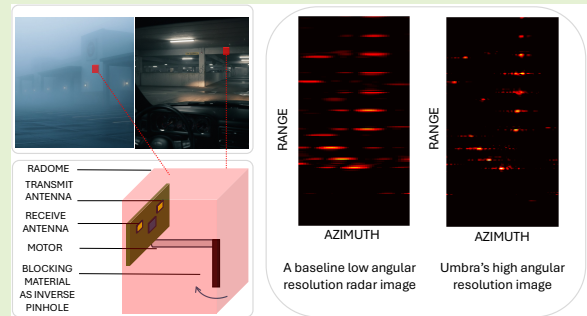


Millimeter Wave Inverse Pinhole Imaging

Akarsh Prabhakara, Yawen Liu, Aswin C. Sankaranarayanan, Anthony Rowe, Swarun Kumar

Abstract—High-resolution radio frequency (RF) imaging relies on large antenna arrays or Synthetic Aperture Radar (SAR). It is challenging to enhance aperture without significant hardware complexity and cost of arrays, or high resolution motor encoders of SAR. We introduce Umbra, a high-resolution imaging system built using just a single millimeter-wave (mmWave) and rotating “inverse pinholes” to enhance angular resolution. This paper presents the imaging system model, design, and evaluation of mmWave inverse pinholes. The inverse pinhole is attractive for its lightweight nature, which enables low-power rotation. As they rotate in the field of view, they implicitly leave signatures in RF measurements that can be used to track pinhole location to sub-wavelength precision even on cheap, low end DC motors without encoders. We finally also show how ordinary rotating objects, such as fans can act as natural mmWave inverse pinholes. One can opportunistically upgrade a single antenna system just by carefully positioning such objects in the antenna’s field of view. Our evaluation shows Umbra resolving up to 2.5° with just a single antenna, a $5\times$ improvement compared to 14° from a multi-antenna baseline.

Index Terms—radio frequency imaging, single antenna, high resolution, pinholes



I. INTRODUCTION

*“The eye is always caught by light,
but shadows have more to say.” — Gregory Maguire*

IN this paper, we explore a new method to increase the resolution of radio frequency (RF) imaging using an extremely simple transceiver chain with only a single antenna. Imaging with a single radio antenna is limited in its angular resolving ability due to its poor aperture [1]. RF imaging therefore creates large apertures with multi-antenna arrays [2], or synthesizes virtual apertures with antenna motion via Synthetic Aperture Radar (SAR) [3]–[5]. Despite these solutions being available, it is interesting to ask “Are there alternative ways to increase the angular resolution of a single antenna?”. Answering this could lead to solutions for use cases where we neither want expensive multi-channel ADC arrays nor high resolution, high update-rate motor encoders for SAR. Such alternatives could also broaden our understanding of what contributes towards a radio system’s aperture.

We propose a novel aperture synthesis method that draws on a principle commonly explored in cameras — the pinhole effect. To understand how it is related to resolution enhancement, let us move away from RF and consider a single photodiode. As a thought experiment, let us move a fully opaque sheet, with a tiny hole, in the field of view of the photo

diode. Although a single photodiode does not have any angular information, the light that passes through the tiny hole is only from a specific angular portion of the scene. The movement of the tiny hole focuses on different angular portions over time, temporally encoding spatial angular information. Pinhole and more advanced masks are used in visible light cameras for depth estimation [6], [7] and in X-rays for imaging [8], [9]. With radios traditionally operating on tens of centimeters and meters of wavelength and pinhole sizes scaling with wavelengths, the masks are unreasonably large and bulky. The trend in commodity radios moving to higher frequencies (lower wavelengths) of millimeter wave (mm λ , mmWave) presents a golden opportunity to create masks of dimensions that are practically reasonable. Thus, akin to a single pixel system, from just a single antenna we exploit this effect of modulating the wavefront with mmWave pinholes.

We present Umbra, a pinhole-based single antenna mmWave imaging system that utilizes the rotation of a custom-designed, lightweight 16 cm long strip (10 grams) rapidly rotating near the antenna, driven by a low-cost, encoder-less motor. The simplicity of this technique enables even naturally rotating objects like fans to serve as natural pinholes at mmWave frequencies. One only needs to position such objects creatively with respect to the antenna beam to get a resolution boost. In building Umbra, we come across two key challenges that make mmWave pinholes different from visible light or X-rays.

1) *Inverse pinhole design*: A pinhole has vast amounts of blocking material and a tiny transparent hole. The heavy blocking material is not suitable for moving fast. An inverse pinhole is the opposite of a pinhole, which lets all signals to

Akarsh Prabhakara is with University of Wisconsin - Madison, Madison WI USA 53706. Email: akarsh@cs.wisc.edu. Yawen Liu, Aswin C. Sankaranarayanan, Anthony Rowe, and Swarun Kumar are with Carnegie Mellon University, Pittsburgh PA USA 15213. Email: {yawenl, saswin, agr, swarunk}@andrew.cmu.edu

go through except a tiny blocker. When the blocker's shadow aligns with the object's angle (Fig 1d), we see a dip in energy, as opposed to a peak. The temporal occurrence of the dip encodes angle information. Because of lesser blocking material, the inverse pinhole is an ideal, lightweight option while still having the same resolution as the pinhole. Visible light sensors are limited by shot noise [10]–[12], where noise increases with signal amplitude. Inverse pinholes are flooded with a higher amplitude signal than regular pinholes. As a result, the inverse pinhole, by and large, is not seen in visible light systems except perhaps in “accidental” settings [13]. mmWave not only presents a chance to build reasonably sized RF pinholes, but also presents a different sensor noise regime with signal independent thermal noise [14]. In this new regime, inverse pinhole is feasible and makes the mask extremely lightweight.

2) Bidirectional pass-through system model: Unlike X-ray or vision, typical packaging at mmWave positions the transmit and receive antennas close together. The wave inevitably goes through the inverse pinhole twice — starting from the transmitter, through the pinhole, backscattering from an object, through the pinhole again, to the receiver. We first show how this bidirectional pass-through is a positive phenomenon that improves resolution. Second, our inverse pinhole being near the antenna and practical mask element dimensions being comparable to the wavelength, we take into consideration propagation models that account for diffraction as opposed to a simple shadow approximation that ray optics in a pinhole camera suggests. We then present robust reconstruction algorithms to recover the image under RF noise. We finally also present material choices to practically realize effective inverse pinholes, and also account for all non-idealities via calibration.

Umbra achieves an angular resolution of 2.5° , a $5\times$ improvement over a standard mmWave radar with 8 antennas [15]. We built our system on top of a standard mmWave radar, using only 1 antenna, and on real world scenes up to 20 meters in range. This resolution is on par with an array of 45 antennas. Achieving this with 1 antenna and an RF passive inverse pinhole drastically reduces the need for high RF complexity systems.

Given that Umbra uses motion to synthesize a larger aperture, although in a very unique way, it is interesting to compare it with SAR. We contrast with SAR both from the point of view of resolution and system design. The effect of different diversity enhancing mechanisms on resolution can be analyzed by examining properties of the linear matrix modeling the known measurements and the unknown image. For a fixed size of antenna motion for SAR and inverse pinhole motion for Umbra, we find that the magnitude of dominant singular values of such a matrix are higher for SAR than for Umbra (Sec IV-C). This implies that image reconstruction via SAR can leverage a larger number of singular values and yield superior resolution. Although Umbra does not mimic SAR perfectly, we empirically find that rotating inverse pinholes still have several dominant singular values above the practical noise floor. Our robust algorithms exploit these useful singular values and enhance resolution without severe noise amplification. Despite

resolution improvement being only modest compared to SAR, what uniquely sets us apart is a practical system design advantage. A significant barrier to implementing SAR is the need for high update rate motor encoders to track antenna motion. Most SAR systems relax update rates by opting for a slow scan (lasting several minutes) with extremely long image acquisition times (lasting several minutes) [16]. Umbra on the other hand just needs low cost DC motors with a scan time on the scale of 100 ms. We demonstrate our fast scan benefits by evaluating on dynamic objects. This is possible because of Umbra's arrangement design: close proximity of inverse pinholes to antenna. This allows us to build software techniques to estimate the motor position (and inverse pinhole) purely by analyzing RF backscatter variations in nearby range bins from the rotating inverse pinhole (Sec IV-D). Thus, while SAR remains a superior motion based aperture synthesis technique from the resolution point of view, practical advantages such as sub-wavelength tracking at kilohertz update rate with low cost motors and fast scan times make Umbra a practically interesting alternate aperture synthesis method.

Despite the ease of upgrading a single antenna system, Umbra is still a mechanical solution and requires 32 cm x 32 cm for inverse pinhole rotation. It is best suited for upgrading static radar deployments with wide clearing, such as around-the-corner imaging in parking lots via wall-mount radars [16], pole-mounted surveillance in vision-denied environments such as through-fog imaging in building premises or farms [17], [18], and in low-light occupancy sensing [19]. We demonstrate this with static radar evaluations. It is also valuable in environments where objects like fans already exist and only need to be re-positioned for the antenna beam to pass-through.

Umbra makes four key contributions.

- A design based on mmWave inverse pinholes to enhance angular resolution with lightweight rotating loads.
- A system model dealing with the mmWave setup of bidirectional pass-through, diffraction effects, spherical wavefronts, antenna placement, and radiation pattern.
- mmWave inverse pinhole realization with material choice that just weighs 10 grams and radar-motor self-synchronization technique that can eliminate the need for expensive motor encoders.
- A system evaluation on objects upto 20 meters away with 5x improvement over a multi-antenna baseline.

Scope: Umbra proposes novel RF sensor design and algorithms for improving angular resolution from a single antenna. Namely: (1) design and implementation of an inverse pinhole in close proximity to an antenna, (2) inverse pinhole processing algorithm and evaluation on real-world scenes.

II. BACKGROUND

RF imaging systems such as radars use active illumination to obtain object-specific information such as range, azimuth, elevation, and relative velocity [20]. In the simplest setup (Fig 1a), a colocated transmit (Tx) and receive (Rx) antenna transmits a wave, which reflects off a target and returns. The spherical outgoing wave decays as $1/d$, changes phase as $\frac{2\pi d}{\lambda}$, where d is the one-way range; and λ is mmWave wavelength

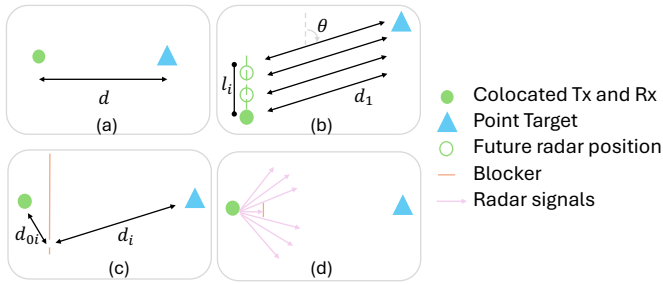


Fig. 1. (a) shows a simple case of antenna and target; (b) shows a synthetic aperture radar case of a moving antenna; (c) shows a pinhole and (d) inverse pinhole imaging case with a static antenna

(we use 77 GHz / 4 mm). The one-way path from Tx to the target is $\frac{1}{d} \exp \frac{j2\pi d}{\lambda}$ where $j = \sqrt{-1}$. The target spherically re-radiates back to the radar. Thus, the measurements received at the radar are of the form $\frac{1}{d^2} \exp \frac{j2\pi 2d}{\lambda}$ [21]. Imaging is effectively locating the target unambiguously in 3D (shown as 2D here). Clearly, a single antenna only has distance information, but no angular information.

Modern compact mmWave radars [15] have multiple non-colocated (8-12) antennas placed on a linear array to get angles. Another approach is to move and measure a single Tx/Rx antenna along a known trajectory (Fig 1b). This is the Synthetic Aperture Radar (SAR) processing [21]. With both approaches, we have $\frac{1}{d_i^2} \exp \frac{j2\pi 2d_i}{\lambda}$, where d_i is the instantaneous/antenna-specific distance. The same target is seen from different (but known) antenna points of view. For a far target (plane wave assumed), the phase of these measurements is $\frac{2\pi 2(d_1 + l_i \cos \theta)}{\lambda}$, where d_1 is the reference distance to the first antenna, l_i is the length from antenna 1 to i and θ is the target azimuth angle for which we want to solve. With known antenna positions, we can solve for θ . The classical bound on the angular resolution is determined by the trajectory length $\lambda / (2 * \text{trajectory length})$. Longer trajectories or array length (for a real antenna array) better is the resolution. The compactness of these commodity radars makes the default resolution poor (14.3°). SAR requires sub-millimeter motion tracking with precise, high update rate, and thus, expensive motor encoders. Thus, we are interested in new methods for synthesizing virtual apertures.

III. RELATED WORK

A. mmWave radar super-resolution

Rotating an antenna can only provide a resolution of its field of view (FoV). For multi-antenna radars, resolution improves as $\propto \lambda / \text{Array Size}$. The TI single-chip radar [15] has 8 antennas (1.5 cm form factor) with a resolution of 14.3° . This results in coarse and blobby images. [2] cascades 4 of these chips to obtain 86 antennas (11 cm form factor) and a resolution of 1.3° . Other single-chip solutions achieve 2° resolution with 48 antennas [22]. Arrays of tens of active antennas come at a significantly high cost ([2] costs \$3500, whereas single antenna radars are $< \$100$). Software techniques to improve resolution on a cheaper, more compact radar [15] are

preferred. (1) Synthetic Aperture Radar: Radar motion along specific trajectories [3], [5], [23], [24]; (2) Sensor Fusion with camera/lidar [25]; (3) Machine Learning / Sparse Processing Super-Resolution [17], [26]–[29]. In contrast, we seek to create dense images with just radar alone, without any training datasets or sparse target assumptions. Rather than moving a radar with precise tracking like in SAR, Umbra moves a lighter inverse pinhole and derives the tracking implicitly without any expensive motors.

B. Pinhole and coded masks in other wavelengths

Pinholes and masks in X-ray [8], [9] are used to estimate angles. In visible light literature, they are used for depth imaging, high speed video, light field imaging, and lensless imaging [6], [30], [31]. [13] explores “accidental” inverse pinholes. For these sub-nanometer or nanometer wavelengths and pinholes of millimeter or centimeter dimensions, systems typically consider ray optics. Operating at millimeter wavelengths, we consider diffraction and also a bidirectional pass.

C. Coded illumination in RF

In radio frequency (RF) literature, a large antenna array is used to generate coded radiation patterns [32], [33], [34], [35] perform unidirectional coded illumination (transmit or receive only). [36]–[38] are bidirectional. However, they use complex multi-hole codes that require active RF elements or moving metal sheets with holes. [39] designs custom metasurfaces for each object being imaged. [40] builds metasurfaces disjoint from the radar to extend coverage. We find that (1) bidirectionality is unique vs. optical and X-ray, (2) RF works [36]–[39] lack an exposition on fundamental effects of bidirectionality, and (3) we can build simple, inverse pinholes that are *both RF passive and lightweight* to enhance resolution.

IV. UMBRA SYSTEM DESIGN

Umbra is motivated by a practical design point with a lightweight inverse pinhole near a single transmit and receive antenna rotating at high speeds. This enables high resolution imaging with a low power motor. We choose rotation over linear actuation (that goes back and forth) to avoid losing momentum, minimizing acquisition time. This section first intuitively shows how we synthesize virtual apertures with pinhole motion, then presents an imaging system model accounting for mmWave bidirectional pass-through, then introduces the inverse design which makes the system lightweight, and finally discusses other practical system challenges. As we build it up, we contrast our aperture synthesis against SAR.

A. Pinhole motion to synthesize virtual aperture

Our idea is to borrow and adapt the pinhole effect from vision or X-ray. Consider a pinhole sheet that is scanned in front of the radar (Fig 1c), with an infinitesimally small pinhole (ignoring diffraction). This blocks all radar signals except where the pinhole opens. Effectively, this creates a virtual radar source at the pinhole location. Moving the pinhole sheet can be viewed as mimicking the motion of the radar that




Fig. 2. Shows the pinhole system model with nearby Transmit (Tx) and Receive (Rx) antennas, a pinhole mask and a target reflector.

creates different points of view of the antenna. At any scan position, for a known d_{0i} dependent on the location of the pinhole and the antenna, from $\approx \frac{1}{d_{0i}^2 d_i^2} \exp \frac{j2\pi 2(d_{0i} + d_i)}{\lambda}$, we can compensate for d_{0i} . The rest of the terms look similar to the SAR setup seen in Sec II. Can we leverage this to do mmWave imaging with only a single antenna? Effective pinhole sizes increase on the wavelength scale. At lower RF frequencies, a tens of centimeter or meter scale pinhole system would be impractical to scan. We leverage smaller mmWave wavelengths and move modest size (millimeter or centimeter scale) pinholes.

B. mmWave Pinhole System Model

This subsection builds an accurate imaging system model and simulator for mmWave pinholes to rapidly prototype pinhole designs and reconstruction algorithms. We want these unique attributes at mmWave to be captured:

- *Bidirectional pass-through:* Due to closely separated Tx/Rx antennas (1 cm), a pinhole mask placed in the field of view of either of the antennas inevitably experiences both outgoing and incoming wavefronts.
- *Wavelength-scale pinhole width:* A finite width pinhole of a modest size (a few mm) would be comparable to the wavelength. Rather than simply ray tracing, we need to account for diffraction effects in the shadow region.
- *Spherical wavefronts:* We envision positioning the pinhole mask a few centimeters away. At this proximity, the wavefront is spherical when it impinges on the mask.

System Model: To achieve these requirements, we build a system model as follows (Fig 2):

- Wave propagation from the transmit antenna (Tx) to the mask plane as c_n . This is of the form $\frac{1}{d_n} \exp \frac{j2\pi d_n}{\lambda}$, where n indexes the discrete position on the mask plane. This takes into account spherical wave radiation.
- Modulation with the pinhole mask Ac , where A is a diagonal matrix indicating the effect of each element in a mask. In an ideal case, “1” indicates pass-through, “0” indicates a block. We take into account the finite pinhole width with a group of matrix elements marked as “1” simultaneously. The pinhole could also induce phase modulations which will be estimated and calibrated for in practice as described in Sec IV-D.
- Wave propagation ZAc to obtain the electric field at the target position. Z captures the effects of radiating from all discrete positions on the mask plane to the target.
- Each target reflects with a certain reflectivity (that we have to solve as the radar image of the scene), resulting in $XZAc$, X being a diagonal matrix capturing multiple target reflectivities.

- Next, the reflected wave goes through similar transformations, but in reverse - first wave propagation from the target to the mask plane, a second pass of the pinhole mask, and finally wave propagation to the receive antenna (Rx). This is akin to a point source at each target position re-radiating back to the radar. $\tilde{c}^T A^T Z^T X Z A c$ is a scalar measurement for each stationary position of the mask. \tilde{c} accounts for the backscattered wave reaching the Rx antenna, which is near the Tx but physically separated. We note that the effects of A are identical for both directions because the time interval between the outgoing and incoming waves is at nanosecond time scales, whereas the mask only moves at microsecond scales.
- When the pinhole sheet is moved, only the elements of A shift, with everything else being the same. Thus, we have measurements $y_k = \tilde{q}_k^T X q_k$, where $q_k = Z A_k c$ and $\tilde{q}_k = Z A_k \tilde{c}$, for k indexing the pinhole sheet motion.
- Given X is diagonal, $y_k = (\tilde{q}_k * q_k)^T x$, where x is the target reflectivities, $\text{diag}(X)$
- Thus, we have $y = (\tilde{Q} * Q)^T x = Bx$, where each column in \tilde{Q} and Q contain different mask position measurements \tilde{q}_k and q_k respectively, a linear measurement model.

Alternate expression for Q: We can further simplify this model to decouple the effects of wave propagation and the mask. q_k is a column in Q that encapsulates a one way pass through the pinhole mask $q_k = Z A_k c$. Z is a $M \times N$ matrix that captures the propagation from N points on the pinhole plane to M points on the target plane, and c is a $N \times 1$ vector that captures the electric field at N points on the pinhole plane. Let us ignore the mask for a moment and define a single term F that captures both these terms, $F^T = Z * \text{repmat}(c^T, [M, 1])$. With this, the instantaneous effect of the mask can be written as $q_k = F^T a_k$, where $a_k = \text{diag}(A_k)$. As the mask moves, we measure the effect of multiple a_k . Putting together all a_k , we define a new term H , that captures a_k along its rows. This gives us $Q = F^T H^T$, where Q is $M \times K$, F^T is $M \times N$, H^T is $N \times K$, K is the total number of scan positions or measurements. Thus, for a bidirectional pass, we have:

$$y = (\tilde{Q} * Q)^T x = (\tilde{F}^T H^T * F^T H^T)^T x = (H \tilde{F} * H F) x = Bx$$

Rayleigh-Sommerfield wave propagation: While the matrix H , the effects induced by the mask, is relatively straightforward to understand and create in a simulator, F and \tilde{F} are complicated propagation related terms. The first step from Tx to mask plane is a simple spherical model. But how should we account for other propagation terms analytically? A naive ray tracing beyond the mask plane would strictly only illuminate a discrete zone in the scene. To build a realistic model, diffraction induced by wavelength-scale pinhole must be considered. This captures the critical contribution from targets outside the pinhole’s direct ray-traced illumination. To achieve this, the next wave propagation step in the system model is performed using the Rayleigh-Sommerfield model [41]. We choose this because it is more comprehensive than diffraction models with Fresnel or Fraunhofer approximations. The forward propagation term is

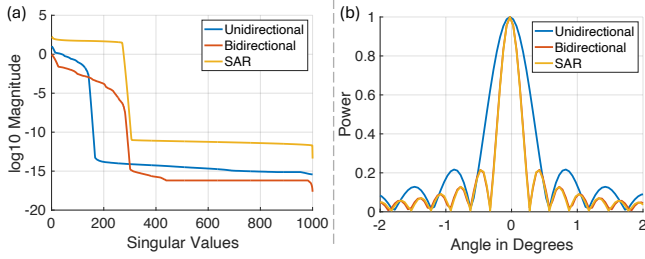


Fig. 3. Shows the positive impact, on (a) singular values (b) angular resolution, of bidirectional pass in contrast to unidirectional pass.

$$F(m, n) \propto \frac{c_n}{d_n d_{mn}} \exp \frac{j2\pi(d_n + d_{mn})}{\lambda} \cos\theta \left(1 + \frac{j\lambda}{2\pi d_{mn}}\right)$$

where

c_n : is the spherical field incident on the mask plane
 d_n : distance from the antenna to n^{th} index in the mask plane
 d_{mn} : distance from n index in the mask plane to the m^{th} target
 λ : the center wavelength at 77 GHz, 3.89 mm
 θ : the angle between the wave propagation vector at n and the position vector between m and n . The summation of multiple $F(m, n)$ with a weight as dictated by the mask H and using sub-millimeter grid sizes in the simulator thoroughly models diffraction. The reverse pass expression, \tilde{F} , is also derived by considering spherical re-radiation from the target to the mask plane and Rayleigh-Sommerfeld diffraction is used from the mask plane to the receiving antenna.

Analysis of bidirectional pass-through: In contrast to a one-way pass, the key difference is in the overall measurement model:

$$y = (H\tilde{F} * HF)x + n = Bx + n,$$

where y is the measurements, H encodes the time-variant blocking from a pinhole, F and \tilde{F} are the one-way propagation matrix for Tx to target and target to Rx, x is the reflectivity of targets in the scene that we want to recover, n is additive noise, $*$ is element wise multiplication.

A unidirectional pass-through only captures $H\tilde{F}x$. For bidirectional pass, let us simplify and assume colocated Tx and Rx ($\tilde{F} = F$). Let each element in HF be of the form of $\approx \frac{1}{d} \exp \frac{j2\pi d}{\lambda}$, where d is an equivalent distance term that captures one direction from antenna to target. Then, each element in the bidirectional pass is $\approx \frac{1}{d^2} \exp \frac{j2\pi 2d}{\lambda}$. To illustrate the significance of this, we draw an analogy from classical antenna array and SAR without any pinholes.

Consider an antenna array with one static Tx and a static linear receiver array. We have $\frac{1}{d_0 d_i} \exp \frac{j2\pi(d_0 + d_i)}{\lambda}$, where d_0 is the constant distance from Tx to the target, d_i is the distance from each array element to the target. Compensating for constants and looking at the relative phases between the array elements, we end up with the form $\propto \frac{1}{d_i} \exp \frac{j2\pi d_i}{\lambda}$. Next, consider a SAR approach that moves both a colocated Tx and Rx resulting in $\frac{1}{d_i^2} \exp \frac{j2\pi 2d_i}{\lambda}$, where d_i is the instantaneous distance from Tx/Rx to the target. Observing these two, we can draw equivalences between (1) unidirectional pinhole system \equiv single Tx & multiple Rx array (2) bidirectional pinhole system \equiv single colocated Tx/Rx SAR. The $2d$ term is the main

difference between these two forms, and this implies that phase shifts are twice as sensitive with a bidirectional pass / SAR. In other words, SAR of length $L/2$ achieves identical phase shifts from a certain target as a receive-only array of length L , and thereby the same resolution performance. Similar to this analogy, bidirectional pass-through should be more effective in enhancing the angular resolution than unidirectional. Thus, Tx and Rx sharing the pinhole mask turns out to be a positive phenomenon.

We test out our analysis in simulation and compare the singular values of the matrices $H\tilde{F}$ and $H\tilde{F} * HF$ (Fig 3a). The singular value curve stretching as far to the right at high magnitudes is ideal as it points to greater diversity in measurements and better resolution. We find that the singular values drop off in significance much slower in the bidirectional case, thus resulting in more diversity. This leads to a $\approx 2x$ improvement in angular resolution (full width at half power) as shown in Fig 3b that matches our analogy explanation above.

Take Away 1: Bidirectional passthrough may seem harder to model vis-a-vis unidirectional passthrough but is actually beneficial and doubles angular resolution.

Comparison against SAR: Although the measurement forms appear to be similar between the bidirectional pinhole and SAR, we characterize our azimuth resolution against circular SAR. In this, the antenna moves along the circumference of the same circle (dimensions same, since this dictates resolution), that the pinhole motion creates. Using *all* singular values, we achieve a resolution on par with SAR (Fig 3b).

Take Away 2: In the absence of noise, pinhole motion achieves resolution on par with SAR.

Robust reconstruction algorithm: Ideally, simply inverting the matrix B should reconstruct x . But in high noise, B 's pseudo inverse will amplify noise, and reconstruction will fail. Depending on the radar configurations, we may be in a high noise setting. To avoid this, we resort to truncated SVD and then pseudo inverting B . Truncated SVD restricts the singular values to only dominant ones and reduces radar noise amplification. Restricting the maximum number of singular values (σ_{MAX}) adversely affects system resolution. The singular value decay pattern as seen in Fig 3a is much sharper than the flat curve seen with SAR. Thus, while SAR gives good resolution even in high noise settings, restricting the number of singular values loses resolution for pinholes. Fortunately, as seen in our evaluations, in practice, we are in a noise regime where several singular values (upto 40) can be used. This makes even simple pinhole designs very effective.

Take Away 3: Pinholes are unfortunately less resilient to noise vis-a-vis SAR, nevertheless offers significant resolution improvements against compact multi-antenna baselines.

Note on Coherent Processing: At first glance, an inverse pinhole based reconstruction seems to only consider amplitude




Fig. 4. (a) shows pinhole setup. (b) & (c) contrasts singular values by varying pinhole width and pinhole's location.

variations of the blocker. This might give the impression that Umbra runs an incoherent processing algorithm by design. Although we used a convenient “1” or “0” modulation by the pinhole in deriving our model, the model itself can support both amplitude and phase variations due to the blocker. In fact, in practice, real blockers lead to an amplitude change (shadowing and ringing in Fig 6a) and a phase change (the orange line in Fig 12a). An amplitude-only algorithm would simply ignore these valuable phase changes and yield inferior angular resolution improvements. Thus, Umbra indeed runs a coherent processing algorithm. Sec IV-D shows how we calibrate and measure the real world B that contains both amplitude and phase variations. Our evaluation also shows the sensitivity of our system to phase changes due to undesirable motion (vibration, object motion etc.).

C. Pinhole Design

We now know what a pinhole mask does to a single antenna to enhance resolution. This subsection uses the realistic simulator for iterating over the pinhole design.

Pinhole Geometry: The simplest structure is literally a hole in the FoV. An infinite and zero-width pinhole either lets all or no waves through, lacking any angle-resolving capability. As we vary hole width (Fig 4b), we find that as size increases (1) a flat (SAR-like) curve becomes less flat (2) more energy passes through and singular values' magnitude increases. For noise resilience, we strike a balance between curve flatness decrease and magnitude increase, choosing pinhole width of 4λ for our simulations and evaluations.

The pinhole center distance from the motor axis (Z) decides the size of the circle that the motion creates. From bidirectional pinhole and SAR analogy, larger circle, better resolution. Thus, the larger the pinhole center's distance, the better the resolution. We verify in simulation (Fig 4c) and choose a radius of 16 cm for our evaluations with a SAR resolution of 0.35° balancing size practicality and resolution.

The depth between the pinhole plane and the antenna plane together with the radiation pattern determine the effective field of view. That is, $\tan^{-1} \frac{16\text{cm}}{\text{Depth}}$. Outside this FoV, the pinhole's illumination will not directly fall, only weakly coupling the spatial location. To achieve a modest FoV of $\pm 53^\circ$, we use simulation and determine 12 cm as the offset.

Inverse Pinhole: The pinhole is built from blocker material. Emulating a tiny 4λ hole requires a lot of material weight. Inverse pinhole is a concept that lets all waves go through except where the pinhole is. For this, we only need a blocker at the hole location. This drastically reduces the weight. It can be held to the rotating motor with a lightweight support.




Fig. 5. (a) Shows noise in mmWave radar isn't signal dependent. (b) & (c) contrasts the singular values and angular resolution of regular pinhole and inverse pinhole respectively.

This can be as simple as a single-blade propeller made of mmWave-transparent material and a blocker at the pinhole location. Remember, mmWave degrades little with thin, cheap, light materials such as plastic. But first, why are they not as popular as regular pinholes?

Visible light systems are dominated by signal dependent shot noise [10]–[12]. The inverse pinhole lets all light through except for a small portion. This floods the visible light detector and enhances the noise dramatically. However, radio systems are dominated instead by signal independent thermal noise [14], [42], [43]. We verify the overall noise in our system from measurements. We experiment by moving a highly reflective target to different ranges from the radar. A closer target reflects more energy, creating a higher amplitude signal at the receiver. However, irrespective of the target range (signal power), the noise power remains low and orders of magnitude lesser than signal power (Fig 5a). The few outliers (4%) that exist are due to signal dependent non-idealities of the radar that only show up when targets are extremely close and the signal power is high. This indicates that we are dominated by thermal noise.

Take Away 4: Inverse pinholes are attractive at mmWave due to signal independent noise regime.

Regular vs Inverse Pinhole: Moving to inverse pinholes only causes the 0s and 1s in H to interchange. For a unidirectional pass, HF and $(1-H)F$ are the models for regular and inverse pinholes. An inverse pinhole is equivalent to a regular pinhole upon background subtraction (F is the background measurement without any pinhole). Thus, the resolution abilities of both are similar. For a bidirectional pass (assuming colocated Tx/Rx), we have inverse pinhole as: $(1-H)F * (1-H)F$, where, $OF * OF$ is the background, $HF * HF$ is regular pinhole both ways, $OF * HF$ is regular pinhole for outgoing but open for incoming wave, $HF * OF$ is vice versa, and O is a matrix full of 1s. To understand the impact of these new terms, we use our simulator. We see that regular and inverse pinholes are not identical (Fig 5b & c) in a bidirectional setup.

First of all, unlike SAR, we can see that both the pinhole and the inverse pinhole have problems with noise (due to the dropping singular value curve). However, if you pick a noise floor (for example: $-5 \log_{10}$ magnitude), horizontal line cutting through Fig 5b, we find that it intersects the pinhole curve first and only later the inverse pinhole curve. This is because the singular value curve for the inverse pinhole lies strictly above the regular pinhole. The corresponding singular

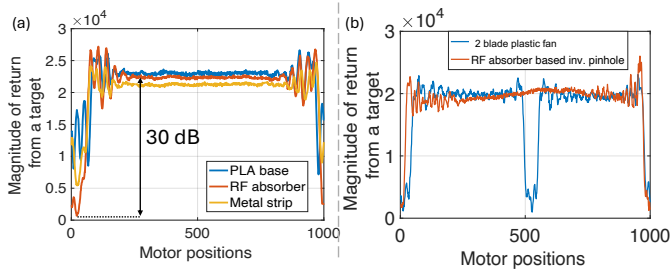


Fig. 6. Shows return from a point target as an inverse pinhole completes 1 rotation (1000 positions). (a) Inverse pinhole material design choices. RF Absorber on a PLA base gives 30 dB attenuation (shadow). (b) Natural inverse pinhole effect from a 2-blade fan yields two nulls.

values, σ_{MAX} , at these intersections tell us how many singular values are above noise and are useful for reconstruction. Thus, given a noise setting, the inverse pinhole can leverage a larger number of singular values above the noise magnitude and offer better resolution.

Take Away 5: For a bidirectional case, inverse pinholes are lighter and offer better performance compared to pinholes.

D. Practical System Challenges

Blocker Material Selection: So far, we have considered an inverse pinhole to be literally a circular hole at a radius of 16 cm with a 4λ width. In practice, we need to attach it to a rotating motor axle with a thin single-blade propeller-like supporting structure. Rather than simply attaching the blocker inverse pinhole at the tip of such a structure, a natural design is to fill the entire support with blocker material. Note that this is only along the radial direction, and is still much less material than emulating a regular pinhole. We build such inverse *pinstrips* and use them for our evaluations. But, we need them to be effective at blocking. This comes down to the selection of the blocker material. A PLA plastic material causes a 9dB power attenuation (shadow) from a point target (see Fig 6a). On top of a PLA substrate, we experiment by adding (1) absorbing material, or (2) reflector (such as a metal strip) that is facing towards the radar. The RF absorber creates the sharpest decrease in power (30dB) between the two. The metal strip only creates a 12dB power attenuation. We select the absorber on PLA as our design. Fig 6a shows that in addition to the deep nulls as expected for shadows, there is ringing throughout that is typical of setups where diffraction is prevalent. Our system model accounts for this by baking in the geometry setup, inverse pinstrip width and spherical illumination – all in matrix B .

Non-idealities and Calibration: Our simulator and model are quite comprehensive. It is also important to consider some missing points. (1) Sensitivity of the antenna in different directions. For each element $F(m, n)$, the antenna sensitivity derived from the position of the mask grid element and the target relative to the antenna is estimated and multiplied. Given the proximity of the pinhole to the antennas and the sufficiently wide beamwidth of $\pm 50^\circ$ of the patch antennas on [15], the entire pinhole motion can be used effectively. (2) A non-binary matrix H . In Sec IV-B, we introduced an ideal binary matrix

H . But, to allow for each practical implementation’s nuances: degree of blockage, phase change induced by the strip, exact strip shape, etc. a one time calibration is performed with a point target, thin metal rods. A point target is swept angularly using a motion stage and recorded measurements are used as practically measured B for reconstruction.

High Speed Motor Tracking: Just like antenna motion tracking in SAR, a pre-requisite for us is tightly synchronizing motor positions and radar measurements. Motor encoders are a natural solution, but a high rate encoder to keep up with radar’s 10 kHz drives up the motor’s cost. Our solution is to use radar’s high resolution (4 cm) ranging ability itself for self-synchronization. Placing the pinhole mask near the radar affects the first few range bins (bin number 6 in our implementation). The rotation creates a unique signature at these bins that periodically repeats and is scene-independent. If the motor rotates at non-uniform rates, this shows up in the signature. We continuously use Dynamic Time Warping [44] on this signature to align multiple rotations. This eliminates the need for expensive motor encoders that SAR relies on.

Natural Inverse Pinhole Effect: Fig 6a already shows that even a base plastic creates a shadow (inverse pinhole) effect. The biggest difference between an ordinary plastic object like a fan and a custom-made inverse pinhole is the number of blades. Multi-blades create multiple nulls in one rotation (2 blade fan shown in Fig 6b). The null location directly indicates the target angle. But, if there are two nulls for a single target, this results in ambiguity (or aliasing). We can overcome this ambiguity by ensuring that the beam’s spotlight passes through only one blade at any time. We can achieve this by building antennas with narrow elevation FoV (such as [15]) and offset from the center (such as Fig 4a). Extending this to 4 or 6 blades involves shrinking the azimuth spotlight, giving a reduced system FoV, a smaller “useful” circle that is scanned, and thus poorer resolution. With the elevation offset design, two blade fans are the best suited and enjoy the natural inverse pinhole effect with the same FoV and resolution as a single blade.

V. IMPLEMENTATION

Umbra uses AWR1843BOOST [15], single-chip mmWave radar at 77-81 GHz that has 3 Tx and 4 Rx antennas. We just use a *single Tx/Rx antenna pair*. This radar board is attached to DCA1000EVM for raw I/Q. We mount a cheap and low-end motor without any motor encoder reading. The radar transmits a Frequency Modulated Continuous Wave (FMCW) chirp at 10 kHz rate.

Arguments in Sec IV-C and Sec IV-D guide our hardware design. We place the inverse pinhole plane at a depth of 12 cm (for desired azimuth FoV) from the antenna plane, the antenna is 12 cm vertically offset from the rotation axis, and the motor rotates at 600 rpm. Inverse pinhole is implemented as an inverse pinstrip of length 16 cm and width 4λ (chosen from simulations for a desirable singular value curve). Strips with a length other than 16 cm can also be built. The length sets the size of the scanned circle and the maximum achievable resolution. A blocking structure that uses layers of absorbing material [45], with a thin 3D printed PLA structure




Fig. 7. Umbra moves a strip of radio absorbing material that just weighs 10 grams in front of a single antenna from a mmWave radar.

for support (Fig 7) is used. The thickness of the PLA was chosen to bare the weight of the absorber. The thickness of the absorber was default from the absorber manufacturer. This material was chosen following absorption analysis experiments mentioned in Fig 6a. The RF absorber demonstrated upto 30dB attenuation in received strength from a point target — mimicking a deep shadow. We calibrate for the amount of blockage and phase change due to blocker, by measuring the response for a point target swept angularly. We measure the practical matrix B empirically for our implementation. Once B is obtained, the reconstruction algorithm follows TSVD as mentioned earlier.

The total weight of our inverse pinhole (that is, the radial pinstrip) is 10 grams while the radar board is 120 grams. To avoid exposing the spinning pinhole, for aesthetics or harsh environments, we can also build a radome. In order to demonstrate the natural inverse pinhole effect, we use a plastic fan with 2 blades — diametrically opposite to each other. We use a drone propeller for this. The propeller dimensions are 16 cm long and 1.5 cm wide – similar to (4λ) our custom designed inverse pinhole.

VI. EVALUATION

Ground Truth and Baselines: We build a rig with a co-located multi-chip and large form factor 1.3° azimuth resolution radar as ground truth [2]. This is same as performing SAR over 8.5 cm trajectory length. Our baseline uses all 8 antennas on [15]. The co-location ensures overlap between Ground Truth (GT) radar and Umbra. All radars use Frequency Modulated Continuous Wave. Given wide ($\pm 50^\circ$) FoV in azimuth, smaller elevation ($\pm 20^\circ$) FoV and poor elevation resolution (18°) in the GT, we mostly focus on range-azimuth (2D overhead view) imaging. Intensity indicates object reflectivities. We fix a maximum range of 20 m. This can be changed for long range imaging.

Metrics: To compare resolution, we use a high frequency content estimator as a “sharpness metric”. Sharper the image, higher the energy in high frequency bands along the azimuth direction. For an image I with $D(\text{range}_{\text{freq}}, \text{azim}_{\text{freq}})$ as the 2D-FFT output, sharpness is

Method	Sharpness	MSE	SSIM	CD
Umbra	0.80	13.11	0.85	0.19
Multi-antenna baseline [15]	0.64	21.20	0.82	0.25
Umbra with $\sigma_{\text{MAX}} = 10$	0.70	18.69	0.83	0.20
Umbra with $\sigma_{\text{MAX}} = 15$	0.73	17.48	0.84	0.20
Umbra with $\sigma_{\text{MAX}} = 20$	0.73	18.61	0.82	0.21
Umbra with $\sigma_{\text{MAX}} = 30$	0.77	15.63	0.84	0.21
Umbra with $\sigma_{\text{MAX}} = 40$	0.80	13.11	0.85	0.19
Umbra with $\sigma_{\text{MAX}} = 60$	0.75	18.69	0.82	0.31
Umbra with $\sigma_{\text{MAX}} = 80$	0.60	30.12	0.72	0.37

TABLE I

WE EVALUATE ON “SHARPNESS” (HIGHER THE BETTER), MSE (LOWER THE BETTER), SSIM (CLOSER TO 1 IS BETTER), CD (IN METERS).

$S(D(I)) = \sum |D(0, \text{azim}_{\text{freq}} \geq 0.1)|$. The baseline’s sharpness is finally defined as $\frac{S(D(\text{baseline image}))}{S(D(\text{ground truth}))}$. Umbra’s sharpness is defined as $\frac{S(D(\text{Umbra’s image}))}{S(D(\text{ground truth}))}$. Higher the value, closer is the sharpness to the ground truth. For imaging accuracy, we use a pixel-wise mean squared error (MSE) on wide range of intensities in $[0.01, 1]$ that covers objects 40dB weaker in power. We also threshold images to convert to 2D point cloud and measure Chamfer Distance. Structural Similarity Index Measure (SSIM) shows perceptual quality.

Resolution Tests: For azimuth resolution tests, we use two point targets (6mm diameter metal rods). We move them on a precision stage from 17.3° until the two targets appear as one. Fig 8a shows that we can resolve targets up to 2.5° . For each measurement, the estimated angles align with the ground truth from the precision stage. Moving closer, at 1.5° , we are unable to resolve the two targets. From our simulation, the resolution limit is 2.2° for our measured radar noise levels. This matches well with our measurements in reality.

Accuracy Tests in Natural Scenes: We evaluate in office spaces, lobbies, and garages (open and cluttered spaces). Fig 8b and 10 shows our superior imaging quality. It is easy to identify objects like metal posts and pillars. In the second scene in Fig 8b, the wall appears discretized, as the studs behind the wall are strong reflectors. We used a stud finder to verify that the discretized points are indeed reflections from studs. We see the high quality reflected in all quantitative metrics, outperforming the baseline (see Table I). Higher the σ_{MAX} in Truncated SVD, better the resolution. We perform an ablation to analyze the trade-off between resolution and noise (as seen in Fig 9). We find a sweet spot at $\sigma_{\text{MAX}} = 40$ and use it for all evaluations.

Motor Dynamics: We evaluate the impact of increasing motor speeds from the default 600 rpm for static objects. As we increase speed, for the same rate of radar data, the number of samples per rotation decreases, and the sensor body vibration increases. Vibration affects the phase of mmWave radars [46]. Here, we empirically measure the MSE between the base image at 600 rpm and at higher speeds (Fig 11a). Cross comparing from qualitative data, we find that the mild variance in MSE for points < 1200 rpm, is due to lesser samples, and the high MSE for > 1200 rpm is when the body vibration severely affects. Increasing radar’s rate from 10 kHz, proofing motor vibration, and compensating for vibration in

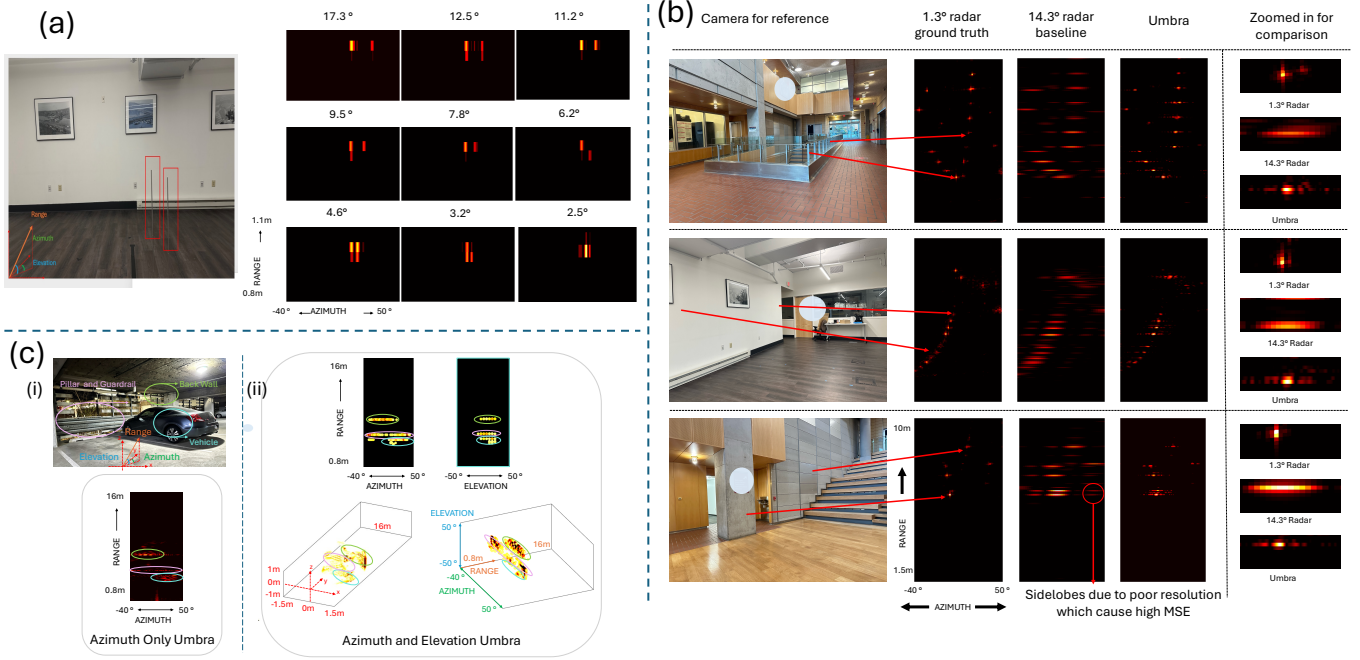


Fig. 8. (a) We characterize the azimuth resolution using two thin metal rods (in red boxes) as point targets. (b) Qualitative results show performance in static scenes. Intensity from [0, 1] is mapped to colorscale. The rightmost column shows that we can generate sharper, high resolution images compared to our baseline. (c) Full 3D reconstruction of objects in a parking lot. Objects are marked with same color across different views.

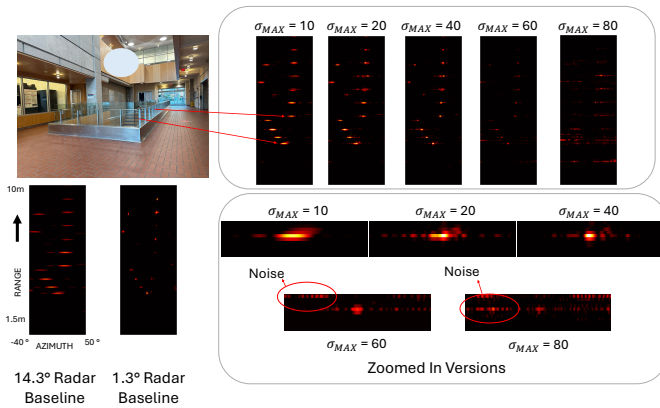


Fig. 9. Results showing the effect of varying the number of singular values used for reconstruction.

software will make Umbra robust. Our goal is to show the sensitivity for a certain implementation.

Object Dynamics: We use a car driving into a parking lot at 5-10 mph as a moving test object. Note that other use cases in Sec I involve human motion which is <3-5mph. Over one spin of the inverse pinhole, variance due to angle-dependent pinhole positions and object motion both influence measurements. Fig 11b shows the increasing phase change as an object moves away added on top of the base phase variations due to object’s angle. The same phenomena also affects SAR systems. However, if the object’s doppler velocity is known, these effects can be easily compensated [47]. The propagation model in Sec IV-B can be adapted to consider both object angle and velocity, to perform joint estimation. Focusing on a specific doppler velocity, and performing our algorithm, Fig 11a shows that as the car approaches the radar,

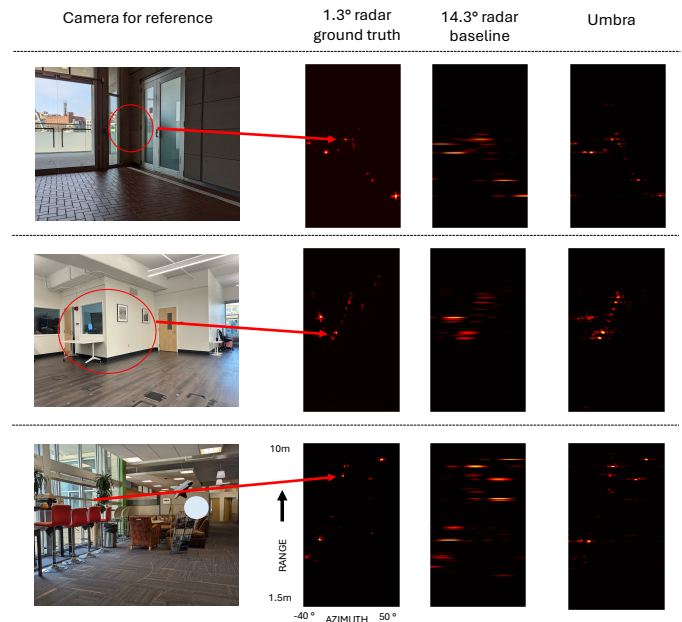


Fig. 10. More qualitative results in diverse scenes.

the object location approaches 0 m. The image quality holds good across frames. These images are generated at 10 fps with a 600 rpm motor. Higher object speeds demand dedicated motion compensation algorithms for inverse pinhole context which we leave for future work. Object phase shifts affects both SAR and Umbra, we just show we can achieve imaging by moving a lighter object.

Occlusion Tests: Radar passes through occluders like fog with little effect (Fig 11d). Mean difference in MSE with and without fog across images is 0.63.




Fig. 11. (a) Error with motor speed. (b) Despite increasing phase shift due to object motion, the phase variations contributing to angle estimation (in red) are still encapsulated. (c) Qualitative result of a dynamic object like a car approaching a radar. The car's range decreases with time and the image quality is retained. (d) Umbra images show little error in seeing a wall with and without fog occlusion.




Fig. 12. Umbra shows how commodity objects like plastic fans can serve as inverse pinhole for resolution enhancement.

Full 3D Reconstruction: Because we rotate as opposed to linear scan, we can obtain both azimuth and elevation resolution enhancements. In Fig 8c(i), we first show range-azimuth images. Next, we modify the bidirectional propagation matrix to account for azimuth and elevation. Using this, we show range-azimuth slices, range-elevation slices, and range-azimuth-elevation images (Fig 8c(ii)). We can see the car, back wall, and guard rails resolved in this 3D space. We resort to qualitative study as the ground truth radar has a poor 18° elevation resolution. Owing to specularly of metal objects, the images don't show the 3D shape well and are similar in quality to other azimuth-elevation SAR images [17]. Perhaps a better comparison would involve moving a single-antenna radar with circular SAR. This needs expensive motors, precise μs radar-motor sync (we achieve effortlessly via self-sync in Sec IV-D), and a complex radar to computer setup without wire entanglement. Since our primary focus (also seen in the community [48]) is azimuth, we emphasize azimuth improvements and show a 3D proof of concept using real data.

Natural Inverse Pinhole Effect: Naturally rotating objects (e.g: fans etc.) can easily take advantage of Umbra. We show this with a 2 blade plastic fan with aerofoil shape, that is different from our radial strips. We perform a 1 time calibration of matrix B as mentioned in Sec IV-D. A key challenge with multi blades is a systematic phase change that arises due to the twin blade motion in addition to the inverse pinhole effect (Fig 12a shows the phase change for a point target). We estimate and cancel out this phase change before imaging. Fig 12b shows our performance using a twin blade fan. MSE for baseline and Umbra are 10.4 and 5.6 respectively.

VII. DISCUSSION

A. Multi-antenna vs SAR vs Umbra

For 2.5° resolution, a multi-antenna array needs at least 45 antennas over 8.5 cm ($\frac{\lambda}{D}$ with $D = 8.5$ cm). Umbra significantly reduces the required complexity and cost of RF. Virtual MIMO antennas can reduce the complexity needed

for emulating 45 antennas but are still far more expensive than a single antenna. Building a SAR rotator at 600 rpm while timestamping radar measurements obtained at 10 kHz, would need an ultra precise, high update rate, expensive motor encoder. Umbra implicitly tracks inverse pinholes at sub-millimeter precision, using only radar signatures, without any motor encoder.

B. Umbra with multi-antenna radars

We built Umbra with just 1 antenna. However, it is interesting to think about the implications for back-compatibility with an antenna array. If a 4-8 element array is available anyway at the radar, theoretically we could also use more than 1 antenna. The positions of all available antennas can be mapped in F and F' . We performed a singular value analysis and found that the singular value curve with 8 antennas is only slightly better than that for 1 antenna. Given the marginal improvement in the singular value curve, a single antenna radar remains sufficient and most cost-effective.

VIII. CONCLUSION

This paper presents a new aperture synthesis technique for upgrading a single-antenna radar to perform high resolution imaging using pinhole motion. Umbra's design point is attractive because of the ease with which a single antenna radar can be upgraded, that is, without additional RF complexity, and complicated precision motion tracking. Umbra leverages properties unique to mmWave setups such as bidirectional pinhole passes and the noise regime, where lightweight inverse pinhole designs are feasible. Beyond evaluating with a custom designed inverse pinhole, Umbra inspires creative exploitation of natural mmWave inverse pinholes like fan blades and broadens our understanding of what can be construed as enhancing a radio system's aperture.

REFERENCES

- [1] S. Zhou and L. Jiang, "Modern description of rayleigh's criterion," *Physical Review A*, vol. 99, no. 1, p. 013808, 2019.
- [2] "Awr cascade," <https://www.ti.com/tool/MMWCAS-RF-EVM>, 2024.
- [3] 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.
- [4] L. Dodds, T. Boroushaki, C. Ham, and F. Adib, "Mito: A millimeter-wave dataset and simulator for non-line-of-sight perception," *arXiv preprint arXiv:2502.10259*, 2025.
- [5] K. Qian, Z. He, and X. Zhang, "3d point cloud generation with millimeter-wave radar," *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.*, vol. 4, no. 4, 2020.

- [6] A. Levin, R. Fergus, F. Durand, and W. T. Freeman, "Image and depth from a conventional camera with a coded aperture," *ACM transactions on graphics (TOG)*, vol. 26, no. 3, pp. 70–es, 2007.
- [7] C. Zhou, S. Lin, and S. Nayar, "Coded aperture pairs for depth from defocus," in *2009 IEEE 12th international conference on computer vision*. IEEE, 2009, pp. 325–332.
- [8] H. Bradt, G. Garmire, M. Oda, G. Spada, B. Sreekantan, P. Gorenstein, and H. Gursky, "The modulation collimator in x-ray astronomy," *Space Science Reviews*, vol. 8, no. 4, pp. 471–506, 1968.
- [9] Z. Lu, J. Wang, Y. Li, P. Shen, Z. Zhu, J. Yang, H. Yang, B. Lin, X. Zhang, B. Liu *et al.*, "Hard x-ray imaging with a slit collimated telescope," *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 362, no. 2-3, pp. 551–555, 1995.
- [10] O. Cossairt, M. Gupta, and S. K. Nayar, "When does computational imaging improve performance?" *IEEE transactions on image processing*, vol. 22, no. 2, pp. 447–458, 2012.
- [11] K. Mitra, O. S. Cossairt, and A. Veeraraghavan, "A framework for analysis of computational imaging systems: role of signal prior, sensor noise and multiplexing," *IEEE transactions on pattern analysis and machine intelligence*, vol. 36, no. 10, pp. 1909–1921, 2014.
- [12] N. Ratner and Y. Y. Schechner, "Illumination multiplexing within fundamental limits," in *2007 IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, 2007, pp. 1–8.
- [13] A. Torralba and W. T. Freeman, "Accidental pinhole and pinspeck cameras: Revealing the scene outside the picture," *International Journal of Computer Vision*, vol. 110, pp. 92–112, 2014.
- [14] S. V. Hum, "Noise in radio systems," <https://www.waves.utoronto.ca/prof/svhum/ece422/notes/21-noise.pdf>, 2024.
- [15] "Awr1843boost," <https://www.ti.com/tool/AWR1843BOOST>, 2024.
- [16] L. Dodds, H. Shanbhag, J. Guan, S. Gupta, and H. Hassanieh, "Around the corner mmwave imaging in practical environments," in *Proceedings of the 30th Annual International Conference on Mobile Computing and Networking*, 2024, pp. 953–967.
- [17] J. Guan, S. Madani, S. Jog, S. Gupta, and H. Hassanieh, "Through fog high-resolution imaging using millimeter wave radar," in *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020, pp. 11 461–11 470.
- [18] Y. Golovachev, A. Etinger, G. A. Pinhasi, and Y. Pinhasi, "Millimeter wave high resolution radar accuracy in fog conditions—theory and experimental verification," *Sensors*, vol. 18, no. 7, p. 2148, 2018.
- [19] H. Xue, Y. Ju, C. Miao, Y. Wang, S. Wang, A. Zhang, and L. Su, "mmesh: towards 3d real-time dynamic human mesh construction using millimeter-wave," in *Proceedings of the 19th Annual International Conference on Mobile Systems, Applications, and Services*, 2021, p. 269–282.
- [20] C. Iovescu and S. Rao, "The fundamentals of millimeter wave sensors," 2017.
- [21] M. A. Richards, J. Scheer, W. A. Holm, and W. L. Melvin, "Principles of modern radar," 2010.
- [22] "Vayyar radar," <https://vayyar.com/auto/technology/79ghz/>, 2024.
- [23] B. Mamandipoor, G. Malysa, A. Arbabian, U. Madhow, and K. Noujeim, "60 ghz synthetic aperture radar for short-range imaging: Theory and experiments," in *2014 48th Asilomar Conference on Signals, Systems and Computers*. IEEE, 2014, pp. 553–558.
- [24] A. Prabhakara, V. Singh, S. Kumar, and A. Rowe, "Osprey: a mmwave approach to tire wear sensing," in *Proceedings of the 18th International Conference on Mobile Systems, Applications, and Services*, 2020, pp. 28–41.
- [25] X. Shuai, Y. Shen, Y. Tang, S. Shi, L. Ji, and G. Xing, "millieye: A lightweight mmwave radar and camera fusion system for robust object detection," in *Proceedings of the International Conference on Internet-of-Things Design and Implementation*, 2021, pp. 145–157.
- [26] H. Lai, G. Luo, Y. Liu, and M. Zhao, "Enabling visual recognition at radio frequency," in *Proceedings of the 30th Annual International Conference on Mobile Computing and Networking*, 2024, p. 388–403.
- [27] A. Prabhakara, T. Jin, A. Das, G. Bhatt, L. Kumari, E. Soltanaghahi, J. Bilmes, S. Kumar, and A. Rowe, "High resolution point clouds from mmwave radar," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 4135–4142.
- [28] S. Haykin, "Array signal processing," *Englewood Cliffs*, 1985.
- [29] S. U. Pillai, *Array signal processing*. Springer Science & Business Media, 2012.
- [30] D. Liu, J. Gu, Y. Hitomi, M. Gupta, T. Mitsunaga, and S. K. Nayar, "Efficient space-time sampling with pixel-wise coded exposure for high-speed imaging," *IEEE transactions on pattern analysis and machine intelligence*, vol. 36, no. 2, pp. 248–260, 2013.
- [31] A. Veeraraghavan, R. Raskar, A. Agrawal, A. Mohan, and J. Tumblin, "Dappled photography: Mask enhanced cameras for heterodyned light fields and coded aperture refocusing," *ACM Trans. Graph.*, vol. 26, no. 3, p. 69, 2007.
- [32] J. Gollub, O. Yurduseven, K. P. Trofatter, D. Arnitz, M. F. Imani, T. Sleasman, M. Boyarsky, A. Rose, A. Pedross-Engel, H. Odabasi *et al.*, "Large metasurface aperture for millimeter wave computational imaging at the human-scale," *Scientific reports*, vol. 7, no. 1, p. 42650, 2017.
- [33] O. Yurduseven, D. L. Marks, T. Fromenteze, J. N. Gollub, and D. R. Smith, "Millimeter-wave spotlight imager using dynamic holographic metasurface antennas," *Optics Express*, vol. 25, no. 15, pp. 18 230–18 249, 2017.
- [34] N. Gopalsami, S. Liao, T. W. Elmer, E. R. Koehl, A. Heifetz, A. C. Raptis, L. Spinoulas, and A. K. Katsaggelos, "Passive millimeter-wave imaging with compressive sensing," *Optical Engineering*, vol. 51, no. 9, pp. 091 614–091 614, 2012.
- [35] J. Han, L. Li, S. Tian, G. Liu, H. Liu, and Y. Shi, "Millimeter-wave imaging using 1-bit programmable metasurface: Simulation model, design, and experiment," *IEEE Journal on Emerging and Selected Topics in Circuits and Systems*, vol. 10, no. 1, pp. 52–61, 2020.
- [36] R. Sharma, R. Hussung, A. Keil, F. Friederich, T. Fromenteze, M. Khalily, B. Deka, V. Fusco, and O. Yurduseven, "Coded-aperture computational millimeter-wave image classifier using convolutional neural network," *IEEE Access*, vol. 9, pp. 119 830–119 844, 2021.
- [37] R. Sharma, J. Zhang, R. Kumar, B. Deka, V. Fusco, and O. Yurduseven, "3-d super-resolution of coded aperture millimeter-wave images using complex-valued convolutional neural network," *IEEE Sensors Journal*, vol. 22, no. 21, pp. 20 921–20 936, 2022.
- [38] D. R. Smith, M. S. Reynolds, J. N. Gollub, D. L. Marks, M. F. Imani, O. Yurduseven, D. Arnitz, A. Pedross-Engel, T. Sleasman, P. Trofatter *et al.*, "Security screening via computational imaging using frequency-diverse metasurface apertures," in *Passive and Active Millimeter-Wave Imaging XX*, vol. 10189. SPIE, 2017, pp. 76–82.
- [39] Y. Wang, Y. Lu, Y. Zhou, Y. Shen, L. Qiu, Z. Lai, Y.-C. Chen, H. Pan, J. Zhou, D. Ding *et al.*, "High-resolution mmwave imaging using metasurface and diffusion," 2025.
- [40] T. Woodford, K. Qian, and X. Zhang, "Metasight: High-resolution nlos radar with efficient metasurface encoding," in *ACM SenSys*, 2023.
- [41] J. W. Goodman, *Introduction to Fourier optics*. Roberts and Company publishers, 2005.
- [42] H. Amca, "The presence of thermal, shot, impulsive and flicker noise at millimeter-wave frequencies," in *2019 International Symposium on Networks, Computers and Communications (ISNCC)*. IEEE, 2019, pp. 1–6.
- [43] ITU-R, "Radio noises," 2024.
- [44] M. Müller, "Dynamic time warping," *Information retrieval for music and motion*, pp. 69–84, 2007.
- [45] "Mwt-td-mas-300," <https://www.mwtmaterials.com/wp-content/uploads/2020/04/MWT-TD-MAS-300-Series-Ver-4.1.pdf>, 2024.
- [46] C. Jiang, J. Guo, Y. He, M. Jin, S. Li, and Y. Liu, "mmvib: micrometer-level vibration measurement with mmwave radar," in *Proceedings of the 26th Annual International Conference on Mobile Computing and Networking*, 2020, pp. 1–13.
- [47] M. Arii, "Efficient motion compensation of a moving object on sar imagery based on velocity correlation function," *IEEE transactions on geoscience and remote sensing*, vol. 52, no. 2, pp. 936–946, 2013.
- [48] Y. Liu, F. Wang, N. Wang, and Z. Zhang, "Echoes beyond points: Unleashing the power of raw radar data in multi-modality fusion," *Neurips 2024*.



Akarsh Prabhakara is an Assistant Professor in the Department of Computer Sciences at the University of Wisconsin–Madison. He earned his Ph.D. from Carnegie Mellon University in 2024. He builds wireless systems that perceive the world and communicate with high fidelity unlocking new application potentials.



Yawen Liu is a Ph.D. candidate in the Electrical and Computer Engineering department at Carnegie Mellon University. Her research centers on mmWave sensing systems, ubiquitous computing, and smart metamaterial surfaces, with a particular interest in advancing mobile and pervasive wireless technologies.



Aswin C. Sankaranarayanan is a professor in the ECE department at CMU. His research interests are broadly in compressive sensing, computational photography, signal processing and machine vision. His doctoral research was in the University of Maryland where his dissertation won the distinguished dissertation award from the ECE department in 2009. Aswin is the recipient of the multiple best paper awards at top vision and graphics venues including CVPR 2019, SIGGRAPH 2023 and a best paper honorable

mention in ICCV 2025, the 2018 CIT Dean's Early Career Fellowship, the 2018 Spira Teaching award, 2017 NSF CAREER award, the Eta Kappa Nu (CMU Chapter) Excellence in Teaching awards, and the Herschel Rich Invention award from Rice University in 2016.



Anthony Rowe is the Siewiorek and Walker Family Professor in the Electrical and Computer Engineering Department at Carnegie Mellon University as well as a Chief Scientist at Bosch Research. His research interests are in networked real-time embedded systems with a focus on wireless communication. He earned a Ph.D in Electrical and Computer Engineering from CMU in 2010.



Swarun Kumar is the Sathaye Family Foundation Professor at Carnegie Mellon University's ECE department, with an affiliate appointment in the CS department, the HCI Institute and the Cylab Security and Privacy Institute. Swarun heads the laboratory for emerging wireless technologies (WiTech lab). He designs and builds novel systems to enable faster wireless networks and new services. His research has impacted a wide-range of domains ranging from low-power wide-area networking, wireless localization,

RFID systems, mmWave sensing and satellite networking. Swarun is a recipient of the 2024 ACM SIGMOBILE Rockstar award, the 2021 ACM SIGBED Early Career Researcher Award, the NSF CAREER award and the Google Faculty Research Award. Swarun received the George Sprows Award for best Ph.D thesis in Computer Science at MIT in 2015 and the President of India gold medal at IIT Madras in 2010.