

# UAV-SAR: Simultaneous Radar-Based Odometry and Synthetic-Array Sensing for Unmanned Aerial Vehicles

David Hunt, Shaocheng Luo, Samuel Rivera, Aarav Prakash<sup>†</sup>, Cameron Morris<sup>†</sup>,  
Tingjun Chen, and Miroslav Pajic

**Abstract**—Unmanned aerial vehicles (UAVs) require accurate odometry—i.e., estimating the position and velocity of the vehicle over time—as well as high-resolution sensing to safely and effectively operate in complex environments. Traditionally, GPS, cameras, and/or lidar sensors have been used to perform these functions. However, GPS can be jammed in contested environments while cameras and lidars fail in visually degraded conditions, limiting UAV operations in these scenarios. In this work, we present UAV-SAR, a unified architecture that utilizes mmWave radars to *simultaneously* achieve precise odometry measurements and perform high-resolution synthetic-array sensing. Here, UAV-SAR measures a UAV’s altitude and velocity from downward- and outward-facing radars and fuses these measurements within a commercially available flight controller to produce accurate odometry estimates. These odometry estimates are then used to dynamically construct synthetic arrays by coherently integrating multiple radar frames together over a duration of 0.5 s, improving the angular resolution by an order of magnitude compared to the physical array alone. Finally, a lightweight deep learning model is utilized to convert high-resolution range-angle responses into 2D point clouds suitable for downstream perception tasks. UAV-SAR is validated on a custom UAV prototype where it is integrated with ROS2 and the PX4 autopilot to demonstrate stable flight, reliable odometry, and high-resolution radar sensing in indoor environments.

## I. INTRODUCTION

Unmanned aerial vehicles (UAVs) are increasingly deployed across diverse domains, including emergency response, military reconnaissance, logistics, and package delivery [1]. Effective operation in these domains requires accurate odometry—i.e., precise position and velocity estimates—in addition to high-resolution sensing, typically obtained via cameras, light detection and ranging (lidar), and/or GPS. Yet, lidar and camera sensors fail in low-visibility conditions (e.g., smoke, fog, dust) and GPS sensors can be jammed in contested environments [2]–[4].

Radio detection and ranging (radar) sensors provide a promising alternative, as they can directly measure an object’s velocity and achieve reliable sensing regardless of lighting or visibility. Compact and lightweight, millimeter-wave (mmWave) radars offer UAVs high range resolutions (4 cm) and fine velocity resolution ( $\leq 0.1$  m/s). Nevertheless, commercial mmWave radar sensors typically use small

antenna arrays whose limited angular resolution ( $14.3^\circ$ ) constrains high-resolution angular sensing [5].

Prior works have separately addressed radar-based high-resolution sensing and odometry. To start, deep learning approaches [6]–[13] have generated high-resolution point clouds from low-resolution radar frames; however, their performance is inherently limited by low-resolution single-frame input and requires training on massive datasets (often  $>50,000$  samples). On the other hand, synthetic aperture radar (SAR) techniques leverage UAV motion to synthesize a large aperture array from a single onboard sensor [14]–[25]. By coherently combining returns along the UAV’s trajectory, SAR may achieve finer spatial resolution than the physical array alone. Yet, these approaches typically rely on GPS or camera-based localization, limiting their applicability in visually degraded or GPS-denied environments.

Alternatively, radar-based UAV odometry leverages the sensor’s ability to measure the velocities of static environmental objects [26]–[32]. To ensure stable flight, commercial controllers recommend odometry updates at least every 10 Hz [33], constraining frame durations to 100 ms. This temporal constraint hinders synthetic array beamforming, which requires longer integration periods to improve angular resolution. Consequently, using a single radar configuration for both *radar-only* synthetic-array sensing and odometry is challenging due to these conflicting requirements. To our knowledge, no prior work has simultaneously achieved high-resolution synthetic-array sensing and accurate radar-based odometry on a UAV.

In this work, we introduce UAV-SAR, a novel architecture for jointly achieving precise odometry and high-resolution synthetic-array sensing using UAV-mounted mmWave radar. We first enhance radar-based odometry techniques to obtain accurate altitude and velocity estimates from downward- and outward-facing radars. These measurements are fed into an industry-standard flight controller to generate precise UAV state estimates. Using this odometry, we dynamically compute the geometry of synthetic arrays formed by UAV motion, coherently integrating 5 radar frames to improve angular resolution by an order of magnitude. Finally, we demonstrate how lightweight deep learning (DL) models can then generate dense, high-resolution point clouds suitable for downstream mapping tasks while requiring  $310\times$  fewer training samples compared to previous works such as [11].

To demonstrate the real-world feasibility of UAV-SAR, we implemented it on a custom UAV prototype, executing as a real-time ROS 2 package fully integrated with the PX4

This work is sponsored in part by the ONR under agreement N00014-23-1-2206, AFOSR award number FA9550-19-1-0169, and by the NSF under NAIAD Award 2332744 and the National AI Institute for Edge Computing Leveraging Next Generation Wireless Networks, Grant CNS-2112562.

The authors are with the Department of Electrical and Computer Engineering, Duke University, Durham, NC 27708 USA (e-mail: {david.hunt, shaocheng.luo, samuel.rivera, tingjun.chen, miroslav.pajic}@duke.edu).  
<sup>†</sup>These authors are with the Durham Academy, Durham, NC 27705 USA (e-mail: {26prakash,26morris}@duke.edu).




Fig. 1: Overview of UAV-SAR: ADC data cubes ( $\mathbf{X}_{\text{ADC}}$ ) and point clouds ( $\mathbf{P}$ ) from two mmWave radar sensors are used to estimate UAV odometry ( $\hat{\xi}$ ) and construct synthetic array geometries from platform motion. These geometries enable high-resolution beamformed responses, which a lightweight U-Net converts into dense point clouds.

autopilot; enabling seamless coordination between sensing, estimation, and control. We show that UAV-SAR supports precise flight, accurate odometry, and high-resolution synthetic array-based sensing in indoor environments.

The key contributions of this work are as follows:

- We present a unified architecture, UAV-SAR, that jointly achieves accurate UAV odometry and high-resolution synthetic-array sensing using mmWave radar.
- We develop methods to leverage UAV motion and environmental features for dynamically constructing synthetic radar arrays, improving angular resolution beyond the limits of the physical sensor.
- We demonstrate the real-world feasibility of UAV-SAR through deployment on a custom UAV prototype, integrated with ROS 2 and the PX4 autopilot, enabling precise flight control, reliable state estimation, and high resolution radar sensing in indoor environments.

## II. RELATED WORKS

**Radar localization and odometry.** Radar-based localization (estimating position in a known environment) and odometry (tracking position relative to a starting point in an unknown environment) have been demonstrated across domains including automotive [34]–[37], ground robots [27], [38]–[40], and UAVs [26]–[32], [41]. Automotive systems typically employ *high-resolution* radars such as CTS350-X [42], which achieve angular resolutions of  $1^\circ$ – $4^\circ$  but are impractical for UAVs due to their weight (6 kg) and power draw (24 W). UAV-focused methods instead leverage lightweight, low-power mmWave radars and have demonstrated accurate (i.e., with velocity and position errors of  $>0.06$  m/s and  $>0.5$  m [28]) odometry [26], [28]–[32], [41]. However, these methods require radar frame rates exceeding 10 Hz to ensure stable flight, leaving a single frame’s duration too short for effective synthetic aperture beamforming. As a result, these prior works achieve odometry but NOT high-resolution sensing. In contrast, UAV-SAR is the first to simultaneously deliver accurate odometry AND high-resolution sensing on a UAV, using radar-based altitude and velocity measurements at 10 Hz to align frames and coherently integrate them for synthetic array imaging.

**Synthetic Aperture Radar (SAR).** SAR—where platform motion is exploited to synthesize a large effective aperture

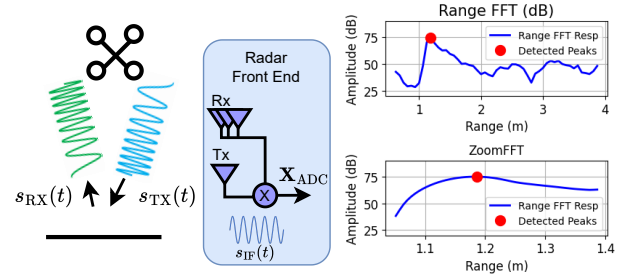


Fig. 2: UAV altitude estimation pipeline.

from a single onboard sensor—has enabled high-resolution radar imaging in applications ranging from concealed item detection to ground mapping [14]–[18]. UAV-based SAR has also been demonstrated for tasks such as landmine detection [19]–[25]. However, these systems rely on external positioning sources, such as GPS, or vision-based methods, which are unreliable in contested or low-visibility environments. In contrast, in this work, we demonstrate, to our knowledge, the first UAV SAR using *radar-only* odometry, enabling synthetic array imaging from 500 ms of integrated radar data without relying on GPS or cameras.

**Deep learning (DL) models for mmWave-based situational awareness.** DL has been applied to enhance low-resolution radar measurements into higher-resolution representations [7]–[13], [27]. For example, [27] employed graph neural networks (GNNs) to refine radar point clouds, while [10] used a U-Net architecture to enhance range–azimuth responses. Yet, most of these methods operate on low-resolution inputs and do not perform odometry. The few that do [12], [13] are limited to 2D odometry and are challenging to deploy on UAVs due to weight and compute constraints. Moreover, [13] required a mechanically rotating radar to enable high-resolution sensing on a UGV, increasing system weight and complexity. By contrast, UAV-SAR uniquely achieves *simultaneous 3D odometry and high-resolution sensing on a UAV using only lightweight mmWave radars*. Rather than relying on external sensors or rotating arrays, we directly exploit UAV motion to synthesize a virtual aperture, enabling efficient radar-only operation in 3D flight.

## III. UAV-SAR SYSTEM DESIGN

UAV-SAR uses downward- and outward-facing mmWave

radars to estimate a UAV's odometry and perform synthetic array processing (Fig. 1). The synthetic array responses are then fed into a U-Net-based ML architecture to obtain an accurate and dense point cloud in real-time, which can then be used by downstream sensing tasks (e.g., mapping).

### A. MmWave Radar Processing

For each frame, a frequency-modulated continuous-wave (FMCW) mmWave radar transmits a chirp signal,  $s_{TX}(t)$ , whose frequency increases linearly with time (Fig. 2). Signals reflected from objects in the environment are received as  $s_{RX}(t)$  at the radar's antennas. The transmitted and received signals are then mixed to produce the intermediate frequency (IF) signal,  $s_{IF}(t) = s_{TX}(t) \cdot s_{RX}^*(t)$ , which is sampled by the radar's analog-to-digital converter (ADC).

Each radar frame consists of multiple chirps, and its reflections from the environments are captured across multiple receive antennas. The sampled IF signals are stored in an *ADC data cube*, denoted by  $\mathbf{X}_{ADC} \in \mathbb{C}^{N_C \times N_S \times N_A}$ , where  $N_S$  is the number of ADC samples per chirp,  $N_C$  is the number of chirps per frame, and  $N_A$  is the number of receive antennas. From this data cube, a radar's on-device processor generates a radar point cloud,  $\mathbf{P} \in \mathbb{R}^{N_p \times 4}$ , where  $N_p$  is the number of detections, and each row in  $\mathbf{P}$  corresponds to one detection  $\mathbf{d}_i = [\mathbf{r}_i, d_{v,i}]$ , where the vector  $\mathbf{r}_i = [(d_{x,i}, d_{y,i}, d_{z,i})]$  denotes the Cartesian coordinates of the  $i$ -th detection relative to the UAV, and the scalar  $d_{v,i}$  is its radial velocity, i.e., the detection's velocity directly toward or away from the UAV.

### B. UAV Odometry Estimation

In GPS-denied environments, such as indoor or contested spaces, UAVs typically estimate odometry by combining measurements from a downward-facing optical flow sensor for planar velocity estimates,  $\hat{\mathbf{v}} = [\hat{v}_x, \hat{v}_y]$ , with altitude measurements from a laser range finder. These measurements are fused with data from the onboard inertial measurement unit (IMU) using an extended Kalman filter (EKF) to provide the flight controller with reliable odometry estimates. We define the resulting state as  $\hat{\xi}(t) = [\hat{\xi}_x(t), \hat{\xi}_y(t), \hat{\xi}_z(t), \hat{\xi}_\phi(t), \hat{\xi}_\theta(t), \hat{\xi}_\psi(t)]$ , corresponding to the estimated  $x$ ,  $y$ ,  $z$  positions and roll, pitch, yaw orientations of the UAV at time  $t$ .

However, in visually degraded conditions such as darkness, smoke, dust, or occlusions, both optical flow and laser sensors can fail. To overcome these limitations, we replace these sensors with a downward- and an outward-facing mmWave radar; this provides both altitude (from the downward-) and planar velocity measurements (from the outward-facing radar) used for robust odometry estimation.

**UAV altitude estimation.** As shown in Fig. 2, accurate altitude estimates from the downward-facing radar are obtained in two stages. First, a "coarse" estimate is computed from a single chirp's IF signal ( $s_{IF}(t)$ ) within the ADC data cube ( $\mathbf{X}_{ADC}$ ) by applying a Hamming window and a fast Fourier transform (FFT) to identify the highest-power peak in the rangeFFT, corresponding to the altitude of the UAV. Next,




Fig. 3: Overview of UAV velocity estimation (left), and delay and sum beamforming on the synthetic array (right).

a ZoomFFT refines this estimate by downshifting  $s_{IF}(t)$  to isolate the frequency spectrum within 20 cm of the coarse estimate. By restricting the FFT to this narrow bandwidth, the ZoomFFT effectively interpolates the signal, increasing frequency resolution for precise peak localization [43]. As shown in Sec. V, this approach achieves  $1.5\times$  more accurate altitude estimates than rangeFFT-only methods (e.g., [26]), yielding an *average absolute measurement error of 1.46 cm*.

**UAV velocity estimation.** Next, we estimate the UAV's velocity,  $\hat{\mathbf{v}} = [\hat{v}_x, \hat{v}_y, \hat{v}_z]$ , using the Cartesian positions,  $\mathbf{r}_i$ , and Doppler velocity measurements,  $d_{v,i}$ , from all detections in the front radar's point cloud ( $\mathbf{P}$ ) (as illustrated in Fig. 3). For a UAV flying at velocity  $\mathbf{v} = [v_x, v_y, v_z]$  in a static environment, the measured Doppler velocity of a target at position  $\mathbf{r}$  can be expressed as

$$-d_{v,i} = \frac{\mathbf{r}_i}{\|\mathbf{r}_i\|} \cdot \mathbf{v} = \mathbf{u}_i \cdot \mathbf{v}, \quad (1)$$

where  $\mathbf{u}_i = \frac{\mathbf{r}_i}{\|\mathbf{r}_i\|}$  is a unit vector pointing in the direction of the  $i$ th detection and the measured Doppler velocity is negated because the object appears to move in the direction opposite to the UAV.

Given  $N$  detections from the point cloud, stacking (1) across all targets yields:

$$\begin{bmatrix} -d_{v,1} \\ -d_{v,2} \\ \vdots \\ -d_{v,N} \end{bmatrix} = \begin{bmatrix} d_{x,1} & d_{y,1} & d_{z,1} \\ d_{x,2} & d_{y,2} & d_{z,2} \\ \vdots & \vdots & \vdots \\ d_{x,N} & d_{y,N} & d_{z,N} \end{bmatrix} \begin{bmatrix} v_x \\ v_y \\ v_z \end{bmatrix} \quad (2)$$

$$y = \mathbf{H} \mathbf{v}.$$

The UAV's velocity can then be estimated using least squares regression.

To mitigate the effect of false detections from clutter, multipath, and noise in radar point clouds, we extract robust velocity estimates using a RANSAC regressor, which iteratively fits the model to random subsets of the data and selects the fit with the largest consensus of inliers. Here, we implement this via Scikit-learn's `LinearRegression` and `RANSACRegressor` [44], parameterized with a 10-sample minimum per iteration and a 0.15 inlier residual threshold.

Finally, although our radar enables full 3D velocity estimation, optical flow sensors inherently provide only planar velocity measurements ( $\hat{v}_x$ , and  $\hat{v}_y$ ). Therefore, we specifically replace the optical flow inputs with the radar's  $\hat{v}_x$ , and  $\hat{v}_y$ , velocity estimates. Here, we note that the UAV estimates

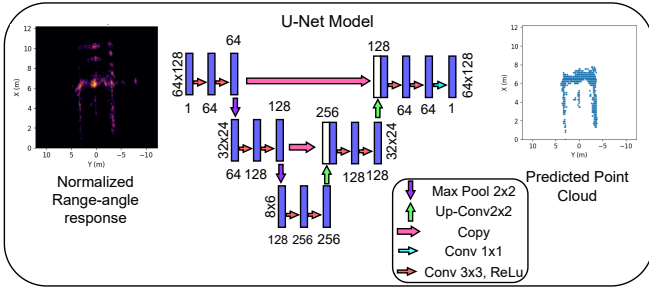


Fig. 4: UAV-SAR U-Net model.

its  $z$ -axis velocity by fusing the flight controller’s IMU data with altitude measurements from the downward-facing radar.

**EKF state estimation.** Finally, accurate UAV state estimates,  $\hat{\xi}$ , are obtained using the onboard flight controller’s EKF. The EKF fuses radar-derived altitude and velocity measurements, provided at a rate of 10 Hz, with data from the onboard IMU to estimate both the physical state,  $\xi$ , and the overall velocity,  $\hat{v}$ , of the UAV. In this work, we employ the PX4 Autopilot’s `ekf2` implementation of the EKF. To integrate with the existing flight controller interface, the radar altitude and velocity measurements are converted into equivalent optical flow sensor readings—effectively replacing the UAV’s onboard optical flow sensor [33].

### C. Synthetic Array Beamforming

While our formulation supports arbitrary 3D trajectories, we focus on optimizing azimuth (cross-range) resolution by leveraging horizontal UAV motion along the  $y$ -axis (orthogonal to the sensor’s look direction). We dynamically compute the 3D position of the  $i$ -th synthetic array element,  $\mathbf{a}_i$ , using UAV EKF velocity estimates,  $\hat{v}$ . Once the array geometry is established, delay-and-sum beamforming is applied to generate the synthetic response.

**Dynamic array geometry computation.** We approximate the UAV’s trajectory as linear within each radar frame, computing element positions as  $\mathbf{a}_i \approx \mathbf{v} \cdot i \cdot T_{\text{chirp}}$ , where  $T_{\text{chirp}}$  is the inter-chirp time. To maximize the effective aperture and resolution, we integrate these positions across five consecutive frames as shown in Fig. 1 and Fig. 3. This 500 ms duration is the empirically determined upper limit for maintaining sufficiently accurate array geometries. To guarantee reliable array formation, we adopt an opportunistic filtering strategy. Geometries are only accepted when lateral velocity ( $\hat{v}_y$ ) is between 0.30 and 0.70 m/s—the upper limit for sustaining controlled indoor flight—while non-lateral velocities ( $\hat{v}_x$ ,  $\hat{v}_z$ ) remain below 0.1 m/s. Additionally, the standard deviation of all velocity components across the integration window must not exceed 0.02 m/s.

**Delay-and-sum beamforming.** Once the synthetic array geometry is established, we apply delay-and-sum beamforming to compute the spatial response over a set of azimuth steering vectors. For a steering direction  $\theta$ , we define the unit steering vector as  $\mathbf{k}(\theta) = [\cos(\theta), \sin(\theta), 0]$ , and the corresponding

beamformed response is given by

$$B(\mathbf{k}, t) = \sum_{i=1}^M s_{\text{IF}}^i(t) \exp\left(j 2\pi \frac{2 \mathbf{a}_i \cdot \mathbf{k}}{\lambda}\right); \quad (3)$$

here,  $M$  is the number of synthetic aperture elements,  $\lambda$  is the transmit wavelength, and the factor of two accounts for the round-trip motion of both transmitter and receiver.

The beamformed signal is then windowed (Hamming) and processed with a RangeFFT to obtain the final range–angle response. In this work, angular responses are computed from  $-90^\circ$  to  $+90^\circ$  in  $2^\circ$  increments, producing 90 angle bins, and the RangeFFT spans 0 m to 12.29 m with a resolution of 0.1365 m, producing 90 range bins. Thus, the final range–angle response is represented as a  $90 \times 90$  matrix.

Although platform motion introduces non-uniform sampling and deviations from an ideal uniform linear array (ULA), the achievable angular resolution can be approximated by the half power beamwidth (HPBW) of a ULA of equivalent length:

$$\text{HPBW} \approx 0.886 \cdot \frac{\lambda}{L} \text{ (radians)}, \quad (4)$$

where  $L$  is the effective aperture size [45]. For context, the TI-IWR1843’s physical array (3 Transmitters–Tx, 4 Receivers–Rx) provides a minimum azimuth HPBW of  $14.5^\circ$ . In contrast, a synthetic aperture generated by moving at 0.3 m/s for 0.5 s ( $L \approx 0.15$  m) narrows the HPBW to approximately  $1.3^\circ$ . Matching this performance physically would require roughly 78 receivers at standard  $\lambda/2$  spacing. Thus, synthetic aperture processing improves angular resolution by *more than an order of magnitude* compared to the sensor’s native capabilities.

### D. U-Net-based Point Cloud Enhancement

To obtain accurate and high resolution point clouds from the range-angle responses, UAV-SAR utilizes the lightweight U-Net architecture featured in Fig. 4. Here, the ground truth training labels are derived from regions in lidar point cloud observed in the radar’s range-angle response.

**Model input.** The input to the U-Net model used by UAV-SAR is derived from the range-angle response, initially stored in dB. To start, the response is interpolated onto a Cartesian grid with 0.1 m resolution, spanning from (0 m, 12.2 m) in the  $x$  (forward) direction and from (-12.2 m, 12.2 m) in the  $y$  (left/right) dimension, resulting in a  $123 \times 246$  tensor. To focus on higher-SNR reflections, while retaining some weaker features, we apply a thresholding step to set an upper and lower bound as 50 dB and 20 dB above an approximate noise level, respectively, estimated from the 30th percentile of the response. The resulting tensor is then normalized to the  $[0, 1]$  range. Finally, for better deployment on resource-constrained platforms and to be compatible with the U-Net architecture, the input is resized to a  $64 \times 128$  tensor as shown in Fig. 4. Compared to prior lightweight models such as in [11], which used a  $40 \times 64 \times 48$  tensor, this represents a roughly  $15\times$  reduction in input features.

TABLE I: UAV-SAR radar configurations.

Specification	TI-IWR6843-ODS (down)	TI-IWR1843 (front)
$N_S$	63	90
$N_{\text{Chirps}}$	100	100
$f_{\text{samp}}$	2.1 MHz	2.05 MHz
$d_{\text{res}}$	6.3 cm	12.1 cm
$d_{\text{max}}$	3.9 m	15.4 m
$v_{\text{res}}$	0.034 m/s	0.028 m/s
$v_{\text{max}}$	1.7 m/s	1.44 m/s

**Model output.** To obtain ground truth data for model training, we use a co-located lidar sensor that captures high-resolution 3D point clouds. We begin by extracting a slice of the lidar point cloud, retaining only points within the radar’s  $x$ - $y$  field of view and with  $z$  coordinates in the range (-0.25 m, 0.5 m), corresponding to the elevation beamwidth of the radar antenna. The segmentation mask is defined as a  $64 \times 128$  Cartesian grid with  $x$  values from 0 m to 12.2 m and  $y$  values from -12.3 m to 12.2 m, using a resolution approximately 19.3 cm.

Next, we construct a ground truth mask to align lidar measurements with the radar input. Cells in the normalized radar input that have values above 0.05 and lie within 1 m of a point in the filtered lidar slice are marked as valid. This mask provides the target labels for training the model to detect reflections corresponding to real objects. At runtime, the model’s predicted segmentation grid is converted back into Cartesian coordinates using the same  $x$ - $y$  grid, producing a high-resolution 2D point cloud.

**Model architecture and loss function.** We employ a 3-stage U-Net architecture [46] to generate high-resolution point clouds from the input range-angle response. Given the input’s already high resolution and to optimize performance on compute-constrained platforms, this 3-stage model is used instead of the 4- and 5-stage U-Nets in [11] and [10], reducing the number of parameters by  $4.1\times$  and  $9\times$ , respectively, and enabling low-latency inference. Training uses a weighted combination of Binary Cross Entropy (BCE, weight 0.9) and Dice loss (weight 0.1) [10], [11], where the BCE term encourages pixel-wise agreement with the ground truth mask, and the Dice term sharpens predicted features.

#### IV. EVALUATION

We implement and evaluate UAV-SAR on a real-world UAV prototype. The UAV-SAR architecture is deployed as a set of real-time ROS2 nodes and integrated with a commercially available flight controller running the PX4 autopilot. This platform enables end-to-end evaluation of both the UAV-SAR odometry estimation and high-resolution sensing pipelines, which are quantitatively benchmarked against ground-truth measurements obtained from commercial optical flow and lidar sensors.

##### A. Experimental Setup

**Platform.** The real-world UAV prototype used to evaluate the UAV-SAR architecture is shown in Fig. 5. Selected

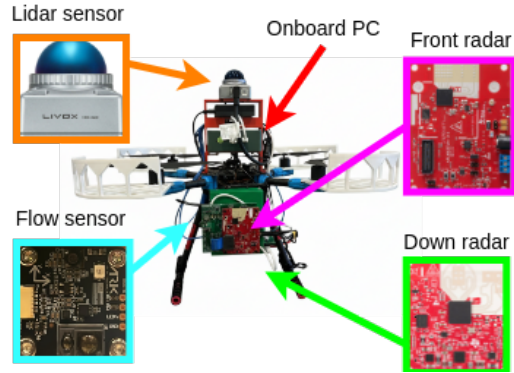


Fig. 5: Experimental platform

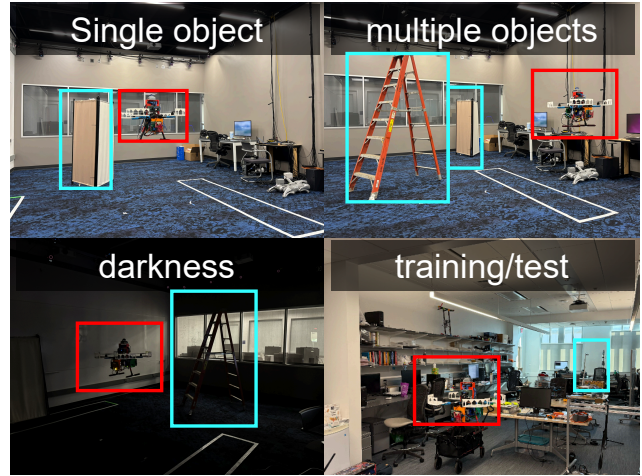


Fig. 6: UAV-SAR is evaluated in a complex campus laboratory (bottom right) and indoor flight arena with obstacles placed in various configurations.

specifically for their low size, weight, and power (SWaP) profile, the sensing suite comprises a downward-facing TI-IWR6843-ODS radar for altitude estimation and a forward-facing TI-IWR1843 for velocity and synthetic aperture sensing. Notably, the forward-facing radar’s configuration was chosen to achieve the best balance between fine velocity resolution and optimal synthetic array sensing performance. Table I summarizes the key parameters for each radar.

Onboard processing is handled by a GMKtec Mini PC with an Intel N100 CPU (i.e., without a GPU), running the UAV-SAR architecture as a set of real-time ROS2 nodes at 10 Hz. All sensors are integrated on a modified Holybro X500 V2 UAV equipped with a Pixhawk 6X flight controller running the PX4 autopilot firmware. Radar-derived altitude and velocity estimates are streamed to the flight controller over Ethernet via the PX4 XRCE-DDS protocol. For ground-truth benchmarking *only*, the platform is additionally equipped with a Livox Mid-360 lidar and an Ark Flow optical-flow and distance sensor module (accuracy of 15 mm).

**Evaluation environments.** Experiments were conducted at two indoor locations due to the limited availability of suitable UAV testing facilities. Training datasets were collected in a

cluttered campus laboratory environment containing tables, chairs, and other obstacles. After training, odometry and high-resolution sensing experiments were carried out in a dedicated indoor UAV flight arena. In each trial, targets were placed at random locations within the environment, and additional experiments were performed under near-darkness conditions to demonstrate UAV operation in scenarios where vision-based sensing is ineffective.

### B. Datasets

To evaluate the UAV-SAR’s performance, we collected three unique datasets.

**Model training and testing datasets.** We collected 213 time-synchronized synthetic array scans and lidar point clouds in a campus laboratory environment across three UAV flights. UAV odometry during these flights relied solely on radar-based altitude and velocity estimates, without optical flow or laser distance sensors. To increase scan diversity for model training, the UAV was manually flown at different altitudes and velocities in lateral left–right trajectories.

**Odometry evaluation.** To evaluate radar-based velocity and altitude estimation, we used the UAV’s optical flow sensor as ground truth. A total of 600 time-synchronized radar frames were recorded, with ground truth estimates refined using the optical flow sensor, distance sensor, and the flight controller’s EKF. For these experiments, the UAV followed varied trajectories including rotations, angled translations, and altitude changes.

**Unseen validation dataset.** Finally, to also evaluate the performance of our architecture in *unseen environments*, we captured a set of 730 synthetic array scans in the UAV flight arena across six unique flights. For each flight, the scene was randomly changed to have different lighting conditions and a combination of objects randomly placed within the radar’s field of view.

### C. Odometry evaluation

Radar altitude accuracy is evaluated using the absolute error,  $|\hat{z} - z_{gt}|$ , after correcting for the bias introduced by the optical flow sensor being mounted slightly above the downward-facing radar. Velocity accuracy is assessed as the mean absolute error relative to the fused optical flow-based velocity estimates produced by the flight controller’s EKF.

### D. Generated point cloud quality

The quality of the UAV-SAR model’s predicted point cloud with respect to the derived ground truth mask is evaluated using the common Chamfer and Modified Hausdorff metrics [47], [48]. Here, we define the Chamfer distance (CD) as

$$\text{CD}(S_1, S_2) = \frac{1}{2|S_1|} \sum_{x \in S_1} \min_{y \in S_2} d(x, y) + \frac{1}{2|S_2|} \sum_{y \in S_2} \min_{x \in S_1} d(x, y), \quad (5)$$

and Modified Hausdorff distance (MHD) as

$$\text{MHD}(S_1, S_2) = \max \left\{ \text{med} \min_{y \in S_2} d(x, y), \text{med} \min_{x \in S_1} d(x, y) \right\},$$

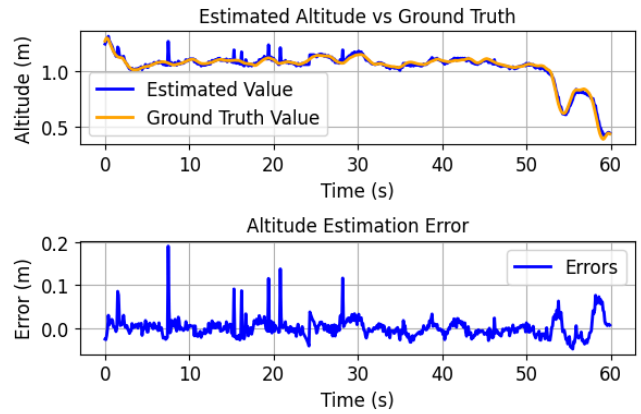


Fig. 7: UAV-SAR’s altitude estimates compared to ground truth from an ARK Flow’s distance sensor.

TABLE II: Mean and standard deviation of altitude estimation error with and without ZoomFFT.

Method	Mean Absolute Error	Std. Dev.
Without ZoomFFT	2.29 cm	2.57 cm
With ZoomFFT	1.46 cm	1.68 cm

where  $d(x, y)$  denotes the Euclidean distance, i.e.,  $\|x - y\|_2^2$ , and  $S_1, S_2$  represent the predicted radar point cloud and the derived ground-truth point cloud based on the lidar data.

## V. RESULTS

We now present the results of the UAV-SAR evaluations, starting from the altitude and velocity estimation, and followed by an analysis of the generated point cloud accuracy. Video of the experiments is added as part of the submission.

### A. Altitude estimation

Fig. 7 compares the estimated altitude against the ground truth obtained from the Ark Flow’s laser distance sensor, while Table II summarizes the key statistics across the dataset. The results show that UAV-SAR achieves highly accurate altitude estimation using the downward-facing radar sensor. Moreover, applying the ZoomFFT further reduces the mean absolute error by a factor of  $1.56\times$ . These findings highlight the practicality of radar-based altitude estimation, enabling robust operation in environments where optical sensors may be unreliable.

### B. Velocity estimation

Fig. 8 compares the radar-based velocity estimates against the ground truth from the Ark Flow optical flow sensor, while Fig. 9 shows the distribution of velocity errors. The results demonstrate that the estimated UAV velocities closely match the optical flow measurements across a variety of translational and rotational trajectories. With a mean absolute error of only 0.063 m/s, and the vast majority of estimates within 0.2 m/s of the ground truth, UAV-SAR achieves highly accurate velocity estimation. This level of accuracy is essential not only for reliable synthetic array generation but also

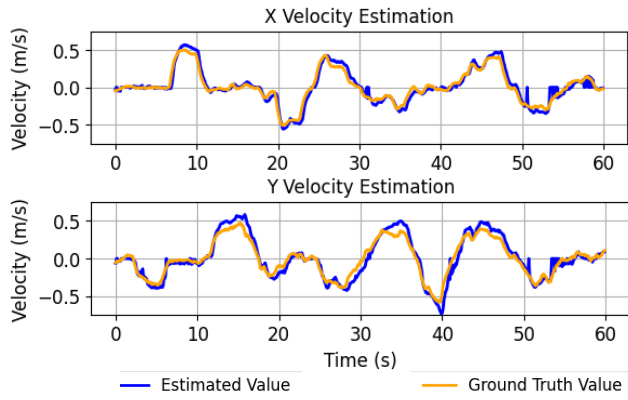


Fig. 8: UAV-SAR’s radar estimated velocity versus ground truth optical flow-based measurements.




Fig. 9: Distribution of velocity estimation errors as compared to optical flow sensor.

TABLE III: Summary statistics for the predicted point cloud errors as compared to the ground-truth point cloud.

Metric	Modified Hausdorff	Chamfer
Mean	0.19 m	0.36 m
Median	0.09 m	0.26 m
90th percentile	0.55 m	0.80 m

for ensuring robust UAV navigation in environments where optical flow sensing may be unreliable.

### C. High resolution sensing

Finally, we present the results of the high-resolution sensing experiments conducted in an unseen validation environment. As shown in Fig. 10, synthetic array processing produces a range-angle response with substantially higher resolution compared to conventional range-angle imaging. This enhanced response allows the lightweight U-Net model to generate point clouds that closely match the original lidar measurements. Notably, the radar response was even able to penetrate curtains and capture the layout of a back room that was completely occluded in the lidar data, underscoring the unique sensing advantages of radar.

Quantitative results, summarized in Fig. 11 and Table III, further confirm that the predicted point clouds align well with the lidar-derived ground truth. Taken together, these findings demonstrate that our fully radar-based approach enables high-resolution environmental sensing, powered by accurate radar odometry, and offers robust perception capabilities in conditions where lidar and camera may be limited.

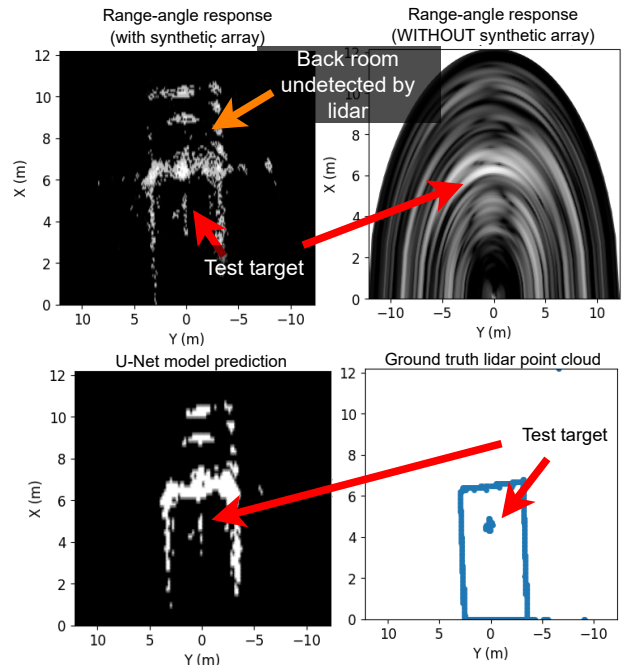


Fig. 10: Overview of high resolution sensing results using the UAV-SAR architecture in the unseen validation environment




Fig. 11: Distribution of predicted point cloud errors as compared to the ground truth segmentation mask.

## VI. CONCLUSION

In this work, we introduced UAV-SAR, a radar-based architecture that jointly enables accurate UAV odometry and high-resolution synthetic-array sensing. Using a custom UAV prototype, we demonstrated that UAV-SAR achieves precise altitude and velocity estimation, high-resolution synthetic-array beamforming, and dense point cloud generation with a lightweight U-Net model trained on significantly fewer samples than prior approaches. By unifying odometry and high-resolution sensing, UAV-SAR overcomes key limitations of optical and GPS-based methods, enabling UAV operation in visually degraded and GPS-denied environments. Future work will extend this framework to more complex outdoor settings and platform motions, while exploring advanced synthetic aperture radar techniques such as adaptive beamforming and autofocus algorithms.

## REFERENCES

- [1] B. Staff. (2025, Apr) Zipline has beaten amazon and ups in drone delivery. now it’s gunning for more. Accessed: 2025-09-09. [Online]. Available: <https://www.barrons.com>

- [2] M. Lam, L. Dodds, A. Eid, J. Hester, and F. Adib, "6d self-localization of drones using a single millimeter-wave backscatter anchor," in *IEEE INFOCOM 2025-IEEE Conference on Computer Communications*. IEEE, 2025, pp. 1–10.
- [3] D. B. Staff. (2025, Feb) U.s. air force eyes atomic clock tech for drone swarms. Accessed: 2025-09-09. [Online]. Available: <https://defence-blog.com>
- [4] J. Epstein. (2025, May) Nato allies are racing russia to develop arctic drones that can survive the cold and operate where gps doesn't work well. Accessed: 2025-09-09. [Online]. Available: <https://www.businessinsider.com>
- [5] S. Rao, "Introduction to mmwave Sensing: FMCW Radars," 2017.
- [6] P. Hügler, T. Grebner, C. Knill, and C. Waldschmidt, "Uav-borne 2-d and 3-d radar-based grid mapping," *IEEE Geosci. and Remote Sensing Letters*, vol. 19, pp. 1–5, 2020.
- [7] C. X. Lu, S. Rosa, P. Zhao, B. Wang, C. Chen, J. Stankovic, N. Trigoni, and A. Markham, "See Through Smoke: Robust Indoor Mapping with Low-cost mmWave Radar," May 2020, arXiv:1911.00398.
- [8] P. Cai and S. Sur, "MilliPCD: Beyond Traditional Vision Indoor Point Cloud Generation via Handheld Millimeter-Wave Devices," *Proc. of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, vol. 6, no. 4, pp. 1–24, Dec. 2022.
- [9] R. Geng, Y. Li, D. Zhang, J. Wu, Y. Gao, Y. Hu, and Y. Chen, "Dream-pcd: Deep reconstruction and enhancement of mmwave radar pointclouds," *IEEE Transactions on Image Processing*, 2024.
- [10] A. Prabhakara, T. Jin, A. Das, G. Bhatt, L. Kumari, E. Soltanaghai, J. Bilmes, S. Kumar, and A. Rowe, "High Resolution Point Clouds from mmWave Radar," in *2023 IEEE Int. Conf. on Robot. and Automat. (ICRA'23)*, May 2023, pp. 4135–4142.
- [11] D. Hunt, S. Luo, A. Khazraei, X. Zhang, S. Hallyburton, T. Chen, and M. Pajic, "RadCloud: Real-Time High-Resolution Point Cloud Generation Using Low-Cost Radars for Aerial and Ground Vehicles," in *IEEE ICRA*, 2024, pp. 12 269–12 275.
- [12] E. Sie, X. Wu, H. Guo, and D. Vasisht, "Radarize: Enhancing radar slam with generalizable doppler-based odometry," in *Proc. 22nd Int. Conf. on Mob. Systems, Applications and Services*, 2024, pp. 331–344.
- [13] H. Lai, G. Luo, Y. Liu, and M. Zhao, "Enabling visual recognition at radio frequency," in *ACM Int. Con. on Mob. Comp. (MobiCom)*, 2024.
- [14] M. Hoffmann, T. Noegel, C. Schüßler, M. Vossiek, M. Schütz, and P. Gulden, "Filter-based segmentation of automotive sar images," in *2022 IEEE Radar Conference (RadarConf22)*. IEEE, 2022, pp. 1–6.
- [15] M. E. Yanik and M. Torlak, "Near-field 2-d sar imaging by millimeter-wave radar for concealed item detection," in *2019 IEEE radio and Wireless Symposium (RWS)*. IEEE, 2019, pp. 1–4.
- [16] M. E. Yanik, D. Wang, and M. Torlak, "Development and demonstration of mimo-sar mmwave imaging testbeds," *IEEE Access*, vol. 8, pp. 126 019–126 038, 2020.
- [17] X. Gao, S. Roy, and G. Xing, "MIMO-SAR: A Hierarchical High-Resolution Imaging Algorithm for mmWave FMCW Radar in Autonomous Driving," *IEEE Trans. on Veh. Technol.*, vol. 70, no. 8, pp. 7322–7334, Aug. 2021.
- [18] S. Y. Mahbub, M. Dawood, and E. Shareef, "Inexpensive sar testbed for 3d mmwave imaging applications," *IET Radar, Sonar & Navigation*, vol. 19, no. 1, p. e12686, 2025.
- [19] E. Schreiber, A. Heinzl, M. Peichl, M. Engel, and W. Wiesbeck, "Advanced buried object detection by multichannel, uav/drone carried synthetic aperture radar," in *2019 13th European Conference on Antennas and Propagation (EuCAP)*. IEEE, 2019, pp. 1–5.
- [20] A. Meta, "Signal processing of fmcw synthetic aperture radar data," 2006.
- [21] J. J. de Wit, A. Meta, and P. Hoogetboom, "Modified range-doppler processing for fm-cw synthetic aperture radar," *IEEE Geoscience and Remote Sensing Letters*, vol. 3, no. 1, pp. 83–87, 2006.
- [22] A. Meta and P. Hoogetboom, "Signal processing algorithms for fmcw moving target indicator synthetic aperture radar," in *Proceedings. 2005 IEEE International Geoscience and Remote Sensing Symposium, 2005. IGARSS'05.*, vol. 1. IEEE, 2005, pp. 4–pp.
- [23] J. Yan, Z. Peng, H. Hong, H. Chu, X. Zhu, and C. Li, "Vital-sar-imaging with a drone-based hybrid radar system," *IEEE Trans. Microw. Theory Techn.*, vol. 66, no. 12, pp. 5852–5862, 2018.
- [24] A. Meta, P. Hoogetboom, and L. P. Lighthart, "Correction of the effects induced by the continuous motion in airborne fmcw sar," in *2006 IEEE Conference on Radar*. IEEE, 2006, pp. 8–pp.
- [25] M. Caris and S. Stanko, "Millimeter wave sar for uav-based remote sensing," in *Workshop on UAV-based Remote Sensing Methods for Monitoring Vegetation 2013*, 2014.
- [26] E. Sie, Z. Liu, and D. Vasisht, "Batmobility: Towards flying without seeing for autonomous drones," in *Proc. 29th Annu. Int. Conf. Mobile Comput. Netw.*, 2023, pp. 1–16.
- [27] D. Hunt, S. Luo, S. Hallyburton, S. Nillongo, Y. Li, T. Chen, and M. Pajic, "Ragnarok: A light-weight graph neural network for enhancing radar point clouds on unmanned ground vehicles," *arXiv preprint arXiv:2507.00937*, 2025.
- [28] C. Doer and G. Trommer, "x-rio: Radar inertial odometry with multiple radar sensors and yaw aiding," *Gyroscopy and Navigation*, vol. 12, no. 4, pp. 329–339, 2021.
- [29] J. Michalczyk, R. Jung, and S. Weiss, "Tightly-coupled ekf-based radar-inertial odometry," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2022, pp. 12 336–12 343.
- [30] E. B. Quist, P. C. Niedfeldt, and R. W. Beard, "Radar odometry with recursive-ransac," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 52, no. 4, pp. 1618–1630, 2016.
- [31] M. Mostafa, S. Zahran, A. Moussa, N. El-Sheimy, and A. Sesay, "Radar and visual odometry integrated system aided navigation for uavs in gnss denied environment," *Sensors*, vol. 18, no. 9, p. 2776, 2018.
- [32] P. Stockel, P. Wallrath, R. Herschel, and N. Pohl, "Motion compensation in six degrees of freedom for a mimo radar mounted on a hovering uav," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 59, no. 5, pp. 5791–5801, 2023.
- [33] P. D. Team and Community, "Px4 autopilot," <https://github.com/PX4/PX4-Autopilot>, accessed: 2025-09-14.
- [34] R. Aldera, D. D. Martini, M. Gadd, and P. Newman, "What Could Go Wrong? Introspective Radar Odometry in Challenging Environments," in *IEEE Intell. Transportation Sys. Conf. (ITSC)*, 2019, pp. 2835–2842.
- [35] R. Aldera, D. D. Martini, M. Gadd, and P. Newman, "Fast radar motion estimation with a learnt focus of attention using weak supervision," in *2019 ICRA*, 2019, pp. 1190–1196.
- [36] D. Lius, K. Burnett, D. J. Yoon, R. Poulton, J. Marshall, and T. D. Barfoot, "Are doppler velocity measurements useful for spinning radar odometry?" *IEEE Robot. and Automat. Letters*, vol. 10, no. 1, pp. 224–231, 2025.
- [37] F. Rennie, D. Williams, P. Newman, and D. De Martini, "Doppler-aware odometry from fmcw scanning radar," in *26th IEEE ITSC*, 2023, pp. 5126–5132.
- [38] H. Chen, Y. Liu, and Y. Cheng, "Drio: Robust radar-inertial odometry in dynamic environments," *IEEE Robot. and Automat. Letters*, 2023.
- [39] C. X. Lu, M. R. U. Saputra, P. Zhao, Y. Almalioglu, P. P. De Gusmao, C. Chen, K. Sun, N. Trigoni, and A. Markham, "milliego: single-chip mmwave radar aided egomotion estimation via deep sensor fusion," in *18th Conf. on Embedded Networked Sensor Syst.*, 2020, pp. 109–122.
- [40] Y. Almalioglu, M. Turan, C. X. Lu, N. Trigoni, and A. Markham, "Milli-rio: Ego-motion estimation with low-cost millimetre-wave radar," *IEEE Sensors J.*, vol. 21, no. 3, pp. 3314–3323, 2020.
- [41] E. B. Quist and R. W. Beard, "Radar odometry on fixed-wing small unmanned aircraft," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 52, no. 1, pp. 396–410, 2016.
- [42] "Clearway technical specifications, [www.navtechradar.com](http://www.navtechradar.com)," 2023.
- [43] B. Al-Qudsi, N. Joram, A. Strobel, and F. Ellinger, "Zoom fit for precise spectrum calculation in fmcw radar using fpga," in *Proceedings of the 2013 9th Conference on Ph.D. Research in Microelectronics and Electronics*. IEEE, 2013, pp. 337–340.
- [44] F. e. a. Pedregosa, "Scikit-learn: Machine learning in python," *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.
- [45] C. A. Balanis, *Antenna Theory: Analysis and Design*, 3rd ed. Hoboken, NJ: John Wiley & Sons, 2005.
- [46] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional Networks for Biomedical Image Segmentation," May 2015, arXiv:1505.04597 [cs]. [Online]. Available: <http://arxiv.org>
- [47] P. Contributors, "Chamfer," Sept. 2023. [Online]. Available: <https://pdal.io>
- [48] M.-P. Dubuisson and A. Jain, "A modified Hausdorff distance for object matching," in *Proc. of 12th Int. Conf. on Pattern Recognition*, vol. 1, 1994, pp. 566–568.