

Relative Information Loss in the PCA

Bernhard C. Geiger*, Gernot Kubin*

*Signal Processing and Speech Communication Laboratory, Graz University of Technology, Austria
{geiger,gernot.kubin}@tugraz.at

Abstract—In this work we analyze principle component analysis (PCA) as a deterministic input-output system. We show that the relative information loss induced by reducing the dimensionality of the data after performing the PCA is the same as in dimensionality reduction without PCA. Furthermore, we analyze the case where the PCA uses the sample covariance matrix to compute the rotation. If the rotation matrix is not available at the output, we show that an infinite amount of information is lost. The relative information loss is shown to decrease with increasing sample size.

I. INTRODUCTION

Principle component analysis (PCA) is a powerful tool for both linear decorrelation and dimensionality reduction, and is thus widely used in machine learning, neural networks, and feature extraction [1], [2]. A vast literature proves the optimality of PCA in an information theoretic sense for specific cases: Linsker [3] proved that for N -dimensional Gaussian data corrupted by Gaussian noise with diagonal covariance matrix, the PCA maximizes the mutual information between the data and a one-dimensional output random variable. Plumbley argued in [4] that for non-Gaussian data the PCA is the linear transform which minimizes an upper bound on the information lost due to dimensionality reduction. Furthermore, the authors of [1] show that under some circumstances the PCA minimizes redundancy in the output data, i.e., is an optimal linear independent component analysis. While this list of previous works is clearly not complete, it highlights the utility of information-theoretic measures to characterize the performance of PCA algorithms.

In the present work, we pursue a slightly different approach: We view the PCA as a multivariate, vector-valued input-output system and analyze it in terms of information loss. In this context, *information* means the *total information* available at the input of the system, in contrast to [3], [4], which considered only information which is *relevant*. As a consequence, while our analysis may superficially appear to contradict the results available in the literature, it rather provides a different view on PCA as an information-processing system.

In its standard notation, PCA is a linear system

$$\mathbf{Y} = \mathbf{W}^T \mathbf{X} \quad (1)$$

where \mathbf{X} and \mathbf{Y} are the N -dimensional input and output vectors, respectively, and \mathbf{W} is an orthogonal matrix. In this sense, the PCA is a bijective transform and thus invertible. Typically, however, the output vector will be truncated after performing the PCA in order to reduce the dimensionality of the data. In this projection to a lower-dimensional subspace

information is lost, and we intend to quantify the loss in this work.

Let us be more specific: Assuming continuous-valued random variables (RV), the information available at the input is infinite. In particular, assuming joint continuity, each component of the multi-dimensional input RV contains an infinite amount of information; thus, if the dimension is reduced, i.e., if some components are dropped, an infinite amount of information is lost (cf. [5]).

In case the orthogonal matrix is not known a priori but has to be estimated from a set of input data vectors collected in the matrix $\underline{\mathbf{X}}$, the PCA becomes a nonlinear operation:

$$\underline{\mathbf{Y}} = \underline{\mathbf{w}}(\underline{\mathbf{X}})\underline{\mathbf{X}}. \quad (2)$$

Here, $\underline{\mathbf{w}}$ is a matrix-valued function which computes the orthogonal matrix required for rotating the data (e.g., using the QR algorithm). If this orthogonal matrix is not stored and made available at the output, one will readily agree that information is lost even if – at a first glance – the dimension of the data is not reduced.

Our notion of *relative* information loss captures the ratio of input information which cannot be retrieved at the output. We make this statement precise in Section II. Before analyzing the information loss in PCA in Section V, we present a general theorem about the loss in systems which reduce the dimensionality of the data (Section III). Section IV acts as a bridge, containing some toy examples which should give an intuitive understanding of relative information loss. In Section VI, the PCA using an input data matrix (cf. (2)) to perform the rotation is shown to destroy information even if the dimensionality is not reduced. Since none of our findings explains the usefulness of the PCA, we finally give an outlook of how to corroborate its optimality using a different notion of information loss in Section VII.

An extended version of this paper, including some proofs and additional examples, can be found in [6].

II. RELATIVE INFORMATION LOSS – A QUICK OVERVIEW

In this section we will briefly present the basic properties of relative information loss. To this end, we introduce

Definition 1. Let \mathbf{X} be an N -dimensional RV valued in \mathcal{X} , and let \mathbf{Y} be obtained by transforming \mathbf{X} with a static function g , i.e., $\mathbf{Y} = g(\mathbf{X})$. We define the *relative information loss* induced by this transform as

$$l(\mathbf{X} \rightarrow \mathbf{Y}) = \lim_{n \rightarrow \infty} \frac{H(\hat{\mathbf{X}}_n | \mathbf{Y})}{H(\hat{\mathbf{X}}_n)} \quad (3)$$

where $\hat{\mathbf{X}}_n = \frac{\lfloor n\mathbf{X} \rfloor}{n}$ (element-wise). The quantity on the left is defined if the limit on the right-hand side exists.

Loosely speaking, this quantity represents the percentage of input information which is lost in the system, with every bit weighted equally. Taking, for example, an 8 bit number, losing the most significant or the least significant bit both amounts to a relative information loss of $l(X \rightarrow Y) = \frac{1}{8}$. The advantage of this definition is its independence of application-specific aspects (where, e.g., the most significant bit may be more important than the least significant one; cf. Section VII for a discussion).

The motivation for introducing this quantity is to complement the absolute notion of information loss, given by $L(\mathbf{X} \rightarrow \mathbf{Y}) = H(\mathbf{X}|\mathbf{Y})$, which suffers from being infinite in many practically relevant cases (cf. [5]).

Among other things, we are able to show that $l(\mathbf{X} \rightarrow \mathbf{Y}) \in [0, 1]$ and that, for continuous input RVs \mathbf{X} , a nonzero relative loss corresponds to an infinite absolute loss (see [6]). Furthermore, we can show a tight connection to the information dimension introduced by Rényi [7]:

Theorem 1. *Let \mathbf{X} be an RV with positive information dimension $d(\mathbf{X}) = \lim_{n \rightarrow \infty} H(\hat{\mathbf{X}}_n) / \log n$. Then, if $d(\mathbf{X}|\mathbf{Y} = \mathbf{y})$ exists $P_{\mathbf{Y}}$ -a.s., the relative information loss equals*

$$l(\mathbf{X} \rightarrow \mathbf{Y}) = \frac{d(\mathbf{X}|\mathbf{Y})}{d(\mathbf{X})} \quad (4)$$

where $d(\mathbf{X}|\mathbf{Y}) = \int_{\mathcal{Y}} d(\mathbf{X}|\mathbf{Y} = \mathbf{y}) dP_{\mathbf{Y}}(\mathbf{y})$.

Proof: See [6]. ■

III. RELATIVE INFORMATION LOSS FOR FUNCTIONS WHICH REDUCE DIMENSIONALITY

We now proceed to analyzing the relative information loss for measurable functions $\mathbf{g}: \mathcal{X} \rightarrow \mathcal{Y}$ which map subsets of $\mathcal{X} \subseteq \mathbb{R}^N$ to sufficiently well-behaved submanifolds of \mathbb{R}^N . In particular, let $\{\mathcal{X}_i\}$, $i = 1, \dots, L$, denote a finite partition of \mathcal{X} and let $P_{\mathbf{X}} \ll \mu^N$, where μ^N is the N -dimensional Lebesgue measure. Here and throughout the remainder of this paper we assume that all involved information dimensions exist and are finite. The latter restriction is fulfilled under the mild condition that $H(\hat{\mathbf{X}}_1) < \infty$ [7], which for scalar (one-dimensional) RVs X translates to $E\{|X|^\epsilon\} < \infty$ for some $\epsilon > 0$ [8].

We maintain

Theorem 2. *Let \mathbf{X} be such that $P_{\mathbf{X}} \ll \mu^N$ is supported on $\mathcal{X} \subseteq \mathbb{R}^N$ and let $\{\mathcal{X}_i\}$ be a partition of \mathcal{X} such that each of its L elements is a smooth N -dimensional manifold. Let \mathbf{g} be such that $\mathbf{g}_i = \mathbf{g}|_{\mathcal{X}_i}$ are submersions to disjoint smooth n_i -dimensional manifolds \mathcal{Y}_i . Then, the relative information loss is*

$$l(\mathbf{X} \rightarrow \mathbf{Y}) = \sum_{i=1}^L P_{\mathbf{X}}(\mathcal{X}_i) \frac{N - n_i}{N}. \quad (5)$$

Proof: By the submersion theorem (e.g., [9]) and the fact that the images of \mathcal{X}_i are disjoint, the preimage of every point

in \mathcal{Y}_i is a closed submanifold of \mathcal{X}_i with codimension n_i . We can thus use Theorem 1 together with $d(\mathbf{X}) = N$ to get

$$\begin{aligned} l(\mathbf{X} \rightarrow \mathbf{Y}) &= \frac{1}{N} \sum_{i=1}^L \int_{\mathcal{Y}_i} d(\mathbf{X}|\mathbf{Y} = \mathbf{y}) dP_{\mathbf{Y}}(\mathbf{y}) \\ &= \frac{1}{N} \sum_{i=1}^L \int_{\mathcal{Y}_i} (N - n_i) dP_{\mathbf{Y}}(\mathbf{y}) = \sum_{i=1}^L \frac{N - n_i}{N} P_{\mathbf{Y}}(\mathcal{Y}_i) \end{aligned} \quad (6)$$

which completes the proof by noticing that $P_{\mathbf{Y}}(\mathcal{Y}_i) = P_{\mathbf{X}}(\mathcal{X}_i)$. ■

Before proceeding, two interesting facts are worth mentioning: First of all, as a submersion is a smooth mapping between smooth manifolds, we can state

Corollary 1. *If $P_{\mathbf{X}} \ll \mu^N$ and if \mathbf{g} is as in Theorem 2, then*

$$l(\mathbf{X} \rightarrow \mathbf{Y}) = 1 - \frac{d(\mathbf{Y})}{d(\mathbf{X})} \quad (7)$$

provided the information dimension of \mathbf{Y} exists.

Proof: We note from the proof of Theorem 2 that

$$l(\mathbf{X} \rightarrow \mathbf{Y}) = 1 - \sum_{i=1}^L \frac{n_i}{N} P_{\mathbf{Y}}(\mathcal{Y}_i) \quad (8)$$

where $N = d(\mathbf{X})$. We now show that $d(\mathbf{Y}|\mathbf{X} \in \mathcal{X}_i) = n_i$. By the fact that $\mathbf{g}|_{\mathcal{X}_i} \in \mathcal{C}^\infty$ is a submersion, the preimages of μ^{n_i} -null sets are μ^N -null sets themselves [10]. Consequently, since $P_{\mathbf{X}} \ll \mu^N$, the conditional probability measure supported on \mathcal{Y}_i is absolutely continuous w.r.t. the n_i -dimensional Lebesgue measure ($P_{\mathbf{Y}|\mathbf{X} \in \mathcal{X}_i} \ll \mu^{n_i}$). Since \mathcal{Y}_i is a smooth n_i -dimensional manifold, we again obtain with Rényi [7] that $d(\mathbf{Y}|\mathbf{X} \in \mathcal{X}_i) = n_i$. The proof is complete with [8, Theorem 2] or [11], noting that

$$d(\mathbf{Y}) = \sum_{i=1}^L d(\mathbf{Y}|\mathbf{X} \in \mathcal{X}_i) P_{\mathbf{X}}(\mathcal{X}_i) = \sum_{i=1}^L n_i P_{\mathbf{Y}}(\mathcal{Y}_i). \quad (9)$$

For more complicated measures (e.g., measures with a non-integer information dimension) or more general functions it might be hard to prove Theorem 2 and Corollary 1. We conjecture, however, that the conditions imposed in Theorem 2 can be loosened such that at least for a larger class of functions (e.g., those for which the images of \mathcal{X}_i are not disjoint) the relative information loss can be evaluated.

Secondly, we believe that one should be able to analyze cascaded systems, at least in a restricted class of cases: To this end, assume that we have a cascade of two projections: one on the first two and then one on the first coordinate, i.e., $\mathbf{g}: [X_1, X_2, X_3] \rightarrow [X_1, X_2]$ and $h: [X_1, X_2] \rightarrow X_1$. Let the RVs X_1 , X_2 , and X_3 have a continuous joint distribution. Then, according to Theorem 2, the function \mathbf{g} destroys one third of the information, while h destroys half of the remaining information. In total, two thirds of the information are lost. Indeed, we obtain

$$l([X_1, X_2, X_3] \rightarrow X_1) = \frac{1}{3} + \frac{1}{2} - \frac{1}{3} \frac{1}{2} = \frac{2}{3}. \quad (10)$$

We are thus lead to the following

Conjecture 1. *Let g be a (measurable) function describing a system with input \mathbf{X} and output \mathbf{Y} . Let further h be another such system function which transforms input \mathbf{Y} to output \mathbf{Z} . There exist conditions under which*

$$l(\mathbf{X} \rightarrow \mathbf{Z}) = l(\mathbf{X} \rightarrow \mathbf{Y}) + l(\mathbf{Y} \rightarrow \mathbf{Z}) - l(\mathbf{X} \rightarrow \mathbf{Y})l(\mathbf{Y} \rightarrow \mathbf{Z}). \quad (11)$$

Clearly, our example of the cascade of two projections fulfills these conditions. Moreover, as we show in [6], this conjecture holds for discrete RVs with finite entropy.

Finally, it is worth mentioning that the shape of the distribution has no effect on the relative amount of information lost. It is essentially this behavior which leads to the somewhat counter-intuitive results presented in the following sections.

IV. TOY EXAMPLES FOR DIMENSIONALITY REDUCTION

With the help of a few simple examples we now try to make the operational meaning of relative information loss intuitive. At the same time we highlight its importance to the development of an information-centered system theory.

Let us introduce a two-dimensional RV \mathbf{X} , $P_{\mathbf{X}} \ll \mu^2$ and $d(\mathbf{X}) = 2$, and a transform simply adding the two vector components, i.e.,

$$Y = g(\mathbf{X}) = X_1 + X_2. \quad (12)$$

We can represent this transform by a cascade of an invertible linear transform (e.g., $\underline{\mathbf{T}}: \mathbf{X} \rightarrow [X_1 + X_2, X_1]$) and a simple projection onto the first component. Let now $\tilde{\mathbf{X}} = [X_1 + X_2, X_1]$. Since $\underline{\mathbf{T}}$ is invertible and linear, it is bi-Lipschitz and thus preserves the dimension of the transformed RVs. Thus,

$$l(\mathbf{X} \rightarrow Y) = l(\tilde{\mathbf{X}} \rightarrow Y) = 0.5 \quad (13)$$

by Theorem 2. Note that this is another example where Conjecture 1 holds.

With this result and the underlying theorems, the theory of information-processing systems can be extended from treating cascade structures as in (11) to also considering parallel structures whose outputs are added. This were not possible by considering only absolute information loss (e.g., [5]), as it would be infinite in these cases.

As a second toy example we consider a center clipper, which is commonly used for noise suppression or residual echo cancellation. We describe the center clipper by the following function:

$$g(x) = \begin{cases} x, & \text{if } |x| > c \\ 0, & \text{else} \end{cases}. \quad (14)$$

Clearly, the domain of this function can be partitioned into three elements, upon two of them the function is the identity function. On the third set, $[-c, c]$, the function is a submersion to a zero-dimensional manifold. We can thus apply Theorem 2 and obtain $l(X \rightarrow Y) = P_X([-c, c])$.

It is interesting to see that in both examples the distribution of the input signal does not have an influence on the relative

information loss, as long as it is continuous. This is counter-intuitive in the sense that adding two strongly correlated RVs should preserve more information than adding two independent RVs, or, in the sense that center clipping a large signal should not hurt too much. This intuition, however, is based on the fact that one tends to attribute unequal importance to different aspects of the information contained in the input signal (e.g., principle direction, magnitude, etc.).

V. PCA WITH POPULATION COVARIANCE MATRIX

In PCA one uses the eigenvalue decomposition (EVD) of the covariance matrix of a multivariate input to obtain a different representation of the input vector. In particular, let \mathbf{X} be an RV with distribution $P_{\mathbf{X}} \ll \mu^N$ and information dimension $d(\mathbf{X}) = N$. We further assume that \mathbf{X} has zero mean and a positive definite population covariance matrix $\underline{\mathbf{C}}_{\mathbf{X}} = \mathbb{E}\{\mathbf{X}\mathbf{X}^T\}$ which is known a priori. The case where $\underline{\mathbf{C}}_{\mathbf{X}}$ is not known but has to be estimated from the data is considered in Section VI.

The EVD of the covariance matrix yields

$$\underline{\mathbf{C}}_{\mathbf{X}} = \underline{\mathbf{W}}\underline{\Sigma}\underline{\mathbf{W}}^T \quad (15)$$

where $\underline{\mathbf{W}}$ is an orthogonal matrix (i.e., $\underline{\mathbf{W}}^{-1} = \underline{\mathbf{W}}^T$) and $\underline{\Sigma}$ is a diagonal matrix consisting of the N eigenvalues of $\underline{\mathbf{C}}_{\mathbf{X}}$. We now can describe the PCA by the following linear transform:

$$\mathbf{Y} = g(\mathbf{X}) = \underline{\mathbf{W}}^T \mathbf{X}. \quad (16)$$

As in Section IV, the linear transform is bi-Lipschitz, and the information loss vanishes¹.

Often, however, the PCA is used for dimensionality reduction, where after the linear transform in (16) the elements of the random vector \mathbf{Y} with the smallest variances are discarded (thus, preserving the subspace with the largest variance). Essentially, the mapping from \mathbf{Y} to, e.g.,

$$\mathbf{Y}_M = [Y_1, \dots, Y_M] \quad (17)$$

is a projection onto the first $M < N$ coordinates, which is a submersion between two smooth manifolds. We can thus apply Theorem 2 and obtain the relative information loss $l(\mathbf{Y} \rightarrow \mathbf{Y}_M) = \frac{N-M}{N}$. In analogy with the example in Section IV we thus get

$$l(\mathbf{X} \rightarrow \mathbf{Y}_M) = \frac{N-M}{N}. \quad (18)$$

We can now extend this analysis to the case where from \mathbf{Y}_M an (N -dimensional) estimate $\tilde{\mathbf{X}}_M$ of the original data \mathbf{X} is reconstructed. This estimate is obtained using a linear transform

$$\tilde{\mathbf{X}}_M = \underline{\mathbf{W}}\underline{\mathbf{I}}_M^T \mathbf{Y}_M \quad (19)$$

where $\underline{\mathbf{I}}_M$ is a rectangular identity matrix with M rows and N columns. The (full-rank) matrix $\underline{\mathbf{I}}_M^T$ is a mapping to a higher-dimensional space (i.e., from \mathbb{R}^M to \mathbb{R}^N) and is thus bi-Lipschitz; so is the rotation with the matrix $\underline{\mathbf{W}}$.

¹Not only the relative information loss $l(\mathbf{X} \rightarrow \mathbf{Y}) = 0$, but also the absolute information loss $L(\mathbf{X} \rightarrow \mathbf{Y}) = 0$.

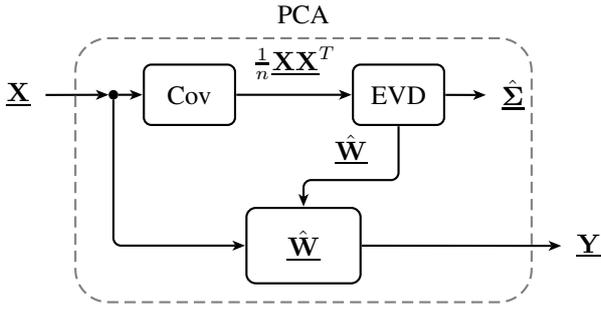


Fig. 1. The PCA as a nonlinear input-output system. “Cov” denotes the computation of the sample covariance matrix and “EVD” stands for eigenvalue decomposition.

Furthermore, the transform from \mathbf{Y}_M to $\tilde{\mathbf{X}}_M$ is invertible and, as a consequence, no additional information is lost. We thus state

$$l(\mathbf{X} \rightarrow \tilde{\mathbf{X}}_M) = \frac{N - M}{N} \quad (20)$$

where, using above notation, $\tilde{\mathbf{X}}_M = \mathbf{W}\mathbf{I}_M^T \mathbf{I}_M \mathbf{W}^T \mathbf{X}$.

Indeed, the same result would have been obtained if the rotation would have been performed using any other orthogonal matrix and regardless which elements of the rotated vector were discarded. In particular, also if just the first M components of \mathbf{X} would have been preserved, we would have $l(\mathbf{X} \rightarrow \mathbf{X}_M) = \frac{N-M}{N}$.

Of course, PCA is known to be optimal in the sense that, by discarding the elements of \mathbf{Y} with the smallest variances, the mean-squared reconstruction error for $\tilde{\mathbf{X}}_M$ is minimized [1]. For this interpretation, measuring the information loss *with respect to a relevant random variable* may do the trick, providing us with a statement about the optimality of PCA in an information-theoretic sense (cf. [1], [4]). Conversely, if one cannot determine which information at the input is relevant, one has no reason to perform PCA prior to reducing the dimension of the data.

VI. PCA WITH SAMPLE COVARIANCE MATRIX

We argued in Section V that the PCA without dimensionality reduction is an invertible transform. This, however, is only true if one has access to the orthogonal matrix \mathbf{W} ; if not, i.e., if one feeds a system with the data matrix \mathbf{X} and just receives the rotated matrix \mathbf{Y} (see Fig. 1), it can be shown that information is lost. In this section we make this statement precise.

We believe that the following analysis can be generalized to all data matrices with a continuous joint distribution, i.e., $P_{\mathbf{X}} \ll \mu^{nN}$. For the sake of simplicity, in this paper, we will focus on a particularly simple scenario: Let \mathbf{X} denote a matrix where each of its n columns represents an *independent* sample of an N -dimensional Gaussian RV \mathbf{X} . Again, let \mathbf{X} have zero mean and positive definite population covariance matrix $\mathbf{C}_{\mathbf{X}}$. As a consequence, the probability distribution of the data matrix \mathbf{X} is absolutely continuous w.r.t. the nN -dimensional Lebesgue measure ($P_{\mathbf{X}} \ll \mu^{nN}$).

The sample covariance matrix $\hat{\mathbf{C}}_{\mathbf{X}} = \frac{1}{n} \mathbf{X} \mathbf{X}^T$ is symmetric and almost surely positive definite. In the usual case where $n \geq N$ one can show that $\frac{N(N+1)}{2}$ entries can be chosen and that the remaining entries depend on these in a deterministic manner. Indeed, since in this case the distribution of $\hat{\mathbf{C}}_{\mathbf{X}}$ possesses a density (the Wishart distribution, cf. [12]), the distribution is absolutely continuous w.r.t. the Lebesgue measure on an $\frac{N(N+1)}{2}$ -dimensional submanifold of the N^2 -dimensional Euclidean space. With some abuse of notation we thus write $P_{\hat{\mathbf{C}}_{\mathbf{X}}} \ll \mu^{\frac{N(N+1)}{2}}$.

The orthogonal matrix $\hat{\mathbf{W}}$ for PCA (see (16); now applied to the matrix \mathbf{X} instead of the vector \mathbf{X}) is obtained from the EVD of the sample covariance matrix, i.e.,

$$\hat{\mathbf{C}}_{\mathbf{X}} = \hat{\mathbf{W}} \hat{\mathbf{\Sigma}} \hat{\mathbf{W}}^T \quad (21)$$

where $\hat{\mathbf{\Sigma}}$ is the diagonal matrix containing the eigenvalues of $\hat{\mathbf{C}}_{\mathbf{X}}$. The joint distribution of the N eigenvalues of $\hat{\mathbf{C}}_{\mathbf{X}}$ possesses a density [12, Ch. 9.4]; thus, the distribution of $\hat{\mathbf{\Sigma}}$ is absolutely continuous w.r.t. the Lebesgue measure on an N -dimensional submanifold of the N^2 -dimensional Euclidean space, or $P_{\hat{\mathbf{\Sigma}}} \ll \mu^N$.

Clearly, the entries of $\hat{\mathbf{C}}_{\mathbf{X}}$ are smooth functions of the eigenvalues and the entries of $\hat{\mathbf{W}}$. Images of Lebesgue null-sets under smooth functions between Euclidean spaces of same dimension are null-sets themselves; were the probability measure $P_{\hat{\mathbf{W}}}$ supported on some set of dimensionality lower than $\frac{N(N-1)}{2}$, the image of the product of this set and \mathbb{R}^N (for the eigenvalues) would be a Lebesgue null-set with positive probability measure. Since this contradicts the fact that $\hat{\mathbf{C}}_{\mathbf{X}}$ is continuously distributed, it follows that

$$P_{\hat{\mathbf{W}}} \ll \mu^{\frac{N(N-1)}{2}}. \quad (22)$$

We now argue (in [6] more rigorously than here) that the rotated data does not tell us anything about the rotation, hence

$$P_{\hat{\mathbf{W}}|\mathbf{Y}=\mathbf{y}} \ll \mu^{\frac{N(N-1)}{2}}. \quad (23)$$

Knowing $\mathbf{Y} = \mathbf{y}$, \mathbf{X} is a linear, bi-Lipschitz function of $\hat{\mathbf{W}}$, and thus the information dimension remains unchanged. Therefore we get with Theorem 1

$$l(\mathbf{X} \rightarrow \mathbf{Y}) = \frac{d(\mathbf{X}|\mathbf{Y})}{d(\mathbf{X})} = \frac{N(N-1)}{2nN} = \frac{N-1}{2n}. \quad (24)$$

We now drop a little of the mathematical rigor to analyze, for the sake of completeness, the less common case where there are less data samples than there are dimensions for each sample ($n < N$). In this case, the sample covariance matrix is not full rank, which means that the EVD yields $N - n$ vanishing eigenvalues. Assuming that still $P_{\hat{\mathbf{C}}_{\mathbf{X}}} \ll \mu^{n \frac{2N-n+1}{2}}$ one finds along the same lines as in the case $n \geq N$ that the loss evaluates to

$$l(\mathbf{X} \rightarrow \mathbf{Y}) = \frac{2N - n - 1}{2N}. \quad (25)$$

The behavior of the relative information loss as a function of n is shown in Fig. 2 for different choices of N .

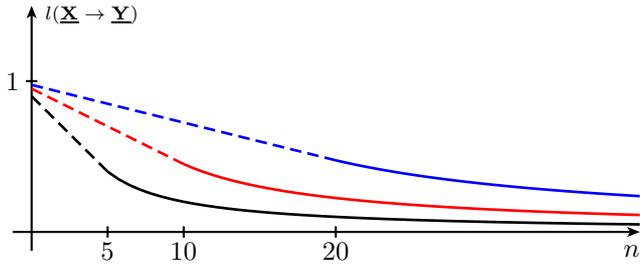


Fig. 2. Relative information loss in the PCA with sample covariance matrix as a function of the number n of independent measurements. The cases $N = 5$ (black), $N = 10$ (red), and $N = 20$ (blue) are shown. The dashed lines indicate the conjectured loss for singular sample covariance matrices.

The relative information loss induced by PCA results from the fact that one cannot know *which rotation* led to the output data matrix. As a consequence, the relative information loss decreases with a larger number of samples: the total information increases while the uncertainty about the rotation remains the same. Note further that the relative information loss cannot exceed $\frac{N-1}{N}$ (for $n = 1$), which is due to the fact that the rotation preserves the norm of the sample.

Finally, it is interesting to observe that in this case the *absolute* information loss in the PCA is infinite [6], even if the dimension of the data is preserved. Moreover, this analysis not only holds for the PCA, but for *any* rotation which depends on the input data in a similar manner – in this sense, the PCA is thus not better than any other rotation.

VII. DISCUSSION

Relative information loss was introduced to cope with the shortcomings of absolute information loss, especially in cases where the information loss $L(\mathbf{X} \rightarrow \mathbf{Y})$ and the information transfer $I(\mathbf{X}; \mathbf{Y})$ are infinite. However, as the examples of Section IV and V showed, even relative information loss not necessarily provides full insight and intuitive interpretations in some cases.

These cases may be characterized by the fact that not all information at the input of the system is *relevant* for a given application, i.e., we are actually not interested in \mathbf{X} itself but in some random variable \mathbf{Z} somehow related to it. \mathbf{Z} , however, is not accessible directly, but only through a function g of the related RV \mathbf{X} . The logical consequence is thus to define a quantity which captures the information loss relevant w.r.t. \mathbf{Z} . We did so in [13], where we analyzed some of this quantity's properties and made a connection to the signal enhancement problem.

A very similar quantity has already been introduced by Plumbley [4] in the context of unsupervised learning. Using this quantity he proved the optimality of the PCA under appropriate constraints, and argued that information loss in some cases is more versatile than mutual information.

In order to be able to minimize this *relevant information loss*, one clearly has to know something about the relevant RV \mathbf{Z} and its relationship to the system input \mathbf{X} . If one does not have this knowledge, all bits of information have to be treated equally, leading to our notion of (relative)

information loss. As a direct consequence, while in some cases the PCA prior to dimensionality reduction might be the optimal solution (cf. [4], [13]), without knowledge about the relevant information one has no reason to apply it to a data set. Even more so, as we showed in Section VI, even without reducing the dimensionality, information can be lost, which should prevent one from unjustified use of PCA.

VIII. CONCLUSION

We have introduced the notion of relative information loss to analyze systems for which the output has a lower dimension than the input, exploiting a tight connection to Rényi's information dimension. As a first result, we showed that the relative information loss for dimensionality reduction is not affected by performing a principle component analysis (PCA) beforehand.

We then showed that even without dimensionality reduction the PCA is not information lossless, given that the sample covariance matrix is used to compute the rotation matrix. There, the relative information loss appears to decrease with increasing sample size.

We proposed to use these somewhat counter-intuitive results as motivation to introduce a notion of information loss which takes the relevance of the data into account, and thus connects to results available in the literature. A detailed analysis of this relevant information loss is within the scope of future work.

REFERENCES

- [1] G. Deco and D. Obradovic, *An Information-Theoretic Approach to Neural Computing*. New York, NY: Springer, 1996.
- [2] J. C. Principe, *Information Theoretic Learning: Rényi's Entropy and Kernel Perspectives*, ser. Information Science and Statistics. New York, NY: Springer, 2010.
- [3] R. Linsker, "Self-organization in a perceptual network," *IEEE Computer*, vol. 21, no. 3, pp. 105–117, Mar. 1988.
- [4] M. Plumbley, "Information theory and unsupervised neural networks," Cambridge University Engineering Department, Tech. Rep. CUED/F-INFENG/TR. 78, 1991.
- [5] B. C. Geiger and G. Kubin, "On the information loss in memoryless systems: The multivariate case," in *Proc. Int. Zurich Seminar on Communications (IZS)*, Zurich, Feb. 2012, pp. 32–35, extended version available: arXiv:1109.4856 [cs.IT].
- [6] —, "Relative information loss in the PCA," Apr. 2012, arXiv:1204.0429 [cs.IT].
- [7] A. Rényi, "On the dimension and entropy of probability distributions," *Acta Mathematica Hungarica*, vol. 10, no. 1-2, pp. 193–215, Mar. 1959.
- [8] Y. Wu and S. Verdú and, "Rényi information dimension: Fundamental limits of almost lossless analog compression," *IEEE Trans. Inf. Theory*, vol. 56, no. 8, pp. 3721–3748, Aug. 2010.
- [9] J. M. Lee, *Introduction to Smooth Manifolds*, ser. Graduate Texts in Mathematics. New York, NY: Springer, 2003.
- [10] S. P. Ponomarev, "Submersions and preimages of sets of measure zero," *Siberian Mathematical Journal*, vol. 28, no. 1, pp. 153–163, Jan. 1987.
- [11] M. Šmíjeja and J. Tabor, "Entropy of the mixture of sources and entropy dimension," *IEEE Trans. Inf. Theory*, vol. 58, no. 5, pp. 2719–2728, May 2012.
- [12] R. J. Muirhead, *Aspects of Multivariate Statistical Theory*, ser. Wiley Series in Probability and Mathematical Statistics. Hoboken, NJ: Wiley Interscience, 1982.
- [13] B. C. Geiger and G. Kubin, "Signal enhancement as minimization of relevant information loss," May 2012, arXiv:1205.6935 [cs.IT].