Journal of *Cybersecurity and Privacy*

MDPI

Article

Metric Differential Privacy on the Special Orthogonal Group SO(3)

Anna Katharina Hildebrandt [1,2], Elmar Schömer [3] and Andreas Hildebrandt [2,3,*]

[1] PRAIVACY UG, Science Park 1, 66123 Saarbrücken, Germany; anna.hildebrandt@praivacy.com
[2] Mondata GmbH, Science Park 1, 66123 Saarbrücken, Germany
[3] Institute for Computer Science, Johannes Gutenberg University Mainz, 55122 Mainz, Germany; schoemer@informatik.uni-mainz.de
* Correspondence: andreas.hildebrandt@uni-mainz.de

Abstract

Differential privacy (DP) is an important framework to provide strong theoretical guarantees on the privacy and utility of released data. Since its introduction in 2006, DP has been applied to various data types and domains. More recently, the introduction of metric differential privacy has improved the applicability and interpretability of DP in cases where the data resides in more general metric spaces. In metric DP, indistinguishability of data points is modulated by their distance. In this work, we demonstrate how to extend metric differential privacy to datasets representing three-dimensional rotations in SO(3) through two mechanisms: a Laplace mechanism on SO(3), and a novel privacy mechanism based on the Bingham distribution. In contrast to other applications of metric DP to directional data, we demonstrate how to handle the antipodal symmetry inherent in SO(3) while transferring privacy from S^3 to SO(3). We show that the Laplace mechanism fulfills $\epsilon\phi$-privacy, where ϕ is the geodesic metric on SO(3), and that the Bingham mechanism fulfills $\tilde{\epsilon}\phi$-privacy with $\tilde{\epsilon} = \frac{\pi}{4}\epsilon$. Through a simulation study, we compare the distribution of samples from both mechanisms and argue about their respective privacy–utility tradeoffs.

Keywords: metric differential privacy; Bingham distribution; Laplace mechanism; 3D rotations

Academic Editors: Giorgio Giacinto and Carlo Blundo

Received: 30 July 2024
Revised: 13 May 2025
Accepted: 21 July 2025
Published: 12 August 2025

Citation: Hildebrandt, A.K.; Schömer, E.; Hildebrandt, A. Metric Differential Privacy on the Special Orthogonal Group SO(3). *J. Cybersecur. Priv.* **2025**, *5*, 57. https://doi.org/10.3390/jcp5030057

1. Introduction

In today's connected world, privacy concerns are becoming an increasingly important aspect of systems design; this is particularly evident in applications involving machine learning or artificial intelligence. In such contexts, sophisticated and careful balancing of privacy against usability and utility of data has been the focus of considerable research efforts. Much of the current work on privacy protection relies on the concept of differential privacy (DP) [1], which allows for the quantification of the maximal privacy loss an individual could experience upon data release through a parameter ϵ. Bounding this parameter then allows the maximal privacy loss to be bounded. In practice, these bounds are usually achieved through noise injection, either modifying the data deposited in the central repository or the response to each query of the data.

The notion of privacy in DP hinges upon the notion of indistinguishability: assume that we have two copies D, D' of a database that only differ by the fact that D contains individual d while D' does not. If all possible queries have a high probability of agreement in their results on D and D', user d has plausible deniability: they can argue convincingly that their data was not contained in the database, and thus, their privacy should not be

affected. The parameter ϵ controls the probability with which all queries have to agree, and hence, the level of afforded privacy. However, it is difficult to interpret, so that setting reasonable ϵ-thresholds is challenging [2]. Differential privacy techniques have been proposed for many settings of interest, such as image data [3,4]; textual data [5,6]; and location and trajectory data, e.g., [7].

In practice, however, data is often sampled from a domain that corresponds to some metric space with a meaningful distance between data points. This metric structure is neglected in classical DP, but can greatly improve both the applicability and interpretability of DP applications if the indistinguishability demands depend on the distance between data points. For temporal data, for instance, an attacker might be able to infer the year or month of a particular event with reasonable confidence, but not the day, hour, minute, or second, depending on the threshold ϵ. Computationally, this is achieved through the ideas of *metric differential privacy*. Metric DP mechanisms have been proposed for many metric spaces in the literature, e.g., for location data [2] and directional data [8].

In this work, we consider the problem of protecting three-dimensional rotation data through metric DP. Such rotations occur in many settings of interest. In protein docking, for instance, actors might be interested in allowing access to broad representations of the docking mode, but not the detailed three-dimensional structure. In computer-aided ergonomics, workers might have an interest in obtaining recommendations on how to improve posture and reduce strain, while simultaneously preventing others identifying that they have a high probability of strain-induced medical problems. In all practical applications, it is crucial to balance privacy with utility [9], which requires the use of a suitable metric.

2. Background

In this section, we recapitulate the mathematical structure of rotations in $3D$-space and important results from the theory of directional differential privacy.

2.1. Rotations in $\mathbb{R}^3$

Rotations in three-dimensional space can be modeled in a number of different ways, each with its own advantages and drawbacks. Possibly the most familiar representation of rotations in $\mathbb{R}^3$ is *special orthogonal* 3×3 *matrices*, i.e., those matrices $R \in \mathbb{R}^{3\times3}$ with $RR^T = \mathbb{1}$ and $\det(R) = 1$ (intuitively, this corresponds to the observation that a rotation only changes the orientation, not the length, of vectors). These matrices form the *special orthogonal group* SO(3). From the matrix representation, it might appear that rotations in $3D$-space are 9-dimensional objects, while in reality, the constraints on these matrices reduce them to a 3-dimensional submanifold of $\mathbb{R}^4$. This can be seen from the *axis–angle* representation: according to Euler's rotation theorem (also known as the *football theorem*), each rotation in three dimensions can be represented by a single rotation axis r and an angle α. The normalization constraint on r effectively reduces the four dimensions (three entries of r and one angle α) to three. Due to normalization, the rotation axis lies on the unit sphere in three dimensions: $r \in S^2$. Since the rotation angle α is constrained to the unit circle, $\alpha \in S^1$, we can represent each rotation as an element of $(S^2 \times S^1)$. Please note that this does not imply that SO(3) is diffeomorphic to $(S^2 \times S^1)$, and in fact, it is not. The reason for this possibly surprising fact is the antipodal symmetry encoded in SO(3): rotations by π and $-\pi$ are identical.

Yet another alternative to represent rotations in $\mathbb{R}^3$ uses Hamilton's *quaternions* $q \in \mathbb{H}$. Quaternions generalize the complex numbers, which have one real and one imaginary axis, to a four-dimensional structure with one real and three imaginary axes, the basic quaternions $\mathbf{i}, \mathbf{j}, \mathbf{k}$. Each quaternion can thus be represented by four real numbers (a, b, c, d),

with $q = a + b\mathbf{i} + c\mathbf{j} + d\mathbf{k}$. Multiplication of two quaternions $p = (p_1, p_2, p_3, p_4)$ and $q = (q_1, q_2, q_3, q_4)$ is performed using the *Hamilton product*:

$$p \odot q = \begin{pmatrix} p_1q_1 - p_2q_2 - p_3q_3 - p_4q_4 \\ p_1q_2 + p_2q_1 + p_3q_4 - p_4q_3 \\ p_1q_3 - p_2q_4 + p_3q_1 + p_4q_2 \\ p_1q_4 + p_2q_3 - p_3q_2 + p_4q_1 \end{pmatrix} = \underbrace{\begin{pmatrix} p_1 & -p_2 & -p_3 & -p_4 \\ p_2 & p_1 & -p_4 & p_3 \\ p_3 & p_4 & p_1 & -p_2 \\ p_4 & -p_3 & p_2 & p_1 \end{pmatrix}}_{=:\tilde{\mathbf{P}}} \cdot \begin{pmatrix} q_1 \\ q_2 \\ q_3 \\ q_4 \end{pmatrix} = \tilde{\mathbf{P}} \cdot q$$

with $\tilde{\mathbf{P}}^T\tilde{\mathbf{P}} = \mathbb{1}_{4\times4}$. Equivalently,

$$p \odot q = \underbrace{\begin{pmatrix} q_1 & -q_2 & -q_3 & -q_4 \\ q_2 & q_1 & q_4 & -q_3 \\ q_3 & -q_4 & q_1 & q_2 \\ q_4 & q_3 & -q_2 & q_1 \end{pmatrix}}_{=:\mathbf{Q}} \cdot \begin{pmatrix} p_1 \\ p_2 \\ p_3 \\ p_4 \end{pmatrix} = \mathbf{Q} \cdot p$$

with $\mathbf{Q}^T\mathbf{Q} = \mathbb{1}_{4\times4}$. Conjugation yields the quaternion $\tilde{q} = (q_1, -q_2, -q_3, -q_4)$, and the norm is defined through $||q|| = \sqrt{q \odot \tilde{q}} = \sqrt{q_1^2 + q_2^2 + q_3^2 + q_4^2}$. Quaternions with unit norm $||q|| = 1$ are known as *versors*. Versors are isomorphic to the *special unitary group* $SU(2)$, the group of complex 2×2 matrices $M \in \mathbb{C}^2$. It can be shown that $SU(2)$ provides a *double cover* of SO(3), i.e., every element of SO(3) is represented by two different elements of $SU(2)$, and hence by two unit quaternions—a result of antipodal symmetry.

Let $q = (w, x, y, z)$ denote a versor. Then, for all *pure* quaternions v, i.e., quaternions with real part 0, the map $v \mapsto q \odot v \odot \tilde{q}$ is a rotation of v. Here, pure quaternions correspond to vectors in $\mathbb{R}^3$. Explicitly computing the Hamilton products in this map yields the rotation matrix:

$$Q = \begin{bmatrix} 1 - 2y^2 - 2z^2 & 2xy - 2zw & 2xz + 2yw \\ 2xy + 2zw & 1 - 2x^2 - 2z^2 & 2yz - 2xw \\ 2xz - 2yw & 2yz + 2xw & 1 - 2x^2 - 2y^2 \end{bmatrix}$$

which is equivalent to a rotation around the vector (x, y, z) with angle $2\arccos(w)$. Note that replacing q by $-q$ in the map yields the same rotation, which shows the two-to-one mapping due to antipodal symmetry. Interestingly, *locally*, in a sufficiently small neighborhood around the identity, SO(3) and $SU(2)$, and hence the unit quaternions, are indeed isomorphic.

2.2. Distances Between 3D Rotations

Arguing about privacy and indistinguishability hinges upon a suitable notion of distance between two instances. To state, e.g., that the orientations of the skeleton of individual A are hard to distinguish from those of individual B requires us to quantify how *close* they are. Since rotations in three dimensions, i.e., members of the special orthogonal group SO(3), can be represented as matrices (c.f. Section 2.1), it might seem reasonable to measure distances in terms of established matrix norms (e.g., the p-norms, the max-norm, or the Frobenius-norm); however, these approaches would not respect the *topology* of SO(3). Topologically, SO(3) is a ball with antipodal symmetry. The effect of the first aspect on distance measurement can be roughly described as follows: imagine a unit ball in three dimensions, and two points on its surface. Measuring distances with a metric taken from $\mathbb{R}^n$ ignores that the points live on a sphere and instead connects them using the shortest line in $\mathbb{R}^n$, which cuts *through* the sphere. If we would, e.g., try to describe an equivalent to a Gaussian distribution of 3D rotations in this way, the resulting function

would decay exponentially in the length of this straight line, neglecting that the two points can only be connected through paths on the surface. To give an even simpler example: the cities of Madrid, Spain, and Wellington, New Zealand, are near antipodes on the surface of the Earth. Measuring in $\mathbb{R}^3$, they are connected by a straight line of roughly 12,740 km in length. To reach Wellington from Madrid, though, will require a much longer distance to be traveled, and the shortest flight path between the cities—a geodesic on the surface of the Earth—has a length of roughly 20,000 km. A metric used for privacy should clearly respect the topology to be of use in practice, and Ref. [8] describes suitable distance metrics and privacy mechanisms for this case. These metrics, however, ignore the second aspect of the topology of SO(3), i.e., the antipodal symmetry. In a simplified picture (ignoring for a moment that Earth is $S(2)$), Madrid and Wellington would denote the same point. Designing metrics that capture those aspects is not trivial, and several alternative approaches have been described in the literature. In [10], Huynh defines and analyzes a variety of such metrics, each highlighting a different aspect of the topology. Importantly, a subset of the metrics defined there—namely, ϕ_2, ϕ_3, ϕ_5, and ϕ_6—are shown to be *boundedly equivalent* and *functionally equivalent*, where functional equivalence of a metric ϕ and ψ implies that there is a positive continuous strictly increasing function h such that

$$h \circ \phi = \psi$$

and bounded equivalence implies that there are positive real numbers $a, b \in \mathbb{R}^+$ such that

$$a\phi(\mathbf{R}_1, \mathbf{R}_2) \leq \psi(\mathbf{R}_1, \mathbf{R}_2) \leq b\phi(\mathbf{R}_1, \mathbf{R}_2)$$

$\forall \mathbf{R}_1, \mathbf{R}_2 \in SO(3)$. For metric differential privacy, this implies that if a mechanism can be shown to fulfill metric differential privacy on one of these metrics, it is private on all of them. We now briefly define the metrics of interest:

$$\phi_2(q_1, q_2) = \min\{||q_1 - q_2||, ||q_1 + q_2||\} \tag{1}$$
$$\phi_4(q_1, q_2) = 1 - |q_1^T q_2| \tag{2}$$
$$\phi_5(\mathbf{R}_1, \mathbf{R}_2) = ||\mathbf{I} - \mathbf{R}_1 \mathbf{R}_2^T||_F \tag{3}$$
$$\phi_6(\mathbf{R}_1, \mathbf{R}_2) = ||\log\left(\mathbf{R}_1 \mathbf{R}_2^T\right)|| \tag{4}$$

ϕ_5 has an alternative definition in terms of versors:

$$\phi_5'(q_1, q_2) = 2\sqrt{2(1 - |q_1^T q_2|^2)}$$

In the following, we call this metric Φ and use it in our proofs. In practice, the most natural norm is ϕ_6, as this corresponds to geodesic distance on SO(3). Since ϕ_6 is boundedly equivalent to Φ, our proof will also show metric differential privacy on ϕ_6, and hence, on geodesic distance.

2.3. Preservation of Privacy

In line with current state-of-the-art research, the kind of privacy afforded by our technique is one of plausible deniability. Hence, the goal of the mechanisms described in this work is to allow an individual to deny that a rotation was generated by them, as every 'similar' individual (subsequently we precisely define the meaning of similar) could just as plausibly have generated the same.

2.4. Random Mechanisms

The concept of a random mechanism is central to our discussion of privacy-preserving data analysis. Let $\mathcal{X}$ and $\mathcal{Y}$ be two sets, and let $\mathbb{D}\mathcal{Y}$ be the set of probability distributions on $\mathcal{Y}$. A **random mechanism** $\mathcal{M}$ from $\mathcal{X}$ to $\mathcal{Y}$ is a probabilistic mapping from $\mathcal{X}$ (the set of inputs) to $\mathcal{Y}$ (the set of outputs), i.e., $\mathcal{M}$ is a function $\mathcal{M} : \mathcal{X} \to \mathbb{D}\mathcal{Y}$. For every $x \in \mathcal{X}$, $\mathcal{M}(x)$ is a probability distribution and $\mathcal{M}(x)(y)$ is the probability that the mechanism assigns output $y \in \mathcal{Y}$ to input $x \in \mathcal{X}$. Random mechanisms are often generated by families of parameterized probability distributions $\mathrm{M}(x)$, where we call the mechanism that assigns $\mathrm{M}(x)$ to x the mechanism **induced** by $\mathrm{M}(x)$. *Applying* or *running* a mechanism on an input x corresponds to sampling a realization z from $\mathrm{M}(x)$, i.e., $z \sim \mathrm{M}(x)$.

The idea behind using a random mechanism for privacy purposes is to only release data perturbed by the mechanism to protect the sensitive inputs. In **local privacy** workflows, each owner of a sensitive data point x (e.g., end users that want to protect the data they share with a cloud service) runs the mechanism on x and only releases the perturbed sample z drawn from $\mathcal{M}(x)$. In a **central privacy** workflow, on the other hand, the data is instead collected without perturbation, but the provider of the database or service runs the mechanism on the result $Q(x)$ of each query Q posed to the data and only releases the perturbed sample Q_z drawn from $\mathcal{M}(Q(x))$. In this model, $\mathcal{M}$ is also called an oblivious mechanism.

2.5. Differential Privacy

Differential privacy is a property of certain random mechanisms that allows the degree by which the mechanism affords plausible deniability to be quantified. Originally, differential privacy was formulated for oblivious mechanisms in a central model. Let $\mathcal{D}$ denote the set of all datasets, and let $\mathcal{Y}$ be the set of possible query results on datasets from $\mathcal{D}$. Two datasets d, d' are called *adjacent* if they differ in at most one record. A mechanism $\mathcal{M} : \mathcal{D} \to \mathbb{D}\,\mathcal{Y}$ is said to satisfy *ϵ-differential privacy (ϵ-DP)*, if for all *adjacent* datasets d, d', and for all sets of possible outputs $Y \in \mathcal{Y}$,

$$\mathcal{M}(d)(Y) \leq e^{\epsilon} \cdot \mathcal{M}(d')(Y) \tag{5}$$

If the *privacy budget* ϵ is sufficiently small, this relation implies that the probability of each possible query result is so similar when running the mechanism on d and d' that an attacker cannot reasonably deduce whether the input dataset was d or d'. Since this holds for each pair of adjacent datasets, and all query results, each user can plausibly deny that their data was contained in the dataset (or that the value of their data was x; or a number of similar statements), since the query result could just as likely have been produced from a dataset that differed from d by only removing (or changing) the data of this individual.

2.6. Differential Privacy on Metric Spaces

While differential privacy is a powerful technique for central privacy workflows, it also has a number of drawbacks. Most importantly, it does not easily translate to local privacy workflows, where each user wants to perturb their data before releasing them into the central database. In addition, the interpretation of the privacy budget is rather unintuitive in the sense that it is often hard for the user to relate the value of ϵ to the privacy of their data points. To overcome these challenges, Chatzikokolakis et al. [11] generalized differential privacy to metric spaces $(\mathcal{X}, d)$, where d is a metric on the set $\mathcal{X}$. The idea is as follows: given a metric space (note: setting $\mathcal{X} = \mathcal{D}$ and $d = d_H$ to the *Hamming distance* on $\mathcal{D}$, we recover the classical ϵ-DP) $(\mathcal{X}, d)$, a set $\mathcal{Y}$, and a privacy budget $\epsilon \geq 0$, a random mechanism $\mathcal{M} : \mathcal{X} \to \mathbb{D}Y$ satisfies ϵd-differential privacy if and only if for all pairs of inputs $x, x' \in \mathcal{X}$, and for all sets of possible outcomes $Y \in \mathcal{Y}$,

$$\mathcal{M}(x)(Y) \leq e^{\epsilon \cdot d(x,x')} \cdot \mathcal{M}(x')(Y) \tag{6}$$

Intuitively, if ϵ is sufficiently small, similar inputs—i.e., inputs with sufficiently small distance d—cannot reasonably be distinguished from each other. In a local privacy model, where each user randomly perturbs their data through $\mathcal{M}$ before releasing it into the dataset, the privacy protection afforded by $\mathcal{M}$ can be imagined as a sphere in the metric space $(\mathcal{X}, d)$ within which the user can claim plausible deniability: all points inside a sphere with radius r around x, i.e., all points in $x' \in \mathcal{X}$ with $d(x, x') \leq r$, have a level of distinguishability of $\epsilon d(x, x') \leq \epsilon r$. We call ϵr the *privacy level* ℓ within the *protection radius* r.

Consider, e.g., that the user wants to protect their location data before releasing them to a cloud service. The degree to which their true location x can then be reasonably distinguished from all other locations x' then depends on the privacy parameter ϵ of the protection mechanism, and on the distance $d(x, x')$. Keeping a fixed privacy level l, shrinking the value of the privacy parameter ϵ increases the protection radius and hence reduces the spatial resolution at which their location can be reasonably inferred.

2.7. Privacy–Utility Tradeoff

Obviously, perturbing the raw data influences the quality or *utility* of downstream algorithms and analyses. Balancing the privacy requirements of the data owners with the utility demands of the application is difficult in practice. It can be shown (c.f. [2]) that, in general, both properties cannot be simultaneously optimized, leading to a difficult *privacy–utility tradeoff*. Here, one advantage of metric DP compared to classic DP is the simplified interpretability of privacy budgets in terms of spheres of a certain radius in $\mathcal{M}$.

2.8. Metric Differential Privacy of Directional Data

In this work, we consider datasets sampled from SO(3), i.e., measurements of three-dimensional rotations. This kind of data falls into the realm of *directional statistics*, where only the orientation of data points, but not their magnitude, is of interest. Recently, Weggenmann and Kerschbaum discussed several metric differential privacy mechanisms for directional datasets in the context of measurements on the $n-1$-dimensional unit sphere S^{n-1} [8]. As discussed there, commonly used approaches are based on extending the classical real-line Laplace or the planar Laplace mechanism to directional data by performing a post-processing to ensure the correct periodicity, but they can perform even worse than simple uniform noise generation. These results, which are demonstrated by simulations in [8], show that while the general Laplace mechanisms as defined in [11] can be used for directional statistics, they often lead to sub-optimal results in terms of the privacy–utility tradeoff. Instead, they propose two privacy mechanisms that are directly defined on the unit $n-1$-sphere: the *VMF* mechanism and the *Purkayastha* mechanism.

2.9. The VMF Mechanism

The VMF mechanism, introduced in [8], samples from the well-known *von Mises–Fisher* distribution on S^{n-1}, which is defined as follows:

$$VMF(\mu, \kappa)[x] = C_{VMF}(n, \kappa) \cdot e^{\kappa \cdot \mu^T x} \tag{7}$$

where, with $\nu(n) := \frac{n}{2} - 1$, and with the *modified Bessel function of the first kind* I_α,

$$C_{VMF}(n, \kappa) = \frac{\kappa^{\nu(n)}}{(2\pi)^{\nu(n)+1} I_{\nu(n)}(\kappa)} \tag{8}$$

Weggemann and Kerschbaum show that the VMF mechanism on S^{n-1}, i.e., the mechanism induced by the VMF distribution, fulfills ϵd_2-privacy, where d_2 is the Euclidean 2-norm. This norm, however, 'uses' the shortest distance in the n-dimensional Euclidean space, in the sense that the privacy guarantee is exponential in this metric. In this norm, the distance of two points $x, y \in S^{n-1}, x \neq y$, is measured from the straight line between x and y, not from the geodesic on S^{n-1}. The geodesic distance between two points on any $n-1$-sphere with radius r can be computed from

$$d_{\angle}(x,y) = r \arccos\left(\frac{x^T y}{r^2}\right) \tag{9}$$

Obviously, the Euclidean distance cannot exceed the geodesic distance, i.e., $d_2(x,y) \leq d_{\angle}(x,y)$, and hence, the VMF mechanism also fulfills $\epsilon d_{\angle}$-privacy, where the privacy parameter ϵ equals the concentration parameter κ.

2.10. The Purkayastha Mechanism

An alternative mechanism, also proposed in [8], is based on the Purkayastha distribution on S^{n-1}, defined as

$$Pur(\mu,\kappa)[x] = C_{\text{Pur}}(n,\kappa) \cdot e^{-\kappa \cdot \arccos\left(\mu^T x\right)} \tag{10}$$

where $C_{\text{Pur}}(n,\kappa) = S_{n-2}^{-1} F_{n-1,-\kappa}^{-1}(\pi)$ with the Riccati–Bessel function of the first kind S_n, and with

$$F_{n-2,-\kappa}^{-1}(\pi) = \begin{cases} \frac{\kappa(\kappa^2+2^2)(\kappa^2+4^2)\cdots(\kappa^2+(n-2)^2)}{(n-2)!(1-e^{-\kappa\pi})} & n \text{ even} \\ \frac{(\kappa^2+1^2)(\kappa^2+3^2)\cdots(\kappa^2+(n-2)^2)}{(n-2)!(1+e^{-\kappa\pi})} & n \text{ odd} \end{cases} \tag{11}$$

The Purkayastha mechanism on S^{n-1} is then defined as the mechanism induced by the Purkayastha distribution.

Weggemann and Kerschbaum also show that the Purkayastha mechanism fulfills $\epsilon d_{\angle}$-privacy, where again the privacy parameter ϵ equals the concentration parameter κ.

The discussion in [8] does not seem to show a clear advantage for either the VMF or Purkayastha mechanisms. However, for equal privacy parameters, the Purkayastha samples seem to be concentrated closer to the mean, at least for larger κ-values, which would indicate a better privacy–utility tradeoff.

3. Results

3.1. Privacy Mechanisms on SO(3)

As the SO(3) is locally isomorphic to SU(2), it might be tempting to use the privacy mechanisms defined in [8] on S^3, which is diffeomorphic to SU(2) and hence locally isomorphic to SO(3) [12]. It thus seems reasonable that, for sufficiently small concentration parameters or, equivalently, privacy parameters, the differential privacy guarantees of the mechanisms on S^3 should carry over to SO(3). But globally, the morphism from S^3 to SO(3) yields a double cover, where each rotation in SO(3) is generated from exactly two points on S^3. The main challenge in transferring the differential privacy guarantee from S^3 to SO(3) thus lies in the different metrics. While for directional data on S^n we are mostly interested in the geodesic distance $d_{\angle}$, the double covering of SO(3) by S^3 would make a direct application of $d_{\angle}$ misleading when representing rotations. Let $q_1, q_2 \in \mathbb{H}, ||q_1|| = ||q_2|| = 1$ be two unit quaternions representing rotations $\in SO(3)$. Taking antipodal symmetry into account leads to the definition of a typical metric on SO(3):

$$d_{\angle}^{SO(3)}(q_1, q_2) = \min\{\arccos(q_1 \cdot q_2), \pi - \arccos(q_1 \cdot q_2)\} = \arccos(|q_1 \cdot q_2|) \quad (12)$$

which corresponds to the metrics $\phi_3(q_1, q_2)$ and $\phi_3'(q_1, q_2)$ in [10]. The flipping of the sign when sampling quaternions from different hemispheres is difficult to incorporate into the VMF or Purkayastha mechanisms.

3.2. The Bingham Mechanism

Instead of pulling mechanisms from S^3 to SO(3), we thus focus on the design of a mechanism specifically adapted to the metric structure of SO(3). The simplest way of constructing such a mechanism is to identify suitable probability distributions on SO(3) that intrinsically respect antipodal symmetry. The obvious candidate for such a distribution is the *Bingham* distribution, which arises when an n-dimensional Gaussian distribution is conditioned to lie on S^{n-1}. It is defined as

$$\mathrm{BD}(\mathbf{M}, \mathbf{Z})[x] = C_{\mathrm{BD}} e^{x^T \mathbf{M}\mathbf{Z}\mathbf{M}^T x} \quad (13)$$

where $\mathbf{M}$ is an orthogonal $n \times n$ matrix, $\mathbf{Z}$ a diagonal matrix with increasing entries ($\mathbf{M}$ and $\mathbf{Z}$ arise from diagonalizing the Gaussian distribution that underlies BD) $z_1 \leq z_2 \leq \cdots \leq z_n$, and the normalization factor C_{BD} can be computed from

$$C_{\mathrm{BD}} = \int_{S^{n-1}} e^{x^T \mathbf{Z} x} dx \quad (14)$$

which can be evaluated to

$$C_{\mathrm{BD}} = |S^{n-1}| \, {}_1F_1\left(\frac{1}{2}, \frac{n}{2}, \mathbf{Z}\right) \quad (15)$$

where ${}_1F_1$ is a hypergeometric function of a matrix argument (c.f. [13]). It is straightforward to see that the Bingham distribution on SO(3) has built-in antipodal symmetry, as, obviously, $\mathrm{BD}(\mathbf{M}, \mathbf{Z})[x] = \mathrm{BD}(\mathbf{M}, \mathbf{Z})[-x]$. Representing quaternions as 4-vectors, we can thus draw versors from a Bingham distribution: the results are constrained to live on S^3 (like versors) and have antipodal symmetry (so we can use them to represent rotations $\in SO(3)$). To see why this latter part is important, imagine drawing versors from a probability distribution that does not account for antipodal symmetry. If these samples were used to ensure privacy through adding noise, the mechanism would consider a quaternion from the other hemisphere to be very different and hence assume a high level of privacy, while the rotation represented by the quaternion might actually be arbitrarily close or even identical to the original rotation due to symmetry.

The orthonormal matrix $\mathbf{M}$ can be used to set the mean of the distribution. This can be seen from the fact that BD has its maxima at the points $\pm\mathbf{M}(0, 0, 0, 1)^T$ [13]. Setting the last column of $\mathbf{M}$ to a versor $\hat{q}$, we can fill the first $n - 1$ columns by orthonormalization, e.g., by computing the null space of the matrix $\hat{q}\hat{q}^T$, and find that

$$\arg\max[\mathrm{BD}(\mathbf{M}, \mathbf{Z})] = \pm\hat{q}$$

This setting of the mean is important for our application in a differential privacy mechanism: if we want to add noise to a given rotation represented by the versor $\hat{q}$, we only need to sample from the Bingham distribution centered around this rotation.

As also noted in [13], working with such Bingham distributions to sample rotations can be further simplified by a change in convention: permuting the order of the vector representing the quaternion such that the real part is stored in the last instead of the first position. I.e., by storing the quaternion in the order (q_2, q_3, q_4, q_1), we find that setting $\mathbf{M}$ to

the identity $\mathbf{M} = \mathbf{I}$ leads to a Bingham distribution with mean zero: $(0, 0, 0, 1)$. For the rest of this section, we follow this convention.

While the matrix $\mathbf{M}$ defines the mean of the Bingham distribution, the diagonal matrix $\mathbf{Z}$ captures its spread. The entries $z_1, \ldots, z_n$ are consequently known as the *concentration parameters* of the Bingham distribution. It can be shown [13] that for all $c \in \mathbb{R}$,

$$\mathrm{BD}(\mathbf{M}, \mathbf{Z}) = \mathrm{BD}(\mathbf{M}, \mathbf{Z} + c\mathbf{I}) \tag{16}$$

This fact is usually used to fix the value of z_n to zero, and hence constrain all concentration parameters to be negative ($z_1 \leq z_2 \leq \cdots \leq z_n = 0$). We also follow this convention.

Definition 1. *Let q be a versor, represented as a 4-dimensional vector in the order (q_2, q_3, q_4, q_1). Compute the null space K of the matrix qq^T and define the orthogonal matrix $\mathbf{M}(q)$ as the column-wise concatenation of the three vectors in K and q. For $\epsilon > 0$, let $\mathbf{Z}(\epsilon) = -\mathrm{Diag}([\epsilon, \ldots, \epsilon, 0]) =: -\mathbf{D}(\epsilon)$. Then, the Bingham mechanism on SO(3) is the induced mechanism of $BD(\mathbf{M}(q), \mathbf{Z}(\epsilon))$.*

As a metric on SO(3), we choose (with $q^T \cdot p$ representing the dot-product of the quaternions q and p, taken as regular 4-vectors)

$$\Phi(q, p) := \sqrt{2(1 - |q^T \cdot p|)}$$

which corresponds to the metric ϕ_5' in [10].

Theorem 1. *The Bingham mechanism on SO(3) fulfills $\sqrt{2}\epsilon\Phi$-privacy.*

Proof of Theorem. To simplify the proof, we fix $\mathbf{M} = \mathbf{I}$ by first sampling $p \sim \mathrm{BD}(\mathbf{I}, \mathbf{Z}(\epsilon))$ and then rotating p by q to yield $r := q \odot p$. Then, $p = \tilde{q} \odot r$. We thus find

$$\mathcal{M}(q, \epsilon)[r] = \frac{1}{N(\mathbf{Z}(\epsilon))} e^{(\tilde{q} \odot r)^T \mathbf{Z}(\epsilon)(\tilde{q} \odot r)}$$

and thus, for two versors, $q_1, q_2 \in \mathbb{H}, ||q_1|| = ||q_2|| = 1$,

$$\begin{aligned}
\ln\left[\frac{\mathcal{M}(q_1, \epsilon)[r]}{\mathcal{M}(q_2, \epsilon)[r]}\right] &= (\tilde{q}_1 \odot r)^T \mathbf{Z}(\epsilon)(\tilde{q}_1 \odot r) - (\tilde{q}_2 \odot r)^T \mathbf{Z}(\epsilon)(\tilde{q}_2 \odot r) \\
&= (\tilde{q}_1 \odot r - \tilde{q}_2 \odot r)^T \mathbf{Z}(\epsilon)(\tilde{q}_1 \odot r + \tilde{q}_2 \odot r) && \text{[P1]} \\
&= \{(\tilde{q}_1 - \tilde{q}_2) \odot r\}^T \mathbf{Z}(\epsilon)\{(\tilde{q}_1 + \tilde{q}_2) \odot r\} && [\odot \text{ distributive over } +, -] \\
&= -\{(\tilde{q}_1 - \tilde{q}_2) \odot r\}^T \mathbf{D}(\epsilon)\{(\tilde{q}_1 + \tilde{q}_2) \odot r\} && \text{[Definition of } \mathbf{D}\text{]} \\
&\leq \left|\{(\tilde{q}_1 - \tilde{q}_2) \odot r\}^T \mathbf{D}(\epsilon)\{(\tilde{q}_1 + \tilde{q}_2) \odot r\}\right| \\
&\leq ||(\tilde{q}_1 - \tilde{q}_2) \odot r|| \cdot ||\mathbf{D}(\epsilon)(\tilde{q}_1 + \tilde{q}_2) \odot r|| && \text{[Cauchy-Schwarz]} \\
&\leq ||(\tilde{q}_1 - \tilde{q}_2) \odot r|| \cdot ||(\tilde{q}_1 + \tilde{q}_2) \odot r|| \cdot ||\mathbf{D}(\epsilon)||_2 && \text{[Cauchy-Schwarz]} \\
&= ||\tilde{q}_1 - \tilde{q}_2|| \cdot ||\tilde{q}_1 + \tilde{q}_2|| \cdot ||\mathbf{D}(\epsilon)||_2 && \text{[P2]} \\
&= 2\sqrt{(1 - |\tilde{q}_1^T \tilde{q}_2|^2)} \cdot ||\mathbf{D}(\epsilon)||_2 && \text{[P3]} \\
&= 2\sqrt{(1 - |\tilde{q}_1^T \tilde{q}_2|^2)} \cdot \sqrt{\epsilon^2} && \text{[Definition of } \mathbf{D}\text{]} \\
&= \frac{2\epsilon}{\sqrt{2}}\sqrt{2(1 - |\tilde{q}_1^T \tilde{q}_2|^2)} \\
\Rightarrow \frac{\mathcal{M}(q_1, \epsilon)[r]}{\mathcal{M}(q_2, \epsilon)[r]} &\leq \exp\left[\sqrt{2}\epsilon\Phi(q_1, q_2)\right] \quad \square && \text{[Definition of } \Phi\text{, P2]}
\end{aligned}$$

where $||\mathbf{A}||_2$ is the *spectral norm* of the matrix $\mathbf{A}$, i.e., $||\mathbf{A}||_2 := \sqrt{\mu_{\max}}$, where $\mu_{\max}$ is the largest eigenvalue of $\mathbf{A}^t\mathbf{A}$. □

Corollary 1. *The Bingham mechanism fulfills ϵd-privacy for $d \in \{\phi_2, \phi_3, \phi_5, \phi_6\}$ as defined in [10].*

Proof of Corollary. As shown in [10], the norms ϕ_2, ϕ_3, ϕ_5, and ϕ_6 defined there are boundedly equivalent. The norm Φ used in this manuscript corresponds to the norm ϕ_5' defined in [10], which is merely a reformulation of ϕ_5 in terms of versors. □

To sample from the Bingham distribution, we follow the rejection scheme described in [14], which we reproduce in Appendix D. Figure 1 demonstrates the effect of the privacy parameter ϵ. In each subfigure, the action of the unmodified quaternion on the standard coordinate system is depicted in bold, while the actions of the samples from the corresponding Bingham distribution are shown as thinner arrows. Colors label corresponding axes. As expected, smaller values of the privacy parameter correspond to higher levels of noise, and hence to increased privacy.

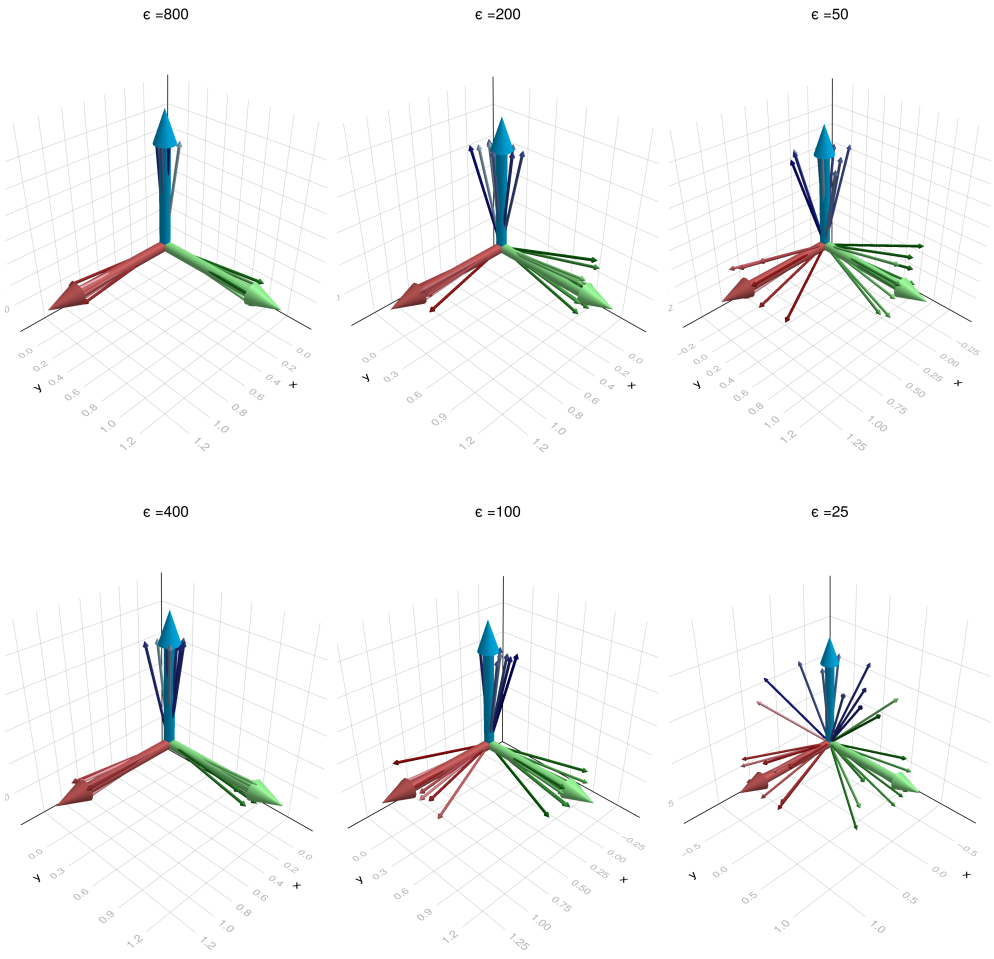

Figure 1. Samples drawn from Bingham distributions for different values of the concentration parameter ϵ. As expected, smaller concentration values correspond to a wider spread of the data, and correspondingly to a higher level of privacy when used as a DP mechanism.

3.3. A Laplace Mechanism on SO(3)

Instead of designing a mechanism specifically adapted to the intricacies of the special orthogonal group SO(3), we can also attempt to design a Laplace mechanism for our

purposes. In classic differential privacy, one of the most popular mechanisms to ensure a privacy budget of ϵ works by adding noise drawn from a suitable Laplace distribution with zero mean and scale parameter b: $\mathrm{Lap}(x|b) = \frac{1}{2b}\exp\left(\frac{-|x|}{b}\right)$. If we denote the query as f, setting $b = \frac{\Delta f}{\epsilon}$, where Δf is the L_1-sensitivity of f, achieves the desired ϵ-privacy. In the case of *metric* differential privacy, we typically want to argue about the intrinsic properties of the mechanism, not about properties of the query or of any associated sensitivities. In [11], Chatzikolakis et al. generalize the Laplace mechanism to arbitrary metric spaces. Here, the key construction is a metric-dependent scaling function λ. The idea given in [11] is as follows:

Definition 2. *Let $\mathcal{Y}, \mathcal{Z}$ be two sets and $d_{\mathcal{Y}}$ be a metric on $\mathcal{Y} \times \mathcal{Z}$. Let $\lambda : \mathcal{Z} \to [0, \infty)$ be a scaling function that normalizes the exponential distribution with respect to $d_{\mathcal{Y}}$, i.e., $\int_{\mathcal{Z}} \lambda(z)\exp[-d_{\mathcal{Y}}(y,z)]\nu(dz) = 1 \quad \forall y \in \mathcal{Y}$, where $\nu(z)$ is a suitable base measure for integration on $\mathcal{Z}$. Obviously, $D(y)(z) := \lambda(z)\exp[-d_{\mathcal{Y}}(y,z)]\nu(z)$ is a probability density function, and the mechanism $\mathcal{L} : \mathcal{Y} \to \mathscr{P}(\mathcal{Z})$ induced by D is called a **Laplace mechanism** from $(\mathcal{Y}, d_{\mathcal{Y}})$ to $\mathcal{Z}$.*

Crucially, Ref. [11] shows that all such mechanisms $\mathcal{L}$ fulfill $d_{\mathcal{Y}}$-differential privacy. Accordingly, designing a differentially private mechanism for an arbitrary metric might seem a simple task—we merely need to determine the scaling and the normalization, and directly obtain the corresponding Laplace mechanism. In the Euclidean case, this simplicity arguably holds. In our case, though, we have $\mathcal{Y} = \mathcal{Z} = SO(3)$, which is a non-Euclidean manifold with non-trivial topology, and we need to carefully consider how to choose the base measure ν, how to sample from D, and which metric $d_{\mathcal{Y}}$ to choose. If we pick an unsuitable metric, the privacy guarantee has no immediate meaning in terms of the data. Choosing the wrong base measure will locally undercount or overcount the *amount* of rotations per volume element and—just like an unsuitable sampling mechanism—lead to noisy samples that cannot guarantee the claimed privacy protection.

Trivially, one can attempt to embed the data in a fitting Euclidean space and use some kind of post-processing to ensure that the resulting noisy samples lie on the manifold of interest. In that case, Euclidean distance between the vectors representing the parameters of the rotation—say, Euler angles—could be used as a metric. This would lead to a trivial Laplace mechanism, but as described above and demonstrated in [8], the resulting mechanisms would be unsuitable, as they measure distance "through" the manifold, not "on" the manifold. The obvious choice is thus to use the geodesic distance on the manifold—the length of the shortest path on the manifold connecting the two rotations. The geodesic can be derived from the natural Riemannian metric of SO(3), i.e., from an inner product on the tangent bundle. It is straightforward to show [10] that for SO(3), the resulting geodesic distance of two rotations $\mathbf{R}_1$ and $\mathbf{R}_2$ is the magnitude of the angle of the rotation turning $\mathbf{R}_1$ onto $\mathbf{R}_2$, and that this magnitude is given by

$$\phi_6(\mathbf{R}_1, \mathbf{R}_2) = ||\log\left(\mathbf{R}_1\mathbf{R}_2^T\right)||$$

as used in our previous derivations. This motivates the representation of the rotations in SO(3) in axis–angle form.

A suitable base measure should allow us to sample *uniformly* from SO(3). In the Euclidean case, the concept of uniformity is immediately apparent. Even in the case of 2D rotations, it seems logical that a uniform rotation can be found by sampling the rotation angle uniformly from $[0, 2\pi)$. For higher-dimensional rotations, however, this intuition breaks down. On the SO(3), in particular, if we write the rotation in axis–angle form,

choosing the angle uniformly from $[0, 2\pi)$ does **not** lead to uniformly distributed rotations. As SO(3) forms a connected and locally compact *Lie group*, the usual way to generalize uniformity is to demand that the volume element defined by the measure is invariant with respect to actions of members of the group, i.e., the distribution should not change after performing an arbitrary rotation. The corresponding measure is called the **Haar measure** of the Lie group [15].

As derived in [16], the Haar measure on SO(3) in axis–angle form with axis $\mathbf{n}$ and angle θ is given by

$$d\mu_{\text{Haar}} := 4 \sin^2\left(\frac{\theta}{2}\right) d\theta d\Omega(\mathbf{n})$$

which yields a total volume of SO(3) of $8\pi^2$. For our probability density function, we need to normalize this to 1 and thus arrive at the normalized Haar measure:

$$d\mu := \frac{1}{2\pi^2} \sin^2\left(\frac{\theta}{2}\right) d\theta d\Omega(\mathbf{n})$$

To sample uniformly on SO(3) using the Haar measure in axis–angle form, we can consider the axis and angle separately. Obviously, we need to generate the *axis* of the noise rotation uniformly to ensure isotropy of our sampling [16]—otherwise, we would prefer certain directions over others. This corresponds to drawing a random unit vector uniformly on S^2 and can be achieved in several different ways, e.g., by using spherical coordinates with $r = 1$ and $z \sim \mathcal{U}[-1, 1]$ and $\psi \sim \mathcal{U}[0, 2\pi)$, or by drawing three independent normal random variables and normalizing the resulting vector.

To sample a random rotation angle θ uniformly using the Haar measure, we would then have to draw

$$\theta \sim \mathcal{P}_\mu(\theta) \propto \sin\frac{\theta}{2}$$

Finally, to produce a Laplace mechanism, we need to change the uniform distribution to include a factor that decays exponentially with the distance. Since we use an axis–angle parametrization, and since the geodesic distance reduces to the angle $\theta \in [0, \pi]$ between the rotations, we can still draw the axis uniformly and need to add the exponential decay to the sampling of noise angles. To include the privacy parameter ϵ, we additionally change the metric from the geodesic distance $d_{\mathcal{Y}}(\mathbf{R}_1, \mathbf{R}_2) = \theta$ to the **scaled geodesic distance** $\hat{d}_{\mathcal{Y}}(\mathbf{R}_1, \mathbf{R}_2) := \epsilon\theta$. This finally yields the following form for the Laplace mechanism on SO(3):

$$\mathcal{L}(y)(dz) = \lambda(z) \exp[-\epsilon\theta] \mu(dz) = \frac{\lambda(z)}{2\pi^2} \exp[-\epsilon\theta] \sin^2\left(\frac{\theta}{2}\right) d\theta d\Omega(\mathbf{n})$$

where $\lambda(z)$ normalizes the distribution. Note that in full generality, the normalization λ could also depend on the input rotation y, which would not fit the definition of Chatzikolakis et al. However, for compact groups with bi-invariant distances—such as SO(3)—one can show that the normalization does **not** depend on y. In fact, it is straightforward to see that the normalization is also independent of z and resolves to a constant that only depends on ϵ, i.e., $\lambda(z) = \frac{1}{\Gamma(\epsilon)}$ for some $\Gamma(\epsilon) \in \mathbb{R}$. Thus, choosing the base measure μ to be the Haar measure μ, we finally arrive at the following sampling procedure for our Laplace mechanism on SO(3): first, we draw a random axis according to

$$z \sim \mathcal{U}[-1, 1] \tag{17}$$

$$\psi \sim \mathcal{U}[0, 2\pi) \tag{18}$$

$$n := \left(\sqrt{1 - z^2} \cos\psi,\ \sqrt{1 - z^2} \sin\psi,\ z\right). \tag{19}$$

In principle, $\theta|y$ can again be drawn using a simple rejection scheme: we draw a candidate $\theta \sim \mu$ from the Haar measure by sampling $u \sim \mathcal{U}(0,1)$ and setting $\theta := 2\arccos\left(\sqrt{1-u}\right)$ (yielding a θ distributed $\propto \sin^2\left(\frac{\theta}{2}\right)$) and accept the proposal with probability $\exp[-\epsilon\theta]$. In practice, this simple scheme only works for small values of ϵ; with increasing ϵ, samples have to be concentrated more tightly around the original rotation. As the acceptance probability decays exponentially in ϵ, the sampler will reject almost all proposals and become pathologically inefficient, leading either to extreme computational cost until convergence or highly skewed sampling results. We have thus developed a stabilized and much more efficient sampling procedure, which works as follows.

For values of ϵ smaller than a threshold t_ϵ (in our implementation, we use $t_\epsilon = 3.5$), we continue to use the simple rejection scheme described above. For larger values of ϵ, we approximate the target distribution $\sin^2\left(\frac{x}{2}\right)\exp[-\epsilon\theta]$. Since in this region θ will be small, we expand the target into a Taylor series, where we use the identity $\sin^2(\frac{\theta}{2}) = \frac{1}{2}(1-\cos(\theta))$ to find

$$\sin^2\left(\frac{\theta}{2}\right) = \frac{1}{2}(1-\cos(\theta)) = \frac{1}{2}\left(1-\left(1-\frac{\theta^2}{2!}+\mathcal{O}(\theta^4)\right)\right)$$

and thus

$$p(\theta) \propto \theta^2 \exp[-\epsilon\theta]$$

Remembering the definition of the *Gamma distribution* with shape parameter α and rate parameter λ,

$$\mathrm{Gamma}(x;\alpha,\lambda) = \frac{\lambda^\alpha}{\Gamma(\alpha)} x^{\alpha-1}\exp[-\lambda x]$$

we see that for a large ϵ, this converges to $\mathrm{Gamma}(\theta;\alpha=3,\lambda=\epsilon)$, which we use to approximate the Laplace distribution on SO(3) for $\epsilon > t_\epsilon$. The importance of choosing an efficient sampling scheme is demonstrated in Figure 2, which shows distance histograms for 50,000 noise rotations sampled from the efficient and inefficient sampler implementations for $\epsilon = 8$. Clearly, the inefficient sampler fails to converge, leading to a drastic overrepresentation of larger distances.

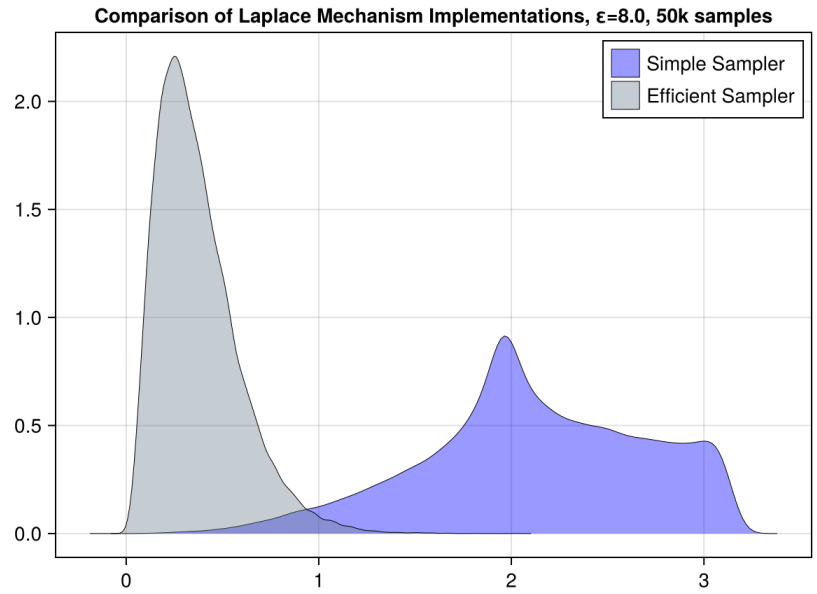

Figure 2. Histograms of the distances for 50,000 noise samples generated by the inefficient and efficient Laplace sampler implementations. The distributions clearly behave entirely differently.

3.4. Comparing the Mechanisms

Despite the popularity of the Laplace mechanism, it is not a priori clear which of the two mechanisms proposed in this work is *better* for any given task. Even the question of what constitutes "better" can be challenging, as discussed in detail in [2]. As the data we want to protect is supposedly used for some downstream analysis, it is natural to think about the *utility* of the released (i.e., noisy) data and to prefer mechanisms that achieve better utility for the same privacy budget. Following [11], we take a utility-focused Bayesian perspective and model the *consumer* of the data by a prior π on the set of secrets, and a loss function ℓ, which is monotone with respect to $d_{\mathbb{R}}$. The *utility* $\mathscr{U}(\mathcal{M}, \pi, \ell)$ of mechanism $\mathcal{M}$ is then given by the expected loss. We denote the metric under consideration by $d_{\mathscr{X}}$. Let f be an arbitrary query on the released data, and let $\mathscr{H}_f(d_{\mathscr{X}})$ denote the set of all mechanisms $\mathcal{M}$ such that $\mathcal{M} \circ f$ fulfills $d_{\mathscr{X}}$-privacy.

Definition 3. *A mechanism $\mathcal{M} \in \mathscr{H}_f(d_{\mathscr{X}})$ is called $f - d_{\mathscr{X}}$-optimal, if and only if $\mathscr{U}(\mathcal{M}, \pi, \ell) \geq \mathscr{U}(\mathcal{M}', \pi, \ell) \quad \forall \mathcal{M}' \in \mathscr{H}_f(d_{\mathscr{X}})$ for all priors π and all loss functions ℓ.*

While there are some theoretical results on the optimality of the Laplace mechanism for selected queries [11] and metrics [2], there is—to the best of our knowledge—no optimality result that covers our use case. In addition, other optimality definitions exist (e.g., risk-averse instead of Bayesian consumers). We thus turn to a simulation-based comparison of both approaches. In order to compare the Bingham and Laplace mechanisms on SO(3), we must first ensure that we evaluate them for an equivalent privacy budget, since both have been defined with respect to different metrics. Let ϵ be the privacy budget of the Laplace mechanism and $\tilde{\epsilon}$ the budget of the Bingham mechanism. In Appendix F, we show that

$$\tilde{\epsilon} = \frac{\pi}{4}\epsilon$$

To compare the suitability of the mechanisms, we set up a simulation study. We draw a random rotation on SO(3) to represent the data (say, the measurement of a sensor registering the orientation of a joint in an ergonomic study). Using both mechanisms, we draw a large number of samples for different privacy budgets and compute histograms for the geodesic distance of each sample to the original rotation. The farther the sampled points are from the original data, the better the privacy, but at the cost of lower utility. For a utility-focused mechanism in a Bayesian perspective, if the samples provided by a mechanism are closer, on average, to the original data point for the same privacy budget, we should prefer this mechanism. To pick a set of suitable privacy budgets for this simulation, we define a *radius of indistinguishability* for the Laplacian mechanism as follows.

Definition 4. *Let $p \in [0, 1]$. We call ρ the **radius of indistinguishability at level** p, if and only if for any rotation q and for any noise sample s we have*

$$P(\phi_6(q, s) \leq \rho) = p$$

where again ϕ_6 denotes the geodesic distance on SO(3).

In other words, p percent of the noise samples will be contained within a radius of ρ (have an angular displacement less than ρ) from q. Remembering that for the Laplace mechanism on SO(3), $p(\theta) \propto \sin^2(\theta/2) \exp[-\epsilon\theta]$, we define the cumulative distribution function

$$F(\rho) := \alpha \int_0^{\rho} \sin^2(\theta/2) \exp[-\epsilon\theta] d\theta$$

with normalization

$$\alpha := \int_0^\pi \sin^2(\theta/2)\exp[-\epsilon\theta]d\theta$$

Finding ρ for a given p then amounts to solving $F(\rho) = p$. While the CDF has no elementary closed form, finding ρ is straightforward using numerical integration and a root-finding method.

Since we convert the privacy budget of the Bingham mechanism to the same scale, we can use the same privacy budget for both mechanisms. Figure 3 shows the results, where the values of the indistinguishability radii ρ were computed with respect to the Laplace mechanism at a level of $p = 0.683$ (chosen to correspond to one standard deviation of a normal distribution). Both mechanisms were implemented in Julia. In our experiments, the Laplace mechanism was about a factor of 2–3 faster to evaluate than the Bingham mechanism for $\epsilon < t_\epsilon$, and up to two orders of magnitude faster for $\epsilon > t_\epsilon$, owing to the simpler sampling procedure and higher acceptance rates.

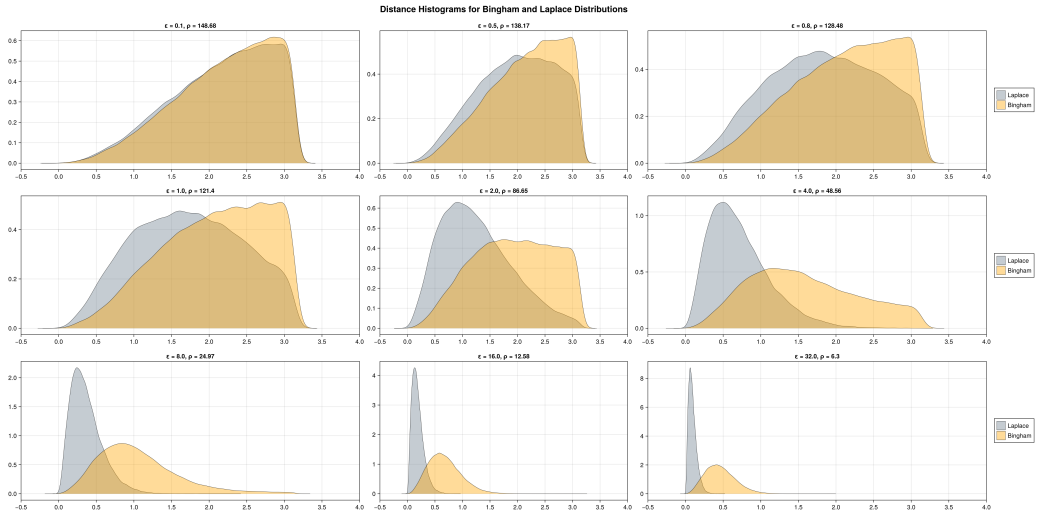

Figure 3. Histograms of the distance distributions of samples generated from the Bingham and Laplace mechanisms for different privacy budgets. Indistinguishability radii were computed with respect to the Laplace mechanism at a level of $p = 0.683$.

The simulation study demonstrates that, for equivalent privacy budgets (after accounting for the difference in metrics), the Laplace mechanism tends to produce noise samples that are more tightly concentrated around the original rotation than those generated by the Bingham mechanism. This greater concentration translates into improved utility in the average case, making the Laplace mechanism an attractive choice for scenarios where downstream tasks benefit from low-distortion data. In particular, the sharper peak of the Laplace distribution around the mean ensures that most samples lie closer to the original rotation, leading to lower expected loss under standard Bayesian utility models with typical monotonic loss functions.

However, the Bingham mechanism, while exhibiting a broader spread in the simulations, has a fundamentally different distributional profile and thus might be advantageous under different utility models or for different types of consumers.

4. Discussion

The privacy of datasets containing three-dimensional rotations, i.e., data sampled from SO(3), has to our knowledge not been previously addressed. In this work, we have derived two privacy mechanisms specifically tailored for this task: a mechanism induced by the

Bingham distribution, and a Laplace mechanism adapted to the geodesic metric on SO(3). We have shown that the Laplace mechanism fulfills $\epsilon\phi_6$-privacy, where, using the notation from [10], ϕ_6 corresponds to the geodesic (intrinsic) metric on SO(3), and that the Bingham mechanism fulfills $\sqrt{2}\epsilon\phi_5'$-privacy. We have further shown that this privacy guarantee corresponds to a $\frac{\pi}{4}\epsilon\phi_6$-privacy guarantee of the Bingham mechanism with respect to the geodesic metric on SO(3), i.e., it fulfills $\tilde{\epsilon}\phi_6$-privacy where $\tilde{\epsilon} = \frac{\pi}{4}\epsilon$, and ϵ is the privacy guarantee for the Laplace mechanism on SO(3).

Both mechanisms have their own rationale: the Bingham mechanism is particularly appealing, as it inherently incorporates the antipodal symmetry on SO(3) and is by design compatible with the space of interest. The Laplace mechanism is a very generally applicable concept that has been used successfully in many other application scenarios, i.e., on other metric spaces. As demonstrated by our numerical experiments (c. f. Figure 3), both mechanisms behave very similarly in a high privacy setting, i.e., for small values of ϵ and large values of ρ. With relaxed privacy, we find that the samples generated by the Laplace mechanism are more concentrated around the original sample—corresponding to the sharper peak of the Laplacian distribution—and hence, should lead to increased utility, albeit at the cost of smaller privacy, but still within the guarantee of the privacy budget.

But both mechanisms are based on different base distributions—the Bingham mechanism arises from a $4D$-Gaussian distribution constrained to reside on S^3, while the Laplace mechanism is based on an exponential distribution. These distributions have decidedly different characteristics—the Gaussian features narrow tails and a smooth, broader peak around the mean, while the Laplace decays significantly more slowly, but has a sharp peak at the mean. These different shapes should in turn lead to differently concentrated noise samples at the same privacy level. Thus, both mechanisms have different privacy–utility tradeoffs. In particular, it might be tempting to assume that the Bingham mechanism could be more appealing in settings where large deviations from the original data are particularly undesirable—for instance, in risk-averse applications or under loss functions that heavily penalize outliers. Our experiments, however, show a different picture. In the low-to-medium-privacy regime, the tails of the Bingham distribution are always heavier than those of the corresponding Laplace mechanism. This surprising fact is explained by the fact that the geometry and topology of SO(3) change the behavior of the distributions in unexpected ways: the tails of the Gaussian become thicker, while the tails of the exponential distribution become slimmer when mapping the distributions to SO(3). Furthermore, the proof for the privacy of the Bingham mechanism was tailored to a different metric and then converted to the geodesic distance. The proof itself also uses a number of inequality constraints that might not be tight. It is likely that the Bingham mechanism actually fulfills even higher privacy guarantees than the ones proved in our manuscript. We thus cannot yet prove a general superiority of the Laplace mechanism over the Bingham one for all values of ϵ. At this point, however, for utility-focused applications, we can unconditionally recommend the Laplace mechanism over the Bingham one, as it has better utility for mid-to-low-privacy scenarios and has a much more efficient implementation. However, as Figure 2 clearly shows, it is crucial to use a stable Laplace sampler implementation with high sampling efficiency.

To allow even more choices from the design space of privacy mechanisms on SO(3), future work might investigate whether the Purkayastha or VMF mechanisms from [8] can be adopted to the group of three-dimensional rotations, and how they compare with the mechanisms described in this work.

Author Contributions: Conceptualization, A.K.H. and A.H.; methodology, A.K.H., A.H. and E.S.; validation, A.H. and E.S.; formal analysis, A.H. and E.S.; writing—original draft preparation, A.K.H.; writing—review and editing, project administration, A.H. and E.S.; funding acquisition, A.H. All authors have read and agreed to the published version of the manuscript.

Funding: Anna Hildebrandt acknowledges support of German BMBF through the grants for the Ergobest Project (16SV8671) and the Primflow Project, Andreas Hildebrandt the support of the Carl-Zeiss project TOPML (P2021-02-014).

Data Availability Statement: No new data were created or analyzed in this study. Data sharing is not applicable to this article.

Conflicts of Interest: Author A.K.H. was employed by the companies PRAIVACY UI and Mondata GmbH, author A.H. was employed by the company Mondata GmbH. Author E.S. declares that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Appendix A. P1: Third Binomial for Symmetric Metrics

$$\begin{aligned}\forall \mathbf{A} = \mathbf{A}^T : (x-y)^T\mathbf{A}(x+y) &= x^T\mathbf{A}x - y^T\mathbf{A}y + x^T\mathbf{A}y - y^T\mathbf{A}x \\ &= x^T\mathbf{A}x - y^T\mathbf{A}y + x^T\mathbf{A}y - x^T\mathbf{A}^T y \\ &= x^T\mathbf{A}x - y^T\mathbf{A}y \qquad [\mathbf{A} = \mathbf{A}^T]\end{aligned}$$

Appendix B. P2: Invariance of Quaternion Scalar Product to Rotations

Theorem A1. *Let* $q, p, r \in \mathbb{H}, ||q|| = ||p|| = ||r|| = 1$.

Proof of Theorem. Denote the Hamilton product of quaternions p and q by $p \odot q$ and their scalar product by $p^T \cdot q$. Then $(p \odot r)^T \cdot (q \odot r) = p^T \cdot q$, i.e., the quaternion scalar product is invariant with respect to rotations. By definition of the Hamilton product,

$$r \odot p = \tilde{\mathbf{R}} \cdot p$$

with $\tilde{\mathbf{R}}^T\tilde{\mathbf{R}} = \mathbb{1}_{4\times 4}$. It follows that

$$\begin{aligned}(p \odot r)^T \cdot (q \odot r) &= \left(\tilde{\mathbf{R}} \cdot p\right)^T \cdot \left(\tilde{\mathbf{R}} \cdot q\right) \\ &= p^T\tilde{\mathbf{R}}^T\tilde{\mathbf{R}}q \\ &= p^T \cdot q\end{aligned}$$

□

Appendix C. P3: Implication of Third Binomial for Norms

Theorem A2. *Let* $a, b \in \mathbb{R}^n$ *with* $||a|| = ||b|| = 1$. *Then,*

$$||a-b||^2 \cdot ||a+b||^2 = 4\left(1 - \left(a^T \cdot b\right)^2\right)$$

Proof of Theorem.

$$\begin{aligned}||a-b||^2 \cdot ||a+b||^2 &= \left(||a||^2 + ||b||^2 - 2a^T \cdot b\right)\left(||a||^2 + ||b||^2 + 2a^T \cdot b\right) \\ &= \left(2 - 2a^T \cdot b\right)\left(2 + 2a^T \cdot b\right) \\ &= 4\left(1 - \left(a^T \cdot b\right)^2\right)\end{aligned}$$

□

Appendix D. The Rejection Sampling Scheme of Kent et al.

In our work, we need to sample noise from a Bingham distribution. In [14], Kent et al. give a simple acceptance–rejection scheme that is efficient and simple to implement. In general, such acceptance–rejection methods work as follows: we intend to sample from a density $f(x)$—in our case, the Bingham distribution—which can be written as

$$f(x) = c_f f^*(x)$$

with known functional form $f^*(x)$ but with a normalization constant c_f that is challenging or even infeasible to compute. Now assume further that there is a second density

$$g(x) = c_g g^*(x)$$

for which we already have a sampling method. If we can find a bound of the type

$$f^*(x) \leq M^* g^*(x) \quad \forall x$$

for some constant M^*, then we can use samples of $f(x)$ to simulate samples from $g(x)$ through the following acceptance–rejection Algorithm A1:

Algorithm A1 General acceptance–rejection scheme to sample $X \sim g$

while true **do**
  Sample $X \sim g$ and, independently, $W \sim \text{Unif}(0,1)$
  if $W \leq \frac{f^*(x)}{M^* g^*(X)}$ **then**
    return X ▷ *Accept proposal*

For the Bingham distribution, $f(\mathbf{x}) = c_f \exp\left(-\mathbf{x}^T \mathbf{A}\mathbf{x}\right)$, where the minus sign—which can be absorbed into $\mathbf{A}$—is unconventional but simplifies further notation. Following Kent et al., we use the angular central Gaussian distribution ACG $(\mathbf{\Omega})$ as the second density $g(x)$, with $g(\mathbf{x}) = c_g\left(\mathbf{x}^t \mathbf{\Omega}\mathbf{x}\right)^{-q/2}$, where $\Omega \in \mathbb{R}^{q\times q}$. In our case, we are interested in the Bingham distribution on SO(3), and hence, $q = 4$. Drawing samples from ACG $(\mathbf{\Omega})$ is simple: for a sample $y \sim \mathcal{N}\left(0, \mathbf{\Omega}^{-1}\right)$, we have $\frac{y}{\|y\|} \sim$ ACG $(\mathbf{\Omega})$.

Kent et al. show that, for $b > 0$ and $\mathbf{\Omega}(b) = \mathbf{I} + 2\mathbf{A}/b$, we find

$$M^*(b) = \exp\left(-(q-b)/2\right)(q/b)^{q/2}$$

and that the bound can be optimized by the solution b^* to

$$\sum_{i=1}^{q} \frac{1}{b + 2\lambda_i} = 1$$

where λ_i are the eigenvalues of $\boldsymbol{A}$.

Appendix E. Table of Definitions Related to SO(3)

Term	Definition
S^2	The unit sphere in 3D.
SO(3)	The special orthogonal group of rotations in three dimensions.
3D rotation matrix	Representation of a member of SO(3) as a matrix $\mathbf{R} \in \mathbb{R}^{3\times 3}$ with $\mathbf{R}\mathbf{R}^T = \mathbb{1}$ and $\det(\mathbf{R}) = 1$.

Term	Definition
H	The space of quaternions—a generalization of complex numbers with one real and three imaginary axes
$p \odot q$ with $p, q \in \mathbb{H}$	The Hamilton product of two quaternions.
Versor	A quaternion q with unit norm $\|q\| = 1$. This forms an alternative representation of members of SO(3).
SU(2)	The special unitary group of complex unitary 2×2 matrices with $\mathbf{U}^\dagger\mathbf{U} = \mathbf{U}\mathbf{U}^\dagger = \mathbb{1}$ and $\det(\mathbf{U}) = 1$, where † denotes Hermitian conjugation. SU(2) is isomorphic to the space of versors, but provides a double cover of SO(3), i.e., every rotation in SO(3) corresponds to exactly two elements of SU(2), and hence, exactly two versors.

Appendix F. Comparing Privacy Budgets Across Mechanisms

Let $\mathcal{M}_\mathcal{B}$ be the Bingham mechanism and $\mathcal{M}_\mathcal{L}$ the Laplace mechanism on SO(3). We have shown that $\mathcal{M}_\mathcal{B}$ fulfills $\sqrt{2}\epsilon_\mathcal{B}$-privacy according to the metric Φ, and that $\mathcal{M}_\mathcal{L}$ fulfills ϵ-privacy according to the geodesic metric $\phi_6 = \theta$. From [10], we can also establish a relationship between the norms $\Phi(q_1, q_2)$ and $\phi_6(q_1, q_2)$ as follows:

$$\Phi(q_1, q_2) = 2\sqrt{2}\sin\left(\frac{\phi_6(q_1, q_2)}{2}\right) = 2\sqrt{2}\sin\left(\frac{\theta}{2}\right)$$

and thus $\theta = 2\arcsin\left(\frac{\Phi(q_1,q_2)}{2\sqrt{2}}\right)$. Our differential privacy guarantee for the Bingham mechanism amounts to

$$\frac{\mathcal{M}_\mathcal{B}(q_1)[r]}{\mathcal{M}_\mathcal{B}(q_2)[r]} \leq \exp\left[\sqrt{2}\epsilon_\mathcal{B} \cdot \Phi(q_1, q_2)\right] \quad \forall q_1, q_2 \tag{A1}$$

Our goal is to map this to a privacy guarantee according to the geodesic norm ϕ_6, i.e., we want to find an $\epsilon'_\mathcal{B}$ such that

$$\frac{\mathcal{M}_\mathcal{B}(q_1)[r]}{\mathcal{M}_\mathcal{B}(q_2)[r]} \leq \exp[\epsilon\phi_6(q_1, q_2)] \leq \exp\left[\sqrt{2}\epsilon'_\mathcal{B} \cdot \Phi(q_1, q_2)\right] \quad \forall q_1, q_2 \tag{A2}$$

i.e., with privacy protections at least as strong as those of the original guarantee. Since the exponential is monotonic, we can equate the exponents:

$$\epsilon\phi_6(q_1, q_2) \leq \sqrt{2}\epsilon'_\mathcal{B} \cdot \Phi(q_1, q_2) \quad \forall q_1, q_2 \tag{A3}$$

If $q_1 = q_2$, this holds trivially. If $q_1 \neq q_2$, $\Phi(q_1, q_2) > 0$ and we find for the new bound $\epsilon'_\mathcal{B}$:

$$\epsilon'_\mathcal{B} \geq \epsilon \cdot \frac{\phi_6(q_1, q_2)}{\sqrt{2}\Phi(q_1, q_2)} \quad \forall q_1, q_2$$

Since this inequality has to hold for **all** pairs q_1, q_2, the new budget $\epsilon'_\mathcal{B}$ must be at least as large as the maximum value of this factor. In other words, we choose $\epsilon'_\mathcal{B}$ as

$$\epsilon'_\mathcal{B} = \epsilon \cdot \sup_{q_1 \neq q_2} \frac{\phi_6(q_1, q_2)}{\sqrt{2}\Phi(q_1, q_2)}$$

Denoting the angle between rotations q_1 and q_2 by θ, we have $\phi_6(q_1, q_2) = \theta$ and $\Phi(q_1, q_2) = 2\sqrt{2}\sin(\theta/2)$ and find

$$\epsilon'_{\mathcal{B}} = \epsilon \cdot \sup_{q_1 \neq q_2} \frac{\theta}{4\sin(\theta/2)} = \epsilon \cdot \sup_{q_1 \neq q_2} \frac{\theta/2}{2\sin(\theta/2)}$$

For $\theta \in (0, \pi]$, we have $0 < \sin(\theta/2) < \theta/2$ and hence $\frac{\theta/2}{2\sin(\theta/2)} > \frac{1}{2} \quad \forall q_1, q_2$. And since $\theta/2$ grows faster than $\sin(\theta/2)$ on $\theta > 0$, we can conclude that the supremum occurs at the *maximum possible geodesic distance* allowed by the space. For the SO(3), this maximum geodesic distance is $\theta = \pi$ and we finally find

$$\epsilon'_{\mathcal{B}} = \epsilon \cdot \sup_{q_1 \neq q_2} \frac{\phi_6(q_1, q_2)}{\sqrt{2}\Phi(q_1, q_2)} = \epsilon \cdot \frac{\pi}{4\sin(\pi/2)}$$

which, with $\sin(\pi/2) = 1$, resolves to

$$\epsilon'_{\mathcal{B}} = \frac{\pi}{4} \cdot \epsilon$$

References

1. Dwork, C.; Roth, A. The algorithmic foundations of differential privacy. *Found. Trends Theor. Comput. Sci.* **2014**, *9*, 211–407. [CrossRef]
2. Fernandes, N. Differential Privacy for Metric Spaces: Information-Theoretic Models for Privacy and Utility with New Applications to Metric Domains. Ph.D. Thesis, École Polytechnique de Paris, Paris, France, and Macquarie University, Sydney, Australia, 2021.
3. Fan, L. Image pixelization with differential privacy. In Proceedings of the Data and Applications Security and Privacy XXXII: 32nd Annual IFIP WG 11.3 Conference, DBSec 2018, Bergamo, Italy, 16–18 July 2018; Proceedings 32; Springer: Berlin/Heidelberg, Germany, 2018; pp. 148–162.
4. Chamikara, M.A.P.; Bertok, P.; Khalil, I.; Liu, D.; Camtepe, S. Privacy preserving face recognition utilizing differential privacy. *Comput. Secur.* **2020**, *97*, 101951. [CrossRef]
5. Weggenmann, B.; Kerschbaum, F. Syntf: Synthetic and differentially private term frequency vectors for privacy-preserving text mining. In Proceedings of the 41st International ACM SIGIR Conference on Research & Development in Information Retrieval, Ann Arbor, MI, USA, 8–12 July 2018; pp. 305–314.
6. Fernandes, N.; Dras, M.; McIver, A. Generalised differential privacy for text document processing. In Proceedings of the Principles of Security and Trust: 8th International Conference, POST 2019, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2019, Prague, Czech Republic, 6–11 April 2019; Proceedings 8; Springer International Publishing: Berlin/Heidelberg, Germany, 2019; pp. 123–148.
7. Primault, V.; Boutet, A.; Mokhtar, S.B.; Brunie, L. The Long Road to Computational Location Privacy: A Survey. *IEEE Commun. Surv. Tutor.* **2019**, *21*, 2772–2793. [CrossRef]
8. Weggenmann, B.; Kerschbaum, F. Differential privacy for directional data. In Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security, Virtual Event, Republic of Korea, 15–19 November 2021; pp. 1205–1222.
9. Imola, J.; Kasiviswanathan, S.; White, S.; Aggarwal, A.; Teissier, N. Balancing Utility and Scalability in Metric Differential Privacy. In Proceedings of the Thirty-Eighth Conference on Uncertainty in Artificial Intelligence, PMLR, Eindhoven, The Netherlands, 1–5 August 2022; pp. 885–894.
10. Huynh, D.Q. Metrics for 3D Rotations: Comparison and Analysis. *J. Math. Imaging Vis.* **2009**, *35*, 155–164. [CrossRef]
11. Chatzikokolakis, K.; Andrés, M.E.; Bordenabe, N.E.; Palamidessi, C. Broadening the Scope of Differential Privacy Using Metrics In Proceedings of the Privacy Enhancing Technologies Symposium (PETS'13), Vigo, Spain, 11–13 July 2012; LNCS 7981, pp. 82–102. [CrossRef]
12. Cornwell, J.F. *Group Theory in Physics: An Introduction, Volume 1*; Academic Press: London, UK, 1984; ISBN 978-0121898007.
13. Gilitschenski, I.; Kurz, G.; Julier, S.J.; Hanebeck, U.D. Unscented Orientation Estimation Based on the Bingham Distribution. *IEEE Trans. Autom. Control* **2016**, *61*, 172–177. [CrossRef]
14. Kent, J.T.; Ganeiber, A.M.; Mardia, K.V. A New Method to Simulate the Bingham and Related Distributions in Directional Data Analysis with Applications. *arXiv* **2013**, arXiv:1310.8110. [CrossRef]

15. Haar, A. Der Massbegriff in Der Theorie Der Kontinuierlichen Gruppen. *Ann. Math.* **1933**, *34*, 147–169. [CrossRef]
16. Miles, R.E. On Random Rotations in $\mathbb{R}^3$. *Biometrika* **1965**, *52*, 636–639. [CrossRef]