they are able to achieve state-of-the-art performance on a range of image classification tasks with fewer parameters and computations compared to other networks. This is particularly relevant in our application, as clients are also not overloaded by the image-based face verification task, as seen in Section 4. Accordingly, the client application can be run on any target CPU, as specified initially. It is worth mentioning that we cannot resort to GPU-accelerated CNN inference due to limitations of running the application on unknown hardware.

Table 1: Face verification computational performance (First test).

Target CPU	Face Detection Runtime (ms)	Inference Runtime (ms)
Intel® Core™ i5-9300H CPU @ 2.40 GHz	251.24 ± 48.20	27.49 ± 6.96
Intel® Core™ i7-6700HQ CPU @ 2.60 GHz	267.36 ± 13.30	24.73 ± 3.37
Intel® Core™ i9-10980HK CPU @ 2.40 GHz	190.10 ± 9.87	16.45 ± 1.97
Intel® Core™ i7-10750H CPU @ 2.60 GHz	189.42 ± 9.73	16.38 ± 1.99
Intel® Core™ i7-11370H @ 3.30 GHz	161.59 ± 18.70	17.39 ± 3.66

4 PRELIMINARY RESULTS AND DISCUSSION

To validate the proposed system, we have prepared two different tests, one for assessing the computational performance of the face verification method, and a second one for evaluating the accuracy of the method.

The first test consists of running the face verification system on a range of different processor architectures. For this, face detection and inference is run on 100 pre-recorded 720p webcam images and run time per image is recorded. Information about the computer's capabilities is also extracted for analysis. We summarize the computational results in Table 1 in terms of mean face detection time and mean inference time for the different user processors tested.

Results show that even in lower-end CPUs, the system is able to verify the identity of the student in real-time, taking as much as 400 ms on an Intel® Core™ i5-9300H CPU (around 2.5 images per second). Typically, image-based face authentication occurs only when logging into the system, which means that most computers will verify the user with our system under a single second. Yet, for continuous student identity confirmation scenarios, the verification process can still be regularly called (e.g. every 15 seconds) without compromising overall system usability. We can also observe that face detection is the most time-consuming component of the algorithm, accounting for 90% of the total runtime, according to the test results. This means the detector is currently the bottleneck of our approach's performance.

The second test consists of a user study, with 133 volunteers (39 female and 94 male, aged between 18 to 50 years old) from three European countries (Portugal, Greece, Cyprus), whom have tested the current prototype of the TRUSTID client. During the study, we have adopted a protocol that considers users' privacy, confidentiality, anonymity, and the right to opt-out from the study at any time.

Participants were instructed to donwload and install the TRUSTID software client, and logged in using credentials received by email. In a first phase, participants have enrolled into the system by registering their face samples (user-specific training data) using their computer's web camera. Secondly, they completed the image-based authentication process, which runs the face verification pipeline described in Section 3. Thirdly, they have interacted with a mock online examination and were instructed to perform impersonation attacks as they wished, e.g. switching seats with another person. Aiming to control and evaluate the resilience of the system, participants had to inform the system about their actions through a feedback mechanism. This allowed us to compare what the system captured with what the participants attempted to do, aiming to evaluate the effectiveness of the implemented identification mechanism.

Table 2: Face verification accuracy (Second test).

Identification Case	Success Rate
Facial authentication for examination access	100%
Continuous Identification prior to impersonation	94.80%
Continuous Identification while impersonating	76.57%

To measure accuracy, we divide the number of correct identifications by the total number of identification attempts to obtain the success rate at the different stages of user interaction. Results in Table 2 show that the system successfully verified all registered users in order to join the examination. In addition, while performing the examination, the mechanism continuously verified the users 94.80% of the time, by analyzing face images randomly every 5 to 8 seconds. We observe that failures mainly occur due to face occlusion, inappropriate lighting conditions and specific head poses. Also, the system successfully detected impersonation attacks 76.57% of the time, conducted by a subset of 56 participants.

Also worth mentioning, feedback from 102 participants yielded a usability score of TRUSTID of 78.5% based on the System Usability Scale (SUS) [1].

5 CONCLUSION

In this paper, we present an image classification framework that utilizes a ResNet CNN for accurate and efficient face verification. Our framework supports any target CPU and web camera under Windows and MacOS, providing high usability and accessibility. The system has been designed to deliver a trade-off between computational complexity and reliability, while preserving user privacy and adhering to proper regulations in biometric data collection.

This framework is a crucial part of the TRUSTID scalable student identity management architecture to be implemented in various educational institutions, and the system's ability to integrate with other identification modalities such as voice, eye-gaze, keystroke, and interaction-based recognition can further be explored to enhance the current threat detection capabilities. Future work will also focus on further boosting the accuracy and performance of the system by: *i)* improving robustness to head pose variations, which often times leads to facial occlusion; ii) refining the image warp step by exploring piecewise affine warping; *iii)* adjusting the system to estimate similarity using a different distance metric instead of Euclidean distance; *iv)* integrating liveness and eye blinking detection; and *iv)* optimizing the face detector's performance.

ACKNOWLEDGMENTS

This work is partially supported by TRUSTID – Intelligent and Continuous Online Student Identity Management for Improving Security and Trust in European Higher Education Institutions (Grant Agreement No: 2020-1-EL01-KA226- HE-094869), which is funded by the European Commission within the Erasmus+ 2020 Program, https://trustid-project.eu/.

REFERENCES

[1] John Brooke et al. 1996. SUS - A quick and dirty usability scale. *Usability evaluation in industry* 189, 194 (1996), 4–7.

[2] Navneet Dalal and Bill Triggs. 2005. Histograms of oriented gradients for human detection. In *2005 IEEE computer society conference on computer vision and pattern recognition (CVPR'05)*, Vol. 1. Ieee, 886–893.

[3] Mohammad Sadegh Ebrahimi and Hossein Karkeh Abadi. 2021. Study of residual networks for image recognition. In *Intelligent Computing: Proceedings of the 2021 Computing Conference, Volume 2*. Springer, 754–763.

[4] Gianni Fenu, Mirko Marras, and Ludovico Boratto. 2018. A multi-biometric system for continuous student authentication in e-learning platforms. *Pattern Recognition Letters* 113 (2018), 83–92.

[5] Christos Fidas, Marios Belk, David Portugal, and Andreas Pitsillides. 2021. Privacy-preserving biometric-driven data for student identity management: challenges and approaches. In *Adjunct Proceedings of the 29th ACM Conference on User Modeling, Adaptation and Personalization*. 368–370.

[6] Christos A. Fidas, Marios Belk, Argyris Constantinides, David Portugal, Pedro Martins, Anna Maria Pietron, Andreas Pitsillides, and Nikolaos Avouris. 2023. Ensuring Academic Integrity and Trust in Online Learning Environments: A Longitudinal Study of an AI-centered Proctoring System in Tertiary Educational Institutions. *Education Sciences* (2023).

[7] Evelyn Fix and Joseph Lawson Hodges. 1951. Discriminatory analysis, nonparametric estimation: consistency properties. *Report 4, Project nº 21-49* 4 (1951).

[8] Asep Hadian Sudrajat Ganidisastra and Yoanes Bandung. 2021. An incremental training on deep learning face recognition for m-learning online exam proctoring. In *2021 IEEE Asia Pacific Conference on Wireless and Mobile (APWiMob)*. IEEE, 213–219.

[9] 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.

[10] Vahid Kazemi and Josephine Sullivan. 2014. One millisecond face alignment with an ensemble of regression trees. In *2014 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 1867–1874.

[11] Davis E King. 2009. Dlib-ml: A machine learning toolkit. *The Journal of Machine Learning Research* 10 (2009), 1755–1758.

[12] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. 2017. Imagenet classification with deep convolutional neural networks. *Commun. ACM* 60, 6 (2017), 84–90.

[13] Mikel Labayen, Ricardo Vea, Julián Flórez, Naiara Aginako, and Basilio Sierra. 2021. Online student authentication and proctoring system based on multimodal biometrics technology. *IEEE Access* 9 (2021), 72398–72411.

[14] Pedro Martins and Jorge Batista. 2008. Accurate single view model-based head pose estimation. In *2008 8th IEEE International Conference on Automatic Face & Gesture Recognition*. IEEE, 1–6.

[15] Hong-Wei Ng and Stefan Winkler. 2014. A data-driven approach to cleaning large face datasets. In *2014 IEEE International Conference on Image Processing (ICIP)*. IEEE, 343–347.

[16] Omkar M Parkhi, Andrea Vedaldi, and Andrew Zisserman. 2015. Deep face recognition. (2015).

[17] David Portugal, José N. Faria, Marios Belk, Pedro Martins, Argyris Constantinides, Anna Pietron, Andreas Pitsillides, Nikolaos Avouris, and Christos A. Fidas. Springer, 2023. Continuous User Identification in Distance Learning: A Recent Technology Perspective. *Smart Learning Environments* (Springer, 2023). Under Submission.

[18] Florian Schroff, Dmitry Kalenichenko, and James Philbin. 2015. Facenet: A unified embedding for face recognition and clustering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 815–823.

[19] Taoufik Sousak, Argyris Constantinides, José Nuno Faria, Anna Pietron, Pedro Martins, David Portugal, Marios Belk, Andreas Pitsillides, and Christos Fidas. 2022. Towards Intelligent and Continuous Online Student Identity Management. In *Invited Talk at UMAP '22: 30th ACM Conference on User Modeling, Adaptation and Personalization, Workshop on Adaptive and Personalized Privacy and Security (APPS)*. Barcelona, Spain.

[20] Potluri Tejaswi, Sista Venkatrama Phani Kumar, and Kolli Venkata Krishna Kishore. 2023. Proctor net: An AI framework for suspicious activity detection in online proctored examinations. *Measurement* 206 (2023), 112266.