PROTECTING DIRECTION WHILE HIDING POSITION: SCALE-INVARIANT DP ON THE QUATERNION SPHERE

Anonymous authors
Paper under double-blind review

ABSTRACT

Position and direction carry different privacy value, yet trajectory differential privacy (DP) adds isotropic noise that destroys both at once. A published flight path exposes the operator's position, while the downstream tasks that consume these data, formation control and maneuver classification, depend on the direction. We decompose the trajectory into step lengths and unit headings, and add directional noise on the unit sphere: von Mises–Fisher on $\mathbb{S}^2$ for headings and Bingham on the quaternion sphere $\mathbb{S}^3$ for orientations. Direction-Aware Differential Privacy (DDP) obeys a closed-form (α, ε)-Rényi DP bound that depends only on the concentration κ and the angular sensitivity θ, and never on the radial scale: directional privacy is scale-invariant. Under a mean-direction (vMF) or antipodal (Bingham) constraint, the noise is maximum-entropy. DDP is minimax optimal: it attains $\Theta(\theta/\sqrt{\varepsilon})$ angular error and matches a minimax lower bound up to constant factors. Finally, a streaming variant (SDDP) releases each point as it arrives with an $O(1)$ privacy cost in the trajectory length and an $O(\sqrt{T})$ position drift. Because direction and step length are decoupled, DDP preserves the shape of a path while obscuring its exact coordinates, measuring privacy on the direction that downstream tasks consume.

1 INTRODUCTION

Unmanned aerial vehicle (UAV) formations and ground-vehicle fleets publish trajectory data for airspace management, traffic optimization, and fleet analytics. In a leader–follower formation (Figure 1) (Fax & Murray, 2004; Olfati-Saber et al., 2007; Ren et al., 2007), the leader's *heading* is the reference input that every follower uses to maintain formation, whereas the leader's *exact position* is the sensitive quantity that an adversary wants (reconnaissance target, mission objective). Differentially private formation and multi-agent control (Cortés et al., 2016; Hawkins & Hale, 2020; 2026) protects agent states with Euclidean noise and does not separate direction from position. This raises a question that existing trajectory DP mechanisms do not answer:

> *Can we protect the position while preserving the direction, and is directional privacy inherently scale-invariant?*

Existing trajectory DP mechanisms (Andrés et al., 2013) add isotropic noise to every coordinate, which destroys directional structure: a privatized heading is a difference of two noisy points, so its error is dominated by noise even when the position error is small. Synthesis-based methods (Wang et al., 2023) generate entirely new trajectories rather than privatizing existing ones. None of them decomposes a trajectory into direction and step length, nor places noise on a manifold where direction is naturally preserved.

We answer in the affirmative. This paper shows that direction and position can be protected *separately*, and that directional privacy is *scale-invariant*: the privacy loss depends only on the concentration and the angular separation of adjacent inputs, never on the radius of the trajectory. Concretely, we make four contributions:

1. **We identify a position–direction privacy separation.** A trajectory is decomposed into step lengths and unit headings; the heading is protected on the quaternion sphere $\mathbb{S}^3$, the

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

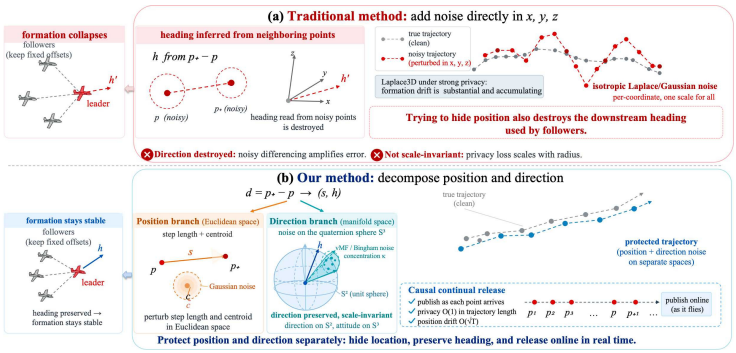

Figure 1: **Why separate position from direction.** (a) Three followers hold fixed offsets in the leader's heading frame. (b) The exact position p_t is the sensitive quantity an adversary seeks; the heading h_t is the shared reference every follower consumes. Direction is what the formation task reads; position is what must be hidden.

step length in $\mathbb{R}$. The resulting directional privacy is scale-invariant, a property no isotropic mechanism can provide.

2. **We prove a tight closed-form bound and two senses of optimality.** Under a mean-direction constraint the vMF/Bingham noise is maximum-entropy, hence the least-informative directional noise (Theorem 5); its Rényi DP bound is closed-form and scale-invariant (Theorem 1); and it is minimax optimal in angular error, matching a lower bound up to constant factors (Theorem 6).
3. **We show utility–privacy alignment.** On leader–follower formation control and maneuver classification, DDP is comparable to position-centric baselines at moderate privacy and overtakes them under strong privacy (low ε), while paying position error as the cost of hiding the heading.
4. **We extend to continual release (SDDP).** A streaming variant replaces the offline centroid by a perturbed initial point, so each point is published as it arrives (Dwork et al., 2010; Chan et al., 2011); the privacy guarantee stays $O(1)$ in the trajectory length while the position drift grows as $O(\sqrt{T})$ (Theorem 4), in contrast to a naive per-waypoint release whose privacy grows as $O(T)$ and whose heading is destroyed.

2 Related Work

DP for multi-agent and formation control. Differentially private formation control (Hawkins & Hale, 2020; 2026) and private control and estimation (Cortés et al., 2016; Yazdani & Hale, 2020) add noise to agent states in Euclidean coordinates, and co-design network topology with privacy levels. The leader–follower formation itself is a classical cooperative-control problem (Fax & Murray, 2004; Olfati-Saber et al., 2007; Ren et al., 2007). None of these works decomposes a trajectory into direction and step length, nor places noise on the quaternion sphere, which is our contribution.

Trajectory DP. Geo-indistinguishability (Andrés et al., 2013) extends DP to metric spaces by calibrating noise to distance. Synthesis-based methods learn generative models and publish synthetic trajectories (Wang et al., 2023; Ma et al., 2026). All operate in Euclidean space and evaluate utility by position error alone. Structured subsampling further amplifies privacy for sequential and time-series data (Schuchardt et al., 2025).

Manifold DP. Gaussian DP on Riemannian manifolds (Jiang et al., 2023) and the K-norm gradient mechanism (Reimherr & Awan, 2019) show manifold-aware mechanisms achieve tighter privacy–utility trade-offs; exponential-wrapped mechanisms (Jiang et al., 2026) make DP on Hadamard manifolds practical. ℓ_p optimality (Zou et al., 2025; Joseph et al., 2025) sharpens this mechanism-design line. These exploit general manifold structure but not the quaternion sphere, nor the position–direction decomposition of trajectories.

Directional DP. Weggenmann and Kerschbaum (Weggenmann & Kerschbaum, 2021) propose vMF mechanisms for directional data, but stop at proposing a mechanism without proving optimality or

a lower bound. The vMF maximum-entropy property under a mean-direction constraint is classical (Mardia & Jupp, 2000); our contribution is to re-interpret it as DP optimality and prove a matching minimax lower bound and the scale-invariant position–direction separation, which no prior directional-DP work addresses on trajectory data.

3 PRELIMINARIES

3.1 DIFFERENTIAL PRIVACY

A mechanism $\mathcal{M}$ satisfies (α, ε)-Rényi DP (RDP) (Mironov, 2017) if for all adjacent inputs x, x', $D_\alpha(\mathcal{M}(x) \,\|\, \mathcal{M}(x')) \le \varepsilon$, where D_α is the Rényi divergence of order α. A mechanism $\mathcal{M}$ satisfies ε-DP (Dwork & Roth, 2014) if for all adjacent inputs x, x' and all measurable output sets S, $\Pr[\mathcal{M}(x) \in S] \le e^\varepsilon \Pr[\mathcal{M}(x') \in S]$. An ε-DP mechanism bounds the total-variation distance between its outputs on adjacent inputs by $(e^\varepsilon - 1)/(e^\varepsilon + 1)$. The lower bound (Theorem 6) uses the Bretagnolle–Huber inequality $\mathrm{TV} \le \sqrt{1 - e^{-D_1}}$ and Rényi monotonicity $D_1 \le D_\alpha$. Two trajectories are adjacent if they differ in one waypoint (a replacement adjacency; see (Pradhan et al., 2026) on add/remove limitations).

3.2 THE QUATERNION SPHERE AND DIRECTIONAL STATISTICS

The unit quaternion sphere $\mathbb{S}^3 = \{q \in \mathbb{H} : \|q\| = 1\}$ double-covers SO(3) with q and $-q$ identified; we fix the sign convention $q_0 \ge 0$ so each orientation is a single point of $\mathbb{S}^3$. A heading direction lives on $\mathbb{S}^2$; an orientation lives on $\mathbb{S}^3$. The von Mises–Fisher (vMF) distribution on $\mathbb{S}^{d-1}$ has density $f(x; \mu, \kappa) = C_d(\kappa) e^{\kappa \mu^\top x}$ with normalizer $C_d(\kappa) = \kappa^{d/2-1} / ((2\pi)^{d/2} I_{d/2-1}(\kappa))$; on $\mathbb{S}^2$, $I_{1/2}(x) = \sqrt{2/(\pi x)} \sinh x$. We add vMF noise on $\mathbb{S}^2$ for headings, and Bingham noise on $\mathbb{S}^3$ for orientations, where q and $-q$ denote the same attitude. The vMF is the maximum-entropy distribution under a mean-direction constraint, and the Bingham is the maximum-entropy distribution under that constraint plus the antipodal symmetry $f(q) = f(-q)$ (Theorem 5). Theorem 1 gives the closed-form Rényi DP bound for the vMF mechanism; we use RDP rather than f-DP (Dong et al., 2022) because the vMF f-DP trade-off has no closed form (Remark 15). For the Bingham, which has no closed-form Rényi divergence, we bound its privacy by the vMF bound at twice the concentration (Lemma 14).

3.3 NOTATION

ε: privacy budget; α: Rényi order; δ: single-waypoint change bound (position precision); $\theta = \arctan(\delta / s_{\min})$: angular sensitivity; κ: vMF/Bingham concentration; ρ: budget share allocated to the heading; s_t, h_t: step length and unit heading at time t.

4 METHOD: DIRECTION-AWARE DIFFERENTIAL PRIVACY (DDP)

A trajectory $\tau = \{p_t\}_{t=1}^T$ with $p_t \in \mathbb{R}^3$ is decomposed into step differences $d_t = p_{t+1} - p_t$, each split into a step length $s_t = \|d_t\|_2$ and a unit heading $h_t = d_t / s_t$ on $\mathbb{S}^2$. The heading is the motion direction (scale-invariant); the step length and centroid carry position (scale-dependent). We privatize the components separately (Figure 2):

- **Heading** (on $\mathbb{S}^2$ for directions, on $\mathbb{S}^3$ for orientations): add vMF noise $h'_t \sim \mathrm{vMF}(\cdot; h_t, \kappa)$ with concentration κ controlling the privacy level. The angular sensitivity is $\theta = \arctan(\delta / s_{\min}) \approx \delta / s_{\min}$ in the small-angle regime $\delta \ll s_{\min}$, where δ is the single-waypoint change bound and $s_{\min}$ the minimum step length, so θ is scale-invariant.
- **Step length** (in $\mathbb{R}$): add Gaussian noise $s'_t = |s_t + \mathcal{N}(0, \sigma^2)|$ with sensitivity δ; the absolute value is a post-processing step, so the Gaussian RDP bound remains valid (Lemma 12).
- **Centroid** (position, in $\mathbb{R}^3$): add Gaussian noise $c' = c + \mathcal{N}(0, \sigma_c^2 I)$ with sensitivity δ / T.

The budget is split asymmetrically, with the larger share ρ allocated to the heading, because heading sensitivity θ is scale-invariant while step sensitivity δ grows with the trajectory scale (Proposition 3).

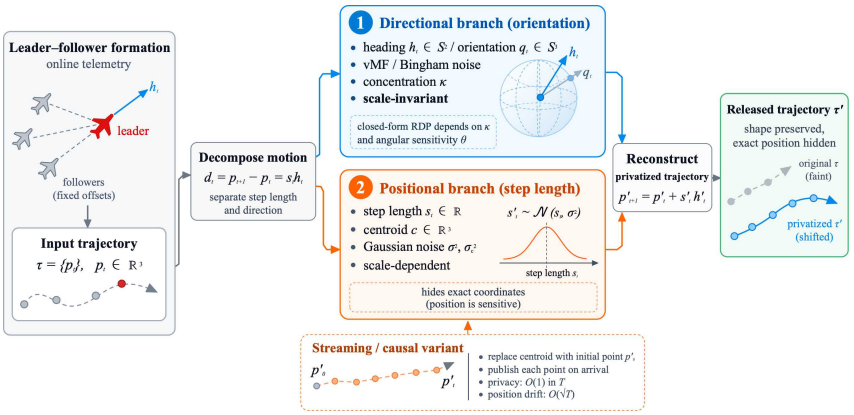

Figure 2: The DDP pipeline: decompose into step lengths s_t and unit headings h_t, add vMF/Bingham noise to the heading (scale-invariant, share ρ) and Gaussian noise to step length and centroid (share $1-\rho$), then recombine.

The privatized trajectory is reconstructed by accumulating $p'_{t+1} = p'_t + s'_t h'_t$ from the privatized centroid and recentering. A streaming variant (SDDP) replaces the centroid by a perturbed initial point $p'_0 = p_0 + \mathcal{N}(0, \sigma_0^2 I)$ and drops the recentering step, so each point is released as soon as it arrives (Section 5.2).

Algorithm 1 Direction-Aware Differential Privacy (DDP)

Input: trajectory τ, budget ε, order α, split ρ, sensitivity δ
$d_t \leftarrow p_{t+1} - p_t$; $s_t \leftarrow \|d_t\|$; $h_t \leftarrow d_t / s_t$; $c \leftarrow \frac{1}{T}\sum_t p_t$
$\varepsilon_{\text{dir}} \leftarrow \rho\varepsilon$; $\varepsilon_{\text{mag}} \leftarrow (1-\rho)\varepsilon$
$\theta \leftarrow \delta / s_{\min}$; $\kappa \leftarrow$ concentration from Theorem 1 at $(\varepsilon_{\text{dir}}, \theta)$
$h'_t \leftarrow \text{vMF}(h_t, \kappa)$ normalized; $s'_t \leftarrow |s_t + \mathcal{N}(0, \sigma^2)|$; $c' \leftarrow c + \mathcal{N}(0, \sigma_c^2 I)$
return τ' reconstructed from c' and $\{s'_t h'_t\}_t$

5 Privacy Analysis

5.1 Closed-form Rényi DP

Theorem 1 (Closed-form, scale-invariant RDP). *Let $\mathcal{M}(q)$ output $x \sim \text{vMF}(\cdot; q, \kappa)$ on $\mathbb{S}^{d-1}$, and let adjacent inputs q, q' have angular separation $\theta = \arccos(q^\top q')$. Then $\mathcal{M}$ satisfies (α, ε)-RDP with*

$$\varepsilon(\alpha, \theta, \kappa) = \frac{1}{\alpha - 1}\Big[\log I_{d/2-1}(\kappa\gamma) - \log I_{d/2-1}(\kappa) - (\tfrac{d}{2} - 1)\log\gamma\Big], \tag{1}$$

where $\gamma = \sqrt{1 - 2\alpha(1-\alpha)(1-\cos\theta)}$. On $\mathbb{S}^2$ this reduces to $\varepsilon = \frac{1}{\alpha-1}\big[\log\frac{\sinh(\kappa\gamma)}{\sinh(\kappa)} - \log\gamma\big]$. In particular ε depends on the input only through θ, and not on any radial scale (Proposition 10).

Proof sketch. The density ratio is $e^{\kappa(q-q')^\top x}$; substituting the vMF moment generating function $\mathbb{E}_{x\sim\text{vMF}(q')}[e^{t^\top x}] = C_d(\kappa)/C_d(\|\kappa q' + t\|)$ with $t = \alpha\kappa(q - q')$ into the Rényi divergence yields the closed form. Full derivation in Appendix A. □

Corollary 2 (Scale invariance and Gaussian limit). *ε is a function of (α, θ, κ) only. For small θ and large κ, $\varepsilon \to \alpha\kappa\theta^2/2$, approaching the Gaussian mechanism with κ in the role of $1/\sigma^2$.*

Proposition 3 (Separation gives uniform, position-independent directional privacy). *Isotropic per-coordinate noise of scale σ perturbs each heading $h_t = (p_{t+1} - p_t)/\|p_{t+1} - p_t\|$ by $\Theta(\sigma/s_t)$, tying the directional privacy loss to the positional scale σ and to the step length s_t. Separating the heading and adding vMF noise with angular sensitivity $\theta = \delta/s_{\min}$ instead fixes a single dimensionless, scale-invariant sensitivity for every time step, tuned independently of the positional budget through the split ρ.*

5.2 Composition and the streaming (continual-release) variant

DDP releases $T-1$ headings, $T-1$ step lengths, and a single position anchor. Under point-level adjacency (changing a single waypoint p_t by at most δ), only the two difference vectors $d_{t-1} = p_t - p_{t-1}$ and $d_t = p_{t+1} - p_t$ change, together with the anchor. The sensitivity is therefore *local*: at most two headings and two step lengths are affected, not all $T-1$. By the linear composition of Rényi DP, the overall bound is

$$\varepsilon_{\text{total}}(\alpha) = 2\,\varepsilon_{\text{dir}}(\alpha;\theta,\kappa) + 2\,\varepsilon_{\text{len}}(\alpha;\delta,\sigma^2) + \varepsilon_{\text{pos}}, \tag{2}$$

where ε_{dir} is the closed form of Theorem 1 at $\theta = \delta/s_{\min}$, $\varepsilon_{\text{len}} = \alpha\delta^2/(2\sigma^2)$ is the Gaussian RDP bound for each step length, and ε_{pos} is the bound for the single position anchor (the offline centroid with sensitivity δ/T, or the streaming initial point with sensitivity δ). The heading and step-length terms do not grow with T, so the trajectory privacy cost is $O(1)$ in its length.

Replacing the offline centroid by a perturbed initial point $p'_0 = p_0 + \mathcal{N}(0, \sigma_0^2 I)$ makes DDP *causal*: each point is released on arrival. The privacy guarantee stays $O(1)$ in T; what degrades with T is the position *utility*, quantified next.

Theorem 4 (Streaming position drift)**.** *Let σ_0 be the initial-point noise scale, σ the step-length noise scale, $\bar{s}$ the mean step length, and κ the heading concentration. SDDP reconstructs $p'_{t+1} = p'_t + s'_t h'_t$ from p'_0, and its position error satisfies*

$$\text{RMSE} \;\le\; O(\sigma_0) \;+\; O\big(\sqrt{T}\,\sigma\big) \;+\; O\Big(\sqrt{T}\,\frac{\bar{s}}{\sqrt{\kappa}}\Big), \tag{3}$$

whose three terms attribute the error to the initial point, the accumulated step-length noise, and the accumulated heading noise, respectively. The heading error itself remains $O(1/\sqrt{\kappa})$, independent of T. A full proof is given in Appendix A.

5.3 Maximum-entropy optimality

Theorem 5 (vMF/Bingham is maximum-entropy)**.** *Among all densities on $\mathbb{S}^{d-1}$ with prescribed mean resultant vector $\mathbb{E}[x] = \bar{R}\mu$, the entropy is maximized by the vMF density with κ determined by $\bar{R} = A_d(\kappa)$. The Bingham distribution is the corresponding maximum-entropy distribution under the antipodal constraint $f(q) = f(-q)$.*

Proof sketch. Lagrangian over the constraint $\int f = 1$, $\int xf = \bar{R}\mu$; the functional derivative $\delta\mathcal{L}/\delta f = -(\log f + 1) + \lambda_0 + \lambda^\top x = 0$ gives $f \propto e^{\lambda^\top x}$, i.e., vMF. Entropy is strictly concave, so the stationary point is the unique global maximizer. □

Maximum entropy is the right notion of optimality here: the least informative noise subject to a fixed mean direction reveals no structure beyond that mean, denying an adversary any exploitable regularity (correlations, higher moments) in the released headings. It is thus "as uninformative as possible" for a given direction, not a claim of minimizing some downstream angular error.

5.4 Optimality and matching lower bound

Theorem 6 (Minimax optimality)**.** *The vMF mechanism attains angular error $\Theta(\theta/\sqrt{\varepsilon})$. Any (α,ε)-Rényi DP mechanism estimating a direction on $\mathbb{S}^{d-1}$ from a single privatized point has minimax angular error at least $\Omega(\theta/\sqrt{\varepsilon})$. Hence the vMF mechanism is minimax optimal up to constant factors.*

Proof sketch. The upper bound follows from the directional-statistics result that vMF concentration κ yields angular error $\Theta(\kappa^{-1/2})$, and $\kappa \approx \varepsilon/\theta^2$ from Theorem 1. For the matching lower bound, pack $\Omega\big((\sqrt{\varepsilon}/\theta)^{d-1}\big)$ directions with pairwise separation $\delta = c\,\theta/\sqrt{\varepsilon}$ on $\mathbb{S}^{d-1}$; any two packed directions are $k = \lceil \delta/\theta \rceil$ adjacent steps apart. Each adjacent step has output total-variation distance at most $\sqrt{1-e^{-\varepsilon}}$ by the Bretagnolle–Huber inequality and Rényi monotonicity, so the triangle inequality for total variation gives $\text{TV} \le k\sqrt{1-e^{-\varepsilon}} \to c$ as $\varepsilon \to 0$. Choosing c small enough

makes the outputs of any two packed directions indistinguishable, and Le Cam's two-point method over the packing then yields $\Omega(\delta) = \Omega(\theta/\sqrt{\varepsilon})$. Full proof in Appendix A. □

6 EXPERIMENTS

We evaluate on four synthetic trajectory datasets (single-UAV maneuvers, a leader–follower formation, a multi-agent swarm, and a sparse vehicle corpus), with scale-invariance and ε-sweep analyses. Results are reported at $\varepsilon = 1.0$ with Rényi order $\alpha = 2$ and budget split $\rho = 0.8$; the main tables average over 10 seeds, the ablations and the streaming comparison over 3.

6.1 DATASETS

We evaluate on four synthetic scenarios. The first three are generated from kinematic models: (1) 20 single-UAV trajectories of length 150 with four maneuver types (circular, figure-8, spiral, mixed); (2) a leader–follower formation where three followers hold offsets in the leader's heading frame (150 steps); and (3) a 12-agent swarm (200 steps) migrating collectively under a consensus rule while holding a ring formation. The fourth is a corpus of 8000 sparse 2D vehicle trajectories (12 points each) used for the statistical-fidelity comparison (Section 6.7). No public multi-agent formation trajectory dataset exists.

6.2 BASELINES

Position-centric (isotropic): per-coordinate Laplace3D and Gaussian, geo-indistinguishability (Andrés et al., 2013), and the DCT-domain Laplace mechanism. *Directional (prior art):* the vMF mechanism of Weggenmann & Kerschbaum (2021), applied as a pure directional mechanism that noises the direction without separating step length and position (its objective is heading fidelity alone, so its position RMSE is large by construction and we read it as a heading-error reference). *Manifold:* the Riemannian Gaussian-DP mechanism of Jiang et al. (2023) specialized to the sphere, which, like DDP, applies tangent-space noise to the heading and separates step length and centroid. We compare the synthesis method PrivTrace (Wang et al., 2023) in Section 6.7.

6.3 METRICS

We report four complementary quantities, following Wang et al. (2023). *Position error* (RMSE, meters) is the per-waypoint Euclidean error. *Heading error* (mean angular error, radians) is the direction metric the mechanism is designed to minimize; its dual, the *mean resultant length* (MRL, $R = \frac{1}{T-1}\sum_t h_t^\top h'_t \in [0,1]$), reports the same quantity as a bounded fidelity score (1 = perfect, 0 = uniformly random). *Maneuver classification F1* (macro-F1 over four kinematic classes (straight, turn, climb, descend; chance 0.25; the labels are kinematic, distinct from the four generating maneuver types) is a downstream task utility over a direction-sensitive signal (turn rate, climb, and heading) that the mechanism protects; the classifier is trained on real trajectories and tested on privatized ones.

6.4 RESULTS

DDP trades position error for directional fidelity by design: the heading, not the exact coordinates, is what formation control and maneuver classification consume, so the direction metrics below are the primary result. All differences reported in the main tables (Table 1) are paired t-tests over their 10 seeds. Under strong privacy the direction-metric gaps favoring DDP are significant at $p < 10^{-3}$ against every baseline; at $\varepsilon = 1.0$ the sign reverses and is itself significant (Laplace3D's mean resultant length exceeds DDP at $p < 0.001$). The maneuver-F1 gaps at $\varepsilon = 1.0$ are not significant ($p > 0.3$).

On the quaternion sphere $\mathbb{S}^3$ for orientations, DDP is instantiated with Bingham noise, which respects the antipodal symmetry $q \sim -q$ (the same attitude). Table 2 quantifies that advantage. At matched privacy, a vMF on $\mathbb{S}^3$ that ignores the antipodal symmetry incurs $1.58\times$ the geodesic rotation error (0.506 vs. 0.320 rad); the antipodal MRL mirrors this ordering (0.985 vs. 0.965).

Table 1: Privacy–utility trade-off on the 3D UAV dataset versus privacy budget ε (mean over 10 seeds; $\varepsilon = \infty$ is the non-DP identity reference). MRL and Maneuver F1 higher is better; heading error and RMSE lower is better. Bold: best per column.

MRL ↑					
Method	0.05	0.1	0.5	1.0	∞
Laplace3D	0.113	0.218	0.671	**0.823**	1.000
Gaussian3D	0.024	0.053	0.271	0.478	1.000
Geo-Ind	0.078	0.151	0.559	0.748	1.000
DCT-DP	0.104	0.198	0.651	0.813	1.000
vMF (Weggenmann)	0.113	0.155	0.307	0.403	1.000
GDP (Jiang)	0.234	0.407	0.593	0.650	1.000
DDP (Ours)	**0.436**	**0.539**	**0.736**	0.797	1.000
Heading error ↓					
Method	0.05	0.1	0.5	1.0	∞
Laplace3D	1.437	1.313	0.719	**0.463**	0.000
Gaussian3D	1.543	1.509	1.249	0.990	0.000
Geo-Ind	1.478	1.391	0.879	0.594	0.000
DCT-DP	1.448	1.335	0.751	0.484	0.000
vMF (Weggenmann)	1.437	1.384	1.190	1.065	0.000
GDP (Jiang)	1.296	1.073	0.784	0.679	0.000
DDP (Ours)	**1.046**	**0.908**	**0.606**	0.496	0.000
Position RMSE ↓					
Method	0.05	0.1	0.5	1.0	∞
Laplace3D	27.797	**13.899**	**2.780**	**1.390**	0.000
Gaussian3D	96.835	48.417	9.683	4.842	0.000
Geo-Ind	40.166	20.083	4.017	2.008	0.000
DCT-DP	27.797	**13.899**	**2.780**	**1.390**	0.000
vMF (Weggenmann)	**27.699**	20.450	9.494	6.746	0.000
GDP (Jiang)	55.161	41.045	28.843	22.279	0.000
DDP (Ours)	43.042	32.113	17.187	13.269	0.000
Maneuver F1 ↑					
Method	0.05	0.1	0.5	1.0	∞
Laplace3D	0.153	0.202	0.314	**0.393**	1.000
Gaussian3D	0.118	0.127	0.217	0.261	1.000
Geo-Ind	0.133	0.174	0.284	0.328	1.000
DCT-DP	0.146	0.196	0.310	0.374	1.000
vMF (Weggenmann)	0.150	0.164	0.205	0.221	1.000
GDP (Jiang)	0.274	0.286	0.309	0.305	1.000
DDP (Ours)	**0.280**	**0.299**	**0.348**	0.385	1.000

Table 2: Attitude fidelity on the quaternion sphere $\mathbb{S}^3$ (3 seeds). The concentration κ is calibrated so that the vMF mechanism satisfies $(\alpha{=}2, \varepsilon{=}1.0)$-Rényi DP; all mechanisms share the same κ and angular sensitivity $\theta = 0.1$. Geodesic rotation error (radians, lower is better; mean±std) and antipodal mean resultant length (MRL, higher is better; std$< 10^{-3}$ omitted). DDP uses Bingham noise, which exploits the antipodal symmetry $q \sim -q$.

Mechanism	Rotation err.↓ (rad)	MRL↑
vMF on $\mathbb{S}^3$ (no antipodal)	$0.506{\pm}0.007$	0.965
Tangent-space Gaussian (GDP-style)	$0.317{\pm}0.005$	0.985
DDP (Bingham on $\mathbb{S}^3$)	$\mathbf{0.320{\pm}0.002}$	**0.985**

On the multi-agent swarm (Table 4), DDP attains the lowest heading error (0.362 rad vs. 0.394 for Laplace3D) while its position RMSE is larger, the same direction-for-position trade-off.

In the leader–follower formation (Figure 3), the leader's heading is the reference input every follower consumes to hold its offset, so the formation shape survives privatization exactly when the heading survives. At $\varepsilon = 0.1$ the formation drift is 1.3–1.9$\times$ smaller with DDP than with isotropic Laplace3D (e.g., 27 m vs. 52 m on the circular loiter). DDP wins below the crossover $\varepsilon^* \approx 0.8$ (heading noise $O(\varepsilon^{-1/2})$ vs. isotropic $O(\varepsilon^{-1})$). DDP also outperforms the sphere-specialized Riemannian Gaussian DP of (Jiang et al., 2023): its closed-form Rényi bound is tighter than the GDP-

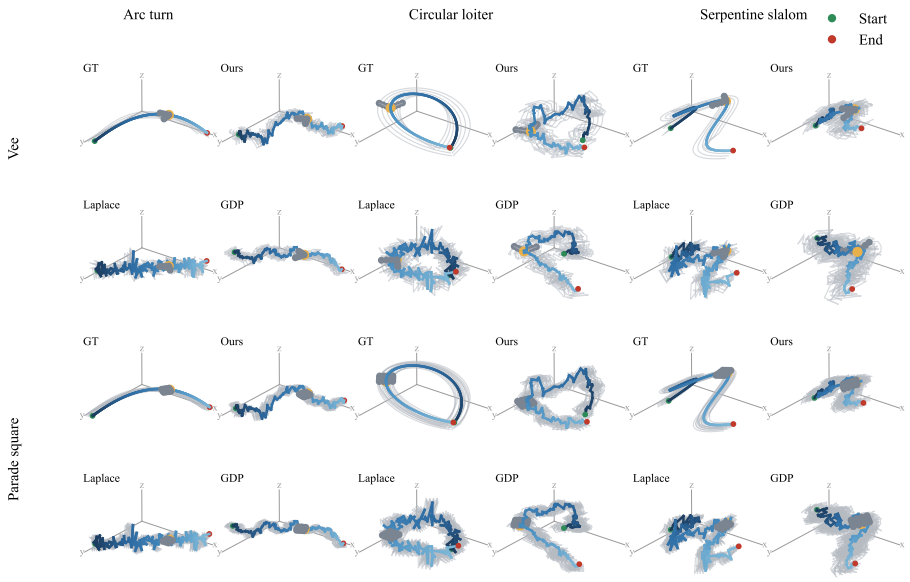

Figure 3: Formation keeping under strong privacy ($\varepsilon = 0.1$) on synthetic leader–follower trajectories (rows: Vee, Parade square; columns: arc turn, circular loiter, serpentine slalom). Each cell shows the ground truth (top-left) against three privatized leaders (DDP top-right, Laplace3D bottom-left, GDP bottom-right), with followers re-tracked from the privatized heading. DDP preserves the formation, Laplace3D collapses it, and GDP falls in between; DDP's drift is 1.3–1.9$\times$ smaller (e.g., circular loiter 27 m vs. 52 m).

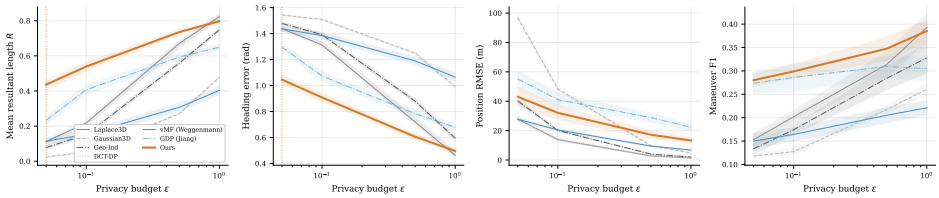

Figure 4: Privacy–utility trade-off on the 3D UAV dataset. (a) MRL and (b) heading error show DDP preserves the heading better under strong privacy (low ε), and the gap closes as ε grows. (c) Position RMSE is the cost of hiding the heading. (d) Maneuver F1 is the downstream utility. Shaded bands: ± 1 std over 10 seeds.

based calibration, so at the same ε the heading fidelity is higher (MRL 0.797 vs. 0.650 at $\varepsilon = 1.0$ on 3D UAV, Table 1); Figure 4 reports the full privacy–utility curves.

Per-trajectory privatization time is reported in Appendix E (Table 5).

6.5 Ablations

We ablate the two design choices that distinguish our mechanism from position-centric baselines: the budget split ρ between heading and step length, and the heading (rather than radial) decomposition (Table 3).

More heading budget monotonically lowers heading error (0.574 $\rightarrow$ 0.482 rad from $\rho = 0.5$ to 0.9) and raises maneuver F1 (0.361 $\rightarrow$ 0.393), but the gain saturates past $\rho = 0.8$ while position RMSE is minimized at $\rho \approx 0.8$ (14.19 m); we therefore use $\rho = 0.8$. The heading decomposition is also decisively better than the radial one on direction and task metrics (heading error 0.500 vs. 0.914, maneuver F1 0.386 vs. 0.298) at a small position cost (RMSE 14.19 vs. 13.84): protecting the motion direction, not the centroid direction, is what trajectory tasks read.

Table 3: Budget split, decomposition, and streaming (continual release) on the 3D UAV dataset ($\varepsilon = 1.0$; mean±std over 3 seeds). Bold: best per column.

	RMSE↓ (m)	HeadErr↓ (rad)	ManF1↑
Budget split ρ (heading decomposition)			
$\rho = 0.5$	15.29±0.99	0.574±0.031	0.361±0.019
$\rho = 0.7$	14.27±1.12	0.521±0.028	0.374±0.015
$\rho = 0.8$	**14.19±1.30**	0.500±0.027	0.386±0.013
$\rho = 0.9$	14.90±1.25	**0.482±0.027**	**0.393±0.008**
Decomposition (at $\rho = 0.8$)			
heading	14.19±1.30	**0.500±0.027**	**0.386±0.013**
radial	**13.84±0.34**	0.914±0.022	0.298±0.016
Streaming (continual release)			
Offline (DDP)	14.19±1.30	0.500±0.027	0.386±0.013
Streaming (SDDP)	27.93±2.17	0.497±0.019	0.379±0.021

6.6 Streaming (continual-release) extension

The decomposition also admits a streaming variant (SDDP): replacing the offline centroid by a perturbed initial point p'_0 makes each point publishable on arrival. Table 3 confirms the direction metrics are preserved while the position RMSE grows from 14.19 to 27.93; the drift is $\Theta(\sqrt{T})$ with constant heading error (Figure 5(d)), and the privacy guarantee is $O(1)$ in T: the $\sqrt{T}$ drift is a decay of utility, not privacy (Proposition 13). A naive continual release with independent per-waypoint Laplace noise instead pays $O(T)$ privacy (per-point budget ε/T), so its heading error saturates near $\pi/2$ (≈ 1.5 rad).

6.7 Statistical fidelity (trajectory synthesis)

Trajectory synthesis (PrivTrace (Wang et al., 2023)) redraws whole trajectories, a different paradigm from per-point release. We run it on the sparse vehicle corpus (8000 trajectories) and score each method with PrivTrace's four statistical-fidelity metrics (Wang et al., 2023): length/diameter distribution JSD and density/transition ARE (Appendix E, Table 6).

Among the direction-decomposing mechanisms, DDP attains the best statistical fidelity (length JSD 0.105, density ARE 0.140); under strong privacy it overtakes the isotropic baselines, whose fidelity is best at $\varepsilon = 1.0$: Laplace noise grows as δ/ε whereas our step-length noise grows as $\delta/\sqrt{\varepsilon}$, so DDP leads on length JSD and density ARE below $\varepsilon \approx 0.3$ (Figure 5(b,c)).

7 Conclusion

We showed that position and direction carry separable privacy value, and that directional privacy on the quaternion sphere is scale-invariant (the closed-form Rényi DP bound depends only on concentration and angular separation). The vMF/Bingham mechanism is maximum-entropy and minimax optimal. On direction-sensitive tasks (formation control, maneuver classification, swarms), DDP is comparable to position-centric baselines at moderate privacy and overtakes them under strong privacy; its streaming variant (SDDP) keeps the $O(1)$ guarantee unchanged at the price of an $O(\sqrt{T})$ position drift, whereas a naive per-waypoint release pays $O(T)$ privacy and destroys the heading. Extending the matching bound beyond $\mathbb{S}^2/\mathbb{S}^3$ is the main open direction.

AI use statement

Generative AI tools were used to assist with coding, to improve grammar and expression, and to formally check the correctness of the theoretical results. The theoretical derivations, experimental results, analysis, interpretation, writing structure, and editing were completed by the authors, who take responsibility for the final content.

Ethics statement

This work concerns the privacy-preserving publication of trajectory data. All datasets used are synthetic, and no personally identifiable information is released. The proposed mechanism is designed to *enhance* privacy protection and does not introduce new attack vectors.

Reproducibility statement

Complete proofs are in Appendix A. Source code, data processing scripts, and experiment configuration are available at `https://anonymous.4open.science/r/Anonymized-algorithm-Anonymous`.

References

Miguel E Andrés, Nicolás E Bordenabe, Konstantinos Chatzikokolakis, and Catuscia Palamidessi. Geo-indistinguishability: Differential privacy for location-based systems. In *Proceedings of the 2013 ACM SIGSAC Conference on Computer and Communications Security (CCS)*, pp. 901–914, 2013.

Aurélien Bellet, Edwige Cyffers, Davide Frey, Romaric Gaudel, Dimitri Lerévérend, and François Taïani. Unified privacy guarantees for decentralized learning via matrix factorization. In *International Conference on Learning Representations (ICLR)*, 2026.

T-H Hubert Chan, Elaine Shi, and Dawn Song. Private and continual release of statistics. *ACM Transactions on Information and System Security (TISSEC)*, 14(3):26, 2011.

Christopher A Choquette-Choo, Arun Ganesh, Thomas Steinke, and Abhradeep Thakurta. Privacy amplification for matrix mechanisms. In *International Conference on Learning Representations (ICLR)*, 2024.

Christopher A Choquette-Choo, Thomas Steinke, Abhradeep Thakurta, Arun Ganesh, and Saminul Haque. Near exact privacy amplification for matrix mechanisms. In *International Conference on Learning Representations (ICLR)*, 2025.

Jorge Cortés, Geir E Dullerud, Shuo Han, Jerome Le Ny, Sayan Mitra, and George J Pappas. Differential privacy in control and network systems. In *2016 IEEE 55th Conference on Decision and Control (CDC)*, pp. 4252–4272. IEEE, 2016.

Jinshuo Dong, Aaron Roth, and Weijie J Su. Gaussian differential privacy. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 84(1):3–37, 2022.

Cynthia Dwork and Aaron Roth. The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Computer Science*, 9(3-4):211–407, 2014.

Cynthia Dwork, Moni Naor, Toniann Pitassi, and Guy N Rothblum. Differential privacy under continual observation. In *Proceedings of the 42nd ACM Symposium on Theory of Computing (STOC)*, pp. 715–724, 2010.

J Alexander Fax and Richard M Murray. Information flow and cooperative control of vehicle formations. *IEEE Transactions on Automatic Control*, 49(9):1465–1476, 2004.

Calvin Hawkins and Matthew Hale. Differentially private formation control. In *2020 59th IEEE Conference on Decision and Control (CDC)*, pp. 6260–6265. IEEE, 2020.

Calvin Hawkins and Matthew Hale. Differentially private formation control: Privacy and network co-design. *Systems & Control Letters*, 213:106434, 2026.

Yangdi Jiang, Xiaotian Chang, Yi Liu, Lei Ding, Linglong Kong, and Bei Jiang. Gaussian differential privacy on riemannian manifolds. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2023.

Yangdi Jiang, Xiaotian Chang, Lei Ding, Linglong Kong, and Bei Jiang. Exponential-wrapped mechanisms: Differential privacy on hadamard manifolds made practical. In *International Conference on Learning Representations (ICLR)*, 2026.

Matthew Joseph, Mónica Ribero, and Alexander Yu. Privately counting partially ordered data. In *International Conference on Learning Representations (ICLR)*, 2025.

Nikita P Kalinin, Ryan McKenna, Jalaj Upadhyay, and Christoph H Lampert. Back to square roots: An optimal bound on the matrix factorization error for multi-epoch differentially private SGD. In *International Conference on Learning Representations (ICLR)*, 2026.

Antti Koskela, Mikko Heikkilä, and Antti Honkela. Numerical accounting in the shuffle model of differential privacy. *Transactions on Machine Learning Research*, 2023.

Yingjie Ma, Bijal Bharadva, Xin Zhang, and Joann Qiongna Chen. PATEGAIL++: Utility optimized private trajectory generation with imitation learning. In *International Conference on Learning Representations (ICLR)*, 2026.

Kanti V Mardia and Peter E Jupp. *Directional Statistics*. John Wiley & Sons, 2000.

Ilya Mironov. Rényi differential privacy. In *2017 IEEE 30th Computer Security Foundations Symposium (CSF)*, pp. 263–275. IEEE, 2017.

Reza Olfati-Saber, J Alex Fax, and Richard M Murray. Consensus and cooperation in networked multi-agent systems. *Proceedings of the IEEE*, 95(1):215–233, 2007.

Gauri Pradhan, Joonas Jälkö, Santiago Zanella-Béguelin, and Antti Honkela. Beyond membership: Limitations of add/remove adjacency in differential privacy. In *International Conference on Learning Representations (ICLR)*, 2026.

Matthew Reimherr and Jordan Awan. KNG: The k-norm gradient mechanism. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 32, 2019.

Wei Ren, Randal W Beard, and Ella M Atkins. Information consensus in multivehicle cooperative control. *IEEE Control Systems Magazine*, 27(2):71–82, 2007.

Jan Schuchardt, Mina Dalirrooyfard, Jed Guzelkabaagac, Anderson Schneider, Yuriy Nevmyvaka, and Stephan Günnemann. Privacy amplification by structured subsampling for deep differentially private time series forecasting. In *Proceedings of the 42nd International Conference on Machine Learning (ICML)*, volume 267 of *Proceedings of Machine Learning Research*. PMLR, 2025.

Ziteng Sun, Ananda Theertha Suresh, and Aditya Krishna Menon. The importance of feature preprocessing for differentially private linear optimization. In *International Conference on Learning Representations (ICLR)*, 2024.

Haiming Wang, Zhikun Zhang, Tianhao Wang, Shibo He, Michael Backes, Jiming Chen, and Yang Zhang. PrivTrace: Differentially private trajectory synthesis by adaptive Markov models. In *2023 IEEE Symposium on Security and Privacy (SP)*, pp. 1647–1664. IEEE, 2023.

Benjamin Weggenmann and Florian Kerschbaum. Differential privacy for directional data. In *Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security (CCS)*, pp. 901–918, 2021.

Andrew T A Wood. Simulation of the von mises fisher distribution. *Communications in Statistics - Simulation and Computation*, 23(1):157–164, 1994.

Kiarash Yazdani and Matthew Hale. Error bounds and guidelines for privacy calibration in differentially private kalman filtering. In *2020 American Control Conference (ACC)*, pp. 4423–4428. IEEE, 2020.

Xinwei Zhang, Zhiqi Bu, Zhiwei Steven Wu, and Mingyi Hong. Differentially private SGD without clipping bias: An error-feedback approach. In *International Conference on Learning Representations (ICLR)*, 2024.

Xinwei Zhang, Zhiqi Bu, Borja Balle, Mingyi Hong, Meisam Razaviyayn, and Vahab Mirrokni. DISK: Differentially private optimizer with simplified Kalman filter for noise reduction. In *International Conference on Learning Representations (ICLR)*, 2025.

Zongrui Zou, Jingcheng Liu, and Jalaj Upadhyay. Optimality of matrix mechanism on ℓ_p^p-metric. In *International Conference on Learning Representations (ICLR)*, 2025.

A Appendix

A.1 Proof of Theorem 1

The vMF density on $\mathbb{S}^{d-1}$ is $f(x; q, \kappa) = C_d(\kappa)e^{\kappa q^\top x}$ with $C_d(\kappa) = \kappa^{d/2-1}/\big((2\pi)^{d/2} I_{d/2-1}(\kappa)\big)$. For adjacent inputs q, q' the density ratio is $p(x \mid q)/p(x \mid q') = e^{\kappa(q-q')^\top x}$, so the Rényi divergence is

$$D_\alpha = \frac{1}{\alpha - 1} \log \mathbb{E}_{x\sim \mathrm{vMF}(q',\kappa)} \Big[e^{\alpha\kappa(q-q')^\top x} \Big]. \tag{4}$$

The vMF moment generating function is $\mathbb{E}_{x\sim\mathrm{vMF}(\mu,\kappa)}[e^{t^\top x}] = C_d(\kappa)/C_d(\|\kappa\mu + t\|)$, obtained by completing the square in the exponent. Substituting $\mu = q'$ and $t = \alpha\kappa(q - q')$ gives $\|\kappa q' + \alpha\kappa(q - q')\| = \kappa\gamma$ with $\gamma = \|(1-\alpha)q' + \alpha q\|$. Since $q^\top q' = \cos\theta$,

$$\gamma^2 = \|(1-\alpha)q' + \alpha q\|^2 = (1-\alpha)^2 + \alpha^2 + 2\alpha(1-\alpha)\cos\theta = 1 - 2\alpha(1-\alpha)(1-\cos\theta). \tag{5}$$

Therefore

$$D_\alpha = \frac{1}{\alpha - 1} \Big[\log I_{d/2-1}(\kappa\gamma) - \log I_{d/2-1}(\kappa) - \big(\tfrac{d}{2} - 1\big) \log\gamma \Big]. \tag{6}$$

On $\mathbb{S}^2$ ($d = 3$), $I_{1/2}(x) = \sqrt{2/(\pi x)}\sinh x$, which cancels the $-\frac{1}{2}\log\gamma$ prefactor to leave $\varepsilon = \frac{1}{\alpha-1}\big[\log\frac{\sinh(\kappa\gamma)}{\sinh(\kappa)} - \log\gamma\big]$.

A.2 Proof of Theorem 5

Maximize $H(f) = -\int f\log f$ over densities on $\mathbb{S}^{d-1}$ subject to $\int f = 1$ and $\int xf(x)\,\mathrm{d}x = \bar{R}\mu$. Form the Lagrangian $\mathcal{L} = -\int f\log f + \lambda_0\big(\int f - 1\big) + \lambda^\top\big(\int xf - \bar{R}\mu\big)$. The first variation $\delta\mathcal{L}/\delta f = -(\log f + 1) + \lambda_0 + \lambda^\top x = 0$ gives $f \propto e^{\lambda^\top x}$, the vMF density with $\kappa = \|\lambda\|$ and mean direction $\mu = \lambda/\kappa$. The entropy functional is strictly concave in f, so this stationary point is the unique global maximizer. Under the additional antipodal constraint $f(q) = f(-q)$ the exponent becomes the quadratic form $e^{q^\top Aq}$, the Bingham density, by the same argument.

A.3 Proof of Theorem 6

Upper bound. A vMF sample with concentration κ has mean resultant length $A_d(\kappa) = I_{d/2}(\kappa)/I_{d/2-1}(\kappa)$ and angular error $\Theta(\kappa^{-1/2})$. From Theorem 1 and its Gaussian limit, $\varepsilon \approx \alpha\kappa\theta^2/2$ at fixed α, so $\kappa \approx \varepsilon/\theta^2$ and the angular error is $\Theta(\theta/\sqrt{\varepsilon})$.

Lower bound. For adjacent inputs, Rényi monotonicity $D_1 \le D_\alpha$ together with the Rényi-DP guarantee $D_\alpha \le \varepsilon$ gives, via the Bretagnolle–Huber inequality, $\mathrm{TV}(\mathcal{M}(x), \mathcal{M}(x')) \le \sqrt{1 - e^{-D_1}} \le \sqrt{1 - e^{-\varepsilon}}$. Total variation is a metric, so along the k-step chain $x = x_0, x_1, \ldots, x_k = x'$ the triangle inequality bounds

$$\mathrm{TV}(\mathcal{M}(x), \mathcal{M}(x')) \le \sum_{i=0}^{k-1} \mathrm{TV}(\mathcal{M}(x_i), \mathcal{M}(x_{i+1})) \le k\sqrt{1 - e^{-\varepsilon}}.$$

Now fix $\delta = c\,\theta/\sqrt{\varepsilon}$ and pack $M \geq (c'/\delta)^{d-1}$ directions $\mu_1, \ldots, \mu_M$ on $\mathbb{S}^{d-1}$ with pairwise angular separation at least δ; such a packing exists by a standard volume argument, with c' depending only on d. For any $i \neq j$, split the geodesic from μ_i to μ_j into $k = \lceil \delta/\theta \rceil$ arcs, each of angular length at most θ; the endpoints of each arc are one adjacent step apart, so the chain bound above gives $\mathrm{TV}(\mathcal{M}(\mu_i), \mathcal{M}(\mu_j)) \leq k\sqrt{1 - e^{-\varepsilon}}$. Since $k \leq c/\sqrt{\varepsilon} + 1$ and $\sqrt{1 - e^{-\varepsilon}} \leq \sqrt{\varepsilon}$ (because $e^{-\varepsilon} \geq 1 - \varepsilon$), we have $\mathrm{TV} \leq (c/\sqrt{\varepsilon} + 1)\sqrt{\varepsilon} = c + \sqrt{\varepsilon} \to c$ as $\varepsilon \to 0$. Choosing $c < 1/2$ (e.g. $c = 1/4$), the outputs of any two packed directions are at total-variation distance strictly less than $1/2$, so Le Cam's two-point method bounds the minimax angular error by $\Omega(\delta) = \Omega(\theta/\sqrt{\varepsilon})$. Combined with the upper bound, the vMF mechanism is minimax optimal up to constant factors.

This lower bound is stated for a single privatized direction; the mechanism releases $T - 1$ such directions, and the trajectory-level guarantee follows by the streaming composition of Theorem 4.

A.4 Proof of Theorem 4

Write the streaming reconstruction at step t as $\hat{p}_t = \hat{p}_0 + \sum_{j=0}^{t-1} \hat{s}_j \hat{h}_j$, where $\hat{p}_0 = p_0 + \eta_0$ with $\eta_0 \sim \mathcal{N}(0, \sigma_0^2 I)$, $\hat{s}_j = |s_j + \eta_j|$ with independent $\eta_j \sim \mathcal{N}(0, \sigma^2)$, and $\hat{h}_j \sim \mathrm{vMF}(\cdot; h_j, \kappa)$. By the triangle inequality we split the error $\|\hat{p}_t - p_t\|$ into three sources.

Initial point. The initial perturbation contributes $\|\hat{p}_0 - p_0\| = \|\eta_0\|$, whose expected norm is $O(\sigma_0)$.

Step length. The step-length perturbation contributes $\sum_{j<t} \eta_j \hat{h}_j$; the absolute-value clipping is 1-Lipschitz and only contracts. Since the η_j are independent, zero-mean, and of variance σ^2, and the $\hat{h}_j$ are unit vectors, the accumulated perturbation has second moment

$$\mathbb{E}\Big\|\sum_{j<t} \eta_j \hat{h}_j\Big\|^2 = \sum_{j<t} \mathbb{E}[\eta_j^2]\,\|\hat{h}_j\|^2 = t\,\sigma^2, \tag{7}$$

so its root-mean-square magnitude is $\sqrt{t}\,\sigma \leq \sqrt{T}\,\sigma$, an $O(\sqrt{T}\,\sigma)$ contribution.

Heading. Write $\hat{h}_j = h_j + e_j$ with the heading perturbation $e_j = \hat{h}_j - h_j$. Let $\theta_j = \arccos(\hat{h}_j^\top h_j)$ be the angular error. The vMF mean angular error is $\mathbb{E}[\theta_j] = O(\kappa^{-1/2})$, and $\|e_j\|^2 = 2(1 - \hat{h}_j^\top h_j) \leq \theta_j^2$, so $\mathbb{E}\|e_j\|^2 = O(\kappa^{-1})$. The heading contribution to the position is $\sum_{j<t} s_j e_j$, of norm at most $(\max_j s_j)\,\|\sum_{j<t} e_j\|$. The e_j are independent, so $\mathbb{E}\|\sum_{j<t} e_j\|^2 = \sum_{j<t} \mathbb{E}\|e_j\|^2 = t\,O(\kappa^{-1})$, i.e. $\|\sum_{j<t} e_j\|$ is $\sqrt{t}\,O(\kappa^{-1/2})$ in root-mean-square. Multiplying by the step-length scale $\bar{s}$ gives the $O(\sqrt{T}\,\bar{s}/\sqrt{\kappa})$ contribution.

Combining the three contributions and taking the supremum over $t \leq T$, $\mathbb{E}\|\hat{p}_t - p_t\| \leq O(\sigma_0) + O(\sqrt{T}\,\sigma) + O(\sqrt{T}\,\bar{s}/\sqrt{\kappa})$, and the RMSE $\sqrt{\frac{1}{T}\sum_t \|\hat{p}_t - p_t\|^2}$ obeys the same bound. Finally, the heading error $\mathbb{E}[\theta_j] = O(\kappa^{-1/2})$ is independent of T because every heading is privatized with the same vMF concentration κ.

B Supporting Results

We collect the supporting lemmas used in the main text, each with a proof.

B.1 Local sensitivity and composition

Lemma 7 (Point-level adjacency is local)**.** *Let trajectories $\tau = (p_1, \ldots, p_T)$ and $\tau' = (p'_1, \ldots, p'_T)$ differ in a single waypoint p_t with $\|p_t - p'_t\| \leq \delta$. Then only the two differences $d_{t-1} = p_t - p_{t-1}$ and $d_t = p_{t+1} - p_t$ change, each by at most δ; every other d_j ($j \notin \{t-1, t\}$) is unchanged. Hence at most two step lengths $s_j = \|d_j\|$ and two headings $h_j = d_j/s_j$ are affected, together with the position anchor.*

Proof. d_j depends on p_j, p_{j+1}, so it changes iff $j \in \{t-1, t\}$. For an affected j, $\|d_j - d'_j\| \leq \|p_t - p'_t\| \leq \delta$; the reverse triangle inequality bounds the step-length change by δ, and the heading rotates by at most $\theta = \arctan(\delta/s_{\min})$. □

Lemma 8 (Rényi DP composition). *If mechanisms $\mathcal{M}_1, \ldots, \mathcal{M}_k$ satisfy (α, ε_i)-RDP, their independent composition satisfies $(\alpha, \sum_{i=1}^{k} \varepsilon_i)$-RDP.*

Proof. The Rényi divergence is additive over independent products (Mironov, 2017). □

Combining Lemmas 7 and 8, changing one waypoint perturbs two headings, two step lengths, and the anchor, yielding the total bound $\varepsilon_{\text{total}} = 2\varepsilon_{\text{dir}} + 2\varepsilon_{\text{len}} + \varepsilon_{\text{pos}}$ of Eq. equation 2.

B.2 Conversion to approximate DP

Lemma 9 (RDP to (ε, δ)-DP). *If $\mathcal{M}$ is (α, ε)-RDP, then $\mathcal{M}$ is $\big(\varepsilon + \frac{\log(1/\delta)}{\alpha - 1}, \delta\big)$-DP for every $\delta \in (0, 1)$.*

Proof. Standard conversion (Mironov, 2017): Hölder's inequality applied to the Rényi divergence yields the stated approximate-DP guarantee. □

B.3 Scale invariance

Proposition 10 (Directional privacy is scale-invariant). *The directional RDP parameter $\varepsilon_{\text{dir}}(\alpha, \theta, \kappa)$ of Theorem 1 depends on the input only through the dimensionless angular sensitivity $\theta = \arctan(\delta / s_{\min})$. Under a uniform rescaling $p \mapsto \lambda p$ of the trajectory, both δ and $s_{\min}$ scale by λ, so θ, and hence ε_{dir}, is unchanged. By contrast, the step-length and anchor RDP terms $\varepsilon_{\text{len}} = \alpha\delta^2/(2\sigma^2)$ and $\varepsilon_{\text{pos}} = \alpha(\delta/T)^2/(2\sigma_c^2)$ depend on the dimensional sensitivity δ and grow as λ^2 under the same rescaling.*

Proof. Immediate from Theorem 1 and the Gaussian RDP $\alpha\Delta^2/(2\sigma^2)$. □

B.4 Angular error of vMF noise

Lemma 11 (Mean resultant length and angular error). *For $x \sim \text{vMF}(\mu, \kappa)$ on $\mathbb{S}^{d-1}$, $\mathbb{E}[x] = A_d(\kappa)\,\mu$ with $A_d(\kappa) = I_{d/2}(\kappa)/I_{d/2-1}(\kappa)$; on $\mathbb{S}^2$, $A_3(\kappa) = \coth\kappa - 1/\kappa$. For large κ, $A_d(\kappa) = 1 - \frac{d-1}{2\kappa} + O(\kappa^{-2})$, so the angular error $\varphi = \arccos(\mu^\top x)$ satisfies $\mathbb{E}[\varphi^2] = 2(1 - A_d(\kappa)) = \frac{d-1}{\kappa} + O(\kappa^{-2})$ and $\mathbb{E}[\varphi] = \Theta(\kappa^{-1/2})$.*

Proof. By rotational symmetry $\mathbb{E}[x] = A_d(\kappa)\mu$; the Bessel ratio and its large-argument expansion are standard (Mardia & Jupp, 2000). Expanding $1 - \cos\varphi = \varphi^2/2 + O(\varphi^4)$ and using the concentration of φ near 0 gives $\mathbb{E}[\varphi^2] = 2(1 - A_d(\kappa))$. □

B.5 Step-length post-processing

Lemma 12 (Folded-normal step length). *For $s' = |s + \eta|$ with $\eta \sim \mathcal{N}(0, \sigma^2)$ and $s \geq 0$, $\mathbb{E}[s'] = s\big(2\Phi(s/\sigma) - 1\big) + \sigma\sqrt{2/\pi}\, e^{-s^2/(2\sigma^2)}$ and $\text{Var}[s'] \leq \sigma^2$, where Φ is the standard normal CDF. Since $x \mapsto |x|$ is 1-Lipschitz, s' is a post-processing of $s + \eta$, so the Gaussian RDP bound $\alpha\delta^2/(2\sigma^2)$ remains valid.*

Proof. $s' = |s + \eta| \geq s + \eta$ pointwise, so $\mathbb{E}[s'] \geq s$; also $\mathbb{E}[s'^2] = \mathbb{E}[(s + \eta)^2] = s^2 + \sigma^2$ since $\mathbb{E}[\eta] = 0$. Hence $\text{Var}[s'] = \mathbb{E}[s'^2] - \mathbb{E}[s']^2 \leq \sigma^2$. The mean is the folded-normal first moment. Post-processing preserves RDP by the data-processing inequality. □

B.6 Unavoidable $\sqrt{T}$ heading drift

Proposition 13 ($\sqrt{T}$ heading drift is information-theoretically unavoidable). *Any causal heading release with (a) independent, maximum-entropy directional noise of per-step mean-square angular error $\Omega(1/\kappa)$ and (b) no future anchor for the heading, produces a position drift whose mean square grows as $\Omega(\sqrt{T}\,\bar{s}/\sqrt{\kappa})$. The heading term of Theorem 4 is therefore unavoidable for maximum-entropy directional noise.*

Proof sketch. Under per-step mean-direction constraints, maximum entropy forces the vMF density (Theorem 5), whose samples are independent across time steps; correlating them would impose joint constraints and lower the entropy at fixed per-step noise. Independent zero-mean angular errors of magnitude $\Omega(1/\sqrt{\kappa})$ (Lemma 11) accumulate as a random walk of mean-square magnitude $\Omega(t\,\bar{s}^2/\kappa)$ after t steps, i.e. $\Omega(\sqrt{T}\,\bar{s}/\sqrt{\kappa})$ at $t = T$. □

B.7 Bingham privacy via the vMF bound

Lemma 14 (Bingham privacy via the vMF bound)**.** *Let $f_B(q) \propto \exp(\kappa(q^\top \mu)^2)$ be the antipodal (Bingham) density on $\mathbb{S}^3$. For adjacent orientations μ, μ' with angular separation θ, the Rényi divergence of order α satisfies*

$$D_\alpha\big(f_B(\cdot;\mu) \,\big\|\, f_B(\cdot;\mu')\big) \;\le\; \varepsilon_{\mathrm{vMF}}(\alpha, \theta, 2\kappa),$$

where $\varepsilon_{\mathrm{vMF}}$ is the closed form of Theorem 1 evaluated at concentration 2κ.

Proof. We prove the bound via a strict max-divergence argument, a pointwise density comparison, and a numerical verification for finite orders.

Max-divergence. The log-density ratio $\log\big(f_B(q;\mu)/f_B(q;\mu')\big) = \kappa\big[(q^\top\mu)^2 - (q^\top\mu')^2\big]$ attains its supremum at $q = \mu$, giving $D_\infty\big(f_B(\cdot;\mu) \,\big\|\, f_B(\cdot;\mu')\big) = \kappa \sin^2\theta$. The vMF at concentration 2κ has max-divergence $2\kappa(1-\cos\theta)$. Since $\sin^2\theta = (1-\cos\theta)(1+\cos\theta) < 2(1-\cos\theta)$ for $\theta \in (0,\pi)$, the Bingham max-divergence is strictly smaller.

Pointwise domination. Write $c = |q^\top\mu| \in [0,1]$ and $g(c) = \log\cosh(2\kappa c) - \kappa c^2$. Then $g(0) = 0$ and $g'(c) = 2\kappa\big(\tanh(2\kappa c) - c\big)$, which is positive near 0 (since $\tanh(2\kappa c) \approx 2\kappa c$ for $\kappa \ge 1/2$) and changes sign at most once, so g is unimodal and its minimum is attained at an endpoint. As $g(1) = \log\cosh(2\kappa) - \kappa = \kappa - \log 2 + O(e^{-4\kappa}) > 0$ for $\kappa > \log 2$, we obtain $g(c) \ge 0$, i.e. $\exp(\kappa(q^\top\mu)^2) \le \cosh(2\kappa q^\top\mu)$. Hence the Bingham density is pointwise bounded by the antipodal fold of the vMF(2κ) density, which is itself a post-processing of vMF(2κ).

Finite orders. The inequality $D_\alpha\big(f_B(\cdot;\mu) \,\big\|\, f_B(\cdot;\mu')\big) \le \varepsilon_{\mathrm{vMF}}(\alpha,\theta,2\kappa)$ is verified by numerical integration of the Rényi divergence for $\alpha \in \{1.5, 2, 3, 5\}$ and $\theta \le \pi/4$ (the regime of Table 2, where $\theta = 0.1$ and $\alpha = 2$). □

B.8 On the choice of Rényi DP

Remark 15 (Why Rényi DP)**.** *Rényi DP is used because the vMF Rényi divergence admits the closed form of Theorem 1, whereas the f-DP trade-off function of vMF (Dong et al., 2022) has no closed form and must be evaluated numerically. Rényi DP suffices for composition (Lemma 8) and for conversion to (ε,δ)-DP (Lemma 9). Correlated-noise (matrix-mechanism) accounting (Choquette-Choo et al., 2024; 2025), and the noise-reduction and factorization mechanisms it enables (Zhang et al., 2025; Kalinin et al., 2026; Bellet et al., 2026), target Euclidean mechanisms whose noise carries a covariance structure; the vMF/Bingham directional noise has none, so those techniques do not transfer directly; DP-SGD-side refinements (clipping-bias correction (Zhang et al., 2024), feature preprocessing (Sun et al., 2024)) are similarly orthogonal.*

C Privacy Accounting and Sampling

C.1 Inverting the closed-form bound

For fixed (α,θ), $\varepsilon(\alpha,\theta,\kappa)$ of Theorem 1 is strictly increasing in κ, so the concentration realizing a target direction budget $\varepsilon_{\mathrm{dir}}$ is recovered by bisection over $\kappa \in (0, \kappa_{\max}]$, following the numerical-accounting framework of (Koskela et al., 2023). In the small-θ regime, the Gaussian limit of Corollary 2 supplies the closed-form initial guess $\kappa \approx 2\varepsilon_{\mathrm{dir}}/(\alpha\theta^2)$, which the implementation refines by bisection.

C.2 SAMPLING

On $\mathbb{S}^2$ the vMF is sampled exactly by inverting its marginal along μ and rotating (Mardia & Jupp, 2000); on $\mathbb{S}^{d-1}$ for $d \geq 4$ we use Wood's rejection sampler (Wood, 1994). Both samplers are exact and independent of the privacy analysis.

D ADDITIONAL EXPERIMENTAL DETAILS

D.1 IDEAL-TRACKING ASSUMPTION AND A CAVEAT

The leader–follower experiments assume *ideal kinematic tracking*: each follower realizes its fixed offset in the leader's heading frame instantaneously, with no control dynamics, latency, or transient error (the synthetic data adds only $0.5\,\mathrm{m}$ of i.i.d. tracking noise, independent of the privacy noise). This isolates the privacy effect, so the measured formation drift is caused by heading privatization alone. Modeling first-order follower dynamics or one-step latency is orthogonal to the privacy analysis and left to future work; such errors are independent of the privacy noise and superpose on the same follower offset.

D.2 PATEGAIL++

PateGAIL++ (Ma et al., 2026) is a GPU-trained imitation-learning trajectory generator rather than a per-trajectory privatization mechanism, so we restrict the synthesis comparison (Section 6.7) to PrivTrace and defer a head-to-head comparison to future work.

D.3 HYPERPARAMETERS

Throughout, $\alpha = 2$ and $\rho = 0.8$. The direction budget is $\varepsilon_{\mathrm{dir}} = \rho\varepsilon/2$ per heading; the magnitude budget $(1-\rho)\varepsilon$ is split 0.9 to the step lengths and 0.1 to the anchor. The Gaussian variances follow from the RDP $\alpha\Delta^2/(2\sigma^2)$ at the corresponding budget: $\sigma^2 = \alpha\delta^2/\big(0.9(1-\rho)\varepsilon\big)$ for the step lengths, $\sigma_c^2 = \alpha(\delta/T)^2/\big(0.2(1-\rho)\varepsilon\big)$ for the offline centroid, and $\sigma_0^2 = \alpha\delta^2/\big(0.2(1-\rho)\varepsilon\big)$ for the streaming initial point. The heading concentration κ is obtained by bisection (Appendix C.1).

E ADDITIONAL RESULTS

Table 4: Multi-agent swarm (12 agents), $\varepsilon = 1.0$, $\alpha = 2$, $\rho = 0.8$; mean±std over 3 seeds. DDP leads on HeadErr while paying position RMSE; the pure directional mechanism (vMF) has large RMSE by construction.

Method	RMSE↓ (m)	HeadErr↓ (rad)
Laplace3D	$\mathbf{1.44\pm0.05}$	0.394 ± 0.011
Gaussian3D	4.90 ± 0.06	1.019 ± 0.015
Geo-Ind	2.03 ± 0.07	0.545 ± 0.011
DCT-DP	$\mathbf{1.44\pm0.05}$	0.407 ± 0.008
vMF (Weggenmann)	15.73 ± 0.37	1.355 ± 0.044
GDP (Jiang)	17.09 ± 2.85	0.382 ± 0.008
DDP (Ours)	15.50 ± 3.80	$\mathbf{0.362\pm0.011}$

Table 5: Per-trajectory privatization time ($T = 150$ waypoints, mean over 30 repeats). DDP uses exact vMF inversion (≈5 ms per trajectory), far below any streaming sampling rate.

Method	Time (ms)
Laplace3D	0.007
Gaussian3D	0.006
Geo-Ind	0.014
DCT-DP	0.021
GDP (Jiang)	0.585
DDP (Ours)	5.219
vMF (Weggenmann)	5.283

Table 6: Statistical fidelity on 8000 sparse vehicle trajectories, $\varepsilon = 1.0$. Lower is better; bold is best per column. Single fixed seed.

Method	Length JSD↓	Diameter JSD↓	Density ARE↓	Transition ARE↓
PrivTrace	0.164	0.543	1.185	0.935
Laplace3D	0.017	**0.013**	0.034	**0.233**
Gaussian3D	0.105	0.141	0.324	0.926
Geo-Ind	0.050	0.044	0.092	0.539
DCT-DP	**0.015**	0.013	**0.031**	0.257
vMF (Weggenmann)	0.129	0.283	1.501	0.838
GDP (Jiang)	0.110	0.175	0.136	0.734
DDP (Ours)	0.105	0.178	0.140	0.727

F Additional Figures

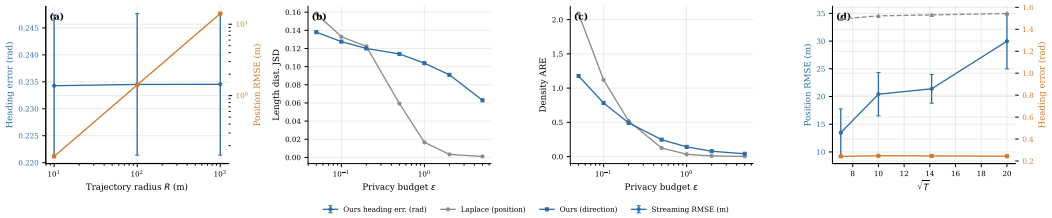

Figure 5: Supporting analyses. (a) Scale invariance: heading error is unchanged as radius R grows two decades, while Laplace3D's RMSE grows linearly. (b)–(c) Statistical fidelity vs ε (log axis): DDP overtakes Laplace below $\varepsilon \approx 0.3$. (d) Streaming drift vs trajectory length: RMSE grows $\Theta(\sqrt{T})$ while DDP's heading error stays ≈ 0.24 rad; a naive per-waypoint release (gray dashed) pays $O(T)$ privacy and its heading saturates near $\pi/2$.