

# FlexLoc: Conditional Neural Networks for Zero-Shot Sensor Perspective Invariance in Object Localization with Distributed Multimodal Sensors

Jason Wu<sup>\*1</sup>, Ziqi Wang<sup>\*1</sup>, Xiaomin Ouyang<sup>1</sup>, Ho Lyun Jeong<sup>1</sup>, Colin Samplawski<sup>2</sup>,  
Lance M. Kaplan<sup>3</sup>, Benjamin Marlin<sup>2</sup>, and Mani Srivastava<sup>†1,4</sup>

<sup>1</sup>University of California, Los Angeles

<sup>2</sup>University of Massachusetts Amherst

<sup>3</sup>US DEVCOM Army Research Laboratory

<sup>4</sup>Amazon

**Abstract**—Localization is a critical technology for various applications ranging from navigation and surveillance to assisted living. Localization systems typically fuse information from sensors viewing the scene from different perspectives to estimate the target location while also employing multiple modalities for enhanced robustness and accuracy. Recently, such systems have employed end-to-end deep neural models trained on large datasets due to their superior performance and ability to handle data from diverse sensor modalities. However, such neural models are often trained on data collected from a particular set of sensor poses (i.e., locations and orientations). During real-world deployments, slight deviations from these sensor poses can result in extreme inaccuracies. To address this challenge, we introduce FlexLoc, which employs *conditional neural networks* to inject node perspective information to adapt the localization pipeline. Specifically, a small subset of model weights are derived from node poses at run time, enabling accurate generalization to unseen perspectives with minimal additional overhead. Our evaluations on a multimodal, multi-view indoor tracking dataset showcase that FlexLoc improves the localization accuracy by almost 50% in the zero-shot case (no calibration data available) compared to the baselines. The source code of FlexLoc is available in <https://github.com/nesl/FlexLoc>.

## I. INTRODUCTION

Modern robotic and automation applications are heavily dependent on the ability to localize and track a target with a high degree of accuracy. For example, an automated warehouse relies on accurate localization of robots, human workers, and packages to maximize productivity with intelligent scheduling [1]. A delivery or cleaning robot within a smart building can also rely on the building’s localization infrastructure to guide it through a complex floor map [2]. Finally, various healthcare applications are also greatly interested in the location of a human subject, deriving behavioral biomarkers from subject location to provide early warnings for chronic diseases [3].

With these numerous groundbreaking applications all depending heavily on accurate localization, recent works have

explored the utilization of both *localization infrastructure* and *egocentric* localization technologies. Localization infrastructure integrates sensors into the environment, contrasting with egocentric localization that relies upon an on-board suite of sensors to perform techniques such as SLAM [4]. Egocentric localization techniques have been well-studied and achieve exemplary performance [5]. Nevertheless, egocentric sensing is susceptible to occlusions and SWaP constraints (size, weight, and power) that can limit the ability of robots to carry sensors. In this paper, we focus on infrastructure-based localization, which does not impose requirements upon the targets and enables ubiquitous localization for non-cooperative human and robotic targets.

In this work, we employ a set of distributed (multi-view) multimodal sensors for object localization. Existing localization infrastructures typically incorporate multiple sensor nodes with partially overlapping sensor perspectives as a **multi-view** system to improve localization accuracy and expand sensing coverage. These systems estimate the target’s location in a global coordinate system using camera [6], UWB radar [7], or acoustic sensors [8]. Moreover, in real-world scenarios, the perception of a target is heavily influenced by environmental factors, and a single sensor modality is often inadequate in capturing sufficient information to enable robustness. For example, RGB cameras fail to provide relevant information in low-lighting conditions. Therefore, several new **multimodal** sensing systems have been proposed to enhance the reliability of indoor localization by fusing information across multiple distinct sensors [9], [10], [11]. For instance, Zhao et al. [10] enhances camera-based indoor localization with inertial measurement units (IMUs) and WiFi CSI signals.

To handle multi-view multimodal sensor data, localization systems use large-scale end-to-end deep learning models to predict target locations as these data-driven models can handle various sensor modalities and yield superior performance [12]. These models can be divided into **early-fusion** (feature-level fusion) and **late-fusion** (result-level fusion) approaches [13]. Early fusion approaches extract intermediate features from each sensor and fuse them together to generate one prediction. Meanwhile, late fusion approaches predict the target location individually using each sensor, generating

\*The first two authors contributed equally to this paper. Contact: jaysunwu@g.ucla.edu, wangzq312@g.ucla.edu, mbs@ucla.edu

†Mani Srivastava holds concurrent appointments as a Professor of ECE and CS (joint) at the University of California, Los Angeles and as an Amazon Scholar. This paper describes work performed at the University of California, Los Angeles and is not associated with Amazon.




Fig. 1. Performance degradation due to sensor perspective shift.

a prediction by reconciling all the results (e.g., using a Kalman Filter). While late-fusion approaches can minimize the communication overhead in a distributed sensing system, early-fusion approaches often achieve superior performance, as the model facilitates the early collaboration of sensors with different strengths.

**Problem and Motivation:** Despite the increased robustness towards environmental factors provided by learning-enabled multimodal systems, current approaches fail to address the critical issue of **sensor perspective shift**. Deep learning localization models are often trained on datasets collected using a fixed set of sensor positions and orientations (i.e., sensor perspectives), and thus the correspondence between sensor data and the target global coordinates is inherently embedded in the neural network weights. This correspondence breaks down when the model is applied to data collected from an unseen configuration of sensor placements. We illustrate this problem here through a motivational study. We leverage GDTM [14], a multi-hour dataset we collected using distributed multimodal sensor nodes to track a remote controlled car in indoor environments. In this dataset, we introduce diversity in sensor perspectives by altering the position and orientation of sensor nodes between trials. We train both an early-fusion and a late-fusion model on data from one sensor perspective, with the results shown in Figure 1. When we evaluate both models on new data from the *same* viewpoint, we achieve reasonable localization performance. However, both models fail spectacularly when evaluated on sensor data collected from a new sensor placement, with average distance errors of  $>1$  m.

This challenge of sensor perspective shift poses a significant barrier to the practicality of localization systems. Not only can sensor pose gradually shift over time, factors such as relocation after maintenance are inevitable during the deployments. Attempting to train a new system from scratch for each new set of sensor perspectives is highly impractical due to the large amount of data required. Thus, we must **design models that can automatically adapt to sensor perspective shift without any training on the new, unseen perspective** (i.e., in a zero-shot manner). GDTM [14] attempted a solution with more than 30 cm of error on new perspectives with a suboptimal late-fusion approach: localizing the target at each distributed sensor within its local coordinate system and applying a hard-coded coordinate transform into global coordinates.

**Proposed Method:** To address this challenge, we propose FlexLoc, a new multimodal multi-view sensor fusion approach that is robust to the perspective shift of nodes within a distributed multimodal sensor network, and is compatible with the end-to-end early fusion approach. As shown in Figure 2, our main idea involves injecting




Fig. 2. Illustration of our key idea to enable perspective invariant object localization. Unlike existing approaches that are unaware of sensor perspectives, we inject node pose information into the network through conditional neural networks.

node orientation and position (pose) information into the localization network. During the training process, the model is exposed to a large variety of viewpoints with associated pose information to learn how the sensor’s pose affects the prediction of global coordinate location from multimodal sensor data. Then, during test time, the system measures the current pose (through self-localization [15] or marker-based [16] techniques) and utilizes that information to make accurate predictions on unseen sensor viewpoints.

Specifically, FlexLoc builds upon a feature-level (early) fusion architecture primarily consisting of backbones, adapters, a transformer encoder, and an output head. From a high level, the backbones (implemented using ResNets and ViTs [17]) extract features from raw sensor data. The features are mapped to fixed-sized vectors by the adapters and then aggregated. A multi-layer transformer encoder then processes the aggregated features. Finally, the output head yields the predicted global coordinates. To inject the node pose information, we design and insert **Conditional Neural Networks** into the aforementioned architecture. The key idea is to augment the main neural network with a small subset of additional weights conditioned on auxiliary inputs (sensor pose in our scenario). These dynamic weights are generated from the auxiliary input using a lightweight neural network during both training and runtime. After model training, the conditional neural networks can leverage supplementary auxiliary information during inference at runtime, thereby achieving greater robustness over the unconditional models relying solely on the original sensor data input.

Conditional neural networks have elevated the performance of applications such as Visual Question Answering or Speech Enhancement [18], [19], [20]. With FlexLoc, we seek to transfer this idea to the sensor perspective shift problem, which contains increased complexity due to the fragility of the learned 3D spatial relationships. Specifically, we create two variants of conditional neural networks. First, drawing inspiration from the multiplicative nature of coordinate transforms, we propose a *Conditional 1D Convolution* (CondConv) whose convolutional weights are conditioned on the sensor pose. These CondConv layers are inserted after the adapter to transform the extracted sensor features with the sensor pose information. Second, we propose a lightweight design for injecting the pose information through *Conditional Layer Normalization* (CLN), given the pervasive existence of normalization layers in the backbone models. We replace the learnable parameters in the normalization layers, i.e., scale ( $\gamma$ ) and offset ( $\beta$ ), with values derived from the pose. The lightweight CLN approach is the better choice when computing resources are limited, while the CondConv

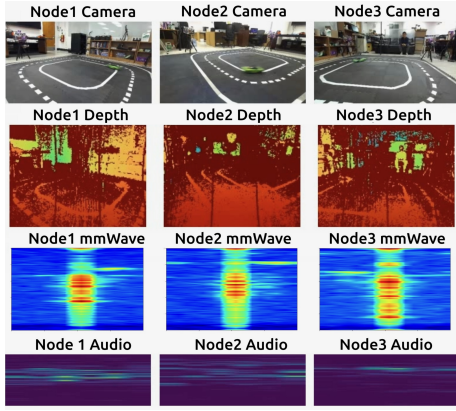


Fig. 3. Examples of multi-modal sensor data collected by three nodes.

method performs better in various complex scenarios.

To evaluate FlexLoc, we leverage our previously collected multi-modal multi-view dataset for indoor geospatial vehicle tracking [14]. Our dataset contains a total of nine-hours of data collected from three nodes that are equipped with four sensor modalities (RGB camera, depth camera, mmWave radar, and microphone), with 22 groups of unique sensor node perspectives (one perspective is an arrangement of 3 nodes). Our results in Section III show that our approach can achieve good zero-shot performance on data collected in unseen perspectives while also exhibiting test-time robustness to different combinations of sensor modalities.

In summary, we make the following key contributions:

- We propose two methods for improving test-time performance of localization models on unseen sensor perspectives: Conditional 1D Convolution and a lightweight design based on Conditional Layer Normalization. This work constitutes the first exploration into applying conditional neural networks to address test-time robustness in localization models.
- We independently integrate the proposed two conditional methods into an early fusion architecture for multimodal multi-view learning. We then perform a detailed study investigating the performance of the two methods.
- FlexLoc involves a minimal number of additional parameters (at most 0.2% of model parameters), and its simple implementation allows for easy integration within existing neural networks
- FlexLoc shows good zero-shot generalization performance to various perspectives while also retaining good performance in the absence of various sensor modalities.

## II. METHODOLOGY

### A. Multimodal Multiview Dataset

Public multimodal, multiview datasets containing shifting perspectives of sensors are scarcely available. We collected a nine-hour multimodal indoor tracking dataset, namely GDTM [14], to research the sensor perspective shift problem. GDTM is composed of three data collection nodes placed at various locations around an indoor race track to localize a remote-controlled car. Each node is composed of camera, depth, mmWave, and audio modalities (Fig. 3). Additionally,

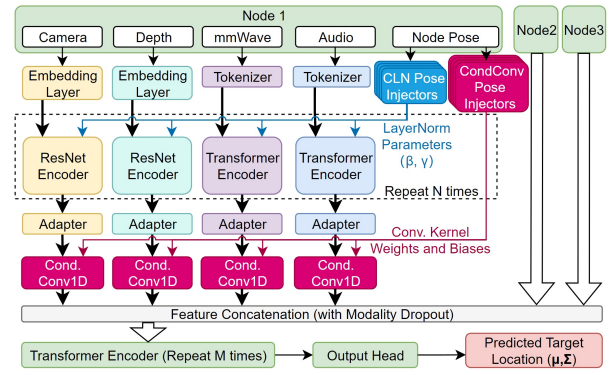


Fig. 4. Complete FlexLoc architecture containing Conditional Convolution and Conditional Layer Normalization. Conditional Convolution derives its 1D convolutional kernel weights from the sensor pose and transforms the extracted features of the sensor data. Conditional Layer Normalization is a more lightweight design integrated into the backbones, where we replace the learnable parameters  $\gamma$  and  $\beta$  with values derived from sensor pose.

the position and orientation of these nodes are adjusted during data collection to generate a large diversity of sensor viewpoints (perspectives of the three nodes). We utilize this variety of viewpoints to develop our perspective invariant model. GTDM contains 22 unique viewpoints, allowing us to construct a dataset containing 13 views in the training dataset (105 minutes), 4 views in the validation set (35 minutes), and 5 views for testing (55 minutes). All three sets contain both circular motions and random trajectories of the robot car moving in the scene.

We quantified the similarity between perspectives  $V_1$  and  $V_2$  by computing the *structural similarity index measure* (SSIM) of the camera data: (1) we take a frame of camera data from each sensor node in  $V_1$  and  $V_2$ , resulting in three images in each view. (2) We compute pair-wise SSIM  $\mathbf{S} \in \mathbb{R}^{3 \times 3}$ . (3) We take  $\max(\mathbf{S})$ , the SSIM of the most similar image pair, as the similarity between  $V_1$  and  $V_2$ . We introduce a two-stage procedure to ensure that the viewpoints used for testing are different from those for training and validation. First, we calculated the similarity score between  $V_{train}$ ,  $V_{val}$ , and  $V_{test}$  to ensure they are all smaller than 0.60. Moreover, we manually performed a visual inspection to double-check.

### B. Overall Model Architecture

Figure 4 showcases the overall architecture of FlexLoc, which outputs predictions of the target location in pre-established global coordinates with the help of sensor pose information. First, FlexLoc utilizes *backbones* shared across all the nodes to process each modality independently with the goal of extracting key information (embeddings) pertinent to the localization task. The *adapters* (shared among nodes) accept these embeddings from the backbones as input and perform modality-specific processing to transform the distinct feature vectors into vectors of the same dimension. The output of the adapters are then processed by *Conditional Convolution* to introduce pose information. Afterwards, a *transformer encoder* performs cross-modal information fusion. Finally, an *output head* utilizes the information in the joint embedding to make a unified prediction of the car's

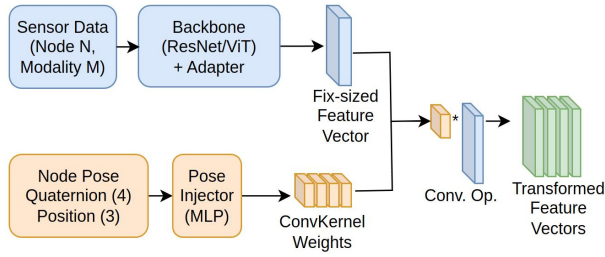


Fig. 5. Implementation and Integration of Conditional 1D Convolution.

location in the pre-established global coordinate system. As an alternative to Conditional Convolution, we also propose a lightweight conditional neural network based on Conditional Layer Normalization, which is integrated into backbones to incrementally introduce pose information into the final embeddings of different sensor modalities (drawn in the same figure to save space).

Specifically, our backbones are composed of either Residual Convolutional Networks (ResNet) [21], or transformer encoders following the work on Vision Transformers [17]. ResNet backbones are a good candidate for processing visual data (camera and depth information) due to their ability to extract information about key objects in the scene while ignoring irrelevant background information. In contrast, when processing mmWave and audio sensor data, we employ a simpler stack of transformer encoders due to the relative lack of background content.

Moreover, Conditional Convolution or Conditional Layer Normalization are strategically inserted into this architecture to incorporate pose information. Conditional Convolution operates on the extracted embeddings containing dense information regarding the object’s position, injecting the sensor pose information through a series of 1D convolutional kernels. In contrast, Conditional Layer Normalization serves as a lightweight method of incorporating pose information into the early stages of data processing. We replace the traditional Batch Normalization and Layer Normalization components within each backbone with Conditional Layer Normalization that derives its parameters from the sensor pose information. We evaluate the performance of the two conditional methods within Section III.

### C. Conditional Convolution

To effectively inject the sensor pose information, we explore transforming the feature vectors of a particular node (outputs of the adapters in Fig. 4) with a *conditional 1D convolution*. Given an input vector  $\mathbf{x} \in \mathbb{R}^N$  from the intermediate features of  $M$  sensor modalities  $\mathbf{X} = \{\mathbf{x}_0, \dots, \mathbf{x}_{M-1}\} \in \mathbb{R}^{M \times N}$ , 1D convolution with  $K$  kernels of size  $S$  generates an output  $\mathbf{y} \in \mathbb{R}^{K \times N}$  according to Equation 1, where  $*$  refers to the convolutional operator,

$$\mathbf{y}_k = b_k + \mathbf{x} * \mathbf{w}_k, \quad \mathbf{w}_k \in \mathbf{W}, b_k \in \mathbf{b}. \quad (1)$$

*Conditional* 1D convolution replaces the learnable weights  $\mathbf{W} \in \mathbb{R}^{K \times S}$  and biases  $\mathbf{b} \in \mathbb{R}^K$  with parameters generated from the auxiliary input (sensor pose), fusing the information within  $\mathbf{X}$  with additional context. Figure 5 provides a detailed snapshot of how conditional convolution is integrated

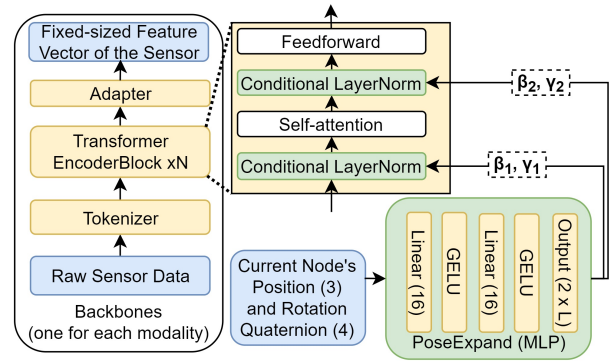


Fig. 6. Implementation and Integration of Conditional Layer Normalization.

into the overall system. The node pose is represented by a set of 7 values, including four quaternion values representing rotation and  $(x, y, z)$  position. These values are expanded by a series of linear layers into a higher dimension, from which we extract the weights and biases for the  $K$  convolutional kernels of length  $S$ . These adaptive kernels are convolved with the feature vector to generate  $K$  new feature vectors with position information.

### D. Conditional Layer Normalization

Layer Normalization [22] is a feature normalization technique that reduces training time and improves generalization of neural networks [23]. It performs normalization with the mean and standard deviation computed over elements of the same batch, while also introducing a learnable *per element scale and bias* to improve performance. Equation 2 outlines applying Layer Normalization to a single vector  $\mathbf{x} \in \mathbb{R}^N$  from a batch of intermediate feature vectors  $\mathbf{X} = \{\mathbf{x}_0, \dots, \mathbf{x}_{B-1}\} \in \mathbb{R}^{B \times N}$ :

$$\mathbf{y} = \gamma \cdot \frac{\mathbf{x} - \mu}{\sigma} + \beta, \quad \mu = \frac{1}{N} \sum_{j=0}^{N-1} x_j, \quad \sigma = \frac{1}{N} \sqrt{\sum_{j=0}^{N-1} (x_j - \mu)^2} \quad (2)$$

with  $\gamma \in \mathbb{R}^N$  and  $\beta \in \mathbb{R}^N$  acting as *per element* scale and shift operations.

We implement *Conditional Layer Normalization* (CLN) by replacing the learnable  $\gamma$  and  $\beta$  parameters with values derived from the node’s pose information. Figure 6 illustrates how the layer normalization blocks contained within each transformer encoder of the backbone are replaced with their adaptive counterparts. This same process occurs with the Batch Normalization layers in the ResNet. These adaptive weights are derived from a *separate* set of linear layers (PoseExpand). PoseExpand encodes the sensor pose into an vector of dimension 16, from which the unique  $\gamma$  and  $\beta$  value of each Layer Normalization block are obtained. Note that to reduce the number of additional parameters, we use *one single*  $\gamma$  scalar and *one single*  $\beta$  scalar to modify the entire embedding instead of performing per element scale and shifts. Furthermore, in vanilla LayerNorm, the  $\gamma$  and  $\beta$  parameters are shared across the batch  $B$ . In CLN, however, data entries within a batch are not guaranteed to come from the same sensor viewpoint. As a result, we compute one  $\gamma_i$  and one  $\beta_i$  scalar for each entry  $i$  in the batch. Thus, with

Models Architecture	View 1		View 2		View 3		View 4		View 5		Average	
	Mean	Std	Mean	Std	Mean	Std	Mean	Std	Mean	Std	Mean	Std
FlexLoc (CondConv)	<b>17.10</b>	0.96	<b>14.45</b>	0.68	<b>13.30</b>	1.06	<b>11.89</b>	1.07	<b>14.67</b>	2.03	<b>14.28</b>	1.16
FlexLoc-light (CLN)	20.50	0.55	24.88	2.78	21.98	4.04	24.38	11.90	18.45	2.33	22.04	4.32
Baseline-early (Brute Force)	32.61	3.57	24.66	3.27	26.52	3.40	24.25	7.38	26.29	5.78	26.86	4.68
Baseline-late (Hard-coded Transform)	33.69	4.97	25.36	1.16	20.93	0.05	32.97	1.23	24.71	1.94	27.53	1.87

TABLE I

EUCLIDEAN DISTANCE ERROR (IN CM, AVERAGED ACROSS 3 SEEDS), FLEXLOC VS BASELINES

an input of  $\mathbf{X} \in \mathbb{R}^{B \times N}$ , every element of  $x_{ij}$  from  $\mathbf{X}$  will be normalized by  $\mu_i$  and  $\sigma_i$ , then scaled and shifted by  $\gamma_i$  and  $\beta_i$ .

### E. Training Strategies

1) *Backbone Pretraining*: We utilize a set of pre-trained weights to imbue our backbones with an understanding of how to process the input sensor data. In the camera and depth ResNet-18 backbones, we load a checkpoint pre-trained on ImageNet1K. As we designed the Transformer Encoder backbones ourselves, we chose to pre-train those backbones on 5 minutes of data from a viewpoint in the training dataset. The pre-training process involved separately training each modality’s backbone to ensure it has an understanding of how to identify the car from sensor data.

2) *Sensor Dropout*: Furthermore, we also utilized a *modality dropout* process during model training to fully realize the robustness of the rich multimodal input. We discovered that training the model with all modalities active resulted in an over-reliance on the most informative modalities. To combat this, we established a 50% chance for any modality to be “dropped-out” in a given batch during training, effectively forcing the model to learn with various subsets of modalities. We analyze this in greater detail within Section III-C.

## III. EVALUATIONS

We evaluate a variety of models on the multiview dataset (Section II-A) to showcase the performance benefit of our conditional networks. We are only concerned with predicting the car’s 2-dimensional coordinate, as the car’s elevation within the global coordinate system is unchanged. Thus, the primary metric for quantifying the performance is *Average Euclidean Distance* between the predicted location and the ground truth. We evaluated all our models on five *unseen sensor perspectives that were not present in training*.

### A. Comparison with Existing Methods

We test our two conditional models and compare them against two baseline models:

- 1) *Unconditional Early Fusion*: This is essentially FlexLoc with no conditional layers.
- 2) *Local Coordinate Late Fusion* [14]: Each modality and node pair localizes the object in the node’s local coordinate system, after which the predictions are transformed into a global coordinate system via a known transform matrix and merged with a Kalman Filter.

Table I showcases that all the conditional models perform better against the baseline models. In particular, Conditional Convolution (CondConv) achieves the best performance with




Fig. 7. Cumulative distribution function (CDF) plot of frame-by-frame localization error of different methods.

at least 46% improvement over the baseline methods. In fact, the localization error of CondConv is comparable to models that are *trained and tested on the same viewpoint*. In Section I, we trained and tested an early-fusion and late-fusion model on one single sensor viewpoint, achieving 12.91 cm and 20.98 cm of error, respectively. CondConv achieves an impressive 32% improvement over the single-perspective late fusion model and comes within 10% of the single-perspective early fusion model. These results exemplify the robustness of CondConv towards completely unseen sensor perspectives. As this is our best performing model, we term FlexLoc as the network containing only CondConv.

While Conditional Layer Normalization (CLN) does not achieve the same level of performance as Conditional Convolution, it can still provide a modest 17.9% improvement over the baseline early fusion model lacking conditional layers. We hypothesize that its worse performance compared to CondConv is due to the more challenging task of incorporating pose information into the backbone itself. One key advantage of CLN lies in its lightweight nature, where it employs far fewer parameters than CondConv (analyzed in Section III-D). Thus, we term CLN as FlexLoc-Light. It is also worth noting that FlexLoc-Light encounters an unlucky seed where the model performance degrades, which contributed to the high standard deviation, highlighting that CLN may exhibit less stable training than CondConv.

Figure 7 showcases a visual depiction of each method’s localization error in a CDF. We observe a significant boost in performance between FlexLoc and the other methods, while FlexLoc-Light maintains a smaller yet notable improvement over the two baseline methods. Another sizable advantage of FlexLoc lies in its low 90% error at only 25 cm, which is 20 cm less error than the next closest baseline method.

### B. Different Methods of Fusing Conditional Information

We also compare FlexLoc’s performance against various existing methods of integrating sensor pose information to showcase the performance benefit of the conditional layers. As the techniques we utilize require transformer encoders, we shift to an alternate model where the ResNets are removed and all the backbones are composed of transformer




Fig. 8. Comparisons of different methods to inject pose information.

encoders. We implement and evaluate two networks inspired from adjacent works targeting perspective shift in other contexts, while stressing that these works do not attempt to address the issue of sensor pose shift for *localization*:

- 1) *Cross Attention Pose Fusion* [24]. This work incorporates camera pose into a diffusion model through cross-attention for improved video generation. We utilize this concept to fuse camera pose information through cross attention layers integrated into the transformer backbones.
- 2) *Pose Concatenation with Input* [25]. This work concatenates camera pose information into the input to improve monocular depth map prediction on unseen camera viewpoints. We implement their technique by encoding the camera pose as a single token that is concatenated to the sensor data tokens to be fed into the transformer backbones.

In Figure 8, we observe that integrating pose information through FlexLoc (CondConv) and FlexLoc-Light (CLN) achieves approximately 20% improvement over the next best method of concatenating prompts to the input. Furthermore, the Cross Attention method performs poorly with an error of 56 cm. The inferior performance of the other methods exemplifies the critical impact of carefully selecting the right method for incorporating node pose information.

### C. Impact of different training strategies

In Table II, we analyze the impact of loading pre-trained weights to initialize the network. In Section II-E, we introduced a pretraining strategy where we loaded pre-trained weights into the backbone to aid in identifying the car. When training CLN without this strategy, the model performance degrades by 7.1 cm (24%), highlighting the importance of starting from a good set of weights for CLN. In contrast, CondConv was not adversely impacted with a performance loss of only 0.7 cm. The difference between these two models is likely due to CLN’s position within the backbone, where it is closely coupled with the backbone weights and thus highly sensitive to their values. CondConv is placed *after* the backbone layers and is more insulated from the backbone weights.

Figure 9 illustrates the impact of the *sensor dropout* training method outlined in Section II-E. Without sensor dropout, the model over-relies on the camera modality and largely ignores other modalities, incurring catastrophic error when testing on combinations of modalities without camera (150 cm for depth and 60 cm for mmWave). Applying sensor dropout increases the robustness of the model to




Fig. 9. FlexLoc performance with and without dropout when only a subset of sensing modalities are available.

	From Pretrained Backbones	From Scratch
FlexLoc (ConcConv)	<b>14.28</b>	14.92
FlexLoc-light (CLN)	<b>22.04</b>	29.14

TABLE II  
EFFECT OF PRETRAINING

missing modalities, allowing it to maintain good localization performance despite the loss of the critical camera modality.

### D. Number of Extra Network Parameters

Both of the conditional components utilize a set of linear layers to derive their weights from the provided pose information. CondConv utilizes a wider set of linear layers with a total of 52992 parameters, while CLN is more lightweight with only 3328 parameters. However, CLN **replaces** the BatchNorm2D layers within the ResNet-18 architecture that each have more learnable  $\gamma$  and  $\beta$  parameters (Section II-D), resulting in *fewer model parameters compared to a network without conditional layers*. Table III illustrates that CondConv results in 0.19% more model parameters and a computational overhead of 0.57% more Multiply-and-Accumulate operations (MACs). In contrast, CLN removes parameters from the network, containing 0.04% *fewer* parameters and 0.32% *fewer* MACs. In smaller networks where these parameters and operations are more significant, CLN will be a better choice than CondConv despite the worse performance. Interestingly, unifying both CondConv and CLN into the same network resulted in worse performance than CondConv alone. This is likely because CLN and CondConv integrate similar information into the network, causing a saturation of performance, and the increased difficulty of learning with two conditional components hurts performance. Due to the similar overhead and worse performance of the unified model compared to CondConv, we choose to keep CLN and CondConv as separate architectures.

Models	Performance		Additional Overhead	
	Localization Error	Parameters	Parameters	MACS (M)
FlexLoc (CondConv)	14.28 (-46.84%)	53904 (+0.19%)	95.46 (+0.57%)	
FlexLoc-light (CLN)	22.04 (-17.96%)	-11788 (-0.04%)	-54.04 (-0.32%)	
CondConv + CLN	16.11 (-40.02%)	42116 (+0.15%)	41.28 (+0.24%)	

TABLE III  
ADDITIONAL OVERHEAD INCURRED BY FLEXLOC IN TERMS OF  
PARAMETERS AND COMPUTATIONS.

## IV. RELATED WORK

### A. Multimodal Learning for Localization and Tracking

Machine learning algorithms based on handcrafted features and deep neural networks have been applied to indoor localization and tracking applications. For example, *Samplawski et al.* [6] used three different cameras to predict the distribution of the object’s location, and designed a *multi-view* late fusion neural network to fuse the prediction of different cameras. Although they demonstrated excellent tracking results, the solution utilizes only RGB cameras and will fail under low-lighting conditions. Therefore, several new multimodal sensing systems have been proposed to leverage the strengths of different sensor modalities to enhance the reliability of indoor localization and tracking [9], [10], [11]. For instance, [9] proposed a method to locate an indoor mobile phone user based on multi-view and multimodal measurements, but incurs a 1-meter error.

However, aside from [14], there exists minimal work in the field of multimodal and multiview localization through localization infrastructures. Existing work typically revolves around either unimodal multi-view localization infrastructures [26], [27], multimodal single-view localization infrastructures [28], [29], or localization through egocentric sensors [30], [31]. The few multimodal and multiview works [32], [33] do not contain shifting sensor viewpoints or a large diversity of multimodal sensors. Thus, we were unable to perform evaluations on different datasets or find similar architectures that attempted to address the challenge of perspective shift in multimodal multi-view localization.

### B. Conditional Neural Networks

Conditional neural networks incorporate additional auxiliary information during the training of the network to enhance the robustness or accuracy of predictions, demonstrating excellent performance in applications such as Visual Question Answering or Speech Enhancement [18], [19], [20]. For example, De Vries et. al [18] introduces conditional batch normalization to modulate convolutional feature maps of visual processing with a linguistic embedding, which significantly improves the performance of visual question-answering tasks. Zhang et. al [19] integrates a deep auto-encoder with neural noise embedding for speech enhancement through conditional layer normalization, which improves the generalization of the trained model to various noisy speech inputs. Therefore, by training with a small portion of additional weights conditioned on the auxiliary input, the conditional neural networks improve the model accuracy and generalization ability. FlexLoc is the only work that applies conditional neural networks to the task of perspective invariant localization.

## V. LIMITATIONS, FUTURE WORK, AND DISCUSSIONS

1) *Sensor Pose Estimation from Sensing Data:* We currently measure the node pose with an OptiTrack motion capture system and provide this ground truth measurement at test time for a new configuration of node viewpoints.

Unfortunately, such motion capture systems are not readily available in practical deployment scenarios due to cost and calibration requirements. We can address this lack of external pose information by implementing self-localization techniques using FlexLoc’s on-board sensor suite to derive its own pose information by employing on-board IMUs, AprilTag markers [16] or through markerless techniques such as iNERF [15]. The choice of self-localization technique will be influenced by hardware costs, available compute, and control over the environment. Additionally, it is important to understand how errors in pose estimation impact the final result. The addition of these techniques will create a fully end-to-end localization system for perspective invariance.

### 2) *Extension to More Comprehensive Environments:*

While we achieved exemplary localization results in the presence of sensor perspective shift, we acknowledge that the setting of the dataset (single car driving around an indoor track) is fairly simplistic. Future work can explore training conditional neural networks on datasets containing more complex scenes, from indoor environments with diverse settings and occlusions, outdoor environments with various weather conditions or landscapes, or even computer-simulated environments. Furthermore, it will be interesting to evaluate FlexLoc’s performance in the case of *multi-object localization*, which can provide critical information to applications such as robotic swarms [34].

3) *Comparison to Late Fusion Architectures:* While we primarily employed an early-fusion strategy due to superior performance, late fusion can be a superior option in certain scenarios. In a distributed computing setting, late fusion has its traditional advantage of flexible composition (independent predictions enable a dynamic number of sensors) and minimal communication overhead by exchanging only the final predictions. We sought to absorb these advantages of late fusion into the design of FlexLoc. First, we used a transformer encoder architecture to enable flexible composition, as transformers are structured to process arbitrary lengths of features. Second, FlexLoc leverages a backbone-adaptor-fusion architecture. Each node can choose to either stream the raw sensor data to the central server or compute the features (256 length embeddings) locally and then stream the features to reduce communication overhead. We envision future work to investigate optimal strategies based on the capacity of the nodes and servers.

## VI. CONCLUSIONS

In this paper, we propose FlexLoc, a multi-view multimodal object localization system that is robust to sensor perspective shifts. FlexLoc employs *conditional neural networks* to inject node perspective information into the localization pipeline, which can automatically adapt to unseen sensor perspectives without additional calibration data, achieving zero-shot generalization. Our evaluations on a self-collected multimodal, multi-view indoor tracking dataset show that FlexLoc outperforms existing baselines by almost 50% in localization accuracy while incurring minimal additional overhead.

## ACKNOWLEDGMENT

The research reported in this paper was sponsored in part by: the IoBT REIGN Collaborative Research Alliance funded by the Army Research Laboratory (ARL) under Cooperative Agreement W911NF-17-2-0196; the National Science Foundation (NSF) under award #1822935; and the Air Force Office of Scientific Research (AFOSR) under award FA9550-22-1-0193. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the funding agencies.

## REFERENCES

- [1] C. Rohrig and S. Spieker, "Tracking of transport vehicles for warehouse management using a wireless sensor network," in *IEEE/RSJ IROS 2008*. IEEE, 2008, pp. 3260–3265.
- [2] K. Lin, M. Chen, J. Deng, M. M. Hassan, and G. Fortino, "Enhanced fingerprinting and trajectory prediction for iot localization in smart buildings," *IEEE Transactions on Automation Science and Engineering*, vol. 13, no. 3, pp. 1294–1307, 2016.
- [3] X. Ouyang, X. Shuai, Y. Li, L. Pan, X. Zhang, H. Fu, X. Wang, S. Cao, J. Xin, H. Mok, *et al.*, "Admarker: A multi-modal federated learning system for monitoring digital biomarkers of alzheimer's disease," *arXiv preprint arXiv:2310.15301*, 2023.
- [4] H. Durrant-Whyte and T. Bailey, "Simultaneous localization and mapping: part i," *IEEE Robotics & Automation Magazine*, vol. 13, no. 2, pp. 99–110, 2006.
- [5] Y. Chen, J. Tang, C. Jiang, L. Zhu, M. Lehtomäki, H. Kaartinen, R. Kaijaluoto, Y. Wang, J. Hyppä, *et al.*, "The accuracy comparison of three simultaneous localization and mapping (slam)-based indoor mapping technologies," *Sensors*, vol. 18, no. 10, p. 3228, 2018.
- [6] C. Samplawski, S. Fang, Z. Wang, D. Ganesan, M. Srivastava, and B. M. Marlin, "Heteroskedastic geospatial tracking with distributed camera networks," in *Proceedings of the Thirty-Ninth Conference on Uncertainty in Artificial Intelligence*, ser. UAI '23. JMLR.org, 2023.
- [7] M. J. Bocus and R. J. Piechocki, "Passive unsupervised localization and tracking using a multi-static uwb radar network," in *2021 IEEE Global Communications Conference*. IEEE, 2021, pp. 01–06.
- [8] T. De Groot, "Localization and classification using an acoustic sensor network," Ph.D. dissertation, Delft University of Technology Delft, The Netherlands, 2010.
- [9] Z. Liu, L. Cheng, A. Liu, L. Zhang, X. He, and R. Zimmermann, "Multiview and multimodal pervasive indoor localization," in *Proceedings of the 25th ACM international conference on Multimedia*, 2017, pp. 109–117.
- [10] Y. Zhao, J. Xu, J. Wu, J. Hao, and H. Qian, "Enhancing camera-based multimodal indoor localization with device-free movement measurement using wifi," *IEEE Internet of Things Journal*, vol. 7, no. 2, pp. 1024–1038, 2019.
- [11] M.-H. Amri, Y. Becis, D. Aubry, and N. Ramdani, "Indoor human/robot localization using robust multi-modal data fusion," in *2015 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2015, pp. 3456–3463.
- [12] A. Vakil, J. Liu, P. Zulch, E. Blasch, R. Ewing, and J. Li, "A survey of multimodal sensor fusion for passive rf and eo information integration," *IEEE Aerospace and Electronic Systems Magazine*, vol. 36, no. 7, pp. 44–61, 2021.
- [13] K. Gadzicki, R. Khamsehashari, and C. Zetsche, "Early vs late fusion in multimodal convolutional neural networks," in *2020 IEEE 23rd international conference on information fusion*. IEEE, 2020, pp. 1–6.
- [14] H. L. Jeong, Z. Wang, C. Samplawski, J. Wu, S. Fang, L. Kaplan, D. Ganesan, B. Marlin, and M. Srivastava, "Gtdm: An indoor geospatial tracking dataset with distributed multimodal sensors," in *arXiv preprint arXiv:2402.14136*, 2024.
- [15] L. Yen-Chen, P. Florence, J. T. Barron, A. Rodriguez, P. Isola, and T.-Y. Lin, "Inerf: Inverting neural radiance fields for pose estimation," in *IEEE/RSJ IROS 2021*, 2021, pp. 1323–1330.
- [16] E. Olson, "Apriltag: A robust and flexible visual fiducial system," in *2011 IEEE International Conference on Robotics and Automation*, 2011, pp. 3400–3407.
- [17] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, *et al.*, "An image is worth 16x16 words: Transformers for image recognition at scale," *arXiv preprint arXiv:2010.11929*, 2020.
- [18] H. De Vries, F. Strub, J. Mary, H. Larochelle, O. Pietquin, and A. C. Courville, "Modulating early visual processing by language," *Advances in neural information processing systems*, vol. 30, 2017.
- [19] Z. Zhang, X. Li, Y. Li, Y. Dong, D. Wang, and S. Xiong, "Neural noise embedding for end-to-end speech enhancement with conditional layer normalization," in *ICASSP 2021 - 2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2021, pp. 7113–7117.
- [20] E. Perez, F. Strub, H. de Vries, V. Dumoulin, and A. Courville, "Film: Visual reasoning with a general conditioning layer," *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 32, no. 1, Apr. 2018.
- [21] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [22] J. L. Ba, J. R. Kiros, and G. E. Hinton, "Layer normalization," *arXiv preprint arXiv:1607.06450*, 2016.
- [23] K. Lyu, Z. Li, and S. Arora, "Understanding the generalization benefit of normalization layers: Sharpness reduction," *Advances in Neural Information Processing Systems*, vol. 35, pp. 34 689–34 708, 2022.
- [24] S. Yang, L. Hou, H. Huang, C. Ma, P. Wan, D. Zhang, X. Chen, and J. Liao, "Direct-a-video: Customized video generation with user-directed camera movement and object motion," *arXiv preprint arXiv:2402.03162*, 2024.
- [25] Y. Zhao, S. Kong, and C. Fowlkes, "Camera pose matters: Improving depth prediction by mitigating pose distribution bias," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 15 759–15 768.
- [26] M. Salimibeni, Z. Hajiakhondi-Meybodi, P. Malekzadeh, M. Atashi, K. N. Plataniotis, and A. Mohammadi, "Iot-td: Iot dataset for multiple model ble-based indoor localization/tracking," in *2020 28th European Signal Processing Conference (EUSIPCO)*, 2021, pp. 1697–1701.
- [27] K. Ngamakeur, S. Yongchareon, J. Yu, and S. Islam, "Passive infrared sensor dataset and deep learning models for device-free indoor localization and tracking," *Pervasive and Mobile Computing*, vol. 88, p. 101721, 2023.
- [28] Z. Kandyllakis, K. Vasili, and K. Karantzalos, "Fusing multimodal video data for detecting moving objects/targets in challenging indoor and outdoor scenes," *Remote Sensing*, vol. 11, no. 4, p. 446, 2019.
- [29] K. Nakamura, K. Nakadai, F. Asano, and G. Ince, "Intelligent sound source localization and its application to multimodal human tracking," in *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2011, pp. 143–148.
- [30] H. Caesar, V. Bankiti, A. H. Lang, S. Vora, V. E. Liong, Q. Xu, A. Krishnan, Y. Pan, G. Baldan, and O. Beijbom, "nuscnets: A multimodal dataset for autonomous driving," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 11 621–11 631.
- [31] P. Sun, H. Kretzschmar, X. Dotiwalla, A. Chouard, V. Patnaik, P. Tsui, J. Guo, Y. Zhou, Y. Chai, B. Caine, *et al.*, "Scalability in perception for autonomous driving: Waymo open dataset," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 2446–2454.
- [32] M. J. Bocus, W. Li, S. Vishwakarma, R. Kou, C. Tang, K. Woodbridge, I. Craddock, K. McConville, *et al.*, "Operanet, a multimodal activity recognition dataset acquired from radio frequency and vision-based sensors," *Scientific data*, vol. 9, no. 1, p. 474, 2022.
- [33] C. Torres, J. C. Fried, K. Rose, and B. S. Manjunath, "A multiview multimodal system for monitoring patient sleep," *IEEE Transactions on Multimedia*, vol. 20, no. 11, pp. 3057–3068, 2018.
- [34] M. Schranz, M. Umlauf, M. Sende, and W. Elmenreich, "Swarm robotic behaviors and current applications," *Frontiers in Robotics and AI*, vol. 7, p. 36, 2020.