

Automated Visual Attention Detection using Mobile Eye Tracking in Behavioral Classroom Studies

Efe Bozkir
Technical University of Munich
Arcisstr. 21
80333, Munich, Germany
efe.bozkir@tum.de

Christian Kosel
Technical University of Munich
Arcisstr. 21
80333, Munich, Germany
christian.kosel@tum.de

Tina Seidel
Technical University of Munich
Arcisstr. 21
80333, Munich, Germany
tina.seidel@tum.de

Enkelejda Kasneci
Technical University of Munich
Arcisstr. 21
80333, Munich, Germany
enkelejda.kasneci@tum.de

ABSTRACT

Teachers' visual attention and its distribution across the students in classrooms can constitute important implications for student engagement, achievement, and professional teacher training. Despite that, inferring the information about where and which student teachers focus on is not trivial. Mobile eye tracking can provide vital help to solve this issue; however, the use of mobile eye tracking alone requires a significant amount of manual annotations. To address this limitation, we present an automated processing pipeline concept that requires minimal manually annotated data to recognize which student the teachers focus on. To this end, we utilize state-of-the-art face detection models and face recognition feature embeddings to train face recognition models with transfer learning in the classroom context and combine these models with the teachers' gaze from mobile eye trackers. We evaluated our approach with data collected from four different classrooms, and our results show that while it is possible to estimate the visually focused students with reasonable performance in all of our classroom setups, U-shaped and small classrooms led to the best results with accuracies of approximately 0.7 and 0.9, respectively. While we did not evaluate our method for teacher-student interactions and focused on the validity of the technical approach, as our methodology does not require a vast amount of manually annotated data and offers a non-intrusive way of handling teachers' visual attention, it could help improve instructional strategies, enhance classroom management, and provide feedback for professional teacher development.

Keywords

educational technologies, eye tracking, computer vision, machine learning, deep learning

Efe Bozkir, Christian Kosel, Tina Seidel, and Enkelejda Kasneci. Automated Visual Attention Detection using Mobile Eye Tracking in Behavioral Classroom Studies. In Caitlin Mills, Giora Alexandron, Davide Taibi, Giosuè Lo Bosco, and Luc Paquette (eds.) Proceedings of the 18th International Conference on Educational Data Mining, Palermo, Italy, July, 2025, pp. 104–118. International Educational Data Mining Society (2025).

© 2025 Copyright is held by the author(s). This work is distributed under the Creative Commons Attribution NonCommercial NoDerivatives 4.0 International (CC BY-NC-ND 4.0) license.

<https://doi.org/10.5281/zenodo.15870187>

1. INTRODUCTION

The integration of computer vision and machine learning techniques into the educational context has led to a new era of data analysis due to the likelihood of automating various time-consuming and non-trivial tasks that require manual human inputs and labeling [18, 35]. Traditional educational research methods that focus on exploring teacher and student behavior in the classroom, relying heavily on human observation and self-reports such as surveys and interviews, have provided valuable insights into teacher and student behavior [7]. However, they face challenges in scalability, especially when analyzing behaviors across multiple classrooms in various settings and institutions. Moreover, even with trained human observers manually coding the data, it is possible to miss subtle behavioral nuances among diverse student groups. The advent of computer vision and machine learning techniques can address some of these challenges if such techniques work accurately. Some previous examples to this end include employing machine learning models to predict students' cognitive load [6] from physiological and behavioral data in classrooms [79] and analyzing students' learning-related behaviors [1]. However, most approaches focus on one-sided attention, such as understanding students' behavioral patterns. These aspects can be combined and fused with teachers' attention and visual behavior simultaneously, providing the opportunity to have a fine-grained understanding of what happens in the classroom, and equipping teachers with eye trackers can enable this.

Eye tracking has been utilized in different domains, such as education [68, 19], medicine [41, 16], marketing [78, 9], entertainment [2, 56], and driving [11, 10] and for different tasks including reading [71, 13], problem-solving [75, 48], and expertise assessment [20, 43]. Yet, many of these setups explicitly use remote and screen-based eye trackers, requiring the users to sit and remain stable during data collection, which might affect ecological validity negatively. When classroom context and especially teachers are considered, using mobile eye tracking (MET) might solve this issue to a certain extent while providing fine-grained visual attention data of teachers during the classroom discourse [23].

In a broad sense, non-invasive eye tracking is widely recognized as one of the powerful methods in behavioral research for studying human cognition and perception [29], while its popularity has grown in cognitive-oriented studies within educational psychology, including investigations into research on classroom observation and teacher expertise [45, 66]. As eye-tracking methods mostly capture gaze patterns and pupillometry, this provides insights into how teachers orient themselves in school classrooms that are characterized by a high level of social interaction and close proximity among students, resulting in a focused social environment that is highly relevant for study attention processes. To this end, previous research has focused on how teachers allocate their attention in the classroom, interact with students, and create effective learning environments [36]. By examining teachers’ gaze patterns and visual attention, studies aimed to understand the dynamics of teacher-student interactions, student engagement, and classroom management [47, 45, 70].

Despite the aforementioned potentials in the classroom context, much of what is currently known about teachers’ perception in the classroom is through lab-based and on-action eye-tracking studies in which participants observe classroom video sequences of others [36]. As a result, there is a need for systematic studies of teachers’ perceptions in the classroom while teaching, not merely while observing [45]. MET is a powerful way of capturing teacher attention during in-action real classroom teaching without disrupting the instruction flow [70]. However, in many studies, MET data analysis is grounded in in-depth area-of-interest (AOI)-based analysis, either using time-consuming and error-prone frame-by-frame manual gaze mapping [22, 60] or dynamic AOIs using either keyframes or different types of markers [31]. Both methods have drawbacks, especially when participants move and interact with objects and other people, which is often the case in the classroom context, making MET studies with several hours of recordings difficult to analyze quantitatively. As a result, researchers are reluctant to incorporate MET in their study designs [60].

To overcome these methodological limitations and provide ecological validity, we developed a specialized processing pipeline to automatically map teacher attention on the students during classroom interactions using computer vision and machine learning techniques, to answer the research question (RQ) of “How can novel face detection and recognition methods be integrated with mobile eye tracking to map teacher attention in conventional classroom settings?”. Answering this RQ, our work aims to empirically show that such a pipeline can accurately work and support further research on teacher-student interactions and behaviors, while supporting the research on teacher-student interactions. To this end, we first train face recognition models for the students in our classrooms using the very initial part of the scene videos of egocentric eye trackers by extending state-of-the-art face detection and recognition methods with transfer learning. Then, in the second part, using the estimated gaze point of the teacher, we identify the closest face patch to the teachers’ gaze in the scene and feed this patch to our face recognition models to identify visually attended students for each frame in the egocentric video streams. As we manually annotated student face patches only from the first minute of

the video data for our face recognition models, due to the effectiveness of our approach, in practice, practitioners would only need a minimal amount of student face data to automatically allocate the attention in video streams, without heavily relying on manual annotations. We show the effectiveness and accuracy of our technical solution in four different classroom situations, and to the best of our knowledge, our work is the first to conduct such analyses in authentic classrooms. Considering all, our major contributions are as follows:

- We propose a novel, automated computational processing pipeline to detect the student that the teacher visually focuses on in the conventional classroom context, achieving substantially better performance compared to chance by utilizing and further extending state-of-the-art face detection and recognition models (i.e., SCRFD-10GF [38], ResNet50 and ResNet100 models trained on WebFace600K [85] and Glint360K [4] datasets) with conventional classifiers like random forest, support vector machine, and k-nearest neighbor by transfer learning.
- We show that it is possible to accurately identify visually attended students in challenging setups, such as when eye-tracker scene cameras are used, where the teacher is non-stationary, and when the real-world environment includes partial occlusions and varying illumination conditions.
- Our approach can provide further insights to facilitate research on the relationships between teachers’ visual attention and different student characteristics, such as success, self-concept, and personality traits, while allowing more time for researchers to focus on the interpretation instead of dealing with manual annotations.

2. RELATED WORK

In this section, we discuss the related works in teacher attention studies using MET in classroom observation, and face processing and eye tracking in the wild in Sections 2.1 and 2.2, respectively.

2.1 Eye-tracking Data in Classroom Studies

Examining the quality of learning-relevant interactions between teachers and students in the classroom makes it possible to predict student achievement [30]. This relationship is supported by a large body of observational studies highlighting the critical role of pedagogical and emotional guidance in developing a student’s cognitive, social, and emotional skills [61]. Previous works often employed various classroom observation methods (e.g., protocols, videotaped lessons, etc.) to explore the relationship between learning outcomes and classroom interactions. In this process, specially trained coders assess classroom dynamics, either live or through video recordings [69]. These coders assess various aspects, such as the teachers’ responsiveness or the quality of feedback [8, 42]. In addition to traditional and high-inferential data obtained through manual coding of classroom interactions, researchers have begun collecting low-inferential data, such as eye tracking, to identify qualitative

differences in teacher behaviors and cognition [36]. However, most studies in this area relied on on-action and lab-based eye-tracking data [36]. These studies often show participants video sequences from authentic teaching situations and analyze their attention regarding different aspects of classroom interactions, such as student engagement and disruptions [47].

In recent years, the availability of cost-effective glasses-based MET has enhanced the study of gaze behavior beyond the lab. MET glasses include two types of integrated cameras: an eye-tracking camera positioned under the eye and a scene camera typically mounted above the eyebrows on the frame [73]. Through precise calibration, the glasses combine the measured gaze point captured by the eye-tracking camera with the broader perspective provided by the scene camera. This integration allows for a comprehensive analysis of an individual's visual attention and facilitates a deeper understanding of their gaze behavior in the wild [64, 70]. However, only few studies conducted in classrooms with teachers have used MET [24, 44, 55, 35, 62].

To this end, Goldberg et al. [35] analyzed which learners' behavior attracted pre-service teachers' attention during instruction using MET. The authors synchronized eye-tracking data with a continuous annotation of observable learning behavior and found that pre-service teachers' attention was mainly attracted by students' salient on-task learning-related behaviors, such as asking questions, raising their hands, or reflecting on something out loud. Pouta et al. [62] went beyond merely recording frontal teaching, a common focus in previous studies [24]. Instead, they honed in on specific episodes of student-teacher interactions. The authors analyzed the teachers' visual perception during these interactions, examining their gaze directions. This approach allowed them to distinguish between passive observation and active engagement.

Teachers' visual attention can rest on students without immediate engagement, serving as a background monitoring process [47]. Research showed that these situations, in which teachers monitor students, are used for collecting information about students' levels of interest and motivation and constitute a large part of their visual behavior in the classroom. Researchers are particularly interested in this monitoring behavior because it involves more top-down driven gaze processes, reflecting deliberate and goal-oriented attention allocation by teachers. At other times, teachers direct their gaze in response to specific student behaviors, such as a raised hand or a verbal question, engaging with students in meaningful ways [35]. Understanding how teachers navigate between these different attentional states underscores the complexity of their interactions with students and highlights the importance of context in interpreting gaze data. Previous research also recognizes this complexity for future studies [36]. In the context of automated attention understanding, a system should also work accurately in situations of passive observation, where the face of the student is not always perfectly visible from the front, as is more often the case in situations of direct interaction or hand-raising. When such a system can accurately detect and interpret visual attention in various classroom scenarios, it will likely provide a more comprehensive understanding of teacher-student in-

teractions. It is then also possible to further extend such analyses towards a multi-modal fashion in more meaningful ways, as previous works partly addressed by extracting teaching activities from the sensor data [63], analyzing expertise for teaching situations [33], and capturing the processes of social interaction [39].

While there has been some progress in understanding teachers' attention processes in the wild, our knowledge remains limited [45]. The studies in this context suffer from significant constraints that hinder a comprehensive understanding, as many of them are characterized by a small sample size, which can reduce the generalizability of the findings [35]. Additionally, the limited recording time in these investigations may not capture the full spectrum of attentional behaviors [47]. These limitations highlight the need for more comprehensive and rigorous research to delve deeper into teachers' attentional processes [45]. There is a growing consensus, as highlighted by previous studies [35, 72, 18], on the increasing need for automated procedures, especially in data handling and analysis. Combining computer vision and machine learning techniques in the context of face detection, face recognition, and eye tracking could offer valuable solutions for this purpose.

2.2 Face Processing and Eye Tracking In the Wild

Face detection and recognition tasks, particularly in the wild, are often considered challenging due to varying environmental factors such as illumination conditions, partial occlusions, and different camera viewing angles [81, 50]. Despite these, deep learning-based solutions provide decent performance even in such challenging conditions, providing the possibility of carrying out further analyses taking face processing results as bases [34, 77]. The importance of face detection and processing in the wild was constantly emphasized [54, 50, 40, 83, 46], which is also essential for accurate educational data mining in the classroom context.

To this end, different approaches exist in the literature. Zhang et al. [82] proposed a face detector consisting of rapidly digested and multiple-scale convolutional layers that can run on CPU in real time, whose speed is invariant of detected faces. In another work, Li et al. [49] introduced a novel face detection network with three key contributions, including a feature enhancer, adoption of a progressive anchor loss, and integration of an anchor assign strategy into data augmentation, achieving superior performance in several face detection benchmarks. Later, Deng et al. [26] proposed RetinaFace, a single-shot and multi-level face localization method that unifies face box prediction, 2D facial landmark detection, and 3D vertices under the target of point regression on the image plane, achieving a stable face detection performance while keeping the efficiency through single-shot inference. Liu et al. [52] studied anchors in the context of face detection and observed that some unmatched anchors have a significant regression ability and further proposed an anchor mining strategy to sample high-quality anchors for the training phase. Such a mining strategy outperformed state-of-the-art methods in multiple datasets.

Recently, Zhu et al. [84] treated the face detection problem as a generic object detection problem by using a ResNet50

backbone and showed in a single-model and single-scale setting that their method, TinaFace, outperforms the state-of-the-art, especially with test-time data augmentation. Guo et al. [38] identified efficient face detection with low computation cost and high precision as an unsolved problem and focused on training data sampling and computation distribution strategies by introducing two methods, including a data augmentation technique for training samples in the most needed stages and a computation redistribution strategy by reallocating the computation between the backbone and head of the model. The proposed method, SCRFD, outperformed TinaFace [84] while being three times faster.

Another essential step, especially for studies in the wild in the classroom context, is facial recognition in challenging conditions and with noisy data. To this end, Deng et al. [27] introduced an additive margin loss (i.e., ArcFace) and further extended ArcFace by sub-centering to deal with faces under real-world noise. The authors showed that without training any additional generator or discriminator network, the ArcFace can provide discriminative feature embeddings. While ArcFace provides effective feature embeddings, it assumes clean training data, and to address this issue and enable robustness to noise, Deng et al. [25] further iterated ArcFace by introducing sub-classes for each identity and showed with large-scale raw web faces that their method achieves state-of-the-art face recognition performance in several benchmarks. In another work, Deng et al. [28] proposed representing every class as a distribution instead of points in the latent space for face recognition, and the authors showed that their approach could simulate sample-to-sample comparisons for classification while being computationally efficient and memory-saving. Other work [85, 3] also tried to unveil the power of large-scale noisy datasets to improve the performance of face recognition models, yet acknowledged the rise in computing costs when training sets include a significant number of identities.

While face detection and recognition tasks can be combined semantically to understand what is going on in the real-world environment, to understand the activities of users (i.e., teachers or students in the educational context) and how they react to a particular situation or stimulus, eye-tracking data can be helpful [45]. Fine-grained eye-tracking data requires dedicated hardware, such as MET devices, to estimate participants' gaze. The MET algorithms are different from the remote eye tracking setups, as in MET, eye cameras are in close proximity to the eyes. To tackle this challenge, Feng et al. [32] focused on a real-time gaze tracking algorithm that can operate on a mobile processor by continuously utilizing predicted regions of interests of near-eye images, and the authors achieved accuracies lower than one degree. Angelopoulos et al. [5] also focused on high-speed gaze tracking utilizing an event-based method offering eye-tracking frequencies beyond 10KHz by using an online 2D pupil fitting method. In another work, Lu et al. [53] proposed a pupil localization method with a novel deep learning-based corneal refraction correction and argued that pupil localization in 3D space is an essential intermediate algorithmic step for accurate gaze tracking and showed that their method helps significantly decrease gaze estimation error.

Ranjan et al. [65] focused on an appearance-based and head-

pose invariant gaze estimation using convolutional neural networks and showed that their approach is ten times faster than the state-of-the-art approaches. Similarly, Palmero et al. [58] and Stojanov et al. [67] focused on the benefits of temporal and depth information on the gaze estimation tasks and found that incorporating both temporal and depth information helps to improve the performance of gaze estimation tasks. Other work focused on unsupervised representation learning to learn low-dimensional gaze representations without gaze annotations [80] and the notion of two-eye asymmetry to utilize an asymmetric regression network to predict 3D gaze directions [21]. While various works focused on different aspects of gaze estimation, face recognition, and face detection, they constantly emphasize the importance of real-time working capability, robustness to noise, and the efficacy of the methods for the real-world use of algorithms. We took these recommendations into account while selecting the methods we use for the education context in the classrooms and utilized SCRFD-10GF [38] for the face detection and the ResNet50 architecture that was trained on the WebFace600K dataset (ResNet50@WebFace600K) [85] for face recognition. We benchmarked facial embeddings of ResNet50@WebFace600K against facial embeddings that come from a ResNet100 trained on the Glint360K dataset (ResNet100@Glint360K) [4], which is considered one of the largest and cleanest face recognition datasets in the literature. Upon obtaining optimal facial embeddings, we utilized transfer learning with several classifiers to identify visually attended students by leveraging facial embeddings and teachers' gaze in the classroom context.

3. METHODOLOGY

Our methodology has two main pillars, including the MET data collection and automated teacher attention mapping. We discuss these separately in Sections 3.1 and 3.2, respectively.

3.1 Experimental Design

The following subsections introduce our experimental design, particularly our participants, data types, and data collection procedure. The Bavarian Ministry of Education and Cultural Affairs reviewed and approved our data protection measures, the design and necessity of the survey, and ethical points.

3.1.1 Participants and Data

Four in-service mathematics teachers (two women and two men) participated in the data collection. Teachers were between 27 and 62 years old ($M = 37.25$, $SD = 16.64$). The class sizes ranged from 14 to 25 students (69 students total). The collected data included video streams from the classrooms and eye-tracking data (i.e., gaze data) recorded from the mobile eye trackers. The resolution of the videos recorded from the eye trackers' scene camera is 1920×1080 .

3.1.2 Data Collection Procedure

The mobile eye-tracking recordings took place during a regular class period, chosen to interfere as little as possible with the regular lesson plan. We used a Tobii Pro Glasses 2 [74] with an eye-tracking sampling frequency of 50 Hz to collect eye movement data. Recordings started after a successful one-point eye-tracking calibration. If an interruption due to

a technical error happened, we recalibrated the eye-tracking system in the same way. During the data collection process, we connected the eye-tracking system to a computer via a wireless connection, where we controlled the recordings with the Tobii Pro Glasses Controller ($\times 64$) [73].

Before data collection began, all participants, including students, teachers, and the students’ parents (or legal guardians), provided written informed consent. The consent forms explicitly outlined the scope and purpose of the study, ensuring that all parties were fully aware that the sessions would involve video recordings and eye-tracking technology. Participants also agreed to appear in academic publications or presentations in disguised forms, such as blurred or altered screenshots.

After the consent phase, each teacher gave a lesson ranging between 60 and 70 minutes in four different higher secondary schools (ninth grade) in Germany. All participating teachers taught similar content (i.e., matrix calculus) during the data collection. In addition, the sampled lesson was minimally pre-determined to allow for consistency across teachers and their lessons. The teachers had five minutes to recap the topic and tasks of the last lesson and the remaining time to introduce new content. We advised the participating teachers to keep their eye-tracking glasses stable throughout the lesson. At the end of the data collection procedure, we thanked teachers and students for participating in our study. Figure 1 depicts different classrooms from which we collected the data.

3.2 Automated Attention Mapping of Teacher

We first introduce our general methodology on automated attention mapping of teachers in Section 3.2.1. We then provide implementation and evaluation details in Section 3.2.2.

3.2.1 Method

The primary purpose of the proposed automated attention mapping pipeline is to automatically identify which student teachers visually focus on at each time point, and this processing pipeline includes two main components. We explain those in the following and summarize the overall process in Figure 2. The results of this attention mapping pipeline can be used as a baseline for further student-teacher interaction analysis or educational data mining in the classroom context.

Face detection and recognition. To identify which student is visually focused on by the teacher, as a first step (i), the use of face detection and recognition models is necessary. For our pipeline, we selected the SCRFD-10GF [38] and the ResNet50@WebFace600K models [85] for face detection and recognition, respectively. SCRFD-10GF includes a substantially smaller number of model parameters and comes with a shorter inference time compared to various other methods for face detection [84, 26]. Considering that a smaller number of parameters and shorter inference times for face detection will also help with real-time face detection through the scene cameras of the eye tracker if deployed, the use of SCRFD-10GF is reasonable, as it also achieves good performance in the face detection tasks, even in challenging setups. As for the face recognition part, to obtain facial embeddings, we use the ResNet50 architecture that was trained on the

WebFace600K dataset, which was cleaned in an automated fashion and could include noisy samples, similar to our face instances from classroom environments. We also apply face alignment with five facial key points (i.e., eye centers, nose, and mouth corners) between our face detection and facial embedding components.

While the aforementioned publicly available machine learning models were trained and evaluated with many users, they do not include students from our classroom setups. This is not an issue for face detection tasks, but for face recognition, we need to incorporate our individuals and their data from the classrooms into our face recognition pipeline. To this end, instead of carrying out the face recognition task with the ResNet50@WebFace600K model, we use its facial feature embeddings of size 512 for transfer learning and further train five conventional classifiers, including the random forest, k-nearest neighbor, support vector machine, gradient boosting, and decision tree classifiers. These models serve as baselines for the teachers’ attention mapping.

Teacher gaze mapping. As the second step in our pipeline (ii), mapping teachers’ gaze to particular students is necessary to complete the attention mapping pipeline. In this step, we rely on the estimated gaze from the eye tracker on video recordings. In parallel, we apply face detection via SCRFD-10GF on the frames coming from the eye tracker’s scene camera. For each frame, once our processing pipeline provides detected faces and 2D coordinates of the gaze point of the teacher, we identify the closest detected face patch to the teacher’s gaze. To identify the closest face patch, we use the Euclidean distance between the teacher’s gaze and the center point of each face patch on the image frames. We do not consider teachers’ eye movement type (e.g., fixation, saccade) and keep this aspect exploratory. Following, we send the feature embedding of the closest face patch, which is extracted from the pre-trained ResNet50@WebFace600K, to our face recognition model (i.e., random forest, support vector machine, etc.) to identify which student the teacher visually focuses on. We carry out this process continuously during the lesson while the teacher wears MET.

3.2.2 Implementation and Evaluation

We implemented our attention mapping processing pipeline using InsightFace [37] and scikit-learn [59] libraries in Python and evaluated it for each classroom video separately, in an offline fashion. To train the face recognition models, we first applied our face detection pipeline frame by frame to each video. We randomly annotated 30 detected face instances per student from the first minute of the videos on the detected faces, and we decided on this number empirically from our data. We used the facial feature embeddings extracted from those 30 face instances from the ResNet50@WebFace600K model to train the random forest, k-nearest neighbor, support vector machine, gradient boosting, and decision tree classifiers to facilitate face recognition in the classroom context. For each of these classifiers, we applied 5-fold cross-validation to tune the hyperparameters of the classifiers by using a grid search approach. During the hyperparameter tuning phase, we preserved the percentages of samples for each class when splitting the data. We provide the complete set of hyperparameters in Table 1. We used the tuned hyperparameters to train the final models for each



(a) A view from the first classroom.



(b) A view from the second classroom.



(c) A view from the third classroom.



(d) A view from the fourth classroom.

Figure 1: Views from the video recordings from eye trackers that teachers wear in each of our four classrooms. We provide blurred versions of the images to protect students’ privacy.

classroom video for the evaluations. For feature normalization, we standardized features by removing the mean and scaling them to unit variance using the training data statistics to avoid data leakage in the validation and test phases.

We used the first minute of the videos to annotate students’ faces because, at the beginning of the lessons, teachers’ visual focus covers a greater field of view, making it possible to identify many student faces. In addition, while the activities that would happen in the later phases of the lessons can differ based on the classroom setup and lesson type, the first several minutes of the lessons are often generic in terms of activities that take place, which helps systematically apply and evaluate the designed processing pipeline. However, any ground truth annotation for the students’ faces should help generate a similar processing pipeline.

To implement gaze mapping and identify the visually attended student, while we relied on the gaze point Tobii eye trackers provided, we carried out some approximations in our automated pipeline to evaluate whether our pipeline works accurately. Firstly, for the gaze signal, if signals from both eyes were valid (i.e., gaze direction is detected accurately, no blinks occurred, etc.), we averaged the gaze locations of both eyes to have the final single point. If only one eye’s signal was valid, we assumed that such information would also provide valuable insights into the teacher’s attention and took the final gaze point on the classroom video from the single eye. Secondly, we calculated the center points of each face patch coming from the SCRFD-10GF face detection model and measured the Euclidean distances between these center points and the final gaze location on the same frame. If multiple faces were detected, we identified the focused student’s face patch based on the shortest distance. We then fed the facial feature embeddings of the identified face patch based on the shortest distance to our

Table 1: Hyperparameter search grids for machine learning models. RF, SVM, k-NN, GB, and DT indicate random forest, support vector machine, k-nearest neighbor, gradient boosting, and decision tree classifiers, respectively. If we found the same best working hyperparameter for all the classrooms for each facial feature embedding type, we report them with bold values for ResNet50@WebFace600K, whereas with underlines for ResNet100@Glint360K.

Model	Hyperparameter grid
RF	Maximum depth: [10, 20, 40, 80, 160] Maximum features: [sqrt, log2] Minimum samples leaf: [<u>1</u> , 2, 3, 4, 5] Minimum samples split: [2, 4, 8, 16] Number of estimators: [25, 50, 100, 200, 400, 600, 800, 1000, 1200, 1400, 1600]
SVM	C: [0.001, 0.005, 0.01, 0.05, 0.1, 0.5, 1] Gamma: [<u>auto</u> , scale] Kernel: [rbf, <u>linear</u> , poly, sigmoid]
k-NN	k: [5 , 7, 9, 11, 13, 15, 17, 19, 21]
GB	Number of estimators: [100, <u>200</u> , 400, 600] Minimum samples leaf: [1, 2] Minimum samples split: [2, 4, 8]
DT	Maximum depth: [10, 20, 40, 80, 160] Maximum features: [sqrt, log2] Minimum samples leaf: [1, 2, 3, 4, 5] Minimum samples split: [2, 4, 8, 16] Criterion: [gini, entropy]

classification models (e.g., random forest) and obtained the final result on the student ID. If no gaze signal was available due to both eyes giving invalid signals, we did not evaluate

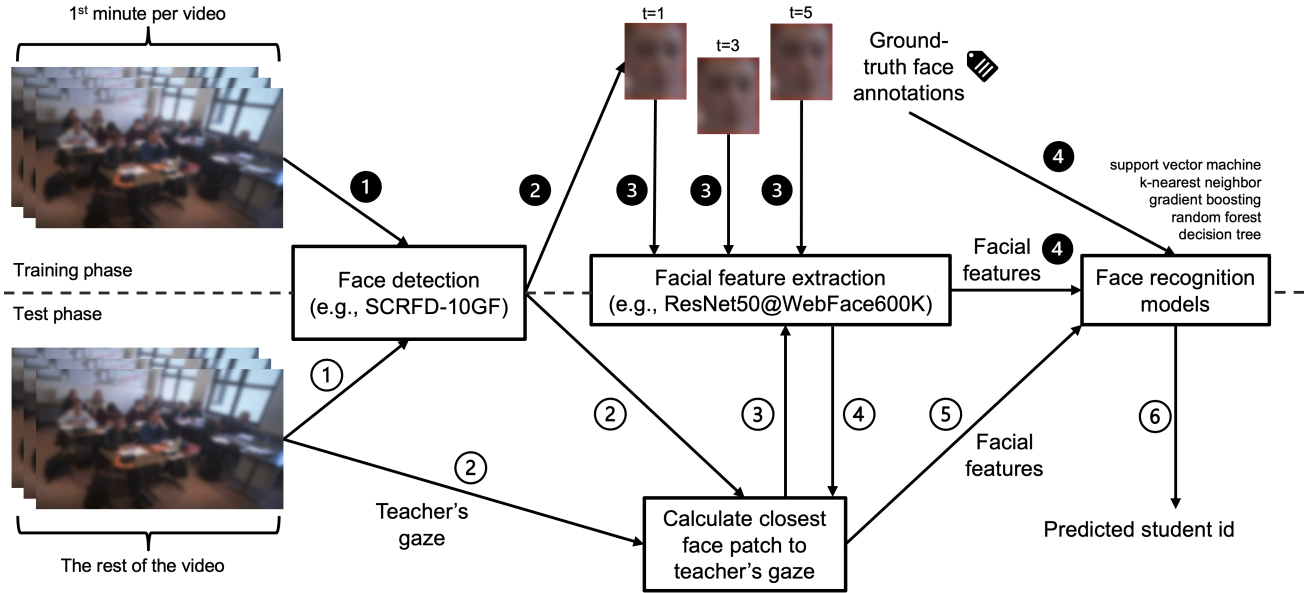


Figure 2: Overview of our automated teacher attention mapping pipeline. We provide blurred versions of the images to protect student privacy. Dashed lines separate the training and testing processes, whereas numbers represent the operations’ order. We use Euclidean distance when calculating the closest patch to the teacher’s gaze.

these cases, as a missing gaze signal would directly lead to a wrong result.

To evaluate the gaze mapping part of our pipeline, we needed to have annotated ground truth data about where the teachers focus in the scene, considering the detected faces. Similar to face annotations from the first minute, we annotated approximately the second minute of each video with the information on which student teachers focused and conducted our evaluations according to this manually annotated ground-truth data. Our annotation process involved two trained annotators, and to ensure the reliability of the annotations, we established inter-rater reliability by calculating the Cohen’s kappa coefficient, which resulted in a high level of agreement ($\kappa = .94$), indicating consistency among annotators. We performed double coding for 25% of the data.

In evaluations, for facial feature embeddings, apart from the ResNet50@WebFace600K, we took ResNet100@Glint360K [4] as a baseline since the Glint360K dataset includes cleaner face instances than our dataset, as well as the WebFace600K. The baseline evaluations with ResNet100@Glint360K only replace the ResNet50@WebFace600K component and keep the rest of the pipeline identical. Considering all, we evaluated our automated processing pipeline using data from four authentic classrooms, and in Table 2, we report the number of training and test instances.

4. RESULTS

We report the results separately as we apply our analyses to each classroom video. For each video, we first report accuracy, precision, recall, and F₁ scores of automated teacher attention mapping results for random forest, support vector machine, k-nearest neighbor, gradient boosting, and decision tree classifiers using ResNet50@WebFace600K

Table 2: Number of training and test samples for each classroom.

Classroom ID	# Training Samples	# Test Samples
1	720	459
2	510	502
3	420	190
4	390	935

and ResNet100@Glint360K facial feature embeddings in Table 3. For our specific task, random forests, support vector machines, and k-nearest neighbor classifiers have great potential for use in practice. In addition, facial feature embeddings from ResNet50@WebFace600K outperform the baseline from ResNet100@Glint360K. Furthermore, we report confusion matrices for the best-performing models for each classroom video in Figure 3 to show how the samples for each student correspond to the overall performance.

As depicted in Table 3, random forest classifiers outperform the others in almost all classroom setups, namely support vector machines, k-nearest neighbors, gradient boosting, and decision trees. Comparing the k-nearest neighbors and support vector machines with random forests, while there is a directly observable performance difference in the first and second classrooms, the performance of these classifiers is comparable in the third and fourth classrooms, and they are greater than the performance from the first two classrooms. When the number of students and seating arrangements are considered in each of these classrooms, as depicted in Figure 1, fewer students exist in the third and fourth classrooms compared to the first two. In addition, the third classroom’s seating structure is U-type, whereas the fourth classroom is smaller in terms of number of students, and the teacher has a


Table 3: Results for automated teacher attention mapping. We indicate the best results and settings in bold for each classroom.

Classroom ID	Classifier	Facial Feature Embeddings	Accuracy	Precision	Recall	F ₁ score
1	Random Forest	ResNet50@WebFace600K ResNet100@Glint360K	0.50 0.46	0.68 0.64	0.50 0.46	0.56 0.52
	Support Vector Machine	ResNet50@WebFace600K ResNet100@Glint360K	0.44 0.43	0.65 0.65	0.44 0.43	0.51 0.50
	k-Nearest Neighbor	ResNet50@WebFace600K ResNet100@Glint360K	0.40 0.32	0.69 0.61	0.40 0.32	0.45 0.40
	Gradient Boosting	ResNet50@WebFace600K ResNet100@Glint360K	0.24 0.25	0.42 0.46	0.24 0.25	0.28 0.29
	Decision Tree	ResNet50@WebFace600K ResNet100@Glint360K	0.11 0.19	0.25 0.33	0.11 0.19	0.14 0.23
2	Random Forest	ResNet50@WebFace600K ResNet100@Glint360K	0.58 0.42	0.71 0.67	0.58 0.42	0.61 0.45
	Support Vector Machine	ResNet50@WebFace600K ResNet100@Glint360K	0.49 0.40	0.72 0.73	0.49 0.40	0.54 0.47
	k-Nearest Neighbor	ResNet50@WebFace600K ResNet100@Glint360K	0.44 0.18	0.73 0.59	0.44 0.19	0.48 0.23
	Gradient Boosting	ResNet50@WebFace600K ResNet100@Glint360K	0.18 0.27	0.53 0.47	0.18 0.27	0.21 0.28
	Decision Tree	ResNet50@WebFace600K ResNet100@Glint360K	0.17 0.12	0.42 0.38	0.17 0.12	0.23 0.15
3	Random Forest	ResNet50@WebFace600K ResNet100@Glint360K	0.68 0.60	0.74 0.62	0.68 0.60	0.69 0.59
	Support Vector Machine	ResNet50@WebFace600K ResNet100@Glint360K	0.63 0.58	0.72 0.63	0.63 0.58	0.65 0.58
	k-Nearest Neighbor	ResNet50@WebFace600K ResNet100@Glint360K	0.63 0.57	0.70 0.60	0.63 0.57	0.64 0.57
	Gradient Boosting	ResNet50@WebFace600K ResNet100@Glint360K	0.32 0.38	0.63 0.55	0.32 0.38	0.36 0.41
	Decision Tree	ResNet50@WebFace600K ResNet100@Glint360K	0.33 0.16	0.46 0.19	0.33 0.16	0.35 0.16
4	Random Forest	ResNet50@WebFace600K ResNet100@Glint360K	0.91 0.92	0.93 0.93	0.92 0.92	0.92 0.92
	Support Vector Machine	ResNet50@WebFace600K ResNet100@Glint360K	0.94 0.93	0.95 0.94	0.94 0.93	0.94 0.93
	k-Nearest Neighbor	ResNet50@WebFace600K ResNet100@Glint360K	0.95 0.89	0.95 0.92	0.95 0.89	0.95 0.90
	Gradient Boosting	ResNet50@WebFace600K ResNet100@Glint360K	0.61 0.72	0.79 0.83	0.61 0.72	0.67 0.76
	Decision Tree	ResNet50@WebFace600K ResNet100@Glint360K	0.51 0.31	0.68 0.64	0.51 0.31	0.56 0.39


better vision of students due to the proximity. However, the first two classrooms are rather conventional, and especially for the faces of the students, a lot of occlusion situations happen. Considering all the results, we found that our automated processing pipeline works accurately when the faces of the students are visible without too many occlusions.

In addition to the performance analyses, we conducted more fine-grained analyses considering the results for each visually


attended student individually by using the precision metric and the confusion matrices provided in Figure 3. We found that in the first classroom, the most successful attention mapping results came from one of the students who sat in the front (i.e., OOU86), whereas the second and third best results came from students who sat in the third row (i.e., TAE96 and FLE76). In the second classroom, there is a more apparent trend that the first and second most successful attention mapping results came from the students who




(a) Confusion matrix of the random forest classifier for the first classroom.



(b) Confusion matrix of the random forest classifier for the second classroom.



(c) Confusion matrix of the random forest classifier for the third classroom.



(d) Confusion matrix of the k-nearest neighbor classifier for the fourth classroom.

Figure 3: Confusion matrices of the best-performing models. Labels correspond to different students in the classrooms.

were seated in the front part of the classroom (i.e., PIR92 and KAN63), whereas the third most successful recognition results came from a student who sat in the second row (i.e., DEN41).

When we conducted the same analyses for the third and fourth classrooms, we also observed a clear trend, similar to the results for the second classroom. For instance, for the third classroom, the first and the third most successful attention-mapping results came from the students who sat in the frontal part of the classroom (i.e., RUG03 and SAI23), whereas the second most performant results came from the students who were sitting in the visible parts of the classroom (i.e., WFG73). Similarly, the samples of the front-sitting students (e.g., KOG75, WAG15, and RLG15) in the fourth classroom contributed the most to the success of our approach, indicating the importance of the visibility of the student faces for our automated attention mapping pipeline.

Lastly, across all the classroom scenarios, we found that experiments that utilized facial feature embeddings from the ResNet50@WebFace600K outperformed those that utilized the ResNet100@Glint360K, which we explored as a baseline. WebFace600K [85] dataset applies an automated cleaning pipeline for filtering the high-quality images, whereas Glint360K [4] is considered one of the cleanest datasets in the face recognition area. As the face patches we extracted from our classroom videos can be classified as noisy images, it is reasonable that facial feature embeddings from the ResNet50@WebFace600K work more accurately than our baseline.

5. DISCUSSION

In eye-tracking research, the post-processing step of mapping AOIs is essential in correlating eye movement metrics with specific parts of the stimulus, such as the duration of gaze on a particular object within that stimulus. The analysis of human attention in the wild, as captured by MET, is a time-consuming process requiring detailed mapping of

AOIs in each video frame. The primary objective of this study was to develop a processing pipeline to automate this manual process using computer vision and machine learning techniques. While previous studies have primarily focused on the automated detection of object classes like vehicles [76] or animals [17], only a few have addressed the complexities of human-to-human interactions, where face recognition becomes pivotal for automated AOI mapping [70]. We particularly designed our processing pipeline to efficiently map AOIs in egocentric videos from an eye tracker, focusing on sequences with (multiple) faces. Our goal is to assist researchers who examine human-to-human attention dynamics in natural, real-world settings, with a special focus on challenging environments. To this end, we have pinpointed traditional classroom settings in schools as an ideal context for testing this automated mapping. Classroom environments typically feature dense and complex social interactions, making them well-suited for advancing our understanding of attentional behaviors in situ. The key result of this study indicates that our approach can accurately detect the student receiving the teacher’s visual attention across various classroom settings, particularly when employing techniques like random forests and k-nearest neighbors, in conjunction with SCRFD-10GF [38] and ResNet50@WebFace600K [85]. This finding and our processing pipeline can facilitate follow-up work on studying teacher-student interactions, assessing teaching quality, and understanding student learning.

5.1 Implications for Automated Attention Mapping in Educational Research

Our automated teacher attention mapping solution provides many opportunities for educational data analysis and mining, as if it is deployed in the real world; it only requires several face images for each student, just like modern smartphone or tablet facial authentication processes. Once these are in place and the machine learning classifiers are trained, teachers’ valid eye-tracking signal on the video stimulus is adequate for its use. Previous research [18] empirically showed that the correlations between automated processes such as hand-raising detection and learning-related student behaviors are comparable to those obtained from the manually labeled data. While we did not explicitly analyze similar correlations with and without using our automated attention mapping pipeline using the teachers’ gaze, as our method works accurately across different classrooms, educational researchers can likely utilize our processing pipeline for their research, similar to the previous work, to understand how the teachers’ attention is related to different student characteristics and whether manual and automated approaches provide correlated results.

While we utilized such a processing pipeline for the first time in the literature using the videos from conventional classrooms, previous research focused on smaller and standardized seminar-like classroom settings where the number of students is few, and it is less likely to have significant issues related to occlusions [70]. Similar to previous research, we also show that teachers’ attention could be accurately mapped even when a greater number of students and issues regarding illuminations and partial occlusions exist. In addition, we found that random forests, k-nearest neighbors, and support vector machines perform well for such purposes, the latter overlapping with the findings of the pre-

vious work [70]. One of the main caveats is that our automated pipeline works more accurately if students’ faces are more visible in the classroom, because it depends on face detection and recognition accuracy (i.e., third and fourth classrooms). Therefore, it is more reasonable to utilize it in classrooms where students’ faces are visible during the classroom discourse, especially if researchers plan any further analyses based on the results of our automated pipeline.

Despite the decent performance, our approach should not be treated as a classroom monitoring system. Due to ethical and privacy reasons, we argue that practitioners should utilize it in real time to avoid recording any personally identifiable image or video data. When it is used in real time, apart from the benefits from manual data annotations point of view, one can also implement additional systems highlighting specific regions in the classrooms with visual or auditory cues, complementing our automated pipeline so that if a teacher does not pay attention to the particular student(s) throughout the lesson, such a feedback system can warn teacher about this to help improve teaching quality and further support developing students’ self-concepts.

5.2 Ethical Considerations

We process video and eye-tracking data in our methods, and as video data includes the faces of the minors, which are directly identifiable data, ethical considerations are necessary. In addition, it is possible to extract sensitive personal information using eye-tracking data in different setups [51, 15]. Therefore, responsible data processing from the teachers’ eye-tracking data perspective is also essential.

As we utilize face detection and recognition models capable of running in real time, when our method is employed, it is possible to process such data without recording it in the first place, ensuring data privacy to a significant extent. In case such video recordings are stored for further offline analyses, anonymization techniques from the classroom context [72] would not only be suitable but also fit well with our approach, as such methods directly work at the face patch level without distorting the remaining parts of the video data. Comparably, as almost all commercial eye-tracking devices support real-time gaze processing, we encourage practitioners to apply these analyses in real time. If storing or sharing eye-tracking data is necessary, especially in the classroom context, computational approaches that can help preserve data privacy should be in place [14, 12, 57].

5.3 Limitations and Future Works

Our approach has great potential in automating teacher attention mapping processes in classrooms; however, we identified certain limitations when practitioners actively utilize such a pipeline. Firstly, we used the videos from the first two minutes of the classroom instruction to design and assess our method. While such an amount of data is sufficient as it already includes video frames and eye-tracking data from different angles in the classroom due to the physical movement of the teacher, as this data corresponds to the beginning of the recordings, the likelihood of having drifted eye-tracking signal is low as the eye tracker is newly calibrated at the beginning of the data collection. Therefore, if our approach is in place for long durations, it is essential for future works

to keep an eye on the additional time fractions regarding the eye-tracking data quality to avoid any inaccuracies in attention mapping. For data quality purposes, teachers can recalibrate the device after a certain time threshold to keep the attention mapping accuracy high. In addition, wearing such an eye tracker might also negatively affect the ecological validity, especially if students and teachers are not familiar with classroom environments with such intelligent systems. However, improving the technology literacy of teachers and students may help with this, especially since it is already needed, considering the fast-growing landscape of machine learning and artificial intelligence.

In our work, we attempted to standardize the classroom situation by using the first minutes of a lesson during which the teacher conducts whole-group instruction, meaning they stand in front of the students, and each student is mainly looking at the teacher. This standardization helps to maintain a consistent and clear view of both the teacher and the students, optimizing the performance of our system. However, in other pedagogical practices like group work, where teachers walk through the room, or students look away more often, the performance of our system might decrease. These dynamic and less structured interactions challenge our current setup as face detection and recognition become more complex with increased movement and significant occlusions. Future works should address these scenarios to ensure the system’s robustness across various teaching methods. It may be interesting to dive deeper into the relationship with occlusions and different video segments; however, complete occlusions likely lead to faces not being detected at all. Therefore, either manual occlusion annotations or student detections using body movements may be needed to analyze such a relationship.

The performance of our automated approach heavily depends on face detection and face recognition results. While we utilized SCRFD-10GF and ResNet50@WebFace600K due to their efficiency and real-time capabilities for our pipeline and ResNet100@Glint360K for baseline comparisons, other models can be interchangeably used. However, if students’ faces are occluded due to their seating positions in the classroom or if there are significant illumination changes in the environment, the performance of our pipeline would likely drop. Due to these, we deliberately evaluated our approach when a face or faces are detected in the scene to avoid any performance decrease due to any issue in the face detection task. Future research may especially focus on the relationship between seating positions and classification accuracy for each student, especially when the sample size is bigger than ours.

Furthermore, researchers following our approach should be aware that the stimulus material to be automatically mapped must include all faces visible for the ground truth annotations. If data from any individual is unavailable (e.g., due to the face detector missing the patches) in different parts of the egocentric video material (i.e., training or test data for the particular machine learning task), one might consider manually including these data samples. Otherwise, these individuals will not be included in the training and testing phases of the experiments, as we also observed in our results. Especially when the training dataset misses some of

the individuals, the visual attention mapping performance will automatically drop if these particular individuals are available in the test set.

Lastly, despite the reasonable performance, we did not utilize the proposed processing pipeline to study student-teacher interactions, teaching quality, or student learning. These remain the focus of our future research. In addition, future research can build on our findings to develop real-time dashboards that inform teachers, visualize engagement levels, highlight areas needing attention, and provide data-driven recommendations for instructional strategies.

6. CONCLUSION

In this paper, we proposed an automated attention mapping pipeline to identify which student teachers focus on without needing a vast amount of manually labeled data in eye-tracked classroom setups by using and extending state-of-the-art face detection, face recognition, and gaze tracking methods. Our findings imply that when students’ faces are not completely occluded and visible from the teachers’ field of view, it is possible to use such an automated pipeline to recognize the visually attended students accurately. As our approach requires minimal manually labeled classroom data, it will likely enrich further analyses in classrooms and educational contexts, such as relationships between teachers’ visual attention and different learner profiles and characteristics. In future work, we plan to evaluate our approach in more diverse classrooms with varying seating arrangements and study the relationship between teachers’ predicted visual attention and teaching effectiveness.

Acknowledgments

Parts of this research were conducted while E.B. and E.K. were affiliated with the University of Tübingen. This research was partly funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany’s Excellence Strategy - EXC number 2064/1 - Project number 390727645.

7. REFERENCES

- [1] T. B. Abdallah, I. Elleuch, and R. Guermazi. Student behavior recognition in classroom using deep transfer learning with VGG-16. *Procedia Computer Science*, 192:951–960, 2021. Knowledge-Based and Intelligent Information & Engineering Systems: Proceedings of the 25th International Conference KES2021.
- [2] S. Almeida, Óscar Mealha, and A. Veloso. Video game scenery analysis with eye tracking. *Entertainment Computing*, 14:1–13, 2016.
- [3] X. An, J. Deng, J. Guo, Z. Feng, X. Zhu, Y. Jing, and L. Tongliang. Killing two birds with one stone: Efficient and robust training of face recognition CNNs by partial FC. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4032–4041, 2022.
- [4] X. An, X. Zhu, Y. Xiao, L. Wu, M. Zhang, Y. Gao, B. Qin, D. Zhang, and Y. Fu. Partial FC: Training 10 million identities on a single machine, 2021.
- [5] A. N. Angelopoulos, J. N. Martel, A. P. Kohli, J. Conradt, and G. Wetzstein. Event-based near-eye gaze tracking beyond 10,000 hz. *IEEE Transactions*

- on *Visualization and Computer Graphics*, 27(5):2577–2586, 2021.
- [6] T. Appel, P. Gerjets, S. Hoffmann, K. Moeller, M. Ninaus, C. Scharinger, N. Sevchenko, F. Wortha, and E. Kasneci. Cross-task and cross-participant classification of cognitive load in an emergency simulation game. *IEEE Transactions on Affective Computing*, 14(2):1558–1571, 2023.
 - [7] S. Apter, Arnold. A mass observation study of student and teacher behaviour in british primary classrooms. *Educational Psychology in Practice*, 26(2):151–171, 2010.
 - [8] G. Blomberg, A. Renkl, M. G. Sherin, H. Borko, and T. Seidel. Five research-based heuristics for using video in pre-service teacher education; fünf evidenzbasierte heuristiken für den einsatz von video in der universitären lehrerbildung. *Journal for educational research online*, 5(1):90–114, 2013.
 - [9] S. C. Boerman and C. M. Müller. Understanding which cues people use to identify influencer marketing on instagram: an eye tracking study and experiment. *International Journal of Advertising*, 41(1):6–29, 2022.
 - [10] E. Bozkir, D. Geisler, and E. Kasneci. Assessment of driver attention during a safety critical situation in VR to generate VR-based training. In *ACM Symposium on Applied Perception 2019*, pages 1–5. ACM, 2019.
 - [11] E. Bozkir, D. Geisler, and E. Kasneci. Person independent, privacy preserving, and real time assessment of cognitive load using eye tracking in a virtual reality setup. In *2019 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, pages 1834–1837, 2019.
 - [12] E. Bozkir, O. Günlü, W. Fuhl, R. F. Schaefer, and E. Kasneci. Differential privacy for eye tracking with temporal correlations. *PLOS ONE*, 16(8):1–22, 2021.
 - [13] E. Bozkir, G. Kasneci, S. Utz, and E. Kasneci. Regressive saccadic eye movements on fake news. In *2022 Symposium on Eye Tracking Research and Applications*, pages 1–7. ACM, 2022.
 - [14] E. Bozkir, A. B. Ünal, M. Akgün, E. Kasneci, and N. Pfeifer. Privacy preserving gaze estimation using synthetic images via a randomized encoding based framework. In *ACM Symposium on Eye Tracking Research and Applications*, pages 1–5. ACM, 2020.
 - [15] E. Bozkir, S. Özdel, M. Wang, B. David-John, H. Gao, K. Butler, E. Jain, and E. Kasneci. Eye-tracked virtual reality: A comprehensive survey on methods and privacy challenges, 2023.
 - [16] J. Bradshaw, F. Shic, A. N. Holden, E. J. Horowitz, A. C. Barrett, T. C. German, and T. W. Vernon. The use of eye tracking as a biomarker of treatment outcome in a pilot randomized clinical trial for young children with autism. *Autism Research*, 12(5):779–793, 2019.
 - [17] P. Buehler, B. Carroll, A. Bhatia, V. Gupta, and D. E. Lee. An automated program to find animals and crop photographs for individual recognition. *Ecological Informatics*, 50:191–196, 2019.
 - [18] B. Bühler, R. Hou, E. Bozkir, P. Goldberg, P. Gerjets, U. Trautwein, and E. Kasneci. Automated hand-raising detection in classroom videos: A view-invariant and occlusion-robust machine learning approach. In N. Wang, G. Rebolledo-Mendez, N. Matsuda, O. C. Santos, and V. Dimitrova, editors, *Artificial Intelligence in Education*, pages 102–113, Cham, 2023. Springer Nature Switzerland.
 - [19] T. Busjahn, C. Schulte, B. Sharif, Simon, A. Begel, M. Hansen, R. Bednarik, P. Orlov, P. Ihantola, G. Shchekotova, and M. Antropova. Eye tracking in computing education. In *Proceedings of the Tenth Annual Conference on International Computing Education Research*, page 3–10. ACM, 2014.
 - [20] N. Castner, E. Kasneci, T. Kübler, K. Scheiter, J. Richter, T. Eder, F. Hüttig, and C. Keutel. Scanpath comparison in medical image reading skills of dental students: Distinguishing stages of expertise development. In *Proceedings of the 2018 ACM Symposium on Eye Tracking Research & Applications*, pages 1–9. ACM, 2018.
 - [21] Y. Cheng, F. Lu, and X. Zhang. Appearance-based gaze estimation via evaluation-guided asymmetric regression. In *Proceedings of the European Conference on Computer Vision (ECCV)*, page 105–121, 2018.
 - [22] L. Chukoskie, S. Guo, E. Ho, Y. Zheng, Q. Chen, V. Meng, J. Cao, N. Devgan, S. Wu, and P. C. Cosman. Quantifying gaze behavior during real-world interactions using automated object, face, and fixation detection. *IEEE Transactions on Cognitive and Developmental Systems*, 10(4):1143–1152, 2018.
 - [23] K. S. Cortina, K. F. Miller, R. McKenzie, and A. Epstein. Where low and high inference data converge: Validation of class assessment of mathematics instruction using mobile eye tracking with expert and novice teachers. *International Journal of Science and Mathematics Education*, 13(2):389–403, 2015.
 - [24] K. S. Cortina, K. Müller, J. Häusler, K. Stürmer, T. Seidel, and K. F. Miller. Feedback mit eigenen augen: Mobiles eyetracking in der lehrerinnen- und lehrerbildung; feedback through one’s own eyes: Mobile eye tracking in teacher education. *Beiträge zur Lehrerinnen- und Lehrerbildung*, 36(2):208–222, 2018.
 - [25] J. Deng, J. Guo, T. Liu, M. Gong, and S. Zafeiriou. Sub-center arcfacel: Boosting face recognition by large-scale noisy web faces. In *Proceedings of the IEEE Conference on European Conference on Computer Vision*, page 741–757, 2020.
 - [26] J. Deng, J. Guo, E. Ververas, I. Kotsia, and S. Zafeiriou. Retinaface: Single-shot multi-level face localisation in the wild. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5202–5211. IEEE/CVF, 2020.
 - [27] J. Deng, J. Guo, N. Xue, and S. Zafeiriou. ArcFace: Additive angular margin loss for deep face recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4685–4694. IEEE/CVF, 2019.
 - [28] J. Deng, J. Guo, J. Yang, A. Lattas, and S. Zafeiriou. Variational prototype learning for deep face recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11906–11915, June 2021.
 - [29] M. K. Eckstein, B. Guerra-Carrillo, A. T. Miller

- Singley, and S. A. Bunge. Beyond eye gaze: What else can eyetracking reveal about cognition and cognitive development? *Developmental Cognitive Neuroscience*, 25:69–91, 2017.
- [30] J. Englehart. Teacher–student interaction. *Science*, 177:711–722, 01 2009.
- [31] K. M. Evans, R. A. Jacobs, J. A. Tarduno, and J. B. Pelz. Collecting and analyzing eye-tracking data in outdoor environments. *Journal of Eye Movement Research*, 5(2):1–19, 2012.
- [32] Y. Feng, N. Goulding-Hotta, A. Khan, H. Reyserhove, and Y. Zhu. Real-time gaze tracking with event-driven eye segmentation. In *2022 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, pages 399–408, 2022.
- [33] H. Gao, E. Bozkir, P. Stark, P. Goldberg, G. Meixner, E. Kasneci, and R. Göllner. Detecting teacher expertise in an immersive vr classroom: Leveraging fused sensor data with explainable machine learning models. In *2023 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*, pages 683–692, 2023.
- [34] M. M. Ghazi and H. K. Ekenel. A comprehensive analysis of deep learning based representation for face recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, pages 102–109, 2016.
- [35] P. Goldberg, Ö. Sümer, K. Stürmer, W. Wagner, R. Göllner, P. Gerjets, E. Kasneci, and U. Trautwein. Attentive or not? toward a machine learning approach to assessing students’ visible engagement in classroom instruction. *Educational Psychology Review*, 33:27–49, 2021.
- [36] A.-S. Grub, A. Biermann, D. Lewalter, and R. Brünken. Professional vision and the compensatory effect of a minimal instructional intervention: A quasi-experimental eye-tracking study with novice and expert teachers. *Frontiers in Education*, 7:1–17, 2022.
- [37] J. Guo and J. Deng. Insightface: 2d and 3d face analysis project. <https://github.com/deepinsight/insightface>, 2025. last access 07/02/2025.
- [38] J. Guo, J. Deng, A. Lattas, and S. Zafeiriou. Sample and computation redistribution for efficient face detection. In *Tenth International Conference on Learning Representations*, pages 1–10, 2022.
- [39] M. S. Hannula, E. Haataja, E. Löfström, E. Garcia Moreno-Esteva, J. F. A. Salminen-Saari, and A. Laine. Advancing video research methodology to capture the processes of social interaction and multimodality. *ZDM – Mathematics Education*, 54(2):433–443, 2022.
- [40] M. U. Haq, M. A. J. Sethi, S. Ahmad, M. A. ELAffendi, and M. Asim. Automatic player face detection and recognition for players in cricket games. *IEEE Access*, 12:41219–41233, 2024.
- [41] K. Harezlak and P. Kasprowski. Application of eye tracking in medicine: A survey, research issues and challenges. *Computerized Medical Imaging and Graphics*, 65:176–190, 2018. *Advances in Biomedical Image Processing*.
- [42] T. Hofkens, R. C. Pianta, and B. Hamre. Teacher-student interactions: Theory, measurement, and evidence for universal properties that support students’ learning across countries and cultures. In R. Maulana, M. Helms-Lorenz, and R. M. Klassen, editors, *Effective Teaching Around the World: Theoretical, Empirical, Methodological and Practical Insights*, pages 399–422, Cham, 2023. Springer International Publishing.
- [43] B. W. Hosp, F. Schultz, O. Höner, and E. Kasneci. Soccer goalkeeper expertise identification based on eye movements. *PLoS ONE*, 16(5):1–22, 2021.
- [44] Y. Huang, K. F. Miller, K. S. Cortina, and D. Richter. Teachers’ professional vision in action. *Zeitschrift für Pädagogische Psychologie*, 37:122–139, 2021.
- [45] H. Jarodzka, I. Skuballa, and H. Gruber. Eye-tracking in educational practice: Investigating visual perception underlying teaching and learning in the classroom. *Educational Psychology Review*, 33(1):1–10, 2021.
- [46] R. Jurevičius, N. Goranin, J. Janulevičius, J. Nugaras, I. Suzdalev, and A. Lapusinskij. Method for real time face recognition application in unmanned aerial vehicles. *Aviation*, 23(2):65–70, 2019.
- [47] C. Kosel, R. Böheim, K. Schnitzler, D. Holzberger, J. Pfeffer, M. Bannert, and T. Seidel. Keeping track in classroom discourse: Comparing in-service and pre-service teachers’ visual attention to students’ hand-raising behavior. *Teaching and Teacher Education*, 128:1–12, 2023.
- [48] J. Y. Lee, J. Donkers, H. Jarodzka, and J. J. van Merriënboer. How prior knowledge affects problem-solving performance in a medical simulation game: Using game-logs and eye-tracking. *Computers in Human Behavior*, 99:268–277, 2019.
- [49] J. Li, Y. Wang, C. Wang, Y. Tai, J. Qian, J. Yang, C. Wang, J. Li, and F. Huang. DSFD: Dual shot face detector. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5055–5064. IEEE/CVF, 2019.
- [50] P. Li, L. Prieto, D. Mery, and P. J. Flynn. On low-resolution face recognition in the wild: Comparisons and new techniques. *IEEE Transactions on Information Forensics and Security*, 14(8):2000–2012, 2019.
- [51] D. J. Liebling and S. Preibusch. Privacy considerations for a pervasive eye tracking world. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing: Adjunct Publication*, page 1169–1177. ACM, 2014.
- [52] Y. Liu, X. Tang, J. Han, J. Liu, D. Rui, and X. Wu. HAMBox: Delving into mining high-quality anchors on face detection. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 13043–13051. IEEE/CVF, 2020.
- [53] C. Lu, P. Chakravarthula, K. Liu, X. Liu, S. Li, and H. Fuchs. Neural 3D gaze: 3D pupil localization and gaze tracking based on anatomical eye model and neural refraction correction. In *2022 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*, pages 375–383, 2022.
- [54] I. Masi, S. Rawls, G. Medioni, and P. Natarajan. Pose-aware face recognition in the wild. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4838–4846. IEEE,

- 2016.
- [55] N. A. McIntyre, H. Jarodzka, and R. M. Klassen. Capturing teacher priorities: Using real-world eye-tracking to investigate expert teacher priorities across two cultures. *Learning and Instruction*, 60:215–224, 2019.
 - [56] A. M. Möller, S. E. Baumgartner, R. Kühne, and J. Peter. The effects of social information on the enjoyment of online videos: An eye tracking study on the role of attention. *Media Psychology*, 24(2):214–235, 2021.
 - [57] S. Özdel, E. Bozkir, and E. Kasneci. Privacy-preserving scanpath comparison for pervasive eye tracking. *Proceedings of the ACM on Human-Computer Interaction*, 8(ETRA):1–28, 2024.
 - [58] C. Palmero Cantarino, O. V. Komogortsev, and S. S. Talathi. Benefits of temporal information for appearance-based gaze estimation. In *ACM Symposium on Eye Tracking Research and Applications*, pages 1–5. ACM, 2020.
 - [59] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
 - [60] R. P. Pieter Vansteenkiste, Greet Cardon and M. Lenoir. Measuring dwell time percentage from head-mounted eye-tracking data – comparison of a frame-by-frame and a fixation-by-fixation analysis. *Ergonomics*, 58(5):712–721, 2015.
 - [61] D. V. Poling, C. L. Van Loan, J. D. Garwood, S. Zhang, and D. Riddle. Enhancing teacher-student relationship quality: A narrative review of school-based interventions. *Educational Research Review*, 37:1–23, 2022.
 - [62] M. Pouta, E. Lehtinen, and T. Palonen. Student teachers’ and experienced teachers’ professional vision of students’ understanding of the rational number concept. *Educational Psychology Review*, 33:109–128, 2021.
 - [63] L. Prieto, K. Sharma, L. Kidzinski, M. Rodríguez-Triana, and P. Dillenbourg. Multimodal teaching analytics: Automated extraction of orchestration graphs from wearable sensor data. *Journal of Computer Assisted Learning*, 34(2):193–203, 2018.
 - [64] L. P. Prieto, K. Sharma, L. Kidzinski, and P. Dillenbourg. Orchestration load indicators and patterns: In-the-wild studies using mobile eye-tracking. *IEEE Transactions on Learning Technologies*, 11(2):216–229, 2018.
 - [65] R. Ranjan, S. De Mello, and J. Kautz. Light-weight head pose invariant gaze tracking. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, pages 2237–22378, 2018.
 - [66] T. Seidel, K. Schnitzler, C. Kosel, K. Stürmer, and D. Holzberger. Student characteristics in the eyes of teachers: Differences between novice and expert teachers in judgment accuracy, observed behavioral cues, and gaze. *Educational Psychology Review*, 33(1):69–89, 2021.
 - [67] S. Stojanov, S. S. Talathi, and A. Sharma. The benefits of depth information for head-mounted gaze estimation. In *2022 Symposium on Eye Tracking Research and Applications*, pages 1–7. ACM, 2022.
 - [68] A. R. Strohmaier, K. J. MacKay, A. Obersteiner, and K. M. Reiss. Eye-tracking methodology in mathematics education research: A systematic literature review. *Educational Studies in Mathematics*, 104(2):147–200, 2020.
 - [69] Ö. Sümer, P. Goldberg, S. D’Mello, P. Gerjets, U. Trautwein, and E. Kasneci. Multimodal engagement analysis from facial videos in the classroom. *IEEE Transactions on Affective Computing*, 14(2):1012–1027, 2023.
 - [70] Ö. Sümer, P. Goldberg, K. Stürmer, T. Seidel, P. Gerjets, U. Trautwein, and E. Kasneci. Teachers’ perception in the classroom. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, pages 1–10, 2018.
 - [71] Ö. Sümer, E. Bozkir, T. Kübler, S. Grüner, S. Utz, and E. Kasneci. Fakenewsperception: An eye movement dataset on the perceived believability of news stories. *Data in Brief*, 35:1–9, 2021.
 - [72] Ö. Sümer, P. Gerjets, U. Trautwein, and E. Kasneci. Automated anonymisation of visual and audio data in classroom studies. In *The Workshops of the AAAI Conference on Artificial Intelligence*, pages 1–7, 2020.
 - [73] Tobii AB. Tobii pro glasses 2 user manual v = 1.1.3, 2016. <https://www.tobiipro.com/siteassets/tobii-pro/user-manuals/tobii-pro-glasses-2-user-manual.pdf?v=1.1.3>.
 - [74] Tobii AB. Tobii pro glasses 2 user manual, 2023. Last accessed 24/10/2023.
 - [75] M.-J. Tsai, H.-T. Hou, M.-L. Lai, W.-Y. Liu, and F.-Y. Yang. Visual attention for solving multiple-choice science problem: An eye-tracking analysis. *Computers & Education*, 58(1):375–385, 2012.
 - [76] M. Vasu, N. Abreu, R. Vásquez, and C. López. Vehicle-counting with automatic region-of-interest and driving-trajectory detection. In *Workshops of International Conference on Machine Learning*, pages 1–5, 2021.
 - [77] M. Wang and W. Deng. Deep face recognition: A survey. *Neurocomputing*, 429:215–244, 2021.
 - [78] M. Wedel and R. Pieters. Eye tracking for visual marketing. *Foundations and Trends® in Marketing*, 1(4):231–320, 2008.
 - [79] Z. Yong, G. Su, X. Li, L. Sun, Z. Li, and L. Liu. Recognizing cognitive load by a hybrid spatio-temporal causal model from multivariate physiological data. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, page 318–334, 2022.
 - [80] Y. Yu and J.-M. Odobez. Unsupervised representation learning for gaze estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7312–7322, 2020.
 - [81] S. Zafeiriou, C. Zhang, and Z. Zhang. A survey on face detection in the wild: Past, present and future.

- Computer Vision and Image Understanding*, 138:1–24, 2015.
- [82] S. Zhang, X. Zhu, Z. Lei, H. Shi, X. Wang, and S. Z. Li. FaceBoxes: A CPU real-time face detector with high accuracy. In *2017 IEEE International Joint Conference on Biometrics (IJCB)*, pages 1–9. IEEE, 2017.
 - [83] J. Zhao, Y. Cheng, Y. Xu, L. Xiong, J. Li, F. Zhao, K. Jayashree, S. Pranata, S. Shen, J. Xing, S. Yan, and J. Feng. Towards pose invariant face recognition in the wild. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2207–2216. IEEE, 2018.
 - [84] Y. Zhu, H. Cai, S. Zhang, C. Wang, and Y. Xiong. Tinaface: Strong but simple baseline for face detection, 2021.
 - [85] Z. Zhu, G. Huang, J. Deng, Y. Ye, J. Huang, X. Chen, J. Zhu, T. Yang, J. Lu, D. Du, and J. Zhou. Webface260m: A benchmark unveiling the power of million-scale deep face recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10487–10497, 2021.