



Research article

Curvature-aware inverse specular geometry for camera-assisted mmWave NLoS localization: identifiability, stability, and numerical analysis

Shih-Lin Lin*

Graduate Institute of Vehicle Engineering, National Changhua University of Education, Changhua City, Taiwan

* **Correspondence:** Email: lin040@cc.ncue.edu.tw.

Abstract: This paper develops a curvature-aware specular-reflection inverse-geometry model for camera-assisted millimeter-wave non-line-of-sight localization with an unknown curved reflector and a hidden target. In a calibrated two-dimensional setting, a finite-dimensional C^3 regular curve, a locally selected nondegenerate Fermat branch, radar range/azimuth measurements, and camera-derived reflector constraints form a whitened augmented forward operator. Implicit differentiation yields a reduced Jacobian for analyses of local identifiability, Fisher-Schur information, and small-perturbation stability. The tangent-plane approximation induces a second-order range bias, but the azimuth and generic localization biases remain first order. A scaled Levenberg-Marquardt solver is evaluated using 400 paired synthetic trials. Under nominal noise, the curved fusion model attains a localization root-mean-square error (RMSE) of 0.145 m versus 0.199 m for the planar model; the paired 95% bootstrap improvement interval is [0.0444, 0.0625] m. Initialization-scale tests and a deterministic 12-geometry, two-level Plackett-Burman screen further characterize empirical local recovery and within-design geometric sensitivity. These synthetic results are limited to calibrated two-dimensional geometry, a finite-dimensional reflector, known single-bounce association, a nongrazing local branch, and branch-valid initialization; they do not establish global path uniqueness, convergence from arbitrary initial values, or vehicle-mounted experimental validity.

Keywords: non-line-of-sight localization; mmWave radar; camera-radar fusion; specular-reflection inverse problem; curvature-aware modeling; local identifiability; Levenberg-Marquardt method

1. Introduction

When autonomous vehicles operate at urban intersections, buildings, parked vehicles, and roadside infrastructure often create line-of-sight occlusions, preventing cameras and direct-line-of-sight radar from observing crossing vehicles or pedestrians in a timely manner. Edge-assisted Internet-of-Vehicles cooperation can reduce risk through multisource sensing, but its effectiveness remains constrained by communication quality, sensing uncertainty, and traffic-flow stability [1,2]. Accordingly, extracting positional information from interpretable physical propagation paths without requiring direct visibility of the target has become an important problem in safety-critical perception for cooperative autonomous driving.

Frequency-modulated continuous-wave (FMCW) millimeter-wave (mmWave) radar provides range and azimuth information from beat frequencies and array phases, while offering all-weather operation and stable ranging performance. Existing studies have applied deep learning to non-line-of-sight (NLoS) echo reconstruction and angle estimation around urban street corners and have developed multichannel echo-separation and preceding-vehicle ranging methods [3–5]. These results demonstrate that interference suppression and feature extraction can improve observation quality. However, when echoes are produced by an unknown reflecting vehicle body or curved surface, measurement precision alone does not ensure that the target position can be recovered uniquely and stably.

Around-the-corner radar studies have demonstrated that multipath reflections from building walls can be used to detect occluded moving targets and perform single-sensor localization [6–8]. Subsequent work has developed multipath imaging, joint detection–localization algorithms, and geometric models of L-shaped street corners [9–11]. Beyond single-target formulations, methods for multitarget localization, single-channel imaging, and driver-assistance radar have also been proposed [12–14], establishing the practical feasibility of NLoS localization with automotive mmWave radar [15]. Nevertheless, most methods either assume that the reflecting wall is a known plane or focus primarily on echo detection and image focusing; the inverse-problem structure that arises when the reflector geometry is both unknown and curved has received limited attention.

Cameras and radar provide complementary information: A camera can supply priors on the contour, position, and orientation of the reflector, whereas radar provides propagation-range information associated with the occluded target. Camera-radar fusion and camera-assisted mmWave imaging have shown that pose constraints can improve the robustness of localization and phase calibration [16,17]. Meanwhile, multipath-assisted simultaneous localization and mapping (SLAM) and the fusion of propagation-path data have treated the reflective environment as a set of estimable parameters [18–20]. These studies indicate that reflection paths need not be uniformly discarded as noise and can instead provide exploitable geometric information. Whether such information resolves the coupling between target and reflector parameters, however, must be determined from the Jacobian rank and the effective information matrix.

Localization theory commonly uses Fisher information and the Cramér-Rao lower bound to quantify the effects of geometric configuration and noise on estimation accuracy [21,22], while optical NLoS imaging has likewise shown that an accurate propagation model is essential for successful inversion [23]. In the present problem, the reflection point is implicitly determined by the Fermat

stationarity condition, and the forward mapping can lose regularity because of surface curvature, branch switching, or grazing incidence. Even when a Levenberg-Marquardt method is used to solve the nonlinear least-squares problem [24,25], numerical convergence alone is insufficient evidence of solution uniqueness or reliability unless conditions for local identifiability and stability have first been established.

Recent studies [26–28] have investigated real-time tracking of microwave scattering matrices, two-dimensional inverse-scattering imaging, and the identification of small anomalies in unknown backgrounds. Other works [29,30] have used synthetic and experimental data for small-target detection and combined linear sampling with neural networks to reconstruct the shapes of acoustic obstacles. Collectively, these studies highlight the shared importance of an accurate forward model, finite observations, and robust inversion, but they do not address camera-constrained automotive-radar localization based on specular reflection from curved surfaces.

In addition, array-based direction-of-arrival estimation and cross-view street-scene data fusion have demonstrated the respective values of angular information and heterogeneous data for traffic perception [31,32]. Modified and inexact Levenberg-Marquardt methods, together with inverse studies of time-varying reaction coefficients, provide mathematical context for local iterative methods, regularization, and discrete stability [33–35]. Plackett-Burman designs enable efficient multifactor screening [36]. Overall, the literature still lacks an analytical framework that simultaneously accommodates curvature, implicitly defined specular points, camera priors, joint range/azimuth observations, and an explicit distinction among local identifiability, stability, and model bias.

Table 1 clarifies the contribution relative to the closest method families. The novelty lies in implicitly differentiating the branch-selected specular point on a finite-dimensional curved reflector and using the resulting reduced derivatives to analyze joint target-reflector recovery, local stability, and tangent-plane bias. The Levenberg-Marquardt method and Fisher information are established tools used, respectively, for numerical inversion and local information assessment; neither is claimed as a standalone contribution.

Table 1. Comparison with representative NLoS localization and sensor-fusion frameworks.

Framework	Reflector model	Unknown parameters	Camera constraint	Reflection-point treatment	Theoretical guarantee
Around-the-corner radar [6–15]	Known planar wall or L-shaped corner	Target position(s)	None	Fixed or derived from planar geometry	Detection/localization formulations; no joint curved-reflector analysis
Camera-radar fusion and imaging [16,17]	Target/scene geometry or calibrated imaging model	Vehicle pose or calibration variables	Yes	Not jointly estimated on an unknown curved reflector	Empirical fusion and calibration performance
Multipath positioning and SLAM [18–20]	Virtual anchors or estimable propagation map	Position and map/path states	No camera contour prior	Mapped paths or virtual anchors	Information or SLAM formulations; no curvature-bias order

Continued on next page

Framework	Reflector model	Unknown parameters	Camera constraint	Reflection-point treatment	Theoretical guarantee
Virtual-station weighted least squares (WLS) [37]	Planar image/virtual-station geometry	Target position; mapped stations supplied	External camera-map input in this study	Eliminated by virtual-station representation	WLS localization; no joint branch-local curved analysis
Present study	Finite-dimensional C^3 curved reflector	Target position and reflector parameters	Yes	Implicit Fermat point on a selected nondegenerate branch	Joint local identifiability; Fisher-Schur position information; local stability; tangent-plane bias orders

To address these gaps, this study proposes a curvature-aware specular-reflection inverse-geometry framework for camera-assisted mmWave NLoS localization. In a calibrated two-dimensional setting, the reflector is represented by a finite-dimensional C^3 regular curve. For each radar pose, one associated and nondegenerate single-bounce Fermat branch is fixed, and radar range/azimuth observations are combined with camera constraints on the reflector parameters to jointly estimate the position of the occluded target and the reflector geometry.

In this study, a model is termed curvature-aware if its range/azimuth predictions and Jacobian are evaluated using the local curved-reflector geometry at the selected Fermat point, without substituting the reference-point tangent line; curvature-induced changes in the reflection point and local normal are therefore retained.

In contrast to planar-reflector and virtual-station formulations, the proposed curvature-aware inverse model evaluates and differentiates both range and azimuth on the curved reflector at the selected Fermat point. It therefore retains curvature-induced changes in the reflection point and local normal while jointly estimating the target and the camera-constrained reflector parameters.

The main contributions of this study are as follows:

(1) A curvature-aware reduced forward model is established that explicitly distinguishes globally multivalued reflections from a locally selected branch. Accordingly, this study neither assumes that an arbitrary curved surface can be reconstructed completely from finite observations nor claims global path uniqueness.

(2) The sensitivities of range and azimuth to the unknown state are derived through implicit differentiation, and a whitened augmented Jacobian, Fisher information matrix, and Schur-complement effective position information matrix are used to formulate conditions for regular local identifiability and small-perturbation stability.

(3) The distinct bias orders of the tangent-plane approximation are identified: The range bias is $O(H^2)$, whereas the azimuth bias and the bias in generic joint range-azimuth localization are $O(H)$. In 400 paired synthetic trials, the localization RMSEs of the curved and planar fusion models are 0.145 m and 0.199 m, respectively; however, these results provide only synthetic numerical evidence.

Section 2 establishes the geometric and measurement models. Section 3 analyzes local identifiability, Fisher-Schur information, and local stability. Section 4 describes the numerical algorithm and experimental design; Section 5 presents the results and discussion; and Section 6

summarizes the limitations and directions for future work. The conclusions are restricted to a calibrated two-dimensional setting, known single-bounce reflection associations, a finite-dimensional family of regular curves, nongrazing local branches, and initializations within a small-perturbation regime. Unknown multipath associations, diffuse reflection, three-dimensional free-form surfaces, and validation using vehicle-mounted hardware are beyond the scope of this study.

2. Curvature-aware specular-reflection forward model and inverse problem

This section formulates camera-assisted NLoS localization as a finite-dimensional curved-reflector inverse problem on a locally selected Fermat branch; globally multivalued reflection paths are identified without undertaking a global analysis. Figure 1 illustrates the single-bounce geometry and its tangent-plane comparator.




Figure 1. Curvature-aware single-bounce specular-reflection NLoS geometry. Four radar poses observe the occluded target via a curved reflector; the gray dashed line denotes the tangent-plane approximation at the reference point, and the dotted lines indicate direct paths blocked by the occluder.

2.1. Geometric configuration and unknown variables

Consider a two-dimensional ground-plane domain $\Omega \subset \mathbb{R}^2$. Let

$$a_i \in \Omega, \quad i = 1, \dots, m \quad (2.1)$$

denote the known equivalent monostatic radar phase center associated with measurement i . Distinct a_i may represent different radar poses, motion-compensated scan locations, or virtual sensing locations.

Let

$$x = \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} \in X \subset \Omega \quad (2.2)$$

denote the position of the occluded target, where X is a prescribed compact search region. All measurements used in a single inversion are assumed to correspond to the same target position or to have been compensated to a common reference time.

The reflector boundary of the parked vehicle is represented by the regular parametric curve

$$\Gamma_\theta = \{r(s; \theta) : s \in I\}, \quad I = [-h, h], \quad (2.3)$$

where $\theta \in \Theta \subset \mathbb{R}^p$ is a finite-dimensional parameter vector and $h > 0$ is the half-width of the candidate reflecting segment. The joint unknown state is

$$z = \begin{bmatrix} x \\ \theta \end{bmatrix} \in \mathbb{R}^{2+p}. \quad (2.4)$$

The finite-dimensional assumption is essential. The measurements considered here cannot uniquely reconstruct an arbitrary C^3 curve. Instead, camera observations provide geometric priors that constrain the reflector to a low-dimensional family of curves.

In practice, calibrated contour or depth landmarks, such as those represented in Eq (2.62), constrain the reflector reference point, orientation, and selected shape coefficients. The inversion therefore searches over a finite-dimensional family rather than arbitrary C^3 curves, while C^3 regularity supplies the smooth tangent, curvature, and branch sensitivities required by the subsequent analysis.

2.1.1. Gauge-fixed curve representation

A convenient parameterization is

$$r(s; \theta) = q + Q(\psi)\gamma(s; \beta), \quad \theta = (q_1, q_2, \psi, \beta^\top)^\top, \quad (2.5)$$

where $q \in \mathbb{R}^2$ is a reference point, ψ is the reflector orientation angle, $\beta \in \mathbb{R}^k$ contains the shape parameters, and $p = k + 3$. The rotation matrix is

$$Q(\psi) = \begin{bmatrix} \cos\psi & -\sin\psi \\ \sin\psi & \cos\psi \end{bmatrix}. \quad (2.6)$$

The following normalization conditions fix the reference point, tangent direction, and local parameter scale, thereby removing the corresponding local gauge redundancies:

$$\gamma(0; \beta) = \mathbf{0}, \quad \frac{\partial \gamma}{\partial s}(0; \beta) = \begin{bmatrix} 1 \\ 0 \end{bmatrix}. \quad (2.7)$$

Consequently, $q = r(0; \theta)$ is the actual physical reference point, whereas $Q(\psi)e_1$ is the tangent direction at $s = 0$. The shape parameters β therefore describe curvature variations only and do not duplicate the rigid position and orientation variables.

Assume that the curve satisfies

$$r(\cdot; \theta) \in C^3(I; \mathbb{R}^2), \quad \|r_s(s; \theta)\| \geq v_0 > 0 \quad (2.8)$$

and that this condition holds for all (s, θ) in the admissible neighborhood.

Define

$$v(s; \theta) = r_s(s; \theta), \quad w(s; \theta) = r_{ss}(s; \theta) \quad (2.9)$$

and define the unit tangent and unit normal vectors by

$$t(s; \theta) = \frac{v(s; \theta)}{\|v(s; \theta)\|}, \quad n(s; \theta) = Jt(s; \theta), \quad (2.10)$$

where

$$J = \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix}. \quad (2.11)$$

The signed curvature is

$$\kappa(s; \theta) = \frac{\det(v(s; \theta), w(s; \theta))}{\|v(s; \theta)\|^3}. \quad (2.12)$$

When s is an arc-length parameter, $\|v\| = 1$, and

$$r_{ss}(s; \theta) = \kappa(s; \theta)n(s; \theta), \quad \|v(s; \theta)\| = 1. \quad (2.13)$$

2.1.2. Reflector models

The primary analytical example is a circular arc with constant signed curvature κ :

$$\gamma_{\text{arc}}(s; \kappa) = \begin{cases} \left[\frac{\sin(\kappa s)}{\kappa}, \frac{1 - \cos(\kappa s)}{\kappa} \right]^T, & \kappa \neq 0, \\ [s, 0]^T, & \kappa = 0. \end{cases} \quad (2.14)$$

As $\kappa \rightarrow 0$, this parameterization converges continuously to the planar model and remains a unit-speed curve with constant curvature κ .

A local polynomial model can be written as

$$\gamma_{\text{poly}}(s; \beta) = \left[\sum_{\ell=2}^K \beta_{\ell} \varphi_{\ell}(s) \right], \quad (2.15)$$

where $\ell=2, \dots, K$, the shape-parameter dimension is $k=K-1$, and the basis functions satisfy

$$\varphi_{\ell}(0) = \varphi'_{\ell}(0) = 0. \quad (2.16)$$

For example, setting $\varphi_2(s) = s^2/2$ yields the quadratic approximation

$$\gamma_{\text{quad}}(s; \kappa) = \begin{bmatrix} s \\ \frac{\kappa s^2}{2} \end{bmatrix}. \quad (2.17)$$

A cubic B-spline curve may also be used in the numerical experiments:

$$\gamma_{\text{spl}}(s; \beta) = \sum_{j=1}^K B_{j,3}(s) P_j(\beta), \quad (2.18)$$

subject to the normalization constraints in Eq (2.7). The theoretical results below apply to any finite-dimensional family of C^3 curves satisfying the regularity assumptions and are independent of the chosen curve basis.

2.2. Fermat paths and specular reflection

For radar location i , define the distance from the reflector to the radar as

$$d_{a,i}(s; \theta) = \|a_i - r(s; \theta)\|. \quad (2.19)$$

The distance from the reflector to the target is

$$d_{x,i}(s; x, \theta) = \|x - r(s; \theta)\|. \quad (2.20)$$

The corresponding one-way equivalent broken-ray path length is

$$L_i(s; x, \theta) = d_{a,i}(s; \theta) + d_{x,i}(s; x, \theta). \quad (2.21)$$

For a candidate reciprocal monostatic propagation path, the round-trip path length and delay are, respectively,

$$\ell_i^{\text{rt}} = 2L_i, \quad \tau_i = \frac{2L_i}{c_0}, \quad (2.22)$$

where c_0 is the propagation speed. Because a conventional monostatic radar reports the range $c_0\tau_i/2$, its range output is equal to L_i . Accordingly, L_i is used throughout this paper as the radar range quantity.

Define the unit vectors directed from the candidate reflection point toward the radar and the target, respectively:

$$e_{a,i}(s; \theta) = \frac{a_i - r(s; \theta)}{d_{a,i}(s; \theta)}, \quad (2.23)$$

$$e_{x,i}(s; x, \theta) = \frac{x - r(s; \theta)}{d_{x,i}(s; x, \theta)}. \quad (2.24)$$

Differentiating Eq (2.21) with respect to s gives the Fermat stationarity function

$$F_i(s; x, \theta) = \partial_s L_i(s; x, \theta) = -(e_{a,i} + e_{x,i})^\top v. \quad (2.25)$$

A candidate specular reflection point is therefore a solution of

$$F_i(s; x, \theta) = 0. \quad (2.26)$$

Because $v = \|v\| t$, Eq (2.26) is equivalent to

$$t^\top (e_{a,i} + e_{x,i}) = 0. \quad (2.27)$$

Assume that the radar and target lie on the same reflective side of the local surface, i.e.,

$$e_{a,i}^\top n > 0, \quad e_{x,i}^\top n > 0. \quad (2.28)$$

Then Eq (2.27), together with the unit-norm constraints on $e_{a,i}$ and $e_{x,i}$, implies

$$e_{x,i} = -e_{a,i} + 2(e_{a,i}^\top n)n. \quad (2.29)$$

Equation (2.29) expresses the standard law of specular reflection: A ray incident on the surface from a_i is reflected toward x , with the angle of incidence equal to the angle of reflection.

For either propagation segment, differentiation gives $\partial_s d_{a,i} = -e_{a,i}^\top v$ and $\partial_s e_{a,i} = -(I - e_{a,i} e_{a,i}^\top)v/d_{a,i}$.

Consequently, $\partial_{ss} d_{a,i} = \frac{\|v\|^2 - (e_{a,i}^\top v)^2}{d_{a,i}} - e_{a,i}^\top w$. Applying the same calculation to $d_{x,i}$ and then summing the two second derivatives yields Eq (2.30).

The second derivative of the path length is

$$\partial_{ss} L_i = \frac{\|v\|^2 - (e_{a,i}^\top v)^2}{d_{a,i}} + \frac{\|v\|^2 - (e_{x,i}^\top v)^2}{d_{x,i}} - (e_{a,i} + e_{x,i})^\top w. \quad (2.30)$$

For a unit-speed curve, Eq (2.30) simplifies to

$$\partial_{ss} L_i = \frac{(e_{a,i}^\top n)^2}{d_{a,i}} + \frac{(e_{x,i}^\top n)^2}{d_{x,i}} - \kappa(e_{a,i} + e_{x,i})^\top n, \quad \|v\| = 1. \quad (2.31)$$

Equation (2.31) explicitly shows how curvature affects the local convexity of a Fermat path. A small $\partial_{ss} L_i$ indicates a nearly degenerate Fermat branch, making the reflection point highly sensitive to geometric perturbations.

Global stationary set:

For a given state $z = (x, \theta)$, define the physically admissible stationary set

$$S_i(z) = \{s \in \text{int}(I) : F_i(s; z) = 0, \partial_{ss} L_i(s; z) > 0, V_i(s; z) = 1\}, \quad (2.32)$$

where V_i is a visibility indicator. Its value is 1 if and only if all of the following conditions hold:

- (1) Neither propagation segment is occluded, except by the designated reflector.
- (2) Both the radar and the target lie on the reflective side of the surface.
- (3) The reflection is not at grazing incidence.
- (4) The path conforms to the assumed single-bounce reflection topology.

The mapping $z \mapsto S_i(z)$ may be multivalued globally. Consequently, in the absence of a path-association rule, the complete forward problem is set-valued.

This study assumes that preprocessing or temporal association has identified a reference interval

$$I_i = [s_{i,-}, s_{i,+}] \subset I \quad (2.33)$$

that contains the Fermat branch of interest. At the true state z_* , let

$$s_{i,*} \in \text{int}(I_i) \quad (2.34)$$

satisfy

$$F_i(s_{i,*}; z_*) = 0, \quad \partial_{ss} L_i(s_{i,*}; z_*) > 0. \quad (2.35)$$

Denote this local branch by

$$s_i^* = s_i^*(x, \theta). \quad (2.36)$$

Section 3 shows that the nondegeneracy condition in Eq (2.35) yields a locally unique and

differentiable branch. This study does not assume global uniqueness over all possible specular-reflection paths.

2.3. Radar measurement model

Let

$$r_i^*(x, \theta) = r(s_i^*(x, \theta); \theta) \quad (2.37)$$

denote the selected specular reflection point. The reduced equivalent range is

$$\rho_i(x, \theta) = L_i(s_i^*(x, \theta); x, \theta). \quad (2.38)$$

If the raw round-trip path length is used instead of the monostatic radar range, the corresponding observable is $2\rho_i$.

The radar direction of arrival associated with the reflection point is

$$p_i(x, \theta) = \frac{r_i^*(x, \theta) - a_i}{\|r_i^*(x, \theta) - a_i\|}, \quad (2.39)$$

and the local azimuth angle is

$$\phi_i(x, \theta) = \text{atan2}(p_{i,2}, p_{i,1}). \quad (2.40)$$

This angle is evaluated on a locally unwrapped branch so that it is differentiable in a neighborhood of the true state. Depending on the available radar output, the measurement function i may contain range only:

$$h_i(x, \theta) = \rho_i(x, \theta) \quad (2.41)$$

or both range and azimuth:

$$h_i(x, \theta) = \begin{bmatrix} \rho_i(x, \theta) \\ \phi_i(x, \theta) \end{bmatrix}. \quad (2.42)$$

Stacking all scalar radar observations gives

$$h_r(x, \theta) \in \mathbb{R}^{m_r}. \quad (2.43)$$

The radar data are modeled as

$$y_r = h_r(x_*, \theta_*) + \varepsilon_r + \delta_{r,r}, \quad (2.44)$$

where ε_r denotes measurement noise, and $\delta_{r,r}$ represents forward-model errors caused by reflector-model mismatch, unmodeled calibration errors, or weak propagation effects not captured by the idealized single-bounce reflection model.

For the probabilistic analysis, the nominal radar noise is assumed to satisfy

$$\varepsilon_r \sim \mathcal{N}(0, \Sigma_r), \quad \Sigma_r \succ 0. \quad (2.45)$$

The covariance matrix may include distinct range and angular variances and correlations between them. Because range and angle have different physical units, whitening is required.

2.4. Differential form of the reduced radar mapping

The derivatives obtained in this subsection are used to construct the reduced Jacobian needed to analyze identifiability, stability, and numerical convergence.

Let

$$P(e) = I_2 - ee^\top \quad (2.46)$$

denote the orthogonal projection operator onto the subspace perpendicular to the unit vector e . Differentiating Eq (2.25) at the selected reflection point gives

$$F_{i,x} = -\frac{v^\top P(e_{x,i})}{d_{x,i}} \quad (2.47)$$

and

$$F_{i,\theta} = v^\top \left[\frac{P(e_{a,i})}{d_{a,i}} + \frac{P(e_{x,i})}{d_{x,i}} \right] r_\theta - (e_{a,i} + e_{x,i})^\top r_{s\theta}. \quad (2.48)$$

All quantities in Eqs (2.47) and (2.48) are evaluated at $s = s_i^*(x, \theta)$. If

$$\partial_{ss} L_i(s_i^*; x, \theta) \neq 0, \quad (2.49)$$

then implicit differentiation of $F_i(s_i^*; x, \theta) = 0$ yields

$$D_x s_i^* = -\frac{F_{i,x}}{\partial_{ss} L_i(s_i^*; x, \theta)}, \quad D_\theta s_i^* = -\frac{F_{i,\theta}}{\partial_{ss} L_i(s_i^*; x, \theta)}. \quad (2.50)$$

Thus, the sensitivity of the reflection point is inversely proportional to the local Fermat curvature $\partial_{ss} L_i$.

For the reduced range, the envelope property eliminates derivatives of s_i^* :

$$D_x \rho_i = e_{x,i}^\top, \quad (2.51)$$

$$D_\theta \rho_i = -(e_{a,i} + e_{x,i})^\top r_\theta. \quad (2.52)$$

Indeed,

$$D_z \rho_i = \partial_z L_i + \partial_s L_i D_z s_i^* = \partial_z L_i \quad (2.53)$$

because $\partial_s L_i = 0$ holds at the specular reflection point.

By contrast, the azimuth derivative explicitly depends on the displacement of the reflection point. Define

$$p_i^\perp = J p_i. \quad (2.54)$$

Then

$$D_x \phi_i = \frac{(p_i^\perp)^\top v}{d_{a,i}} D_x s_i^* \quad (2.55)$$

and

$$D_{\theta} \phi_i = \frac{(p_i^{\perp})^{\top}}{d_{a,i}} (r_{\theta} + v D_{\theta} s_i^*). \quad (2.56)$$

Equations (2.50), (2.55), and (2.56) show that azimuth observations become increasingly sensitive as the Fermat branch approaches degeneracy. The denominator $\partial_{ss} L_i$ does not appear in the first derivative of the reduced range, but it enters the azimuth derivative directly.

Let

$$H_x(x, \theta) = D_x h_r(x, \theta), \quad H_{\theta}(x, \theta) = D_{\theta} h_r(x, \theta). \quad (2.57)$$

Then the full radar Jacobian matrix is

$$J_r(x, \theta) = [H_x(x, \theta) \quad H_{\theta}(x, \theta)]. \quad (2.58)$$

2.5. Camera-derived reflector observations

Assume that the camera can observe the vehicle serving as the reflector but not necessarily the occluded target. Its output is represented as

$$y_c = g(\theta_{\star}) + \varepsilon_c, \quad (2.59)$$

where the camera-geometry mapping

$$g: \Theta \rightarrow \mathbb{R}^{m_c} \quad (2.60)$$

is continuously differentiable.

Direct parameter observations can be written as

$$g(\theta) = \theta, \quad y_c = \theta_{\star} + \varepsilon_c. \quad (2.61)$$

More generally, suppose that the camera extracts N_c contour or depth landmarks from the vehicle. Then

$$g(\theta) = \text{vec} \left[\pi_c(r(u_1; \theta)), \dots, \pi_c(r(u_{N_c}; \theta)) \right], \quad (2.62)$$

where π_c denotes a calibrated ground-plane or image projection and u_j is the associated curve coordinate.

The nominal camera error is modeled as

$$\varepsilon_c \sim \mathcal{N}(0, \Sigma_c), \quad \Sigma_c > 0. \quad (2.63)$$

Depth uncertainty, contour-extraction errors, and small camera-to-radar extrinsic calibration errors can be incorporated into Σ_c . If these extrinsic parameters are estimated jointly rather than absorbed into the error term, they must be included explicitly in θ .

Define the camera Jacobian matrix

$$G_{\theta}(\theta) = D_{\theta} g(\theta). \quad (2.64)$$

The camera need not independently identify all reflector parameters. Rank deficiency in G_{θ} is permitted, provided that the joint radar-camera mapping contains sufficient independent information.

2.6. Augmented forward operator

Let

$$W_r = \Sigma_r^{-1}, \quad W_c = \Sigma_c^{-1}. \quad (2.65)$$

The whitened augmented forward operator is

$$\mathcal{H}(x, \theta) = \begin{bmatrix} W_r^{1/2} h_r(x, \theta) \\ W_c^{1/2} g(\theta) \end{bmatrix} \in \mathbb{R}^{m_r+m_c}. \quad (2.66)$$

The corresponding whitened observation vector is

$$d = \begin{bmatrix} W_r^{1/2} y_r \\ W_c^{1/2} y_c \end{bmatrix}. \quad (2.67)$$

The forward operator's Jacobian matrix has the following block structure:

$$D\mathcal{H}(x, \theta) = \begin{bmatrix} W_r^{1/2} H_x & W_r^{1/2} H_\theta \\ 0 & W_c^{1/2} G_\theta \end{bmatrix}. \quad (2.68)$$

This matrix is central to the subsequent analysis of regular local identifiability. A necessary dimensional condition is

$$m_r + m_c \geq 2 + p. \quad (2.69)$$

Equation (2.69) is necessary but not sufficient; the columns in Eq (2.68) must also be linearly independent.

Under the Gaussian model, the local Fisher information matrix is

$$\mathcal{I}(z) = D\mathcal{H}(z)^\top D\mathcal{H}(z). \quad (2.70)$$

Partitioning according to $z = (x, \theta)$ gives

$$\mathcal{I} = \begin{bmatrix} \mathcal{I}_{xx} & \mathcal{I}_{x\theta} \\ \mathcal{I}_{\theta x} & \mathcal{I}_{\theta\theta} \end{bmatrix}. \quad (2.71)$$

When θ is treated as a nuisance parameter, the Schur complement further quantifies its effect on the position information:

$$\mathcal{I}_x^{\text{eff}} = \mathcal{I}_{xx} - \mathcal{I}_{x\theta} \mathcal{I}_{\theta\theta}^{-1} \mathcal{I}_{\theta x}. \quad (2.72)$$

This expression applies provided that $\mathcal{I}_{\theta\theta}$ is nonsingular. If it is singular, the Moore-Penrose pseudoinverse should be used, and the corresponding expression should be interpreted on the identifiable subspace.

2.7. Inverse problem

In the ideal noiseless setting, the inverse problem seeks, from

$$\mathcal{H}(z_*) = d_* \quad (2.73)$$

to recover $z_* = (x_*, \theta_*)$.

For noisy observations, the unregularized weighted least-squares estimator is

$$\hat{z} \in \arg \min_{z \in \mathcal{Z}_{\text{ad}}} \Phi_0(z), \quad (2.74)$$

where the objective function is defined as

$$\Phi_0(x, \theta) = \frac{1}{2} \|W_r^{1/2}(h_r(x, \theta) - y_r)\|^2 + \frac{1}{2} \|W_c^{1/2}(g(\theta) - y_c)\|^2. \quad (2.75)$$

The admissible state set is

$$\mathcal{Z}_{\text{ad}} = \left\{ (x, \theta) \in X \times \Theta, \quad \begin{array}{ll} i = 1, \dots, m, \\ d_{a,i}(s_i^*; \theta) \geq d_0, & d_{x,i}(s_i^*; x, \theta) \geq d_0, \\ s_i^* \in \text{int}(I_i), & V_i(s_i^*; x, \theta) = 1 \end{array} \right\}, \quad (2.76)$$

where $d_0 > 0$ excludes collisions among the target, radar, and reflector.

When the camera geometry is poorly conditioned or a higher-dimensional spline model is used, the regularized problem can be written as

$$\hat{z}_\lambda \in \arg \min_{z \in \mathcal{Z}_{\text{ad}}} \left\{ \Phi_0(z) + \frac{\lambda}{2} \|R(\theta - \theta_{\text{ref}})\|^2 \right\}, \quad (2.77)$$

where θ_{ref} is a camera-derived reference estimate, R penalizes implausible shape variations, and $\lambda \geq 0$ is the regularization parameter. The numerical solution of Eq (2.77) is described in Section 4.

To characterize reflector-model mismatch, let the true reflector boundary be

$$r_{\text{true}}(s) = r(s; \theta_*) + \Delta r(s). \quad (2.78)$$

Its local discrepancy is measured by

$$\|\Delta r\|_{C^1(I)} = \sup_{s \in I} (\|\Delta r(s)\| + \|\Delta r_s(s)\|). \quad (2.79)$$

To preserve a nondegenerate Fermat branch, the C^2 perturbation is also assumed to be sufficiently small. The C^1 discrepancy in Eq (2.79) is the principal geometric-error term in the stability bound of Section 3.

2.8. Basic assumptions

The theoretical results that follow require the following assumptions.

Assumption 2.1. Calibrated two-dimensional geometry.

The radar poses and the camera-to-radar coordinate transformation are known. Any extrinsic parameters not treated as known are included in θ .

Assumption 2.2. Finite-dimensional regular reflector.

The reflector belongs to a gauge-fixed finite-dimensional C^3 family of curves satisfying Eqs (2.7) and (2.8). The parameter set Θ is locally bounded.

Assumption 2.3. Positive separation distance.

There exists $d_0 > 0$ such that, throughout the neighborhood under consideration, all radar-to-reflector and target-to-reflector distances are no smaller than d_0 .

Assumption 2.4. Known single-bounce reflection associations.

Each radar observation is associated with a specified reflector and a reference interval I_i . Multiple reflections, diffuse scattering, and unknown path associations are excluded from the core theory.

Assumption 2.5. Interior visible reflection.

The true specular reflection point lies in the interior of I_i , satisfies the same-side condition in Eq (2.28), and has two unobstructed propagation segments.

Assumption 2.6. Nongrazing incidence.

There exists $\nu_0 > 0$ such that, at the true reflection point,

$$\begin{aligned} e_{a,i}(s_{i,*}; \theta_*)^\top n(s_{i,*}; \theta_*) &\geq \nu_0, \\ e_{x,i}(s_{i,*}; x_*, \theta_*)^\top n(s_{i,*}; \theta_*) &\geq \nu_0, \quad i = 1, \dots, m. \end{aligned} \quad (2.80)$$

The two inner products are the normal-direction cosines of the incident and outgoing rays. The bound $\nu_0 > 0$ keeps both rays uniformly away from tangency, a regime in which small geometric perturbations can cause large bearing changes or switches between reflection branches.

Assumption 2.7. Nondegenerate Fermat branch.

There exists $\mu_0 > 0$ such that

$$\partial_{ss} L_i(s_{i,*}; z_*) \geq \mu_0, \quad i = 1, \dots, m. \quad (2.81)$$

This condition establishes the local uniqueness and differentiability of the selected Fermat branch.

Assumption 2.8. Regular observation mappings.

The radar and camera observation mappings are at least C^2 on a locally admissible neighborhood, and Σ_r and Σ_c are positive definite.

Assumption 2.9. Small local perturbations.

The measurement noise, camera error, and reflector-model mismatch are all sufficiently small that the selected Fermat branch neither leaves its assigned interval nor switches to another stationary path.

Assumption 2.10. Local scope.

All identifiability, stability, and convergence conclusions apply only within a neighborhood of the true state. No claim is made regarding the global injectivity of $\mathcal{H}$, the global uniqueness of specular reflection paths, or convergence from arbitrary initial values.

2.9. Mathematical scope

The proposed formulation is a high-frequency, single-bounce geometric-optics model rather than a full-wave electromagnetic scattering model. It represents one detected and correctly associated specular component as a broken ray whose reflection point satisfies the Fermat stationarity condition;

the modeled radar observables are determined by the corresponding geometric path length and arrival bearing. The formulation does not solve Maxwell's equations or explicitly predict material-dependent reflection coefficients, radar cross section, complex field amplitude and phase, polarization response, diffraction, surface-roughness effects, diffuse scattering, coherent interference, or additional multiple-bounce paths. Accordingly, the estimated reflector is a geometric boundary within the prescribed curve family, rather than an electromagnetic reconstruction of the reflecting material.

Small omitted propagation effects may be represented by the forward-model discrepancy term in Eq (2.44) only when they remain perturbative and preserve the selected reflection branch and path association. Effects that create, suppress, or switch propagation paths cannot generally be treated as small measurement errors and fall outside the local theory. The geometric-optics approximation is therefore most applicable when the wavelength is small relative to the relevant surface scale and a dominant specular component from a smooth reflector can be isolated.

Sections 3–5 address:

- (1) selected-branch existence and local uniqueness;
- (2) joint identifiability of the target and reflector;
- (3) propagation of sensing and reflector-model errors;
- (4) tangent-plane bias in the data and recovered position; and
- (5) local convergence of curvature-aware iteration.

These analyses remain local under Assumptions 2.1–2.10.

3. Local identifiability, stability, and curvature-model bias

This section develops a local theory for the selected Fermat branches introduced in Section 2. This analysis does not claim uniqueness for every possible reflection path. Under Assumptions 2.1–2.10, it instead determines whether an associated, nondegenerate, and nongrazing local branch admits regular recovery from radar and camera observations and quantifies how measurement perturbations and reflector-model errors affect the estimate.

3.1. Selected Fermat branches and regular local identifiability

For each true reflection coordinate $s_{i,*}$, Eq (2.81) shows that the partial derivative of the Fermat stationarity equation with respect to s is nonzero. By the implicit function theorem, there exists a neighborhood of the true state z_* in which $s_i^*(z)$, selected by Eq (2.32), is a unique continuously differentiable local branch. Consequently, the augmented forward operator in Eq (2.66) can be treated as an ordinary finite-dimensional smooth mapping in that neighborhood.

To account for the different physical units of position, orientation, and curvature on the singular values, let $S > 0$ be a positive definite diagonal state-scaling matrix, and define the normalized perturbation $u = S^{-1}(z - z_*)$ and the scaled Jacobian matrix at the true state as

$$\tilde{J}_* = D\mathcal{H}(z_*)S, \quad \alpha = \sigma_{\min}(\tilde{J}_*). \quad (3.1)$$

Theorem 3.1 (Regular local identifiability). Suppose Assumptions 2.1–2.10 hold, every selected Fermat branch satisfies Eq (2.81), and $\alpha > 0$. Then the Jacobian of $\mathcal{H}$ has full column rank in a neighborhood of z_* , and z_* is locally identifiable in the regular sense on that branch. That is, there exists a neighborhood U such that $\mathcal{H}(z) = \mathcal{H}(z_*)$ and $z \in U$ together imply $z = z_*$.

Proof. Let $n = 2 + p$. Equation (2.81) and the implicit function theorem establish a common neighborhood in which all selected reflection branches, and hence $\mathcal{H}$, are continuously differentiable. Because S is nonsingular, $\alpha > 0$ implies $\text{rank} D\mathcal{H}(z_*) = n$. Therefore, some $n \times n$ minor of $D\mathcal{H}(z_*)$ has a nonzero determinant. Continuity of $D\mathcal{H}$ allows the neighborhood to be shrunk so that this determinant remains nonzero; thus, $\text{rank} D\mathcal{H}(z) = n$ throughout that neighborhood. The rank is therefore constant in that neighborhood. The constant-rank theorem provides local coordinates in which $\mathcal{H}$ has the normal form $u \mapsto (u, 0)$, which is injective. Hence, after possibly shrinking the neighborhood U , $\mathcal{H}(z) = \mathcal{H}(z_*)$ with $z \in U$ implies $z = z_*$. This argument is branch-local and does not exclude remote stationary paths with identical data.

The role of camera observations in establishing full rank can be expressed directly in terms of nullspaces. When the weighting matrices are positive definite, the augmented Jacobian matrix has full column rank if and only if

$$\ker(J_r) \cap (\mathbb{R}^2 \times \ker(G_\theta)) = \{0\}. \quad (3.2)$$

This condition shows that the camera need not identify all reflector parameters independently. If the geometric information retained by the camera eliminates the residual nullspace of the radar Jacobian matrix, the joint radar-camera observations establish regular local identifiability.

3.2. Fisher-Schur effective position information and accuracy lower bound

In Eqs (2.70)–(2.72), the reflector parameters are treated as nuisance parameters, and the Schur complement isolates the Fisher-Schur effective position information. If $\mathcal{J}_x^{\text{eff}} > 0$, every locally unbiased position estimator satisfies the corresponding Cramér-Rao lower bound, and the root-mean-square lower bound for two-dimensional position can be written as

$$\text{RMSE}_{\text{CRLB}} = \sqrt{\text{tr}[(\mathcal{J}_x^{\text{eff}})^{-1}]}. \quad (3.3)$$

The reported Cramér-Rao lower bound (CRLB) is a local theoretical lower bound for locally unbiased estimators under the stated Gaussian noise model and Assumptions 2.1–2.10. It does not guarantee a practical error level for arbitrary geometries, biased estimators, or branch-changing regimes.

As the camera-noise level decreases, the information in the reflector-parameter block increases, thereby reducing the coupling penalty between the target position and the reflector. Conversely, if the radar poses are collocated or the measurement directions are repeated, the radar measurements may still fail to provide sufficiently independent geometric directions for identifying the target position, even when the camera accurately constrains the reflector. Thus, camera-geometric accuracy and radar spatial baselines are complementary rather than interchangeable.

3.3. Small-perturbation stability and model mismatch

Let $\tilde{J}(u) = D\mathcal{H}(z_* + Su)S$. Suppose that, within a normalized neighborhood of radius r , the scaled Jacobian matrix satisfies

$$\sup_{\|u\|_2 \leq r} \|\tilde{J}(u) - \tilde{J}_*\|_2 \leq \frac{\alpha}{2}. \quad (3.4)$$

For $u = S^{-1}(z - z_*)$, let $z(t) = z_* + tSu$ and $\tilde{J}(tu) = D\mathcal{H}(z(t))S$. The fundamental theorem of calculus gives $\mathcal{H}(z) - \mathcal{H}(z_*) = \tilde{J}_*u + \int_0^1 [\tilde{J}(tu) - \tilde{J}_*]u dt$. Equation (3.4) therefore implies $\|\mathcal{H}(z) - \mathcal{H}(z_*)\|_2 \geq \alpha/2 \|u\|_2$, which is Eq (3.5):

$$\|\mathcal{H}(z) - \mathcal{H}(z_*)\|_2 \geq \frac{\alpha}{2} \|S^{-1}(z - z_*)\|_2. \quad (3.5)$$

Write the whitened observation as $d = \mathcal{H}(z_*) + \eta + b_r$ and set $\delta = \eta + b_r$. If the local estimate $\hat{z}$ remains in this neighborhood and its data residual does not exceed the residual at z_* , then the triangle inequality gives $\|\mathcal{H}(\hat{z}) - \mathcal{H}(z_*)\|_2 \leq 2\|\delta\|_2 \leq 2(\|\eta\|_2 + C_r \|\Delta r\|_{C^1(I)})$. Combining this inequality with Eq (3.5) yields Eq (3.6):

$$\|S^{-1}(\hat{z} - z_*)\|_2 \leq \frac{4}{\alpha} (\|\eta\|_2 + C_r \|\Delta r\|_{C^1(I)}). \quad (3.6)$$

Corollary 3.1 (Local Lipschitz stability). Under the preceding conditions, the recovery error varies to first order with the whitened measurement perturbation and the small reflector-model error. The factor $4/\alpha = (2/\alpha) \times 2$ combines the inverse lower slope in Eq (3.5) with the factor of two from the residual comparison; the local bound ceases to apply when the Jacobian approaches rank deficiency, the Fermat branch switches, or the geometry approaches grazing incidence.

These results establish a sequence of distinct local implications. First, the nondegeneracy condition in Assumption 2.7 permits the implicit function theorem to be applied, ensuring that each selected reflection coordinate varies smoothly with the state and that the reduced forward operator is locally well defined. Second, full column rank of the augmented Jacobian, equivalently $\alpha > 0$ in Eq (3.1), provides regular local identifiability on the selected branch, but not global uniqueness. Third, the Fisher-Schur complement in Eq (2.72) quantifies the local positional information remaining after the reflector parameters are treated as nuisance variables and yields the corresponding CRLB under the stated Gaussian and local-unbiasedness assumptions. Local stability requires the additional Jacobian-variation condition in Eq (3.4) and sufficiently small perturbations that preserve the selected branch, leading to the error bound in Eq (3.6). These rank, information, and stability results do not by themselves guarantee convergence of the Levenberg-Marquardt solver, which also depends on initialization, damping and step acceptance, and continued satisfaction of the branch-admissibility conditions.

3.4. Bias orders of the tangent-plane approximation

For the circular-arc model in Eq (2.14), when $|s| \leq H$ and $|\kappa|H \ll 1$, the difference from the reference-point tangent-plane parameterization $\gamma_0(s) = [s, 0]^\top$ has the Taylor expansion

$$\gamma_{\text{arc}}(s; \kappa) - \gamma_0(s) = \begin{bmatrix} -\kappa^2 s^3 / 6 \\ \kappa s^2 / 2 \end{bmatrix} + \begin{bmatrix} O(\kappa^4 s^5) \\ O(\kappa^3 s^4) \end{bmatrix}. \quad (3.7)$$

For a bounded-curvature family, Eq (3.7) gives an $O(s^2)$ position discrepancy and an $O(s)$ tangent-direction discrepancy relative to the tangent line. To define the limit $H \rightarrow 0$ and pair the curved and planar reflection branches, consider the following uniform small-footprint construction.

Let n index a sequence of admissible scenes with true states z_n^* , a fixed number m of radar poses, fixed state and observation dimensions, and uniformly bounded state-scaling and whitening operators. For $0 \leq \lambda \leq 1$, define the circular-arc homotopy

$$r_{\lambda,n}(s) = q_n + Q(\psi_n)\gamma_{\text{arc}}(s; \lambda\kappa_n), \quad 0 \leq \lambda \leq 1,$$

so that $\lambda = 1$ gives the curved reflector and $\lambda = 0$ gives its tangent line. For radar pose i , let $s_{i,n}^\lambda$ be the stationary point obtained by continuous continuation of the selected branch and define

$$H_n = \max_{1 \leq i \leq m} \sup_{0 \leq \lambda \leq 1} |s_{i,n}^\lambda|, \quad H_n \rightarrow 0.$$

Uniformly bounded curvature and Eq (3.7) give constants C_0, C_1 independent of i , λ , and n such that, for $|s| \leq H_n$ and writing $r_{0,n}$ for the tangent-line model,

$$\|r_{\lambda,n}(s) - r_{0,n}(s)\| \leq C_0 H_n^2, \quad \|\partial_s r_{\lambda,n}(s) - \partial_s r_{0,n}(s)\| \leq C_1 H_n.$$

Theorem 3.2 (Uniform small-footprint bias orders of the tangent-plane approximation). Assume that, for every n , pose i , and $0 \leq \lambda \leq 1$, the paired stationary points remain on a common interior, visible, single-bounce branch with a uniform endpoint margin; the propagation distances are bounded below by $d_0 > 0$; the normal-direction cosines are bounded below by $\nu_0 > 0$; and $\partial_{ss} L_{i,n}^\lambda(s_{i,n}^\lambda) \geq \mu_0 > 0$. Use a common local azimuth unwrap. Then, writing $H = H_n$, the curved and tangent-plane radar data satisfy the first two estimates in Eq (3.8). For the localization estimate, let $\mathcal{H}_{\lambda,n}$ denote the whitened augmented map with the same state, camera mapping, weights, and observation dimension, and let $d_n = \mathcal{H}_{1,n}(z_n^*)$ and $b_n = d_n - \mathcal{H}_{0,n}(z_n^*)$. If $\alpha_n = \sigma_{\min}(D\mathcal{H}_{0,n}(z_n^*)S) \geq \alpha_0 > 0$, the planar analogue of Eq (3.4) holds uniformly, and an unregularized local planar least-squares estimate $\hat{z}_{0,n}$ remains in that neighborhood with residual no larger than that at z_n^* , then its position component satisfies the final estimate in Eq (3.8):

$$\Delta\rho = O(H^2), \quad \Delta\phi = O(H), \quad \|\hat{x}_{\text{plane}} - x_\star\|_2 = O(H). \quad (3.8)$$

Proof. Set $\dot{r}_{\lambda,n} = \partial_\lambda r_{\lambda,n}$. Equation (3.7), bounded curvature, and $|s| \leq H_n$ imply uniformly

$$\|\dot{r}_{\lambda,n}\| = O(H_n^2), \quad \|\partial_s \dot{r}_{\lambda,n}\| = O(H_n).$$

Let $F_{i,\lambda,n}(s) = \partial_s L_{i,n}^\lambda(s)$. The uniform nondegeneracy condition and the implicit function theorem yield a differentiable paired branch with

$$\frac{ds_{i,n}^\lambda}{d\lambda} = -\frac{\partial_\lambda F_{i,\lambda,n}}{\partial_s F_{i,\lambda,n}} = O(H_n),$$

because $\partial_\lambda F_{i,\lambda,n}$ contains terms of order $O(H_n)$ from $\partial_s \dot{r}_{\lambda,n}$ and terms of order $O(H_n^2)$ divided by distances bounded below by d_0 , whereas the denominator is at least μ_0 . Let $e_{a,i}^\lambda$ and $e_{x,i}^\lambda$ denote the unit vectors from the reflection point toward the radar and target. Along the stationary branch, Fermat stationarity gives the envelope identity

$$\frac{d}{d\lambda} L_{i,n}^\lambda(s_{i,n}^\lambda) = \partial_\lambda L_{i,n}^\lambda = -(e_{a,i}^\lambda + e_{x,i}^\lambda)^\top \dot{r}_{\lambda,n} = O(H_n^2).$$

Integration over $0 \leq \lambda \leq 1$ proves the $O(H_n^2)$ range bound. For azimuth, let $p_{i,n}^\lambda = r_{\lambda,n}(s_{i,n}^\lambda) -$

$a_{i,n}$ and $(p_{i,n}^\lambda)^\perp = \left(-(p_{i,n}^\lambda)_2, (p_{i,n}^\lambda)_1\right)^\top$. On the common unwrap,

$$\frac{d\phi_{i,n}^\lambda}{d\lambda} = \frac{\left((p_{i,n}^\lambda)^\perp\right)^\top}{\|p_{i,n}^\lambda\|^2} \left(\dot{r}_{\lambda,n} + \partial_s r_{\lambda,n} \frac{ds_{i,n}^\lambda}{d\lambda}\right) = O(H_n).$$

The positive-distance bound and integration in λ give the $O(H_n)$ azimuth estimate. Because m , the observation dimension, and the whitening operators are fixed or uniformly bounded, the augmented data discrepancy satisfies $\|b_n\|_2 = O(H_n)$. Applying the planar local lower-slope and residual-comparison argument used for Eq (3.6) yields

$$\|S^{-1}(\hat{z}_{0,n} - z_n^*)\|_2 \leq \frac{4}{\alpha_n} \|b_n\|_2 = O(H_n).$$

Projection onto the position components proves the final estimate in Eq (3.8).

The range result is an envelope consequence of Fermat stationarity: Although the paired reflection coordinate changes by $O(H)$, that branch displacement does not enter the first variation of the stationary path length. By contrast, the $O(H)$ tangent-direction perturbation changes the stationarity equation and moves the reflection point by $O(H)$, which enters the arrival bearing at first order. Equation (3.8) states uniform upper-order bounds; special symmetries may yield higher-order cancellation. These conclusions remain branch-local and asymptotic and do not imply global branch uniqueness or convergence from arbitrary initial values.

4. Numerical solution method and numerical experiment design

This section describes how the forward model is evaluated within the admissible set defined in Section 2, how the prescribed reflection branch is tracked, and how the curvature-aware inverse problem is solved using a scaled Levenberg-Marquardt method. Figure 2 summarizes the complete workflow from multimodal observations to localization outputs and theoretical diagnostics.




Figure 2. Branch-aware inversion and numerical validation workflow for curvature-aware radar-camera NLoS localization.

4.1. Branch-aware forward evaluation

For each candidate state and radar pose, the algorithm first solves the Fermat stationarity equation over the prescribed interval I_i and selects the continuous local branch using a reference reflection coordinate. A candidate point must simultaneously satisfy the interior-point, same-side reflection, path-visibility, positive separation-distance, and positive Fermat-second-derivative conditions; if any condition fails, the state is deemed not to belong to the current admissible branch. After these checks are passed, the radar range and azimuth are assembled according to Eqs (2.38) and (2.40), and the reduced Jacobian is constructed using the implicit derivatives in Eqs (2.50)–(2.56). This procedure prevents arbitrary switching to another stationary point after each state update. An initial state is termed an admissible start if every radar pose passes these checks on its prescribed branch before LM iteration. Admissibility is a geometric pre-solver condition and does not imply numerical termination or recovery below the error threshold.

4.2. Scaled Levenberg-Marquardt inversion

At iteration k , let the whitened residual and scaled Jacobian be, respectively,

$$r_k = \mathcal{H}(z_k) - d, \quad J_k = D\mathcal{H}(z_k)S. \quad (4.1)$$

The Levenberg-Marquardt step Δu_k is obtained from

$$(J_k^T J_k + \mu_k D_k) \Delta u_k = -J_k^T r_k, \quad D_k = \text{diag}[\max(\text{diag}(J_k^T J_k), 10^{-8})], \quad \mu_k > 0. \quad (4.2)$$

The candidate state is

$$z_k^{\text{trial}} = \Pi_{Z_{\text{ad}}}(z_k + S\Delta u_k), \quad (4.3)$$

where $\Pi_{Z_{\text{ad}}}$ denotes projection onto the search domain and the physically admissible set, with an additional branch-validity check. If the candidate state remains on the prescribed reflection branch and yields a lower objective value, the step is accepted and the damping parameter is reduced; otherwise, the step is rejected and the damping parameter is increased. Numerical termination is determined by the gradient infinity norm, the normalized step length, and the maximum number of iterations. This procedure is a local solver and provides no guarantee of global convergence from arbitrary initial values.

4.3. Simulation design, comparison methods, and performance metrics

The baseline scenario comprises four effective radar poses, a single circular-arc reflector, and one fixed occluded target. The random seed, ground-truth geometry, sensor standard deviations, state scales, solver tolerances, and numbers of trials are listed in Table 2. All results were generated by the same computational pipeline; no plotting values were entered manually.

The baseline ground truth consists of a target position of (4.00, 4.20) m, a reflector orientation of 4° , and a curvature of 0.060 m^{-1} . All prescribed reflection points lie within the parameter interval. The minimum Fermat second derivative along the baseline paths is 0.063736 m^{-1} , and the minimum same-side reflection indicator is 0.63095; thus, all four paths satisfy the nondegeneracy and same-side reflection checks in this local scenario.

Table 2. Simulation, sensor, and numerical-solver parameters.

Parameter	Symbol	Value	Unit
Random seed	s_{rng}	20260724	—
Number of radar poses	m	4	—
Reflector half-width	h	3.00	m
Target position	x_*	(4.00, 4.20)	m
Reflector reference point	q_*	(0.00, 0.00)	m
Reflector orientation	ψ_*	4.00	°
Reflector curvature	κ_*	0.060	m ⁻¹
Radar range standard deviation	σ_ρ	0.040	m
Radar azimuth standard deviation	σ_ϕ	0.40	°
Camera position standard deviation	σ_q	0.040	m
Camera orientation standard deviation	σ_ψ	0.60	°
Camera curvature standard deviation	σ_κ	0.0080	m ⁻¹
Local-initialization standard deviation of target position	$\sigma_{x,0}$	(0.20, 0.18)	m
Local-initialization standard deviation of reflector position	$\sigma_{q,0}$	0.060	m
Local-initialization standard deviation of orientation	$\sigma_{\psi,0}$	1.00	°
Local-initialization standard deviation of curvature	$\sigma_{\kappa,0}$	0.0060	m ⁻¹
State-scaling vector	$\text{diag}(S)$	[1, 1, 0.4, 0.4, 0.0873, 0.04]	Mixed
Noise standard-deviation multiplier	η	0.5, 1.0, 1.5, 2.0, 3.0	—
Local-recovery error threshold	e_{loc}	1.0	m
Monte Carlo LM gradient tolerance	g_{tol}	2.0×10^{-6}	—
Monte Carlo LM step tolerance	x_{tol}	2.0×10^{-9}	—
Maximum number of Monte Carlo LM iterations	K_{max}	28	—
Number of Monte Carlo trials per noise multiplier	N_{MC}	400	—
Number of bootstrap resamples	N_{boot}	1,200	—

Note: The state-scaling vector contains scale factors for quantities with different physical dimensions: position, orientation, and curvature. The tabulated values define a representative synthetic scenario and are not calibration results for any specific radar or camera hardware. Figure 5 uses a curvature of 0.040 m^{-1} for the asymptotic-order experiment.

The local-identifiability experiment varies the horizontal radar-pose baseline and the camera-noise multiplier. The stability experiment uses 7 perturbation levels and 64 recovery trials per level, with 25 points sampled along each line segment joining the ground truth and the estimate. The tangent-plane bias experiment performs deterministic log–log regression over a small-footprint range. The noise-free trajectory is used solely to verify consistency among the forward model, Jacobian, and solver implementation.

The Monte Carlo comparison includes four estimators: curved radar-camera, planar radar-camera, curved range-only+camera, and an oracle reflector with known geometry. At each noise multiplier, 400 paired trials are performed. For a given trial index, all methods share the same standardized radar noise, camera noise, and local initial value; different noise multipliers are also obtained by scaling the same standardized noise realizations, thereby reducing the sampling variance of the estimated differences between methods.

Candidate initial states for the principal paired comparison were generated as $z_0 = \text{clip}(z_* + D_0\xi, \ell, u)$, where $\xi \sim N(0, I_6)$ has independent components and $D_0 = \text{diag}(0.20, 0.18, 0.060, 0.060, \pi/180, 0.0060)$ in the state order $(x_1, x_2, q_1, q_2, \psi, \kappa)$. Of 430 sequentially generated candidates, the first 400 satisfying the prescribed curved- and planar-branch validity checks were retained. Neither measurement noise nor any solver outcome was used in this screening; the main paired comparison is therefore conditional on initial branch validity.

Initialization sensitivity was assessed with 400 common, unclipped starts at each initialization-scale multiplier $c_{\text{init}} \in \{0.5, 1, 2, 4, 8\}$ under noiseless and nominal-noise observations. No start was resampled or removed; branch-invalid states remained in the unconditional denominator and were not passed to the local solver.

A nominal-noise cross-scenario screen used 12 geometries from a six-factor, two-level Plackett-Burman design [36]: $x_1 \in \{3.7, 4.3\}$ m, $x_2 \in \{3.9, 4.5\}$ m, $q_2 \in \{-0.1, 0.1\}$ m, $\psi \in \{3^\circ, 9^\circ\}$, $\kappa \in \{0.040, 0.075\} \text{ m}^{-1}$, and baseline $B \in \{1.6, 2.7\}$ m. All passed the interior, nondegeneracy, same-side, and full-rank checks; 100 trials per geometry compared the two radar-camera models with a camera-map VS-WLS reference adapted from [37]. VS-WLS used equal-weight ranges, a linear LS start, and bounded TRF. Common-start RMSE retained IDs admitted by both joint solvers; recovery retained all starts. Timings exclude data generation and I/O; interactions are not identified. All runtimes were measured within the same implementation and computing environment and are reported only for relative within-pipeline comparison, not as hardware-independent benchmarks.

For N trials, the positional root-mean-square error is defined as

$$\text{RMSE} = \sqrt{\frac{1}{N} \sum_{j=1}^N \|\hat{x}_j - x_*\|_2^2}. \quad (4.4)$$

Pointwise 95% intervals for the RMSE and paired between-method differences are constructed from the 2.5th and 97.5th percentiles of 1,200 bootstrap resamples; pointwise Wilson intervals are used for the solver failure rate. Strict solver success requires numerical termination under either the gradient or step criterion, a finite estimate, positive values of both the Fermat second derivative and the same-side indicator on the prescribed reflection branch, and no state component lying within an absolute tolerance of 10^{-8} of a search boundary. Local recovery additionally requires a position error below 1 m. The all-trial RMSE includes every finite output within the prescribed search bounds. Numerical termination, local recovery, and estimation error are therefore reported separately, so solver failures are neither interpreted directly as localization errors nor discarded.

Individual RMSE intervals are obtained by resampling trial IDs within each estimator. For comparisons between estimators, the same trial IDs are resampled jointly, and the resulting interval

is computed from the paired RMSE difference; it is not an RMSE interval for either individual estimator.

5. Numerical results and discussion

Using the same curvature-aware forward model throughout, this section examines geometric nondegeneracy, regular local identifiability, local stability under measurement perturbations, tangent-plane approximation bias, consistency of the local solver, and localization performance under different observation and geometric models. All figures and tables were generated using a fixed random seed and a single simulation pipeline, and no plotting values were entered manually. The analysis is restricted to the selected and associated single-specular-reflection branch and makes no claim of global path uniqueness.

5.1. Local identifiability and sensor complementarity

Figure 3 uses the range-only+camera configuration to illustrate how the horizontal radar-pose baseline B and the camera-noise multiplier jointly affect the regularity of the augmented observation model and the effective positional information. To ensure that a zero baseline represents genuinely colocated geometry, all four radar poses coincide exactly when $B=0$. In this case, the corresponding rows of the range Jacobian are repeated, the full Jacobian has rank 5 for every camera-noise multiplier, and the positional CRLB diverges. When $B>0$, pose diversity restores full local column rank in this scenario. Increasing the baseline or reducing the camera noise generally increases the minimum singular value and lowers the positional CRLB.




Figure 3. Effects of the horizontal radar-pose baseline and camera noise on local identifiability. (a) Phase diagram of the minimum singular value of the whitened augmented Jacobian; $B=0$ is exactly rank deficient. (b) CRLB RMSE associated with the effective positional information.

Figure 3 samples 41 equally spaced baselines over $0 \leq B \leq 3$ m, with $\Delta B = 0.075$ m. At each B , the four phase centers are generated about $(-4.02, 3.06)$ m using normalized horizontal offsets $(-1/2, -1/6, 1/6, 1/2)$ and a small B -scaled vertical dither. This range includes exact colocation,

the nominal horizontal aperture of 2.15 m, and a modest wider-aperture regime; it is a geometry-sensitivity range rather than a hardware requirement.

This trend demonstrates how the radar spatial aperture and camera geometric accuracy complement one another: Camera observations provide independent constraints along the reflector-parameter directions, whereas a nonzero radar baseline increases the geometric diversity of the target-dependent measurement directions. Nevertheless, the absolute magnitude of the minimum singular value depends on the whitening scheme and the choice of state scales. The contours in the figure are therefore conditional diagnostics for the present setting and should not be interpreted as universal physical thresholds.

To compare discrete observation configurations, Table 3 summarizes the number of scalar observations, full-state Jacobian rank, positional-information rank, minimum singular value, condition number, and positional-information diagnostic.

Table 3. Local-identifiability and information diagnostics for different observation configurations.

Observation configuration	Scalar observations	Jacobian rank	Position-information rank	$\sigma_{\min}$	Condition number	Positional information/CRLB RMSE (m)
Range-only (one pose)	1	1	0	0	∞	∞
Range-only (four poses)	4	4	1	0	∞	∞
Range+azimuth (four poses)	8	5	2	0	∞	34.671 [†]
Range-only+camera	8	6	2	3.691	15.376	0.269
Range+azimuth+camera	12	6	2	5.683	13.727	0.142

[†] This finite entry is a local positional-information diagnostic, not a full-state CRLB. As specified after Eq (2.72), a singular nuisance-information block is handled using the Moore-Penrose pseudoinverse, with interpretation restricted to the identifiable subspace; Eq (3.3) applies only when the corresponding positional information is nonsingular. The rank-5 Jacobian does not establish identifiability or a finite unconstrained CRLB for all six target-reflector parameters.

The discrete-configuration comparison shows that the Jacobian rank is only 1 for a single-pose range observation. Even when four poses are used, the rank of the range-only configuration remains 4. After radar azimuth is added, the effective positional-information rank reaches 2, but the full-state Jacobian rank remains 5, indicating one reflector-parameter direction that cannot be identified from radar data alone. Adding the camera yields a full column rank of 6 for the range-only+camera configuration, with a minimum singular value of 3.691 and a positional CRLB RMSE of 0.269 m. Adding azimuth further increases the minimum singular value to 5.683, reduces the condition number to 13.727, and lowers the positional CRLB RMSE to 0.142 m, an improvement of approximately 47.0% over the range-only+camera configuration. Thus, the camera primarily resolves the local coupling ambiguity between reflector geometry and target position, whereas azimuth further strengthens the positional information.

Four poses are sufficient only for the present six-parameter, camera-constrained configurations because the augmented Jacobian has full column rank. There is no geometry-independent pose-count

guarantee: Additional nonredundant poses may improve conditioning and robustness, whereas coincident or repeated poses do not increase rank.

5.2. Numerical evidence of local stability

Figure 4 examines the local scaling relationship between whitened measurement perturbations and inversion error. The simulation uses 7 perturbation levels and 64 trials at each level; for every trial, it samples 25 points along the line segment between the ground truth and the estimate to assess Jacobian variation. The empirical slope of the normalized state error with respect to perturbation magnitude is 0.9995, close to the first-order relationship expected under local Lipschitz stability.




Figure 4. Finite-sample assessment of local stability. (a) Empirical scaling of normalized state error with whitened perturbation magnitude and the conservative local upper bound. (b) Local-condition diagnostics based on the error-to-bound ratio, Jacobian-variation ratio, and normalized chord ratio.

Using $\alpha = \sigma_{\min}(\tilde{J}_*) = 5.6827$, the conservative bound constant $4/\alpha$ is 0.7039, and the maximum ratio of normalized state error to the upper bound over all trials is 0.1613. The maximum Jacobian-variation ratio within the sampled region is

$$\max\|\tilde{J} - \tilde{J}_*\|_2/\alpha = 0.4335,$$

which is below the adopted local-condition threshold of 0.5. The minimum segmentwise singular-value ratio and minimum mean-Jacobian singular-value ratio are 0.9965 and 0.9982, respectively, and the minimum normalized chord ratio is 1.1765. These results show that, within the sampled small-perturbation neighborhood, the finite samples are consistent with the theoretically derived local-stability bound. The coefficient $4/\alpha$, the normalized error-to-bound ratio, and the Jacobian-variation ratio are sufficient, neighborhood-dependent diagnostics under the selected branch, state scaling, and whitening; they are neither necessary conditions nor universal thresholds across geometries or sensors. Likewise, the positional CRLB is a local lower bound under the stated Gaussian model and local-unbiasedness assumptions, rather than a universal error threshold. These numerical results therefore provide only empirical evidence of local stability; they neither replace a theoretical proof nor extrapolate to arbitrarily large perturbations, grazing geometries, or reflection-branch switching.

5.3. Scaling order of tangent-plane approximation bias

Figure 5 compares the forward-data and localization biases induced by replacing the true circular-arc reflector with its tangent-plane approximation at the reference point. Because Figure 5 is an asymptotic order test rather than a repetition of the nominal scenario, $\kappa = 0.040 \text{ m}^{-1}$ was used so that the dimensionless expansion parameter satisfies $|\kappa|H \leq 0.020$ over the fitted interval $H \leq 0.50 \text{ m}$, thereby suppressing higher-order terms in Eq (3.7); the performance experiments retain the nominal $\kappa = 0.060 \text{ m}^{-1}$. For $H \leq 0.50 \text{ m}$, the log-log ordinary-least-squares regression orders for the range discrepancy, azimuth discrepancy, and joint localization bias are, respectively,

$$p_\rho = 2.0047, \quad p_\phi = 0.9940, \quad p_x = 0.9881.$$




Figure 5. Scaling order of tangent-plane approximation bias. (a) The range discrepancy is approximately second order; (b) the azimuth discrepancy is approximately first order; and (c) the joint range-azimuth localization bias is approximately first order. The shaded region denotes the small-footprint range used to estimate the slopes. The slopes are fitted over $H \leq 0.50 \text{ m}$ using $\kappa = 0.040 \text{ m}^{-1}$.

The coefficients of determination for the three regressions are 0.999999, 0.999996, and 0.999982, respectively. These values support the following local scaling relations:

$$\Delta\rho = O(H^2), \quad \Delta\phi = O(H), \quad \|\hat{x}_{\text{plane}} - x_\star\|_2 = O(H).$$

This distinction shows that the first-order geometric term in the range measurement can cancel under the Fermat stationarity condition, causing the tangent-plane range bias to exhibit approximately second-order scaling. By contrast, azimuth remains directly affected by changes in the surface normal; consequently, a first-order term ultimately dominates the general joint range-azimuth localization bias. Comparing range residuals alone may therefore underestimate the effect of curvature-model mismatch on final localization. The slope intervals shown in the figure are regression intervals for deterministic numerical points, not confidence intervals from repeated measurements or Monte Carlo trials. The inferred orders are also restricted to small H , a fixed nondegenerate reflection branch, and nongrazing geometry.

5.4. Local-solver consistency and initialization sensitivity

Figure 6 presents a Levenberg-Marquardt (LM) solution trajectory for noise-free, perfectly model-matched data when the initial value lies on the prescribed local reflection branch. The objective function, normalized state error, and gradient infinity norm all decrease over the accepted iterations. After 7 accepted iterations, the algorithm yields

$$\|S^{-1}(\hat{z} - z_*)\|_2 = 4.85 \times 10^{-12}.$$

The final gradient infinity norm is 1.15×10^{-10} . This trajectory confirms the numerical consistency of the forward model, Jacobian, and local solver for this test case.




Figure 6. A local LM trajectory for noise-free, model-matched data. (a) Objective function; (b) normalized state-recovery error; and (c) gradient infinity norm and damping parameter. This figure serves only as a local numerical correctness check.

A single noise-free trajectory is insufficient to establish a general convergence order, a global basin of attraction, or robustness to arbitrary initial values. Figure 6 therefore provides a local numerical check on the code and derivative implementation; it is not an experimental substitute for a general theory of convergence.

Figure 7 shows that unconditional recovery falls from 100% to 8% as c_{init} increases, whereas all initially admissible starts converge with nearly unchanged strict-success RMSE; the experiment therefore characterizes a local admissible region rather than a global basin of attraction.

Table 4 summarizes the principal theoretical and numerical diagnostics corresponding to Figures 1 and 3–6.

Table 4 presents mutually consistent results across geometric nondegeneracy, local information, compatibility with the stability bound, tangent-plane model bias, and solver implementation. Its primary value lies in providing cross-checks based on a common forward model: a positive Fermat second derivative ensures local nondegeneracy of the selected branch; a positive minimum singular value of the scaled augmented Jacobian provides a diagnostic of regular local identifiability; the near-first-order stability scaling and the error-to-bound ratio below unity support the local perturbation analysis; and the bias orders clarify that range and azimuth errors induced by planarization do not have the same scaling.




Figure 7. Empirical initialization-sensitivity region as c_{init} varies. (a) Unconditional local recovery; (b) strict-success RMSE and 95th-percentile error; and (c) mutually exclusive outcomes.

Table 4. Principal theoretical and numerical diagnostics.

Diagnostic metric	Symbol	Value	Unit
Minimum Fermat second derivative along the baseline paths	$\min_i \partial_{ss} L_i$	0.063736	m^{-1}
Minimum singular value of the scaled augmented Jacobian	$\alpha = \sigma_{\min}(\tilde{J}_*)$	5.682686	—
Jacobian rank in the zero-baseline phase diagram	$\text{rank}(\tilde{J}_*) _{B=0}$	5	—
Maximum sampled Jacobian-variation ratio	$\max \ \tilde{J} - \tilde{J}_*\ _2 / \alpha$	0.433507	—
Empirical stability-scaling order	p_{stab}	0.999545	—

Continued on next page

Diagnostic metric	Symbol	Value	Unit
Maximum normalized state-error-to-bound ratio	$\max\ u\ _2/(4\ \delta\ _2/\alpha)$	0.161270	—
Order of tangent-plane range bias	p_ρ	2.004726	—
Order of tangent-plane azimuth bias	p_ϕ	0.993966	—
Order of joint planar-localization bias	p_x	0.988131	—
Final normalized state error of LM	$\ S^{-1}(\hat{z} - z_*)\ _2$	4.85×10^{-12}	—

Note: The stability indices are numerical diagnostics based on finite local sampling. The bias orders are slope estimates obtained by deterministic log–log ordinary least squares over $H \leq 0.50$ m.

5.5. Localization performance in paired Monte Carlo trials

Figure 8 compares four estimators: joint curved radar-camera estimation, joint planar radar-camera estimation, a curved range-only+camera ablation, and an oracle reflector with known geometry. The Jacobian rank of the full-state radar-only configuration is 5, so the state is not uniquely identifiable without an added prior; this configuration is therefore excluded from the RMSE ranking. By contrast, the curved range-only+camera configuration is a full-rank, identifiable ablation rather than a radar-only method.




Figure 8. Overall localization results for four estimators. (a) Sensitivity to paired noise with 95% bootstrap intervals; (b) strict solver failure rate with 95% Wilson intervals; and (c) empirical cumulative distributions of position error for 400 paired trials at the nominal noise level.

At each noise multiplier, the experiment includes 400 paired trials. For each trial index, all methods share the same standardized radar noise, camera noise, and local initial value; the different noise multipliers are generated by scaling the same standardized noise realizations. The 400 initial

values were selected from 430 candidates according to local branch-validity conditions common to the curved and planar models, giving an acceptance rate of approximately 93.0%. This screening used neither measurement errors nor method outcomes.

Figure 8(a) shows that the all-trial RMSE of every method increases with the noise multiplier. For the curved radar–camera method, the RMSE increases from 0.072 m at $0.5\times$ noise to 0.434 m at $3\times$ noise. At $3\times$ noise, the RMSE values for planar radar–camera, curved range-only + camera, and the oracle reflector are 0.477, 0.521, and 0.085 m, respectively. With common random numbers and nested noise multipliers, the oracle error also increases monotonically with noise, eliminating the small-sample artifact in which performance could appear to improve at a higher noise level.

In terms of numerical reliability, the strict solver failure rate of the curved radar–camera method does not exceed 0.25% at any noise multiplier. By contrast, the failure rate of the curved range-only + camera method rises to 18.0%, 35.25%, and 56.75% at noise multipliers of 1.5, 2.0, and 3.0, respectively. Thus, although the camera information gives the range-only configuration a full-rank Jacobian near the ground truth, radar azimuth remains important for numerical stability at high noise levels. At nominal noise, the empirical cumulative distributions also show that the curved radar-camera error distribution lies generally to the left of those for the planar and range-only ablation methods. The oracle reflector provides an ideal reference with completely known reflector geometry and is not directly deployable. The strict solver failure rate and local recovery rate are not equivalent. For example, at $3\times$ noise, the planar method has a solver failure rate of 0% but a local recovery rate of 96.25%; numerical termination must therefore not be equated with successful localization.

Localization errors, interval estimates, and numerical reliability at the nominal noise level are reported in Table 5.

Table 5. Localization performance at the nominal noise level (400 paired trials).

Method	N	All-trial RMSE [95% CI] (m)	Conditional RMSE over converged trials (m)	Median (m)	P95 (m)	Solver failure (%)	Local recovery rate (%)
Curved radar-camera	400	0.145 [0.137, 0.153]	0.145	0.112	0.266	0.0	100.0
Planar radar-camera	400	0.199 [0.187, 0.211]	0.199	0.145	0.368	0.5	99.5
Curved range-only+camera	400	0.262 [0.246, 0.279]	0.261	0.187	0.479	3.0	97.0
Oracle reflector	400	0.028 [0.027, 0.030]	0.028	0.024	0.047	0.0	100.0

Note: The all-trial RMSE includes all finite outputs within the search bounds; the conditional RMSE over converged trials includes only runs that terminate numerically, remain on a physically valid branch, and do not reach a boundary. Local recovery additionally requires a position error below 1 m. Solver failures are therefore neither discarded nor conflated with estimation error.

At the nominal noise level, the all-trial RMSE for the curved radar-camera method is 0.145 m, with a 95% bootstrap interval of [0.137, 0.153] m. Its median and 95th-percentile errors are 0.112 and 0.266 m, respectively; the solver failure rate is 0%, and the local recovery rate is 100%. The corresponding RMSE for the planar radar–camera method is 0.199 [0.187, 0.211] m, with median and 95th-percentile errors of 0.145 and 0.368 m. Thus, relative to the planar model, the curved model

reduces the RMSE, median error, and 95th-percentile error by approximately 27.0%, 22.9%, and 27.7%, respectively.

The paired planar-minus-curved RMSE difference is 0.0537 m, with a 95% bootstrap interval of [0.0444, 0.0625] m; the difference is positive in all 1,200 resamples. The RMSE of curved range-only+camera is 0.262 [0.246, 0.279] m. Adding azimuth further reduces the RMSE by 0.1170 m, with a paired interval of [0.1002, 0.1328] m. The oracle reflector RMSE is 0.028 [0.027, 0.030] m. The approximately 0.117 m gap between curved joint estimation and the oracle reflector reflects the performance cost of uncertainty in the reflector estimate and its coupling with the target position.

5.6. Cross-scenario robustness and computational cost

Within the deterministic 12-geometry, two-level Plackett-Burman screen, the curvature-aware method gave the lowest common-start RMSE in every tested geometry: a scenario median of 0.145 m versus 0.364 m for tangent-plane radar-camera and 0.754 m for camera-map VS-WLS (Table 6). This result is descriptive of the tested design points; interactions are confounded, and no extrapolation beyond the tested factor levels is intended.

All 1,023 branch-admissible starts for the curvature-aware method converged to errors below 1 m; thus, the method's recovery range chiefly reflects branch admissibility. Tangent-plane minima of 0% in S1 and S12 arose when finite-difference Jacobians left the branch despite finite estimates below 1 m. VS-WLS omits azimuth; timing is implementation-specific.

Table 6. Cross-scenario localization accuracy, recovery, and computational cost over 12 fixed geometries under nominal noise (100 attempted trials per geometry).

Method	Paired RMSE (m)*	Recovery (%)*	Runtime/attempt (ms)†	Accepted LM steps†
Curvature-aware radar-camera	0.145 [0.132, 0.196]	92.0 [56.0, 99.0]	180.6 [149.4, 212.6]	5.0 [5.0, 6.0]
Tangent-plane radar-camera‡	0.364 [0.195, 1.088]	99.5 [0.0, 100.0]	121.0 [106.7, 151.8]	4.0 [4.0, 5.0]
Camera-map VS-WLS [37]§	0.754 [0.374, 1.957]	80.5 [16.0, 100.0]	9.3 [7.6, 11.6]	—

Note: * Scenario median [minimum, maximum]; † pooled median [IQR]. RMSE uses common trial IDs per geometry; recovery uses all starts and requires numerical termination, a finite physical solution off the boundary, and error below 1 m. Runtime is implementation-specific and is reported only for relative comparison within the present research implementation, not as a hardware-independent benchmark. Accepted steps apply only to LM; VS-WLS used TRF with 12 [10, 15] residual evaluations (median [IQR]), so '—' denotes not applicable. ‡ Matched-input comparator using the same observations, trial IDs, noise realizations, and initial values as the proposed method; this comparison isolates curvature. § Literature-aligned external reference with a different reflector representation and numerical procedure; it provides an end-to-end reference and does not isolate curvature.

5.7. Overall discussion and scope of applicability

Together, the results support three local conclusions. First, the transition from exact rank deficiency at zero baseline to full rank at a positive baseline shows that repeated measurements cannot

replace geometric diversity among poses; camera-based geometric observations and the radar spatial aperture must complement one another to resolve the local ambiguity between target position and reflector parameters. Second, the second-order cancellation of tangent-plane range bias does not automatically extend to azimuth or final localization. The main benefit of curvature modeling therefore extends beyond reducing path-length residuals to suppressing first-order systematic mismatch in directional information. Third, by reporting both estimation error and solver failure rate, the paired Monte Carlo results show that the curved radar-camera estimator provides lower localization error and greater numerical reliability in the synthetic setting considered here.

The sensitivity assessment covers several complementary directions. Figure 3 varies the horizontal radar-pose baseline B together with the camera-noise multiplier and reports their effects on the regularity of the augmented Jacobian and the positional CRLB. Here, B denotes the spatial aperture spanned by the four effective radar phase centers; it is not the physical mounting separation between the radar and camera, whose coordinate transformation is assumed calibrated and fixed under Assumption A1. Figure 7 evaluates initialization sensitivity under both noise-free and nominal-noise observations, whereas Figure 8 varies the common sensor-noise multiplier from 0.5 to 3.0 using 400 paired trials at each level. Table 6 extends the assessment to 12 fixed Plackett-Burman geometries, including two levels of reflector curvature, target position, reflector position and orientation, and radar-pose baseline. These experiments provide local synthetic sensitivity evidence only within the tested ranges. The two-level screen does not separately identify interaction effects, and the physical radar-camera mounting baseline and extrinsic-calibration error are not varied in the present study.

The assumptions delimit both the class of recoverable scenes and the validity of the theoretical conclusions. The calibrated two-dimensional formulation represents a ground-plane cross-section and therefore excludes elevation, out-of-plane propagation, and general three-dimensional surface geometry. The finite-dimensional reflector assumption restricts recovery to the prescribed curve family; it does not imply reconstruction of an arbitrary free-form or nonsmooth vehicle body, and adding shape parameters may alter identifiability and conditioning. Discrepancy between the actual surface and the selected basis instead appears as model bias. Known single-bounce association makes the inversion conditional on successful path preprocessing and excludes unknown associations, additional reflections, diffuse returns, and ghost paths; an incorrect association may change the branch or propagation topology and cannot generally be treated as a small measurement perturbation.

Nongrazing incidence and Fermat-branch nondegeneracy ensure that the selected reflection point is locally well defined and differentiable. Near grazing incidence, branch degeneracy, a segment endpoint, or a visibility transition, the reflection-point sensitivity may become large, and branch switching may invalidate the smooth reduced mapping and its local identifiability and stability bounds. The principal LM comparison is also conditional on branch-valid local initializations. Figure 7 characterizes this dependence but describes a local admissible region rather than a global basin of attraction; branch validity alone does not guarantee numerical convergence.

The 12-scene, six-factor Plackett-Burman screen remains local: Interactions are confounded and extrapolation beyond the tested region is unsupported. The paired-method comparison uses 400 screened starts, whereas the initialization-sensitivity experiment retains all generated starts in its unconditional denominator. Accordingly, this study is a local synthetic geometric benchmark under Assumptions 2.1–2.10, not a general guarantee for unrestricted road scenes, proof of global identifiability, or vehicle-mounted experimental validation.

The comparison methods address different sources of performance gain. The matched-input tangent-plane radar-camera estimator uses the same observations, paired noise realizations, and initial values as the curvature-aware estimator and therefore provides the closest controlled comparison for

isolating the effect of replacing the curved reflector with its tangent-plane approximation. The curved range-only+camera configuration quantifies the contribution of radar azimuth, whereas the oracle-reflector estimator indicates the accuracy attainable when reflector geometry is known and is not directly deployable. Table 6 additionally evaluates a camera-map virtual-station weighted least-squares method adapted from [37] as a literature-aligned NLoS reference. Because VS-WLS omits azimuth and differs in reflector representation and numerical procedure, this comparison assesses end-to-end performance under common scenarios rather than the effect of curvature alone.

The curvature-aware method achieved the lowest common-start RMSE in all 12 tested geometries, but this result does not establish superiority over all NLoS localization methods. In particular, the numerical experiments instantiate circular-arc reflectors only. Although the theoretical formulation in Section 2.1.2 permits other finite-dimensional regular curve families, including local polynomials and constrained B-splines, no empirical superiority over alternative curved-reflector parameterizations is claimed. Systematic comparisons involving different curve bases, model dimensions, regularization strategies, camera observability, and basis mismatch remain subjects for future work.

Practical three-dimensional deployment requires three extensions:

(i) Geometry and regularity: The two-dimensional curve and target state must become local surface patches and a three-dimensional state in a calibrated, time-aligned sensor frame; Fermat stationarity then has two tangential equations, and branch regularity and local identifiability require a positive-definite 2×2 surface Hessian and a full-rank three-dimensional augmented Jacobian [38].

(ii) Path association: Known single-bounce association must be replaced by a multi-hypothesis front end combining camera contours or depth with radar range, angle, Doppler, visibility, and temporal consistency [14,16,18–20,39].

(iii) Nonsmooth surfaces and multipath: C^3 regularity applies only patch-wise to smooth surfaces; edges, rough or diffuse regions, multiple reflections, and branch changes require piecewise models or rejection.

These extensions are not implemented in the present study. Accordingly, this study is a local geometric benchmark, not a deployable system. Randomized three-dimensional simulations and synchronized vehicle tests should next quantify association, calibration/timing sensitivity, localization, failure coverage, and latency.

6. Conclusions

This study developed a curvature-aware inverse-geometric framework for camera-assisted mmWave NLoS localization under a high-frequency, single-bounce specular-reflection model. Its principal methodological contribution is the implicit differentiation of branch-selected reflection points on a finite-dimensional curved reflector and the use of the resulting reduced derivatives to analyze joint target-reflector recovery. The Levenberg-Marquardt algorithm and Fisher information are established tools used for numerical inversion and local information assessment, respectively, and are not claimed as independent innovations. Fermat-branch nondegeneracy makes the reduced map locally smooth; full column rank of its augmented Jacobian gives regular local identifiability; Fisher-Schur information quantifies local position information after eliminating reflector nuisance directions; and local stability additionally requires bounded Jacobian variation and small branch-preserving perturbations. These implications neither establish global uniqueness nor guarantee solver convergence.

The strengthened tangent-plane analysis defines a uniform small-footprint family with paired curved and planar branches and proves an $O(H^2)$ range discrepancy together with $O(H)$ azimuth

and generic localization-bias upper bounds; the localization statement is conditional on local inverse stability. The numerical study examines radar-pose aperture, camera accuracy, initialization scale, sensor-noise level, observation ablations, and 12 fixed geometries. Across 400 paired synthetic trials at nominal noise, curved radar-camera localization attained an RMSE of 0.145 m versus 0.199 m for the planar model, and it had the lowest common-start RMSE in all 12 screened geometries. These findings support explicit curvature modeling within the prescribed local synthetic setting, but they do not establish superiority over every NLoS method or curved-reflector basis, arbitrary-initialization convergence, full electromagnetic validity, or road-vehicle performance.

These conclusions are limited to calibrated two-dimensional geometry, finite-dimensional circular-arc curves, known association of a single reflection, 12 deterministic two-level screening geometries, a nongrazing local branch, and local initial values screened for branch validity. The study does not address full electromagnetic scattering, material and polarization effects, diffuse reflection, multiple reflections, unknown data association, three-dimensional free-form surfaces, dynamic synchronization errors, or non-Gaussian correlated noise. Accordingly, the cross-scenario ranking is descriptive of the tested design points and does not support extrapolation to untested geometries or interaction effects.

Operational validation will require a three-dimensional surface formulation, explicit multipath association, analysis of calibration and timing errors, and synchronized vehicle-mounted camera–radar experiments.

Use of Generative-AI tools declaration

The author declares he has not used Artificial Intelligence (AI) tools in the creation of this article.

Acknowledgments

The author thanks the National Science and Technology Council, Taiwan, for financial support of this research (grant NSTC 115-2221-E-018-009).

Conflict of interest

The author declares no conflict of interest.

References

1. S. L. Lin, An edge-assisted Internet-of-Vehicles computing framework for fair tail-risk allocation in cooperative autonomous driving, *Comput. Mater. Continua*, **88** (2026), 11. <https://doi.org/10.32604/cmc.2026.084571>
2. S. L. Lin, Nonlinear stability and critical penetration analysis of mixed autonomous traffic with anticipative optimal-velocity control, *AIMS Math.*, **11** (2026), 15302–15323. <https://doi.org/10.3934/math.2026629>
3. S. L. Lin, Y. H. Chen, Deep learning–aided frequency-modulated continuous-wave radar for around-the-corner non-line-of-sight perception at urban intersections, *Comput. Model. Eng. Sci.*, **147** (2026), 37. <https://doi.org/10.32604/cmes.2026.078862>

4. S. L. Lin, Advanced multi-channel echo separation techniques for high-interference automotive radars, *Comput. Mater. Continua*, **85** (2025), 1365–1382. <https://doi.org/10.32604/cmc.2025.067764>
5. W. T. Hsu, S. L. Lin, Using FMCW in autonomous cars to accurately estimate the distance of the preceding vehicle, *Int. J. Auto. Tech.*, **23** (2022), 1755–1762. <https://doi.org/10.1007/s12239-022-0153-4>
6. A. Sume, M. Gustafsson, M. Herberthson, A. Jänis, S. Nilsson, J. Rahm, Radar detection of moving targets behind corners, *IEEE Trans. Geosci. Remote Sens.*, **49** (2011), 2259–2267. <https://doi.org/10.1109/TGRS.2010.2096471>
7. P. Setlur, G. E. Smith, F. Ahmad, M. G. Amin, Target localization with a single sensor via multipath exploitation, *IEEE Trans. Aerosp. Electron. Syst.*, **48** (2012), 1996–2014. <https://doi.org/10.1109/TAES.2012.6237575>
8. O. Rabaste, E. Colin-Koeniguer, D. Poullin, A. Cheraly, J. F. Pétex, H. K. Phan, Around-the-corner radar: detection of a human being in non-line of sight, *IET Radar Sonar Nav.*, **9** (2015), 660–668. <https://doi.org/10.1049/iet-rsn.2014.0337>
9. R. Zetik, M. Eschrich, S. Jovanoska, R. S. Thomä, Looking behind a corner using multipath-exploiting UWB radar, *IEEE Trans. Aerosp. Electron. Syst.*, **51** (2015), 1916–1926. <https://doi.org/10.1109/TAES.2015.140303>
10. K. P. H. Thai, O. Rabaste, J. Bosse, D. Poullin, I. D. H. Sáenz, T. Letertre, Detection-localization algorithms in the around-the-corner radar problem, *IEEE Trans. Aerosp. Electron. Syst.*, **55** (2019), 2658–2673. <https://doi.org/10.1109/TAES.2019.2897031>
11. X. Q. Yang, S. H. Fan, S. S. Guo, S. L. Li, G. L. Cui, W. Zhang, NLOS target localization behind an L-shaped corner with an L-band UWB radar, *IEEE Access*, **8** (2020), 31270–31286. <https://doi.org/10.1109/ACCESS.2020.2973046>
12. S. L. Li, S. S. Guo, J. H. Chen, X. Q. Yang, S. H. Fan, C. Jia, Multiple targets localization behind L-shaped corner via UWB radar, *IEEE Trans. Veh. Technol.*, **70** (2021), 3087–3100. <https://doi.org/10.1109/TVT.2021.3068266>
13. Q. Tang, J. Li, L. Y. Wang, Y. Jia, G. L. Cui, Multipath imaging for NLOS targets behind an L-shaped corner with single-channel UWB radar, *IEEE Sens. J.*, **22** (2022), 1531–1540. <https://doi.org/10.1109/JSEN.2021.3131665>
14. J. H. Chen, S. S. Guo, H. L. Luo, N. Li, G. L. Cui, Non-line-of-sight multi-target localization algorithm for driver-assistance radar system, *IEEE Trans. Veh. Technol.*, **72** (2023), 5332–5337. <https://doi.org/10.1109/TVT.2022.3227971>
15. Z. Y. Liu, W. L. Zhang, J. Y. Zheng, S. S. Guo, G. L. Cui, L. J. Kong, Non-LOS target localization via millimeter-wave automotive radar, *J. Syst. Eng. Electron.*, **34** (2023), 1171–1181. <https://doi.org/10.23919/JSEE.2023.000070>
16. D. Kang, D. Kum, Camera and radar sensor fusion for robust vehicle localization via vehicle part localization, *IEEE Access*, **8** (2020), 75223–75236. <https://doi.org/10.1109/ACCESS.2020.2985075>
17. Y. D. Li, D. H. Zhang, R. X. Geng, Z. Lu, Z. Wu, Y. Hu, et al., A high-resolution handheld millimeter-wave imaging system with phase error estimation and compensation, *Commun. Eng.*, **3** (2024), 4. <https://doi.org/10.1038/s44172-023-00156-2>
18. C. Gentner, T. Jost, W. Wang, S. W. Zhang, A. Dammann, U. C. Fiebig, Multipath assisted positioning with simultaneous localization and mapping, *IEEE Trans. Wirel. Commun.*, **15** (2016), 6104–6117. <https://doi.org/10.1109/TWC.2016.2578336>

19. E. Leitinger, P. Meissner, C. Rüdiger, G. Dumphart, K. Witrisal, Evaluation of position-related information in multipath components for indoor positioning, *IEEE J. Sel. Areas Commun.*, **33** (2015), 2313–2328. <https://doi.org/10.1109/JSAC.2015.2430520>
20. E. Leitinger, A. Venus, B. Teague, F. Meyer, Data fusion for multipath-based SLAM: combining information from multiple propagation paths, *IEEE Trans. Signal Process.*, **71** (2023), 4011–4028. <https://doi.org/10.1109/TSP.2023.3310360>
21. H. Godrich, A. M. Haimovich, R. S. Blum, Target localization accuracy gain in MIMO radar-based systems, *IEEE Trans. Inform. Theory*, **56** (2010), 2783–2803. <https://doi.org/10.1109/TIT.2010.2046246>
22. Y. Shen, M. Z. Win, Fundamental limits of wideband localization—Part I: A general framework, *IEEE Trans. Inform. Theory*, **56** (2010), 4956–4980. <https://doi.org/10.1109/TIT.2010.2060110>
23. M. O’Toole, D. B. Lindell, G. Wetzstein, Confocal non-line-of-sight imaging based on the light-cone transform, *Nature*, **555** (2018), 338–341. <https://doi.org/10.1038/nature25489>
24. K. Levenberg, A method for the solution of certain non-linear problems in least squares, *Quart. Appl. Math.*, **2** (1944), 164–168. <https://doi.org/10.1090/qam/10666>
25. D. W. Marquardt, An algorithm for least-squares estimation of nonlinear parameters, *J. Soc. Indust. Appl. Math.*, **11** (1963), 431–441. <https://doi.org/10.1137/0111030>
26. S. H. Son, K. J. Lee, W. K. Park, Real-time tracking of moving objects from scattering matrix in real-world microwave imaging, *AIMS Math.*, **9** (2024), 13570–13588. <https://doi.org/10.3934/math.2024662>
27. S. Kang, W. K. Park, A novel study on the bifocusing method for imaging unknown objects in two-dimensional inverse scattering problem, *AIMS Math.*, **8** (2023), 27080–27112. <https://doi.org/10.3934/math.20231386>
28. W. K. Park, On the identification of small anomaly in microwave imaging without homogeneous background information, *AIMS Math.*, **8** (2023), 27210–27226. <https://doi.org/10.3934/math.20231392>
29. J. Eom, W. K. Park, Real-time detection of small objects in transverse electric polarization: evaluations on synthetic and experimental datasets, *AIMS Math.*, **9** (2024), 22665–22679. <https://doi.org/10.3934/math.20241104>
30. B. W. Tang, X. Y. Yang, L. Su, Shape reconstruction of acoustic obstacle with linear sampling method and neural network, *AIMS Math.*, **9** (2024), 13607–13623. <https://doi.org/10.3934/math.2024664>
31. P. Wang, J. J. Huang, W. J. He, J. Q. Zhang, F. Guo, Maximum likelihood DOA estimation based on improved invasive weed optimization algorithm and application of MEMS vector hydrophone array, *AIMS Math.*, **7** (2022), 12342–12363. <https://doi.org/10.3934/math.2022685>
32. I. M. Nasir, H. Alshaya, S. Tehsin, W. Bouchelligua, Cross-view remote sensing and street-level data fusion for intelligent traffic congestion analysis, *AIMS Math.*, **11** (2026), 1547–1589. <https://doi.org/10.3934/math.2026065>
33. D. Y. Zhu, Y. T. Yang, M. Y. Cao, An accelerated adaptive two-step Levenberg-Marquardt method with the modified Metropolis criterion, *AIMS Math.*, **9** (2024), 24610–24635. <https://doi.org/10.3934/math.20241199>
34. L. Y. Zhao, J. Y. Tang, Convergence properties of a family of inexact Levenberg-Marquardt methods, *AIMS Math.*, **8** (2023), 18649–18664. <https://doi.org/10.3934/math.2023950>

35. E. Alruwaili, Recovery of a time-dependent reaction coefficient in a space-time fractional diffusion model via Nyström discretization, *AIMS Math.*, **11** (2026), 13196–13215. <https://doi.org/10.3934/math.2026544>
36. R. L. Plackett, J. P. Burman, The design of optimum multifactorial experiments, *Biometrika*, **33** (1946), 305–325. <https://doi.org/10.1093/biomet/33.4.305>
37. D. L. Liu, Y. Wang, P. He, Y. Zhai, H. Y. Wang, TOA localization for multipath and NLOS environment with virtual stations, *EURASIP J. Wirel. Commun. Netw.*, **2017** (2017), 104. <https://doi.org/10.1186/s13638-017-0896-1>
38. E. Wise, Q. L. Cheng, J. Kelly, Spatiotemporal calibration of 3-D millimetre-wavelength radar-camera pairs, *IEEE Trans. Robot.*, **39** (2023), 4552–4566. <https://doi.org/10.1109/TRO.2023.3311680>
39. L. Zheng, J. M. Long, M. Lops, F. Liu, X. Y. Hu, C. H. Zhao, Detection of ghost targets for automotive radar in the presence of multipath, *IEEE Trans. Signal Process.*, **72** (2024), 2204–2220. <https://doi.org/10.1109/TSP.2024.3384750>



AIMS Press

© 2026 the Author(s), licensee AIMS Press. This is an open access article distributed under the terms of the Creative Commons Attribution License (<https://creativecommons.org/licenses/by/4.0>)