Approximate Distance-Comparison-Preserving Symmetric Encryption

Georg Fuchsbauer[1] Riddhi Ghosal[2] Nathan Hauke[3] Adam O'Neill[4]

Abstract

We introduce *distance-comparison-preserving* symmetric encryption (DCPE), a new type of property-preserving encryption (PPE) that preserves relative distance between plaintext vectors. DCPE is naturally suited for nearest-neighbor search on encrypted data. To achieve meaningful security, we divert from prior work on PPE and ask for *approximate correctness*, which is natural given the prevalence of approximate nearest neighbor (ANN) search. We conduct a thorough study of what security approximate DCPE can provide and how to construct it.

Based on a relation we prove between approximate DCP and *approximate distance-preserving* functions, we design our core approximate DCPE scheme we call *Scale-And-Perturb* (SAP). The encryption algorithm of SAP processes data on-the-fly. To boost security, we also introduce two preprocessing techniques: (1) *normalizing* the plaintext distribution, and (2) *shuffling*, wherein the component-wise encrypted dataset is randomly permuted. We prove (under suitable restrictions) that SAP achieves an indistinguishability-based security notion we call *Real-or-Replaced* (RoR). In particular, our RoR result implies that our scheme prevents *membership inference attacks* by Yeom *et al.* (CSF 2018). Moreover, we show for i.i.d. multivariate normal plaintexts, we get security against *approximate frequency-finding attacks*, the main line of attacks against property-preserving encryption. This follows from a *one-wayness* (OW) analysis. Finally, carefully combining our OW and RoR results, we are able characterize bit-security of SAP.

Our overall findings are that our scheme not only has superior bit-security to OPE but resists specific attacks that even ideal order-revealing encryption (Boneh *et al.*, EUROCRYPT 2015) does not. This suggests it could be sufficient for certain ANN applications, a subject on which we encourage further study.

1 Introduction

We review the problem we address in this paper and then overview our results.

1.1 Background and Motivation

The paradigm of *secure outsourced databases* refers to a setting where a client transmits its database to an untrusted server that hosts it. The goal of such a protocol is to protect information about the database from the server to the extent possible, while maintaining the ability for the client to issue queries. This paradigm was first introduced in the database community by [HILM02] and has received intensive study since then (see [Kam15]) from the database and cryptographic community.

An attractive approach to constructions, which has already seen real-world deployment, is the emerging notion of *function-revealing encryption* (FRE) (*e.g.*, [BBO07, BCLO09, PR12, BLR$^+$15, JP18]). A (private-key) function-revealing encryption scheme for a function f allows anyone from the encryptions $c_1, \ldots, c_n$ of $m_1, \ldots, m_n$ respectively to compute $f(m_1, \ldots, m_n)$. The terminology is important here: *function* may be replaced by *property*, which refers to the case that f is a predicate;[1] *revealing* can be replaced by *preserving*, where the way of computing the above output of f is itself $f(c_1, \ldots, c_n)$. FRE is attractive because it allows the construction of outsourced database protocols that let the server index and process queries almost exactly the same way as for unencrypted data (in fact, exactly in the function-preserving case).

So far, FRE-based protocols have mainly been built for running queries on encrypted SQL databases. The types of FRE used here are *order-revealing* (ORE) and *order-preserving* encryption (OPE) [AKSX04, BCLO09, BCO11, BLR$^+$15], in which plaintexts are numbers and f is the comparison predicate $p_{\mathsf{comp}}(x, y) = 1$ iff $x < y$;

[1]TU Wien, Austria. georg.fuchsbauer@tuwien.ac.at
[2]UCLA. riddhi@cs.ucla.edu
[3]Georgetown University, USA. nah52@georgetown.edu
[4]University of Massachusetts Amherst, USA. adamo@cs.umass.edu

[1] Property-preserving encryption (PPE) is a special case that our construction actually falls into, but we stick with FRE terminology for generality.

and *deterministic encryption*[2] (DE) [ABO07, BBO07], which is the function-preserving case where f is the equality predicate $p_{\mathsf{eq}}(x, y) = 1$ iff $x = y$. Combined with some other schemes and tricks, DE and ORE/OPE give rise to outsourced database protocols for most SQL queries, such as the CryptDB system [PLZ13].

Unfortunately, even when these underlying FRE schemes are ideal, this approach is subject to attacks in the outsourced database setting ([NKW15, DDC16, PW16, GSB+17, BGC+18] in the "snapshot" setting, and [KKNO16, LMP18, GLMP18, GLMP19, KPT20] in the "persistent" setting — see below). This has created a viewpoint in the community that community the FRE approach in such a higher-level applications is inherently insecure. Indeed, existing positive results for FRE-based protocols have major restrictions: they either assume uniform, high-entropy plaintexts [BCO11] or an unknown prior on the data [CLO+18], neither of which seems likely to hold in practice. Other work introduces high overhead for practical datasets [LP18, PGW19].

In this work, we make further progress on studying FRE. We put forth a new type of FRE and a construction, achieving novel security guarantees moving past previous limitations. We aim for guarantees in a "snapshot" attack model where the adversary sees one snapshot of the encrypted database. This is opposed to a "persistent" attack model where the adversary observes the query processing over time.[3]

1.2 Our Results

(Approximate) Distance-Comparison Preserving Encryption. The first step is to identify the "core operation" that nearest neighbor (NN) search algorithms use. For standard NN algorithms [WL83, CD07], this is *distance comparison*. This gives rise to *distance-comparison preserving* and *distance-comparison revealing* encryption (DCPE/DCRE) for the ternary predicate $p_{\mathsf{dist\text{-}comp}}(x, y, z) = 1$ if $\mathsf{dist}(x, y) < \mathsf{dist}(x, z)$.[4] Even on plaintext data, in some applications it is common to return *approximate* nearest neighbors (ANN) instead of the exact values [AMN+98, ML14]. Indeed, ANN is very useful for high-dimensional data due to "curse of dimensionality" [BGRS99, IM98]. (In higher dimensions, the notion of exact distances between points becomes less significant, thereby making it difficult for exact algorithms to converge.) We leverage this feature as a way to move past security limitations of prior work on PPE. We thus consider "approximate" DCPE/DCRE, where the predicate evaluates to 1 if $\mathsf{dist}(x, y)$ and $\mathsf{dist}(x, z)$ are sufficiently far. Approximate DCPE, which we show preserves the approximation factor of the overlying NN search algorithm, is the central notion of our work.

Relation to Approximate Distance-Preserving Encryption. In order to get a handle on a DCPE construction, we would like to understand the "structure" of DCP functions — what do they look like? We prove that approximate DCP is approximate *distance-preserving* (DP) with a related approximation factor. To make this precise, let us call an encryption function β-DCP if

$$\mathsf{dist}(x, y) < \mathsf{dist}(x, z) - \beta \implies p_{\mathsf{dist\text{-}comp}}(c_x, c_y, c_z) = 1 \ ,$$

where c_x denotes a ciphertext of x. Let us call a function (α, β)-DP if

$$\alpha\, \mathsf{dist}(x, y) - \beta \leq \mathsf{dist}(c_x, c_y) \leq \alpha\, \mathsf{dist}(x, y) + \beta.$$

Our main result here is that all functions which are β-DCP satisfy the notion of (α, β')-DP for some β' which depends on α and β. The exact parameters are given in Theorem 15.

The Scale-and-Perturb (SAP*) Construction* Inspired by the above results, our "core" approximate DCPE scheme is called the "Scale-and-Perturb" (SAP) scheme, which works as follows. The encryption algorithm scales the plaintext by a factor held in the secret key followed by adding a random perturbation factor. Note that this does not allow decryption, hence we apply a pseudorandom function to derandomize the perturbation step. The parameters are drawn from a uniform distribution, which is inspired by previous works [GC18, KDWS05, JPW06]. Namely, we sample the perturbation from a uniform distribution within a sphere. The radius of the sphere determines the maximum permissible approximation for distance comparison. The choice of uniform distribution for noise is crucial to prevent averaging out the noise by an adversary, thereby disallowing trivial *known plaintext attacks*.

[2] In our terminology, it could also be called equality-preserving encryption

[3] We note that our analyses do extend to encompass queries as long as they have the same distribution as the data. In future work, this might be ensured via preprocessing as in [MCO+15].

[4] The reason why we specifically consider DCPE/DCRE is, like OPE's claim to fame *they allow the backend of the server to remain essentially unchanged.* For example, suppose we used encryption that allowed comparisons only between query vectors and server vectors This would not allow the client to upload encrypted vectors to the server on-the-fly such that the server can index them. Thus search would require a linear scan.

Preprocessing Techniques. While our core encryption scheme can encrypt any individual message given only the secret key, we propose preprocessing steps which enhance the security guarantees and utility of our scheme. The preprocessing steps assume additional knowledge when encrypting. The first idea is *shuffling*. Here, the entire dataset is encrypted component-wise at once, after the plaintexts are shuffled according to a random permutation. This can be achieved efficiently as shown [BEM+17, KLDF16] in differential privacy. To our knowledge, this technique is new in property-preserving encryption (PPE). The second idea is *normalization*, which converts the plaintexts to a normal distribution. This can be done by applying standard statistical tools like the BoxCox transform [Sak92]. In our results we sometimes assume that such preprocessing techniques have been applied; in particular, our results based on one-wayness apply to plaintexts (*i.e.*, vectors) following a multivariate normal distribution. Many natural statistics already follow a normal distribution as well. Indeed, this is a much better assumption than uniform data as in [BCO11].

Indistinguishability-Based Security and "ideal DCPE" A natural type of security notions to consider here is "indistinguishability-based." Indeed, as DCPE is a special case of property-preserving encryption (PPE), the Left-or-Right (LoR) security notion of [PR12, CD15] applies. Roughly, LoR considers an adversary making queries $(x_0^1, x_1^1), \ldots, (x_0^q, x_1^q)$ to an oracle returning encryptions of either the left messages or the right messages. For the given predicate p, it is further required that $(x_b^1, \ldots, x_b^q)$ have the same "p-equality pattern" for $b \in \{0, 1\}$. This restriction ensures that the functionality of the scheme does not allow the adversary to trivially win the game and that *only* p leaks. We show that it follows from structural results discussed above that this notion is not achievable for approximate $p_{\mathsf{dist\text{-}comp}}$ with practical approximation factors.

Faced with this impossibility result, we move towards using the above mentioned pre-processing techniques along with other security notions to analyze guarantees for our scheme. We conduct a very involved analysis showing our scheme meets several important notions.

First Target: Security against Membership Inference. Our first target is security against *membership inference attacks* [YGFJ18] which determine whether an individual is in the database or not.

Our "real-or-replaced" (RoR) exactly captures these attacks as it deals with indistinguishability of datasets that differ at exactly one point aka. *neighboring* datasets. As mentioned, a crucial technique here is that of a *shuffle* that outputs a random permutation of the ciphertexts. Rather than encrypting a point at a time, RoR captures the scenario where the entire dataset is encrypted at once post preprocessing. *Shuffling* has become a state-of-the-art technique for security amplification in this setting for differential privacy and we observe similar results in our case as well. We note that one could imagine a more practical scenario where instead of uploading the database all at once, uploads are "batched" and a shuffle is applied only to the current batch. We leave an analysis for future work.

In RoR, more specifically the adversary chooses a dataset and sends it to an oracle which does either of the two things with equal probability- (1) Creates a random permutation of the dataset, encrypts it, (2) Chooses one point in the dataset at random, replaces it with another point chosen uniformly within a bounded distance (parameter of the security notion) from the to-be-replaced point, generates a random permutation of the modified dataset and encrypts it. The encrypted points are then returned to the adversary whose goal is to identify which dataset was encrypted.

Second Target: Frequency-Finding Security. Here, we target security against *frequency-finding* (FF) attacks, where the adversary tries to estimate how many times some element appears in the database. In our formulation, the adversary need not even know which element it is and tries to guess an *approximate* frequency. Leakage of frequency information about the plaintext has proved to be the Achilles' heel of previous works in property-preserving encryption schemes (for instance, OPE/ORE and even ideal ORE) [NKW15, GSB+17]. Such attacks have successfully reconstructed (partially or completely) the messages encrypted under the aforementioned schemes by exploiting this leakage. Thus, it is imperative we ensure that similar attacks to do not apply, meaning significantly different techniques would be needed for attacks.

In Theorem 12 we prove that indeed *such a leakage does not occur for* SAP. This is proven by a reduction from a new security notion in the spirit of *one-wayness*. In particularly, we generalize the window one-wayness (WOW) notion of [BCO11] to higher dimensions, calling it *attribute* window one-wayness AWOW and prove it holds for SAP relative to a *message* oracle, which takes as input a distribution chosen by the adversary and outputs the encryption of a randomly sampled message from the said distribution. The reduction to FF exploits the fact that a frequency table (histogram) can be used to construct the Empirical Cumulative Distribution

Function (ECDF), which in turn provides vital information about the high-density points in the support of the underlying distribution. As a conclusion, the FF results say that it is impossible to guess approximate frequencies for any attribute occuring with significantly high probability from a message space sampled from a multivariate gaussian distribution.

In order to help reader have a clearer understanding of the concrete security values that we achieve and show the dependence of various parameters on the security bounds, we have put down some values in Tables 1 and 2 for reference.

Third Target: Bit Security. Finally, we aim to characterize *bit-security* of the plaintexts. This effort, inspired by [BCO11, TYM14], is motivated by the fact that while DCPE inherently cannot hide all partial information, it may still protect some "physical bits" of the plaintexts that represent important partial information in practice. To explain bit-security more precisely, let x be a plaintext in a dataset sampled from some distribution (in our results it is multivariate normal). Fix some stretch of bit-positions of x. We call the stretch OW if given the encryption of the dataset it is hard to compute that stretch of bits of x. We call the stretch pseudorandom if given the encrypted of the dataset one cannot distinguish that stretch of bits from random. Intuitively, this means all partial information about these bits is hidden.

To characterize bit security, we go through several steps. First, we introduce an experiment called *Hardcore Bits* (HCB) which enables us to talk about one-wayness and pseudorandomness of lower order bits for the same message. HCB creates a hybrid to compose the RoR result atop OW. We prove the lower $\log \delta$ bits are pseudorandom, where δ is the distance parameter for RoR. Further, *at least* half the lower half of the bits are one way. Concretely the number of one way bits of a n bit string is $\frac{n}{2} + k$, where k is directly proportional to the approximation factor β. The latter beats the result for OPE [BCO11]; the OPE scheme also does not have hardcore bits.

1.3 Discussion

Setting the Parameters. Overall, the choice of β is pivotal in balancing security and utility. Hence, a natural question to ask is what value of β should be chosen for some application? Unfortunately, there is no single answer, as the parameters needed would vary based on the size of the domain and tolerable error in the application. For instance, if we have a dataset where the message-space for each component is $(-N, N)$, and the ANN can tolerate an error up to $E_{\max} \leq N$, then $\beta \leq E_{\max}$ will ensure that the error is within the specified limit. (Refer to Section 3.2 for details.) From a security perspective, Tables 1 and 2 suggest taking $\beta \geq \sqrt{N}$. Hence, we are looking at the range $\sqrt{N} \leq \beta \leq E_{\max}$. Note, that this is a contradiction if $E_{\max} \leq \sqrt{N}$. This suggests our scheme should not be used in such a case, *i.e.* maximum tolerable error must be at least square-root the domain size.

An Example Application. Based on our results, one example applications where our scheme may be suitable is where the dataset consists of information of people from various demographics. It might be acceptable to differentiate people based on zip code but the identity of a person must be indistinguishable from others living in the same zip code. Our RoR results guarantee strong security while ensuring utility in such setting. More generally, as in the OPE work we do not make any strong claims about when our scheme would be useful. It must be further studied and extended to better understand its security. However, we stress that, promisingly, our results already show that it meets several notions OPE does not.

Fruitful Future Directions: A step forward in this direction would be generalising SAP schemes to make them compatible with other metric spaces. Distance Comparison *Revealing* Encryption (DCRE) is also a compelling subject of study. Recall, a secret key encryption scheme is Distance Comparison Revealing if the scheme takes as input a set of three plaintexts and outputs the pairwise distance comparisons between them. This problem is compelling because it can have security improvements over DCPE, similar to the effect of Order Revealing Encryption(ORE) [CLWW16, BLR$^+$14] over Order Preserving Encrypton(OPE) [BCLO09, BCO11].

Another forward direction is to combine *sketching* algorithms [TYUC17, KKMM12] with our scheme. A sketch of some data set with respect to a function f is a compression (eg. dimension reduction) which allows users to (approximately) compute f by having access to the sketch alone. Such compression approaches are intuitively expected to be effective techniques to significantly improve bit-security.

1.4 Further Related Work

The only previous FRE scheme that works on higher-dimensional vectors is the scheme of [HJL+17] for partial ordering. "Left-or-right" ORE [LW16] is designed to immunize ORE-based protocols against snapshot attacks but suffers the drawback that the support of efficient higher-level protocols afforded by plain ORE is no longer present. Non-FRE based protocols such as Arx [PBP19] or essentially any structured encryption scheme [CK10] are semantically secure in the snapshot attack model, but also do not carry the practical benefits of using FRE, namely avoiding implementing an entirely different backend. Moreover, for these protocols that are secure against snapshot attacks there is evidence that they are not secure given some information about the queries [GRS17]. Non-FRE based protocols for secure nearest-neighbor search have been widely addressed by the database community, although they are rather *ad hoc* or insecure (see [YLX13] and references therein).

2 Preliminaries

2.1 Notation

The following are a list of frequently used notations in the paper which shall remain consistent unless otherwise stated.

- We refer to members of $\{0,1\}^*$ as strings. If x, y are strings then $x\|y$ denotes their concatenation.
- N (N^+) and R (R^+) represent the (non-negative) natural and real numbers respectively.
- For $n \in \mathsf{N}$, $[n]$ is the list $\{1, 2, \ldots, n\}$ and $[a, b]$, $\{a, b\} \in \mathsf{N}$ represents a closed interval consisting of all elements between a and b (inclusive).
- All sets are denoted by sans-serif capital letters (S, C, etc), with the exception of D which is reserved for datasets. Sometimes subscripts like D_1 might be used to uniquely identify multiple datasets. All sets are usually unordered unless specified. Intervals are written as $\mathsf{I}_{[a,b]}$.
- For a finite set S, we denote by $s \xleftarrow{\$} \mathsf{S}$ the process of sampling s uniformly from S.
- Distributions are represented in calligraphic font, like $\mathcal{D}$. Specifically, $\mathcal{U}, \mathcal{N}, \mathcal{MVN}$ typically denote the *Uniform, Normal and Multivariate Normal* distributions respectively.
- $\|\cdot\|$ denotes the length of a vector or size of a set, interval, dataset, *etc.*
- Vectors are denoted in boldface ($\mathbf{v}$, $\mathbf{M}$). A d-dimensional vector is denoted as $\mathbf{v} = (v[1] \ldots, v[d])$ and $v[i]$ refers to the i^{th} component of $\mathbf{v}$.
- Generic domains and range are written as $\mathcal{X}$ and $\mathcal{Y}$, respectively.
- d is the dimension of vector. The dot product between two vectors $\mathbf{u}$ and $\mathbf{v}$ is given by $\mathbf{u} \cdot \mathbf{v}$.
- Calligraphic $\mathcal{M}$ (also called message space) and $\mathcal{C}$ (possibly followed by subscript) represent the plaintext and ciphertext space respectively.
- Mostly plaintext space ($\mathcal{M}$) and ciphertext space ($\mathcal{C}$) are of the form $[-M, M]^d$ and $[-C, C]^d$ respectively where $M, C \in \mathsf{N}$, i.e. a d dimensional space with each dimension being $[-M, M]$. We shall not explicitly mention the final fact in the entire paper.

2.2 Mathematical Background

We recall some mathematical preliminaries that we will use.

Open Ball: The open ball of radius $r > 0$ centered at a point $p \in \mathsf{R}$, denoted by $B(p; r)$, is defined as

$$B(p; r) = \{x \in \mathsf{R} \mid \mathsf{dist}(x, p) < r\},$$

where $\mathsf{dist}(x, y) := \|x - y\|$.

Orthogonal Matrix: A square matrix $\mathbf{Q}$ is said to be an orthogonal matrix if $\mathbf{Q}\mathbf{Q}^T = \mathbf{I}$ (Identity matrix), where $\mathbf{Q}^T$ is the transpose of $\mathbf{Q}$ and $\mathbf{I}$ denotes the identity matrix of the same dimension as $\mathbf{Q}$.

Gaussian Distribution: The Gaussian Distribution is characterized by two parameters, the mean μ and the variance σ^2. The probability density function of a univariate random variable X, following $\mathcal{N}(\mu, \sigma^2)$ is,

$$\phi(x) = \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{1}{2}\frac{(x-\mu)^2}{\sigma^2}}.$$

Multivariate Normal Distribution: This is a distribution defined over a vector, say $\mathbf{x} = (x_1, \ldots, x_d)$ and is characterized by the parameters $\boldsymbol{\mu}, \boldsymbol{\Sigma}$ and denoted by $\mathcal{MVN}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$. Here, $\boldsymbol{\mu} = \mathsf{E}(\mathbf{x})$ is a d-dimensional vector $\boldsymbol{\mu} = (\mu_1, \ldots, \mu_d)$ and $\boldsymbol{\Sigma}$ is a $d \times d$ matrix whose $(i,j)^{th}$ entry is defined as $\boldsymbol{\Sigma}_{i,j} = \mathsf{E}[(x_i - \mu_i)(x_j - \mu_j)]$ The pdf is

$$\phi(\mathbf{x}) = \frac{1}{\sqrt{2\pi|\boldsymbol{\Sigma}|}} e^{-\frac{1}{2}(\mathbf{x}-\boldsymbol{\mu})^T \boldsymbol{\Sigma}^{-1}(\mathbf{x}-\boldsymbol{\mu})}.$$

The discrete versions of the above two distribution are a simple extension.

2.3 Standard Cryptographic Primitives

Symmetric Encryption: A symmetric encryption scheme $\mathcal{SE} = (\mathcal{K}ey\mathcal{G}en, \mathcal{E}nc, \mathcal{D}ec)$ with associated plaintext space $\mathbf{M}$ and ciphertext space $\mathbf{C}$ consists of three algorithms. The randomized key generation algorithm $\mathcal{K}ey\mathcal{G}en$ returns a secret key K. The (possibly randomized) encryption algorithm $\mathcal{E}nc$ takes the secret key K and a plaintext $m \in \mathbf{M}$ to return a ciphertext $c \in \mathbf{C}$. The deterministic decryption algorithm $\mathcal{D}ec$ takes the secret key K and a ciphertext c to return a corresponding plaintext m. We require the usual correctness condition, namely that $\mathcal{D}ec(K, \mathcal{E}nc(K, m)) = m$ for all K output by $\mathcal{K}eyGen$ and all $m \in \mathbf{M}$. Finally, we say that $\mathcal{SE}$ is deterministic if $\mathcal{E}nc$ is deterministic.

Pseudorandom Functions: A function family $F : \{0,1\}^s \times \{0,1\}^n \to \{0,1\}^m$ is called a *pseudorandom function* (PRF) if the following holds:

- There is an efficient algorithm that given a key $K \in \{0,1\}^s$ and an input $x \in \{0,1\}^n$ computes $F_K(x) = F(K, x)$.
- For any algorithm A, its advantage

$$\mathbf{Adv}^{\mathrm{prf}}_F(A) \;=\; \Pr_{K \xleftarrow{\$} \{0,1\}^s}\left[A^{F_K(\cdot)} = 1\right] - \Pr\left[A^{\$} = 1\right]$$

is negligible, where \$ above denotes the oracle that implements a random function from $\{0,1\}^n \to \{0,1\}^m$.

2.4 Specializations

We specialize some notions relevant to our setting. A dataset is composed of a list of vectors, which we refer to as messages $\mathbf{m} = (m[1], \ldots, m[d])$, Messages $\mathbf{m}$ lie in a bounded d-dimensional discrete space $\mathcal{M}$. We denote by $|\mathsf{D}|$ the number of messages in D. Component $m[i]$ is also referred to as the i^{th} *attribute* of $\mathbf{m}$. On some occasions, we might drop the subscript and write only $[-M, M]$ for *attribute*-wise message space when appropriate and obvious from context.

In all our results, messages will be sampled independently from a *multivariate* distribution $\mathcal{MD}$ with support $\mathcal{M}$ (note the attributes in a given message may still depend on one another). If there is a dataset D, where each message $\mathbf{m} \in \mathsf{D} \sim \mathcal{MD}$, at times, we take the liberty to refer to it as a distribution $\mathcal{MD}$ of a dataset D. Furthermore, each attribute $m[i]$ is assumed sampled from a distribution $\mathcal{D}_i$ defined on $[-M, M]$. Thus, $\mathcal{MD} = (\mathcal{D}_1, \ldots, \mathcal{D}_d)$. Again, we might drop the subscript and write only $\mathcal{D}$ when appropriate.

3 Approximate Distance-Comparison-Preserving Functions and Their Properties

Before turning to corresponding encryption schemes, we give definitions of distance-comparison-preserving functions and related notions. Our central notion is that of approximate distance preservation and others will serve as auxiliary notions we relate to it. Note that the definitions have been presented in a generalized form. The domain, range and parameter space can be easily chosen as per need.

3.1 Notions Considered

Below *we allow functions to be randomized.* If $f : \mathcal{X} \to \mathcal{Y}$ is a randomized function then when $f(x)$ occurs in an equation it means the equation should hold *for any possible outcome of the coins.* Note that in this subsection, we use generic symbols, as these notions can be applied on a variety of domains. For instance, x, y can be vectors or numbers.

Distance-Preserving (DP) Function: A function $f : \mathcal{X} \to \mathcal{Y}$ is said to be DP if

$$\forall x, y \in \mathcal{X} : \mathsf{dist}(x, y) = \mathsf{dist}(f(x), f(y)) \ .$$

Approximate-Distance-Preserving ((α, β')-DP) Function: Let $\alpha \in \mathsf{R}, \beta' \in \mathsf{R}^+$. A function $f : \mathcal{X} \to \mathcal{Y}$ is said to be (α, β')-DP if

$$\forall x, y \in \mathcal{X} : \alpha\, \mathsf{dist}(x, y) - \beta' < \mathsf{dist}(f(x), f(y)) < \alpha\, \mathsf{dist}(x, y) + \beta' \ .$$

Distance-Comparison-Preserving (DCP) Function: A function $f : \mathcal{X} \to \mathcal{Y}$ is said to be DCP if

$$\forall x, y, z \in \mathcal{X} : \mathsf{dist}(x, y) < \mathsf{dist}(x, z) \implies \mathsf{dist}(f(x), f(y)) < \mathsf{dist}(f(x), f(z)) \ .$$

Approximate-Distance-Comparison-Preserving (β-DCP) Function: For $\beta \in \mathsf{R}^+$, a function $f \colon \mathcal{X} \to \mathcal{Y}$ is said to be β-DCP if

$$\forall x, y, z \in \mathcal{X} : \mathsf{dist}(x, y) < \mathsf{dist}(x, z) - \beta \implies \mathsf{dist}(f(x), f(y)) < \mathsf{dist}(f(x), f(z)) \ .$$

Note that a function f is 0-DCP $\iff$ f is DCP.

Using encryption systems that are distance-preserving has been proven to be highly insecure [TPS+08, LGK06]. Hence we introduce further notions which help to achieve the necessary security requirements. The reason we concentrate on approximate-distance-comparison-preserving functions over DCP functions is that the former comprises of functions whose formulations are independent of the dataset. This is because we set the approximation fator as a constant independent of the underlying dataset. Exact distance comparison-preserving encryptions need to have parameters that depend on particular datapoints in the message space. The notion of β-DCP does not have any such restrictions as the bounds on the perturbations are independent of the dataset on which it is being applied.

3.2 Accuracy of Nearest Neighbors for β-DCP Functions

When using an existing nearest-neighbour search algorithm with a β-DCP function, our goal is to guarantee some reasonable bounds on the accuracy of the algorithm. The following claim proves that any nearest-neighbour search algorithm run on a set of points after post processing by a β-DCP function returns a point whose plaintext distance from the user query is no more than β larger than the distance to the actual nearest neighbour.

Let NN be a Nearest-Neighbor algorithm that is given query q and a set of points P and $\mathsf{NN}(q, \mathsf{P})$ returns $s \in \mathsf{P}$ if $\forall x \in P : \mathsf{dist}(q, s) \leq \mathsf{dist}(q, x)$, i.e., s is the nearest neighbor for q.

Let f be a β-DCP function. Consider a run of NN with query $f(q)$ and set $f(\mathsf{P})$. Let s^* be such that $\mathsf{NN}(f(q), f(\mathsf{P})) = f(s^*)$ (which exists since NN returns a value in $f(\mathbf{P})$.)

Claim. $\forall x \in \mathsf{P} : \mathsf{dist}(q, s^*) \leq \mathsf{dist}(q, x) + \beta$.

Proof. Assume that for some $x \in \mathsf{P}$ we had $\mathsf{dist}(q, s^*) > \mathsf{dist}(q, x) + \beta$, that is, $\mathsf{dist}(q, x) < \mathsf{dist}(q, s^*) - \beta$. Since f is β-DCP, this implies $\mathsf{dist}(f(q), f(x)) < \mathsf{dist}(f(q), f(s^*))$. But since $\mathsf{NN}(f(q), f(\mathsf{P})) = f(s^*)$, and thus $\mathsf{dist}(f(q), f(s^*)) \leq \mathsf{dist}(f(q), f(x))$, this is a contradiction. ∎

We stress that β is the worst-case error in predicting the nearest neighbours.

4 Relations Between the Notions

4.1 Structure of DP Functions

In order to explain the choice of our encryption scheme, we characterize approximate distance-preserving DP functions. The theorems and corollaries below state that approximate distance-comparison-preserving (DCP) functions are indeed approximate distance-preserving. Further, by definition, approximate DP functions can be formed by adding a certain amount of perturbation to a DP function. Hence, the intuition is to start from a DP function and add a bounded perturbation to it to obtain an approximate DCP function. It is therefore sufficient to characterize DP functions.

We take inspiration from examples of isometries, i.e., distance-preserving transformations [BQ53] to provide an intuition for the choice of encryption functions.

Theorem 1. *(Informal) [Wei00] Isometries of a plane are linear transformations that preserve distance. They are characterized by rotation, translation, reflection and the identity map.*

For the formal version of the above theorem refer to [BQ53]. Although scaling as a transformation does not necessarily preserve distance (except the trivial identity map and reflection at the origin), all mutual distances are scaled by the same factor, thereby keeping relative distance comparisons unchanged. This therefore yields a distance-comparison-preserving transformation.

4.2 DCP Functions vs Approximately DP Functions

To get a handle on designing DCPE, we would like to understand how DCP functions "behave." In this subsection we analyse the relation between 0-DCP and approximately *distance-preserving* functions.

Following is our main result for one dimension (which is a bit simpler to understand) and its generalization to arbitrary dimensions.

Theorem 2. *Let $N \in \mathsf{R}, M \in \mathsf{N}$ and $f\colon [0,M] \to \mathsf{R}$ be such that $f(0) = 0$ and $f(M) = N$. If f is* DCP *then for all $0 < x < M$ we have $\frac{N}{M}(x-1) \le f(x) \le \frac{N}{M}(x+1)$.*

Note that the assumption $f(0) = 0$ is wlog by translation.

Proof. Assume that for some x:

$$f(x) \le \tfrac{N}{M}(x-1)\ . \tag{1}$$

Let x be such that $\frac{N}{M}(x-1) - f(x)$ is maximal; then for all y, we have that $\frac{N}{M}(y-1) - f(y) \le \frac{N}{M}(x-1) - f(x)$ and thus

$$f(y) \ge \tfrac{N}{M}(y-x) + f(x)\ . \tag{2}$$

Suppose $x \le \frac{M}{2}$. Then by (1): $f(x) - f(0) \le \frac{N}{M}(x-1)$ and and by (2), with $y = 2x-1$: $f(2x-1) - f(x) \ge \frac{N}{M}(x-1)$. (Note that $x \le \frac{M}{2}$ ensures that $y \le M$.)

This means that there are three points, 0, x and $2x-1$, such that $\mathsf{dist}(0,x) > \mathsf{dist}(x, 2x-1)$, but $\mathsf{dist}(f(0), f(x)) \le \mathsf{dist}(f(x), f(2x-1))$ and thus f is not DCP.

The proof for $x > \frac{M}{2}$ is similar with the contradicting points being $2x - M + 1$, x and M. This proves the lower bound.

The proof for the upper bound is similar, taking x to be such that $f(x) - \frac{N}{M}(x+1)$ is maximal and proceeding with steps analogous to the above. ∎

Now, we generalise this results to the following theorems that cater to multiple dimensions. We restrict ourselves to stating just the theorems in the main body and send the proofs to Appendix A to not divert from the main attention of this work. Moreover, the proofs for multi-dimensional setting follows a very similar line of argument as the one-dimensional setup.

Theorem 3. *Let U be a subset of R and $[0,M]^d$ denote the d-dimensional Cartesian product of the closed interval $[0,M]$ and $\mathbf{M} = (M, M, \ldots, M)$. Let $f : [0,M]^d \to \mathsf{U}$ with $f(\mathbf{0}) = (\mathbf{0})$ and $\mathbf{N} := f(\mathbf{M})$. If f is* DCP*, then for all $\mathbf{x} \in [0,M]^d \setminus \{0, \mathbf{M}\}$, we have $\frac{\|\mathbf{N}\|}{\|\mathbf{M}\|}(\|\mathbf{x}\| - \sqrt{d}) \le \|f(\mathbf{x})\| \le \frac{\|\mathbf{N}\|}{\|\mathbf{M}\|}(\|\mathbf{x}\| + \sqrt{d})$.*

This theorem tells us that given two fixed points, a distance-comparison-preserving function always maps a point $\mathbf{x} \in [0,M]^d$ to a point $\mathbf{x}' \in \mathsf{U}$ such that $\mathbf{x}'$ lies in a ball of radius $\sqrt{d}$ around $\mathbf{x}$ scaled by a constant. This shows that any DCP function is approximately distance-preserving. In particular, our claim proves a bound on the amount any point in a DCP function can be perturbed. For concreteness, we have chosen euclidean spaces for our results, but they can be easily generalized to any generic metric space.

Corollary 1. *Let U be any subset R and $[0,M]^d$ denote the d-dimensional Cartesian products of the closed interval $[0,M]$ and $\mathbf{M} = (M, M, \cdots, M)$. Let $f : [0,M]^d \to \mathsf{U}$ be β-*DCP*, and $f(\mathbf{0}) = (\mathbf{0}), f(\mathbf{M}) = \mathbf{N} \in \mathsf{U}$. Then $\forall \mathbf{x} \in [0,M]^d - \{\mathbf{0}, \mathbf{M}\}$, $\frac{\|\mathbf{N}\|}{\|\mathbf{M}\|}(\|\mathbf{x}\| - \beta) \le \|f(\mathbf{x})\| \le \frac{\|\mathbf{N}\|}{\|\mathbf{M}\|}(\|\mathbf{x}\| + \beta)$.*

The above corollary follows directly from Theorem 14. The only difference is that the radius of the ball in which the projected point lies is β. This validates our claim that any β-DCP function is also approximately distance preserving with higher perturbations.

4.3 Approximately DCP Functions vs Approximately DP Functions

We study the relation between approximately DCP and approximately DP functions. Our main result is that approximately DCP functions are approximately DP with a large value of β'.

Theorem 4 (1-Dimension). *Let $N \in \mathsf{R}, M \in \mathsf{N} \cup \{0\}$, if $f : \mathsf{N} \to \mathsf{R}$ is a β-DCP, and $f(0) = 0, f(M) = N$, then $\forall x, y$ such that $0 < x, y < M$,*

$$\frac{N}{M}(\mathsf{dist}(x,y) - 2\beta) \leq \mathsf{dist}(f(x), f(y)) \leq \frac{N}{M}(\mathsf{dist}(x,y) + 2\beta) \ .$$

In other words, if f is β-DCP, then f is (α, β')-DP, where $\alpha = \frac{N}{M}$ and $\beta' = 2\frac{N}{M}\beta$.

The main goal of this section is to show that an "ideal" notion of security (cf [BCLO09]) for DCPE is not achievable. The proof of Theorem 4 follows from the proof of Corollary 2.

The following is the converse of the above theorem.

Theorem 5. *For any function $f : \mathsf{N} \to \mathsf{R}$ such that f is (α, β)-DP, f is β'-DCP for all $\beta' \geq \frac{2\beta}{\alpha}$.*

Proof. Let f be an arbitrary α, β-DP function. Let x, y, and z be points such that $\mathsf{dist}(x,y) < \mathsf{dist}(x,z) - \beta'$. By the definition of (α, β)-DP, we have that $\mathsf{dist}(f(x), f(y)) < \alpha\,\mathsf{dist}(x,y) + \beta$ implies $\frac{\mathsf{dist}(f(x),f(y))-\beta}{\alpha} < \mathsf{dist}(x,y)$. Moreover, $\alpha\,\mathsf{dist}(x,z) - \beta < \mathsf{dist}(f(x), f(z))$ implies $\mathsf{dist}(x,z) < \frac{\mathsf{dist}(f(x),f(z))+\beta}{\alpha}$. By our choice of x, y, and z, $\mathsf{dist}(x,y) < \mathsf{dist}(x,z) - \beta'$, so

$$\begin{aligned} \frac{\mathsf{dist}(f(x), f(y)) - \beta}{\alpha} &< \mathsf{dist}(x,y) < \mathsf{dist}(x,z) - \beta' \\ &< \frac{\mathsf{dist}(f(x), f(z)) + \beta}{\alpha} - \beta' \\ \implies \mathsf{dist}(f(x), f(y)) &< \mathsf{dist}(f(x), f(z)) + 2\beta - \alpha\beta'. \end{aligned}$$

Therefore, as long as $2\beta - \alpha\beta' \leq 0$, $\mathsf{dist}(f(x), f(y)) < \mathsf{dist}(f(x), f(z))$, so for $\beta' \geq \frac{2\beta}{\alpha}$, f is β'-DCP. ∎

Like the previous subsection, we again generalize to multiple dimensions. Proofs for the same can be found in Appendix A.

Theorem 6 (n-Dimensions). *If $f : [0, M]^d \to \mathsf{U}$ is β-DCP, and $f(\mathbf{0}) = \mathbf{0}, f(M, \cdots, M) = \mathbf{N}$, then $\forall \mathbf{x}, \mathbf{y} \in [0, M]^d - \{\mathbf{0}, \mathbf{M}\}$,*

$$\frac{\|\mathbf{N}\|}{\|\mathbf{M}\|}(\mathsf{dist}(\mathbf{x},\mathbf{y}) - \gamma - 2\beta) < \mathsf{dist}(f(\mathbf{x}), f(\mathbf{y})) < \frac{\|\mathbf{N}\|}{\|\mathbf{M}\|}(\mathsf{dist}(\mathbf{x},\mathbf{y}) + \gamma + 2\beta)$$

where $\gamma = \sqrt{2\|\mathbf{x}\|\|\mathbf{y}\| + 2|\mathbf{x}\mathbf{y}|}$.

In other words, if f is β-DCP, then f is (α, β')-DP, where

$$\alpha = \frac{\|\mathbf{N}\|}{\|\mathbf{M}\|} \quad \text{and} \quad \beta' = \frac{\|\mathbf{N}\|}{\|\mathbf{M}\|}(\sqrt{2\|\mathbf{x}\|\|\mathbf{y}\| + 2|\mathbf{x}\mathbf{y}|} + 2\beta) \ .$$

(Approximate-)Distance-Comparison-Preserving Encryption ($(\beta$-)DCPE): We say that a symmetric key encryption scheme $\mathcal{SE} = (\mathcal{K}eyGen, \mathcal{E}nc, \mathcal{D}ec)$ with plaintext and ciphertext spaces $\mathcal{X}$ and $\mathcal{Y}$ is (approximate-)distance-comparison-preserving if $\mathcal{E}nc(K, \cdot)$ is a (β-)DCP function from $\mathcal{X}$ to $\mathcal{Y}$ for all K output by $\mathcal{K}eyGen$.

4.4 Impossibility of Ideal Security

As in the study of its predecessor OPE [BCLO09], a first question about β-DCPE is whether it can achieve "ideal" security, meaning it leaks *only* the approximate distance comparisons between the plaintexts. As in the case of OPE, the answer is "no". (However, there is a caveat, hence we refer readers to Remark 1.) Toward this end, we first introduce the relevant definition.

Indistinguishability-based Security of approximate DCPE

Definition 1. *Let $\mathcal{SE} = (\mathcal{K}eyGen, \mathcal{E}nc, \mathcal{D}ec)$ be a β-*DCPE *scheme with message space $\mathcal{M} = [-M, M]^d$ and ciphertext space $\mathcal{C} = [-C, C]^d$. For an adversary A, define its* LoR*-advantage*

$$\mathbf{Adv}^{\mathsf{LoR}}_{\mathcal{SE}}(A) = 2 \cdot \mathbf{Pr}\left[\mathbf{Exp}^{\mathsf{LoR}}_{\mathcal{SE}}(A) = 1\right] - 1$$

where the experiment above is defined as follows:

Experiment $\mathbf{Exp}^{\mathsf{LoR}}_{\mathcal{SE}}(A)$:
$K \xleftarrow{\$} \mathcal{K}eyGen$
$b \xleftarrow{\$} \{0,1\}$
$b' \xleftarrow{\$} A^{\mathsf{LR}(\cdot,\cdot)}$
If $b == b'$
  return 1
Else return 0

Oracle $\mathsf{LR}(\mathbf{m}_0, \mathbf{m}_1)$*:*
$\mathbf{c} \xleftarrow{\$} \mathcal{E}nc_K(\mathbf{m}_b)$
Return $\mathbf{c}$

Let $(\mathbf{m}^0_1, \mathbf{m}^1_1), \ldots, (\mathbf{m}^0_q, \mathbf{m}^1_q)$ be a sequence of queries made by A to its oracle. We call A an LoR*-adversary if for every such sequence the following holds for all $i, j, k \in [q]$:*

$$\mathsf{dist}(\mathbf{m}^0_i, \mathbf{m}^0_j) \leq \mathsf{dist}(\mathbf{m}^0_j, \mathbf{m}^0_k) - \beta \;\Rightarrow\; \mathsf{dist}(\mathbf{m}^1_i, \mathbf{m}^1_j) \leq \mathsf{dist}(\mathbf{m}^1_j, \mathbf{m}^1_k) - \beta.$$

This can be seen as a special case of the notion introduced by [PR12]. We say that $\mathcal{SE}$ is *ideal-secure* if $\mathbf{Adv}^{\mathsf{LoR}}_{\mathcal{SE}}(A)$ is small for every efficient LoR-adversary A.

Impossibility Result

We show that no β-DCPE scheme is ideal-secure unless β is likely too large to be useful in applications. The proof relies on a "Big-jump" style attack as in [BCLO09] that uses only two pairs of oracle queries.

Theorem 7. *Let $\mathcal{SE} = (\mathcal{K}eyGen, \mathcal{E}nc, \mathcal{D}ec)$ be a β-*DCPE *scheme with plaintext space $\mathcal{M} = [0, M]^d$. Let $\mathbf{m} \in \mathcal{M}$ be such that $\forall\, \mathbf{x} \in \mathcal{M}, \|\mathbf{x}\| \leq \|\mathbf{m}\|$. If $\beta < \frac{\|\mathbf{m}\|}{4}$ then there is an LoR-adversary A such that $\mathbf{Adv}^{\mathsf{LoR}}_{\mathcal{SE}}(A) = 1$.*

Proof. Consider the following adversary:

Algorithm 1 Big-Jump Adversary

procedure A_{BigJump}
  $\mathbf{n}$ $\mathsf{LR}(\mathbf{m}, \mathbf{m})$; α $\frac{\|\mathbf{n}\|}{\|\mathbf{m}\|}$; β' $2\alpha\beta$
  $\mathbf{c}^b$ $\mathsf{LR}(\mathbf{0}, \mathbf{m})$
  If $\mathsf{dist}(\mathbf{n}, \mathbf{c}^b) < \beta'$, then return b' 1
  Return b' 0

It is vacuously true that A_{BigJump} is LoR. We now argue that $\mathbf{Adv}^{\mathsf{LoR}}_{\mathcal{SE}}(A_{\mathsf{BigJump}}) = 1$. Theorem 4 of Section 4.3 tells that that any β-DCP function is (α, β')-DP, where $\alpha = \frac{\|\mathbf{n}\|}{\|\mathbf{m}\|}$ and $\beta' = 2\frac{\|\mathbf{n}\|}{\|\mathbf{m}\|}\beta$, i.e. for any $\mathbf{m}_1, \mathbf{m}_2$ and any possible corresponding ciphertexts (for any possible key and randomness of the encryption) $\mathbf{c}_1, \mathbf{c}_2$ we have

$$\frac{\|\mathbf{n}\|}{\|\mathbf{m}\|} \cdot \mathsf{dist}(\mathbf{m}_1, \mathbf{m}_2) - 2\frac{\|\mathbf{n}\|}{\|\mathbf{m}\|}\beta \leq \mathsf{dist}(\mathbf{c}_1, \mathbf{c}_2) \leq \frac{\|\mathbf{n}\|}{\|\mathbf{m}\|} \cdot \mathsf{dist}(\mathbf{m}_1, \mathbf{m}_2) + 2\frac{\|\mathbf{n}\|}{\|\mathbf{m}\|}\beta.$$

If $\mathbf{c}^b$ corresponds to the encryption of $\mathbf{m}$, the above expression gives us $\mathsf{dist}(\mathbf{c}^b, \mathbf{n}) \leq 2\frac{\|\mathbf{n}\|}{\|\mathbf{m}\|}\beta$. Whereas $\mathsf{dist}(\mathbf{c}^b, \mathbf{n}) \geq \|\mathbf{n}\| - 2\frac{\|\mathbf{n}\|}{\|\mathbf{m}\|}\beta$ when $\mathbf{0}$ is encrypted by LR.

If $\beta < \frac{\|\mathbf{m}\|}{4}$, then $2\frac{\|\mathbf{n}\|}{\|\mathbf{m}\|}\beta < \|\mathbf{n}\| - 2\frac{\|\mathbf{n}\|}{\|\mathbf{m}\|}\beta$, and thus $\mathbf{Adv}^{\mathsf{LoR}}_{\mathcal{SE}}(A_{\mathsf{BigJump}}) = 1$. ∎

Remark 1. This proof does not necessarily hold if the data is subjected to preprocessing on the entire database before encryption. One such example of preprocessing is *shuffling* which has been explored in Section 5.3. It must be noted that shuffling does not contradict this proof. However there is no guarantee that there does not exist *any* preprocessing method which can bypass this result. In the case of OPE, with preprocessing ideal security *can* be achieved. We are unsure if this can be achieved for DCPE using some preprocessing and leave this for future work.

5 The Scale-and-Perturb (SAP) Scheme

We first give our core encryption scheme and then discuss additional preprocessing techniques.

5.1 Our Core β-DCPE Scheme

We now propose our core β-DCPE scheme based on our prior characterization of approximation distance-comparison preserving functions. We will suggest data preprocessing techniques in Section 5.3 in addition. Let $\mathcal{M} = [-M, M]^d$ be a *discrete* message space of dimension d. Let $\mathsf{PRF}\colon \{0,1\}^k \times \{0,1\}^\ell \to \{0,1\}^*$ be a function family for some $k, \ell \in \mathsf{N}$. We leave the number of output bits implicit in our algorithms for simplicity. The scheme is also parameterized by β which can take any non negative value less than $2M\sqrt{(}d)$.

The keyspace is denoted by $\mathcal{S}$. Define the "Scale-And-Perturb" (SAP) encryption scheme on $\mathcal{M}$ as

$$\mathsf{SAP} = \mathsf{SAP}[\mathsf{PRF}, \beta, \mathsf{S}] = (\mathcal{K}eyGen_{\mathsf{SAP}}, \mathcal{E}nc_{\mathsf{SAP}}, \mathcal{D}ec_{\mathsf{SAP}})$$

as shown in Algorithm 2.

Algorithm 2 The SAP scheme.

procedure $\mathcal{K}eyGen_{\mathsf{SAP}}()$
- $s \xleftarrow{\$} \mathcal{S}$
- $K \xleftarrow{\$} \{0,1\}^k$
- return (s, K)

procedure $\mathcal{E}nc_{\mathsf{SAP}}((s, K), \mathbf{m})$
- $n \xleftarrow{\$} \{0,1\}^\ell$
- $\mathsf{coins}_1 \| \mathsf{coins}_2 \leftarrow \mathsf{PRF}(K, n)$
- $\mathbf{u} \leftarrow \mathcal{N}(0, I_d; \mathsf{coins}_1)$
- $x' \leftarrow \mathcal{U}(0, 1; \mathsf{coins}_2)$
- $x \leftarrow \frac{s\beta}{4}(x')^{\frac{1}{d}}$; $\lambda_{\mathbf{m}} \leftarrow \frac{\mathbf{u}x}{\|\mathbf{u}\|}$
- $\mathbf{c} \leftarrow s\mathbf{m} + \lambda_{\mathbf{m}}$
- return $(\mathbf{c}, n)$

procedure $\mathcal{D}ec_{\mathsf{SAP}}((s, K), (\mathbf{c}, n))$
- $\mathsf{coins}_1 \| \mathsf{coins}_2 \leftarrow \mathsf{PRF}(K, n)$
- $\mathbf{u} \leftarrow \mathcal{N}(0, I_d; \mathsf{coins}_1)$
- $x' \leftarrow \mathcal{U}(0, 1; \mathsf{coins}_2)$
- $x \leftarrow \frac{s\beta}{4}(x')^{\frac{1}{d}}$
- $\lambda_{\mathbf{m}} \leftarrow \frac{\mathbf{u}x}{\|\mathbf{u}\|}$
- $\mathbf{m} \leftarrow \frac{\mathbf{c} - \lambda_{\mathbf{m}}}{s}$
- return $\mathbf{m}$

Scheme Overview: The coins generated using the PRF are used for decryption. They provide a unique identity to each plaintext-ciphertext pair which makes decryption possible. The TapeGen PRF from [BCLO09] can be a candidate PRF. Of course, security of our scheme also depends on the chosen PRF. For simplicity, we do not

talk about the PRF in the remainder, analyzing the core (no-decrypt) scheme. Our results all then transfer to the scheme described above *mutatis mutandis*.

The scaling factor is selected uniformly at random from the keyspace $\mathcal{S}$ (Line 1 of $\mathcal{K}eyGen_{\mathsf{SAP}}$). The choice of the size of $\mathcal{S}$ does not affect utility but has an influence on the one-wayness bounds. Specific values of the size would vary based on applications and we have tabulated some results in Table 2. Take note that λ (which we sometimes denote as $\lambda_{\mathbf{m}}$ since it is chosen independently for each message $\mathbf{m} \in \mathcal{M}$) is a d-dimensional vector whose norm has an upper bound of $\frac{s\beta}{4}$. We sample it in such a way that λ_m is chosen uniformly from the d-dimensional ball of radius $\frac{s\beta}{4}$ (Line 6 of $\mathcal{E}nc_{\mathsf{SAP}}$). To do so, we first generate a vector from a multivariate normal distribution with mean $\mathbf{0}$ and variance I_d, which is a d-dimensional identity matrix. The uniform point inside the ball is generated by multiplying the standardized version (point divided by its norm) of this point with the d^{th} root of a uniformly generated point from $[0, 1]$ followed by re-scaling with the radius of the ball. This mechanism ensure that the each point inside the ball can be sampled with uniform probability [HL10].

Claim 2: For any scaling factor $s \in \mathcal{S}$, $\mathcal{E}nc_{\mathsf{SAP}}(s, \cdot)$ is β-DCP.

Proof. Denote $\mathcal{E}nc_{\mathsf{SAP}}(s, \cdot)$ by $f(\cdot)$ for notational simplicity. Let $\mathbf{x}, \mathbf{y}, \mathbf{z} \in \mathbf{M}$. Suppose $\|\mathbf{x} - \mathbf{y}\| < \|\mathbf{y} - \mathbf{z}\| - \beta$.
$f(\mathbf{x}) = s\mathbf{x} + \lambda_{\mathbf{x}}$.
$f(\mathbf{y}) = s\mathbf{y} + \lambda_{\mathbf{y}}$.
$f(\mathbf{z}) = s\mathbf{z} + \lambda_{\mathbf{z}}$.
Hence,

$$
\begin{aligned}
\|f(\mathbf{x}) - f(\mathbf{y})\| &\leq \|f(\mathbf{x}) - s\mathbf{x}\| + \|s\mathbf{x} - s\mathbf{y}\| + \|f(\mathbf{y}) - s\mathbf{y}\| && (1)\\
&= \|\lambda_{\mathbf{x}}\| + s\|\mathbf{x} - \mathbf{y}\| + \|\lambda_{\mathbf{y}}\| \\
&< \|\lambda_{\mathbf{x}}\| + \|\lambda_{\mathbf{y}}\| + s(\|\mathbf{y} - \mathbf{z}\| - \beta) && (2)\\
&= \|s\mathbf{y} - s\mathbf{z}\| - s\beta + \|\lambda_{\mathbf{x}}\| + \|\lambda_{\mathbf{y}}\| \\
&< \|s\mathbf{y} - s\mathbf{z}\| - s\beta + s\beta/4 + s\beta/4 && (3)\\
&= \|s\mathbf{y} - s\mathbf{z}\| - s\beta/2 \\
&< \|s\mathbf{y} - s\mathbf{z}\| - (\|\lambda_{\mathbf{z}}\| + \|\lambda_{\mathbf{y}}\|) \\
&\leq \|s\mathbf{y} - s\mathbf{z}\| - (\|\lambda_{\mathbf{z}} - \lambda_{\mathbf{y}}\|) && (4)\\
&\leq \|(s\mathbf{y} - s\mathbf{z}) - (\lambda_{\mathbf{z}} - \lambda_{\mathbf{y}})\| \\
&= \|f(\mathbf{y}) - f(\mathbf{z})\| .
\end{aligned}
$$

where (1) Triangle Inequality, (2) by assumption, (3) $\|\lambda_{\mathbf{x}}\|, \|\lambda_{\mathbf{y}}\| < s\beta/4$ and (4) Triangle Inequality. ∎

Thus, $\|\mathbf{x} - \mathbf{y}\| < \|\mathbf{y} - \mathbf{z}\| - \beta \implies \|f(\mathbf{x}) - f(\mathbf{y})\| < \|f(\mathbf{y}) - f(\mathbf{z})\|$,) f is β-DCP.

5.2 Scale and Perturb — Isometry

SAP-I ("Scale And Perturb-Isometry"): (Generalized Version of SAP) The β-DCP function used in the SAP scheme mentioned above is a particular instance among the family of all possible β-DCP functions.

Our intuition to generate Approximate Distance Comparison Preserving transforms is to use the structureal relation between Distance Preserving and Distance Comparison Preserving transformations, and add some noise to the function. Now, as a linear transformation, an orthogonal matrix [Kea89] preserves the dot product of vectors, and therefore acts as an isometry of Euclidean space which is the requirement for a Distance Preserving transformation. An orthogonal matrix however, involves rotation and reflection only and fails to represent translation which is also a distance distance preserving transformation. In addition, we also note that scaling a distance preserving function transforms it to a distance comparison preserving (DCP) transformation. Therefore, combining all these operation provides a recipe to construct generic DCP functions.

Theorem 8. *Given an orthogonal matrix,* Q, $\forall s \in \mathsf{R}, \mathbf{x}, \gamma_{\mathbf{t}}, \lambda \in \mathsf{U}$, *where* $\|\lambda\| \leq \frac{s\beta}{4}$. *If* $f(\mathbf{x}) = s\mathsf{Q}\mathbf{x} + \gamma_{\mathbf{t}} + \lambda$ *then* f *is a* β-DCP *function.*

Proof. Let $\mathbf{x}, \mathbf{y}, \mathbf{z} \in \mathsf{U}$. Suppose $\|\mathbf{x}-\mathbf{y}\| < \|\mathbf{y}-\mathbf{z}\| - \beta \implies \|\mathsf{Q}\mathbf{x}-\mathsf{Q}\mathbf{y}\| < \|\mathsf{Q}\mathbf{y}-\mathsf{Q}\mathbf{z}\| - \beta$.

Since Q is an orthogonal matrix, it leads to a distance preserving transformation.

$f(\mathbf{x}) = s\mathsf{Q}\mathbf{x} + \gamma_{\mathbf{t}} + \lambda_{\mathbf{x}}$.

$f(\mathbf{y}) = s\mathsf{Q}\mathbf{y} + \gamma_{\mathbf{t}} + \lambda_{\mathbf{y}}$.

$f(\mathbf{z}) = s\mathsf{Q}\mathbf{z} + \gamma_{\mathbf{t}} + \lambda_{\mathbf{z}}$.

Hence,

$$\begin{aligned}
\|f(\mathbf{x}) - f(\mathbf{y})\| &\leq \|f(\mathbf{x}) - s\mathsf{Q}\mathbf{x} - \gamma_{\mathbf{t}}\| + \|s\mathsf{Q}\mathbf{x} - s\mathsf{Q}\mathbf{y}\| \\
&\quad + \|f(\mathbf{y}) - s\mathsf{Q}\mathbf{y} - \gamma_{\mathbf{t}}\| \\
&= \|\lambda_{\mathbf{x}}\| + s\|\mathsf{Q}\mathbf{x} - \mathsf{Q}\mathbf{y}\| + \|\lambda_{\mathbf{y}}\| \\
&< \|\lambda_{\mathbf{x}}\| + \|\lambda_{\mathbf{y}}\| + s(\|\mathsf{Q}\mathbf{y} - \mathsf{Q}\mathbf{z}\| - \beta) \\
&= \|s\mathsf{Q}\mathbf{y} - s\mathsf{Q}\mathbf{z}\| - s\beta + \|\lambda_{\mathbf{x}}\| + \|\lambda_{\mathbf{y}}\| \\
&< \|s\mathsf{Q}\mathbf{y} - s\mathsf{Q}\mathbf{z}\| - s\beta + s\beta/4 + s\beta/4 \\
&= \|s\mathsf{Q}\mathbf{y} - s\mathsf{Q}\mathbf{z}\| - s\beta/2 \\
&< \|s\mathsf{Q}\mathbf{y} - s\mathsf{Q}\mathbf{z}\| - (\|\lambda_{\mathbf{z}}\| + \|\lambda_{\mathbf{y}}) \\
&\leq \|s\mathsf{Q}\mathbf{y} - s\mathsf{Q}\mathbf{z}\| - (\|\lambda_{\mathbf{z}} - \lambda_{\mathbf{y}}\|) \\
&\leq \|(s\mathsf{Q}\mathbf{y} - s\mathsf{Q}\mathbf{z}) - (\lambda_{\mathbf{z}} - \lambda_{\mathbf{y}})\| \\
&= \|f(\mathbf{y}) - f(\mathbf{z})\| \ .
\end{aligned}$$

∎

Challenges with SAP − I: SAP − I would have $s, \mathsf{Q}, \gamma_{\mathbf{t}}$ as the secret keys. Analysing the security gives rise of complex product distributions which often makes it difficult to achieve security bounds in closed form. Hence, as an introductory work in this topic, we concentrate on the simpler SAP scheme, which nevertheless achieves strong security. SAP − I however be a very interesting theoretical problem as a follow up to this work.

5.3 Two Preprocessing Algorithms

To boost security and compatibility for real life application of SAP, we now propose two additional preprocessing algorithms. The first operates on the entire dataset D. The second only needs to know the *distribution* of D and can otherwise operate on the data on-the-fly. Thus, both of the transforms make stronger assumptions about the model.

Shuffle(dataset) : On input dataset D which has n entries $(\mathbf{m}_1, \mathbf{m}_2, \ldots, \mathbf{m}_n)$, sample a random permutation $\Pi\colon [n] \to [n]$. Output the transformed dataset D′ that is $m_{\pi(1)}, m_{\pi(2)}, \ldots, m_{\pi(n)}$.

Such a shuffle can be implemented using a mix network (mixnet). Very efficient implementations of mixnets handling large data exist [BEM+17, KLDF16].

Shuffling enhances the security because it hides the identity of the ciphertext from an adversary. By looking at a set of ciphertexts, the adversary cannot map it to the plaintext even if it knows them in advance. It enables security improvements without adding much computational overheads. Shuffling has been recently employed in differential privacy works [CSU+19, EFM+19, BBGN19] to achieve security enhancements while maintaining utility, which is our goal as well.

Normalize($\mathbf{m}$, $\mathcal{MD}$): On input $\mathbf{m}$ a data point coming from a multivariate distribution $\mathcal{MD}$, apply algorithm BoxCox [Sak92] (state-of-the-art normalization algorithm) to input $\mathbf{m}$ and output the result.

Intuitively, BoxCox is a transformation which takes as input the distribution of the dataset, and makes a transformation using maximum likelihood estimation. This step can be considered as a heuristic as we do not rigorously deal with the error on our analyses.

This preprocessing step will be used because our security analyses assume the data follows a multivariate normal distribution. Such an assumption has practical significance as a large number of data available in practice [Mor16] either follow this distribution or can be easily simulated as per the above if not. Note that if the data is naturally normally distributed this preprocessing step is not needed.

To combine these preprocessing steps with encryption, the idea is that the SAP is then applied to each data point output by the transformation; in the shuffling case this means the dataset is encrypted and sent all together. Naturally, the transforms can also be composed.

6 Real-or-Replaced Indistinguishability for Neighboring Datasets

To allow for the shuffling preprocessing step described in Subsection 5.3, we introduce the notion of security to accommodate the adversary querying an *dataset* rather than many points individually. In essence, we define a "real-or-repliaced" (RoR-type) definition where the oracle either shuffles and encrypts: (1) the dataset D provided by the adversary or (2) the dataset D with a random plaintext resampled uniformly below a certain distance threshold relative to the original. We thus have security wrt. "neighboring databases" (as in [Dwo08, YGFJ18]), speaking to the adversary's ability to infer whether information of a particular individual is present in the dataset. Shuffling plays a key role in our analysis here.

Real-or-Replaced Indistinguishability. Let $\mathcal{SE} = (\mathcal{K}eyGen, \mathcal{E}nc, \mathcal{D}ec)$ be a symmetric key encryption scheme. Let $\mathcal{MD}$ be the distribution from which the plaintext is sampled. In our results, *the adversary makes a single* Swap *oracle query only.*

We say that $\mathcal{SE}$ is (r, ε)-RoR-secure for $\mathcal{MD}$ if for every δ-RoR adversary A against $\mathcal{SE}$, its advantage

$$\mathbf{Adv}^{\delta-\mathsf{RoR}}_{\mathcal{SE},\mathcal{MD}}(\mathsf{A}) = 2 \cdot \mathbf{Pr}\left[\mathbf{Exp}^{\delta-\mathsf{RoR}}_{\mathcal{SE},\mathcal{MD}}(\mathsf{A}) = 1\right] - 1 \leq \varepsilon,$$

where the experiment $\mathbf{Exp}^{\delta-\mathsf{RoR}}_{\mathcal{SE},\mathcal{MD}}(\mathsf{A})$ is defined as follows (where some of its algorithms are defined below it):

Experiment $Exp^{\delta-\mathsf{RoR}}_{\mathcal{SE},\mathcal{MD}}(A)$:
$K \xleftarrow{\$} \mathcal{K}eyGen$
$\mathsf{Ctxt} \leftarrow \emptyset$
$b \xleftarrow{\$} \{0,1\}$
$b' \leftarrow \mathsf{A}^{\mathsf{Swap}(\cdot)}$
If $b == b'$ return 1
Else return 0

Oracle $\mathsf{Swap}(\mathsf{D}_0)$:
$i \xleftarrow{\$} |\mathcal{D}|$; $b \xleftarrow{\$} \{0.1\}$
$\mathsf{D}_1 \xleftarrow{\$} \mathsf{Resamp}(D_0, i, \delta)$
$\mathsf{D}'_b \xleftarrow{\$} \mathbf{shuffle}(\mathsf{D}_b)$
For all $\mathbf{m} \in \mathsf{D}'_b$
  $\mathsf{Ctxt} \leftarrow \mathsf{Ctxt.append}(\mathcal{E}nc(K, \mathbf{m}))$
return Ctxt

We stress that Ctxt is a list because the order in which the ciphertexts are presented to the adversary is important. Above, we define $\mathsf{Resample}(\mathsf{D}, i, \delta)$ to be the algorithm that on input a dataset D, an index i that denoted which message in D must be replaced and parameter δ follows the following steps:(1) Picks up the i^{th} message in D, call it $\mathbf{m}_i$. (2) Construct a d dimensional sphere of radius δ $\mathcal{B}(\mathbf{m}_i; \delta)$ around $\mathbf{m}_i$. (3) Samples a point at random inside $\mathcal{B}(\mathbf{m}_i; \delta)$ and return it. The final step can be done efficiently in the same way in which the perturbation factor is chosen for our encryption scheme.

Additionally, note that the **shuffle** step in the Swap oracle models this preprocessing step being applied.

6.1 δ-RoR Security Bounds

The following is the main result which upper bounds the $\delta - \mathsf{RoR}$ adversary's advantage.

Theorem 9. *Let $\delta \leq \frac{\beta}{2}$. For any δ-RoR adversary A generating its query D of size N according to distribution $\mathcal{MD}$,*

$$\boldsymbol{Adv}^{\delta\text{-}\mathsf{RoR}}_{\mathcal{SE}}(A) \leq \frac{2(2-p)}{N(1-p)}\left(1 - \frac{pT}{2}\right)^{N-1}$$

where $p = \left(\frac{h}{\mathsf{rad}+\frac{\delta}{2}}\right)^d$, $\mathsf{rad} = \frac{\beta}{4}$, $a = \sqrt{\mathsf{rad}^2 - \frac{\delta^2}{4}}$, $\cos(\theta) = \frac{2\mathsf{rad}^2 - 4a^2}{2\mathsf{rad}^2}$, $h = a\tan(\frac{\theta}{4})$ and $T = \mathbf{Pr}\left[\|X - Y\| \leq \delta\right]$, $(X, Y) \sim \mathcal{MD}$.

The proof of this theorem requires a careful and complex analysis. Hence, we split the proof into several sub parts.

On a given run of the experiment, let $\mathbf{c}_k \in \mathsf{C}$ be the correct ciphertext corresponding to the plaintext where D_0 and D_1 differ on. For ease of notation, let us say that $\mathsf{D}_0 = \mathsf{D} \cup \mathbf{m}_0$ and $\mathsf{D}_1 = \mathsf{D} \cup \mathbf{m}_1$, for some $|\mathsf{D}| = N - 1$. Recall b' is the guess of the bit b made by A.

Canonical Adversaries. We say that A_{Can} is *canonical* if it can be written as two algorithms $\mathsf{A}_{Can} = (\mathsf{A}_1, \mathsf{A}_2)$ each of which has the following form. A_1 queries D according to $\mathcal{MD}$, it receives Ctxt and chooses some ciphertext $\mathbf{c} \in \mathsf{Ctxt}$. Then it chooses some messsage from the message space, $\mathbf{m} \in \mathcal{M}$ depending only on $\mathbf{c}$. Following this, A_2 receives $(\mathbf{m}, \mathbf{c})$ as input and makes a guess b' depending only on $\mathbf{m}$ and $\mathbf{c}$. Let DS be the set of all datasets that conform to the distribution $\mathcal{MD}$ (i.e., all elements of DS are sampled from $\mathcal{MD}$). Formally, any canonical adversary $\mathsf{A}_{Can} = (\mathsf{A}_1, \mathsf{A}_2)$ can be written as follows:

Algorithm 3 Canonical Adversary

procedure $\mathsf{A}_{Can} = (\mathsf{A}_1, \mathsf{A}_2)$
$(\mathbf{m}, \mathbf{c}) \quad \mathsf{A}_1^{\mathsf{Swap}(\cdot)}(\mathcal{MD}, \mathsf{DS})$
$b' \quad \mathsf{A}_2(\mathbf{m}, \mathbf{c})$

Defining random variables for the proof: $\mathbf{M}_0, \cdots, \mathbf{M}_N$ represent messages in the dataset. $\mathbf{C}_0, \cdots, \mathbf{C}_N$ represent messages in the corresponding ciphertext set. Hence a particular instance of the dataset that looks like $D \in \mathsf{DS} = (\mathbf{m}_0, \cdots, \mathbf{m}_N)$ means $M_0 = \mathbf{m}_0, \cdots, M_N = \mathbf{m}_N$ and similarly for the ciphertexts. $\mathbf{M}_c$ and $\mathbf{C}_c$ denote the *correct* message and ciphertext respectively, i.e. $\mathbf{M}_c = \mathbf{m}_b$ and $\mathbf{C}_c = \mathbf{c}_k$.

Claim. For every δ-RoR' adversary A, there exists a canonical adversary B such that

$$\mathbf{Adv}^{\delta-\mathsf{RoR}}_{\mathcal{SE}}(\mathsf{A}) \leq \frac{2}{1-p}\mathbf{Adv}^{\delta-\mathsf{RoR}}_{\mathcal{SE}}(\mathsf{B}) \ .$$

where $p = \left(\frac{h}{\mathsf{rad}+\frac{\delta}{2}}\right)^d$, $\mathsf{rad} = \frac{\beta}{4}$, $a = \sqrt{\mathsf{rad}^2 - \frac{\delta^2}{4}}$, $\cos(\theta) = \frac{2\mathsf{rad}^2 - 4a^2}{2\mathsf{rad}^2}$, $h = a\tan(\frac{\theta}{4})$.

Proof. It suffices to prove that for *every* δ-RoR' adversary A there exists a *cannonical* δ-RoR adversary B such that the δ-RoR'-advantage of B is at least that of A,

Adversary B runs A on random coins coins. When A makes a query D_0, B forwards the query to its own oracle, to receive result Ctxt, Adversary B replies to A with Ctxt. A outputs a guess bit to B (denote it by random variable b'). Say B runs A K times and receives K many possible values of b'. The density function of b' can be written as $f(b'|b, \mathbf{m}_b, \mathbf{c}_k)$ (i.e. it solely depends on the correct message and ciphertext pair). After looking at all outcomes, B can estimate the values of $\mathbf{m}_b, \mathbf{c}_k$ from the density function f using the *Maximum Likelihood Estimation (MLE)* [Ros18] procedure. Call the estimates $\mathbf{m}_{ml}, \mathbf{c}_{ml}$. There are several efficient algorithms to compute the MLE using exact or iterative techniques [Fle13, Osb92].

Now, using the *Consistency*[5] [C$^+$17] property of MLE, there exists a constant t, such that for all $K > t$, $\mathbf{m}_{ml} = \mathbf{m}_b, \mathbf{c}_{ml} = \mathbf{c}_k$.

It is important to point out why would the guess b' by A depend solely on the correct message and ciphertext pair. The following set of equations demonstrate why that is the case. Intuitively, all messages and thereby their ciphertext except $\mathbf{m}_b$ and $\mathbf{c}_k$ are independent of bit b.

$$\begin{aligned}
&\mathbf{Pr}\left[b' = b | M_0 = \mathbf{m}_0, \cdots, M_N = \mathbf{m}_N, C_0 = \mathbf{c}_0, \cdots, C_N = \mathbf{c}_N\right] \\
=&\frac{\mathbf{Pr}\left[\mathsf{DS} = \mathsf{D} | b' = b, M_c = \mathbf{m}_b, C_c = \mathbf{c}_k\right] \cdot \mathbf{Pr}\left[b' = b | M_c = \mathbf{m}_b, C_c = \mathbf{c}_k\right]}{\mathbf{Pr}\left[\mathsf{DS} = \mathsf{D} | M_c = \mathbf{m}_b, C_c = \mathbf{c}_k\right]} \quad (1) \\
=&\mathbf{Pr}\left[b' = b | M_c = \mathbf{m}_b, C_c = \mathbf{c}_k\right] \qquad (2)
\end{aligned}$$

as desired.

(1): Bayes Theorem for 3 events.

(2): $\mathbf{Pr}\left[\mathsf{DS} = \mathsf{D}\right]$ is independent of b,M_c and C_c. Thus, the unconditional and conditional probabilities are the same.

[5] Consistency of MLE: If θ_{ml} is the MLE of θ, then there exists n, such that for all $N \quad n$, $\mathbf{Pr}\left[j\theta_{ml} \quad \theta j \quad \varepsilon\right] = 0$, for any $\varepsilon \quad 0$.

Adversary B now guesses the value of b (call it b_{ml}) dependent solely on $\mathbf{m}_{ml}, \mathbf{c}_{ml}$. Thus, we have,

$$\mathbf{Adv}(\mathsf{B}) \geq \mathbf{Pr}\left[b_{ml} = b | \mathbf{m}_{ml} = \mathbf{m}_b, \mathbf{c}_{ml} = \mathbf{c}_k\right] \mathbf{Adv}^{\delta-\mathsf{RoR}}_{\mathcal{SE}}(\mathsf{A})$$

We need a bound on $\mathbf{Pr}\left[b_{ml} = b | \mathbf{m}_{ml} = \mathbf{m}_b, \mathbf{c}_{ml} = \mathbf{c}_k\right]$.

Define $\mathsf{win} = 1$ if B can distinguish between $\mathbf{m}_0, \mathbf{m}_1$, 0 otherwise. Thus the events $\mathsf{win} = 1$ and $b_{ml} = b$ are identical.

We now find $\mathbf{Pr}\left[\mathsf{win} = 1 | \mathsf{C} = \mathbf{c}_k, \mathsf{M} = \mathbf{m}_b\right]$.

First, we provide some definitions.

Definition 2. *For any message m, the "Ciphertext Region" is the convex hull of all the points which could be a possible ciphertext some choice of secret key and randomness of the encryption algorithm. For example, say in 2 dimensions, if $\overrightarrow{m} = (2, 2); s = 1, 2; \beta = 8$, the possible choices of $\overrightarrow{c}$ would be all points that lie in the circle of radius 2 around $(1, 1)$ and $(4, 4)$. All points enclosed within these two circles form the "Ciphertext Region".*

Definition 3. *Let C_1 and C_2 represent the "Ciphertext Regions" for messages m_1 and m_2 respectively. We define "Overlapping Region"(O) corresponding to these sets of messages as $C_1 \cap C_2$.*

Write

$$\mathbf{Pr}\left[\mathsf{win} = 1 | \mathsf{C} = \mathbf{c}_k, \mathsf{M} = \mathbf{m}_b\right]$$

$$= \mathbf{Pr}\left[\mathsf{win} = 1 | \mathsf{C} = \mathbf{c}_k, \mathsf{M} = \mathbf{m}_b, \mathbf{c}_k \in O\right] \mathbf{Pr}\left[c_k \in O\right]$$

$$+ \mathbf{Pr}\left[\mathsf{win} = 1 | \mathsf{C} = \mathbf{c}_k, \mathsf{M} = \mathbf{m}_b, \mathbf{c}_k \notin O\right] \mathbf{Pr}\left[c_k \notin O\right]$$

.

The above expression is upper bounded by,

$$(1 - p) + p \cdot \mathbf{Pr}\left[\mathsf{win} = 1 | \mathsf{C} = \mathbf{c}_k, \mathsf{M} = \mathbf{m}_b, \mathbf{c}_k \in O\right]$$

where $p = \mathbf{Pr}\left[c_k \in O\right]$

Claim. $\mathbf{Pr}\left[\mathsf{win} = 1 | \mathsf{C} = \mathbf{c}_k, \mathsf{M} = \mathbf{m}_b, \mathbf{c}_k \in O\right] = 0.5$

Proof. First, we define the following random variables:

- $\mathcal{S}$: A uniform random variable with support S (keyspace) that represents the secret key (scaling factor).
- Ctxt Represents the ciphertext returned after encryption.
- λ_m : Represents the perturbation factor used to compute the ciphertext.

$$\begin{aligned}
&\Pr[\mathsf{Ctxt} = c | \mathbf{m}, \mathbf{m}', c \in O] \\
&= \sum_s \Pr[\mathsf{Ctxt} = c | \mathbf{m}, \mathbf{m}', c \in O, \mathcal{S} = s] \Pr[\mathcal{S} = s] \\
&= \frac{1}{|\mathsf{S}|} \sum_s \Pr[\lambda_m = \mathbf{c} - s\mathbf{m}_b | \mathbf{m}, \mathbf{m}', c \in O, \mathcal{S} = s] \\
&= \frac{1}{|\mathsf{S}|} \sum_s \Pr[\lambda_m = \mathbf{c} - s\mathbf{m}_b] \\
&= \frac{1}{\mathsf{S}} \sum_s \frac{1}{f(\frac{s\beta}{4}, \delta)}.
\end{aligned}$$

Here, $f(\frac{s\beta}{4}, \delta)$ is a function that denotes the total number of points lying inside O which is independent of all queries. $f(\cdot)$ depends on δ and $\frac{s\beta}{4}$.

N	dim	δ	β	$Adv_{\mathcal{SE}}^{\delta-\mathsf{RoR-D}}(A)$
100	3	2^6	2^{15}	2^{-47}
300	5	2^5	2^{10}	2^{-96}
1000	5	2^6	2^{10}	2^{-116}
1000	10	2^9	2^{14}	2^{-217}
5000	5	2^8	2^{11}	2^{-154}
5000	8	2^8	2^{12}	2^{-160}

Table 1: Some concrete parameters and upper bounds on $Adv_{\mathcal{SE}}^{\delta-\mathsf{RoR-D}}(A)$.

The penultimate line in the equation above holds because λ is chosen uniformly from all possible points in O.

Thus the distribution of Ctxt is independent of the queries and oracle responses. This proves perfect secrecy of SAP constrained in the region O_i. ∎

Thus,

$$(1-p) + p\mathbf{Pr}\left[\mathsf{win}=1|\mathsf{C}=\mathbf{c}_k, \mathsf{M}=\mathbf{m}_b, \mathbf{c}_k \in O\right] = 1 - 0.5p$$

$\mathbf{Pr}\left[b_{ml}=b|\mathbf{m}_{ml}=\mathbf{m}_b, \mathbf{c}_{ml}=\mathbf{c}_k\right] \leq 1-0.5p.$

$$\begin{aligned}
\mathbf{Pr}\left[b_{ml}=b\right] &= \mathbf{Pr}\left[b_{ml}=b|b=b'\right]\mathbf{Pr}\left[b=b'\right] \\
&\quad + \mathbf{Pr}\left[b_{ml}=b|b\neq b'\right]\mathbf{Pr}\left[b\neq b'\right] \\
&\geq (1-0.5p)\mathbf{Pr}\left[b=b'\right] + 0.5\mathbf{Pr}\left[b\neq b'\right] \\
&= (0.5-0.5p)\mathbf{Pr}\left[b=b'\right] + 0.5 \\
\implies \mathbf{Adv}(\mathsf{B}) &= 2\mathbf{Pr}\left[b_{ml}=b\right] - 1 \\
&= (1-p)\mathbf{Pr}\left[b=b'\right] \\
&= (1-p)\frac{\mathbf{Adv}(A)+1}{2} \\
&\geq (1-p)\frac{\mathbf{Adv}(A)}{2}
\end{aligned}$$

Thus,

$$\mathbf{Adv}_{\mathcal{SE}}^{\delta-\mathsf{RoR}}(\mathsf{A}) \leq \frac{2}{1-p}\mathbf{Adv}(\mathsf{B}).$$

∎

It is now left to find the value of p and prove the claimed upper-bound for *canonical* adversaries. The analysis involves a geometric approach along with standard set theoretic and probabilistic arguments. Due to its cumbersome nature, we have shifted the remaining proof to Appendix B.

Note: Consistency of MLE holds under certain regularity conditions [NM94] are satisfied. In order to not distract from the proof, we assume that these conditions are satisfied. A careful analysis can be done to verify that this is indeed true.

Table 1 shows a few values for $\mathbf{Adv}_{\mathcal{SE}}^{\delta-\mathsf{RoR-D}}(\mathsf{A})$.

7 Security against Approximate Frequency-Finding Attacks

Here we define security against an adversary that tries to approximately guess any one element of the *histogram* corresponding to an attribute of the plaintext. We call it the Freq-Find (FF) notion.

We proceed to state some definitions which will be useful in the formal security analysis.

Attribute Histogram: For a list of attributes $\mathsf{L_{attr}} = (a_1, \cdots, a_n)$, let $(a'_1, \cdots, a'_{n'})$, $n' \leq n$ be the set of all unique elements in $\mathsf{L_{attr}}$. We define the *histogram* of $\mathsf{L_{attr}}$ as a list denoted by $\mathsf{Hist}(\mathsf{L_{attr}})$, with each element $\mathsf{Hist}(\mathsf{L_{attr}})[j], j \in [n']$ as,

$$\mathsf{Hist}(\mathsf{L_{attr}})[j] = \sum_{i=1}^{n} \mathbb{1}\{a_i = a'_j\}.$$

Most Likely Attribute Histogram: For a set of ciphertexts $\mathsf{L_{ctxt}} = (\mathbf{c}_1, \cdots, \mathbf{c}_n)$, the most likely attribute histogram for the j^{th} attribute, $\mathsf{j} \in [d]$ is a list denoted by $\mathsf{Hist}(\mathsf{L}^{\mathsf{j,ml}}_{\mathsf{ctxt}})$, where $\mathsf{L}^{\mathsf{j,ml}}_{\mathsf{ctxt}} = (m_1^{\mathsf{j,ml}}, \cdots, m_n^{\mathsf{j,ml}})$. $m_i^{\mathsf{j,ml}}$ is the j^{th} attribute of the most likely guess for message m_i corresponding to c_i, $\forall\ i \in [n]$.

γ-Approximate Histograms: A most likely histogram Hist is called a γ-Approximate Histogram for the actual histogram $\mathsf{Hist}(\mathsf{L})$ if $\forall i$,

$$\mathsf{Hist}[i] \in [\mathsf{Hist}(\mathsf{L})[i] - \gamma, \mathsf{Hist}(\mathsf{L})[i] + \gamma].$$

In our case, the goal of the adversary is to guess *an entry* of *γ-approximate histogram* of the plaintext histogram.

Up next, we take a detour to introduce an intermediate security notion which is pivotal in proving the security against FF attacks.

7.1 Window One-Wayness Security Notion

In this section, $\mathcal{MD}$ is $\mathcal{MVN}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$, such that, $\boldsymbol{\mu} = (\mu[1], \ldots, \mu[d])$ and $\boldsymbol{\Sigma}$ is the $d \times d$ covariance matrix. Naturally, $\mathcal{D}_i$ becomes $\mathcal{N}(\mu[i], \sigma_1^2)$, where σ_i^2 is the i^{th} diagonal entry of $\boldsymbol{\Sigma}$. For some theorems (specifically AWOW) which deal with attribute space, the subscript has been dropped for ease of reading.

We introduce an intermediate *Window One-Wayness* based security notion which was introduced by [BCO11]. It measures the probability that an adversary, given a set of ciphertexts corresponding to messages chosen at random from the underlying plaintext distribution decrypts one of them. The definition considers a general scenario that asks the adversary given some inputs to guess an interval (window) within which the underlying challenge plaintext lies. They do not need to point out which plaintext they intend to guess the window around. The size of the window and the number of challenge ciphertexts are parameters of the definition.

We analyze the security of individual attributes for each plaintext. This is much stronger that than looking at the security of a plaintext as a whole as window one-way security of each attribute implies one-wayness security for the whole point. The converse need not be true.

Attribute Window One-Wayness: Let $\mathcal{SE} = (\mathcal{K}eyGen, \mathcal{E}nc, \mathcal{D}ec)$ be a symmetric key encryption scheme. Let $\mathcal{MD}$ be a stateful "plaintext sampler" that on input $(\mathsf{state}, \mathsf{d}^*)$ (due to ease of notation, we drop state in some function definitions) outputs a plaintext $\mathbf{m}$ whose a^{th} attribute is denoted by $m[a]$ along with the updated state. Let $r \in \mathsf{N}$. In our case, $\mathcal{MD}$ will denote a multivariate distribution sample whose i^{th} attribute follows a univariate gaussian (denoted by $\mathcal{D}_i$).

We say that $\mathcal{SE}$ is $(r, \varepsilon) - \mathsf{AWOW}$-secure for $\mathcal{MD}$ if for every r-AWOW adversary A against $\mathcal{SE}$, *i.e.*, obeying the restrictions given below, its advantage

$$\mathbf{Adv}^{r-\mathsf{AWOW}}_{\mathcal{SE},\mathcal{MD}}(A) = \mathbf{Pr}\left[\mathbf{Exp}^{r-\mathsf{AWOW}}_{\mathcal{SE},\mathcal{MD}}(A) = 1\right] \leq \varepsilon,$$

where the experiment $\mathbf{Exp}^{r-\mathsf{AWOW}}_{\mathcal{SE},\mathcal{MD}}(A)$ is defined as follows:

Experiment $Exp^{r-\mathsf{AWOW}}_{\mathcal{SE},\mathcal{MD}}(A)$:
$K \xleftarrow{\$} \mathcal{K}eyGen$
$\mathsf{S}'_M \leftarrow \emptyset$
$(m_L, m_R) \xleftarrow{\$} A^{\mathsf{Msg}()}$
If $\exists \mathbf{m} \in \mathsf{S}'_M$ such that
for some $a \in [d]$, $m_a \in [m_L, m_R]$
return 1
Else return 0

Oracle $\mathsf{Msg}(d^*)$:
$(state, \mathbf{m}) \xleftarrow{\$} \mathcal{MD}(d^*)$
$\mathsf{S}'_M \leftarrow \mathsf{S}'_M \cup \{\mathbf{m}\}$
$\mathbf{c} \xleftarrow{\$} \mathcal{E}nc_K(\mathbf{m})$
return $\mathbf{c}$

Restrictions on the adversary: An $(r, \varepsilon) - \mathsf{AWOW}$ adversary A must obey the following rules:

– For any output (m_L, m_R), $|m_L - m_R| \leq r$.

We now define an alternate $r - \mathsf{AWOW}$ security experiment. Here, define that on input $(\mathsf{state}, \mathsf{d}^*, N)$, the stateful sampler $\mathcal{MD}$ outputs a dataset D of size N from this distribution.

The experiment $\mathbf{Exp}^{r-\mathsf{AWOW}-1}_{\mathcal{SE},\mathcal{MD}}(A)$ where the adversary has the same restrictions as above is defined as follows:

Experiment $Exp^{r-\mathsf{AWOW}-1}_{\mathcal{SE},\mathcal{MD}}(A)$:
$K \xleftarrow{\$} \mathcal{K}eyGen$
$\mathsf{S}'_M, \mathsf{C}_D \leftarrow \emptyset$
$(m_L, m_R) \xleftarrow{\$} A^{\mathsf{Msg}(\cdot)}$
If $\exists \mathbf{m} \in \mathsf{S}'_M$ such that
for some $a \in [d]$, $\mathbf{m}_a \in [m_L, m_R]$
return 1
Else return 0

Oracle $\mathsf{Msg}(d^*, N)$:
for $i \in [N]$
$(state, \mathbf{m}_i) \xleftarrow{\$} \mathcal{MD}(d^*)$
$\mathsf{S}'_M \leftarrow \mathsf{S}'_M \cup \{\mathbf{m}_i\}$
$\mathbf{c} \xleftarrow{\$} \mathcal{E}nc_K(\mathbf{m})$
$\mathsf{C}_D \leftarrow \mathsf{C}_D \cup \mathbf{c}$
return $\mathbf{shuffle}(\mathsf{C}_D)$

Lemma 1. *If in* $Exp^{r-\mathsf{AWOW}}_{\mathcal{SE},\mathcal{MD}}(A)$ *the adversary makes* N *queries to* Msg *and a single oracle query in* $Exp^{r-\mathsf{AWOW}-1}_{\mathcal{SE},\mathcal{MD}}(A)$*, then we have*

$$\mathbf{Pr}\left[\boldsymbol{Exp}^{r-\mathsf{AWOW}}_{\mathcal{SE},\mathcal{MD}}(A) = 1\right] = \mathbf{Pr}\left[\boldsymbol{Exp}^{r-\mathsf{AWOW}-1}_{\mathcal{SE},\mathcal{MD}}(A) = 1\right] .$$

Note that the Experiment $Exp^{r-\mathsf{AWOW}-1}_{\mathcal{SE},\mathcal{MD}}(A)$ captures the case where **shuffle** has been applied to the dataset whereas Experiment $Exp^{r-\mathsf{AWOW}}_{\mathcal{SE},\mathcal{MD}}(A)$ is the scenario where messages are encrypted on-the-fly without $\mathsf{shuffle}$.

Proof. Since each message belonging to the dataset are independently generated, the messages sampled by Msg oracles for both the experiments follow the same distribution. Moreover, all random permutations are identically distributed so the oracle's output for both experiments will also be identically distributed. Hence the lemma follows. ▌

Thus, we see that the shuffle does not have any influence on the r-AWOW security bounds. From now on, we use $Exp^{r-\mathsf{AWOW}-1}_{\mathcal{SE},\mathcal{MD}}(A)$ and $Exp^{r-\mathsf{AWOW}}_{\mathcal{SE},\mathcal{MD}}(A)$ interchangeably as per convenience.

We say that such $\mathcal{MD}$ is multivariate *Gaussian* if $\forall$ state, d^*, every i^{th} attribute follows a univariate distribution $\mathcal{D}_i$, where $\mathcal{D}_i$ is $\mathcal{N}(\mu_i, \sigma_i^2)$.

7.2 One-Wayness Bounds

Note that we shift the proofs in this subsection to the appendix to keep focus on the main point of the section, i.e. Frequency Finding Attacks.

We pay attention to the case when the adversary looks to decrypt the ciphertext to come with a correct guess for the "*most likely*" plaintext.

Most Likely Plaintext. Fix a symmetric encryption scheme $\mathcal{SE} = (\mathcal{K}eyGen, \mathcal{E}nc, \mathcal{D}ec)$. For given $\mathbf{c} \in \mathbf{C}$, if $\mathbf{m}_c \in \mathbf{M}$ is a message such that

$$\mathop{\mathbf{Pr}}_{K \xleftarrow{\$} \mathcal{K}eyGen}[\mathcal{E}nc(K, \mathbf{m}) = c] .$$

achieves a maximum at $\mathbf{m} = \mathbf{m}_c$, then we call $\mathbf{m}_c$ a (if unique, "the") *most likely* plaintext for $\mathbf{c}$.

An Upper Bound on the r-AWOW Advantage. The following theorem states an upper bound on any AWOW adversary against SAP.

Theorem 10. *For any* r-AWOW *adversary* A *making at most* z Msg *oracle queries*

$$\boldsymbol{Adv}^{r-\mathsf{AWOW}}_{\mathsf{SAP},\mathcal{MVN}(\boldsymbol{\mu},\boldsymbol{\Sigma})}(A) \leq z(1-p)^{|\mathsf{M}|}\frac{|\mathsf{S}|+1}{2|\mathsf{S}|}\sum_m \left(\left(\frac{2rm-\beta m}{m^2-r^2} - \frac{\beta}{m}\right)\binom{|\mathsf{M}|}{m}\right) . \quad (3)$$

where $\mathsf{M} = [-M, M]$ *is the attribute-space of* SAP *and* S *is the keyspace. Here,* p *is the parameter of the Binomial Distribution (denoted by* $Bin(|\mathsf{M}|, p)$*) which is used to approximate* $\mathcal{N}(\mu, \sigma^2)$*, which is the univariate distribution of the chosen attribute. Hence,* $|\mathsf{M}|p = \mu, |\mathsf{M}|p(1-p) = \sigma^2$*. For simplicity, we use* m *to denote an attribute instead of* $m[a]$ *which is an abuse of notation.*

The proof for this theorem is obtained using straightforward algebraic manipulation and probabilistic arguments. Hence, it has been postponed to in Appendix C.1. To help understand the bounds, we present some values in Table 2. In the table, message space $\mathcal{M} = [-2^{80}, 2^{80}]^d$ and $|\mathcal{S}| = 2^{30}$.

d	r	β	r, 1-AWOW Upper Bound	r, 1-AWOW Lower Bound
1	$\frac{\beta}{2^4}$	2^{15}	2^{-32}	2^{-35}
2	$\frac{\beta}{2^5}$	2^{20}	2^{-40}	2^{-43}
4	$\frac{\beta}{2^5}$	2^{20}	2^{-51}	2^{-55}
6	$\frac{\beta}{2^8}$	2^{25}	2^{-62}	2^{-66}
10	$\frac{\beta}{2^{10}}$	2^{30}	2^{-76}	2^{-78}

Table 2: Upper and Lower Bounds on r, 1-AWOW Advantage

A Lower Bound on Large Attribute Window One-Wayness. Here we show that there exists an efficient adversary attacking the window one-wayness of SAP for a sufficiently large window size.

Theorem 11. *For any* r-AWOW *adversary* A*, with* $r \geq \frac{\beta}{2}$*,*

$$\boldsymbol{Adv}^{r-\mathsf{AWOW}}_{\mathsf{SAP},\mathcal{MVN}(\boldsymbol{\mu},\boldsymbol{\Sigma})}(A) \geq \frac{r}{2\sqrt{|\mathsf{M}|}} \ln \frac{1-p}{1-2p} \,. \tag{4}$$

where $\mathsf{M} = [-M, M]$ *is the attribute-space of* SAP *and* S *is the keyspace. Here,* p *is the parameter of the Binomial Distribution (denoted by* $Bin(|\mathsf{M}|, p)$*) which is used to approximate* $\mathcal{N}(\mu, \sigma^2)$*, the univariate normal distribution corresponding to the chosen attribute. Hence,* $|\mathsf{M}|p = \mu, |\mathsf{M}|p(1-p) = \sigma^2$*.*

The bound has been proved in Appendix C.2 due to reasons mentioned in case of the previous theorem.

7.3 Security against Freq-Find adversaries

The adversary wins the game if it can guess *an entry* HistEntry of a γ-approximate histogram, which occurs at most ψ times. We say that $\mathcal{SE}$ is $(\gamma, \psi, \varepsilon) - \mathsf{FF}$ secure for $\mathcal{D}$ if the $(\gamma, \psi, \varepsilon) - \mathsf{FF}$ advantage of an adversary A against $\mathcal{SE}$ is,

$$\mathbf{Adv}^{(\gamma,\psi)-\mathsf{FF}}_{\mathcal{SE},\mathcal{MD}}(A) = \mathbf{Pr}\left[\mathrm{Exp}^{(\gamma,\psi)-\mathsf{FF}}_{\mathcal{SE},,\mathcal{MD}}(A) = 1\right] \leq \varepsilon,$$

where the experiment $Exp^{(\gamma,\psi)-\mathsf{FF}}_{\mathcal{SE},,\mathcal{MD}}(A)$ is defined as:

Experiment $Exp^{(\gamma,\psi)-\mathsf{FF}}_{\mathcal{SE}}$(A):
$K \xleftarrow{\$} \mathcal{K}eyGen$
$h \leftarrow A^{\mathsf{Msg}(\cdot)}$ // a guess for any element of the approx. histogram
$count \leftarrow 0$
for i in 1 to n
  for j in 1 to d
    If $h \in [\mathsf{Hist}(\mathsf{S}'^j_M)[i] - \gamma, \mathsf{Hist}(\mathsf{S}'^j_M)[i] + \gamma]$ // (1)
      $count \leftarrow count + 1$ // (2)
If $0 < count \leq \psi$ // (3)
  Return 1
Else return 0

Above, the Msg oracle is exactly the same as in the r-AWOW experiment.

1. Check if the guess within γ approx. of an histogram entry. $\mathsf{Hist}(\mathsf{S}'^{j}_{M})$: histogram for the list of j^{th} attribute of elements in the list S'_M.
2. Track the number of times the guessed frequency occurs.
3. Make sure that this guessed frequency occurs at most λ times.

Parameter ψ is an essential parameter which depends on the underlying plaintext distribution. It prevents the adversary from trivially winning the game by guessing a frequency value which has a very high number of occurrence in the histogram. For example, a dataset from a well spread distribution will have plenty of points sampled only once. In that case, the adversary can win the game easily by guessing 1.

Upper Bound on the Freq-Find Advantage:

Theorem 12. *(Main result.) Let $\mathcal{SE} = (\mathcal{K}eyGen, \mathcal{E}nc, \mathcal{D}ec)$ be a normed vector-space encryption scheme. Let A be a (γ,ψ)-FF adversary. Then there exists a $\frac{\gamma}{\psi}$-AWOW adversary B making at most q_m queries to the Msg oracle such that and $\mathcal{D} \sim \mathcal{N}(\mu, \sigma^2)$ is the univariate distribution for the chosen attribute.*

$$\boldsymbol{Adv}^{(\gamma,\psi)-\mathsf{FF}}_{\mathcal{SE},\mathcal{MD}}(A) \leq \frac{1}{0.5(\frac{0.39\gamma}{\psi\sigma} - 2e^{\frac{-0.5}{q_m}})} \boldsymbol{Adv}^{\frac{\gamma}{\psi}-\mathsf{AWOW}}_{\mathcal{SE}}(B).$$

Optimal attribute interval. Let $\mathsf{Hist}(\mathsf{S}'^{j}_{M})$ be the histogram for the list of j^{th} attributes of elements in the list S'_M. Let X be the random variable used to denote an attribute following distribution $\mathcal{D}$ ($\mathcal{D}$ is a Normal Distribution in our case) over the attribute space. Let $\mathsf{HistEntry}$ be any arbitrary guess by the $\mathsf{Freq-Find}$ adversary.

The optimal attribute interval for such guess for a γ-approximate histogram *entry*, is denoted by $\mathsf{OptInt}(\mathsf{HistEntry}, \mathsf{S}'_M) = [m_{L_{opt}}, m_{R_{opt}}]$ such that $m_{R_{opt}} - m_{L_{opt}} = r$ and $\left| \underset{X \xleftarrow{\$} \mathcal{D}}{\mathbf{Pr}} [l_i \leq X \leq m_i] - \frac{\mathsf{HistEntry}}{|\mathsf{S}'_M|} \right|$ $= \min_{j:m_j - l_j = r} \left| \underset{X \xleftarrow{\$} \mathcal{D}}{\mathbf{Pr}} [l_j \leq X \leq m_j] - \frac{\mathsf{HistEntry}}{|\mathsf{S}'_M|} \right|$. where r is the attribute window length defined in the AWOW experiment.

Intuitively, the above equation selects the interval of length r among all possible intervals whose sampling probability is closest to the guess of the FF adversary.

Lemma 2. *Assume that $X_1, X_2, \ldots$ are independent and identically distributed random variables in R with cumulative distribution function $F(x)$. The empirical distribution function for $X_1, \ldots, X_n$ is defined by*

$$F_n(x) = \frac{1}{n} \sum_{i=1}^{n} \mathbb{1}\{X_i \leq x\}.$$

Glivenko Cantelli Theorem (GCT) [Tuc59] states that:

$$\sup_{x \in \mathsf{R}} |F_n(x) - F(x)| \xrightarrow{n \to \infty} 0.$$

GCT can be strengthened using the Dvoretzky–Kiefer–Wolfowitz (DKW) inequality which quantifies the rate of convergence.

Lemma 3 (Dvoretzky–Kiefer–Wolfowitz (DKW) inequality [DKW56]).

$$\Pr[\sup_{x \in \mathsf{R}} |F_n(x) - F(x)| > z] \leq e^{-2nz^2}.$$

Algorithm 4 $\frac{\gamma}{\psi}$-AWOW Adversary

procedure $B^{\mathsf{Msg}}(\cdot)$
  Run A
    On Message oracle query
      $(state, \mathbf{x}) \xleftarrow{\$} \mathcal{MD}(d^*)$
      $\mathsf{Sim} \leftarrow \mathsf{Sim} \cup \{\mathbf{x}\}$
      Return Msg
  Until A outputs $\mathsf{HistEntry}^*$
  $count \leftarrow 0$
  If $\exists i, j$ $\mathsf{HistEntry}^* \in [\mathsf{Hist}(\mathsf{Sim}^j)[i] - \gamma, \mathsf{Hist}(\mathsf{Sim}^j)[i] + \gamma]$
    $count \leftarrow count + 1$
  If $0 < count \leq \psi$
    Return $\mathsf{OptInt}(\mathsf{HistEntry}^*, \mathsf{Sim})$ // As calculated in 7.3
  Else Return $\perp$

Special Case of DKW. We present a particular instance of DKW which is significant for our analysis. Let $z = \frac{1}{n}$ (This is the case where both the ecdfs converge to the same place.) and ecdfs F_n and G_n. DKW gives,

$$\Pr[\sup_{x \in \mathsf{R}} |F_n(x) - G_n(x)| > \frac{1}{n}] \leq \Pr[\sup_{x \in \mathsf{R}} |F_n(x) - F(x)| > \frac{0.5}{n}] + \Pr[\sup_{x \in \mathsf{R}} |G_n(x) - F(x)| > \frac{0.5}{n}] \leq 2e^{\frac{-0.5}{n}} .$$

We now define an AWOW adversary B that simulates adversary A.

Notation used in the proof below. Let $\mathbf{m}[a]$ denote any attribute of any arbitrary message $\mathbf{m}$ sampled from $\mathcal{MD}$. S'_M and Sim define the two ecdfs F_n and G_n in DKW. The histograms defined by the two sets are same if $z < 1$ in DKW. Standard Normal CDF and PDF are denoted by Φ and ϕ respectively.

The following reduction directly gives the upper bound for the (γ, ψ)-FF advantage of any (γ, ψ)-FF adversary A.

$$\begin{aligned}
&\mathbf{Adv}^{\frac{\gamma}{\psi}-\mathsf{AWOW}}_{\mathcal{SE},\mathcal{D}}(B) \\
\geq& 0.5\mathbf{Pr}\left[\mathrm{Exp}^{(\gamma,\psi)-\mathsf{FF}}_{\mathcal{SE},\mathcal{D}}(A) = 1\right] \cdot \left(\mathbf{Pr}\left[\exists\, \mathbf{m}, a,\ \mathbf{m}[a] \in \mathsf{OptInt}(\mathsf{HistEntry}^*, \mathsf{S}'_M)\right]\right) && (1) \\
\geq& 0.5\mathbf{Pr}\left[\mathrm{Exp}^{(\gamma,\psi)-\mathsf{FF}}_{\mathcal{SE},\mathcal{D}}(A) = 1\right] \cdot \left(\mathbf{Pr}\left[\exists\, \mathbf{m}, a,\ \mathbf{m}[a] \in \mathsf{OptInt}(\mathsf{HistEntry}^*, \mathsf{Sim})\right] - 2e^{\frac{-0.5}{qm}}\right) && (2) \\
=& 0.5\left(1 - [1 - \mathop{\mathbf{Pr}}_{X \xleftarrow{\$} \mathcal{D}}[X \in \mathsf{OptInt}(\mathsf{HistEntry}^*, \mathsf{Sim})]]^{qmd} - 2e^{\frac{-0.5}{qm}}\right) \cdot \mathbf{Pr}\left[\mathrm{Exp}^{(\gamma,\psi)-\mathsf{FF}}_{\mathcal{SE},\mathcal{D}}(A) = 1\right] \\
\geq& 0.5\left(\mathop{\mathbf{Pr}}_{X \xleftarrow{\$} \mathcal{D}}[X \in \mathsf{OptInt}(\mathsf{HistEntry}^*, \mathsf{Sim})] - 2e^{\frac{-0.5}{qm}}\right) \mathbf{Adv}^{(\gamma,\psi)-\mathsf{FF}}_{\mathcal{SE},\mathcal{D}}(A) \\
=& 0.5\left(\mathop{\mathbf{Pr}}_{X \xleftarrow{\$} \mathcal{D}}\left[X \in (m_{L_{opt}}, m_{R_{opt}})\right]\right] - 2e^{\frac{-0.5}{qm}}\right) \mathbf{Adv}^{(\gamma,\psi)-\mathsf{FF}}_{\mathcal{SE},\mathcal{D}}(A) \\
=& 0.5\left(\mathop{\mathbf{Pr}}_{X \xleftarrow{\$} \mathcal{N}(0,1)}\left[X \in \left(\frac{m_{L_{opt}} - \mu}{\sigma}, \frac{m_{R_{opt}} - \mu}{\sigma}\right)\right] - 2e^{\frac{-0.5}{qm}}\right) \mathbf{Adv}^{(\gamma,\psi)-\mathsf{FF}}_{\mathcal{SE},\mathcal{D}}(A) \\
=& 0.5\left(\Phi\left(\frac{m_{R_{opt}} - \mu}{\sigma}\right) - \Phi\left(\frac{m_{L_{opt}} - \mu}{\sigma}\right) - 2e^{\frac{-0.5}{qm}}\right) \mathbf{Adv}^{(\gamma,\psi)-\mathsf{FF}}_{\mathcal{SE},\mathcal{D}}(A) \\
=& 0.5\left(\frac{m_{R_{opt}} - \mu}{\sigma} \cdot \phi\left(\frac{m_{R_{opt}} - \mu}{\sigma}\right) - \frac{m_{L_{opt}} - \mu}{\sigma} \cdot \phi\left(\frac{m_{L_{opt}} - \mu}{\sigma}\right) - 2e^{\frac{-0.5}{qm}}\right) \mathbf{Adv}^{(\gamma,\psi)-\mathsf{FF}}_{\mathcal{SE},\mathcal{D}}(A) && (3) \\
\geq& 0.5\left(\frac{\gamma}{\psi\sigma} \cdot \phi\left(\frac{m_{R_{opt}} - \mu}{\sigma}\right) - 2e^{\frac{-0.5}{qm}}\right) \mathbf{Adv}^{(\gamma,\psi)-\mathsf{FF}}_{\mathcal{SE},\mathcal{D}}(A) .
\end{aligned}$$

Now, this holds for any arbitrary $m_{R_{opt}}$, thus we have:

$$\mathbf{Adv}^{\frac{\gamma}{\psi}-\mathsf{AWOW}}_{\mathcal{SE},\mathcal{D}}(B) \geq 0.5\left(\frac{0.39\gamma}{\psi\sigma} - 2e^{\frac{-0.5}{q_m}}\right)\mathbf{Adv}^{(\gamma,\psi)-\mathsf{FF}}_{\mathcal{SE},\mathcal{D}}(A)$$

because the standard normal pdf is upper bounded by 0.39.

1. Due to symmetry, we get the correct interval up to reflection, thus with probability at least 0.5, the correct interval is chosen.
2. Using DKW.
3. Taylor Expansion [M$^+$04].

To better understand the bound, we demonstrate a graph in Figure 1 to show how the leading multiplicative constant decays, thus giving a tight bound.

Some Practical Parameters: We present a graph (cf. Figure 1) to demonstrate the trend of $\frac{1}{(\frac{0.39\gamma}{\psi\sigma}-2e^{\frac{-0.5}{q_m}})}$, the multiplicative constant theorem 12 with respect to the parameter. (It is log scaled for better visuals.) This is necessary as the expression is difficult to analyse and very high values of this constant would make our reduction meaningless.

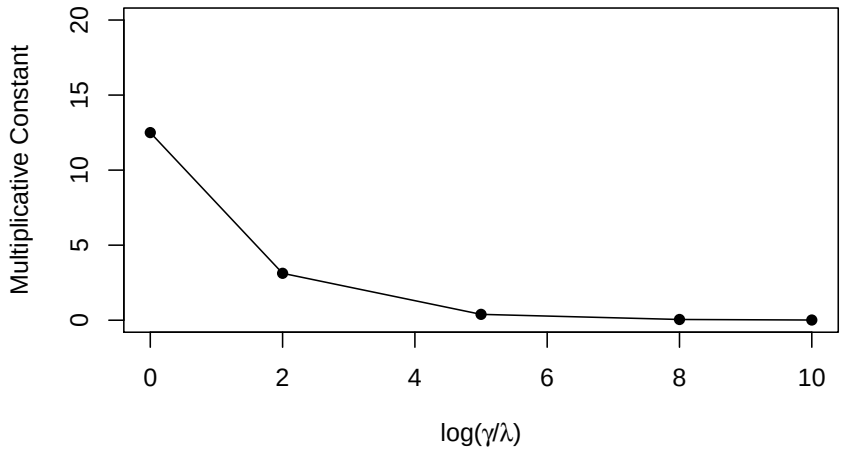

Fig. 1: $\frac{1}{(\frac{0.39\gamma}{\psi\sigma}-2e^{\frac{-0.5}{q_m}})}$ vs $\log_2(\frac{\gamma}{\psi})$, when $\sigma = 1,\ q_m = 2^5$

8 Bit Security

In this section, we are concerned with characterizing *bit security* of the plaintexts. We define a *Hardcore Bits* experiment which enables us to comment on the *one-wayness* and *pseudorandomness* of different bits from the same message. The hardcore bits notion actually considers a specific hardcore function and differs from the classical such notion in that the adversary may request multiple challenges on related messages. We give a reduction from this notion to δ-RoR.

Theorem 13. *Let $\mathcal{SE}$ be a symmetric key encryption scheme as defined previously. For any adversary B, there exists an δ-RoR adversary A such that $\boldsymbol{Adv}^{\mathsf{HCB}}_{\mathcal{SE},\mathcal{MD}}(A) = \mathbf{Pr}\left[\boldsymbol{Exp}^{\mathsf{HCB}}_{\mathcal{SE},\mathcal{MD}}(A) = 1\right] \leq \varepsilon$, and*

$$\boldsymbol{Adv}^{\mathsf{HCB}}_{\mathcal{SE},\mathcal{MD}}(B) \leq \boldsymbol{Adv}^{\delta-\mathsf{RoR}}_{\mathcal{SE},\mathcal{MD}}(A). \tag{5}$$

Experiment HCB:
$K \xleftarrow{\$} \mathcal{K}eyGen$
$\mathrm{S}'_M \leftarrow \emptyset$
$b \xleftarrow{\$} \{0,1\}$
$b' \xleftarrow{\$} A^{Msg()}()$
If $b' == b$ return 1
Else return 0

Oracle Msg(d^*, N):
for $i \in [N]$
$(state, \mathbf{m}_i) \xleftarrow{\$} \mathcal{MD}(d^*)$
$S'_M \leftarrow S'_M \cup \{\mathbf{m}_i\}$
$\mathbf{m}_i \xleftarrow{\$} \mathcal{E}nc_K(\mathbf{m}_i)$
If $b = 0$
bits $\leftarrow$ right most $\log_2 \delta$ bits of m_i
Else
bits $\xleftarrow{\$} \{0,1\}^{\log_2 \delta}$
$C \leftarrow C \cup \mathbf{c}_i \| \mathsf{bits}$
return C

Let A be an adversary taking part in δ-RoR-Experiment. We reduce adversary B to A. Algorithm 5 gives a perfect simulation by A using its oracle for the queries made by B. Thus,

$$\mathbf{Adv}^{\mathsf{HCB}}_{\mathcal{SE}}(B) \leq \mathbf{Adv}^{\mathsf{RoR}}_{\mathcal{SE}}(A).$$

Hence, theorem 13 follows.

Algorithm 5 RoR Adversary A

procedure A^{Msg}
Run B
On Msg oracle query (d^*, N)
$S_D \leftarrow S_D \cup \{d^*\}$
for $i \in [N]$
$(state, \mathbf{m}_i) \xleftarrow{\$} \mathcal{MD}(\cdot, d^*)$
$S'_M \leftarrow S'_M \cup \{\mathbf{m}_i\}$
$\mathsf{bits}_i \leftarrow \mathbf{m}_i[\log_2 \delta \cdots]$ // Right most $\log_2 \delta$ bits
Bits $\leftarrow$ **Bits** $\cup$ bits_i
$C \xleftarrow{\$} \mathsf{Swap}_b(S'_M)$
Return C & Bits to B
Repeat until B outputs guess bit b'
Return guess bit b'

Note that the reduction holds because $\|m - m'\| \leq \delta$. It is clear that A succeeds in breaking the δ-RoR experiment if B breaks the HCB experiment. The ciphertexts generated for both the messages are identically distributed because the lower order bits are masked using a uniformly distributed noise. ∎

Note: It must be pointed out that the HCB experiment has been carefully crafted to ensure that one can comment on the one-wayness and pseudorandomness of bits on the *same* message. To achieve this, the standard Msg oracle has been modified to append the rightmost $\log \delta$ bits of a message (or random $\log \delta$ bits) along with the ciphertext for the particular message. The one-wayness of the bits follows directly from the AWOW results. Experiment HCB allows us to create a hybrid that can process indistinguishability on top of one-wayness.

Claim. For any message $\mathbf{m}$ following $\mathcal{MVN}((\mu, \Sigma))$ whose components are n bit long encrypted by SAP,

1. The lowest $\log_2 \delta$ bits are pseudorandom (*i.e.*, hardcore).
2. The number of left most bits leaked (*i.e.*, efficiently computed) is strictly less than $\frac{\log_2 |\mathsf{M}|}{2}$ (half the higher order bits).
3. If we remove the left-most k bits from the lowest $\frac{\log_2 |\mathsf{M}|}{2}$ bits, the advantage of guessing the remaining lower order bits decreases by a multiplicative factor of 2^k.

Proof. (1) Using the analysis done in [TYM14], we observe that the RoR security notion makes the lower $\log_2 \delta$ bits indistinguishable from random. Let $L = \log_2 \delta, B = \mathcal{B}(; \log_2 \delta)$ (any arbitrary ball of radius δ) and take any

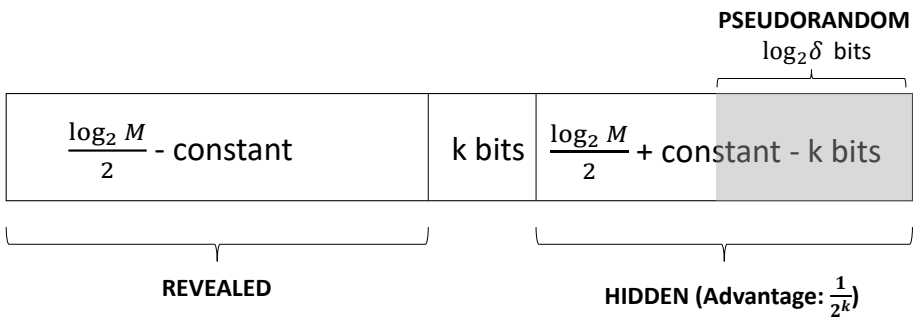

Fig. 2: Demonstrating the bit security.

interval I in such a way that for any two elements of B, all of their bits except the least significant L bits are the same for each attribute. I can thus be written as $I = \{2^L u + x | x \in [2^L - 1]\}$ for some u. By definition the length of the interval I is not more than δ.

Our notion ensures that any element m_0 of B is indistinguishable from that of a uniformly random element m_1 of B because $\|m_0 - m_1\| \leq \delta$. Since the least significant $[2^L - 1]$ bits of each attribute are distributed uniformly at random (This is due to the uniform noise added in our protocol), the right most $\log_2 \delta$ bits being indistinguishable from random follows.

(2) We analyze the number of upper order plaintext bits that an adversary can efficiently compute when $r \geq \frac{\beta}{2}$. Equation 4 is $\mathbf{Adv}_{\mathsf{SAP}}^{r-\mathsf{AWOW}}(A) = c\left(\frac{r}{\sqrt{|M|}}\right) = 1$, when $r \geq \frac{\sqrt{M}}{c}$, where $c = \ln \frac{1-p}{1-2p}$. Note that the result goes through if $p \leq \frac{1}{2}$. Thus, the number of bits leaked is $\log_2 |\mathsf{M}| - \log_2 r = \frac{\log_2 |\mathsf{M}|}{2} - \log_2\left(\frac{1}{c}\right)$.

(3) Start from the window size r above for which the adversarial advantage as 1 and reduce it by k bits, i.e. the new window length is $r' = r - 2^k$, then $\mathbf{Adv}_{\mathsf{SAP}}^{r-\mathsf{AWOW}}(A)$ in equation 4 which is linear in r reduces by a multiplicative factor of 2^k. Therefore, the probability of guessing the lowest $\log_2\{\frac{M}{2}\} - k$ bits is 2^{-k}. ▮

This shows that SAP scheme leaks strictly less than half of the total bits and the number of total bits leaked is a decreasing function of the approximation factor β. More precisely, increasing the approximation factor by k times decreases the number of bits leaked by $\log_2 k$. This is an improvement over at least half bits leaked by its predecessor OPE [BCO11].

References

ABO07. Georgios Amanatidis, Alexandra Boldyreva, and Adam O'Neill. Provably-secure schemes for basic query support in outsourced databases. In *IFIP Annual Conference on Data and Applications Security and Privacy*, pages 14–30. Springer, 2007.

AKSX04. Rakesh Agrawal, Jerry Kiernan, Ramakrishnan Srikant, and Yirong Xu. Order preserving encryption for numeric data. In *SIGMOD*, 2004.

AMN+98. Sunil Arya, David M Mount, Nathan S Netanyahu, Ruth Silverman, and Angela Y Wu. An optimal algorithm for approximate nearest neighbor searching fixed dimensions. *Journal of the ACM (JACM)*, 45(6):891–923, 1998.

BBGN19. Borja Balle, James Bell, Adrià Gascón, and Kobbi Nissim. The privacy blanket of the shuffle model. In Alexandra Boldyreva and Daniele Micciancio, editors, *CRYPTO 2019, Part II*, volume 11693 of *LNCS*, pages 638–667. Springer, Heidelberg, August 2019.

BBO07. Mihir Bellare, Alexandra Boldyreva, and Adam O'Neill. Deterministic and efficiently searchable encryption. In Alfred Menezes, editor, *CRYPTO 2007*, volume 4622 of *LNCS*, pages 535–552. Springer, Heidelberg, August 2007.

BCLO09. Alexandra Boldyreva, Nathan Chenette, Younho Lee, and Adam O'Neill. Order-preserving symmetric encryption. In Antoine Joux, editor, *EUROCRYPT 2009*, volume 5479 of *LNCS*, pages 224–241. Springer, Heidelberg, April 2009.

BCO11. Alexandra Boldyreva, Nathan Chenette, and Adam O'Neill. Order-preserving encryption revisited: Improved security analysis and alternative solutions. In Phillip Rogaway, editor, *CRYPTO 2011*, volume 6841 of *LNCS*, pages 578–595. Springer, Heidelberg, August 2011.

BEM+17. Andrea Bittau, Úlfar Erlingsson, Petros Maniatis, Ilya Mironov, Ananth Raghunathan, David Lie, Mitch Rudominer, Ushasree Kode, Julien Tinnes, and Bernhard Seefeld. Prochlo: Strong privacy for analytics in the crowd. In *Proceedings of the 26th Symposium on Operating Systems Principles*, pages 441–459, 2017.

BGC+18. Vincent Bindschaedler, Paul Grubbs, David Cash, Thomas Ristenpart, and Vitaly Shmatikov. The tao of inference in privacy-protected databases. *Proc. VLDB Endow.*, 11(11):1715–1728, 2018.

BGRS99. Kevin Beyer, Jonathan Goldstein, Raghu Ramakrishnan, and Uri Shaft. When is "nearest neighbor" meaningful? In *International conference on database theory*, pages 217–235. Springer, 1999.

BLR+14. Dan Boneh, Kevin Lewi, Mariana Raykova, Amit Sahai, Mark Zhandry, and Joe Zimmerman. Semantically secure order-revealing encryption: Multi-input functional encryption without obfuscation. Cryptology ePrint Archive, Report 2014/834, 2014. `http://eprint.iacr.org/2014/834`.

BLR+15. Dan Boneh, Kevin Lewi, Mariana Raykova, Amit Sahai, Mark Zhandry, and Joe Zimmerman. Semantically secure order-revealing encryption: Multi-input functional encryption without obfuscation. In Elisabeth Oswald and Marc Fischlin, editors, *EUROCRYPT 2015, Part II*, volume 9057 of *LNCS*, pages 563–594. Springer, Heidelberg, April 2015.

BQ53. FS Beckman and DA Quarles. On isometries of euclidean spaces. *Proceedings of the American Mathematical Society*, 4(5):810–815, 1953.

C+17. Jiahua Chen et al. Consistency of the mle under mixture models. *Statistical Science*, 32(1):47–63, 2017.

CD07. Padraig Cunningham and Sarah Jane Delany. k-nearest neighbour classifiers. *Multiple Classifier Systems*, 34(8):1–17, 2007.

CD15. Sanjit Chatterjee and M. Prem Laxman Das. Property preserving symmetric encryption revisited. In Tetsu Iwata and Jung Hee Cheon, editors, *ASIACRYPT 2015, Part II*, volume 9453 of *LNCS*, pages 658–682. Springer, Heidelberg, November / December 2015.

CK10. Melissa Chase and Seny Kamara. Structured encryption and controlled disclosure. In Masayuki Abe, editor, *ASIACRYPT 2010*, volume 6477 of *LNCS*, pages 577–594. Springer, Heidelberg, December 2010.

CLO+18. David Cash, Feng-Hao Liu, Adam O'Neill, Mark Zhandry, and Cong Zhang. Parameter-hiding order revealing encryption. In Thomas Peyrin and Steven Galbraith, editors, *ASIACRYPT 2018, Part I*, volume 11272 of *LNCS*, pages 181–210. Springer, Heidelberg, December 2018.

CLWW16. Nathan Chenette, Kevin Lewi, Stephen A. Weis, and David J. Wu. Practical order-revealing encryption with limited leakage. In Thomas Peyrin, editor, *FSE 2016*, volume 9783 of *LNCS*, pages 474–493. Springer, Heidelberg, March 2016.

CSU+19. Albert Cheu, Adam D. Smith, Jonathan Ullman, David Zeber, and Maxim Zhilyaev. Distributed differential privacy via shuffling. In Yuval Ishai and Vincent Rijmen, editors, *EUROCRYPT 2019, Part I*, volume 11476 of *LNCS*, pages 375–403. Springer, Heidelberg, May 2019.

DDC16. F. Betül Durak, Thomas M. DuBuisson, and David Cash. What else is revealed by order-revealing encryption? In Edgar R. Weippl, Stefan Katzenbeisser, Christopher Kruegel, Andrew C. Myers, and Shai Halevi, editors, *ACM CCS 2016*, pages 1155–1166. ACM Press, October 2016.

DKW56. Aryeh Dvoretzky, Jack Kiefer, and Jacob Wolfowitz. Asymptotic minimax character of the sample distribution function and of the classical multinomial estimator. *The Annals of Mathematical Statistics*, pages 642–669, 1956.

Dwo08. Cynthia Dwork. Differential privacy: A survey of results. In *International conference on theory and applications of models of computation*, pages 1–19. Springer, 2008.

EFM+19. Úlfar Erlingsson, Vitaly Feldman, Ilya Mironov, Ananth Raghunathan, Kunal Talwar, and Abhradeep Thakurta. Amplification by shuffling: From local to central differential privacy via anonymity. In Timothy M. Chan, editor, *30th SODA*, pages 2468–2479. ACM-SIAM, January 2019.

Fle13. Roger Fletcher. *Practical methods of optimization*. John Wiley & Sons, 2013.

GC18. Riddhi Ghosal and Sanjit Chatterjee. Privacy preserving multi-server k-means computation over horizontally partitioned data. In *International Conference on Information Systems Security*, pages 189–208. Springer, 2018.

GLMP18. Paul Grubbs, Marie-Sarah Lacharité, Brice Minaud, and Kenneth G. Paterson. Pump up the volume: Practical database reconstruction from volume leakage on range queries. In David Lie, Mohammad Mannan, Michael Backes, and XiaoFeng Wang, editors, *ACM CCS 2018*, pages 315–331. ACM Press, October 2018.

GLMP19. Paul Grubbs, Marie-Sarah Lacharité, Brice Minaud, and Kenneth G. Paterson. Learning to reconstruct: Statistical learning theory and encrypted database attacks. In *2019 IEEE Symposium on Security and Privacy*, pages 1067–1083. IEEE Computer Society Press, May 2019.

GRS17. Paul Grubbs, Thomas Ristenpart, and Vitaly Shmatikov. Why your encrypted database is not secure. In Alexandra Fedorova, Andrew Warfield, Ivan Beschastnikh, and Rachit Agarwal, editors, *HotOS 2017*, pages 162–168. ACM, 2017.

GSB+17. Paul Grubbs, Kevin Sekniqi, Vincent Bindschaedler, Muhammad Naveed, and Thomas Ristenpart. Leakage-abuse attacks against order-revealing encryption. In *2017 IEEE Symposium on Security and Privacy*, pages 655–672. IEEE Computer Society Press, May 2017.

HILM02. Hakan Hacigümüş, Bala Iyer, Chen Li, and Sharad Mehrotra. Executing sql over encrypted data in the database-service-provider model. In *SIGMOD*, 2002.

HJL+17. Helene Haagh, Yue Ji, Chenxing Li, Claudio Orlandi, and Yifan Song. Revealing encryption for partial ordering. In Máire O'Neill, editor, *16th IMA International Conference on Cryptography and Coding*, volume 10655 of *LNCS*, pages 3–22. Springer, Heidelberg, December 2017.

HL10. Radoslav Harman and Vladimír Lacko. On decompositional algorithms for uniform sampling from n-spheres and n-balls. *Journal of Multivariate Analysis*, 101(10):2297–2304, 2010.

IM98. Piotr Indyk and Rajeev Motwani. Approximate nearest neighbors: towards removing the curse of dimensionality. In *Proceedings of the thirtieth annual ACM symposium on Theory of computing*, pages 604–613. ACM, 1998.

JP18. Marc Joye and Alain Passelègue. Function-revealing encryption - definitions and constructions. In Dario Catalano and Roberto De Prisco, editors, *SCN 18*, volume 11035 of *LNCS*, pages 527–543. Springer, Heidelberg, September 2018.

JPW06. Geetha Jagannathan, Krishnan Pillaipakkamnatt, and Rebecca N Wright. A new privacy-preserving distributed k-clustering algorithm. In *2006 SIAM International Conference on Data Mining 2006*, pages 494–498. SIAM, 2006.

Kam15. Seny Kamara. How to search on encrypted data, 2015. https://cs.brown.edu/ seny/slides/encryptedsearch-full.pdf.

KDWS05. Hillol Kargupta, Souptik Datta, Qi Wang, and Krishnamoorthy Sivakumar. Random-data perturbation techniques and privacy-preserving data mining. *Knowledge and Information Systems*, 7(4):387–414, 2005.

Kea89. Simon K Kearsley. On the orthogonal transformation used for structural comparisons. *Acta Crystallographica Section A: Foundations of Crystallography*, 45(2):208–210, 1989.

KKMM12. Krishnaram Kenthapadi, Aleksandra Korolova, Ilya Mironov, and Nina Mishra. Privacy via the johnson-lindenstrauss transform. *arXiv preprint arXiv:1204.2606*, 2012.

KKNO16. Georgios Kellaris, George Kollios, Kobbi Nissim, and Adam O'Neill. Generic attacks on secure outsourced databases. In Edgar R. Weippl, Stefan Katzenbeisser, Christopher Kruegel, Andrew C. Myers, and Shai Halevi, editors, *ACM CCS 2016*, pages 1329–1340. ACM Press, October 2016.

KLDF16. Albert Kwon, David Lazar, Srinivas Devadas, and Bryan Ford. Riffle: An efficient communication system with strong anonymity. *PoPETs*, 2016(2):115–134, April 2016.

KPT20. Evgenios M. Kornaropoulos, Charalampos Papamanthou, and Roberto Tamassia. The state of the uniform: Attacks on encrypted databases beyond the uniform query distribution. In *2020 IEEE Symposium on Security and Privacy*, pages 1223–1240. IEEE Computer Society Press, May 2020.

LGK06. Kun Liu, Chris Giannella, and Hillol Kargupta. An attacker's view of distance preserving maps for privacy preserving data mining. In *European Conference on Principles of Data Mining and Knowledge Discovery*, pages 297–308. Springer, 2006.

LMP18. Marie-Sarah Lacharité, Brice Minaud, and Kenneth G. Paterson. Improved reconstruction attacks on encrypted data using range query leakage. In *2018 IEEE Symposium on Security and Privacy*, pages 297–314. IEEE Computer Society Press, May 2018.

LP18. Marie-Sarah Lacharité and Kenneth G. Paterson. Frequency-smoothing encryption: preventing snapshot attacks on deterministically encrypted data. *IACR Trans. Symm. Cryptol.*, 2018(1):277–313, 2018.

LW16. Kevin Lewi and David J. Wu. Order-revealing encryption: New constructions, applications, and lower bounds. In Edgar R. Weippl, Stefan Katzenbeisser, Christopher Kruegel, Andrew C. Myers, and Shai Halevi, editors, *ACM CCS 2016*, pages 1167–1178. ACM Press, October 2016.

M+04. George Marsaglia et al. Evaluating the normal distribution. *Journal of Statistical Software*, 11(4):1–7, 2004.

MCO+15. Charalampos Mavroforakis, Nathan Chenette, Adam O'Neill, George Kollios, and Ran Canetti. Modular order-preserving encryption, revisited. In Timos K. Sellis, Susan B. Davidson, and Zachary G. Ives, editors, *Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data, Melbourne, Victoria, Australia, May 31 - June 4, 2015*, pages 763–777. ACM, 2015.

ML14. Marius Muja and David G Lowe. Scalable nearest neighbor algorithms for high dimensional data. *IEEE transactions on pattern analysis and machine intelligence*, 36(11):2227–2240, 2014.

Mor16. J Toby Mordkoff. The assumption (s) of normality. *Dostupno na: goo. gl/g7MCwK (Pristupljeno 27.05. 2017.)*, 2016.

NKW15. Muhammad Naveed, Seny Kamara, and Charles V. Wright. Inference attacks on property-preserving encrypted databases. In Indrajit Ray, Ninghui Li, and Christopher Kruegel, editors, *ACM CCS 2015*, pages 644–655. ACM Press, October 2015.

NM94. Whitney K Newey and Daniel McFadden. Large sample estimation and hypothesis testing, chapter 36, theorem 2.5. *Handbook of econometrics*, 4:2111–2245, 1994.

Osb92. Michael R Osborne. Fisher's method of scoring. *International Statistical Review/Revue Internationale de Statistique*, pages 99–117, 1992.

PBP19. Rishabh Poddar, Tobias Boelter, and Raluca Ada Popa. Arx: An encrypted database using semantically secure encryption. *PVLDB*, 12(11):1664–1678, 2019.

PGW19. David Pouliot, Scott Griffy, and Charles V. Wright. The strength of weak randomization: Easily deployable, efficiently searchable encryption with minimal leakage. In *Dependable Systems and Networks, DSN 2019,*, pages 517–529. IEEE, 2019.

PLZ13. Raluca A. Popa, Frank H. Li, and Nickolai Zeldovich. An ideal-security protocol for order-preserving encoding. In *2013 IEEE Symposium on Security and Privacy*, pages 463–477. IEEE Computer Society Press, May 2013.

PR12. Omkant Pandey and Yannis Rouselakis. Property preserving symmetric encryption. In David Pointcheval and Thomas Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 375–391. Springer, Heidelberg, April 2012.

PW16. David Pouliot and Charles V. Wright. The shadow nemesis: Inference attacks on efficiently deployable, efficiently searchable encryption. In Edgar R. Weippl, Stefan Katzenbeisser, Christopher Kruegel, Andrew C. Myers, and Shai Halevi, editors, *ACM CCS 2016*, pages 1341–1352. ACM Press, October 2016.

Ros18. Richard J Rossi. *Mathematical statistics: an introduction to likelihood based inference*. John Wiley & Sons, 2018.

Sak92. RM Sakia. The box-cox transformation technique: a review. *Journal of the Royal Statistical Society: Series D (The Statistician)*, 41(2):169–178, 1992.

TPS+08. E Onur Turgay, Thomas B Pedersen, Yücel Saygın, Erkay Savaş, and Albert Levi. Disclosure risks of distance preserving data transformations. In *International Conference on Scientific and Statistical Database Management*, pages 79–94. Springer, 2008.

Tuc59. Howard G Tucker. A generalization of the glivenko-cantelli theorem. *The Annals of Mathematical Statistics*, 30(3):828–830, 1959.

TYM14. Isamu Teranishi, Moti Yung, and Tal Malkin. Order-preserving encryption secure beyond one-wayness. In Palash Sarkar and Tetsu Iwata, editors, *ASIACRYPT 2014, Part II*, volume 8874 of *LNCS*, pages 42–61. Springer, Heidelberg, December 2014.

TYUC17. Joel A Tropp, Alp Yurtsever, Madeleine Udell, and Volkan Cevher. Practical sketching algorithms for low-rank matrix approximation. *SIAM Journal on Matrix Analysis and Applications*, 38(4):1454–1485, 2017.

Wei. Eric W Weisstein. "sphere-sphere intersection." from mathworld–a wolfram web resource. https://mathworld.wolfram.com/sphere-sphereintersection.html.

Wei00. Eric W Weisstein. Isometry. 2000.

WL83. M Anthony Wong and Tom Lane. A kth nearest neighbour clustering procedure. *Journal of the Royal Statistical Society: Series B (Methodological)*, 45(3):362–368, 1983.

YGFJ18. Samuel Yeom, Irene Giacomelli, Matt Fredrikson, and Somesh Jha. Privacy risk in machine learning: Analyzing the connection to overfitting. In *2018 IEEE 31st Computer Security Foundations Symposium (CSF)*, pages 268–282, 2018.

YLX13. Bin Yao, Feifei Li, and Xiaokui Xiao. Secure nearest neighbor revisited. In *29th IEEE International Conference on Data Engineering, ICDE 2013*, pages 733–744, 2013.

A Relation between (Approximate) DCP Functions and (Approximate) DP Functions

Theorem 14. *Let* U *be a subset of* R *and* $[0, M]^d$ *denote the d-dimensional Cartesian product of the closed interval* $[0, M]$ *and* $\mathbf{M} = (M, M, \ldots, M)$*. Let* $f : [0, M]^d \to \mathsf{U}$ *with* $f(\mathbf{0}) = (\mathbf{0})$ *and* $\mathbf{N} := f(\mathbf{M})$*. If* f *is* DCP*, then for all* $\mathbf{x} \in [0, M]^d \setminus \{0, \mathbf{M}\}$*, we have* $\frac{\|\mathbf{N}\|}{\|\mathbf{M}\|}(\|\mathbf{x}\| - \sqrt{d}) \leq \|f(\mathbf{x})\| \leq \frac{\|\mathbf{N}\|}{\|\mathbf{M}\|}(\|\mathbf{x}\| + \sqrt{d})$.

This theorem tells us that given two fixed points, a distance-comparison-preserving function always maps a point $\mathbf{x} \in [0, M]^d$ to a point $\mathbf{x}' \in \mathsf{U}$ such that $\mathbf{x}'$ lies in a ball of radius $\sqrt{d}$ around $\mathbf{x}$ scaled by a constant. This shows that any DCP function is approximately distance-preserving. In particular, our claim proves a bound on the amount any point in a DCP function can be perturbed. For concreteness, we have chosen Euclidean spaces for our results, but they can be easily generalized to any metric space.

Proof (of Theorem 14). Assume that for some $\mathbf{x}$,

$$\|f(\mathbf{x})\| \leq \frac{\|\mathbf{N}\|}{\|\mathbf{M}\|}(\|\mathbf{x}\| - \|\mathbf{1}\|) \tag{6}$$

Let $\mathbf{x}$ be such that $\frac{\|\mathbf{N}\|}{\|\mathbf{M}\|}(\|\mathbf{x}\| - \|\mathbf{1}\|) - \|f(\mathbf{x})\|$ is maximal, then $\forall \mathbf{y}$:

$$\frac{\|\mathbf{N}\|}{\|\mathbf{M}\|}(\|\mathbf{y}\| - \|\mathbf{1}\|) - \|f(\mathbf{y})\| \leq \frac{\|\mathbf{N}\|}{\|\mathbf{M}\|}(\|\mathbf{x}\| - \|\mathbf{1}\|) - \|f(\mathbf{x})\|$$

$$\implies \|f(\mathbf{y})\| - \|f(\mathbf{x})\| \geq \frac{\|\mathbf{N}\|}{\|\mathbf{M}\|}(\|\mathbf{y}\| - \|\mathbf{x}\|) \ . \tag{7}$$

Suppose $\|\mathbf{x}\| \leq \frac{\|\mathbf{M}\|}{2}$.

Then by equation 6, $\|f(\mathbf{x})\| - \|f(\mathbf{0})\| \leq \frac{\|\mathbf{N}\|}{\|\mathbf{M}\|}(\|\mathbf{x}\| - \|\mathbf{1}\|)$.

Let $\mathbf{u}$ be a vector whose each component u_i is defined as, $u_i = \min(x_i, 1)$, where x_i is the i^{th} component of $\mathbf{x}$. Note that $\|\mathbf{u}\| \leq \sqrt{d}$.

Using equation 7, with $\mathbf{y} = 2\mathbf{x} - \mathbf{u}$, $\|f(2\mathbf{x}-\mathbf{u})\| - \|f(\mathbf{x})\| \geq \frac{\|\mathbf{N}\|}{\|\mathbf{M}\|}(\|2\mathbf{x}-\mathbf{u}\| - \|\mathbf{x}\|) \geq \frac{\|\mathbf{N}\|}{\|\mathbf{M}\|}(\|2\mathbf{x}\| - \|\mathbf{u}\| - \|\mathbf{x}\|) \geq \frac{\|\mathbf{N}\|}{\|\mathbf{M}\|}(\|\mathbf{x}\| - \|\mathbf{1}\|)$.

So, we have a set of three points, $(\mathbf{0}, \mathbf{x}, 2\mathbf{x}-\mathbf{u})$, such that $\mathsf{dist}((2\mathbf{x}-\mathbf{u}), \mathbf{x}) < \mathsf{dist}(\mathbf{x}, \mathbf{0})$, but $\mathsf{dist}(f(2\mathbf{x}-\mathbf{u}), f(\mathbf{x})) \geq \mathsf{dist}(f(\mathbf{x}), f(\mathbf{0}))$ that implies f is not DCP, which is a contradiction.

The proof for $\|\mathbf{x}\| > \frac{\|\mathbf{M}\|}{2}$ is similar with the contradicting points being $(2\mathbf{x} - \mathbf{M} + \mathbf{u}, \mathbf{x}, \mathbf{M})$. This proves the lower bound. The proof for the upper bound is similar, taking $\mathbf{x}$ to be such that $\|f(\mathbf{x})\| - \frac{\|\mathbf{N}\|}{\|\mathbf{M}\|}(\|\mathbf{x}\| - \|\mathbf{1}\|)$ is maximal.

The remaining part is similar to what has been done for the Theorem 2. Following those steps, we conclude the proof.

Corollary 2. *Let* U *be any subset* R *and* $[0, M]^d$ *denote the d-dimensional Cartesian products of the closed interval* $[0, M]$ *and* $\mathbf{M} = (M, M, \cdots, M)$*. Let* $f : [0, M]^d \to \mathsf{U}$ *be* β-DCP*, and* $f(\mathbf{0}) = (\mathbf{0}), f(\mathbf{M}) = \mathbf{N} \in \mathsf{U}$*. Then* $\forall \mathbf{x} \in [0, M]^d - \{\mathbf{0}, \mathbf{M}\}$, $\frac{\|\mathbf{N}\|}{\|\mathbf{M}\|}(\|\mathbf{x}\| - \beta) \leq \|f(\mathbf{x})\| \leq \frac{\|\mathbf{N}\|}{\|\mathbf{M}\|}(\|\mathbf{x}\| + \beta)$*.*

The above corollary follows directly from Theorem 14. The only difference is that the radius of the ball in which the projected point lies is β. This validates our claim that any β-DCP function is also approximately distance preserving with higher perturbations.

Proof. (Of Corollary 2) Define a vector $\mathbf{k}$ such that $k_i = \min\left(0, \frac{\beta}{\sqrt{d}}\right)$. (Note that $\|\mathbf{k}\| \leq \beta$.) We get a set of three points, $(\mathbf{0}, 2\mathbf{k})$, such that $\mathsf{dist}((2\mathbf{k}), < \mathsf{dist}(\mathbf{0}) - \beta$, but $\mathsf{dist}(f(2\mathbf{k}), f() \geq \mathsf{dist}(f(, f(\mathbf{0}))$ which would imply that f is not β-DCP, making it a contradiction.

Following up with the d-dimensional generalization of Theorem 2 (cf. Theorem 14) completes the proof. ∎

Theorem 15 (n-Dimensions). *If* $f : [0, M]^d \to \mathsf{U}$ *is* β-DCP*, and* $f(\mathbf{0}) = \mathbf{0}, f(M, \cdots, M) = \mathbf{N}$*, then* $\forall \mathbf{x}, \mathbf{y} \in [0, M]^d - \{\mathbf{0}, \mathbf{M}\}$*,*

$$\frac{\|\mathbf{N}\|}{\|\mathbf{M}\|}(\mathsf{dist}(\mathbf{x}, \mathbf{y}) - \gamma - 2\beta) < \mathsf{dist}(f(\mathbf{x}), f(\mathbf{y})) < \frac{\|\mathbf{N}\|}{\|\mathbf{M}\|}(\mathsf{dist}(\mathbf{x}, \mathbf{y}) + \gamma + 2\beta)$$

where $\gamma = \sqrt{2\|\mathbf{x}\|\|\mathbf{y}\| + 2|\mathbf{x}\mathbf{y}|}$.

In other words, if f *is* β-DCP*, then* f *is* (α, β')-DP*, where*

$$\alpha = \frac{\|\mathbf{N}\|}{\|\mathbf{M}\|} \quad \textit{and} \quad \beta' = \frac{\|\mathbf{N}\|}{\|\mathbf{M}\|}(\sqrt{2\|\mathbf{x}\|\|\mathbf{y}\| + 2|\mathbf{x}\mathbf{y}|} + 2\beta) \ .$$

Proof. (of Theorem 15) Without loss of generalization, let us assume, $\|\mathbf{x}\| \leq \|\mathbf{y}\|$. We know that $\frac{\|\mathbf{N}\|}{\|\mathbf{M}\|}(\|\mathbf{x}\|-\beta) \leq \|f(\mathbf{x})\| \leq \frac{\|\mathbf{N}\|}{\|\mathbf{M}\|}(\|\mathbf{x}\|+\beta)$. Thus,

$$\frac{k\mathbf{N}k}{k\mathbf{M}k}(k\mathbf{x}k \quad k\mathbf{y}k \quad 2\beta) \quad kf(\mathbf{x})k \quad kf(\mathbf{y})k \quad \frac{k\mathbf{N}k}{k\mathbf{M}k}(k\mathbf{x}k \quad k\mathbf{y}k + 2\beta) \ . \tag{8}$$

and

$$\frac{k\mathbf{N}k}{k\mathbf{M}k}(k\mathbf{x}k + k\mathbf{y}k \quad 2\beta) \quad kf(\mathbf{x})k + kf(\mathbf{y})k \quad \frac{k\mathbf{N}k}{k\mathbf{M}k}(k\mathbf{x}k + k\mathbf{y}k + 2\beta) \ . \tag{9}$$

Using triangle inequality, we get,

$$\|f(\mathbf{x})\| - \|f(\mathbf{y})\| \leq \mathsf{dist}(f(\mathbf{x}), f(\mathbf{y})) \leq \|f(\mathbf{x})\| + \|f(\mathbf{y})\| \ .$$

) from equations 8 and 9,

$$\frac{k\mathbf{N}k}{k\mathbf{M}k}(k\mathbf{x}k \quad k\mathbf{y}k \quad 2\beta) \quad \mathsf{dist}(f(\mathbf{x}), f(\mathbf{y})) \quad \frac{k\mathbf{N}k}{k\mathbf{M}k}(k\mathbf{x}k + k\mathbf{y}k + 2\beta) \ . \tag{10}$$

Now,

$$\begin{aligned} \mathsf{dist}(\mathbf{x},\mathbf{y})^2 &= \|\mathbf{x}\|^2 + \|\mathbf{y}\|^2 - 2\mathbf{x}\mathbf{y} \\ &= (\|\mathbf{x}\| + \|\mathbf{y}\|)^2 - 2\|\mathbf{x}\|\|\mathbf{y}\| - 2\mathbf{x}\mathbf{y} \ . \\ \implies \|\mathbf{x}\| + \|\mathbf{y}\| &= \sqrt{\mathsf{dist}(\mathbf{x},\mathbf{y})^2 + 2\|\mathbf{x}\|\|\mathbf{y}\| + 2\mathbf{x}\mathbf{y}} \\ &\leq \mathsf{dist}(\mathbf{x},\mathbf{y}) + \sqrt{2\|\mathbf{x}\|\|\mathbf{y}\| + 2\mathbf{x}\mathbf{y}} \\ &\leq \mathsf{dist}(\mathbf{x},\mathbf{y}) + \sqrt{2\|\mathbf{x}\|\|\mathbf{y}\| + 2|\mathbf{x}\mathbf{y}|} \ . \end{aligned} \tag{11}$$

Similarly,

$$\begin{aligned} \|\mathbf{x}\| - \|\mathbf{y}\| &= \sqrt{\mathsf{dist}(\mathbf{x},\mathbf{y})^2 - 2\|\mathbf{x}\|\|\mathbf{y}\| + 2\mathbf{x}\mathbf{y}} \\ &\geq \mathsf{dist}(\mathbf{x},\mathbf{y}) - \sqrt{2\|\mathbf{x}\|\|\mathbf{y}\| - 2\mathbf{x}\mathbf{y}} \\ &\geq \mathsf{dist}(\mathbf{x},\mathbf{y}) - \sqrt{2\|\mathbf{x}\|\|\mathbf{y}\| + 2|\mathbf{x}\mathbf{y}|} \ . \end{aligned} \tag{12}$$

) applying, equations 11 and 12 to equation 10, the relation changes to,

$$\begin{aligned} \frac{\|\mathbf{N}\|}{\|\mathbf{M}\|}(\mathsf{dist}(\mathbf{x},\mathbf{y}) - \sqrt{2\|\mathbf{x}\|\|\mathbf{y}\| + 2|\mathbf{x}\mathbf{y}|} - 2\beta) &\leq \mathsf{dist}(f(\mathbf{x}), f(\mathbf{y})) \\ &\leq \frac{\|\mathbf{N}\|}{\|\mathbf{M}\|}(\mathsf{dist}(\mathbf{x},\mathbf{y}) + \sqrt{2\|\mathbf{x}\|\|\mathbf{y}\| + 2|\mathbf{x}\mathbf{y}|} + 2\beta) \ . \end{aligned}$$

Hence, the result is proved. ∎

B Analysis of canonical RoR adversary's advantage in the proof of Theorem 9

Event A: $\mathbf{Exp}^{\delta-\mathsf{RoR}}_{\mathcal{SE},\mathcal{MD}}(\mathsf{B}) = 1$.

Event B is a sequence of 3 events.

- B chooses $\mathbf{c}_k$.
- Given B has chosen $\mathbf{c}_k$, it selects $\mathbf{m}_b$ (entry which may have been modified) correctly.
- Given the above two events happen, B distinguishes between $\mathbf{m}_0$ and $\mathbf{m}_1$, *i.e.* outputs the challenge bit b.

Claim. The probability of Event A occurring is strictly greater then 0.5 iff Event B occurs.

Proof. If event B occurs, then the probability of event A happening is 1. Hence, we get one side of the implication. For the other side of the implication, we look a the contrapositive which says that if Event B does not happen then event A occurs with probability exactly 0.5.

If B chooses a $\mathbf{c} \in \mathsf{C} \neq \mathbf{c}_k$ and can still guess b correctly, it means that B can guess b significantly better than random if it was given $\mathsf{C}/\mathbf{c}_k$ because every entry of C are independent of each other. This leads to a contradiction since $\mathsf{C}/\mathbf{c}_k$ are identical for $b \in \{0, 1\}$.

Take the case where B choose $\mathbf{c}_k$ correctly but fails to pick $\mathbf{m}_b$, yet guess b with high probability. b only determines the value of $\mathbf{m}_b$, hence b is independent of any other point. This means B does not use the information that $\mathbf{c}_k$ gives regarding b, hence its probability of distinguishing remains 0.5.

The final case where B chooses the correct points but cannot distinguish them is effectively equivalent to B not being able to guess b significantly better than random. ∎

Using Bayes theorem, a consequence of step 3 is,

$$\mathbf{Pr}\left[\mathbf{Exp}^{\delta-\mathsf{RoR}}_{\mathcal{SE},\mathcal{MD}}(\mathsf{B}) = 1\right] \leq \mathbf{Pr}\left[\mathsf{C} = \mathbf{c}_k, \mathsf{M} = \mathbf{m}_b, \mathsf{win} = 1\right] + 0.5 \ .$$

Now,

$$\begin{aligned}
&\mathbf{Pr}\left[\mathsf{C} = \mathbf{c}_k, \mathsf{M} = \mathbf{m}_b, \mathsf{win} = 1\right] \\
=&\mathbf{Pr}\left[\mathsf{win} = 1 | \mathsf{C} = \mathbf{c}_k, \mathsf{M} = \mathbf{m}_b\right] \mathbf{Pr}\left[\mathsf{M} = \mathbf{m}_b | \mathsf{C} = \mathbf{c}_k\right] \mathbf{Pr}\left[\mathsf{C} = \mathbf{c}_k\right] \\
=&\frac{1}{N}\mathbf{Pr}\left[\mathsf{win} = 1 | \mathsf{C} = \mathbf{c}_k, \mathsf{M} = \mathbf{m}_b\right] \mathbf{Pr}\left[\mathsf{M} = \mathbf{m}_b | \mathsf{C} = \mathbf{c}_k\right] \ . \qquad (1)
\end{aligned}$$

Event C: $\mathsf{M} = \mathbf{m}_b | \mathsf{C} = \mathbf{c}_k$.
Let $\mathsf{win}_i = 1$ if given $\mathbf{c}_k$, B can find whether $\mathbf{c}_k = \mathcal{E}nc(\mathbf{m_i})$ or $\mathcal{E}nc(\mathbf{m}_b)$, 0 otherwise.
Event D: $\mathsf{win}_i == 1, \ \forall \ i \in [N]/j$.

Claim. Event C is a subset of Event D.

Proof. We need to show that Event C implies Event D. Proceed by contrapositive. If there exists i such that $\mathsf{win}_i = 0$, B cannot definitively choose $\mathbf{m}_b$ over $\mathbf{m}_i$. Thus, Event C $\subset$ Event D. ∎

Therefore,

$$\begin{aligned}
&\frac{1}{N}\mathbf{Pr}\left[\mathsf{win} = 1 | \mathsf{C} = \mathbf{c}_k, \mathsf{M} = \mathbf{m}_b\right] \mathbf{Pr}\left[\mathsf{M} = \mathbf{m}_b | \mathsf{C} = \mathbf{c}_k\right] \\
&\qquad = \frac{1}{N}\mathbf{Pr}\left[\mathsf{win} = 1 | \mathsf{C} = \mathbf{c}_k, \mathsf{M} = \mathbf{m}_b\right] \prod^{N-1} \left(\mathbf{Pr}\left[\mathsf{win}_i = 1 | \mathsf{C} = \mathbf{c}_k\right]\right) \ . \qquad (13)
\end{aligned}$$

We now find $\mathbf{Pr}\left[\mathsf{win} = 1 | \mathsf{C} = \mathbf{c}_k, \mathsf{M} = \mathbf{m}_b\right]$.

First, we provide some definitions.

Definition 4. *For any message m, the "Ciphertext Region" is the convex hull of all the points which could be a possible ciphertext some choice of secret key and randomness of the encryption algorithm. For example, say in 2 dimensions, if $\overrightarrow{m} = (2, 2); s = 1, 2; \beta = 8$, the possible choices of $\overrightarrow{c}$ would be all points that lie in the circle of radius 2 around $(1, 1)$ and $(4, 4)$. All points enclosed within these two circles form the "Ciphertext Region".*

Definition 5. *Let C_1 and C_2 represent the "Ciphertext Regions" for messages m_1 and m_2 respectively. We define "Overlapping Region"(O) corresponding to these sets of messages as $C_1 \cap C_2$.*

Write

$$\begin{aligned}
&\mathbf{Pr}\left[\mathsf{win} = 1 | \mathsf{C} = \mathbf{c}_k, \mathsf{M} = \mathbf{m}_b\right] \\
&= \mathbf{Pr}\left[\mathsf{win} = 1 | \mathsf{C} = \mathbf{c}_k, \mathsf{M} = \mathbf{m}_b, \mathbf{c}_k \in O\right] \mathbf{Pr}\left[c_k \in O\right] \\
&+ \mathbf{Pr}\left[\mathsf{win} = 1 | \mathsf{C} = \mathbf{c}_k, \mathsf{M} = \mathbf{m}_b, \mathbf{c}_k \notin O\right] \mathbf{Pr}\left[c_k \notin O\right]
\end{aligned}$$

.

The above expression is upper bounded by,

$$(1-p) + p\mathbf{Pr}\left[\mathsf{win} = 1 | \mathsf{C} = \mathbf{c}_k, \mathsf{M} = \mathbf{m}_b, \mathbf{c}_k \in O\right]$$

where $p = \mathbf{Pr}\left[c_k \in O\right]$

Claim. $\mathbf{Pr}\left[\mathsf{win} = 1 | \mathsf{C} = \mathbf{c}_k, \mathsf{M} = \mathbf{m}_b, \mathbf{c}_k \in O\right] = 0.5$

Proof. First, we define the following random variables:

- $\mathcal{S}$: A uniform random variable with support S (keyspace) that represents the secret key (scaling factor).
- Ctxt Represents the ciphertext returned after encryption.
- λ_m : Represents the perturbation factor used to compute the ciphertext.

$$\begin{aligned}
&\Pr[\mathsf{Ctxt} = c | \mathbf{m}, \mathbf{m}', c \in O] \\
&= \sum_s \Pr[\mathsf{Ctxt} = c | \mathbf{m}, \mathbf{m}', c \in O, \mathcal{S} = s] \Pr[\mathcal{S} = s] \\
&= \frac{1}{|\mathsf{S}|} \sum_s \Pr[\lambda_m = \mathbf{c} - s\mathbf{m}_b | \mathbf{m}, \mathbf{m}', c \in O, \mathcal{S} = s] \\
&= \frac{1}{|\mathsf{S}|} \sum_s \Pr[\lambda_m = \mathbf{c} - s\mathbf{m}_b] \\
&= \frac{1}{\mathsf{S}} \sum_s \frac{1}{f(\frac{s\beta}{4}, \delta)}.
\end{aligned}$$

Here, $f(\frac{s\beta}{4}, \delta)$ is a function that denotes the total number of points lying inside O which is independent of all queries. $f(\cdot)$ depends on δ and $\frac{s\beta}{4}$.

The penultimate line in the equation above holds because λ is chosen uniformly from all possible points in O.

Thus the distribution of Ctxt is independent of the queries and oracle responses. This proves perfect secrecy of SAP constrained in the region O_i. ∎

Thus,

$$(1-p) + p\mathbf{Pr}\left[\mathsf{win} = 1 | \mathsf{C} = \mathbf{c}_k, \mathsf{M} = \mathbf{m}_b, \mathbf{c}_k \in O\right] = 1 - 0.5p$$

Computing the value of p is equivalent to finding the ratio of the volume(area in 2-d) of overlapping region to the total area of occupied by both the "Ciphertext Regions". The following depiction will make it easier to understand the method to estimate this ratio.

The result can be first proved for 2-dimensional cases and can be easily extended to arbitrary dimensions.

In 2-dimensions, for a particular s, the area of overlap is formed by two circular caps [Wei] stuck to each other at points C and E (cf. Figure 3). When calculated over all the choice of keys, p - the probability of intersection becomes $\frac{Area(DFG)}{Area(DCABE)}$ which is the area spanned by the circular caps as the secret key changes (cf. Figure 4).

In figure 4, $Area(DCABE)$ can be approximated by the area of the triangle formed by extending AB to join DC and DE (This procedure overestimates the area by a slight margin, thereby providing an upper bound on p). Due to properties of similar triangles, without loss of generalization, we can assume $s = 1$, i.e $\mathsf{rad} = \frac{\beta}{4}$, where rad is the radius of each circle. The height of each cap is $h = \frac{\overline{DF}}{2}$ (cf. Figure 3 $\overline{DF}$: length of line segment DF). We define $a = \frac{\overline{CE}}{2}$ (radius of the base of the cap). θ is the angle formed at the centre of the circles by the end points of the cap *i.e.* C and E (The lines are not drawn explicitly in figure 3).

It is easy to check that the following relations hold and thus we get the value of h.

$$a = \sqrt{\mathsf{rad}^2 - \frac{\delta^2}{4}}, \cos(\theta) = \frac{2\mathsf{rad}^2 - 4a^2}{2\mathsf{rad}^2}, h = a\tan(\frac{\theta}{4}) \ .$$

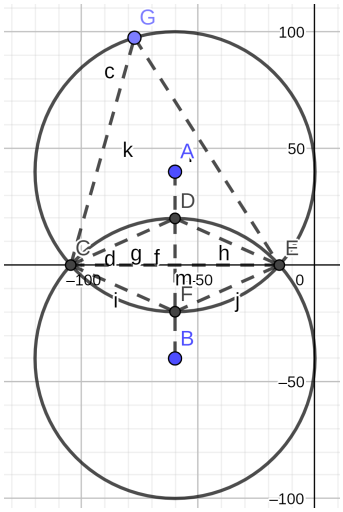

Fig. 3: Depiction for calculating h

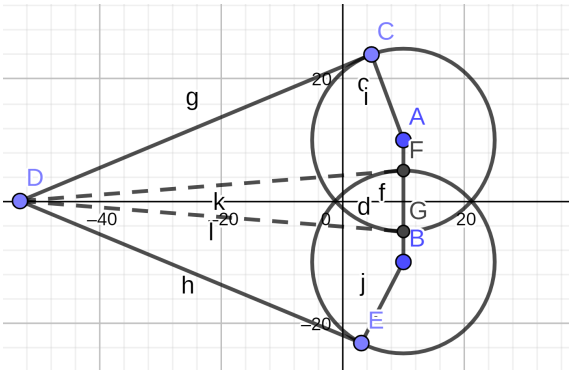

Fig. 4: Cross-section of area of overlap as s changes

By properties of triangles, $p = \frac{h}{\mathsf{rad}+\frac{\delta}{2}}$.

On extension to d-dimensions, the ratio of volume of these d-dimensional cones would be $p = \left(\frac{h}{\mathsf{rad}+\frac{\delta}{2}}\right)^d$.

$$\mathbf{Pr}\left[\mathsf{win} = 1 | \mathsf{C} = \mathbf{c}_k, \mathsf{M} = \mathbf{m}_b\right] = 1 - \left(\frac{h}{2\mathsf{rad}+\delta}\right)^d$$

Now we find $\mathbf{Pr}\left[\mathsf{win}_i = 1 | \mathsf{C} = \mathbf{c}_k\right]$.

$$\mathbf{Pr}\left[\mathsf{win}_i = 1 | \mathsf{C} = \mathbf{c}_k\right]$$

$$= \mathbf{Pr}\left[\mathsf{win}_i = 1 | \mathsf{C} = \mathbf{c}_k, \|\mathbf{m}_i - \mathbf{m}_b\| \leq \delta\right] \mathbf{Pr}\left[\|\mathbf{m}_i - \mathbf{m}_b\| \leq \delta\right]$$

$$+\mathbf{Pr}\left[\mathsf{win}_i = 1 | \mathsf{C} = \mathbf{c}_k, \|\mathbf{m}_i - \mathbf{m}_b\| \geq \delta\right] \mathbf{Pr}\left[\|\mathbf{m}_i - \mathbf{m}_b\| \geq \delta\right]$$

This is upper bounded by $p^*T + (1-T)$, where $T = \mathbf{Pr}\left[\|\mathbf{m}_i - \mathbf{m}_b\| \leq \delta\right]$ and $p^* = 1 - 0.5p$

Combining all the results together we get Equation 13.

$$\frac{1}{N}\mathbf{Pr}\left[\mathsf{win} = 1 | \mathsf{C} = \mathbf{c}_k, \mathsf{M} = \mathbf{m}_b\right] \prod^{N-1} \left(\mathbf{Pr}\left[\mathsf{win}_i = 1 | \mathsf{C} = \mathbf{c}_k\right]\right)$$

$$= \frac{p^*}{N}(p^*T + (1-T))^{N-1} = \frac{1-0.5p}{N}\left(1 - \frac{pT}{2}\right)^{N-1}$$

Thus,

$$\mathbf{Pr}\left[\mathbf{Exp}^{\delta-\mathsf{RoR}}_{\mathcal{SE},\mathcal{MD}}(\mathsf{B})=1\right] \leq \frac{1-0.5p}{N}\left(1-\frac{pT}{2}\right)^{N-1}+0.5$$

which implies,

$$\mathbf{Adv}(B) \leq \frac{2-p}{N}\left(1-\frac{pT}{2}\right)^{N-1} \tag{14}$$

C Bounds on Attribute Window One Wayness (AWOW) advantage

C.1 Upper Bound on AWOW advantage assuming Gaussian distribution on Plaintexts:

Lemma 4. *For any r-*AWOW *adversary A making at most z* Msg *queries,*

$$\boldsymbol{Adv}^{r-\mathsf{SAWOW}}_{\mathsf{TSAP},\mathcal{N}(\mu,\sigma^2)}(A) \leq (1-p)^{|M|}\frac{|\mathsf{S}|+1}{2|\mathsf{S}|}\sum_{m}\left(\left(\frac{2rm-\beta m}{m^2-r^2}-\frac{\beta}{m}\right)\binom{|M|}{m}\right) .$$

Proof. The proof makes the use of two intermediate lemmas. The intuition is to find a relation between Attribute Window One-Wayness and Specified Attribute Window One-Wayness (SAWOW) which is defined below. This is followed by upper bounding the SAWOW advantage.

Lemma 5. *For any* DCPE *scheme with domain and range as $\mathcal{X},\mathcal{Y}$ respectively, window size r, $z \in \mathsf{N}$, and any r-* AWOW *adversary* A, $\exists$ *an equally efficient specified adversary A' such that*

$$\boldsymbol{Adv}^{r-\mathsf{AWOW}}_{\mathsf{DCP}\mathcal{E}_{\mathcal{X},\mathcal{Y}}}(A) \leq z\boldsymbol{Adv}^{r-\mathsf{SAWOW}}_{\mathsf{DCPE}_{\mathcal{X},\mathcal{Y}}}(A') .$$

Proof. Refer of Lemma B.1 of [BCO11] for the proof. ∎

Specified Attribute Window One-Wayness(SAWOW): The specified r,z-attribute window one-wayness is an intermediate security definition where the advantage of an adversary A with respect to $\mathcal{SE}=(\mathcal{K}eyGen,\mathcal{E}nc,\mathcal{D}ec)$, a symmetric encryption scheme, is

$$\mathbf{Adv}^{r-\mathsf{SAWOW}}_{\mathsf{TSAP},\mathcal{D}}(A) = \mathbf{Pr}\left[\mathbf{Exp}^{r-S\mathsf{AWOW}}_{\mathcal{SE}}(A)=1\right]$$

where the experiment $\mathbf{Exp}^{r,z-S\mathsf{AWOW}}_{\mathcal{SE}}(A)$ above is defined as follows.

Experiment $Exp^{r-\mathsf{SAWOW}}_{\mathcal{SE},\mathcal{D}}(A)$:
$K \xleftarrow{\$} \mathcal{K}eyGen$; $S'_M, S_D \leftarrow \emptyset$
$(i, m_L, m_R) \xleftarrow{\$} A^{\mathsf{Enc}(\cdot),\mathsf{Msg}(\cdot)}()$
If $m_R - m_L \leq r$ and for any $a \in [d]$, $m_{ia} \in [m_L, m_R]$ return 1
Else return 0

The only difference between this experiment and the standard $r-$ AWOW is that here, the experiment demands that the adversary return an r-window containing the pre-image of the *specified* ciphertext.

Lemma 6. *For any r-*AWOW *adversary $\mathcal{A}$ making at most z* Msg *queries,*

$$\boldsymbol{Adv}^{r-\mathsf{SAWOW}}_{\mathsf{TSAP},\mathcal{U}(0,|M|)}(A') \leq (1-p)^{|M|}\frac{|\mathsf{S}|+1}{2|\mathsf{S}|}\sum_{m}\left(\left(\frac{2rm-\beta m}{m^2-r^2}-\frac{\beta}{m}\right)\binom{|M|}{m}\right)$$

where M is the attribute space and S *is the keyspace.*

Proof. We compute the upper bound of the attribute window one wayness adversary. Note that we want $m \in [m_L, m_L + r]$. Given a ciphertext$c$, $\underset{k\in\mathcal{K}}{\mathbf{Pr}}\left[Enc(\mathbf{m}) = c\right]$ is same for all m due to uniform distribution of secret key. The decryption to obtain a guess for m, denoted by m^{ml} from c has the transformation, $m^{ml} = \frac{sm+\lambda-\lambda'}{s'}$, where λ', s' are the possible guesses of the scaling and perturbation factors respectively. $\lambda - \lambda' \geq \frac{-s\beta}{2} \implies m_L \geq \frac{sm - s\frac{\beta}{2}}{s'}$ and $\lambda - \lambda' \leq \frac{s\beta}{2} \implies m_U \leq \frac{sm+s\frac{\beta}{2}}{s'}$. We have:

$$
\begin{aligned}
&\Pr[m \in [m_L, m_U]] \\
\leq &\Pr[m - r \leq m_L \leq m] + \Pr[m \leq m_U \leq m + r] \\
= &\Pr[m - r \leq \frac{sm - s\frac{\beta}{2}}{s'} \leq m] + \Pr[m \leq \frac{sm + s\frac{\beta}{2}}{s'} \leq m + r] \\
= &\Pr[\frac{sm - s\frac{\beta}{2}}{m} \leq s' \leq \frac{sm - s\frac{\beta}{2}}{m - r}] + \Pr[\frac{sm + s\frac{\beta}{2}}{m + r} \leq s' \leq \frac{sm + s\frac{\beta}{2}}{m}] \\
= &\frac{1}{|\mathsf{S}|}\left[\frac{sm - s\frac{\beta}{2}}{m - r} - \frac{sm - s\frac{\beta}{2}}{m}\right] + \frac{1}{|\mathsf{S}|}\left[\frac{sm + s\frac{\beta}{2}}{m} - \frac{sm + s\frac{\beta}{2}}{m + r}\right] \\
= &\frac{s}{|\mathsf{S}|}\left(\frac{2rm - \beta m}{m^2 - r^2} + \frac{\beta}{m}\right) .
\end{aligned}
$$

Hence, the adversarial advantage for choice of a secret key is upper bounded by:

$$\frac{s}{|\mathsf{S}|}\left(\frac{2rm - \beta m}{m^2 - r^2} + \frac{\beta}{m}\right) .$$

The average advantage for over the choice of all secret keys is upper bounded,

$$\frac{|\mathsf{S}| + 1}{2|\mathsf{S}|}\left(\frac{2rm - \beta m}{m^2 - r^2} + \frac{\beta}{m}\right) .$$

∎

Since, normal distribution is continuous in nature and in practical scenario, we would to deal with discrete or floating point cases, we perform the approximation of the normal distribution using the Binomial distribution.

The expected advantage assuming a normal distribution approximated by $Bin(M, p)$ is as follows:

$$\frac{|\mathsf{S}| + 1}{2|\mathsf{S}|}\sum_m \left(\frac{2rm - \beta m}{m^2 - r^2} + \frac{\beta}{m}\binom{|M|}{m} p^m (1-p)^{|M|-m}\right) .$$

This expression can be simplified as:

$$
\begin{aligned}
&(1-p)^{|M|}\frac{|\mathsf{S}| + 1}{2|\mathsf{S}|}\sum_m \left(\left(\frac{2rm - \beta m}{m^2 - r^2} - \frac{\beta}{m}\right)\binom{|M|}{m} p^m (1-p)^{-m}\right) \\
\leq &(1-p)^{|M|}\frac{|\mathsf{S}| + 1}{2|\mathsf{S}|}\sum_m \left(\left(\frac{2rm - \beta m}{m^2 - r^2} - \frac{\beta}{m}\right)\binom{|M|}{m}\right) . \\
&\text{(maximum at } p = 0.5)
\end{aligned}
$$

∎

C.2 Lower Bound on AWOW advantage assuming Gaussian Plaintexts:

We prove the following result:

$$\mathbf{Adv}^{r-\mathsf{AWOW}}_{\mathsf{TSAP},\mathcal{U}(0,|M|)}(A) \geq \mathbf{Adv}^{r-\mathsf{SAWOW}}_{\mathsf{TSAP},\mathcal{U}(0,|M|)}(A') \geq \frac{r}{|M|}\sum_m \left(\left(\frac{m - \frac{\beta}{2}}{(m-r)m}\right)\right)$$

The first half of the result,

$$\mathbf{Adv}^{r-\mathsf{AWOW}}_{\mathsf{TSAP},\mathcal{U}(0,|M|)}(A) \geq \mathbf{Adv}^{r-\mathsf{SAWOW}}_{\mathsf{TSAP},\mathcal{U}(0,|M|)}(A')$$

is obvious, as the advantage of the adversary when he has to make one correct guess out of z guesses is obviously more than the advantage when the adversary has to make a particular guess correctly. So, it is sufficient to proof the lower bound on $\mathbf{Adv}^{r-S\mathsf{AWOW}}_{\mathsf{DCPE}_{\mathcal{X},\mathcal{Y}}}(A')$. Consider the following adversary.

Proof. The optimal attack scheme for such an adversary is provided here.

Algorithm 6 Adversary

procedure $A(c_0)$
  $s^* \xleftarrow{\$} [|S|]$
  On $\mathsf{Msg}(\cdot)$ query d^*
    $S_{\mathcal{MD}} \leftarrow S_{\mathcal{MD}} \cup \{(\text{state}, d^*)\}$
    $(\text{state}, m) \xleftarrow{\$} \mathcal{MD}(\text{state}, d^*)$
    $S'_M \leftarrow S'_M \cup \{m\}$
    $c \leftarrow \mathcal{E}nc_K(m)$
    return c
  Repeat until A chooses z ciphertexts
  Fix some $i \in z$, $m_L \leftarrow \frac{c^i - \frac{s\beta}{2}}{s^*}$
  return $[i, m_L, m_L + r]$

We now compute the advantage for such an adversary explicitly which would be the necessary lower bound. Note that we want $m \in [m_L, m_L + r]$. The decryption to obtain a guess for m, denoted by m^{ml} from c has the transformation, $m^{ml} = \frac{sm+\lambda-\lambda'}{s'}$, where λ', s' are the possible guesses of the scaling and perturbation factors respectively. $\lambda - \lambda' \geq \frac{-s\beta}{2} \implies m_L \geq \frac{sm - s\frac{\beta}{2}}{s'}$. For the adversary to succeed, we must have:

$$\begin{aligned} & 0 \leq m - m_L \leq r \\ \implies & 0 \leq m - \frac{sm - s\frac{\beta}{2}}{s'} \leq r \\ \implies & \frac{sm - s\frac{\beta}{2}}{m} \leq s' \leq \frac{sm - s\frac{\beta}{2}}{m - r} \end{aligned}$$

Hence, the adversarial advantage for choice of a secret key is:

$$\frac{r}{|S|}\left(\frac{sm - s\frac{\beta}{2}}{(m-r)m}\right)$$

The average advantage for over the choice of all secret keys is lower bounded by,

$$r\frac{|S|+1}{2|\mathsf{S}|}\left(\frac{m - \frac{\beta}{2}}{(m-r)m}\right)$$

Thus, the expected value of the average advantage is:

$$\mathbf{Pr}\left[\mathbf{Exp}^{r,z-\mathsf{SAWOW}}_{\mathcal{SE}}(A') = 1\right] = r\frac{|S|+1}{2|\mathsf{S}|}\sum_m \left(\frac{m - \frac{\beta}{2}}{(m-r)m}\binom{|M|}{m}\left(\frac{p}{1-p}\right)^m\right)$$

$$r\frac{|S|+1}{2|\mathsf{S}|}\sum_m\left(\frac{m-\frac{\beta}{2}}{(m-r)m}\binom{|M|}{m}\left(\frac{p}{1-p}\right)^m\right)$$

$$=r\frac{|S|+1}{2|\mathsf{S}|\sqrt{|M|}}\sum_m\left(\frac{m-\frac{\beta}{2}}{(m-r)}\left(\frac{1}{m}\right)\left(\frac{M}{m}\right)^m\left(\frac{p}{1-p}\right)^m\right)$$

(By Sterling Approximation)

$$\geq r\frac{|S|+1}{2|\mathsf{S}|\sqrt{|M|}}\sum_m\left(\left(\frac{1}{m}\right)\left(\frac{M}{m}\right)^m\left(\frac{p}{1-p}\right)^m\right)$$

Since, $r \geq \frac{\beta}{2}$

$$=r\frac{|S|+1}{2|\mathsf{S}|\sqrt{|M|}}\ln\frac{1-p}{1-2p}$$

$$\geq\frac{r}{2\sqrt{|M|}}\ln\frac{1-p}{1-2p}\,.$$

∎