

3D Gaze Tracking for Studying Collaborative Interactions in Mixed-Reality Environments

Eduardo Davalos
eduardo.davalos.anaya@vanderbilt.edu
Vanderbilt University
Nashville, TN, USA

Yike Zhang
yike.zhang@vanderbilt.edu
Vanderbilt University
Nashville, TN, USA

Ashwin T. S.
ashwindixit9@gmail.com
Vanderbilt University
Nashville, TN, USA

Joyce H. Fonteles
joyce.h.fonteles@vanderbilt.edu
Vanderbilt University
Nashville, TN, USA

Umesh Timalsina
umesh.timalsina@vanderbilt.edu
Vanderbilt University
Nashville, TN, USA

Gautam Biswas
gautam.biswas@vanderbilt.edu
Vanderbilt University
Nashville, TN, USA




Figure 1: 3 sequential representations of a scene: monocular, 3D skeleton data, and full 3D reconstruction to perform gaze analysis.

ABSTRACT

This study presents a novel framework for 3D gaze tracking tailored for mixed-reality settings, aimed at enhancing joint attention and collaborative efforts in team-based scenarios. Conventional gaze tracking, often limited by monocular cameras and traditional eye-tracking apparatus, struggles with simultaneous data synchronization and analysis from multiple participants in group contexts. Our proposed framework leverages state-of-the-art computer vision and machine learning techniques to overcome these obstacles, enabling precise 3D gaze estimation without dependence on specialized hardware or complex data fusion. Utilizing facial recognition and deep learning, the framework achieves real-time, tracking of gaze patterns across several individuals, addressing common depth estimation errors, and ensuring spatial and identity consistency within the dataset. Empirical results demonstrate the accuracy and reliability of our method in group environments. This provides mechanisms for significant advances in behavior and interaction analysis in educational and professional training applications in dynamic and unstructured environments.

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

ICMIW'24, Nov 04–08, 2024, San José, Costa Rica

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-XXXX-X/18/06...\$15.00

<https://doi.org/XXXXXXX.XXXXXXX>

CCS CONCEPTS

• **Applied computing** → **Collaborative learning**; • **Computer systems organization** → *Real-time system architecture*; • **Human-centered computing** → **Collaborative interaction**.

KEYWORDS

eye-tracking, gaze, collaborative, joint attention, automated

ACM Reference Format:

Eduardo Davalos, Yike Zhang, Ashwin T. S., Joyce H. Fonteles, Umesh Timalsina, and Gautam Biswas. 2024. 3D Gaze Tracking for Studying Collaborative Interactions in Mixed-Reality Environments. In *Proceedings of 26th ACM International Conference on Multimodal Interaction Workshop (ICMIW'24)*. ACM, New York, NY, USA, 9 pages. <https://doi.org/XXXXXXX.XXXXXXX>

1 INTRODUCTION

Eye-tracking technology has significantly advanced our understanding of how humans perceive and process information, particularly in situations where cognitive processing and information acquisition are essential components for completing tasks. This technology is widely used in educational and professional training environments, offering insights into the intricate relationship between cognition and task performance [21].

In educational environments, eye-tracking technology informs researchers and educators about students' visual attention and engagement patterns associated with successful learning and comprehension. Therefore, insights gained from eye-tracking studies contribute to developing educational tools and methodologies that

are more engaging and effective in improving student learning outcomes. Additionally, a significant body of eye-tracking research examines the interplay between individual and collective gaze behaviors within collaborative problem-solving settings. Comprehending these interactions is instrumental in enhancing team-based strategies and communication techniques, as eye-tracking data elucidates the distribution of visual attention among team members and their reactions to shared visual signals during cooperative activities [22].

A lot of current eye-tracking research is predominantly confined to individual settings [9]. Collaborative and team-based studies present many challenges, such as the alignment of data across multiple devices, limited support from eye-tracking manufacturers, and the substantial costs associated with engineering and developing specialized hardware and software to support data collection and analyses. Most eye-tracking devices are tailored for single-user applications and do not easily generalize to studying multi-user interactions. Additionally, the development of new methods and technologies involves significant engineering expenses in terms of human hours and technical expertise. This issue is compounded by the traditionally poor scalability and high costs of eye-tracking equipment.

This research addresses the complexities of collaborative gaze analysis for children enacting scientific processes in mixed-reality environments [12]. Our framework enhances the conventional gaze-tracking approaches by incorporating 3D objects-of-interest (OOI) encoding. This extension progresses from the traditional 2D areas-of-interest (AOI) to a more dynamic 3D space, offering detailed insights into both the social and environmental context of gaze behavior. The key contributions of our work are:

- **Facial Recognition** We implemented face recognition to consistently track individuals as they move around in physical space, even when they occlude one another, thus enabling continuous monitoring of gaze patterns for multiple users.
- **3D Reconstruction and Gaze Reprojection** Our approach includes a sophisticated 3D reconstruction of the environment, allowing us to reproject gaze data accurately within this space. This method provides a realistic depiction of where individuals are looking, enhancing the accuracy of gaze tracking.
- **Social Network Analysis** We use the continuous and simultaneous gaze tracking of multiple individuals to perform social network analysis and visualize the social dynamics among the participating students.

Our framework enhances 3D reconstruction and robustly encodes OOIs in a scene, capturing dynamic user movements and static objects like screens and furniture for in-depth identity-aware gaze tracking, which is essential for deciphering social interactions and engagement by monitoring how participants view each other and interact with room objects [6]. These capabilities ensure precise identification and analysis of gaze targets in complex, noisy environments, such as classrooms. A detailed timeline of OOI gaze events in real-world scenarios demonstrates the framework's practicality. For transparency and to promote further development, we have made the entire codebase for this project open-source, available in «Github» repository.

2 BACKGROUND

The field of eye-tracking and gaze estimation has witnessed significant advances, evolving from hardware-centric approaches, like the ones developed by Tobii and Eyeware, to sophisticated computer vision techniques. The emergence of datasets like GazeFollow [20] has led to the development of algorithms capable of predicting gaze within complex scenes. Technologies like MIDAS [19] and transformer-based approaches [29] have further refined accuracy by integrating depth and object recognition, although challenges like depth noise and computational intensity remain. Pre-processing techniques have been instrumental in translating gaze data into behavioral insights, particularly within social network analysis, thereby enhancing our understanding of interaction dynamics across both digital and physical domains.

2.1 Eye-Tracking & Gaze Estimation

Eye-tracking, an elaborate 3D problem, has traditionally been addressed by eye-tracking manufacturers, such as Tobii, EyeWare, and EyeLink. They include on-screen bars, webcams, and glasses. These plug-and-play solutions have enriched the field of eye-tracking research, opened many new directions, and accelerated eye-tracking research. Within the educational domain, these tools and equipment have been successfully employed to use gaze tracking analysis to understand individuals' learning behaviors [10, 18, 26].

Eye-tracking devices may be limited in their applications, especially when the environment, such as a collaborative learning setting, has requirements beyond their initial design specifications. Although some studies have attempted to use these devices in group environments [13, 17, 31], they have not produced the same depth of analysis as individual-focused research. This shortfall is largely due to the complexities in integrating and synchronizing gaze data from multiple sensors. Such integration typically requires sophisticated engineering solutions that existing eye-tracking and multimodal software, such as iMotions^{TM1}, do not support. Consequently, conducting collaborative gaze studies often incurs significant development costs in hardware and software, which reduces the robustness and scalability of these approaches.

Recent developments in computer vision have produced innovative systems designed to address the challenges of 3D gaze estimation. Notable examples, Gaze360 [14], L2CS-Net [1], and MCGaze [11], represent significant advances in developing robust and scalable approaches to gaze tracking and analysis. These systems harness deep learning algorithms to estimate gaze vectors' pitch and yaw from standard monocular camera feeds. Such advances negate the need for specialized eye-tracking hardware and the need for complex data synchronization algorithms, facilitating the tracking of multiple users' gazes concurrently and making it easier to run education studies in classrooms. However, these gaze estimation methods only provide 3D gaze vector, lacking the necessary 3D gaze origin to determine visual focus within a complex 3D scene.

The field of eye-tracking has recently expanded to include innovative research in *gaze target detection*, an approach aimed at discerning the focal point of a person's gaze within a given scene. This technique involves estimating the gaze-fixed object and the gaze vector using deep learning (DL) methods. The GazeFollow

¹<https://imotions.com/>

dataset [20] has been instrumental in propelling this research area forward by providing a benchmark dataset for the training and evaluation of DL models tailored to this task. For the gaze target detection task, a single monocular image serves as the input. This image is analyzed to predict the exact point within the visual field where the gaze is directed. A pioneering approach by Tonini et al. [28] has furthered this technique by incorporating MIDAS relative depth estimation models. These models enhance the input monocular images with 3D scene information, creating a more comprehensive understanding of the space within the image.

Following the depth estimation, a series of transformer modules are applied to perform a sequence of tasks. Initially, these transformers identify and classify objects within the image. Subsequently, they generate field-of-vision (FOV) cones, which help in visualizing the possible areas within the scene that might be the focus of the gaze. Finally, the transformers link the gaze vector to these objects, effectively predicting the gaze target. However, this approach does come with trade-offs. Many of the gaze target estimation methods are not geared for temporal analysis; therefore, their output is not temporally consistent or congruent. While it bypasses the need for more complex camera setups to achieve a 3D understanding of the scene, it may suffer from high-depth noise. This noise can lead to inaccuracies in scenarios where precise depth information is crucial. Moreover, the computational expense of depth estimation methods needed for accurate 3D scene reconstruction can be substantial, potentially limiting the method's applicability in real-time or on-device applications. Despite these challenges, the advances in gaze target detection exemplify the dynamic nature of eye-tracking research and its potential to enhance our interaction with and understanding of complex visual environments.

2.2 Pre-processing Gaze

The literature on eye-tracking features various processing techniques that correlate gaze data with learning outcomes [18], cognitive workload estimation [16], and behaviors [3]. Eye-tracking devices, such as the EyeLink 2000, often produce data at rates exceeding 150 Hz, leading to the development of numerous pre-processing techniques. These techniques refine the data into more manageable representations and metrics that are easier to analyze and link to educational or training contexts. According to Srivastava et al. [25], eye-tracking features are categorized into three levels: LOW, which includes fixation and saccade metrics; MID, encompassing gaze radial direction and patterns; and HIGH, which involves encoding areas-of-interest (AOIs). These categories are utilized for various applications, including training machine learning algorithms, computing statistics, and generating visualizations.

Each category of eye-tracking features offers specific advantages, limitations, and demands for computational resources. LOW and MID level features do not require additional data and are less computationally intensive to process. Conversely, HIGH-level features, which connect gaze data directly with relevant visual elements within the environment, provide deeper insights but are more complex to compute. AOI encoding, for example, involves defining the geometry of an AOI – typically a rectangle or a circle – and determining whether a gaze point falls within it [8]. AOIs are further

classified as dynamic or static, with dynamic AOIs presenting significant challenges as they can move within the tracking session. Depending on the nature of the AOI – such as a person or an object like a television – and the specific eye-tracking system employed, tracking dynamic AOIs may require advanced deep-learning algorithms for object detection.

2.3 Collaboration & Gaze Tracking

Even with the technical challenges in tracking the eye movements of multiple individuals simultaneously, gaze analysis has become a significant component in understanding the interactive and social dynamics in collaborative learning using micro-gaze movements. Joint and mutual attention are important components of collaboration by aligning information acquisition and team dialog [30]. Social network analysis provides another methodology for measuring team performance by analyzing gaze behaviors [13].

In individualistic learning environments, eye-tracking systems typically interface with digital platforms such as computers or mobile devices, where AOIs are defined by on-screen elements [18]. The digital nature of these AOIs simplifies tracking, as the geometry of on-screen elements can be determined directly from the rendering data, circumventing the need for advanced artificial intelligence (AI) or computer vision techniques. Conversely, collaborative environments present a contrast, often featuring co-located teams interacting with both physical and digital elements. In such settings, tracking of physical AOIs requires sophisticated computer vision algorithms due to the absence of inherent digital geometries to reference.

Using computer vision methods for 3D gaze and gaze target estimation marks a progression in collaborative gaze analysis by enabling the capture of multiple gaze data through a single device. However, this approach has its limitations. Traditional 3D gaze estimation techniques primarily provide the 3D rotation, $R \in \mathbb{R}$, of gaze orientation but lack the 3D translation, t . Moreover, these methods often struggle with targets that move out of frame and face challenges related to video consistency and accurate depth estimation. To overcome these issues, we adopt a method involving 3D reconstruction to maintain spatial consistency across our gaze analysis. This approach ensures that the gaze analysis remains accurate and consistent. Without this additional 3D contextual information, such as the precise locations of the gaze vectors and the environmental layout, the 3D gaze data would typically be projected back onto the 2D image plane, leading to a significant loss of critical spatial information [4, 11, 27, 28].

3 PROPOSED METHOD

Our proposed method is illustrated in Figure 2. Our framework consists of two main sequential components: (1) a face recognition module and (2) a gaze analysis module. We start with a monocular image and perform face detection employing a finely-tuned Multi-task Cascaded Convolutional Network (MTCNN) model [32]. The detected faces are enclosed within bounding boxes, which are then tracked using a custom-designed tracking algorithm. These tracklets serve as input to a re-identification process that utilizes FaceNet [23] to create face embeddings and match them to a pre-established gallery of participant faces.

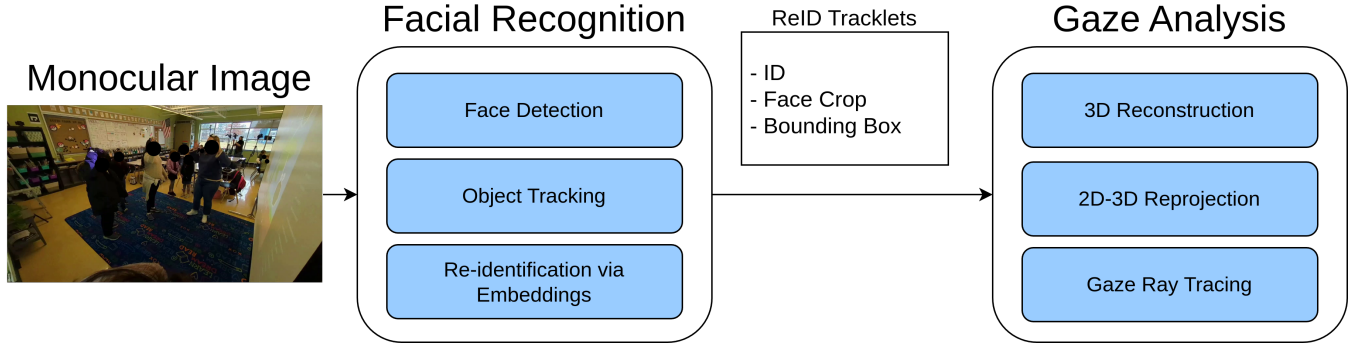


Figure 2: Framework Overview: composed of facial recognition and gaze analysis modules.

This automated detection, tracking, and re-identification process allows us to maintain consistent identification across frames, producing reliable bounding boxes for subsequent analysis. In the gaze analysis module, we perform 3D reconstruction by reprojecting the 2D detected data into 3D space. As a next step, we perform 3D gaze ray tracking, which enables us to determine the focus of each participant’s gaze and to encode this information in relation to OOIs. This comprehensive approach ensures accurate gaze tracking and analysis by integrating advanced face recognition with sophisticated gaze-tracking techniques, allowing for a deep understanding of participant engagement and focus.

3.1 Face Recognition Module

Using off-the-shelf computer-vision models that were fine-tuned to our setting (i.e., a classroom), we track the 2D position of participants (students, teachers, and researchers) in an identify-preserving fashion. The first step of this computational pipeline (see Figure 3) is the 2D bounding box detection of participants’ faces using a fine-tuned MTCNN in a frame-by-frame fashion. To reduce the need to perform a costly re-identification operation, we employed a tracking algorithm that leverages prior knowledge, i.e., bounding boxes of the same object should exhibit small changes from one frame to the next. To implement this, the tracking algorithm uses an Euclidean distance-based approach to match current and prior detections in a computationally efficient manner. Using this tracking strategy we were able to link a large percentage of detections together, except when complete or out-of-frame occlusions caused a failed tracking and a new tracklet ID was generated. For these cases, we had to develop a robust re-identification strategy that used face embeddings and a prior face gallery.

Our study utilized the DeepFace library and its FaceNet model to create vector embeddings from facial images [23, 24]. We established a reference set of face embeddings for each participant, known as anchor embeddings. By leveraging a pre-established set of frontal face images, our facial recognition has the benefit of only prioritizing and processing the tracked faces of consented participants. When the tracking algorithm lost continuity and initiated a new tracklet ID, the corresponding face crop was embedded and compared against the gallery using cosine similarity, adhering to a stringent identity-matching threshold. Unidentifiable face detections, often due to side profiles or distant captures, were excluded

from subsequent analysis to maintain data integrity. The facial recognition module’s output is a re-identified (ReID) tracklet that includes participant ID, bounding box coordinates, and the associated face crop, facilitating participant-specific analysis during the inference phase of the computer vision pipeline.

3.2 Gaze Analysis Module

The gaze module utilizes ReID tracklets for each frame to conduct gaze estimation, 3D reconstruction, 2D-3D reprojection, and OOI encoding (see Figure 4). Gaze estimation is achieved using L2CS-Net [1], which processes face crops to output 3D gaze vectors represented by pitch and yaw angles. These vectors yield a 3D rotation matrix for gaze orientation, but, by itself, it does not fully capture the 3D gaze information to properly place the gaze within the 3D scene. To overcome this, we employed ZoeDepth [2], a metric depth estimation model that provides consistent frame-by-frame depth estimation. By using a metric depth estimation model, we achieve a consistent and properly scaled depth in meters. Unlike other methods that infer depth based on facial or eyeball dimensions, ZoeDepth offers dense pixel-level depth predictions. These predictions facilitate 3D scene reconstruction by creating a mesh from the monocular image, which is updated per frame to accurately represent both dynamic and static elements within the classroom environment.

To effectively reproject our 2D face tracking analyses into 3D gaze tracking, we employ a 2D-3D reprojection method. For each identified feature, such as a bounding box surrounding a face, we project this information into 3D space using a transformation matrix \mathbf{M}_{face} . The matrix’s rotation component \mathbf{R} is an identity matrix, as rotation adjustment is not required for gaze ray tracing. The translation vector t is determined by the bounding box’s centroid $^{pix}\mathbf{P}_{face}$ and the depth z is calculated as follows:

$$t = \begin{bmatrix} X \\ Y \\ Z \end{bmatrix} = z \cdot \begin{bmatrix} \frac{x-c_x}{f_x} \\ \frac{y-c_y}{f_y} \\ 1 \end{bmatrix}$$

This process results in a 3D bounding box that is specific to the face, and a second bounding box for the full body, which is estimated based on fixed dimensions and orientation. We carefully manage potential mislabeling issues by manually marking the floor




Figure 3: Face Recognition Module: With the consequent use of detecting, tracking, and re-identifying faces, ReID tracklets are generated across a video.

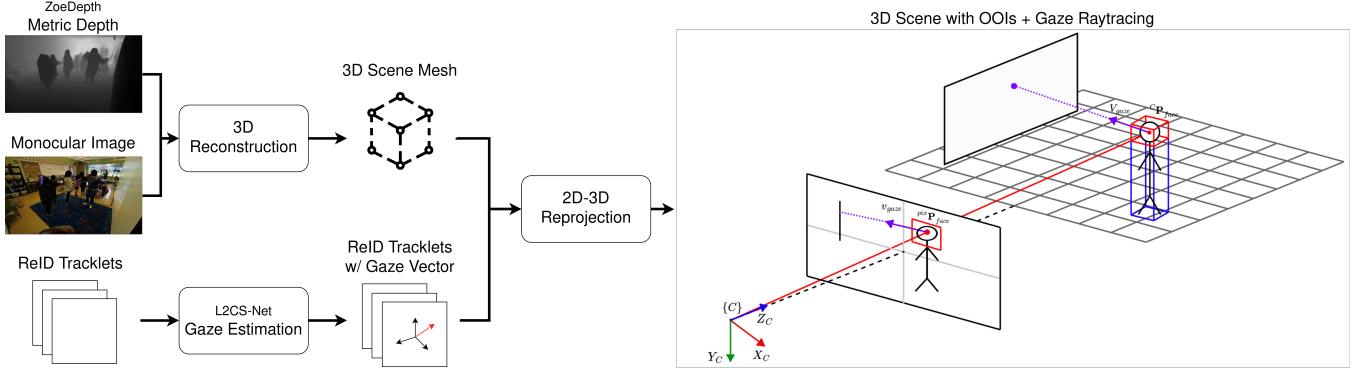


Figure 4: Gaze Analysis Module: Given a monocular image, the depth image is predicted, combined with the ReID tracklets, and then projected into a 3D scene.

plane once per session, especially when participant heights vary. Following this, the gaze vector, originally derived from the L2CS-Net, is transformed into a 3D vector using the established face bounding box as its origin and then converted into a rotation matrix to finalize the 3D placement, denoted as M_{gaze} .

For the 3D scene reconstruction, we manually place static objects such as displays, walls, and floors using a 3D annotation tool, Vision6D. With all essential objects placed, we conduct a comprehensive 3D gaze analysis through gaze ray tracking. Utilizing the transformation matrix M_{gaze} , we determine the object of focus by ray tracing the gaze vector and identifying the first object it intercepts, excluding the individual's own face and body. This process is facilitated by Trimesh ray tracing method, encoding each gaze fixation on dynamic participants or static objects within the scene. In more detail, the Trimesh library is a versatile tool that facilitates the loading and processing of triangular meshes, providing utilities for creating, editing, and analyzing 3D geometry in an efficient and straightforward manner. In the context of ray tracing with Trimesh, rays can be cast from a point in space through the mesh, and the library can calculate the points at which the rays intersect the mesh.

The ray tracing mechanism is mathematically represented by:

$$\mathbf{r}_{\text{transformed}}(t) = \mathbf{M}_{gaze} \begin{bmatrix} \mathbf{o} + t\mathbf{d} \\ 1 \end{bmatrix}$$

where M_{gaze} includes both rotation and translation components. To find the initial point of intersection (t_{\min}) with any mesh:

$$t_{\min} = \min \{ t \mid t > 0, \mathbf{r}_{\text{transformed}}(t) \cap \text{Mesh} \}.$$

Here, \cap denotes the intersection between the ray and a mesh, identifying the closest point of contact. This approach enables 3D-consistent tracking of gaze interactions within a complex 3D environment, enhancing our understanding of participant engagement and interaction dynamics.

4 DATASET

To test our approach in an authentic educational setting, we conducted an empirical study in a middle school classroom in the south-eastern United States. In the mixed-reality learning environment X, a team of (3-5) students participant co-participant in an embodied game where they are represented as 2D virtual avatars (e.g. water) [5]. Their position in the classroom is tracked via a Pozyx system and are superimposed into a top-view 2D arena. Their movements are mirrored by their avatars. By "touching" virtual objects (e.g. plants), scientific process like photosynthesis occur, resulting in changes to the student's avatars (e.g. oxygen). This interactive design grants students the autonomy to make choices regarding their actions and interactions, not only with the simulation but also with their peers, educators, researchers, and the audience of fellow students. Through these interactions, students' explore and learn new scientific concepts. The study was approved by the Institutional Review Board (IRB) at Y University.

The dynamic nature of the X environment yields a complex array of interactions encompassing physical movements, gestures, eye



Figure 5: Data Collection and Instrumentation in the X environment.

contact, emotional expressions, and verbal communication. It is important to collect comprehensive multimodal data to thoroughly understand students' learning behaviors and assess their advances in grasping scientific concepts and processes. This includes video recordings, audio captures, spatial tracking, and detailed logs from the simulation activities conducted within this mixed-reality setting.

Data Collection and Instrumentation. We employed ChimeraPy [7], a distributed multimodal data collection framework, to collect data from various devices and data streams. Figure 5 illustrates our equipment setup from a top-down perspective of the room. For data capture, we used Google Pixel 7a mobile phones to record standard monocular videos at a resolution of 1920x1080. Aruco markers were strategically placed within the room to facilitate 3D calibration and ensure consistent positioning across different cameras. Audio data was captured using wireless microphones and subsequently integrated via an audio interface. We also gathered system logs and screen recordings to enable a comprehensive analysis of students' learning behaviors. For this study, we focused on analyzing a single video to assess our methodology, with plans to incorporate a multi-camera setup in future expansions of the research. Our method processed the video frame-by-frame, allowing for detailed gaze analysis of each participant, and maintaining consistent identity tracking throughout the activity.

5 RESULTS

Recognition Error. As Figure 2 shows, the initial module of the method is facial recognition performed by automatic identify-persistent tracking but errors occurred. Failed re-identification occurs when the top-1 matches face embeddings.

Our facial recognition module showed a decrease in performance relative to the 99.2% accuracy benchmark established by FaceNet on the Labeled Faces in the Wild dataset. Notably, the module's accuracy in identifying distinct individuals within our experimental sessions varied. Recognition of the 'TEACHER' participant was commensurate with FaceNet's reported accuracy. Conversely, the accuracy markedly declined to around 80% for participants 'S1' and 'S2', who are children. This decrease is likely attributed to the challenges associated with accurately identifying younger facial

features, which can differ significantly from adult features that the model is predominantly trained on.

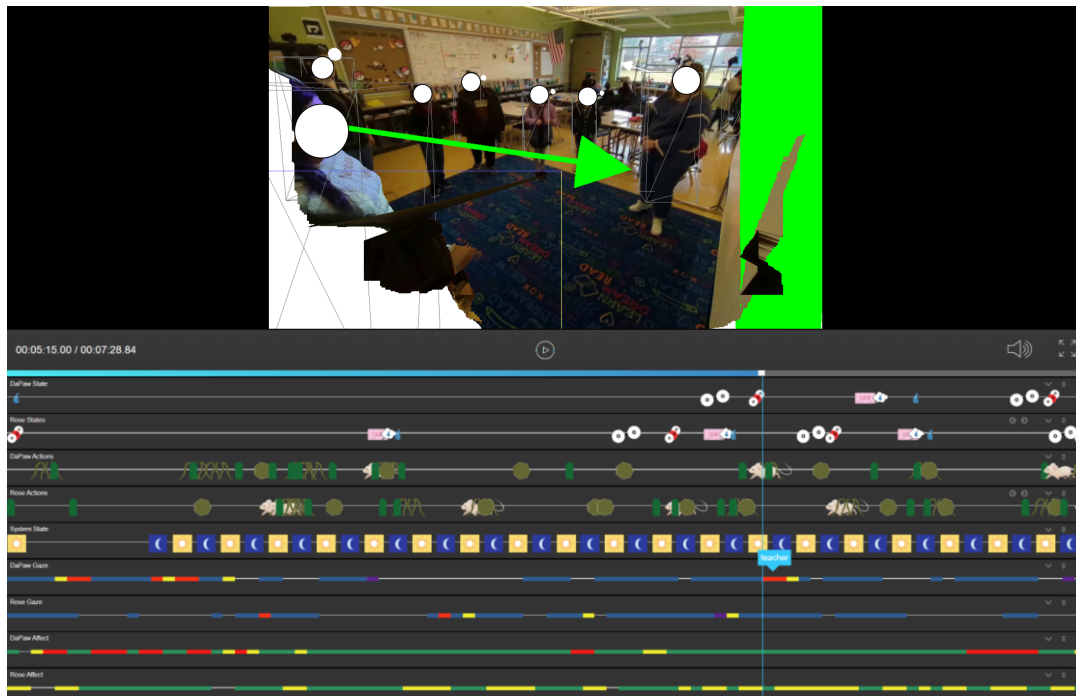
The performance was notably lower for 'S3', at just 44%. This decrease is part of a broader issue related to algorithmic bias in facial recognition technology [15]. For example, individuals many algorithms show lower accuracy in identifying faces with darker skin complexions. This issue arises because many facial recognition models, including the one we employed, were trained on datasets consisting of mostly white adult faces. This lack of diversity in training data leads to biased performance, where the model fails to generalize effectively across different skin tones and facial structures associated with different ethnicities.

This disparity in recognition accuracy underscores the urgent need for developing facial recognition models that are trained on a more diverse range of datasets. These datasets should encompass a wide variety of ages, skin tones, and facial features to ensure more equitable and effective performance across all user demographics. Such improvements in model training are essential to enhance the inclusivity and fairness of facial recognition technology.

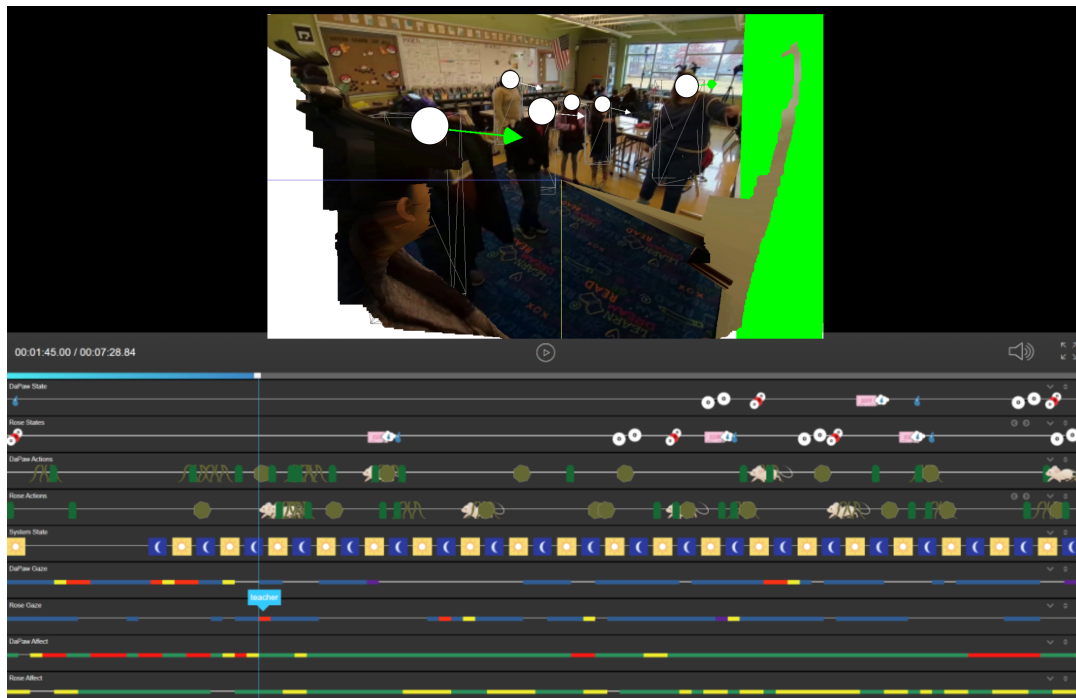
Timeline. We constructed a comprehensive timeline to represent the gaze data throughout the entire video, as depicted at the bottom of Figure 6. This timeline divides the gaze data into 5-second intervals, applying a median pooling strategy with a 2-second threshold to mitigate noise and improve clarity [12]. We leveraged our facial recognition system to generate individual timelines for each participant, distinctly illustrating their gaze directed at the teacher, other participants, and objects within the room. During the X game, our analysis showed that participants predominantly focused on the central display of the game, namely the screen. Our system adeptly captured this focal behavior despite considerable movement by the students and teacher, all without the necessity for calibration. Additionally, the timeline exposed gaps primarily resulting from obstructions in the participants' faces, as movements within the environment intermittently blocked the camera's line of sight to certain individuals.

In our study, we employed an interactive multimodal timeline user interface [12] to display videos augmented with gaze visualizations in conjunction with the timeline. This approach enables learning sciences researchers to explore the data informatics thoroughly, enabling them to gain a deeper understanding of the learning experiences of both students and teachers. The timeline incorporates data from the X-game state that includes students' gaze and affective states. Our methodology's frame-by-frame analysis yields granular temporal gaze data, facilitating the discernment of overarching trends and behaviors. This level of detail represents an advance over prior studies that primarily utilized heat maps and other aggregate visualization methods for human interpretation and analysis. By presenting the data in a continuous, temporal manner, our interface allows for the swift recognition of changes in patterns and salient events, which can be subsequently scrutinized by human evaluators or artificial intelligence systems.

We have identified two key events in the timeline that illustrate the pedagogical interactions between the students and their teacher. In Figure 6a, the teacher is seen engaging multiple students, but one student is focused on the game screen rather than the teacher. While this student may still be listening, the lack of mutual gaze



(a) Timeline Example A: Instance of a single student looking at the classroom projector, with the teacher conversing with all students.



(b) Timeline Example B: Teacher leading instruction by gazing toward the projector and students following her gaze and body language.

Figure 6: Examples of Multimodal Timelines with gaze visualizations.




Figure 7: Gaze Attention Network: Describing the observer-observed relationships via a directed graph. The edge weights reflect the total time fixated on the observed, with the node weights equalling the total incoming edge weights.

may hinder the student’s understanding, especially if the teacher is using body language (e.g., pointing and gestures) or spatial cues to explain science ideas. In another scenario depicted in Fig. 6b, the teacher stands by the projector, using hand gestures to highlight aspects of the game and directing attention to the screen with her gaze, indicating its importance in her instructions.

By incorporating body language and speech with gaze data, we capture the nuances of these instructional interactions and their effectiveness. The interactive timeline interface allows researchers and developers to better analyze and develop support systems that enhance teaching and learning experiences through multimodal data integration.

Network Analysis. By encoding objects of interest (OOIs), tracking participants, and relevant contextual objects, we used social network analysis with a directed graph. Nodes represent OOIs, sized by fixation duration, while edges illustrate gaze relationships, with widths indicating interaction duration. This method helped us construct an accumulative gaze attention network, shown in Figure 7. The ‘DISPLAY’ node, central to the game experience, was the most prominent due to significant fixation time. Student-to-student gaze was minimal and brief, but there was notable gaze directed at the teacher. Student ‘S1’ had the least fixation on both the ‘TEACHER’ and ‘DISPLAY’, suggesting a potential link between lower gaze fixation and reduced learning outcomes, warranting further exploration.

6 CONCLUSIONS

In this research, we introduce a method that integrates 3D reconstruction with gaze estimation for detailed gaze ray tracing. This supports scalable OOI encoding, crucial for contextualized gaze analysis in collaborative settings. Our approach uses advanced computer vision and deep learning to enhance traditional gaze

estimation by incorporating dynamic (users) and static (props) objects within a 3D space, offering a nuanced understanding of gaze dynamics in mixed-reality environments.

The methodology involves two steps: a facial recognition module using a fine-tuned MTCNN model for real-time face detection, an Euclidean-based tracking algorithm, and FaceNet for face re-identification; and a gaze analysis module performing 3D reconstruction with ZoeDepth and 3D gaze estimation using L2CS-Net. We use 3D gaze ray tracing to accurately determine gaze targets, improving our understanding of interactive dynamics in collaborative tasks.

Our approach offers significant improvements over traditional methods by allowing for the analysis of gaze interactions in a spatially aware context, providing insights into how participants in a study focus on and interact with different elements of their environment. This methodology not only advances the field of gaze analysis but also provides a practical framework for analyzing collaborative interactions in educational and professional settings. This research paves the way for novel applications of gaze analysis in intricate, real-world settings where deciphering the nexus between attention and interaction is vital.

Limitations. The framework’s accuracy suffers from the compounding of errors across multiple machine learning models (detection, tracking, re-identification, gaze estimation, and depth measurement). The accumulation of errors degrades the overall system performance, which is particularly noticeable in facial recognition where misidentification of participants can significantly skew results. To address these inaccuracies, manual corrections of re-identification errors were required, indicating a demand for more refined models that are customized for particular domains.

Depth estimation, while crucial for OOI encoding, introduces additional noise and demands high computational resources, achieving only 3 FPS on an RTX 3090 GPU without techniques to reduce computational complexity. This bottleneck hampers the system’s ability to operate in real-time and will mandate high computational resources for real-world deployment.

Future Work. To address these issues, we plan to refine the re-identification process by fine-tuning models specifically for educational domains, using classroom data to enhance the accuracy of participant recognition. Additionally, we aim to incorporate both monocular and stereo cameras to improve the accuracy and reliability of the 3D scene reconstruction. This would potentially allow for computationally efficient and real-time performance by optimizing the depth estimation model/technique. Techniques such as model quantization might be explored to speed up inference, although they typically trade off some accuracy. By implementing these improvements, we hope to minimize the computational load and increase the framework’s robustness and applicability in real-time educational settings.

7 ACKNOWLEDGMENTS

This work has been supported by the National Science Foundation AI Institute Grant No. DRL-2112635.

REFERENCES

- [1] Ahmed A. Abdelrahman, Thorsten Hempel, Aly Khalifa, Ayoub Al-Hamadi, and Laslo Dinges. 2023. L2CS-Net : Fine-Grained Gaze Estimation in Unconstrained Environments. In *2023 8th International Conference on Frontiers of Signal Processing (ICFSP)*. IEEE, Corfu, Greece, 98–102. <https://doi.org/10.1109/ICFSP59764.2023.10372944>
- [2] Shariq Farooq Bhat, Reiner Birkel, Diana Wofk, Peter Wonka, and Matthias Müller. 2023. ZoeDepth: Zero-shot Transfer by Combining Relative and Metric Depth. <https://doi.org/10.48550/ARXIV.2302.12288>
- [3] Xiuge Chen, Namrata Srivastava, Rajiv Jain, Jennifer Healey, and Tilman Dingler. 2023. Characteristics of Deep and Skim Reading on Smartphones vs. Desktop: A Comparative Study. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems* (<conf-loc>, <city>Hamburg</city>, <country>Germany</country>, </conf-loc>) (CHI '23). Association for Computing Machinery, New York, NY, USA, Article 516, 14 pages. <https://doi.org/10.1145/3544548.3581174>
- [4] Eunji Chong, Yongxin Wang, Nataniel Ruiz, and James M. Rehg. 2020. Detecting attended visual targets in video. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*. IEEE, Seattle, WA, USA, 5395–5405. <https://doi.org/10.1109/CVPR42600.2020.00544>
- [5] J Danish, Gabriella Anton, Nitasha Mathayias, Tessaly Jen, Morgan Vickery, Sarah Lee, Xintian Tu, Lana Cosic, Mengxi Zhou, Efrat Ayalon, et al. 2022. Designing for shifting learning activities. *The Journal of Applied Instructional Design* 11, 4 (2022), 169–185.
- [6] Joshua A Danish, Noel Enyedy, Asmalina Saleh, and Megan Humburg. 2020. Learning in embodied activity framework: A sociocultural framework for embodied cognition. *International Journal of Computer-Supported Collaborative Learning* 15 (2020), 49–87.
- [7] E. Davalos, U. Timalisina, Y. Zhang, J. Wu, J. Fonteles, and G. Biswas. 2023. ChimeraPy: A Scientific Distributed Streaming Framework for Real-time Multimodal Data Retrieval and Processing. In *2023 IEEE International Conference on Big Data (BigData)*. IEEE Computer Society, Los Alamitos, CA, USA, 201–206. <https://doi.org/10.1109/BigData59044.2023.10386382>
- [8] Eduardo Davalos, Caleb Vatrak, Clayton Cohn, Joyce Horn Fonteles, Gautam Biswas, Naveeduddin Mohammed, Madison Lee, and Daniel Levin. 2023. Identifying Gaze Behavior Evolution via Temporal Fully-Weighted Scanpath Graphs. In *LAK23: 13th International Learning Analytics and Knowledge Conference* (<conf-loc>, <city>Arlington</city>, <state>TX</state>, <country>USA</country>, </conf-loc>) (LAK2023). Association for Computing Machinery, New York, NY, USA, 476–487. <https://doi.org/10.1145/3576050.3576117>
- [9] Ruiqi Deng and Yifan Gao. 2023. A review of eye tracking research on video-based learning. *Education and Information Technologies* 28, 6 (2023), 7671–7702. <https://doi.org/10.1007/s10639-022-11486-7>
- [10] Paul Fester, Ali Shafit, Alex Harston, Michey Li, Pavel Orlov, and A. Aldo Faisal. 2022. MIDAS: Deep learning human action intention prediction from natural eye movement patterns. *arXiv:2201.09135* [cs.CV]
- [11] Yiran Guan, Zhuoguang Chen, Wenzheng Zeng, Zhiguo Cao, and Yang Xiao. 2023. End-to-End Video Gaze Estimation via Capturing Head-Face-Eye Spatial-Temporal Interaction Context. *IEEE Signal Processing Letters* 30 (2023), 1687–1691. <https://doi.org/10.1109/LSP.2023.3332569>
- [12] Joyce Horn Fonteles, Eduardo Davalos, TS Ashwin, Clayton Cohn, Yike Zhang, ..., and Gautam Biswas. 2024. A First Step in Using Machine Learning Methods to Enhance Interaction Analysis for Embodied Learning Environments. In *25th International Conference on AI in Education*. Association for Computing Machinery, Recife, Brazil.
- [13] Allison Jing, Kieran May, Gun Lee, and Mark Billinghurst. 2021. Eye See What You See: Exploring How Bi-Directional Augmented Reality Gaze Visualisation Influences Co-Located Symmetric Collaboration. *Frontiers in Virtual Reality* 2, June (2021), 1–17. <https://doi.org/10.3389/frvir.2021.697367>
- [14] Petr Kellnhofer, Adria Recasens, Simon Stent, Wojciech Matusik, and Antonio Torralba. 2019. Gaze360: Physically unconstrained gaze estimation in the wild. *Proceedings of the IEEE International Conference on Computer Vision* 2019-Octob (2019), 6911–6920. <https://doi.org/10.1109/ICCV.2019.00701>
- [15] Ashraf Khalil, Soha Glal Ahmed, Asad Masood Khattak, and Nabeel Al-Qirim. 2020. Investigating bias in facial analysis systems: A systematic review. *IEEE Access* 8 (2020), 130751–130761.
- [16] Emmanouil Ktistakis, Vasileios Skaramagkas, Dimitris Manousos, Nikolaos S. Tachos, Evanthis Tripoliti, Dimitrios I. Fotiadis, and Manolis Tsiknakis. 2022. COLET: A dataset for COgnitive workLoad estimation based on eye-tracking. *Computer Methods and Programs in Biomedicine* 224 (2022), 106989. <https://doi.org/10.1016/j.cmpb.2022.106989>
- [17] Joni Lämsä, Jimi Kotkajuuuri, Antti Lehtinen, Pekka Koskinen, Terhi Mäntylä, Jasmin Kilpeläinen, and Raija Hämmäläinen. 2022. The focus and timing of gaze matters: Investigating collaborative knowledge construction in a simulation-based environment by combined video and eye tracking. *Frontiers in Education* 7, October (2022), 1–16. <https://doi.org/10.3389/feeduc.2022.942224>
- [18] Ramkumar Rajendran, Anurag Kumar, Kelly E. Carter, Daniel T. Levin, and Gautam Biswas. 2018. Predicting learning by analyzing eye-gaze data of reading behavior. In *Proceedings of the 11th International Conference on Educational Data Mining, EDM 2018*. International Educational Data Mining Society, Buffalo, NY, USA, 30–32.
- [19] Rene Ranftl, Katrin Lasinger, David Hafner, Konrad Schindler, and Vladlen Koltun. 2022. Towards Robust Monocular Depth Estimation: Mixing Datasets for Zero-Shot Cross-Dataset Transfer. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 44, 3 (2022), 1623–1637. <https://doi.org/10.1109/TPAMI.2020.3019967>
- [20] Adria Recasens, Aditya Khosla, Carl Vondrick, and Antonio Torralba. 2015. Where are they looking?. In *Advances in Neural Information Processing Systems*, C. Cortes, N. Lawrence, D. Lee, M. Sugiyama, and R. Garnett (Eds.), Vol. 28. Curran Associates, Inc., Palais des Congrès de Montréal, Montréal CANADA. https://proceedings.neurips.cc/paper_files/paper/2015/file/ec8956637a99787bd197eacd77acce5e-Paper.pdf
- [21] Jonathan L Rosch and Jennifer J Vogel-Walcutt. 2013. A review of eye-tracking applications as tools for training. *Cognition, technology & work* 15 (2013), 313–327.
- [22] Bertrand Schneider, Kshitij Sharma, Sebastian Cuendet, Guillaume Zufferey, Pierre Dillenbourg, and Roy Pea. 2018. Leveraging mobile eye-trackers to capture joint visual attention in co-located collaborative learning groups. *International Journal of Computer-Supported Collaborative Learning* 13 (2018), 241–261.
- [23] Florian Schroff, Dmitry Kalenichenko, and James Philbin. 2015. FaceNet: A unified embedding for face recognition and clustering. *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition* 07-12-June (2015), 815–823. <https://doi.org/10.1109/CVPR.2015.7298682>
- [24] Sefik Ilkin Serengil and Alper Ozpinar. 2020. LightFace: A Hybrid Deep Face Recognition Framework. In *2020 Innovations in Intelligent Systems and Applications Conference (ASYU)*. IEEE, IEEE, Istanbul, Turkey, 23–27. <https://doi.org/10.1109/ASYU50717.2020.9259802>
- [25] Namrata Srivastava, Joshua Newn, and Eduardo Velloso. 2018. Combining Low and Mid-Level Gaze Features for Desktop Activity Recognition. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 2, 4, Article 189 (dec 2018), 27 pages. <https://doi.org/10.1145/3287067>
- [26] Robert H Tai, John F Loehr, and Frederick J Brigham. 2006. An exploration of the use of eye-gaze tracking to study problem-solving on standardized science assessments. *International journal of research & method in education* 29, 2 (2006), 185–208.
- [27] Francesco Tonini, Cigdem Beyan, and Elisa Ricci. 2022. Multimodal Across Domains Gaze Target Detection. *ACM International Conference Proceeding Series* 1, 1 (2022), 420–431. <https://doi.org/10.1145/3536221.3556624>
- [28] F. Tonini, N. Dall'Asen, C. Beyan, and E. Ricci. 2023. Object-aware Gaze Target Detection. In *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*. IEEE Computer Society, Los Alamitos, CA, USA, 21803–21812. <https://doi.org/10.1109/ICCV51070.2023.01998>
- [29] Danyang Tu, Xiongkuo Min, Huiyu Duan, Guodong Guo, Guangtao Zhai, and Wei Shen. 2022. End-to-End Human-Gaze-Target Detection with Transformers. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, New Orleans, LA, USA, 2192–2200. <https://doi.org/10.1109/CVPR52688.2022.00224>
- [30] Caleb Vatrak, Gautam Biswas, Clayton Cohn, Eduardo Davalos, and Naveeduddin Mohammed. 2022. Using the DiCoT framework for integrated multimodal analysis in mixed-reality training environments. *Frontiers in Artificial Intelligence* 5 (2022), 1–30. <https://doi.org/10.3389/frai.2022.941825>
- [31] Maureen M. Villamor and Ma Mercedes T. Rodrigo. 2019. Gaze collaboration patterns of successful and unsuccessful programming pairs using cross-recurrence quantification analysis. *Research and Practice in Technology Enhanced Learning* 14, 1 (2019), 25. <https://doi.org/10.1186/s41039-019-0118-z>
- [32] Kaipeng Zhang, Zhanpeng Zhang, Zhifeng Li, and Yu Qiao. 2016. Joint Face Detection and Alignment Using Multitask Cascaded Convolutional Networks. *IEEE Signal Processing Letters* 23, 10 (2016), 1499–1503. <https://doi.org/10.1109/LSP.2016.2603342>

Received 10 May 2024; revised 1 June 2024; accepted 1 July 2024