

No Title Given

Ahmad Yehia^{1,*}, Abdualлах Mohamed^{2,*†}, Tianyi Wang¹, Jiseop Byeon¹,
Kun Qian^{3,‡}, Junfeng Jiao¹, and Christian Claudel^{1,*}

¹ The University of Texas at Austin, Austin, TX 78712, USA

² AIDACHip Inc.

³ Unity AI Technologies

{ahmad.yehia, bonny.wang, jsbyeon, christian.claudel}@utexas.edu
jjiao@austin.utexas.edu, kun.qian@unity3d.com, abdu@aidachip.com

Abstract. Accurately forecasting human trajectories from an egocentric perspective plays a central role in applications such as humanoid robotics, wearable sensing systems, and assistive navigation. However, progress in this direction remains limited due to the scarcity of egocentric trajectory datasets collected in real-world environments. Addressing this need, we introduce EgoTraj, an egocentric multimodal open dataset recorded using Meta Quest Pro (MQPro). EgoTraj contains 75 sequences of human navigation collected from multiple MQPro wearers in real-world urban environments. Each recording provides synchronized RGB video along with ground-truth data, including 6-degree-of-freedom head poses, per-frame 3D gaze vectors, and scene annotations. To the best of our knowledge, EgoTraj differs from typical egocentric trajectory datasets by capturing long-horizon, self-chosen pedestrian navigation on consumer AR headset across urban routes with broad participant diversity. We benchmark several state-of-the-art trajectory prediction models and ablate the contributions of gaze, scene, and motion cues. The results highlight the utility of EgoTraj for AR-based perception, navigation, and assistive systems. The EgoTraj dataset, code, and *EgoViz Dashboard* are publicly available at <https://github.com/yehiaahmad/EgoTraj>.

Keywords: Egocentric Trajectory Prediction · Multimodal Dataset · Augmented Reality · Pedestrian Navigation · Gaze Tracking

1 Introduction

Egocentric perception forms the foundation of human navigation behavior [45]. Humans naturally perceive and interpret their surroundings from an egocentric viewpoint to guide decisions about subsequent motions [13]. Replicating this effortless human capability in machines, however, remains a long-standing challenge in embodied AI and computer vision [23]. Learning such perceptual cues

* These authors contributed equally to this work.

† Work done in a personal capacity, outside of AIDACHip Inc. and Unity AI.

‡ Corresponding author: Christian Claudel.



Fig. 1: Overview of EgoTraj. (a) **Protocol design:** the Meta Quest Pro headset records synchronized 6DoF head pose, gaze signals, and egocentric RGB video during real-world outdoor navigation. (b) **Multimodal EgoTraj capture:** representative egocentric frames from a crosswalk navigation scenario with multiple social interactions. (c) **Analysis and applications:** the dataset supports downstream tasks such as assistive navigation for blind and visually impaired users.

plays a pivotal role in forecasting the trajectory of humanoid robots [17, 26, 30], AR-VR systems [4], and autonomous vehicles [10, 11, 46]. At the same time, recent advances in egocentric vision and wearable sensing technologies have facilitated the development of assistive wayfinding systems, including blind navigation [38], object echo-location [37], and personalized object recognition [15].

As augmented reality (AR) technologies evolve, assistive systems increasingly leverage egocentric observations to interpret the user’s surroundings and provide context-aware guidance for subsequent actions. For instance, by modeling navigation behaviors learned from sighted individuals, forecasting the short-term trajectory can help blind or visually impaired users navigate more independently and safely in social spaces [31]. Moreover, using synchronized multimodal sensors, predicted trajectories can further be conveyed to the headset users through visual overlays [45], haptic feedback [39], or auditory instructions [6]. Beyond assistive applications, insights from human navigation behavior provide valuable guidance for social-aware robotic navigation. Humans naturally adhere to implicit social conventions when navigating crowded environments, such as maintaining interpersonal distance, yielding in constrained spaces, and adjusting motion based on gaze and attention cues [14]. By modeling these behaviors, mobile robots [41] and humanoid robots [31] can learn navigation strategies that are not only collision-free but also socially compliant, enabling safer operation in human-centered environments.

Existing human trajectory prediction methods have made remarkable progress; however, this progress largely relies on datasets captured from bird’s-eye viewpoints (BEV) or static cameras [1, 24]. Such datasets capture only externally observable motion and fail to model how humans perceive, plan, and initiate their

motion from a first-person perspective. Consequently, most trajectory forecasting methods [1, 35] rely primarily on motion history and scene context rather than intention-related cues. In contrast, egocentric vision provides direct access to how humans perceive their surroundings and navigate through them. Over the last two decades, egocentric vision research has produced large-scale datasets spanning diverse activities and environments [7, 8, 20, 27]. However, these datasets are primarily designed for action recognition and video understanding, with limited emphasis on forecasting the wearer’s trajectory for navigation. As a result, several works have introduced egocentric datasets for trajectory prediction [26, 32]. Nevertheless, many of these egocentric trajectory datasets remain limited in scale, gathered from a single participant, restricted to indoor environments, or lacking synchronized gaze data necessary for intention-aware trajectory forecasting. Critically, only a few leverage AR-style sensing devices, which introduce realistic challenges such as rapid ego-motion, viewpoint changes, unstable visual observations, and low-power sensing hardware.

Motivated by the lack of real-world egocentric trajectory datasets, we introduce EgoTraj, a multimodal dataset designed to advance research in egocentric trajectory forecasting and assistive AR navigation. To the best of our knowledge, EgoTraj is the first egocentric outdoor trajectory dataset to jointly provide synchronized 6DoF head pose, per-frame 3D gaze vectors, egocentric RGB video, and scene annotations, collected from various participants navigating self-chosen routes across sidewalks, crosswalks, and busy streets.

The main contributions of this work are summarized as follows:

- **EgoTraj Dataset.** We introduce an outdoor multimodal dataset for egocentric trajectory forecasting collected by Meta Quest Pro (MQPro) wearers navigating in real-world traffic environments, containing synchronized 6DoF head pose, gaze tracking data, egocentric RGB video, and scene annotations.
- **Benchmark and Evaluation Protocol.** We establish a standardized benchmark for egocentric trajectory forecasting by evaluating several state-of-the-art prediction models for both qualitative and quantitative analyses.
- **Multimodal Analysis.** Through ablation studies, we analyze the impact of gaze, social interactions, and scene context on trajectory prediction under realistic egocentric conditions.

2 Related Work

2.1 Non-Egocentric Human Trajectory Datasets

Most human trajectory prediction studies rely on a third-person perspective, where scenes are captured either from a BEV or stationary cameras. In particular, the ETH [29] and UCY [19] datasets have become the dominant benchmarks for trajectory prediction, with a total of 1,536 tracked pedestrians covering trajectory patterns such as people walking together, crossing paths, and forming or separating into groups. However, these benchmarks remain limited in scene

Table 1: Comparison of egocentric trajectory datasets across recording settings, scale, and sensing modalities. [†] Uses GPS-only localization. [‡] Reconstructs trajectories offline using structure-from-motion. * Captures activity-level trajectories rather than pedestrian navigation.

Dataset	Year	Setting	Hours	Frames	Subj.	Device	Modalities			Scene Ann.
							Gaze	IMU	6DoF	
KrishnaCam [†] [36]	2016	Outdoor	70.0	7.6M	1	Google Glass	✗	✓	✗	✗
EgoMotion [‡] [28]	2016	In+Out	9.1	65.5K	N/P	GoPro Stereo	✗	✗	✗	✗
FPL* [43]	2018	Outdoor	4.5	162K	N/P	Chest Cam	✗	✗	✗	✗
EgoNav [42]	2024	In+Out	3.3	237.6K	N/P	RealSense	✗	✓	✓	✗
LookOut [26]	2025	In+Out	4.0	288K	N/P	Aria	✓	✓	✓	✗
EgoCogNav [32]	2025	In+Out	6.0	432K	17	Tobii + Aria	✓	✓	✓	✗
EgoTraj (ours)	2026	Outdoor	10.7	1.15M	75	Quest Pro	✓	✓	✓	✓

diversity and restricted to pedestrian-only trajectories. To address these limitations, the Stanford Drone Dataset [34] and inD [2] expanded dataset scale using drone footage to capture the movements of multiple road users, including pedestrians, bicyclists, skateboarders, and vehicles. In parallel, the autonomous driving domain introduced large-scale ego-vehicle datasets such as nuScenes [3], Waymo Open [5], and JRDB [22]. Despite their scale and diversity, these datasets remain fundamentally third-person in nature. While they accurately capture where pedestrians move, they do not model how individuals perceive the scene or how gaze and visual attention shape future actions. This limitation motivates the development of egocentric trajectory datasets that integrate motion with perceptual cues, such as gaze and head pose, for intent-aware AR applications.

2.2 Egocentric Human Trajectory Datasets

Recent advances in wearable sensors have enabled data collection from a first-person perspective, accelerating research in egocentric vision [20, 27]. However, a limited subset of egocentric datasets explicitly addresses pedestrian navigation. KrishnaCam [36] provides egocentric video streams with GPS position, acceleration, and body orientation collected by one participant, while Yagi *et al.* [43] record first-person videos from a chest-mounted camera to predict future 2D locations of nearby pedestrians. However, neither dataset included gaze or ground-truth 6DoF pose of the camera wearer, limiting their behavioral diversity and applications for gaze-informed trajectory prediction. Park *et al.*’s EgoMotion dataset [28] pioneers egocentric future localization using GoPro stereo cameras across 26 indoor and outdoor scenes. However, trajectories are reconstructed offline via structure from motion rather than tracked in real time, and the dataset includes neither gaze nor pose telemetry. These limitations have motivated subsequent efforts that incorporate AR glasses and additional sensing modalities.

Project Aria datasets such as the Aria Digital Twin (ADT) [27], Aria Everyday Activities (AEA) [20], and Nymeria [21] paired egocentric video with 6DoF pose and gaze across various activities. However, ADT was restricted to indoor

activities, AEA focused on single-human scenarios, and Nymeria captured multiple actors only in collaborative settings. LookOut [26] provided 6D head pose trajectories with gaze in real-world pedestrian scenarios, yet was limited to 4 hours of data and lacked scene annotations. The dataset most closely related to our work is the Cognition-Aware Egocentric Navigation (CEN) dataset from EgoCogNav [32], which paired gaze with walking trajectories across 6 hours of real-world navigation. However, CEN is designed for cognition-aware modeling and perceived navigational uncertainty. In addition, the 6-hour recordings are not exclusive to outdoor environments. EgoTraj addresses a key gap by providing synchronized 6DoF head pose, gaze tracking, and egocentric RGB video augmented with scene annotations, collected from 75 participants navigating self-chosen routes in real urban environments. EgoTraj enables the study of gaze-informed egocentric trajectory prediction in naturalistic pedestrian environments. The detailed comparison of the existing egocentric trajectory datasets and the proposed EgoTraj dataset is shown in Table 1.

3 Dataset Description

3.1 Dataset Collection

Hardware. Each participant wears an MQPro headset operating in full-color passthrough mode. MQPro offers a non-intrusive, cost-effective, and readily deployable alternative to prior works that rely on custom-built sensor suites [25, 42] or teleoperated robotic platforms [16]. The headset integrates a passthrough RGB camera, two infrared eye-tracking cameras, four inside-out tracking cameras, and a 6-axis IMU, as illustrated in Fig. 1 (a). A custom Unity application interfaces with the built-in visual-inertial SLAM system and records time-synchronized data at 30 Hz, including 6DoF head pose, 3D gaze origin and direction vectors, and egocentric RGB video. During each session, participants use an MQPro controller to start and terminate data recording.

Recording Protocol. Before each session, participants completed a brief orientation covering the study objectives, experimental procedure, headset fitting, and eye-tracking calibration. During data collection, participants navigated between two of seven predefined outdoor waypoints distributed across urban areas, including sidewalks, crosswalks, and busy streets (Fig. 1 (a)). In contrast to fully scripted protocols, participants were free to choose their own routes between waypoints, enabling naturalistic route-finding behavior, spontaneous interactions with surrounding traffic conditions and realistic decision-making during navigation. Each session was capped at 15 minutes (8 min on average) to promote diverse navigation patterns across participants. Participants were instructed to obey traffic rules, check for passing vehicles and traffic signals before crossing streets, to avoid indoor shortcuts, and terminate the session immediately if they experienced any discomfort. A researcher followed each participant

at a safe distance to monitor the session and verify waypoint completion without influencing navigation behavior. Importantly, the seven waypoints denote landmarks rather than pinned GPS points, yielding 21 origin–destination pairs, and 31 of the 75 participants were unfamiliar with the recording area. To verify that participants did not follow templated routes between the same landmarks, we computed pairwise Dynamic Time Warping (DTW) distances within each of the 21 waypoint pairs ($n = 103$ within-pair comparisons), obtaining a median of 122.8 m with a range of [19.3, 379.4] m. This wide spread indicates that walkers chose meaningfully different paths between identical endpoint pairs rather than retracing a scripted route.

Participants. We recruited $N = 75$ volunteers to capture diverse navigation behaviors within the dataset. Eligibility criteria required adults with normal or corrected-to-normal vision and the ability to walk outdoors unassisted. Each participant contributed exactly one recording session, and no individual appears more than once in the dataset. Participants received monetary compensation upon completion. The cohort was balanced across gender to reduce demographic bias, with participants ranging from 18 to 38 years old and representing 14 nationalities (Fig. 2 (b)).

Privacy and Ethics. This study was approved by an Institutional Review Board (IRB). All volunteers provided written informed consent prior to data collection and retained the right to withdraw from the study, request deletion of their data, or redact specific portions of their recordings. To satisfy de-identification requirements for personally identifiable information (PII), all recordings were processed using EgoBlur [33] to automatically detect and blur faces and vehicle license plates. Session identifiers consist of anonymized date–time codes with no linkage to personal information, and raw unprocessed videos are not publicly distributed.

3.2 Data Processing

To prepare the public release, we implemented a multi-stage processing pipeline. Raw telemetry from the MQPro is recorded as a single CSV file per session at 30 Hz, containing 6DoF head pose (position, orientation quaternion, linear velocity, and angular velocity) together with binocular gaze signals (3D origin and direction vectors). Egocentric RGB video is simultaneously recorded as H.264-encoded MP4 segments at 30 fps. Because certain sensor streams (e.g., IMU) operate at higher frequencies than the RGB camera, all telemetry signals are resampled to align with the 30 Hz video timeline. Synchronized ground truth is obtained by matching each RGB frame to the closest telemetry timestamp. The finalized dataset is packaged as a single HDF5 file per session with three groups: *pose* (timestamp, position, rotation, linear velocity, angular velocity), *gaze* (origin, direction), and *video* (segment index, frame index). This processed format enables efficient frame-level indexing and direct retrieval of each RGB




Fig. 2: Dataset statistics of EgoTraj participants. (a) Nationality breakdown of the 75 recruited participants. (b) Age and gender composition across participants. (c) Walking speed distribution observed during navigation sessions.

frame with its corresponding motion and gaze data. The processed HDF5 files with privacy-blurred RGB frames are included in the public release.

3.3 Dataset Statistics

EgoTraj comprises 75 recording sessions totaling 10.7 accumulated recording hours of synchronized data with 1.15M RGB frames (Tab. 1). Across all sessions, participants walked a cumulative distance of 46.73 km in diverse urban environments, with individual session durations ranging from 5 to 15 minutes. Recordings were collected during both moderate and high pedestrian traffic conditions. Figure 2 (c) shows the per-session speed distribution sorted by median speed. The dataset-wide mean walking speed is 1.25 m/s, consistent with typical pedestrian walking speeds (1.2–1.4 m/s). Lower-speed segments generally correspond to congested areas such as busy intersections or narrow walkways, while higher-speed sessions occur along straight sidewalks or during hurry-up segments before traffic signals change. Across the dataset, 7% of frames exhibit near-stationary motion (<0.3 m/s), capturing stop-and-go pedestrian dynamics.

3.4 Scene Annotation

To augment EgoTraj with high-level contextual cues, we generated structured egocentric scene descriptions using a vision-language model (VLM) pipeline. The annotations target navigation-relevant elements in urban walking environments, including surrounding context, traffic activity, gaze fixation targets, and inferred short-term movement intent. These language descriptions complement the trajectory data by capturing behavioral context, such as following pedestrian flow




Fig. 3: Examples of VLM-generated scene annotations. Each tab shows an egocentric frame with the gaze marker (red dot) alongside the structured annotation produced by Qwen2.5-VL-32B.

or approaching an intersection, while remaining consistent with the underlying multimodal signals. From 253 privacy-blurred video segments, frames were sampled at 1 fps, yielding 38,606 annotated frames. Each frame was processed using Qwen2.5-VL-32B-Instruct [40]. An example-driven prompt enforces a consistent egocentric narration style that prioritizes interaction and navigation cues over generic scene descriptions (Fig. 3). Annotation reliability was evaluated in two stages. First, structural compliance was measured on 100 stratified frames sampled across sessions. Compliance required adherence to the prescribed tag structure and inclusion of an explicit navigation decision based on corresponding motion data. Using chain-of-thoughts (CoT), after up to two prompt retries for non-conforming outputs, the pipeline achieved 96% compliance. Second, two annotators independently reviewed the same sample; we report per-field Cohen’s κ (0.83–0.96) and VLM accuracy against the human reference in Appendix D.2, along with a failure-mode audit. Final annotations were stored as JSON sidecar files indexed by session and frame identifier alongside the HDF5 dataset.

3.5 Gaze Calibration

EgoTraj records per-frame 3D gaze origin and direction vector, capturing where the participant is looking relative to the walking scene. Prior work in visual neuroscience has established that gaze anticipates locomotor actions by 1–2 seconds [18]. For example, pedestrians often fixate on intended turning points, obstacles, or crossing opportunities well before executing the corresponding maneuver. To associate 3D gaze geometry with the egocentric RGB stream, we fit a per-session quadratic calibration model that maps gaze yaw–pitch angles to pixel coordinates (u, v) in the video frame. This calibration enables predictive algorithms to directly project gaze fixation points into image space, as illustrated in Fig. 4. Combined with the scene annotations (Sec. 3.4), the calibrated gaze stream supports joint reasoning about what pedestrians see and where they look, supporting intent-aware trajectory prediction for real world applications.



Fig. 4: Gaze-to-pixel calibration examples. The green dot indicates the projected gaze fixation overlaid on egocentric RGB frames confirming accurate alignment between the collected 3D gaze stream and the egocentric video plane.

3.6 Data Visualization and Quality Control

To ensure synchronization consistency and annotation integrity across modalities, we developed *EgoViz Dashboard*, an interactive inspection interface for the EgoTraj dataset (Fig. 5). The dashboard synchronizes four complementary views: (i) a zoomable 2D trajectory plot showing 30 Hz pose samples with the corresponding gaze-direction vector, (ii) a BEV of the full-session path, (iii) the aligned egocentric RGB frame, and (iv) the associated VLM-generated scene annotation. Users can navigate recordings by timestamp or frame index, enabling frame-level verification of pose-gaze-video alignment. The dashboard also displays session metadata including timestamp, 3D position, walking speed, and frame index to facilitate detailed inspection. During data collection, the dashboard was used to detect synchronization drift, validate gaze reprojection, and identify sessions requiring reprocessing. The *EgoViz Dashboard* will be released alongside the dataset to support data exploration and reproducibility.

4 Benchmarking

To validate EgoTraj as a benchmark for egocentric trajectory prediction, we evaluated several state-of-the-art trajectory prediction baselines, and performed ablation studies to quantify the predictive value of gaze and scene context. All evaluations were conducted on a held-out test split with sessions unseen during training. The results demonstrated that EgoTraj provides a challenging benchmark for egocentric navigation tasks and multimodal trajectory prediction.

4.1 Quantitative Evaluation

Metrics. For trajectory prediction, we reported Average Displacement Error (ADE) and Final Displacement Error (FDE). ADE measures the mean Euclidean distance between predicted and ground-truth positions over all predicted timesteps, while FDE measures the Euclidean distance at the final predicted




Fig. 5: A snapshot of the EgoViz Dashboard showing synchronized trajectory, gaze, video, and annotation streams. **Left:** the 30 Hz MQPro trajectory (green points) with the instantaneous gaze direction vector (white arrow) on a zoomable 2D map. **Top right:** a BEV of the full session path with the current position highlighted. **Middle right:** the corresponding egocentric RGB frame with projected gaze (red dot). **Bottom right:** VLM-generated scene description. The bottom status bar reports timestamp, 3D position, speed, and frame index for the selected sample.

timestep. For head orientation, we calculated the L1 rotation error [26], computed as

$$\mathcal{L}_{\text{rot}} = \frac{1}{T_{\text{pred}}} \sum_{i=1}^{T_{\text{pred}}} \|\hat{R}_{t+i} R_{t+i}^{\top} - I\|_1,$$

where $\hat{R}$ and R denote the predicted and ground-truth rotation matrices, respectively.

Baselines. Egocentric trajectory prediction remains relatively underexplored, with few directly comparable baselines. While single-agent trajectory models such as Social-LSTM [1], and TUTR [35] exist, they are trained on BEV surveillance datasets and are not compatible with egocentric 6DoF pose sequences. In addition, the closest related egocentric trajectory prediction works include EgoNav [42], EgoCast [4], LookOut [26], and EgoCogNav [32]. However, EgoNav, LookOut, and EgoCogNav are currently available only as arXiv preprints without publicly released code. Therefore, we adapted several representative methods that took multiple modalities simultaneously as their inputs:

- **Constant Velocity (Const_Vel)** [26], which extrapolates future body-frame translation and rotation using linear and angular velocities estimated from the last observed steps;

Table 2: Baseline comparison on the EgoTraj test split. ADE and FDE are reported in meters; $L_{1\text{head}}$ denotes the mean L_1 quaternion distance for head rotation prediction. Best results are in **bold**.

Model	ADE ↓	FDE ↓	$L_{1\text{head}}$ ↓
Const_Vel [26]	0.24	0.35	0.82
Lin_Ext [26]	0.26	0.39	1.39
M_Transformer [26]	0.20	0.32	0.74
CXA-Transformer [31]	0.19	0.29	0.69
EgoCast [4]	0.16	0.28	0.78

- **Linear Extrapolation**(Lin_Ext) [26], which fits a per-axis linear regression model to past translation and rotation observations and projects the motion sequence into the future;
- **Multimodal Transformer**(M_Transformer) [26], a standard Transformer baseline that takes the same inputs as the other learned baselines but performs early fusion by concatenating modality embeddings, followed by a temporal Transformer decoder and a linear head that predicts future trajectory;
- **CXA-Transformer** [31], a multi-stream cascaded cross-attention architecture with additional social and scene context streams. We adapt it by feeding the same ego-translation and rotation streams as M_Transformer, but fusing them via cascaded cross-attention rather than simple concatenation, and
- **EgoCast** [4], where we adapt its Transformer-based forecasting module, originally designed for full-body pose prediction, to take past ego-translation and head rotation as input and directly predict future trajectory, removing the pose-estimation stage.

All methods used an observation window of $T_{\text{obs}} = 1.5\text{ s}$ to predict a future horizon of $T_{\text{pred}} = 3.5\text{ s}$, and were trained and evaluated using the same data split. All examined approaches were trained and evaluated using an 80/10/10 split of the EgoTraj sessions into training, validation, and test sets, respectively.

Results. Table 2 reports quantitative results on ego-translation and rotation data modeling using the held-out EgoTraj test split. Among all evaluated methods, the adapted EgoCast model achieves the best trajectory forecasting performance, obtaining the lowest ADE (0.16 m) and FDE (0.28 m). In contrast, CXA-Transformer achieves the lowest head rotation error (0.69), indicating that its cascaded cross-attention mechanism better captures head orientation dynamics. Classical kinematic baselines such as Constant Velocity and Linear Extrapolation perform reasonably for short-term translation but exhibit substantially higher rotation errors. In particular, Linear Extrapolation yields the worst rotation accuracy (1.39 m), suggesting that simple per-axis regression fails to capture the nonlinear nature of head motion. Neural sequence models consistently improve trajectory forecasting performance. The vanilla Transformer baseline reduces ADE to 0.20 m and FDE to 0.32 m while improving rotation accuracy

Table 3: Ablation study on input modalities using CXA-Transformer on the EgoTraj test split. $\mathcal{Y}$: ego-trajectory, $\mathcal{C}/\mathcal{B}/\mathcal{P}$: nearby people (center point / bounding box / pose), $\mathcal{S}$: scene segmentation, $\mathcal{D}$: relative depth, $\mathcal{G}$: gaze (u, v). Best results per group in **bold**.

Modality	ADE ↓	FDE ↓	L_{head} ↓
$\mathcal{Y}$ [31]	0.19	0.29	0.69
$\mathcal{Y} + \mathcal{C}$	0.18	0.29	0.79
$\mathcal{Y} + \mathcal{B}$	0.18	0.30	0.81
$\mathcal{Y} + \mathcal{P}$	0.17	0.27	0.77
$\mathcal{Y} + \mathcal{S}$	0.16	0.26	0.74
$\mathcal{Y} + \mathcal{D}$	0.18	0.29	0.78
$\mathcal{Y} + \mathcal{G}$	0.15	0.26	0.69
$\mathcal{Y} + \mathcal{C} + \mathcal{G}$	0.14	0.25	0.67
$\mathcal{Y} + \mathcal{B} + \mathcal{G}$	0.16	0.26	0.70
$\mathcal{Y} + \mathcal{P} + \mathcal{G}$	0.12	0.24	0.63
$\mathcal{Y} + \mathcal{S} + \mathcal{G}$	0.12	0.25	0.65
$\mathcal{Y} + \mathcal{D} + \mathcal{G}$	0.15	0.27	0.71
$\mathcal{Y} + \mathcal{P} + \mathcal{S} + \mathcal{G}$	0.12	0.23	0.58

to 0.74, demonstrating the benefit of temporal modeling. CXA-Transformer further improves both trajectory and rotation prediction compared to motion-based baselines, highlighting the advantage of cross-stream attention mechanisms in the egocentric setting. Overall, these results indicate that learned temporal representations significantly outperform simple kinematic extrapolation for egocentric trajectory prediction, while multimodal attention mechanisms improve the modeling of head orientation dynamics.

Ablation Study. In addition to ego-translation and rotation, we performed an ablation study to analyze the contribution of different input modalities to egocentric trajectory prediction. We adopted the CXA-Transformer [31] as the base architecture, since its cascaded cross-attention mechanism allowed progressive fusion of multiple modality streams. We considered ego-motion $\mathcal{Y}$ (ego-translation and rotation), social context from nearby people represented as center points $\mathcal{C}$, bounding boxes $\mathcal{B}$, or body poses $\mathcal{P}$, scene understanding via semantic segmentation $\mathcal{S}$, relative depth $\mathcal{D}$, and gaze $\mathcal{G}$ represented as normalized image coordinates (u, v). Body pose $\mathcal{P}$, center points $\mathcal{C}$, and bounding boxes $\mathcal{B}$ were extracted using YOLOv8-Pose [9], with center point $\mathcal{C}$ derived from the torso keypoint centroid. Scene segmentation $\mathcal{S}$ was obtained from OneFormer [12]. Relative depth $\mathcal{D}$ was estimated using Depth Anything V2 [44]. Gaze $\mathcal{G}$ was obtained directly from the calibrated gaze projection (u, v). Table 3 summarizes the results.

Starting from ego-motion only ($\mathcal{Y}$), both scene segmentation ($\mathcal{S}$) and gaze ($\mathcal{G}$) provide the largest individual improvements. Adding scene context reduces ADE from 0.19 to 0.16, while incorporating gaze further reduces it to 0.15. This suggests that semantic layout and visual attention both contribute anticipatory information beyond kinematic motion cues. Among social representations,

Table 4: Generalization across three test splits using CXA-Transformer. Values are ADE/FDE in meters with 95% bootstrap confidence intervals from 1000 resamples. Best per split in **bold**.

Modality	Random Participant ($n=8$)		Waypoint Held-Out ($n=10$)		Unfamiliar ($n=8$)	
	ADE ↓	FDE ↓	ADE ↓	FDE ↓	ADE ↓	FDE ↓
$\mathcal{Y}$	0.19 ± .014	0.29 ± .021	0.21 ± .018	0.32 ± .024	0.23 ± .019	0.34 ± .027
$\mathcal{Y} + \mathcal{P}$	0.17 ± .011	0.27 ± .019	0.19 ± .015	0.29 ± .022	0.20 ± .013	0.31 ± .020
$\mathcal{Y} + \mathcal{S}$	0.16 ± .013	0.25 ± .014	0.18 ± .012	0.28 ± .018	0.18 ± .016	0.29 ± .023
$\mathcal{Y} + \mathcal{G}$	0.15 ± .009	0.26 ± .017	0.16 ± .014	0.26 ± .013	0.16 ± .010	0.29 ± .018
$\mathcal{Y} + \mathcal{P} + \mathcal{S} + \mathcal{G}$	0.12 ± .008	0.23 ± .012	0.14 ± .010	0.25 ± .011	0.14 ± .012	0.26 ± .014

pose ($\mathcal{P}$) consistently outperforms bounding boxes ($\mathcal{B}$) and center points ($\mathcal{C}$), achieving the lowest trajectory error among the three. This finding suggests that fine-grained human motion cues are more informative than coarse spatial localization for modeling social interactions. In addition, adding gaze to social and scene features consistently improves performance across configurations. The full multimodal combination ($\mathcal{Y} + \mathcal{P} + \mathcal{S} + \mathcal{G}$) achieves the best overall performance, reducing ADE to 0.12, FDE to 0.23, and head rotation error to 0.58. These results highlight that visual attention provides additional predictive signal even in the presence of rich scene and social context.

Generalization across splits. To examine how well the multimodal models generalize beyond the random-participant split, we evaluated CXA-Transformer on two additional, stricter splits. The *waypoint-pair held-out split* reserves 3 of the 21 origin–destination pairs entirely for testing ($n=10$ sessions), and the *unfamiliar split* reserves 8 of the 31 participants who reported being unfamiliar with the recording environment ($n=8$ sessions). By construction, each session is contributed by a unique participant (Sec. 3), so all three splits are subject-disjoint. Table 4 reports ADE/FDE on each split with 95% bootstrap confidence intervals computed from 1000 resamples. The full multimodal configuration ($\mathcal{Y} + \mathcal{P} + \mathcal{S} + \mathcal{G}$) remains the strongest across all three splits, with a modest generalization gap (random-participant ADE 0.12 $\rightarrow$ waypoint-held-out ADE 0.14 $\rightarrow$ unfamiliar ADE 0.14). This indicates that the multimodal cues transfer to held-out landmark pairs and to participants unfamiliar with the area, rather than overfitting to specific route templates.

4.2 Qualitative Evaluation

Figure 6 presents representative trajectory forecasting examples across diverse navigation scenarios. Compared to motion-based baselines, multimodal models better capture turning behavior and obstacle avoidance. In particular, CXA-Transformer and EgoCast models produce smoother and more socially compliant trajectories that align closely with ground truth, while kinematic baselines often overshoot or fail to anticipate directional changes.




Fig. 6: Qualitative trajectory forecasting. Baseline predictions from motion data (ego-translation and rotation) on three EgoTraj test scenarios. Dark blue: observed path; green dashed: ground truth; colors: predictions. *Left:* gentle segment. *Center:* moderate turn where attention-based models track the path best. *Right:* sharp $\sim 90^\circ$ intersection turn where all baselines underestimate the turn magnitude.

Figure 7 visualizes the corresponding multimodal cues for CXA-Transformer at consecutive timesteps ($t+1$ to $t+3$), including relative depth, semantic segmentation, detected human pose, and projected gaze. Relative depth encodes the proximity of surrounding objects to the wearer, while scene segmentation provides structural priors such as walkable regions and obstacles. Human pose captures nearby pedestrian motion, and gaze indicates the wearer’s visual attention toward potential navigation targets. The predicted trajectories of full multimodal combination ($\mathcal{V} + \mathcal{P} + \mathcal{S} + \mathcal{G}$) closely follow these attended regions, illustrating the complementary role of visual attention in guiding motion forecasting. Overall, these qualitative results demonstrate that jointly modeling ego-motion, scene structure, social context, and gaze leads to more realistic and anticipatory trajectory predictions in complex urban environments.

Failure Cases Figure 6 (Window-Slice 2512) shows a representative failure, a sharp $\sim 90^\circ$ turn at an intersection. After the traffic light turns green, the participant abruptly changes heading, causing all baselines to deviate from ground truth. Kinematic baselines extrapolate along the pre-turn heading, while learned models partially anticipate the turn but underestimate its curvature. This arises from the limited signal in short-horizon ego-motion history. Before the turn, linear and angular velocities stay within normal walking ranges, giving little indication that a sharp maneuver is imminent. In addition, the dataset did not include enough data representing such abrupt transitions. More broadly, this behavior reflects an inherent challenge of deterministic trajectory forecasting under multimodal intent. Abrupt intersection turns resist prediction without strong anticipatory cues or richer environmental context, motivating future work on intent-aware and uncertainty-aware egocentric prediction.




Fig. 7: Multimodal observations across three consecutive timesteps. Rows are frames at $t+1$, $t+2$, and $t+3$ of a sidewalk sequence. (a) RGB with gaze (red dot), (b) depth (Depth Anything V2), (c) segmentation (OneFormer), (d) pedestrians (YOLOv8-Pose) ranked by depth, and (e) ground-truth (green) vs. predicted (red).

5 Conclusion

In this paper, We introduce EgoTraj, a multimodal egocentric trajectory dataset of 75 sessions recorded with Meta Quest Pro headsets in real-world urban environments. To our knowledge, it is the first outdoor egocentric trajectory dataset to jointly provide synchronized 6DoF head pose, per-frame 3D gaze vectors, egocentric RGB video, and scene annotations. Benchmarking several state-of-the-art forecasting models, we showed that EgoTraj is challenging for egocentric trajectory prediction, with distinctive motion such as frequent turns and abrupt directional changes. Ablations show that gaze and scene cues add complementary signal beyond ego-motion, underscoring the role of visual attention in anticipating navigation intent. EgoTraj dataset supports future work on egocentric trajectory prediction and vision-language understanding of navigation intent.

Acknowledgements

This work was supported by Honda Development & Manufacturing of America, LLC. We thank Jorge Monsivais, Haithi Donahue, and Steven Feit of the Emerging Technology Department for their contributions. We are also grateful to Kristen Grauman at UT Austin for her guidance throughout the data collection.

References

1. Alahi, A., Goel, K., Ramanathan, V., Robicquet, A., Fei-Fei, L., Savarese, S.: Social lstm: Human trajectory prediction in crowded spaces. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 961–971 (2016)
2. Bock, J., Krajewski, R., Moers, T., Runde, S., Vater, L., Eckstein, L.: The ind dataset: A drone dataset of naturalistic road user trajectories at german intersections. In: IEEE Intelligent Vehicles Symposium. pp. 1929–1934. IEEE (2020)
3. Caesar, H., Bankiti, V., Lang, A.H., Vora, S., Liong, V.E., Xu, Q., Krishnan, A., Pan, Y., Baldan, G., Beijbom, O.: nuscenes: A multimodal dataset for autonomous driving. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11621–11631 (2020)
4. Escobar, M., Puentes, J., Forigua, C., Pont-Tuset, J., Maninis, K.K., Arbelaez, P.: Egocast: Forecasting egocentric human pose in the wild. In: IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 5831–5841. IEEE (2025)
5. Ettinger, S., Cheng, S., Caine, B., Liu, C., Zhao, H., Pradhan, S., Chai, Y., Sapp, B., Qi, C.R., Zhou, Y., et al.: Large scale interactive motion forecasting for autonomous driving: The waymo open motion dataset. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 9710–9719 (2021)
6. Google: Project guideline: Enabling those with low vision to run independently. <https://blog.google/outreach-initiatives/accessibility/project-guideline/> (2021)
7. Grauman, K., Westbury, A., Byrne, E., Chavis, Z., Furnari, A., Girdhar, R., Hamburger, J., Jiang, H., Liu, M., Liu, X., et al.: Ego4d: Around the world in 3,000 hours of egocentric video. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 18995–19012 (2022)
8. Grauman, K., Westbury, A., Torresani, L., Kitani, K., Malik, J., Afouras, T., Ashutosh, K., Baiyya, V., Bansal, S., Boote, B., et al.: Ego-exo4d: Understanding skilled human activity from first- and third-person perspectives. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 19383–19400 (2024)
9. Hermens, F.: Automatic object detection for behavioural research using yolov8. Behavior research methods **56**(7), 7307–7330 (2024)
10. Hu, Y., Yang, J., Chen, L., Li, K., Sima, C., Zhu, X., Chai, S., Du, S., Lin, T., Wang, W., et al.: Planning-oriented autonomous driving. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 17853–17862 (2023)
11. Jain, A., Casas, S., Liao, R., Xiong, Y., Feng, S., Segal, S., Urtasun, R.: Discrete residual flow for probabilistic pedestrian behavior prediction. In: Conference on Robot Learning. pp. 407–419. PMLR (2020)
12. Jain, J., Li, J., Chiu, M.T., Hassani, A., Orlov, N., Shi, H.: Oneformer: One transformer to rule universal image segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2989–2998 (2023)
13. John, R., Kesari, A., DiMatteo, V., Dana, K.: Egocampus: Egocentric pedestrian eye gaze model and dataset. arXiv preprint arXiv:2512.07668 (2025)
14. Johnson, F., Cao, B.B., Jain, S., Ashok, A., Dana, K.: A simple framework for visual navigation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3167–3177 (2026)
15. Kacorri, H., Kitani, K.M., Bigham, J.P., Asakawa, C.: People with visual impairment training personal object recognizers: Feasibility and challenges. In: Proceed-

- ings of the CHI Conference on Human Factors in Computing Systems. pp. 5839–5849 (2017)
16. Karnan, H., Nair, A., Xiao, X., Warnell, G., Pirk, S., Toshev, A., Hart, J., Biswas, J., Stone, P.: Socially compliant navigation dataset (scand): A large-scale dataset of demonstrations for social navigation. *IEEE Robotics and Automation Letters* **7**(4), 11807–11814 (2022)
 17. Kim, D., Srouji, M., Chen, C., Zhang, J.: Armor: Egocentric perception for humanoid robot collision avoidance and motion planning. *arXiv preprint arXiv:2412.00396* (2024)
 18. Land, M.F.: Eye movements and the control of actions in everyday life. *Progress in Retinal and Eye Research* **25**(3), 296–324 (2006)
 19. Lerner, A., Chrysanthou, Y., Lischinski, D.: Crowds by example. In: *Computer Graphics Forum*. vol. 26, pp. 655–664. Wiley Online Library (2007)
 20. Lv, Z., Charron, N., Moulon, P., Gamino, A., Peng, C., Sweeney, C., Miller, E., Tang, H., Meissner, J., Dong, J., et al.: Aria everyday activities dataset. *arXiv preprint arXiv:2402.13349* (2024)
 21. Ma, L., Ye, Y., Hong, F., Guzov, V., Jiang, Y., Postyeni, R., Pesqueira, L., Gamino, A., Baiyya, V., Kim, H.J., et al.: Nymeria: A massive collection of multimodal egocentric daily motion in the wild. In: *European Conference on Computer Vision*. pp. 445–465. Springer (2024)
 22. Martin-Martin, R., Patel, M., RezaTofighi, H., Shenoi, A., Gwak, J., Frankel, E., Sadeghian, A., Savarese, S.: JrdB: A dataset and benchmark of egocentric robot visual perception of humans in built environments. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **45**(6), 6748–6765 (2021)
 23. Meegan, R., D’Souza, A., Cao, B.B., Jain, S., Dana, K.: Yopo-nav: Visual navigation using 3dgs graphs from one-pass videos. *arXiv preprint arXiv:2512.09903* (2025)
 24. Mohamed, A., Zhu, D., Vu, W., Elhoseiny, M., Claudel, C.: Social-implicit: Rethinking trajectory prediction evaluation and the effectiveness of implicit maximum likelihood estimation. In: *European Conference on Computer Vision*. pp. 463–479. Springer (2022)
 25. Nguyen, D.M., Nazeri, M., Payandeh, A., Datar, A., Xiao, X.: Toward human-like social robot navigation: A large-scale, multi-modal, social human navigation dataset. In: *IEEE/RSJ International Conference on Intelligent Robots and Systems*. pp. 7442–7447. IEEE (2023)
 26. Pan, B., Harley, A.W., Engelmann, F., Liu, C.K., Guibas, L.J.: Lookout: Real-world humanoid egocentric navigation. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 24977–24988 (2025)
 27. Pan, X., Charron, N., Yang, Y., Peters, S., Whelan, T., Kong, C., Parkhi, O., Newcombe, R., Ren, Y.C.: Aria digital twin: A new benchmark dataset for egocentric 3d machine perception. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 20133–20143 (2023)
 28. Park, H.S., Hwang, J.J., Niu, Y., Shi, J.: Egocentric future localization. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 4697–4705 (2016)
 29. Pellegrini, S., Ess, A., Schindler, K., Van Gool, L.: You’ll never walk alone: Modeling social behavior for multi-target tracking. In: *Proceedings of the IEEE International Conference on Computer Vision*. pp. 261–268. IEEE (2009)
 30. Peng, C., Paredes, V., Castillo, G.A., Hereid, A.: Real-time safe bipedal robot navigation using linear discrete control barrier functions. In: *IEEE International Conference on Robotics and Automation*. pp. 14903–14909. IEEE (2025)

31. Qiu, J., Chen, L., Gu, X., Lo, F.P.W., Tsai, Y.Y., Sun, J., Liu, J., Lo, B.: Egocentric human trajectory forecasting with a wearable camera and multi-modal fusion. *IEEE Robotics and Automation Letters* **7**(4), 8799–8806 (2022)
32. Qiu, Z., Liu, Z., Niu, W., Bhattacharjee, T., Kalantari, S.: Egocognav: Cognition-aware human egocentric navigation. *arXiv preprint arXiv:2511.17581* (2025)
33. Raina, N., Somasundaram, G., Zheng, K., Miglani, S., Saarinen, S., Meissner, J., Schwesinger, M., Pesqueira, L., Prasad, I., Miller, E., Gupta, P., Yan, M., Newcombe, R., Ren, C., Parkhi, O.: Egoblur model (2023)
34. Robicquet, A., Sadeghian, A., Alahi, A., Savarese, S.: Learning social etiquette: Human trajectory understanding in crowded scenes. In: *European Conference on Computer Vision*. pp. 549–565. Springer (2016)
35. Shi, L., Wang, L., Zhou, S., Hua, G.: Trajectory unified transformer for pedestrian trajectory prediction. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 9675–9684 (2023)
36. Singh, K.K., Fatahalian, K., Efros, A.A.: Krishnacam: Using a longitudinal, single-person, egocentric dataset for scene understanding tasks. In: *IEEE Winter Conference on Applications of Computer Vision*. pp. 1–9. IEEE (2016)
37. Tang, T.J., Li, W.H.: An assistive eyewear prototype that interactively converts 3d object locations into spatial audio. In: *Proceedings of the ACM International Symposium on Wearable Computers*. pp. 119–126 (2014)
38. Tian, Y., Liu, Y., Tan, J.: Wearable navigation system for the blind people in dynamic environments. In: *IEEE International Conference on Cyber Technology in Automation, Control and Intelligent Systems*. pp. 153–158. IEEE (2013)
39. Wang, H.C., Katzschnmann, R.K., Teng, S., Araki, B., Giarré, L., Rus, D.: Enabling independent navigation for visually impaired people through a wearable vision-based feedback system. In: *IEEE International Conference on Robotics and Automation*. pp. 6533–6540. IEEE (2017)
40. Wang, P., Bai, S., Tan, S., Wang, S., Fan, Z., Bai, J., Chen, K., Liu, X., Wang, J., Ge, W., et al.: Qwen2-vl: Enhancing vision-language model’s perception of the world at any resolution. *arXiv preprint arXiv:2409.12191* (2024)
41. Wang, T., Byeon, J., Yehia, A., Wang, H., Xu, Y., Zeng, T., Wang, Z., Jiao, J., Claudel, C.: Xr-dt: Extended reality-enhanced digital twin for agentic mobile robots. *arXiv preprint arXiv:2512.05270* (2025)
42. Wang, W., Liu, C.K., Kennedy III, M.: Egonav: Egocentric scene-aware human trajectory prediction. *arXiv preprint arXiv:2403.19026* (2024)
43. Yagi, T., Mangalam, K., Yonetani, R., Sato, Y.: Future person localization in first-person videos. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 7593–7602 (2018)
44. Yang, L., Kang, B., Huang, Z., Zhao, Z., Xu, X., Feng, J., Zhao, H.: Depth anything v2. *Advances in Neural Information Processing Systems* **37**, 21875–21911 (2024)
45. Yehia, A., Byeon, J., Wang, T., Wang, H., Xu, Y., Jiao, J., Claudel, C.: Arcas: An augmented reality collision avoidance system with slam-based tracking for enhancing vru safety. *arXiv preprint arXiv:2512.05299* (2025)
46. Zheng, W., Song, R., Guo, X., Zhang, C., Chen, L.: Genad: Generative end-to-end autonomous driving. In: *European Conference on Computer Vision*. pp. 87–104. Springer (2024)