

OOSTraj: Out-of-Sight Trajectory Prediction With Vision-Positioning Denoising

Haichao Zhang

Northeastern University
 360 Huntington Ave Boston MA 02115
 zhang.haich@northeastern.edu

Hongsheng Lu

Toyota Motor North America
 465 N Bernardo Ave Mountain View CA 94043
 hongsheng.lu@toyota.com

Yi Xu

Northeastern University
 360 Huntington Ave Boston MA 02115
 xu.yi@northeastern.edu

Takayuki Shimizu

Toyota Motor North America
 465 N Bernardo Ave Mountain View CA 94043
 takayuki.shimizu@toyota.com

Yun Fu

Northeastern University
 360 Huntington Ave Boston MA 02115
 yunfu@ece.neu.edu

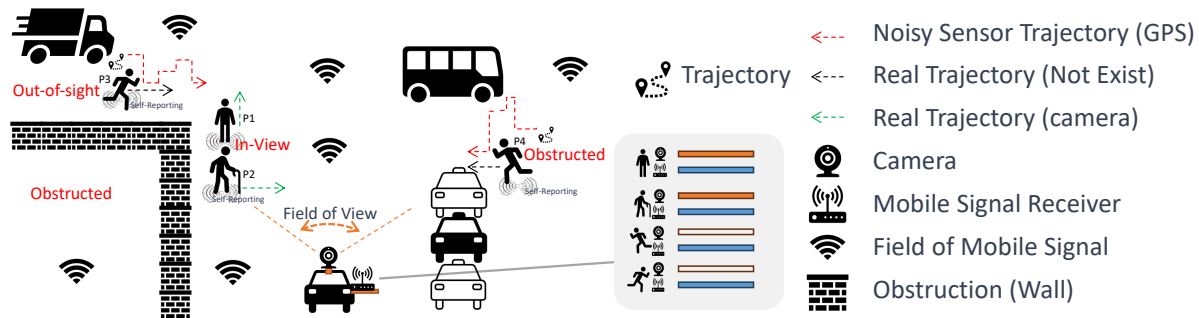


Figure 1. A representative illustration of real-world out-of-sight scenarios in autonomous driving. The autonomous vehicle is equipped with a camera (capturing precise visual trajectories, indicated by green dotted arrows) and a mobile signal receiver (capturing noisy sensor trajectories, represented by red dotted arrows) for tracking pedestrians and other vehicles. Pedestrians P1 and P2 are within the camera’s field of view, while P3 is entirely out of sight and P4 is obscured by other vehicles. Consequently, P3 and P4 lack captured visual trajectories and are positioned dangerously, potentially crossing into the vehicle’s path, posing a risk of collision. The black dotted arrows depict the hypothesized noise-free real trajectories, ideally captured by mobile sensors, contrasting with the actual noisy sensor trajectories (red arrows). The gray area in the figure demarcates the visibility range of the mobile and visual modalities: white indicates no data captured, orange signifies the presence of visual trajectories, and blue represents the availability of mobile trajectories.

Abstract

Trajectory prediction is fundamental in computer vision and autonomous driving, particularly for understanding pedestrian behavior and enabling proactive decision-making. Existing approaches in this field often assume precise and complete observational data, neglecting the challenges associated with out-of-view objects and the noise inherent in sensor data due to limited camera range, physical obstructions, and the absence of ground truth for denoised sensor data. Such oversights are critical safety con-

cerns, as they can result in missing essential, non-visible objects. To bridge this gap, we present a novel method for out-of-sight trajectory prediction that leverages a vision-positioning technique. Our approach denoises noisy sensor observations in an unsupervised manner and precisely maps sensor-based trajectories of out-of-sight objects into visual trajectories. This method has demonstrated state-of-the-art performance in out-of-sight noisy sensor trajectory denoising and prediction on the Vi-Fi and JRDB datasets. By enhancing trajectory prediction accuracy and addressing the challenges of out-of-sight objects, our work signif-

icantly contributes to improving the safety and reliability of autonomous driving in complex environments. Our work represents the first initiative towards Out-Of-Sight Trajectory prediction (OOSTraj), setting a new benchmark for future research.

1. Introduction

Trajectory prediction is pivotal in diverse research domains, including computer vision, virtual reality, robotics, and autonomous driving. This capability is fundamental for understanding human behavior in deep learning models and crucial for enabling autonomous vehicles to anticipate pedestrian movements, thereby enhancing safety and preventing collisions. Real-world applications, however, present complex environments where trajectory prediction remains a formidable challenge. While substantial progress has been made in this field, existing approaches heavily rely on computer vision algorithms under the assumption of noise-free and constantly visible observations. This assumption often overlooks the critical aspects of out-of-sight objects and noisy data, which are prevalent in real-world scenarios. Our research addresses this gap by proposing a vision-positioning denoising method that enhances trajectory prediction in these challenging conditions.

In the context of autonomous driving and robotics, the limitations of camera-based systems pose significant challenges. As shown in Fig. 1, one common scenario involves pedestrians or vehicles being obscured by other vehicles, walls, or falling into the vehicle’s blind spots. In such cases, they may suddenly appear in the camera’s field of view, leaving insufficient time for deep learning algorithms to respond proactively. Additionally, the high speed of vehicles exacerbates this issue. When a camera moves rapidly, pedestrians can emerge from out-of-sight areas too quickly for the camera to capture and process their trajectory information. In both scenarios, the absence of prior visual data on pedestrians or vehicles hinders the trajectory prediction capabilities of autonomous systems. This lack of out-of-sight observations poses a serious safety risk and limits the practical deployment of autonomous driving technology in real-world situations. Our research aims to mitigate these challenges by enhancing trajectory prediction through vision-positioning denoising, enabling more reliable detection and response to such unpredictable occurrences.

Recent studies in trajectory prediction have increasingly acknowledged the impediments posed by incomplete observations, a factor that significantly hinders accurate predictions. While the bulk of existing research presumes the availability of complete and precise observational data, several noteworthy studies, such as those by Xu *et al.* [27] and Fujii *et al.* [6], have begun to address this challenge. These works predominantly focus on imputing missing views or

directly predicting from partial trajectory data. A notable method by Zhang *et al.* [28] demonstrates the prediction of layout trajectories using just a single visual timestamp, supported by mobile sensor modalities. However, these studies share a common limitation: their dependency on a minimum quantum of visual observational data for trajectory prediction. To the best of our knowledge, there is an absence of research dedicated to predicting trajectories solely from out-of-sight observations. This research gap is not only significant in the academic field but also constitutes an urgent safety issue in autonomous driving, a concern that our research directly addresses.

Addressing the out-of-sight problem in trajectory prediction presents two key challenges. The initial challenge involves coping with noisy sensor measurements. In the absence of visual timestamps, our approach relies on commonly used localization sensors, such as GPS and odometers, to provide a basic spatial location. This reliance is necessitated by the nature of the out-of-sight problem, where traditional visual inputs are not available. However, as established in prior research, these mobile sensor modalities are subject to significant noise. GPS, for example, can exhibit errors ranging from 1 to 4 meters [10], while odometers are prone to drift noise accumulation [29]. This noise significantly hampers the accuracy of sensor-derived trajectory predictions. Furthermore, the absence of ground truth for denoised sensor data complicates the situation, as it prevents us from employing supervised learning methods for denoising. Therefore, our approach involves training a denoising model in an unsupervised manner.

To mitigate this issue, we propose the use of visual information for denoising sensor data. Although our scenario lacks direct visual observations, object tracking algorithms can offer highly accurate and noise-free location data in image coordinates. By integrating this visual information, we aim to refine and denoise the trajectory data derived from the noisy sensor inputs. However, direct visual observation of out-of-sight subjects is not possible, we need to establish a connection between sensor data and visual coordinates. This brings us to our second challenge.

However, since the pedestrian is out-of-sight, to do such a denoising method with vision-positioning, another challenge is there is no visual reference even if we assume the localization is precise. To solve this challenge, we have to find a localization-vision mapping to solve the visual reference problem. Luckily, since the camera’s intrinsic matrix is not changing within a frame or is not moving, we can build such localization-vision mapping by estimating the camera’s intrinsic matrix. Since the pedestrians and vehicles shares the same camera’s intrinsic matrix in the camera’s view. To best utilize the information within camera frames, through analysis of other visible pedestrians and vehicle movements between sensor localization trajectories

and visual trajectories, we can build a model to predict the mapping relationship which is highly associated with the camera’s intrinsic matrix. Nevertheless, another benefit of this design is, that if the camera is not moving in some widely used applications like video surveillance [8, 30] or traffic monitoring [11], we can infer the localization-vision mapping through pedestrians and vehicles in previous frames.

The contributions of this paper are as follows:

- We introduce a pioneering task in the field of trajectory prediction - predicting the noise-free visual trajectory of out-of-sight objects using noisy sensor trajectories. This task (OOSTraj) addresses a critical gap in current research and extends the capabilities of trajectory prediction in complex environments.
- Our work proposes an innovative vision-positioning denoising module. This module adeptly leverages precise visual data to effectively denoise sensor data with an unsupervised manner of no need for the ground truth of denoised sensor data through the camera matrix estimator, providing a much-needed visual reference for objects that are out of the camera’s view. This approach marks a significant advancement in combining sensor and visual data for enhanced prediction accuracy.
- Through rigorous ablation studies and plug-and-play experiments, we demonstrate that our method achieves state-of-the-art results on out-of-sight noisy sensor trajectory denoising and prediction over baselines. These studies provide concrete evidence of the effectiveness and efficiency of our proposed approach.

2. Related Works

2.1. Vision-Wireless Fusion

Vision-wireless fusion represents an interdisciplinary approach that synergizes the strengths of visual and wireless (mobile) modalities to tackle inherent challenges in each. The visual modality, while offering detailed and direct observation capabilities, grapples with limitations such as obstructions, out-of-sight issues, computational intensity, and a restricted observational range. On the other hand, the wireless modality, characterized by its extensive range of sensing and communication, encounters significant issues with sensor noise [15] and environmental disturbances [5]. A persistent challenge in wireless data is the lack of ground truth, compounded by the difficulty in human labeling due to the non-visual nature of the signals.

Despite these challenges, the fusion of visual and wireless data presents a powerful tool. The visual modality, with its ease of visualization and labeling, coupled with advancements in computer vision algorithms, offers robust and precise supervision. This fusion has seen limited but impactful applications in recent studies. For instance, Liu

et al. [16] utilized wireless trajectories to address person reidentification problems, accounting for variations in posture and clothing. Alahi et al. [1] (2015) employed RGB-D cameras to enhance indoor localization accuracy by estimating distances using Wi-Fi signals from smartphones. Papaioannou et al. [19] focused on tracking people across visual and wireless modalities in dynamic industrial environments. These pioneering works have significantly advanced the field by fusing visual and mobile modalities to improve task accuracy. However, less attention has been paid to leveraging the visual modality for denoising noisy sensor trajectories, especially for out-of-sight agents, an area our work aims to explore.

2.2. Obstructed Trajectory Prediction

The field of obstructed trajectory prediction, while closely related, is distinct from out-of-sight trajectory prediction, which remains underexplored. Trajectory prediction has widespread applications in robotics [12], autonomous driving [4], and human behavior understanding [21]. Many existing studies [18] operate under the assumption of complete and unobstructed observations. However, recent research has begun to acknowledge and address the challenges posed by observation obstructions. Methods such as data imputation [27] and modality fusion [28] have been proposed to tackle these issues. Nonetheless, these approaches typically rely on having at least some observational data, leaving a gap in the context of completely out-of-sight trajectory prediction. In real-world scenarios, noisy and out-of-sight trajectories are commonplace, presenting a significant challenge for existing systems. Our work seeks to bridge this gap by employing vision positioning and denoising techniques for trajectory prediction, addressing a crucial need in the field.

3. Problem Definition

3.1. Symbol Annotations

Consider a set of N agents, represented as $P_n, n = 1, 2, \dots, N$, observed over timestamps $t_1 : t_2$, starting from t_s and ending at t_e . Agents visible to the camera, P_i , constitute the set \mathbb{P}_{in} , while those out-of-sight, P_o , form the set \mathbb{P}_{out} . The sensor-based noisy trajectory of the n -th agent between t_1 and t_2 is denoted by $S_{P_n}^{t_1:t_2}$, which ideally corresponds to a precise but non-existent real localization trajectory, $\hat{S}_{P_n}^{t_1:t_2}$. All agents possess a noisy sensor trajectory $S_{P_n}^{t_s:t_e}$, significantly influenced by sensor measurement noise. Concurrently, in-view agents $P_i, i \in \mathbb{P}_{in}$, are tracked through visual trajectories $V_{P_i}^{t_s:t_e}$ in camera frames.

3.2. Task Definition

Our primary objective, given the sole input of noisy sensor trajectories $S_{P_o}^{t_s:t_e}$, is twofold:

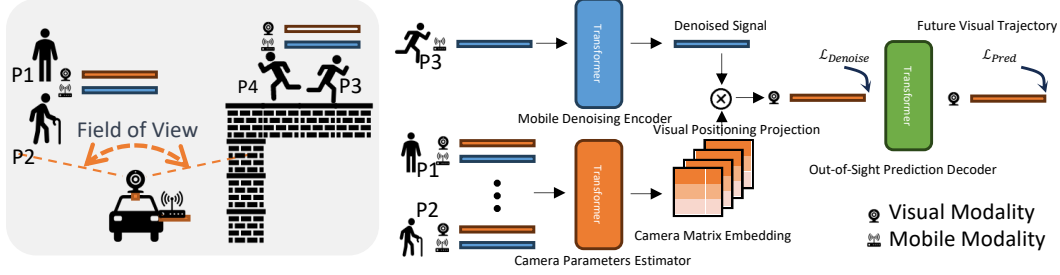


Figure 2. Overview of the Vision-Positioning Denoising and Predicting Model architecture. This illustration highlights the processing of pedestrian data, where pedestrians P3 and P4 are detectable only by mobile receivers, while P1 and P2 are visible to both camera and mobile receivers. The Camera Parameters Estimator Module utilizes the dual-modality trajectories of in-view pedestrians (like P1 and P2) to analyze the relationship between camera and world coordinates, resulting in a camera matrix embedding. For out-of-sight pedestrians (e.g., P3, P4), their noisy mobile trajectories are first refined by the Mobile Denoising Encoder, producing a denoised signal embedding. This embedding is then merged with the matrix embedding in the Visual Positioning Projection Module, facilitating the mapping of data into camera coordinates, with the application of $\mathcal{L}_{Denoise}$. Finally, the Out-of-Sight Prediction Decoder leverages the denoised visual signals to predict the trajectories of pedestrians not captured by the camera.

1. Denoise the noisy sensor trajectory $S_{P_o}^{t_s:t_e}$ without direct supervision, addressing the potential impact of sensor noise on prediction accuracy.
2. Predict noise-free future visual trajectories $V_{P_o}^{t_e:t_p}$ for out-of-sight agents P_o in camera frames, spanning from the observation’s end timestamp t_e to the target prediction timestamp t_p , using only the provided noisy sensor trajectory data.

It is important to note that P_o lacks visual trajectories and references due to being out-of-sight. However, we can leverage other available signals, as outlined in Sec. 3.1, to navigate complex environmental factors.

Our strategy involves initially denoising and projecting the noisy sensor trajectory into a visual trajectory in an unsupervised manner, constructing supervision from the available signals. Subsequently, another module is employed to execute precise trajectory prediction within the visual modality.

4. Methodology

The overall architecture of our proposed method is detailed in Fig. 2. Our approach is specifically designed to predict noise-free future trajectories for out-of-sight agents, utilizing noisy sensor data without relying on direct denoising supervision or visual references.

To accurately predict a noise-free trajectory from noisy sensor data, it is crucial to establish an unsupervised denoising pipeline. This pipeline is essential for filtering out noise before it can impact the trajectory prediction model. Addressing this complex challenge necessitates the integration of visual references and the construction of effective denoising supervision. We have divided this process into several distinct steps to ensure comprehensive coverage of all aspects involved in this sophisticated task.

4.1. Mobile Denoising Encoder (MDE)

At the outset of our methodology, we introduce the Mobile Denoising Encoder (MDE), designed to operate under an ideal scenario where accurate supervision is available. The MDE’s primary function is to denoise noisy sensor trajectories of out-of-sight agents, leveraging appropriate supervisory signals for accuracy.

Structurally, the MDE comprises a Transformer model, flanked by two fully connected layers, both preceding and succeeding the Transformer layers. This architecture is chosen for its efficacy in handling sequential data and its capacity for capturing complex dependencies within the sensor trajectories.

In an optimal setting, the MDE processes the noisy sensor trajectories, denoted by $S_{P_o}^{t_s:t_e}$, as its input. The objective is to refine these inputs by filtering out the noise, thereby producing what we term ‘noise-free real sensor trajectories’, represented by $\hat{S}_{P_o}^{t_s:t_e}$. The operational equation of the MDE is formalized as follows:

$$\hat{S}_{P_o}^{t_s:t_e} = \mathbf{E}_{denoise}(S_{P_o}^{t_s:t_e}), \quad o \sim \mathbb{P}_{out}, \quad (1)$$

where the $\mathbf{E}_{denoise}$ is the MDE, the \mathbb{P}_{out} is the set of out-of-sight agents, and the P_o is an out-of-sight agent sampled from \mathbb{P}_{out} .

However, the practical implementation of the Mobile Denoising Encoder (MDE) faces a significant challenge: the ideal supervision in the form of noise-free real sensor trajectories, $\hat{S}_{P_o}^{t_s:t_e}$, does not exist in real-world scenarios. The inherent nature of sensor measurement noise complicates the acquisition of such ideal data. Theoretically, one could envisage using higher-accuracy sensor devices to collect comparable data under identical conditions. However, this approach is impractical for most datasets and application scenarios. Even if it were feasible, these higher-

accuracy sensors would likely still be subject to some level of measurement noise, albeit reduced. As a result, completely eliminating noise to construct a truly noise-free real sensor trajectory for supervision is not achievable with the current state of sensor technology.

Consequently, we are compelled to explore alternative methods for constructing effective noise-free supervision for the denoising process described in Eq. 1. This involves developing innovative approaches that can approximate the ideal noise-free conditions, despite the inherent limitations of existing sensor technologies.

4.2. Visual-Positioning Denoising Module (VPD)

In the Visual-Positioning Denoising Module (VPD), we leverage visual information to establish noise-free supervision for the Mobile Denoising Encoder (MDE). This approach is a response to the challenges posed by the noisy sensor modality, which is inherently prone to sensor measurement noise and difficult to refine. Visual information, on the other hand, benefits from the advancements in object tracking algorithms, offering significantly higher precision. Additionally, the nature of visual image processing, where images are discretized into integers during the rasterization process, results in visual trajectories that are highly precise, easily labelable, and effectively noise-free.

Despite the inherent noise-free quality of visual trajectories and their ideal use as supervision in trajectory prediction, a major challenge arises when dealing with out-of-sight object trajectories, which lack visual data. To address this, it is necessary to find a method to translate the denoised sensor trajectories into the visual modality effectively. This translation process is broken down into several key components: the Visual Positioning Projection Module (discussed in Sec. 4.3), the Camera Parameters Estimator (outlined in Sec. 4.4), and the Denoising Loss (described in Sec. 4.5). Each of these components plays a crucial role in ensuring the accurate mapping of sensor data into the visual domain, which is vital for the success of our proposed method.

4.3. Visual Positioning Projection Module (VPP)

The Visual Positioning Projection Module (VPP) plays a critical role in transforming noisy sensor trajectories into a visual format. As all visual trajectories are inherently captured by cameras, this module focuses on converting 3D world coordinates, represented by noisy sensor data, into 2D coordinates within the camera’s frame. This conversion is fundamental to aligning sensor-based data with visual data, and it hinges on the principles of geometric camera calibration [24].

The World-To-Camera Transformation, central to this module, is designed to compute a 2D point $p \in \mathbb{R}^{2 \times 1}$ in the camera frame based on a corresponding 3D point $P \in \mathbb{R}^{3 \times 1}$ in world coordinates. To facilitate this process,

the height dimension in these coordinates is usually kept constant, simplifying the transformation from 3D to 2D space. The mathematical formulation of this transformation is as follows:

$$[p, 1]^T = w \cdot K \cdot R_t \cdot [P, 1]^T, \quad (2)$$

where $K \in \mathbb{R}^{3 \times 3}$ denotes the camera’s intrinsic matrix, which includes parameters like the focal length and optical center that are inherent to the camera. The scale factor w adjusts the scale of the transformation to fit the camera’s specific dimensions. Additionally, $R_t \in \mathbb{R}^{3 \times 4}$ represents the extrinsic rotation and translation matrix, defining the camera’s position and orientation relative to the 3D world.

To streamline this transformation process, we introduce a simplification by employing a matrix $M \in \mathbb{R}^{3 \times 4}$. This matrix M is formulated by combining the intrinsic matrix K , scale factor w , and the extrinsic matrix R_t . This consolidation is critical as it encapsulates all the necessary parameters and transformations required to convert a point from world coordinates into camera coordinates. The simplified form of the World-To-Camera Transformation function can be expressed as:

$$M = w \cdot K \cdot R_t, \quad (3)$$

$$[p, 1]^T = M \cdot [P, 1]^T. \quad (4)$$

To adapt the World-To-Camera Transformation Eq. 2 into trajectory denoising, we need to consider the camera status in this process. If the visual trajectory is captured by a surveillance camera, the R_t in M will not change for engaging of timestamps, but considering the scenarios that visual trajectory is captured by a moving camera on a mobile device, we need to consider M is changing for different timestamps. So, we define the sequence of M as Camera Matrix Embedding $M^{t_s:t_e} \in \mathbb{R}^{(t_e-t_s) \times 3 \times 4}$, which can be a formula as,

$$M^{t_s:t_e} = w \cdot K \cdot R_t^{t_s:t_e}, \quad t \in [t_s, t_e], \quad (5)$$

in which $R_t^{t_s:t_e}$ and $M^{t_s:t_e}$ are extrinsic rotation & translation matrix and Camera Matrix Embedding in timestamps between t_s and t_e .

We regarded noise-free real sensor trajectory $\hat{S}_{P_o}^{t_s:t_e}$ as a sequence of 3D world coordinates $P \in \mathbb{R}^{3 \times 1}$. So, combing Eq. 4 and Eq. 5, the Visual Positioning Projection function to map $\hat{S}_{P_o}^{t_s:t_e}$ into visual modality can be defined as,

$$V_{P_o}^{t_s:t_e} = M^{t_s:t_e} \cdot \hat{S}_{P_o}^{t_s:t_e}, \quad t \in [t_s, t_e], o \sim \mathbb{P}_{out}, \quad (6)$$

in which the $V_{P_o}^{t_s:t_e}$ is the visual trajectory for out-of-sight agent P_o .

4.4. Camera Parameters Estimator (CPE)

Although we have the Visual Positioning Projection in Eq. 6 to map noise-free sensor trajectory into the visual trajectory, the camera matrix is still not accessible because most datasets do not include the camera matrix, and it cannot be obtained from images directly. Moreover, considering the camera is moving in many practical scenarios like autonomous driving, mobile robots, and smartphones, the camera matrix may change in this process. That makes obtaining the camera matrix sequence extremely hard.

However, since there are many in-view agents captured by the same camera in the same timestamp and have both visual and sensor trajectories, we propose a Camera Parameter Estimator to analyze the relationship between the visual and sensor trajectories. Because that relationship is highly related to the camera parameters as shown in Eq. 6, we predict camera matrix embedding $M^{t_s:t_e}$ from all the pairs of in-sight visual and sensor trajectories, with the Camera Parameter Estimator. The Camera Parameter Estimator consists of a transformer with a fully connected layer before and after. The prediction process can be denoted as

$$M^{t_s:t_e} = \mathbf{E}_{cpe}(V_{P_0}^{t_s:t_e}, S_{P_0}^{t_s:t_e}, \dots, V_{P_i}^{t_s:t_e}, S_{P_i}^{t_s:t_e}), \quad i \in \mathbb{P}_{in}, \quad (7)$$

where the $V_{P_i}^{t_s:t_e}, S_{P_i}^{t_s:t_e}$ are the visual and sensor trajectories of in-sight agents P_i , respectively. The P_{in} is the set of in-sight agents. \mathbf{E}_{cpe} is the Camera Parameters Estimator.

4.5. Denoising Loss

In the last subsection, we obtain the visual trajectory of the out-of-sight agent with Eq. 7, since we don't have ground truth for noisy sensor trajectories and the visual modality is noise-free and easy to obtain or label as we described in Sec. 4.2, we can use the process to construct a supervision for the Mobile Denoising Encoder and Camera Parameters Estimator. The denoising loss function $\mathcal{L}_{Denoise}$ can be formulated as,

$$\mathcal{L}_{Denoise} = \mathcal{L}_2(V_{P_o}^{t_s:t_e}, \bar{V}_{P_o}^{t_s:t_e}), \quad (8)$$

where \mathcal{L}_2 is the L2 [2] loss. The $V_{P_o}^{t_s:t_e}$ and $\bar{V}_{P_o}^{t_s:t_e}$ are the predicted visual trajectories of out-of-sight agents and the ground truth. Please note that the ground truth of the visual trajectories of out-of-sight agents is only available during training.

4.6. Out-of-Sight Prediction Decoder (OPD)

We propose an Out-of-Sight Prediction Decoder to predict future visual trajectories from the predicted out-of-sight visual trajectories in Eq. 7, to avoid predicting directly from noisy sensor trajectories. We applied a transformer model

Dataset	Vi-Fi Dataset [14]			JRDB Dataset [17]		
	Baselines	SUM↓	MSE-D↓	MSE-P↓	SUM↓	MSE-D↓
ViTag [13]	200.90	100.53	100.37	143.08	71.23	71.85
Vanilla LSTM [9, 23]	116.01	58.31	57.70	56.05	27.98	28.07
Vanilla GRU [3, 25]	57.34	28.69	28.65	71.91	35.83	36.07
Vanilla RNN [20, 22]	31.61	15.92	15.69	112.40	56.00	56.41
Vanilla Transformer [7, 26]	28.33	14.26	14.08	33.37	16.71	16.66
Ours	27.24	13.42	13.83	25.51	10.52	14.99

Table 1. Experiments of Quantitative Comparison of Models

as the prediction module for simplicity. This process can be formulated as

$$V_{P_o}^{t_s:t_p} = \mathbf{D}_{pred}(V_{P_o}^{t_e:t_p}), \quad o \sim \mathbb{P}_{out}, \quad (9)$$

where the $V_{P_o}^{t_e:t_p}$ is the observation of predicted visual trajectory of out-of-sight agent P_o between timestamps $[t_s : t_e]$, and the $V_{P_o}^{t_s:t_p}$ is the predicted future visual trajectory between timestamps $[t_e : t_p]$. \mathbf{D}_{pred} is the Out-of-Sight Prediction Decoder. \mathbb{P}_{out} is the set of out-of-sight agents.

We applied a prediction loss \mathcal{L}_{Pred} for this process,

$$\mathcal{L}_{Pred} = \mathcal{L}_2(V_{P_o}^{t_s:t_p}, \bar{V}_{P_o}^{t_s:t_p}), \quad (10)$$

where the \mathcal{L}_2 is the L2 loss. $V_{P_o}^{t_s:t_p}$ and $\bar{V}_{P_o}^{t_s:t_p}$ are the predicted and ground truth of future trajectories between t_s and t_p timestamp.

4.7. Implementation Details

The training samples of out-of-sight agents are randomly sampled from all pedestrians in a scenario, all other completed visible agents in both visual and noisy modalities are regarded as in-sight agents for the estimating of the camera parameters estimator.

The wireless data referenced in our study pertains to sensor localization signals, such as noisy GPS or odometers, typically collected by mobile phones carried by pedestrians. The vehicle, equipped with both a camera and a wireless receiver, captures sensor and visual trajectories. This data is then fed into our model for the prediction.

5. Experiments

5.1. Datasets

Vi-Fi Multimodal Dataset [14]. The Vi-Fi dataset is a comprehensive multimodal dataset designed for vision-wireless systems, particularly focusing on linking pedestrian identities across visual and wireless modalities. It features wireless data captured from smartphones carried by pedestrians, encompassing technologies like FTM, IMU, and noisy GPS. Additionally, an RGB-D camera surveillance system, accompanied by a wireless receiver, is deployed either on a wall or on a bicycle to record pedestrian bounding boxes. This setup ensures simultaneous acquisition of both visual and wireless data. Given the nature of

the data collection process, the wireless data exhibits considerable noise, a factor that cannot be overlooked. This dataset comprises 90 sequences, each lasting approximately 3 minutes, and includes both indoor and outdoor scenarios. Indoor data were collected with five legitimate users and without any passersby, while outdoor data involved 3 actual users and 12 passersby. For our experiments, we utilized the noisy GPS signals from this dataset as representative of noisy mobile trajectories and the visual points as the visual modality. To simulate out-of-sight scenarios, we selectively obscured one pedestrian in each sequence.

JackRabbit Dataset (JRDB) [17]. The JRDB dataset, collected by the social mobile robot JackRabbit, is designed for autonomous robot navigation and social robotics studies within human environments. It features 60,000 annotated frames, including 2.4 million 2D and 2.8 million 3D human bounding box annotations, captured using 360 degree RGB (5 cameras) and LIDAR point cloud technology. This dataset encompasses both obstructed and visible pedestrians, with the potential for each camera to have blind spots relative to the others. The bounding boxes are manually labeled, albeit at a lower annotation rate, and are subject to noise due to upsampling via linear interpolation. Size and location estimates of obstructed pedestrians in point clouds are manually conducted. For our experiments, we use the center points of 3D bounding boxes as noisy sensor trajectories and visual points as visual trajectories. To simulate out-of-sight scenarios, we randomly obscured a pedestrian, using the trajectories of several other visible pedestrians to estimate camera parameters in our model.

5.2. Evaluation Setup

In both the JRDB and Vi-Fi datasets, we observe 100 timestamps of out-of-sight trajectories and predict the subsequent 100 timestamps in the visual modality. Simultaneously, we input 100 timestamp pairs of in-sight visual trajectories and noisy sensor signals into our camera parameters estimator to obtain the camera matrix embeddings.

5.3. Metrics

As the first study to denoise out-of-sight sensor trajectories into a visual modality for noise-free trajectory prediction, we adapt the Mean Square Error per Timestamp (MSE-T) metric from [28]. This metric calculates the average pixel distance over the timestamp dimension. Given our task’s dual focus on denoising and prediction performance, we apply MSE-T to both projected out-of-sight visual trajectories and predicted visual trajectories, denoting these as **MSE-D** and **MSE-P**, respectively. Additionally, we introduce the **SUM** metric, which aggregates MSE-D and MSE-P, to provide an overall score balancing denoising and prediction efficacy.

5.4. Baselines

To benchmark our method in noisy sensor trajectory denoising and out-of-sight trajectory prediction, we select several baselines, encompassing fundamental models like Vanilla LSTM [23], Vanilla RNN [20], Vanilla GRU [25], and Vanilla Transformer [7]. We also include ViTag [13], the state-of-the-art method on the Vi-Fi dataset, modified to align with our task outputs.

5.5. Quantitative Comparison Experiments

We conducted quantitative experiments to compare our method’s performance against five baselines across two datasets and three metrics, as detailed in Table. 1. Our model outperforms the baselines, showcasing its superiority. Notably, although ViTag is recognized as the state-of-the-art on the Vi-Fi dataset for visual and wireless identity association, it underperforms in our task. A key finding is that our model demonstrates more pronounced superiority in the denoising task than in the prediction task, underscoring the complexity and difficulty in denoising noisy sensor trajectories and mapping out-of-sight objects into the visual modality. Furthermore, a positive correlation between denoising and prediction performance is observed among the baselines. Improved denoising leads to enhanced prediction accuracy, validating our approach of focusing on sensor data denoising to facilitate better out-of-sight visual trajectory prediction.

We also note varying performances of the baselines across different datasets. For instance, Vanilla RNN excels in the Vi-Fi dataset but falls short in the JRDB dataset, likely due to differing sensor noise characteristics in these datasets. Our vision-positioning denoising model, tailored to handle various sensor noises with precise visual modality supervision, consistently achieves robust denoising performance across different datasets and scenarios. This hypothesis is further explored in our Plug-and-Play experiment.

5.6. Plug and Play Experiments

In the Plug-and-Play experiment, outlined in Table. 2, we assess whether our vision-positioning denoising model enhances the performance of baselines in both noisy sensor trajectory denoising and future visual trajectory prediction. Baselines were modified into a 2-stage format for a comparative analysis. The first stage focuses on denoising, followed by prediction in the second stage. The "+ VPD" row features our vision-positioning denoising (VPD) model, as described in Sec. 4.2, concatenated with the baseline to evaluate the added improvement.

Results in Table. 2 indicate that incorporating our VPD model significantly enhances baseline performance in both denoising and prediction tasks. Notably, Vanilla RNN, which struggles with JRDB’s sensor noise, shows marked improvement in the Vi-Fi dataset when paired with our

Dataset		Vi-Fi Dataset [14]			JRDB Dataset [17]		
Baselines	Add Module	SUM↓	MSE-D↓	MSE-P↓	SUM↓	MSE-D↓	MSE-P↓
Vanilla LSTM [9, 23]	+ 2 Stage	118.44	38.85	79.59	81.22	39.93	41.28
	+ VPD (Ours)	30.39	13.76	16.63	53.16	11.36	41.80
Vanilla RNN [20, 22]	+ 2 Stage	32.16	17.51	14.65	138.41	103.97	34.44
	+ VPD (Ours)	27.78	13.56	14.22	31.32	12.31	19.01
Vanilla GRU [3, 25]	+ 2 Stage	40.35	16.88	23.48	56.56	17.90	38.66
	+ VPD (Ours)	28.47	13.69	14.78	31.70	11.53	20.16
Vanilla Transformer [7, 26]	+ 2 Stage	28.87	14.22	14.65	36.99	14.21	22.79
	+ VPD (Ours)	27.24	13.42	13.83	25.51	10.52	14.99

Table 2. Experiment of Plug-and-Play on the Vi-Fi and JRDB Datasets.

Dataset	Vi-Fi Dataset [14]			JRDB Dataset [17]		
	Module Components	SUM↓	MSE-D↓	MSE-P↓	SUM↓	MSE-D↓
w/o CPE in Sec. 4.4	32.12	17.01	15.12	29.66	14.35	15.31
w/o MDE in Sec. 4.1	32.33	17.65	14.68	32.94	17.90	15.03
w/o VPP in Sec. 4.3	28.42	14.05	14.37	32.48	14.45	18.03
w/o OPD in Sec. 4.6	27.33	13.65	13.68	30.01	14.99	15.02
Full Model	27.24	13.42	13.82	25.51	10.52	14.99

Table 3. Ablation Study of Module Components.

VPD. This improvement in denoising subsequently boosts prediction accuracy. These findings affirm the efficacy of our model in augmenting denoising and prediction capabilities of existing models in a plug-and-play approach.

5.7. Ablation Study

An ablation study was conducted to evaluate the contribution of each module component. This was achieved by systematically removing each component, with the results presented in Table. 3. The study reveals that each component plays a vital role in the overall performance of the model.

The absence of the Camera Parameter Estimator (CPE) leads to a significant performance drop. Without the CPE, the model struggles to establish the visual-positioning relationship through camera matrix estimation, reducing it to a mere two-stage encoder-decoder architecture akin to the baselines. Similarly, the removal of the Mobile Denoising Encoder (MDE) results in an even more pronounced decline in performance. Without the MDE, the model faces challenges in denoising the mobile noisy trajectories, inadvertently allowing noise to propagate into the Visual Positioning Projection (VPP) process. The VPP module, responsible for projecting out-of-sight sensor trajectories into visual trajectories, proves to be crucial. The notable performance drop in scenarios without VPP underscores the importance of explicitly incorporating knowledge from geometric camera calibration into the model. Additionally, the absence of an out-of-sight prediction decoder slightly reduces the model’s effectiveness, as it then needs to infer the entire sequence during the denoising process. Overall, this ablation study clearly demonstrates the necessity and significance of

each component within our model for achieving noise-free out-of-sight trajectory prediction.

6. Conclusion

Our study tackles a significant challenge in trajectory prediction, specifically in situations where agents are entirely out of sight, making visual observations unfeasible. We introduce an innovative task that focuses on predicting the visual trajectories of such out-of-sight agents by utilizing noisy sensor data. The cornerstone of this task is our Visual-Positioning Denoising Module, which adeptly addresses the absence of visual references. This is achieved through sophisticated camera calibration techniques, enabling us to estimate camera matrix sequences and establish a vision-positioning mapping. This method effectively denoises sensor trajectories without relying on direct supervision, instead utilizing supervision constructed from the visual positioning projection.

Our experimental results robustly demonstrate the effectiveness of our proposed pipeline and architecture, significantly outperforming existing methods in predicting trajectories of agents that are not visible. The inclusion of visual positioning in the denoising process markedly improves the accuracy of our trajectory predictions. To the best of our knowledge, this is the first study to successfully predict out-of-sight trajectories using noisy sensor data and to employ vision-positioning projection for the purpose of denoising these trajectories. The methodologies and insights gleaned from our work pave the way for further research and development in the trajectory prediction field, especially in real-world scenarios where complete visibility of agents is not guaranteed, or where sensor data are inherently noisy and challenging to denoise.

7. Acknowledgements

We gratefully acknowledge the support and sponsorship provided by Toyota Motor North America for this research.

References

- [1] Alexandre Alahi, Albert Haque, and Li Fei-Fei. Rgb-w: When vision meets wireless. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3289–3297, 2015.
- [2] Peter Bühlmann and Bin Yu. Boosting with the l_2 loss: regression and classification. *Journal of the American Statistical Association*, 98(462):324–339, 2003.
- [3] Junyoung Chung, Caglar Gulcehre, KyungHyun Cho, and Yoshua Bengio. Empirical evaluation of gated recurrent neural networks on sequence modeling. *arXiv preprint arXiv:1412.3555*, 2014.
- [4] Henggang Cui, Vladan Radosavljevic, Fang-Chieh Chou, Tsung-Han Lin, Thi Nguyen, Tzu-Kuo Huang, Jeff Schneider, and Nemanja Djuric. Multimodal trajectory predictions for autonomous driving using deep convolutional networks. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 2090–2096. IEEE, 2019.
- [5] Corrado Di Natale, Eugenio Martinelli, and Arnaldo D’Amico. Counteraction of environmental disturbances of electronic nose data by independent component analysis. *Sensors and Actuators B: Chemical*, 82(2-3):158–165, 2002.
- [6] Ryo Fujii, Jayakorn Vongkulbhisal, Ryo Hachiuma, and Hideo Saito. A two-block rnn-based trajectory prediction from incomplete trajectory. *IEEE Access*, 9:56140–56151, 2021.
- [7] Francesco Giuliari, Irtiza Hasan, Marco Cristani, and Fabio Galasso. Transformer networks for trajectory forecasting. In *2020 25th international conference on pattern recognition (ICPR)*, pages 10335–10342. IEEE, 2021.
- [8] Niels Haering, Péter L Venetianer, and Alan Lipton. The evolution of video surveillance: an overview. *Machine Vision and Applications*, 19(5):279–290, 2008.
- [9] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- [10] Jihua Huang and H-S Tan. A low-order dgps-based vehicle positioning system under urban environment. *IEEE/ASME Transactions on mechatronics*, 11(5):567–575, 2006.
- [11] Neeraj Kumar Jain, RK Saini, and Preeti Mittal. A review on traffic monitoring system techniques. *Soft computing: Theories and applications: Proceedings of SoCTA 2017*, pages 569–577, 2019.
- [12] Nikolay Jetchev and Marc Toussaint. Trajectory prediction: learning to map situations to robot trajectories. In *Proceedings of the 26th annual international conference on machine learning*, pages 449–456, 2009.
- [13] Hansi Liu, Abrar Alali, Mohamed Ibrahim, Hongyu Li, Marco Gruteser, Shubham Jain, Kristin Dana, Ashwin Ashok, Bin Cheng, and Hongsheng Lu. Lost and found! associating target persons in camera surveillance footage with smartphone identifiers. In *Proceedings of the 19th Annual International Conference on Mobile Systems, Applications, and Services*, pages 499–500, 2021.
- [14] Hansi Liu, Abrar Alali, Mohamed Ibrahim, Bryan Bo Cao, Nicholas Meegan, Hongyu Li, Marco Gruteser, Shubham Jain, Kristin Dana, Ashwin Ashok, et al. Vi-fi: Associating moving subjects across vision and wireless sensors. In *2022 21st ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN)*, pages 208–219. IEEE, 2022.
- [15] Sijia Liu, Sundeep Prabhakar Chepuri, Makan Fardad, Engin Maşazade, Geert Leus, and Pramod K Varshney. Sensor selection for estimation with correlated measurement noise. *IEEE Transactions on Signal Processing*, 64(13): 3509–3522, 2016.
- [16] Yiheng Liu, Wengang Zhou, Mao Xi, Sanjing Shen, and Houqiang Li. Vision meets wireless positioning: Effective person re-identification with recurrent context propagation. In *Proceedings of the 28th ACM International Conference on Multimedia*, pages 1103–1111, 2020.
- [17] Roberto Martin-Martin, Mihir Patel, Hamid Rezatofghi, Abhijeet Sheno, JunYoung Gwak, Eric Frankel, Amir Sadeghian, and Silvio Savarese. Jrdb: A dataset and benchmark of egocentric robot visual perception of humans in built environments. *IEEE transactions on pattern analysis and machine intelligence*, 2021.
- [18] Nishant Nikhil and Brendan Tran Morris. Convolutional neural network for trajectory prediction. In *Proceedings of the European Conference on Computer Vision (ECCV) Workshops*, pages 0–0, 2018.
- [19] Savvas Papaioannou, Hongkai Wen, Zhuoling Xiao, Andrew Markham, and Niki Trigoni. Accurate positioning via cross-modality training. In *Proceedings of the 13th ACM Conference on Embedded Networked Sensor Systems*, pages 239–251, 2015.
- [20] Edoardo Mello Rella, Jan-Nico Zaeck, Alexander Liniger, and Luc Van Gool. Decoder fusion rnn: Context and interaction aware decoders for trajectory prediction. In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 5937–5943. IEEE, 2021.
- [21] Andrey Rudenko, Luigi Palmieri, Michael Herman, Kris M Kitani, Dariu M Gavrilă, and Kai O Arras. Human motion trajectory prediction: A survey. *The International Journal of Robotics Research*, 39(8):895–935, 2020.
- [22] David E Rumelhart, Geoffrey E Hinton, and Ronald J Williams. Learning internal representations by error propagation, parallel distributed processing, explorations in the microstructure of cognition, ed. de rumelhart and j. mccllland. vol. 1. 1986. *Biometrika*, 71:599–607, 1986.
- [23] Zhiyuan Shi, Min Xu, Quan Pan, Bing Yan, and Haimin Zhang. Lstm-based flight trajectory prediction. In *2018 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8. IEEE, 2018.
- [24] Dennis Strelow, Jeffrey Mishler, David Koes, and Sanjiv Singh. Precise omnidirectional camera calibration. In *Proceedings of the 2001 IEEE Computer Society Conference on Computer Vision and Pattern Recognition. CVPR 2001*, pages I–I. IEEE, 2001.
- [25] Hung Tran, Vuong Le, and Truyen Tran. Goal-driven long-term trajectory prediction. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pages 796–805, 2021.

- [26] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- [27] Yi Xu, Armin Bazarjani, Hyung-gun Chi, Chiho Choi, and Yun Fu. Uncovering the missing pattern: Unified framework towards trajectory imputation and prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9632–9643, 2023.
- [28] Haichao Zhang, Yi Xu, Hongsheng Lu, Takayuki Shimizu, and Yun Fu. Layout sequence prediction from noisy mobile modality. In *Proceedings of the 31st ACM International Conference on Multimedia*, pages 3965–3974, 2023.
- [29] Hongsong Zhao, Lingjuan Miao, and Haijun Shao. Adaptive two-stage kalman filter for sins/odometer integrated navigation systems. *The Journal of Navigation*, 70(2):242–261, 2017.
- [30] Qihua Zhou, Song Guo, Jun Pan, Jiacheng Liang, Jingcai Guo, Zhenda Xu, and Jingren Zhou. Pass: Patch automatic skip scheme for efficient on-device video perception. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pages 1–18, 2024.