

mmFlux: Crowd Flow Analytics with Commodity mmWave MIMO Radar

Anurag Pallaprolu[†], Winston Hurst, and Yasamin Mostofi

{apallaprolu, winstonhurst, ymostofi}@ucsb.edu

University of California Santa Barbara

Santa Barbara, USA

Abstract

In this paper, we present mmFlux: a novel framework for extracting underlying crowd motion patterns and inferring crowd semantics using mmWave radar. First, our proposed signal processing pipeline combines optical flow estimation concepts from vision with novel statistical and morphological noise filtering. This approach generates high-fidelity mmWave flow fields—compact 2D vector representations of crowd motion. We then introduce a novel approach that transforms these fields into directed geometric graphs. In these graphs, edges capture dominant flow currents, vertices mark crowd splitting or merging, and flow distribution is quantified across edges. Finally, we show that analyzing the local Jacobian and computing the corresponding curl and divergence enables extraction of key crowd semantics for both structured and diffused crowds. We conduct 21 experiments on crowds of up to 20 people across 3 areas, using commodity mmWave radar. Our framework achieves high-fidelity graph reconstruction of the underlying flow structure, even for complex crowd patterns, demonstrating strong spatial alignment and precise quantitative characterization of flow split ratios. Finally, our curl and divergence analysis accurately infers key crowd semantics, e.g., abrupt turns, boundaries where flow directions shift, dispersions, and gatherings. Overall, these findings validate mmFlux, underscoring its potential for various crowd analytics applications.

Keywords

mmWave Radar, Crowd Analytics, Crowd Flow Graph Reconstruction, Crowd Semantics, RF Sensing

1 Introduction

Understanding crowd motion dynamics offers valuable insights into how a space is used. In commercial settings, analyzing customer movement can reveal which items attract attention, informing optimal product placement and store

layout [1, 2]. In security, detecting anomalies in crowd behavior can help identify potential theft or safety threats [3]. Urban planners study crowd flows to design public spaces that minimize bottlenecks and ensure safe evacuations in emergencies [4], and event organizers use similar analyses to improve venue layouts and crowd control at concerts, stadiums, and festivals [5]. More broadly, smart cities leverage crowd flow data to optimize traffic patterns, enhance public transportation, and improve overall urban mobility [6].

Historically, vision-based systems have been the primary method for extracting crowd analytics, due to their high resolution. At the same time, growing privacy concerns [7–11] have fueled interest in RF-sensing based solutions, which are naturally privacy-preserving, with several earlier works using WiFi signals [12–16]. However, the low sensing resolution of WiFi limits its applicability for crowd analytics, despite recent advances in WiFi-based sensing for other tasks [17–21]. On the other hand, in recent years, there has been a shift toward mmWave frequencies, driven by their central role in 5G/6G networks [22] as well as the increasing availability of low-cost commodity mmWave transceivers. As such, recent studies have demonstrated the potential of mmWave systems for RF sensing in general, and crowd analytics in particular [23–25].

In this paper, we push the boundaries of mmWave-based crowd analytics by presenting mmFlux: a novel method to extract high-level motion patterns for large crowds **using a single commodity mmWave radar board**, validated through 21 experiments involving crowds of up to 20 people. To capture a broad set of crowd behavior, we consider two categories of crowd dynamics in this paper. Specifically, we refer to crowds which have dominant flow patterns that can be well-represented by a graph as **structured crowds**. We further consider crowds that are not structured as such, but still exhibit emerging flow behaviors, which we refer to as **diffuse crowds**. See Fig. 1 for a few samples and Fig. 3 for a comprehensive list of considered crowd dynamics. We assume that the underlying flow structure/behavior is time-invariant, whenever there is a flow. We next summarize our key contributions.

a) We introduce **flow fields** as a compact representation of crowd motion dynamics in mmWave crowd analytics.

[†]Corresponding author: apallaprolu@ucsb.edu .




Figure 1

This approach obviates the need for end-to-end tracking of individuals in the crowd, reducing the effects of resolution limitations and occlusion. Building on traditional optical flow methods in the area of vision, our novel signal processing pipeline ensures robust flow field estimation by developing new noise-filtering and interpolation techniques tailored to the sparse and noisy nature of mmWave data.

b) Using the estimated flow field, we then show how to extract the high-level flow topology of structured crowds (denoted by $G(V, E)$) by constructing a directed geometric graph, $\hat{G}(\hat{V}, \hat{E})$, whose nodes represent crowd split or merge points, and edges denote dominant flow currents. Additionally, we propose a method to estimate flow split ratios at each node, quantifying the relative distribution of flows across edges. The flow topology, combined with split ratios, enables a detailed characterization of structured crowd dynamics.

c) For both structured and diffuse crowds, we analyze the flow field itself to extract key semantic insights. Specifically, we present a novel method that computes the curl and divergence of the flow field, using vector field analysis, enabling the identification of flow split and merge regions, sharp turns, and boundaries where dominant flow directions shift abruptly.

d) We validate mmFlux through extensive testing using a commodity TI AWR2243BOOST mmWave radar board [26]. Specifically, we present the results of 21 real-world experiments with crowds of up to (and including) 20 people in three distinct areas that experience significant environmental noise and multipath effects. Our framework achieves high-fidelity graph reconstruction of the underlying flow structure, demonstrating strong spatial alignment even for complex crowd patterns, and further accurately estimates the flow split ratios with a Mean Absolute Error (MAE) of 0.1. We finally demonstrate how to use the curl and divergence of the flow field to partition the space into regions with different behaviors, delineate boundaries, and identify crowd semantics such as sharp turns and splits.

We next start with a comprehensive literature review. Crowd analytics has gained significant attention in recent

years, driven by the need for efficient urban planning [27–29], safety management [30], and behavioral understanding in dynamic environments [31].

Vision-based methods have long dominated crowd analytics, with early work using optical flow and physics-inspired models, such as fluid mechanics and social forces, to analyze crowd dynamics [32–40], and later work applying deep learning methods [41–47]. Despite these advances, vision-based approaches remain vulnerable to adverse environmental conditions such as poor lighting and rain/fog, and they have provoked increasing societal resistance due to mounting privacy concerns from citizens [10], regulatory bodies [48], and advocacy groups [11].

As a privacy-preserving alternative to vision, researchers have explored RF-based modalities for crowd analytics. In particular, WiFi and BLE signals have been widely used for crowd monitoring [12], with prior work mainly focusing on occupancy estimation [13–16, 49–56]. However, the poor sensing resolution of these frequencies greatly limits their applicability and further necessitates additional assumptions, e.g., line-of-sight crossing or use of costly and extensive transceiver networks.

On the other hand, mmWave radars have recently gained widespread attention for their privacy-preserving and weather-resilient sensing capabilities, and their high-resolution spatial sensing has been leveraged for fine-grained human motion analysis [57]. Recent work has also started investigating crowd analytics with mmWave radar, but with most work focusing on occupancy estimation/crowd counting [23–25, 58–62], as opposed to extracting underlying crowd flow patterns and emerging behaviors. In other words, existing mmWave-based crowd analytics studies do not address the specific subject of this paper and their solutions are not adaptable to solve our problem of interest. For instance, several crowd counting works track individuals for counting [63–66], but they suffer from occlusion effects and resolution limits, and can thus only handle small crowd sizes [67]. In [25], we proposed a foundation for crowd counting inspired by concepts from stochastic geometry. However, the approach relied on prior knowledge of the crowd spatial usage. This underscores the need for a new framework for aggregate motion representation and analysis, in order to robustly extract underlying crowd flow patterns, which is the main motivation for this paper.

Specifically, we start by introducing a new flow field model for mmWave crowd sensing, which shall provide a compact and robust representation of crowd dynamics, enabling subsequent crowd flow analysis. While flow field models of crowd behavior have received significant attention in vision, such approaches remain largely unexplored for mmWave radars. A notable exception is [68], which focuses exclusively on anomaly detection and employs a flow field methodology



Figure 2




Figure 3

tailored to that task, thus lacking generalizable analysis and broader applicability. Beyond the domain of crowd analytics, mmWave flow field models have been developed for automotive scene flow [69] and activity recognition for a single individual [70], but their task-specific training limits their applicability to crowd flow analysis.

To the best of our knowledge, no existing RF-based approach has addressed the problem of characterizing aggregate spatial crowd motion patterns. Moreover, existing methods that focus on estimating the crowd count are not applicable for characterizing the underlying crowd flow patterns. In fact, underlying flow patterns are often treated as a prior in crowd counting. As such, the methodology proposed in this paper can also contribute to advancing the state of the art in crowd counting, opening avenues for future research.

2 Results

We next present the results of a comprehensive evaluation of our proposed system for crowd flow analytics through a series of real-world experiments. We first describe our experimental setup and testing areas before presenting empirical results from 21 experiments, with crowds of up to (and including) 20 individuals that validate our system’s capability to reconstruct the crowd flow topology and further extract relevant spatial insights. Specifically, our approach in Sec. 4.6 recovers the flow topology in 17 of 18 (94%) experiments focused on graph reconstruction, with all 18 experiments

showing strong visual alignment with the ground truth. Even in the challenging case, our system successfully infers a majority of the flow structure (see Sec. 3 for discussion). To quantify the accuracy of our reconstruction, we compute the one-sided Chamfer distance [71] between the 2D embeddings of our estimated graph $\hat{G}$ and the ground-truth topology G , achieving an average distance of 0.45 m, well within the width of a typical pedestrian traffic lane [72]. Further, we achieve an edge orientation MAE of 8.8° between the ground-truth and estimated edges, and a split ratio MAE of 0.1. Notably, these results are obtained using only a single mmWave radar board, and in environments with substantial foliage noise. We conclude by showing how the flow field curl and divergence can be used to localize regions that may correspond to boundaries between dominant flow directions or anomalous events within the sensing field of view. **This research has been reviewed and approved by our Institutional Review Board (IRB) committee.**

2.1 Experimental Setup

We assess our proposed system using a TI AWR2243BOOST off-the-shelf mmWave MIMO radar board [26], as shown in Fig. 2 (d). Operating at a base frequency of $f_0 = 76$ GHz, the radar transmits FMCW pulses with a $B = 5$ GHz bandwidth. To leverage its MIMO capabilities, we configure $N_{TX} = 3$ transmitters and $N_{RX} = 4$ receivers, constructing a 2×8 virtual array. However, given the limited elevation resolution (Sec. 4.1), we collapse the vertical dimension, treating the array as a 1×8 virtual uniform linear array. To improve azimuthal resolution, we apply zero-padding before performing angle-of-arrival estimation. The radar’s Channel State Information is recorded via a DCA1000EVM FPGA [73], which streams received chirp data to a laptop over Ethernet.

We conduct experiments in three outdoor environments, illustrated in Fig. 2 (Bottom). Area 1 is a $12 \text{ m} \times 12 \text{ m}$ open space with seating on one side and a trailer on the other. Area 2 is a $21 \text{ m} \times 14 \text{ m}$ semi-enclosed passage between two buildings, featuring outdoor seating, a large tree, and other




Figure 4

structural elements that influence crowd movement. Area 3 is a $14\text{ m} \times 11\text{ m}$ open area surrounded by substantial foliage. **While waypoints on the ground provide destination guidance, participants navigate freely between them without any prescribed paths or lanes, allowing for natural variations in movement while still adhering to the overall flow structure.** This fosters a dynamic and authentic crowd flow, capturing realistic movement patterns during data collection.

2.2 Experimental Results

We next present empirical results that validate the performance of our system across three experimental areas, accommodating crowd flows of up to (and including) 20 individuals, and with varying split ratios (see Sec. 4.8). We model the observed area as a $30\text{ m} \times 15\text{ m}$ rectangle, discretized at 0.25 m intervals along each axis to align with the average biacromial width of an individual [74]. Our implementation leverages OpenCV's Lucas-Kanade algorithm (Sec. 4.3) [75], with $N_o(\mathbf{x})$ set as a $2\text{ m} \times 2\text{ m}$ square surrounding each point within the sensing field of view, based on maximum expected displacement between consecutive frames. For noise filtering, we set the Kolmogorov p-value (Sec. 4.4) threshold at $p_{\text{th}} = 0.15$. Additionally, we apply the following filtering parameters: minimum area of connected flow component (Sec. 4.5): $\mu_S = 2\text{ m}^2$, corresponding to the spatial footprint

of small pedestrian groups [76], and median filter window (Sec. 4.5): $\mu_M = 0.25\text{ m}^2$, matching the approximate area occupied by an individual [77]. For constructing $\hat{G}$ (Sec. 4.7), we set $\Upsilon = 1\text{ m}$, and for split ratio calculation, we set the width of the edge bounding polygons to $\rho = 1\text{ m}$, both aligned with standard lateral clearance maintained during pedestrian movement [78]. To compute the curl and divergence fields (Sec. 4.9), we perform semantic flow analysis within a $|N_L(\mathbf{x})| = \mu_M = 0.25\text{ m}^2$ square around each point $\mathbf{x}$.

We conducted 18 spatially structured crowd flow experiments across three areas, with crowd sizes of up to (and including) 20 participants, and with a variety of flow split ratios. The ground-truth flow topology graphs, $G(V, E)$, for these experiments correspond to configurations C1 through C12 in Fig. 3. Among these, topologies C4 and C9 are used in two additional experiments each, with distinct flow-splitting ratios to evaluate the sensitivity of our pipeline to varying crowd distribution patterns. C11 is used in two additional experiments where both flows move in the same direction, one with simultaneous and one with staggered starts, to test occlusion. These topologies were designed to emulate pedestrian movement in environments where structured flows naturally emerge due to underlying spatial constraints, while accommodating flow splitting and merging at points within the sensing field of view. In each experiment, the radar was located at the center of the bottom periphery of the area. We

Config.	C10		C9	C9	C9	C8	C7	C6		C5		C4	C4	C3	C2			C1			
Vertex	v_1^{in}	v_1^{out}	v_1	v_1	v_1	v_1	v_1	v_1	v_2	v_1	v_2	v_1	v_1	v_1	v_1	v_2	v_3	v_1	v_2	v_3	v_4
True	0.50	0.50	0.20	0.50	0.80	0.25	0.25	0.33	0.50	0.53	0.53	0.20	0.50	0.25	0.5	0.66	0.75	0.5	0.5	0.5	0.5
Est.	0.47	0.47	0.23	0.38	0.70	0.20	0.20	0.26	0.56	0.38	0.32	0.29	0.47	0.16	0.68	0.63	0.55	0.61	0.69	0.51	0.33

Table 1

processed 30 seconds of radar data and applied our pipeline (Sec. 4.1, 4.2 and 4.7) to produce an estimate, $\hat{G}(\hat{V}, \hat{E})$ of the underlying flow topology, $G(V, E)$.

Our pipeline successfully reconstructed flow topologies, demonstrating strong visual alignment with the ground-truth topologies across 17 experiments. Fig. 4 illustrates samples of these high-fidelity reconstructions, which accurately capture the true crowd flow patterns, even for complex flow patterns, such as the branching structures in configurations C1, C2, and C6. However, our flow recovery fails in one unique case within C4. We discuss this scenario in more detail in Sec. 3, while for other C4 experiments, we successfully recover the flow topology (see Fig. 4).

To further assess our proposed approach, we evaluate the accuracy of the graph's embedding in 2D space. Determining a suitable metric for this task requires careful consideration, as participants move freely between nodes to maintain realistic crowd behavior, rather than following strict straight-line paths defined by the flow topologies. We therefore use a metric that is robust to small local shifts, while still penalizing large structural discrepancies. To this end, we evaluate the one-sided Chamfer distance d_{avg} [71, 79], given by the average minimum distance between a point on the estimated graph, $\hat{G}$, and any point on the true graph, G . More specifically, let $\mathcal{P}$ and $\hat{\mathcal{P}}$ denote the set of points that constitute the 2D embeddings of G and $\hat{G}$, respectively. We then calculate $d_{\text{avg}} = (1/|\hat{\mathcal{P}}|) \int_{\hat{\mathcal{P}}} \min_{p \in \mathcal{P}} \|\hat{p} - p\|_2 d\hat{p}$, where $|\hat{\mathcal{P}}|$ is the sum of the edge lengths of $\hat{G}$. We calculate d_{avg} for each of the 17 successfully reconstructed graphs, obtaining a mean d_{avg} of 0.45 m. We emphasize that although the ground-truth edge is represented as a 1D curve, in reality, the width of a typical pedestrian traffic lane is 0.76 m [72]. Thus, our estimated edges remain well within this natural width.

To further quantify accuracy, we associate edges in G and $\hat{G}$, and characterize errors in flow direction estimation by calculating the angular orientation error between the edges. For the 62 edges present in the successfully reconstructed graphs, we achieve an edge orientation MAE of 8.8° . Overall, these results demonstrate the robustness of our approach both in reconstructing the abstract flow graph and reliably localizing its embedding in 2D space.

We next demonstrate an application of this extracted directed geometric graph, $\hat{G}$, to develop an understanding of

the relationships between the regions of the observed space. Specifically, we estimate the flow split ratios into (out of) a node in $\hat{G}$ along each of the incoming (outgoing) edges using our proposed pipeline from Sec. 4.8. In each experiment that featured a splitting or merging point (configurations C1 through C10 of Fig. 3), we set flow split ratios beforehand and assigned participants to specific flow edges accordingly to ensure that the real crowd flows reflected these ratios. We note that the participants were not pre-grouped by their assigned directions so that they remained naturally mixed until they reached the splitting point. Table 1 compares our estimated flow split ratios to the true values, highlighting a great performance, with a Mean Absolute Error (MAE) of 0.10 across all experiments. Furthermore, when considering the simpler branching topologies in configurations C7-C9, the MAE drops to 0.05. For configurations C9 and C4, we conducted multiple experiments with different split ratios. Our estimation method not only accurately captured these variations but also effectively distinguished relative changes in the flow split ratio within the same topology. Importantly, flow split ratio estimation relies on an accurate $\hat{G}$, as its topology is essential for linking inflow with outflow directions.

2.3 Semantic Insights from Flow Fields

In results above, we have shown how to use the estimated flow field, $\hat{\mathbf{v}}_{\text{unit}}(\mathbf{x})$, to reconstruct the flow topology. Complementary to this reconstruction process, we next demonstrate the effectiveness of our semantic flow analysis pipeline (Sec. 4.9) in methodically discovering regions in the flow that exhibit significant local divergence or curl properties. These properties are important as they correspond to key crowd behaviors: high positive (negative) divergence is associated with flow splitting (merging), while high positive (negative) curl is associated with acute counterclockwise (clockwise) turns in the flow. We note that an additional low-pass filter is applied to $\hat{\mathbf{v}}_{\text{unit}}(\mathbf{x})$, to enhance the visualization of curl and divergence fields. We first present this analysis in the context of structured crowds before showing that this approach generalizes to diffuse crowds as well.

We next demonstrate how divergence and curl fields provide semantic information in structured crowd flows. To illustrate, Fig. 5 (Top) shows curl and divergence fields for configurations C3 and C12 in Fig. 3. In Fig. 5 (a), the curl field for C3 shows intense negative curl in region R1, where




Figure 5

the crowd flow makes an abrupt clockwise (CW) turn, with less intense positive curl around R2 and R3 indicating less acute counterclockwise (CCW) motion. Fig. 5 (b) shows the divergence field for C3, which indicates strong positive divergence at R4, where the flow splits, and strong negative divergence at R5, where the flow merges. Finally, Fig. 5 (c) shows the curl field for configuration C12, with negative and positive curl naturally segmenting the flow into regions of CW and CCW movement. Thus, our pipeline provides semantic understanding of the flows to complement the reconstruction of the flow topology discussed earlier, and we next illustrate how this analysis generalizes to diffuse flows.

The semantic flow analysis pipeline (Sec. 4.9) also provides useful insights into the structure of more diffuse flows. For such crowds, the KS denoising step of Sec. 4.4 requires parameter adjustment given the more random nature of the true flows. We therefore set the threshold parameter, $p_{th} = 1$ for these scenarios, which allows us to recover estimated flow fields rich enough to provide structural insights into crowd behavior. One can easily determine when to use each parameter, as we discuss in Sec. 3.

Fig. 5 (Bottom) presents the curl and divergence fields for configurations C13, C14, and C15 in Fig. 3, respectively. In C13, participants move in two diffuse lane-like flows separated by a boundary (red dotted line in Fig. 3, C13, see also Fig. 2 (b)), where the flow direction undergoes an abrupt 180-degree reversal. In C14, upon receiving a signal, participants evacuate radially outward from a designated starting position, simulating a crowd panic scenario. Conversely, C15 captures participants converging toward a central destination from multiple directions, modeling a gathering scenario.

As observed in Fig. 5 (d), the curl field of C13 distinctly highlights the flow separator where the two lane-like streams undergo direction reversal. Similarly, the location of the highest intensity peak in the divergence field of C14 lies close to the source of the crowd panic event, aiding in localizing anomaly origins in real-world scenarios (Fig. 5 (e)). In C15, a region of strong negative divergence corresponds to the collective gathering location (Fig. 5 (f)). Thus, our estimated flow fields provide sufficient detail to recover structural information via semantic analysis, even for diffuse flows.

3 Discussion

We next benchmark our approach against relevant methods and discuss directions for future work.

Since there is no prior work on mmWave-based crowd flow analytics, as discussed earlier, we next compare specific components of the pipeline to the related work in literature. For instance, direct application of classical optical flow methods from the area of vision produces visibly noisy flows, as shown in Fig. 6. Furthermore, mmWave flow extraction pipelines developed in other domains, e.g., [70], do not generalize for crowd flows, and processing our point clouds through [70], for instance, produced unusable flow fields due to the scale disparity between single human activity and crowd motion. These comparisons highlight our proposed flow field analysis as a significant advancement in the emerging domain of mmWave crowd flow analytics.

As stated in Sec. 2.2, our pipeline did not completely recover the flow topology in one of the 18 structured-crowd experiments. Specifically, we ran 3 experiments on configuration C4 with 80%, 50%, and 20% of the crowd taking the upper route at v_1 , and in the last of these, flow topology recovery




Figure 6

was partially successful (missed one edge), as shown in Fig. 7. Two key factors caused this incomplete recovery, both stemming from the highly unequal split ratio. First, significant occlusion, caused by the majority of the crowd traveling along the lower route, reduced visibility of the upper route. Second, a tree along the upper route produced persistent noisy returns that dominated the sparse crowd-related signals. Consequently, points pertaining to the true crowd flow were discarded as noise. For such cases, a more granular analysis over smaller temporal windows is an avenue for future work.

Our proposed graph construction framework can be further used to differentiate between structured and diffuse crowds, as graphs constructed for diffuse crowd flows exhibit distinguishing characteristics. One such indicator is the inconsistency between the directions of the inferred graph edges and those of the underlying estimated flow field, as measured by, *e.g.*, the average cosine similarity. Structured crowds show much higher consistency, as expected. A comprehensive exploration of such distinctive factors is also a future work direction.

In conclusion, this work represents a significant advancement in crowd flow analytics using commodity mmWave MIMO radar. Our method successfully bridges the gap between raw radar data and meaningful crowd behavior analysis by transforming radar point clouds into directed geometric graphs through interpretable flow fields. The high accuracy achieved across three areas validates the robustness of our approach, with high-fidelity graph reconstruction of the underlying flow structure in considerably complex crowd patterns, and accurate flow split ratio estimation with an MAE of 0.1. Key contributions of our work include a novel framework for the high-fidelity generation of mmWave crowd flow fields, a new approach for extracting the underlying crowd topology, and further quantifying flow ratios, as well as the application of vector field analysis to extract meaningful semantic information from crowd flows. By analyzing the Jacobian of the flow field, particularly through the divergence and curl, we provide deeper insights into crowd behavior, applicable to both spatially structured and diffused crowds. This is particularly valuable for characterizing spatially diffuse




Figure 7

flows where topology reconstruction alone proves insufficient. Our extensive experimental validation confirms the effectiveness of our methodology, establishing a solid foundation for future advancements in mmWave-radar-based crowd flow analysis.

4 Methods

We next present our methodology for crowd flow analysis, which begins with a novel point cloud extraction method, followed by our proposed flow field estimation procedure and concluding with the crowd graph reconstruction and semantic flow analysis pipeline.

4.1 Radar Point Cloud Generation

We start by summarizing mmWave propagation in FMCW MIMO radar systems, followed by presenting a robust methodology for generating point clouds corresponding to individuals in crowded environments. A point cloud is a collection of spatial data points representing reflective surfaces in the sensing field of view, and it serves as a key intermediate representation in several radar-based sensing works. Traditional approaches for point cloud generation rely on range-angle 2D-FFT, combined with Constant False Alarm Rate (CFAR) detection [63, 80]. However, they require careful parameter tuning and are sensitive to environmental variations. To improve robustness, we instead propose a modified approach in this section by performing multi-scale range-only peak detection over time before applying targeted Angle-of-Arrival (AoA) estimation at identified sensing depths. This method can reduce false detections while providing precise spatial representations of crowd dynamics. We next begin with a summary of FMCW radar transmission, followed by presenting the details of our approach to point cloud generation.

Consider a mmWave FMCW MIMO radar system consisting of N_{TX} transmitters (TX) and N_{RX} receivers (RX), where each TX periodically emits a chirp sinusoid with a frequency that increases linearly from f_0 . The transmitted frequency is given by $f_{TX}(t) = f_0 + St$, $0 \leq t < T_c$, where $S = B/T_c$ denotes the chirp slope, B is the bandwidth, and T_c is the chirp duration. The complex baseband signal of the m^{th} chirp of the i^{th}

TX is then given by $s_{\text{TX}}^{(i)}(m, t) = \sqrt{P_{\text{TX}}} \exp(j\{2\pi f_0 t + \pi S t^2\})$, where P_{TX} is the transmit power.

Next, consider an object located at a distance $d_j[m]$ from the j^{th} RX. The signal received at the j^{th} RX after scattering is given by $s_{\text{RX}}^{(i,j)}(m, t) = s_{\text{TX}}^{(i)}(m, t - \tau_j)/d_j[m]^2$, where $\tau_j = 2d_j[m]/c$ is the round-trip time of flight, and c is the speed of light. The received signal is then mixed with the transmitted chirp at each RX to yield the intermediate frequency (IF) signal [81, 82]: $s_{\text{IF}}^{(i,j)}(m, t) = s_{\text{RX}}^{(i,j)*}(m, t) \times s_{\text{TX}}^{(i)}(m, t)$. To this end, we use a TDM-MIMO setup where TXs emit chirps sequentially, avoiding the need to resolve signal arrivals from different TXs.

To improve angular resolution, we construct a virtual array by pairing each TX-RX element in order to synthesize an extended rectangular aperture, as is commonly done in prior work [83]. Defining virtual indices (p, q) with $p \in \{1, \dots, N_x\}$ and $q \in \{1, \dots, N_y\}$, we map each TX-RX pair (i, j) to a virtual array element (p, q) via a convolution of the TX array pattern over the RX array [84]. Reordering the received data accordingly forms a planar array, effectively increasing spatial samples from N_{RX} to $N_x N_y = N_{\text{TX}} N_{\text{RX}}$ and refining angular resolution beyond a single TX-RX pair.

Let $s_{\text{IF}}^{(p,q)}[m, n]$ denote the discrete IF signal, where n is the Analog-to-Digital Converter (ADC) bin index, $N_r = T_c/T_r$ represents the number of discrete range bins with T_r denoting the ADC sampling duration, N_c denotes the chirp duration, and $T_f = N_c T_c$ denotes the total time duration. Traditional point cloud generation methods [80, 85, 86] employ a range-Doppler-angle FFT followed by CFAR thresholding in order to extract point clouds over time. However, these methods exhibit several limitations. CFAR assumes a predefined radar noise distribution and requires the inversion of complex nonlinear operations to determine adaptive threshold values [87, 88]. Moreover, noise characteristics are environment-dependent, necessitating recalibration for each deployment scenario [89]. Lastly, CFAR applies a fixed-size sliding window [90] for threshold estimation, lacking a multi-scale perspective. This limitation reduces its effectiveness in detecting targets with varying spatial extents or in cases where noise characteristics change across scales.

In this paper, we then adopt a modified approach that enhances the quality of point cloud generation by first performing range detection and then applying targeted AoA estimation only at the detected ranges, thereby minimizing the false positives typically encountered in direct detection over 2D range-angle spectra. Specifically, we leverage the pipeline described in [25] to first generate Binary Trace Maps, which encode the occupancy of visible agents across range and time. This method involves adaptive multi-scale range peak detection that automatically adjusts to varying target dimensions and spatial characteristics, enabling more effective

denoising across different spatial scales. For completeness, we next outline the prerequisite steps.

First, we compute the Fourier transform along the ADC axis, *i.e.*, the range-FFT:

$$R_{\text{IF}}[m, r] = \sum_{n=0}^{N_r-1} s_{\text{IF}}[m, n] \exp\left(-j \frac{2\pi n r}{N_r}\right). \quad (1)$$

We note that range localization of visible agents does not utilize the virtual array, and thus we drop the (p, q) indices on $s_{\text{IF}}[m, n]$. Next, instead of the traditional Doppler Fourier transform step, we compute the Wrapped Phase Spectrum:

$$\zeta_{\text{IF}}[z, w, r] = \sum_{i=w}^{w+W} \text{Arg}(R_{\text{IF}}[i, r]) \exp\left(-j \frac{2\pi z i}{W}\right), \quad (2)$$

where $\text{Arg}(Ae^{j\gamma}) = \gamma \bmod 2\pi$ denotes the wrapped phase of a complex number, and w represents the index of the starting chirp in a sliding window of W chirps. We then compute the spectral bandwidth along the z -dimension of $\zeta_{\text{IF}}[z, w, r]$ as:

$$\mathcal{H}[w, r] = \Gamma_z(\zeta_{\text{IF}}[z, w, r]), \quad (3)$$

where $\Gamma_z(f(z)) \in \mathbb{R}$ is a suitable measure of the bandwidth of $f(z)$. Subsequently, we binarize $\mathcal{H}$ using an ensemble of pre-trained Gradient Boosted Machines [25] to obtain the Binary Trace Map $\mathcal{H}_{\text{b, vis}}[w, r]$, resulting in a set of detections:

$$\mathcal{D}_R[w] = \{r_i \mid \mathcal{H}_{\text{b, vis}}[w, r_i] = 1, 1 \leq i \leq K[w]\}, \quad (4)$$

where $K[w] = \sum_{r=0}^{N_r} \mathcal{H}_{\text{b, vis}}[w, r]$ represents the number of agents visible to the radar at chirp window w .

Having established the range detection approach, which localizes activity only along the range dimension over time, we next extend it by reintroducing the virtual array to perform AoA estimation only at the sensing depths detected in $\mathcal{D}_R[w]$ for each w , thereby robustly constructing point clouds of crowd activity over time. First, we denote the range-Doppler transform for the array element (p, q) over chirp window w as

$$Q_{\text{IF}}^{(p,q)}[w, v, r] = \mathcal{F}_{m,n}\{s_{\text{IF}}^{(p,q)}[m, n] \mathbb{1}_{w \leq m < w+W}\}.$$

For each range detection in Eq. (4), we then perform a 2D angle-FFT along the aperture dimensions (p, q) . Since Doppler information is not utilized in this work, we average over the Doppler axis. Furthermore, given that off-the-shelf radar boards inherently have poor elevation resolution, we sum over the elevation dimension after the angle-FFT step to obtain the effective range-azimuthal angle spectrum:

$$\mathcal{A}[w, \theta, r] = \sum_{\phi} A_{\text{IF}}[w, \theta, \phi, r], \text{ where} \\ A_{\text{IF}}[w, \theta, \phi, r] = \frac{1}{N_c} \sum_{v=0}^{N_c-1} \mathcal{F}_{p,q}\{Q_{\text{IF}}^{(p,q)}[w, v, r]\}. \quad (5)$$

Finally, we construct our point cloud by iterating over each detected range $r_i \in \mathcal{D}_R[w]$ to calculate the corresponding




Figure 8

bearing angle, θ_i , that maximizes the azimuthal spectrum $\mathcal{A}$ of Eq. (5) at range r_i :

$$\mathcal{D}[w] = \{(r_i, \theta_i = \arg \max_{\theta} \mathcal{A}[w, \theta, r_i]) \mid r_i \in \mathcal{D}_R[w]\}, \quad (6)$$

which results in the final point cloud representation. While this representation does not fully resolve individual trajectories due to occlusions, noise, and overlapping detections, it effectively captures the overall crowd dynamics, enabling the inference of collective movement trends, as we shall see. We next introduce the concept of a flow field, which helps transform these discrete, sporadic, and noisy observations into efficient spatial representations of crowd motion patterns, enabling accurate reconstruction of crowd flow topologies and inference of underlying flow behaviors and semantics.

4.2 Flow Field Estimation Methodology

The core step between point cloud generation and graph reconstruction is the estimation of the flow field, a 2D vector field, $\mathbf{v}(\mathbf{x}) : \mathbb{R}^2 \rightarrow \mathbb{R}^2$, which characterizes the aggregate movement of individuals in the crowd. We next describe our method for estimating $\mathbf{v}(\mathbf{x})$. We first leverage a classical optical flow method to extract instantaneous flows at each time step from the sequence of point clouds (Sec. 4.1). We then present novel filtering and interpolation techniques tailored for the sparse, noisy nature of mmWave data to construct an estimated flow field, $\hat{\mathbf{v}}(\mathbf{x})$, which is then used in Sec. 4.6 as part of our proposed methodology to uncover crowd flow topologies and extract crowd flow semantics. Fig. 8 outlines our approach.

4.3 Pairwise Flow Estimation

As illustrated in Fig. 8, the first step towards estimating the flow field, $\mathbf{v}(\mathbf{x})$, is to transform the sequence of point clouds into a sequence of sparse vector fields. More specifically, for

each point $\mathbf{x}_{w,i} = (r_i[w] \cos(\theta_i[w]), r_i[w] \sin(\theta_i[w]))$ in a point cloud $\mathcal{D}[w]$, we estimate the flow at that point, $\hat{\mathbf{v}}_{w,i}$, by comparison to the subsequent point cloud, $\mathcal{D}[w+1]$. If $\mathbf{x}_{w,i} \in \mathcal{D}[w]$ evolves according to a flow field $\mathbf{v}(\mathbf{x})$, then at the next radar sampling window, $w+1$, we will find a radar detection at a location $\mathbf{x}_{w+1,i'} = \mathbf{x}_{w,i} + \mathbf{v}(\mathbf{x}_{w,i})$, wherein we assume that time in the flow field is normalized with respect to the radar sampling rate. Thus, if there were no errors in observation and if point association from one point cloud to the next were known, we could precisely recover the flow field at the observed point as $\mathbf{v}(\mathbf{x}_{w,i}) = \mathbf{x}_{w+1,i'} - \mathbf{x}_{w,i}$.

However, associating points between point clouds is challenging [67, 91] due to the large number of matches required, compounded by missed detections and false alarms. Furthermore, the presence of noise implies that even if associations are known, the resulting flow estimation would be inexact.

To address these challenges, we leverage well-established optical flow methods from computer vision, which are designed to infer motion between successive image frames. Specifically, we employ the classical Lucas-Kanade algorithm [75, 92] on successive point clouds, $\mathcal{D}[w]$ and $\mathcal{D}[w+1]$ to find the pairwise flow (PWF) $\hat{\mathbf{v}}_{w,i}$ for all $\mathbf{x}_{w,i} \in \mathcal{D}[w]$. As Lucas-Kanade requires inputs in image format, we first discretize the space and create a binary occupancy map, $\mathcal{H}[w](\mathbf{x}) \in \{0, 1\}$, where $\mathcal{H}[w](\mathbf{x}) = 1$ if there exists a point in $\mathcal{D}[w]$ at location $\mathbf{x}$. For each point $\mathbf{x}_{w,i} \in \mathcal{D}[w]$, the Lucas-Kanade algorithm then estimates the corresponding flow, $\hat{\mathbf{v}}_{w,i}$, by solving the following least-squares optimization problem:

$$\hat{\mathbf{v}}_{w,i} = \arg \min_{\mathbf{h}} \sum_{\mathbf{y} \in \mathcal{N}_o(\mathbf{x}_{w,i})} \|\mathcal{H}[w+1](\mathbf{y} + \mathbf{h}) - \mathcal{H}[w](\mathbf{y})\|^2. \quad (7)$$

Here, $\mathcal{N}_o(\mathbf{x}_{w,i})$ denotes a local neighborhood around $\mathbf{x}_{w,i}$, and $\hat{\mathbf{v}}_{w,i} \in \mathbb{R}^2$ is the displacement vector that minimizes the intensity difference between successive frames.

Due to the sparsity of point clouds, each point in space will not have an estimated flow associated with it at every time step, but after computing the PWFs for all consecutive pairs of point clouds, we can define a (possibly empty) set of PWFs at each point, $\psi(\mathbf{x}) = \{\hat{\mathbf{v}}_{w,i} | \mathbf{x}_{w,i} = \mathbf{x}\}$, which gives all estimated flows at $\mathbf{x}$ over time. We next propose a method of statistical analysis on these sets to discard unreliable flows.

4.4 Flow Field Denoising with Kolmogorov-Smirnov Statistical Tests

To produce the best estimate of the flow field at a point, $\mathbf{x}$, we must incorporate information from all relevant PWFs contained in the set $\psi(\mathbf{x})$. A straightforward approach is to directly estimate $\mathbf{v}(\mathbf{x})$ by averaging over $\psi(\mathbf{x})$:

$$\mathbf{v}_{\text{TAF}}(\mathbf{x}) = \begin{cases} \frac{1}{|\psi(\mathbf{x})|} \sum_{\hat{\mathbf{v}}_{w,i} \in \psi(\mathbf{x})} \hat{\mathbf{v}}_{w,i} & \text{If } |\psi(\mathbf{x})| > 0 \\ 0 & \text{Otherwise} \end{cases}, \quad (8)$$

and we refer to this initial estimate as the Time-Averaged Flow Field (TAF). However, this direct approach is ineffective as it fails to distinguish between valid flows generated by crowd movement and spurious returns due to environmental factors [93, 94] or multipath effects.

To address these challenges, we propose a novel denoising strategy that allows us to generate high-integrity flow fields. Our key insight is that a truly noisy flow exhibits no consistent directional preference, resulting in flow angles that are uniformly distributed across all possible directions. In contrast, legitimate flow patterns demonstrate consistent directionality, even in the presence of measurement noise.

We formalize this analysis by modeling each PWF in $\psi(\mathbf{x})$ as the sum of a true crowd flow component plus a noise term: $\hat{\mathbf{v}}_{w,i} = \mathbf{v}(\mathbf{x}) + \eta$. Here, η is the realization of a zero-mean isotropic 2D Gaussian random vector with covariance $\sigma^2(\mathbf{x})I_2$, where I_2 is the 2×2 identity matrix, and the noise power at $\mathbf{x}$, $\sigma^2(\mathbf{x})$, is determined by spatially varying environmental factors. If the noise power is much greater than the magnitude of the true flow, $\mathbf{v}(\mathbf{x})$, then the PWFs in $\psi(\mathbf{x})$ are approximately drawn from a zero-mean 2D Gaussian random variable, and consequently, the empirical distribution of the angles of the flows is consistent with a uniform distribution over $[0, 2\pi)$ [95]. When this is the case, the PWFs at $\mathbf{x}$ are discarded since they do not carry reliable information about the true flow field, $\mathbf{v}(\mathbf{x})$.

To this end, we employ a Kolmogorov-Smirnov (KS) statistical hypothesis test [96]. At each spatial location $\mathbf{x}$, we calculate the empirical CDF of the PWF angles:

$$\Omega_{\mathbf{x}}(\alpha) = \frac{1}{|\psi(\mathbf{x})|} \sum_{\hat{\mathbf{v}}_{w,i} \in \psi(\mathbf{x})} \mathbb{1}_{\angle \hat{\mathbf{v}}_{w,i} \leq \alpha}, \quad (9)$$

where $\angle \hat{\mathbf{v}}_{w,i}$ denotes the orientation of the flow vector. We then test whether the empirical CDF of flow angles, $\Omega_{\mathbf{x}}(\alpha)$, differs significantly from the CDF of a uniform distribution on $[0, 2\pi)$, denoted by $F_{\text{unif}}(\alpha)$. The KS test computes the test statistic, $D_{\mathbf{x}} = \sup_{\alpha \in [0, 2\pi)} |\Omega_{\mathbf{x}}(\alpha) - F_{\text{unif}}(\alpha)|$, and then uses the complementary CDF of the Kolmogorov distribution, κ , to compute the p -value, $p_{\mathbf{x}} = \kappa(D_{\mathbf{x}})$ [97].

We next establish a threshold, p_{th} , for classification: points with $p_{\mathbf{x}} > p_{\text{th}}$ are classified as having noisy flow, as we cannot reject the hypothesis that their flow directions follow a uniform distribution. Conversely, points with $p_{\mathbf{x}} \leq p_{\text{th}}$ exhibit a statistically significant directional preference, indicating meaningful flow. By applying this test to each point in our field of view, we generate a binary mask defined as $M(\mathbf{x}) = \mathbb{1}_{p_{\mathbf{x}} \leq p_{\text{th}}}$. We use this mask to filter out noisy points from the TAF, resulting in a Denoised Flow Field (DFF), $\mathbf{v}_{\text{DFF}}(\mathbf{x}) = M(\mathbf{x}) \odot \mathbf{v}_{\text{TAF}}(\mathbf{x})$, where $\odot$ indicates the Hadamard product. The threshold, p_{th} , represents a tunable parameter that balances between retaining meaningful flow information and eliminating noise.

Overall, our proposed KS filtering significantly reduces noise by testing for the temporal consistency of flows at each point. To further refine our flow field estimate, we next introduce complementary filtering steps based on the spatial correlations naturally present in crowd movement.

4.5 Flow Pruning and Smoothing

As a final step in the flow field estimation process, we exploit the spatial correlation present in natural crowd flows to refine $\mathbf{v}_{\text{DFF}}(\mathbf{x})$. In particular, points along dominant movement paths form contiguous regions, so that points with valid flows which are possibly erroneously rejected as noise are often embedded in regions that, overall, show strong, well-defined flows. To account for this spatial correlation, we perform morphological pruning and spatial smoothing.

We employ morphological pruning to account for small clusters of spurious flows that remain detached from the dominant flow topology. These artifacts typically arise due to transient reflections, multipath effects, or local inconsistencies in radar detections, and they do not meaningfully contribute to the flow structure. To address this, we employ a connected-component analysis [98] to filter out small, disjoint flow regions. Given a predefined size threshold of μ_S m², this operation removes any contiguous clusters of flow vectors that fall below μ_S , preserving only the large-scale flow structures. Fig. 8 illustrates the impact of this procedure.

As a final filtering step, we eliminate erratic flow variations that are atypical in real crowds through median filtering, a widely recognized technique for improving optical flow estimation [99]. We apply two consecutive rounds of median filtering with a window size of μ_M m². Specifically, in the




Figure 9

first round, we compute the median considering only neighboring non-zero flow values, while in the second, we apply standard median filtering. Beyond smoothing the flow, this two-step process has an overall effect similar to morphological closure [100], a classical image processing technique for noise removal that combines dilation and erosion. The first median filter acts as a dilation, interpolating flow at zero-flow points with at least one non-zero neighbor. The second functions similarly to erosion, setting the flow to zero if more than half of the neighboring flows are zero.

By combining statistical KS-based filtering, morphological pruning, and spatial smoothing, we obtain our final estimated flow field, $\hat{\mathbf{v}}(\mathbf{x})$, that is significantly more robust to spurious noise while retaining meaningful motion patterns (see Fig. 6 and related discussion in Sec. 3). This refined flow field provides the basis for constructing a compact graph representation of crowd movement, capturing dominant flow pathways and connectivity patterns, which we explore next.

4.6 Crowd Flow Analysis Methodology

We next introduce our proposed methodology for extracting the crowd flow topology, $\hat{G}(\hat{\mathbf{V}}, \hat{\mathbf{E}})$ from the estimated flow field, $\hat{\mathbf{v}}(\mathbf{x})$, as illustrated in Fig. 9. We first present our topology reconstruction pipeline, specifically designed for structured flows, followed by a robust estimation framework for flow split ratios at the detected vertices, $\hat{\mathbf{V}}$. We then derive the divergence and curl fields of $\hat{\mathbf{v}}(\mathbf{x})$, which provide complementary semantic insights into accumulation, dispersion, and rotation in the underlying flow.

4.7 Flow Topology Reconstruction

We infer the flow topology from the final flow field by employing a systematic approach that extracts the core structure and identifies critical nodes and edges, as shown in Fig. 9. We begin by refining our final flow field, $\hat{\mathbf{v}}(\mathbf{x})$, through a normalization process. Specifically, for each point $\mathbf{x}$, we normalize the flow vector to unit magnitude: $\hat{\mathbf{v}}_{\text{unit}}(\mathbf{x}) = \hat{\mathbf{v}}(\mathbf{x}) / \|\hat{\mathbf{v}}(\mathbf{x})\|$

if $\hat{\mathbf{v}}(\mathbf{x}) \neq \mathbf{0}$, else $\hat{\mathbf{v}}_{\text{unit}}(\mathbf{x}) = \mathbf{0}$. We then consider the scalar field $s(\mathbf{x}) = \|\hat{\mathbf{v}}_{\text{unit}}(\mathbf{x})\|$, which effectively represents a binary image where the non-zero regions correspond to the presence of flow. To extract the dominant flow structure, we use the Zhang-Suen skeletonization algorithm [101, 102] on $s(\mathbf{x})$, which iteratively applies local connectivity rules to erode boundary pixels while preserving the topology (Fig. 9). This process yields a flow skeleton $s_{\text{skel}}(\mathbf{x})$, which is a single-pixel wide binary representation that captures the essential flow pathways. We next implement a path tracing algorithm that traverses the skeleton, extracting these pathways which then serve as precursors to the edges in our flow graph.

To this end, we utilize an 8-connected neighborhood convolution kernel [103] to examine the local pixel neighborhood, enabling the systematic identification of critical points on the skeleton. Specifically, we classify branch points as pixels with more than two neighbors, and start/end points as pixels with exactly one neighbor. The algorithm then constructs skeletal traces, defined as connected paths within the skeleton, and represented as an ordered sequence of coordinates $T_i = \{\mathbf{t}_1, \mathbf{t}_2, \dots, \mathbf{t}_m\}$. We employ a systematic traversal methodology by iteratively exploring neighboring pixels, starting from each start/end point and subsequently from any unvisited branch points. Thus, $\mathbf{t}_1$ and $\mathbf{t}_m$ are either branch points or start/end points, whereas any other $\mathbf{t}_j \in T_i$ has exactly two neighbors. A global visited coordinate set is maintained throughout to ensure that each pixel is processed only once, preventing redundant traversal. We denote the set of all such skeletal traces as $\mathcal{T} = \{T_1, T_2, \dots, T_n\}$, forming the basis for building our directed geometric graph.

Constructing a directed geometric graph from skeletal traces presents two significant challenges: establishing connectivity at flow junctions and ensuring that the edge directionality aligns with the flow field. We achieve this through a structured approach for graph construction, as detailed in Fig. 9. First, the vertex set, $\hat{\mathbf{V}}$, is initialized using the endpoints of skeletal traces in $\mathcal{T}$, and each trace T is added as an edge,

$e = (t_s, t_e, T)$, where t_s and t_e denote the start and end points of the trace. Next, a vertex-trace proximity analysis is performed, where each vertex is examined for close proximity to existing traces; if a vertex lies within a distance threshold, Υ , the corresponding trace is partitioned, and new edges are created to incorporate the vertex. Finally, a trace direction validation step ensures consistency with the underlying flow field by reversing the orientation of edges misaligned with the flow direction. The resulting graph, $\hat{G} = (\hat{V}, \hat{E})$, encodes the topological structure of the skeletal traces while integrating spatial and directional constraints. We note that points on edges in the geometric graph are naturally assigned a direction by referencing $\hat{\mathbf{v}}_{\text{unit}}(\mathbf{x})$.

We next show the use of the reconstructed geometric graph, $\hat{G}$, in conjunction with the point clouds, $\mathcal{D}$ (see Eq. (6)), to determine the flow split ratios at each vertex in $\hat{V}$.

4.8 Flow Split Ratio Estimation

In this subsection, we augment our estimated crowd flow topology, $\hat{G}$, by finding flow split ratios, $\lambda_e \in [0, 1]$, for each edge $e \in \hat{E}$. These ratios represent the relative flow distribution along the edges of $\hat{G}$. To achieve this, we employ poly-line buffering [104] with a width parameter ρ to generate a bounding polygon, $\mathcal{B}_e$, around the skeletal trace T corresponding to the edge e . To find the relative occupancy of each edge, we then determine the number of points over all point clouds, $\mathcal{D}[w]$, that lie within the bounding polygon, and we divide this number by the total area of the polygon to get the number of returns per unit area associated with each edge. Formally, we find unnormalized weights $\tilde{\lambda}_e = \sum_{\mathbf{x}_{w,i} \in \mathcal{D}} \mathbb{1}_{\mathbf{x}_{w,i} \in \mathcal{B}_e} / |\mathcal{B}_e|$, where $\mathcal{D} = \cup_{w=0}^W \mathcal{D}[w]$ and $|\mathcal{B}_e|$ denotes the area of the bounding polygon $\mathcal{B}_e$. Finally, for each vertex $q \in \hat{V}$, we identify the set of outgoing edges, $\hat{V}_{\text{out}}(q)$, and calculate the flow split ratios as $\lambda_e = \tilde{\lambda}_e / \sum_{e' \in \hat{V}_{\text{out}}(q)} \tilde{\lambda}_{e'}$, $e \in \hat{V}_{\text{out}}(q)$, ensuring that $\sum_{e \in \hat{V}_{\text{out}}(q)} \lambda_e = 1$.

The flow split ratios provide insight into where the crowd flows after passing through a vertex. By normalizing with respect to incoming rather than outgoing edges, analogous ratios can be computed to describe where the crowd was before arriving at the vertex. With the weighted directed graph, $\hat{G}$, fully constructed, we next analyze the local flow Jacobian to extract semantic insights, offering a deeper understanding of the underlying dynamics and local features.

4.9 Jacobian-Based Semantic Flow Analysis

While the flow topology reconstruction of Sec. 4.7 effectively captures spatially structured flows, it can be complemented with additional analysis to reveal semantic information about the flow field, such as inflow (convergence), outflow (divergence), or rotational tendencies (clockwise/counterclockwise

curl), even at points that are not vertices of $\hat{G}$. Moreover, for spatially diffuse flows, the extracted topology may lack meaningful structure. To this end, we next introduce an approach inspired by dynamical systems that identifies key regions without relying on topology reconstruction.

We can describe the dynamics of a particle at a location $\mathbf{x} = (x, y)$ in the flow field $\hat{\mathbf{v}}_{\text{unit}}(\mathbf{x}) = [\hat{v}_x(x, y), \hat{v}_y(x, y)]^T$ using a continuous-time equation $\frac{d\mathbf{x}}{dt} = \hat{\mathbf{v}}_{\text{unit}}(\mathbf{x})$. This governing equation can be expanded as a Taylor series [105] around a suitable point $\mathbf{x}_0 = (x_0, y_0)$ near $\mathbf{x}$, yielding $\frac{d\mathbf{x}}{dt} = \hat{\mathbf{v}}_{\text{unit}}(\mathbf{x}) \approx \hat{\mathbf{v}}_{\text{unit}}(\mathbf{x}_0) + \mathcal{J}(\mathbf{x}_0) \cdot (\mathbf{x} - \mathbf{x}_0)$, where

$$\mathcal{J}(\mathbf{x}_0) = \begin{bmatrix} \frac{\partial \hat{v}_x}{\partial x} & \frac{\partial \hat{v}_x}{\partial y} \\ \frac{\partial \hat{v}_y}{\partial x} & \frac{\partial \hat{v}_y}{\partial y} \end{bmatrix} \bigg|_{(x_0, y_0)}. \quad (10)$$

The Jacobian, $\mathcal{J}(\mathbf{x}_0)$, describes how the flow varies around $\mathbf{x}_0$ [106, 107], and it includes all information needed to calculate the 2D curl and divergence at $\mathbf{x}_0$. Specifically, we have [108]

$$\nabla \cdot \hat{\mathbf{v}}_{\text{unit}} = \frac{\partial \hat{v}_x}{\partial x} + \frac{\partial \hat{v}_y}{\partial y} \text{ and } \nabla \times \hat{\mathbf{v}}_{\text{unit}} = \left(\frac{\partial \hat{v}_y}{\partial x} - \frac{\partial \hat{v}_x}{\partial y} \right), \quad (11)$$

for the divergence and curl, respectively. Intuitively, divergence quantifies the extent of gathering (negative) or dispersing (positive), while curl measures clockwise (negative) or counterclockwise (positive) rotation.

To characterize the local flow dynamics, we compute the curl and divergence scalar fields of the vector flow field. To do so, we must obtain the Jacobian, $\mathcal{J}(\mathbf{x})$, at each point, which in practice is typically computed using finite difference methods. However, this approach is not robust to localized errors in the estimated flow and further suffers at the boundary of regions with zero flow (where $\hat{\mathbf{v}}_{\text{unit}}(\mathbf{x}) = \mathbf{0}$). We therefore propose a novel procedure for estimating $\mathcal{J}$ by solving a simple least-squares convex program.

Specifically, recall that for a point, $\mathbf{y}$, in a neighborhood $\mathcal{N}_L(\mathbf{x})$ around $\mathbf{x}$, we have $\hat{\mathbf{v}}_{\text{unit}}(\mathbf{y}) \approx \hat{\mathbf{v}}_{\text{unit}}(\mathbf{x}) + \mathcal{J}(\mathbf{x}) \cdot (\mathbf{y} - \mathbf{x})$. The estimated Jacobian, $\hat{\mathcal{J}}(\mathbf{x})$, is then given as

$$\hat{\mathcal{J}}(\mathbf{x}), \mathbf{b}^* = \arg \min_{J, \mathbf{b}} \sum_{\mathbf{y} \in \mathcal{N}_L(\mathbf{x})} \|\hat{\mathbf{v}}_{\text{unit}}(\mathbf{y}) - (\mathbf{b} + J \cdot (\mathbf{y} - \mathbf{x}))\|^2, \quad (12)$$

with $\mathbf{b}^*$ giving a smoothed estimate of $\hat{\mathbf{v}}_{\text{unit}}(\mathbf{x})$. We then use the estimated Jacobian, $\hat{\mathcal{J}}(\mathbf{x})$, to find the divergence and curl field at the point $\mathbf{x}$, as shown in Eq. (11). We note that if $\hat{\mathbf{v}}_{\text{unit}}(\mathbf{y}) = \mathbf{0}$, we do not include $\mathbf{y}$ in the neighborhood, as a flow of $\mathbf{0}$ indicates the point is not part of the region of interest.

5 Data Availability

We plan to make our data and codes available once the paper clears the review and tech transfer process.

6 Acknowledgements

This work is funded in part by ONR award N00014-23-1-2715 and in part by NSF CNS award 2215646.

7 Author Contributions

A.P., W.H., and Y.M. jointly contributed to the development of the research, including the conception of ideas, the design of the experimental framework, and the drafting and revision of the manuscript. All authors (A.P., W.H., and Y.M.) have read and approved the final version of the manuscript.

8 Competing Interests

Dr. Yasamin Mostofi serves as a Guest Editor for this collection. Anurag Pallaprolu and Winston Hurst declare no competing interests.

References

- [1] Greco, A., Saggese, A. & Vento, B. A Robust and Efficient Overhead People Counting System for Retail Applications. In *Image Analysis and Processing – ICIAP 2022: 21st International Conference*, 139–150 (Springer-Verlag, Berlin, Heidelberg, 2022). URL https://doi.org/10.1007/978-3-031-06430-2_12.
- [2] Nguyen, K., Le, M., Martin, B., Cil, I. & Fookes, C. When AI meets store layout design: a review. *Artificial Intelligence Review* **55**, 5707–5729 (2022).
- [3] Charankumar, P., Sunil, B., Nishant, Y., Vamsi, M. S. & Hemalatha, M. Using Existing CCTV Networks for Crowd Management and Crime Prevention using Deep Learning. In *2024 15th International Conference on Computing Communication and Networking Technologies (ICCCNT)*, 1–10 (IEEE, 2024).
- [4] Aldahlawi, R. Y., Akbari, V. & Lawson, G. A systematic review of methodologies for human behavior modelling and routing optimization in large-scale evacuation planning. *International Journal of Disaster Risk Reduction* **110**, 104638 (2024). URL <https://www.sciencedirect.com/science/article/pii/S221242092400400X>.
- [5] Suraj, B., Jaiwanth, T., Ghei, A., Srinivas, K. *et al.* Dynamic Crowd Management Through Potential Field-Based Simulations. In *2024 8th International Conference on Smart Grid and Smart Cities (ICSGSC)*, 341–346 (IEEE, 2024).
- [6] Kujareanpaisal, P. *et al.* BMA for the BMA: Boosting Mobility Analysis for the Bangkok Metropolitan Administration via Automated Pedestrian Counting from CCTV. In *2024 Joint 13th International Conference on Soft Computing and Intelligent Systems and 25th International Symposium on Advanced Intelligent Systems (SCIS&ISIS)*, 1–6 (IEEE, 2024).
- [7] Guardian, T. Is it time to ring the alarm on internet door cameras? *The Guardian* (2025). URL <https://www.theguardian.com/world/2025/feb/09/is-it-time-to-ring-the-alarm-on-internet-door-cameras>.
- [8] Monde, L. Let's beware of a post-Olympic drift in the use of AI-powered video surveillance. *Le Monde* (2024). URL https://www.lemonde.fr/en/opinion/article/2024/09/26/let-s-beware-of-a-post-olympic-drift-in-the-use-of-ai-powered-video-surveillance_6727335_23.html.
- [9] Chronicle, S. F. Here's why I am opting out of TSA facial recognition - and you should, too. *San Francisco Chronicle* (2024). URL <https://www.sfchronicle.com/opinion/openforum/article/facial-recognition-tsa-19998085.php>.
- [10] News.com.au. No warning: Disturbing find in Aussie tree sparks mass surveillance outrage. *News.com.au* (2024). URL <https://www.news.com.au/technology/environment/no-warning-disturbing-find-in-aussie-tree-sparks-mass-surveillance-outrage/news-story/02ad245ff18c8924437549971edad206>.
- [11] Birnbaum, E. Privacy activists launch campaign against airlines using facial recognition technology. *The Hill* (2019). URL <https://thehill.com/policy/technology/447066-privacy-activists-launch-campaign-against-airlines-using-facial-recognition/>.
- [12] Shahbazian, R. & Trubitsyna, I. Human sensing by using radio frequency signals: A survey on occupancy and activity detection. *IEEE Access* **11**, 40878–40904 (2023).
- [13] Korany, B. & Mostofi, Y. Counting a stationary crowd using off-the-shelf wifi. In *Proceedings of the 19th annual international conference on mobile systems, applications, and services*, 202–214 (2021).
- [14] Depatla, S., Muralidharan, A. & Mostofi, Y. Occupancy estimation using only WiFi power measurements. *IEEE Journal on Selected Areas in Communications* **33**, 1381–1393 (2015).
- [15] Yang, Y., Cao, J., Liu, X. & Liu, X. Wi-Count: Passing people counting with COTS WiFi devices. In *2018 27th International Conference on Computer Communication and Networks (ICCCN)*, 1–9 (IEEE, 2018).
- [16] Sharma, A., Li, J., Mishra, D., Batista, G. & Seneviratne, A. Passive WiFi CSI sensing based machine learning framework for COVID-Safe occupancy monitoring. In *2021 IEEE international conference on Communications workshops (ICC workshops)*, 1–6 (IEEE, 2021).
- [17] Pallaprolu, A., Korany, B. & Mostofi, Y. Wiffract: a new foundation for RF imaging via edge tracing. In *Proceedings of the 28th annual international conference on mobile computing and networking*, 255–267 (2022).
- [18] Cai, H., Korany, B., Karanam, C. R. & Mostofi, Y. Teaching rf to sense without rf training measurements. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* **4**, 1–22 (2020).
- [19] Karanam, C. R., Korany, B. & Mostofi, Y. Tracking from one side: Multi-person passive tracking with WiFi magnitude measurements. In *Proceedings of the 18th International Conference on Information Processing in Sensor Networks*, 181–192 (2019).
- [20] Miao, F. *et al.* Wi-Fi sensing techniques for human activity recognition: Brief survey, potential challenges, and research directions. *ACM Computing Surveys* **57**, 1–30 (2025).
- [21] Tan, S., Ren, Y., Yang, J. & Chen, Y. Commodity WiFi sensing in ten years: Status, challenges, and opportunities. *IEEE Internet of Things Journal* **9**, 17832–17843 (2022).
- [22] Semaan, E., Tejedor, E., Kochhar, R. K., Magnusson, S. & Parkvall, S. 6G spectrum - enabling the future mobile life beyond 2030. Tech. Rep., Ericsson (2024).
- [23] Li, S. & Hishiyama, R. An Indoor People Counting and Tracking System using mmWave sensor and sub-sensors. *IFAC-PapersOnLine* **56**, 7096–7101 (2023).
- [24] Vaidya, G. & Zuniga, M. Exploiting mmWave and Deep-Learning Models to Estimate People Count in Urban Scenarios. In *2024 IEEE/ACM Ninth International Conference on Internet-of-Things Design and Implementation (IoTDI)*, 73–84 (IEEE, 2024).
- [25] Pallaprolu, A., Peng, P., Sandhu, S., Hurst, W. & Mostofi, Y. Crowd Analytics with a Single mmWave Radar. In *Proceedings of the 30th Annual International Conference on Mobile Computing and Networking*, 830–844 (2024).
- [26] Texas Instruments. *AWR2243 Evaluation Module (AWR2243BOOST) mmWave Sensing Solution* ([n. d.]). Available online at <https://www.ti.com/lit/ug/spruit8d/spruit8d.pdf>.

- [27] Bendali-Braham, M., Weber, J., Forestier, G., Idoumghar, L. & Muller, P.-A. Recent trends in crowd analysis: A review. *Machine Learning with Applications* **4**, 100023 (2021).
- [28] Celes, C., Boukerche, A. & Loureiro, A. A. Crowd management: A new challenge for urban big data analytics. *IEEE Communications Magazine* **57**, 20–25 (2019).
- [29] Liao, P. *et al.* Applying crowdsourcing techniques in urban planning: A bibliometric analysis of research and practice prospects. *Cities* **94**, 33–43 (2019).
- [30] Helbing, D. *et al.* Saving human lives: What complexity science and information systems can contribute. *Journal of statistical physics* **158**, 735–781 (2015).
- [31] Adrian, J., Seyfried, A. & Sieben, A. Crowds in front of bottlenecks at entrances from the perspective of physics and social psychology. *Journal of the Royal Society Interface* **17**, 20190871 (2020).
- [32] Alfarano, A., Maiano, L., Papa, L. & Amerini, I. Estimating optical flow: A comprehensive review of the state of the art. *Computer Vision and Image Understanding* 104160 (2024).
- [33] Zhai, M., Xiang, X., Lv, N. & Kong, X. Optical flow and scene flow estimation: A survey. *Pattern Recognition* **114**, 107861 (2021).
- [34] Ali, S. & Shah, M. A lagrangian particle dynamics approach for crowd flow segmentation and stability analysis. In *2007 IEEE conference on computer vision and pattern recognition*, 1–6 (IEEE, 2007).
- [35] Ali, S. & Shah, M. Floor fields for tracking in high density crowd scenes. In *Computer Vision—ECCV 2008: 10th European Conference on Computer Vision, Marseille, France, October 12–18, 2008, Proceedings, Part II* **10**, 1–14 (Springer, 2008).
- [36] Mehran, R., Moore, B. E. & Shah, M. A streakline representation of flow in crowded scenes. In *Computer Vision—ECCV 2010: 11th European Conference on Computer Vision, Heraklion, Crete, Greece, September 5–11, 2010, Proceedings, Part III* **11**, 439–452 (Springer, 2010).
- [37] Solmaz, B., Moore, B. E. & Shah, M. Identifying behaviors in crowd scenes using stability analysis for dynamical systems. *IEEE transactions on pattern analysis and machine intelligence* **34**, 2064–2070 (2012).
- [38] Nayak, N. M., Zhu, Y. & Roy-Chowdhury, A. K. Vector field analysis for multi-object behavior modeling. *Image and Vision Computing* **31**, 460–472 (2013).
- [39] Mehran, R., Oyama, A. & Shah, M. Abnormal crowd behavior detection using social force model. In *2009 IEEE conference on computer vision and pattern recognition*, 935–942 (IEEE, 2009).
- [40] Zhao, J., Xu, Y., Yang, X. & Yan, Q. Crowd instability analysis using velocity-field based social force model. In *2011 Visual Communications and Image Processing (VCIP)*, 1–4 (IEEE, 2011).
- [41] Shi, X. *et al.* Flowformer++: Masked cost volume autoencoding for pretraining optical flow estimation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 1599–1610 (2023).
- [42] Lu, Y. *et al.* Optical Flow as Spatial-Temporal Attention Learners. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2024).
- [43] Bartoli, F., Lisanti, G., Ballan, L. & Del Bimbo, A. Context-aware trajectory prediction. In *2018 24th international conference on pattern recognition (ICPR)*, 1941–1946 (IEEE, 2018).
- [44] Alahi, A. *et al.* Social LSTM: Human trajectory prediction in crowded spaces. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 961–971 (2016).
- [45] Farooq, M. U., Saad, M. N. M. & Khan, S. D. Motion-shape-based deep learning approach for divergence behavior detection in high-density crowd. *The Visual Computer* **38**, 1553–1577 (2022).
- [46] Habara, T. & Kojima, R. Floor-Field-Guided Neural Model for Crowd Counting. *IEEE Access* (2024).
- [47] Luo, L., Xie, S., Yin, H., Peng, C. & Ong, Y.-S. Detecting and quantifying crowd-level abnormal behaviors in crowd events. *IEEE Transactions on Information Forensics and Security* (2024).
- [48] Politico. Labor’s New Front in the AI Fight. *Politico* (2025). URL <https://www.politico.com/news/2025/03/04/california-labor-federation-ai-bills-00211296>.
- [49] Ling, S., Zhao, J., Lu, Z., Wen, X. & Xiao, Z. EasyCount: Crowd Counting Based on Easy Deployment Using Commodity Wi-Fi. In *2024 IEEE Wireless Communications and Networking Conference (WCNC)*, 1–6 (IEEE, 2024).
- [50] De Sanctis, M. *et al.* RF-based device-free counting of people waiting in line: A modular approach. *IEEE Transactions on Vehicular Technology* **71**, 10471–10484 (2022).
- [51] Demrozi, F., Turetta, C., Chiarani, F., Kindt, P. H. & Pravadelli, G. Estimating indoor occupancy through low-cost BLE devices. *IEEE Sensors Journal* **21**, 17053–17063 (2021).
- [52] Liu, S., Zhao, Y., Xue, F., Chen, B. & Chen, X. DeepCount: Crowd counting with WiFi via deep learning. *arXiv preprint arXiv:1903.05316* (2019).
- [53] Denis, S., Bellekens, B., Weyn, M. & Berkvens, R. Sensing thousands of visitors using radio frequency. *IEEE Systems Journal* **15**, 5090–5093 (2020).
- [54] Kulshrestha, T., Saxena, D., Niyogi, R. & Cao, J. Real-time crowd monitoring using seamless indoor-outdoor localization. *IEEE Transactions on Mobile Computing* **19**, 664–679 (2019).
- [55] Andión, J., Navarro, J. M., López, G., Álvarez-Campana, M. & Dueñas, J. C. Smart Behavioral Analytics over a Low-Cost IoT Wi-Fi Tracking Real Deployment. *Wireless Communications and Mobile Computing* **2018**, 3136471 (2018).
- [56] Manjappa, R., Rao, V., Prasad, R. V. & Sinitsyn, A. Estimating crowd distribution using smart bulbs. In *2019 IEEE Global Communications Conference (GLOBECOM)*, 1–6 (IEEE, 2019).
- [57] Kong, H., Huang, C., Yu, J. & Shen, X. A survey of mmwave radar-based sensing in autonomous vehicles, smart homes and industry. *IEEE Communications Surveys & Tutorials* (2024).
- [58] Hsu, P., Liu, G., Fang, S.-H., Wu, H.-C. & Yan, K. Novel robust online indoor occupancy counting system using mmWave radar. *IEEE Sensors Journal* **23**, 12159–12170 (2023).
- [59] Li, S., Luo, H., Shao, T. & Hishiyama, R. A mmWave Sensor and Camera Fusion System for Indoor Occupancy Detection and Tracking. *IEEE TRANSACTIONS on Information and Systems* **107**, 1192–1205 (2024).
- [60] Ren, L., Yarovoy, A. & Fioranelli, F. Grouped People Counting Using mm-wave FMCW MIMO Radar. *IEEE Internet of Things Journal* (2023).
- [61] Hu, K., Liao, H., Li, M. & Wang, F. Mmcount: Stationary crowd counting system based on commodity millimeter-wave radar. In *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 56–60 (IEEE, 2024).
- [62] Pallaprolu, A. *et al.* Crowd Size Estimation for Non-Uniform Spatial Distributions with mmWave Radar. *arXiv preprint arXiv:2410.09200* (2024).
- [63] Liu, H., Liu, X., Xie, X., Tong, X. & Li, K. PmTrack: enabling personalized mmWave-based human tracking. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* **7**, 1–30 (2024).
- [64] Chen, Y. *et al.* mmTAI: Biometrics-assisted Multi-person Tracking with mmWave Radar. In *2024 IEEE 30th International Conference on Parallel and Distributed Systems (ICPADS)*, 26–33 (IEEE, 2024).
- [65] Shen, Z., Nunez-Yanez, J. & Dahnoun, N. Advanced millimeter-wave radar system for real-time multiple-human tracking and fall detection. *Sensors* **24**, 3660 (2024).

- [66] Zhu, J. *et al.* Joint Multitarget Detection and Tracking with mmWave Radar. *arXiv preprint arXiv:2412.17211* (2024).
- [67] Rakai, L., Song, H., Sun, S., Zhang, W. & Yang, Y. Data association in multiple object tracking: A survey of recent techniques. *Expert systems with applications* **192**, 116300 (2022).
- [68] Zhang, L., Lin, C., Zhao, Z., Wang, D. & Fu, C. A Crowd Movement Analysis Method Based on Radar Particle Flow. *Sensors* **24** (2024).
- [69] Ding, F., Pan, Z., Deng, Y., Deng, J. & Lu, C. X. Self-Supervised Scene Flow Estimation With 4-D Automotive Radar. *IEEE Robotics and Automation Letters* **7**, 8233–8240 (2022).
- [70] Ding, F., Luo, Z., Zhao, P. & Lu, C. X. milliFlow: Scene Flow Estimation on mmWave Radar Point Cloud for Human Motion Sensing. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 1–14 (2024).
- [71] Liu, M.-Y., Tuzel, O., Veeraraghavan, A. & Chellappa, R. Fast directional chamfer matching. In *2010 IEEE Computer Society conference on computer vision and pattern recognition*, 1696–1703 (IEEE, 2010).
- [72] Fruin, J. J. *Designing for pedestrians a level of service concept* (Polytechnic University, 1970).
- [73] Texas Instruments. *DCA1000EVM Quick Start Guide* ([n. d.]). Available online at <https://www.ti.com/lit/ml/spruik7/spruik7.pdf>.
- [74] US Department of Health, E. & Welfare. Skinfolts, Body Girths, Biacromial Diameter, and Selected Anthropometric Indices of Adults. Tech. Rep., National Center for Health Statistics (1970). URL https://www.cdc.gov/nchs/data/series/sr_11/sr11_035acc.pdf.
- [75] OpenCV Contributors. OpenCV: Object Tracking (2024). URL https://docs.opencv.org/3.4/dc/d6b/group_video_track.html. Accessed: 2025-02-17.
- [76] Costa, M. Interpersonal distances in group walking. *Journal of Nonverbal Behavior* **34**, 15–26 (2010).
- [77] Fruin, J. J. & Strakosch, G. R. *Pedestrian planning and design* (Elevator World, 1987).
- [78] Gérin-Lajoie, M., Richards, C. L. & McFadyen, B. J. The negotiation of stationary and moving obstructions during walking: anticipatory locomotor adaptations and preservation of personal space. *Motor control* **9**, 242–269 (2005).
- [79] Chen, Q. *et al.* UNIST: unpaired neural implicit shape translation network. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 18614–18622 (2022).
- [80] Canil, M., Pegoraro, J., Shastri, A., Casari, P. & Rossi, M. Oracle: Occlusion-resilient and self-calibrating mmwave radar network for people tracking. *IEEE Sensors Journal* (2023).
- [81] Stove, A. G. Linear FMCW radar techniques. In *IEE Proceedings F (Radar and Signal Processing)*, 5, 343–350 (IET, 1992).
- [82] Rao, S. Introduction to mmWave sensing: FMCW radars. *Texas Instruments (TI) mmWave Training Series* 1–11 (2017).
- [83] Kong, H. *et al.* m3track: mmWave-Based Multi-User 3D Posture Tracking. In *Proceedings of the 20th Annual International Conference on Mobile Systems, Applications and Services*, 491–503 (2022).
- [84] Sandeep Rao. *MIMO Radar Application Note* (2018). URL <https://www.ti.com/lit/an/swra554a/swra554a.pdf>. Accessed: 2025-02-11.
- [85] Pegoraro, J. & Rossi, M. Real-time people tracking and identification from sparse mm-wave radar point-clouds. *IEEE Access* **9**, 78504–78520 (2021).
- [86] Yan, B. *et al.* mmGesture: Semi-supervised gesture recognition system using mmWave radar. *Expert Systems with Applications* **213**, 119042 (2023).
- [87] Rohling, H. Radar CFAR thresholding in clutter and multiple target situations. *IEEE transactions on aerospace and electronic systems* 608–621 (1983).
- [88] Kronauge, M. & Rohling, H. Fast two-dimensional CFAR procedure. *IEEE Transactions on Aerospace and Electronic Systems* **49**, 1817–1823 (2013).
- [89] Hong, S. W. & Han, D. S. Performance analysis of an environmental adaptive CFAR detector. *Mathematical Problems in Engineering* **2014**, 615704 (2014).
- [90] Rohling, H. Ordered statistic CFAR technique-an overview. In *2011 12th International Radar Symposium (IRS)*, 631–638 (IEEE, 2011).
- [91] Pearce, A., Zhang, J. A., Xu, R. & Wu, K. Multi-object tracking with mmwave radar: A review. *Electronics* **12**, 308 (2023).
- [92] Lucas, B. D. & Kanade, T. An iterative image registration technique with an application to stereo vision. In *IJCAI'81: 7th international joint conference on Artificial intelligence*, vol. 2, 674–679 (1981).
- [93] Zavorka, R. *et al.* Characterizing the 80 GHz Channel in Static Scenarios: Diffuse Reflection, Scattering and Transmission through Trees under Varying Weather Conditions. *IEEE Access* (2024).
- [94] Cardamis, M. *et al.* Leafeon: Towards Accurate Sensing of Leaf Water Content for Protected Cropping, with mmWave Radar. *IEEE Internet of Things Journal* (2025).
- [95] Mardia, K. V. & Jupp, P. E. *Directional statistics* (John Wiley & Sons, 2009).
- [96] Massey Jr, F. J. The Kolmogorov-Smirnov test for goodness of fit. *Journal of the American statistical Association* **46**, 68–78 (1951).
- [97] Marsaglia, G., Tsang, W. W. & Wang, J. Evaluating Kolmogorov's distribution. *Journal of statistical software* **8**, 1–4 (2003).
- [98] Stockman, G. & Shapiro, L. G. *Computer vision* (Prentice Hall PTR, 2001).
- [99] Sun, D., Roth, S. & Black, M. J. Secrets of optical flow estimation and their principles. In *2010 IEEE computer society conference on computer vision and pattern recognition*, 2432–2439 (IEEE, 2010).
- [100] Soille, P. *et al.* *Morphological image analysis: principles and applications*, vol. 2 (Springer, 1999).
- [101] Zhang, T. Y. & Suen, C. Y. A fast parallel algorithm for thinning digital patterns. *Communications of the ACM* **27**, 236–239 (1984).
- [102] Chen, W., Sui, L., Xu, Z. & Lang, Y. Improved Zhang-Suen thinning algorithm in binary line drawing applications. In *2012 International Conference on Systems and Informatics (ICSAI2012)*, 1947–1950 (IEEE, 2012).
- [103] Toussaint, G. Grids, Connectivity and Contour Tracing (2008). URL <https://jeff.cs.mcgill.ca/~godfried/teaching/mir-reading-assignments/Chapter-2-Grids-Connectivity-Contour-Tracing.pdf>. Course notes for MIR Reading Assignments, McGill University.
- [104] Shapely Community. Shapely 2.0.6 Documentation: Buffer Reference (2024). URL <https://shapely.readthedocs.io/en/2.0.6/reference/shapely.buffer.html>. Accessed: 2025-03-07.
- [105] Apostol, T. M. *Mathematical Analysis* (Addison-Wesley, Reading, MA, 1974), 2nd edn.
- [106] Perry, A. E. & Fairlie, B. Critical points in flow patterns. In *Advances in geophysics*, vol. 18, 299–315 (Elsevier, 1975).
- [107] Sanchez, D. A. *Ordinary Differential Equations and Stability Theory: An Introduction* (Dover Publications, 2012).
- [108] Prieve, D. C. A Course in Fluid Mechanics with Vector Field Theory. *Carnegie Mellon University* **4**, 17 (2000).

Figure and Table Legends

Figure 1: Sample *structured crowds* showing two dominant flows (a) intersecting (b) splitting at the red star. Sample *diffuse crowds* showing individuals (c) fleeing from (d) converging to red star.

Figure 2: Experimental setup and environments: (Top) A mmWave MIMO radar transmits FMCW pulses, which reflect off individuals moving in (a, c) structured and (b) diffuse crowd flows (white arrows). Reflected signals are logged via Ethernet. (d) Front-facing view of the radar. (Bottom) Testing areas include open spaces with foliage (Areas 1 & 3) and a courtyard with structural constraints (Area 2), where participants move freely without predefined paths or controlled pacing, mimicking real-world pedestrian dynamics. See color PDF for best viewing.

Figure 3: Graph representations of various crowd flow topologies tested using our proposed pipeline. Graphs C1 to C12 represent spatially structured flows, whereas C13 to C15 represent spatially diffuse flows. Arrows denote the flow paths and circles denote waypoints.

Figure 4: Sample estimated flow topology, $\hat{G}(\hat{V}, \hat{E})$, obtained from our proposed pipeline, compared with the true flow topology, $G(V, E)$. Our results capture the spatial distribution, connectivity patterns, and flow characteristics of the ground truth, and demonstrate the high-fidelity reconstruction of crowd flow topologies. The estimated flow field, $\hat{\mathbf{v}}_{\text{unit}}(\mathbf{x})$, has been downsampled for clarity of presentation. See color PDF for best viewing.

Table 1: Estimated flow split ratios at vertices of flow bifurcation across configurations C1–C10. For each vertex, we report the ratio for one of the edges, λ_e , with the ratio for the other edge given by $1 - \lambda_e$. Our method achieves a Mean Absolute Error (MAE) of 0.1, highlighting its efficacy in capturing details on the flow dynamics.

Figure 5: Sample curl and divergence fields for structured (top) and diffuse (bottom) crowd flows. (a) High negative curl in R1 indicates an abrupt clockwise turn. (b) High positive (negative) divergence at R4 (R5) indicates a split (merge) in the flow. (c) Distinct curl regions correspond to clockwise and counterclockwise motion. (d) A narrow region of strong curl corresponds to the line dividing the two flow directions (orange: true, black: estimated). (e) Strong positive divergence (yellow star) aligns with the true divergence point (orange circle). (f) Strong negative divergence (yellow star) aligns with the true dispersion point (orange circle). See color PDF for best viewing.

Figure 6: Motivating the need for new flow field modeling for mmWave crowd sensing: (a) our proposed flow field model of Sec. 4.2 results in significantly better performance as compared to (b) the well-known Lucas-Kanade optical flow algorithm in vision. Both fields are downsampled for visibility.

Figure 7: Case of incomplete graph reconstruction on C4. Despite significant noise and sparse flows on the upper route,

nearly all of the flow topology is recovered, except for one edge.

Figure 8: Illustration of our proposed flow field generation pipeline. We extract flows between consecutive point clouds and filter locations with inconsistent flows over time using a Kolmogorov-Smirnov test. We then compute the average flow, apply morphological pruning to remove isolated flow regions, and perform local spatial smoothing to obtain the final estimated flow field, which serves as the starting point for the analysis presented in Sec. 4.6.

Figure 9: Illustration of our proposed robust graph reconstruction pipeline. To complement the extracted directed geometric graph, $\hat{G}$, we find edge weights, λ_e , as well as semantic details captured by curl and divergence fields.