Anomaly Detection based on Compressed Data: an Information Theoretic Characterization

This paper was downloaded from TechRxiv (<https://www.techrxiv.org>).

LICENSE

CC BY 4.0

SUBMISSION DATE / POSTED DATE

05-10-2021 / 12-10-2021

CITATION

Marchioni, Alex; Enttsel, Andriy; Mangia, Mauro; Rovatti, Riccardo; Setti, Gianluca (2021): Anomaly Detection based on Compressed Data: an Information Theoretic Characterization. TechRxiv. Preprint.
<https://doi.org/10.36227/techrxiv.16738171.v1>

DOI

[10.36227/techrxiv.16738171.v1](https://doi.org/10.36227/techrxiv.16738171.v1)

Anomaly Detection based on Compressed Data: an Information Theoretic Characterization

Alex Marchioni, *Student Member, IEEE*, Andriy Enttsel, Mauro Mangia, *Member, IEEE*,
Riccardo Rovatti, *Fellow, IEEE*, and Gianluca Setti, *Fellow, IEEE*

Abstract—We analyze the effect of lossy compression in the processing of sensor signals that must be used to detect anomalous events in the system under observation. The intuitive relationship between the quality loss at higher compression and the possibility of telling anomalous behaviours from normal ones is formalized in terms of information-theoretic quantities. Some analytic derivations are made within the Gaussian framework and possibly in the asymptotic regime for what concerns the stretch of signals considered.

Analytical conclusions are matched with the performance of practical detectors in a toy case allowing the assessment of different compression/detector configurations.

Index Terms—Outlier detection, lossy compression, rate-distortion curve Internet of Things, Edge computing.

I. INTRODUCTION

A typical scenario for nowadays massive acquisition systems can be modelled as a large number of sensing units, each transforming some physical unknown quantities into samples of random processes that are then transmitted over a network. To reduce transmission bitrate, signals are often compressed by a lossy mechanism that is theoretically capable of preserving useful information. Before reaching some cloud facility in which they will be ultimately stored or processed, the corresponding bitstreams may traverse several levels of hierarchical aggregation and intermediate devices that are often indicated as the *edge* of the cloud [1]. For latency or privacy reasons, some computational tasks may benefit from their deployment at the edge. One of those tasks is the detection of anomalies/novelty.

This is especially true when dealing, for example, with networks that sensorize plants or structures subject to monitoring as depicted in Fig. 1. The aggregated sensor readings may be processed in the cloud for off-line monitoring relying on long-term historical trends, while the outputs of subsets of sensors may be processed at the edge to give low-latency feedback on possible critical events that require immediate intervention.

A. Marchioni, A. Enttsel, M. Mangia and R. Rovatti are with the Department of Electrical, Electronic, and Information Engineering, University of Bologna, 40136 Bologna, Italy, and also with the Advanced Research Center on Electronic Systems, University of Bologna, 40125 Bologna, Italy (e-mail: alex.marchioni@unibo.it, andriy.enttsel@unibo.it, mauro.mangia@unibo.it, riccardo.rovatti@unibo.it).

R. Rovatti is also with the Alma Mater Research Institute for Human Centered AI University of Bologna, 40015 Bologna, Italy

G. Setti is with the Department of Electronics and Telecommunications, Politecnico di Torino, 10129 Torino, Italy, and also with the Advanced Research Center on Electronic Systems (ARCES), University of Bologna, 40125 Bologna, Italy (e-mail: gianluca.setti@polito.it).

Fig. 1. A sensorized plant whose acquisitions are aggregated at the edge before being sent to the cloud.

Usually, compression schemes applied to sensor data are asymmetric and entail a lightweight encoding performed on very-low complexity devices paired with a possibly expensive decoding stage running on the cloud. In these conditions, it is sensible that anomaly detectors work on compressed data and not on the recovered signal.

Yet, lossy compression bases its effectiveness on neglecting some of the signal details. This translates into a distortion between the original and the recovered signal but also in a loss of details that, in principle, could have been used to tell normal behaviours from anomalous ones.

This sets a trade-off between compression and distinguishability between normal and anomalous signals. Such a trade-off goes in parallel to the one between compression and distortion. Here, we analyze the former with the information theoretic machinery that leads also to the well-known rate-distortion analysis, and implicitly show that the two trade-offs are different.

In a sense, the approach we pursue is somehow similar to the information-bottleneck scheme [2], [3]. In that scheme, distortion is substituted by a very general criterion which identifies the features that should be preserved when compressing with the information that the original signal contains about a second (suitably introduced) signal. Yet, our discussion takes a different direction as, when dealing with anomaly/novelty/outlier, we may completely ignore the statistics of the anomaly and, even if we have priors on that, we need to be able to treat also cases in which the mutual information between normal

and anomalous signals is null.

For the same reason, the analysis we propose is also different from other modifications of classical rate-distortion theory that substitute energy-based distortion with perceptive criteria [4], [5].

Though not overlapping with the problem we address, it is also worthwhile to mention [6], [7], in which the authors assume that the original signal is characterized by some parameters (e.g., their mean) and study how the estimation of such parameters is affected by lossy compression.

Note also that other applications exist in which rate and distortion are paired with additional merit figures taking into account relevant features of the system. As an example [8] adds computational effort considerations to the analysis of rate-distortion of wavelet-based video coding.

Finally, even without emphasis on compression, the relationship between the analysis of suitably defined subcomponents of a signal to detect possible outlier behaviours is a classic theme that is still under investigation [9], [10], [11].

In this paper, we propose two information-theoretic measures of distinguishability to model the potential capabilities of two kinds of anomaly detectors, namely those that know the normal behaviour but are agnostic of the anomaly, and those that have information on both normal and anomalous behaviours. We then study how these measures behave when we change the distortion and thus the compression of the lossy mechanism, by deriving some analytical results when signals are Gaussian, both in the finite and asymptotic regime.

Section II gives the mathematical definition of the signal and quantities we use in the following with special emphasis on the functionals that quantify the effectiveness of compression, the distortion, and the distinguishability between normal and anomalous signals. Section III specializes the model to the case in which signals are Gaussian and revisits well-known results on Gaussian rate-distortion trade-off. Section IV applies the Gaussian assumption to specialize distinguishability both in the pointwise and average case, with an emphasis on the asymptotic characterization of anomalies in the high-dimensional case. Section V reports some numerical evidence analysing the behaviour of some suitably simplified anomaly detection strategies with respect to ideal and suboptimal compression strategies. Theoretical curves anticipate many aspects of practical performance trends and show that compression that optimizes the rate-distortion trade-off is not necessarily addressing at best the compromise with distinguishability. Some conclusion are finally drawn. Proofs of the properties stated in the discussion are reported in the Appendix.

II. MODEL DEFINITION

The signal flow we consider is reported in Fig. 2.

Normal and anomalous behaviours are modelled as two sources producing independent discrete-time, stationary, n -dimensional stochastic processes $x^{\text{ok}}[t] \in \mathbb{R}^n$ and $x^{\text{ko}}[t] \in \mathbb{R}^n$ with different PDFs $f^{\text{ok}} : \mathbb{R}^n \mapsto \mathbb{R}^+$ and $f^{\text{ko}} : \mathbb{R}^n \mapsto \mathbb{R}^+$. At any t , the observable discrete-time process $x[t] \in \mathbb{R}^n$ is either $x[t] = x^{\text{ok}}[t]$ or $x[t] = x^{\text{ko}}[t]$.

The observable $x[t]$ is passed to an encoding stage producing a compressed version $y[t]$ that may then be decompressed

Fig. 2. The signal chain we consider is tuned to the normal signal x^{ok} which is compressed into a smaller representation that can be re-expanded into a distorted version $\hat{x}^{\text{ok}}$ of the original signal. Sometimes, an anomalous signal x^{ko} is presented and the input of the chain. A detector should signal this based on the compressed version of the signal.

into $\hat{x}[t] \in \widehat{\mathbb{R}}^n$, where $\widehat{\mathbb{R}}^n \subset \mathbb{R}^n$ is a finite subset of $\mathbb{R}^n$. The compression mechanism is lossy as the encoding stage is not injective and thus introduces some distortion.

As discussed in the Introduction, the encoded signal $y[t]$ may be used for anomaly detection, i.e., to decide whether the original signal is $x[t] = x^{\text{ok}}[t]$ or $x[t] = x^{\text{ko}}[t]$. We assume the decoding stage to be injective so that, in abstract terms, the detector may be thought to work on $\hat{x}[t]$ even if in practical embodiments decoding is not performed at the edge.

Compression is designed assuming $x[t] = x^{\text{ok}}[t]$ and thus $\hat{x}[t] = \hat{x}^{\text{ok}}[t]$ for every t . It is characterized by the average distortion

$$\begin{aligned} D &= \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^{T-1} \|x^{\text{ok}}[t] - \hat{x}^{\text{ok}}[t]\|^2 \\ &= \mathbf{E} [\|x^{\text{ok}}[t] - \hat{x}^{\text{ok}}[t]\|^2] \end{aligned} \quad (1)$$

Since $\widehat{\mathbb{R}}^n$ is finite, a digital word can be assigned to each of its elements so that $y[t]$ is encoded with $b[t]$ bits. This allows to define the average rate of the stream as $R = \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=0}^{T-1} b[t] = \mathbf{E} [b[t]]$.

Rate and distortion are clearly two elements of a trade-off whose Pareto curve is the so called rate-distortion curve, i.e., the function

$$\rho(\delta) = \inf R \text{ s.t. } D \leq \delta$$

To identify such a function, it is classical [12, Chapter 9] to model the cascade of encoding and decoding as the conditional PDF $f_{\hat{x}|x} : \mathbb{R}^{2n} \mapsto \mathbb{R}^+$ such that the joint probability of $\hat{x}[t]$ and $x[t]$ is given by $f_{\hat{x},x}(\alpha, \beta) = f_{\hat{x}|x}(\alpha, \beta) f_x(\beta)$.

We design the system for normal cases $x[t] = x^{\text{ok}}[t]$, i.e., considering $f_{\hat{x},x}^{\text{ok}}(\alpha, \beta) = f_{\hat{x}|x}(\alpha, \beta) f_x^{\text{ok}}(\beta)$ and the marginal PDF $f_{\hat{x}}^{\text{ok}}(\alpha) = \int_{\mathbb{R}^n} f_{\hat{x},x}^{\text{ok}}(\alpha, \beta) d\beta$ with which distortion is expressed as

$$D = \int_{\mathbb{R}^{2n}} \|\alpha - \beta\|^2 f_{\hat{x},x}^{\text{ok}}(\alpha, \beta) d\alpha d\beta$$

With this, classical rate-distortion theory (see, e.g., [12, Chapter 13]) shows that

$$\rho(\delta) = \inf_{f_{\hat{x}|x}} I(\hat{x}^{\text{ok}}; x^{\text{ok}}) \quad \text{s.t. } D \leq \delta$$

where $I(\hat{x}^{\text{ok}}; \hat{x}^{\text{ok}})$ is the mutual information between $\hat{x}^{\text{ok}}[t]$ and $x^{\text{ok}}[t]$ [12, Chapter 8].

To proceed further it is convenient to define the functional

$$L(x'; x'') = - \int_{\mathbb{R}^n} f_{x'}(\alpha) \log_2 [f_{x''}(\alpha)] d\alpha$$

that is the average coding rate of a source characterized by the PDF $f_{x'}$ with a code optimized for a source with PDF $f_{x''}$, so that $H(x) = L(x; x)$ is equal to the differential entropy of x [12, Chapter 8].

When $n = 1$ and we deal with a quadratic distortion constraint, one may derive a well known Property [12, Chapter 13].

Property 1. *If $x^{\text{ok}}, \hat{x}^{\text{ok}} \in \mathbb{R}$ and we constrain $D \leq \delta$, then the minimum achievable rate has the lower bound*

$$I(\hat{x}^{\text{ok}}; x^{\text{ok}}) \geq H(x^{\text{ok}}) - \frac{1}{2} \log_2 (2\pi e \delta) \quad (2)$$

that can be achieved when the encoding is such that $x^{\text{ok}} - \hat{x}^{\text{ok}}$ is a Gaussian random variable.

Property 1 is the basis of classical developments and indicates that an encoding in which $x^{\text{ok}} - \hat{x}^{\text{ok}}$ distributes as a Gaussian has rate-minimization capabilities.

If anomalies need to be considered, the situation to model becomes more complex. In fact, as rate decreases, distortion increases, and the compressed versions of normal and anomalous signals tend to be less distinguishable. Hence, if the compressed stream $\hat{x}[t]$ is used for anomaly detection, one expects that detector performance degrades for increasing compression. To quantify this effect we need theoretical performance figures that characterize the distinguishability of two signals.

To begin tackling the problem, one can notice that anomalies are distorted with the encoder tuned on normal signals $f_{\hat{x}|x}$. This means that the input and the output of the distortion are characterized by the joint PDF $f_{\hat{x},x}^{\text{ko}}(\alpha, \beta) = f_{\hat{x}|x}(\alpha, \beta) f_x^{\text{ko}}(\beta)$ and the marginal PDF $f_{\hat{x}}^{\text{ko}}(\alpha) = \int_{\mathbb{R}^n} f_{\hat{x},x}^{\text{ko}}(\alpha, \beta) d\beta$.

Anomaly detectors work on the difference between $f_{\hat{x}}^{\text{ok}}$ and $f_{\hat{x}}^{\text{ko}}$ that we quantify with two kinds of information-theoretic measure, which model two distinct scenarios, i.e., one in which the detector knows both $f_{\hat{x}}^{\text{ok}}$ and $f_{\hat{x}}^{\text{ko}}$ and one in which it knows only $f_{\hat{x}}^{\text{ok}}$.

A. Distinguishability in anomaly-aware detection

When both $f_{\hat{x}}^{\text{ok}}$ and $f_{\hat{x}}^{\text{ko}}$ are known, it is most natural to measure their distinguishability as

$$\begin{aligned} \kappa &= L(\hat{x}^{\text{ko}}; \hat{x}^{\text{ok}}) - L(\hat{x}^{\text{ko}}; \hat{x}^{\text{ko}}) \\ &= \int_{\mathbb{R}^n} f_{\hat{x}}^{\text{ko}}(\alpha) \log_2 \left[\frac{f_{\hat{x}}^{\text{ko}}(\alpha)}{f_{\hat{x}}^{\text{ok}}(\alpha)} \right] d\alpha \end{aligned} \quad (3)$$

where one readily identify the latter expression as the Kullback-Leibler divergence of anomalies from normal signals.

This models a detector that knows the optimal codes for both the normal and the anomalous sources. It then observes the distorted stream $\hat{x}[t]$, encodes it with both codes and decides whether there is an anomaly or not depending on which of the two encoding rates is lower. As a result, large κ values correspond to system configurations with high detection capability.

B. Distinguishability in anomaly-agnostic detection

When only the statistical characterization of the normal source is known at the detector, we may measure the distinguishability between $f_{\hat{x}}^{\text{ok}}$ and $f_{\hat{x}}^{\text{ko}}$ with

$$\begin{aligned} \zeta &= L(\hat{x}^{\text{ko}}; \hat{x}^{\text{ok}}) - L(\hat{x}^{\text{ok}}; \hat{x}^{\text{ok}}) \\ &= \int_{\mathbb{R}^n} [f_{\hat{x}}^{\text{ko}}(\alpha) - f_{\hat{x}}^{\text{ok}}(\alpha)] \log_2 f_{\hat{x}}^{\text{ok}}(\alpha) d\alpha \end{aligned} \quad (4)$$

This models a detector using only the code optimized for the normal source. When anomalous signals appear, their encoding causes a difference in average rate that is used to reveal them. Hence, large ζ values stand for high distinguishability.

Unlike κ , ζ is not always positive since there may be anomalies whose encoding yields a lower rate with respect to normal signals. When a positive quantity is needed we will use $|\zeta|$.

Both in the anomaly-aware and in the anomaly-agnostic case, we may want to assess average performance when the anomalies are randomly drawn from a certain set of possible behaviours.

III. OPTIMUM RATE-DISTORTION TRADE-OFF IN THE GAUSSIAN FRAMEWORK

We assume that $x^{\text{ok}}[t]$ are n -dimensional Gaussian vectors such that if $t' \neq t''$, then $x^{\text{ok}}[t']$ and $x^{\text{ok}}[t'']$ are independent and identically distributed $x^{\text{ok}}[t] \sim \mathcal{G}(0, \Sigma^{\text{ok}})$, i.e., zero mean Gaussian vectors with covariance matrix $\Sigma^{\text{ok}} \in \mathbb{R}^{n \times n}$. Anomalous instances $x^{\text{ko}}[t]$ are also independent and identically distributed $x^{\text{ko}}[t] \sim \mathcal{G}(0, \Sigma^{\text{ko}})$.

In general, $\Sigma^{\text{ok}} \neq \Sigma^{\text{ko}}$, but we will assume $\text{tr}(\Sigma^{\text{ok}}) = \text{tr}(\Sigma^{\text{ko}}) = n$ meaning that, on the average, each sample in the vector contributes with a unit energy.

By applying a proper orthonormal change of coordinates to $x[t]$ (and thus to both $x^{\text{ok}}[t]$ and $x^{\text{ko}}[t]$) we may assume $\Sigma^{\text{ok}} = \text{diag}(\lambda_0^{\text{ok}}, \dots, \lambda_{n-1}^{\text{ok}})$ with $\lambda_0^{\text{ok}} \geq \lambda_1^{\text{ok}} \geq \dots \geq \lambda_{n-1}^{\text{ok}} \geq 0$ and $\sum_{j=0}^{n-1} \lambda_j^{\text{ok}} = n$.

These assumptions allow to derive an analytical assessment of the impact of compression on the anomaly detection capabilities. Such derivations rely on the function

$$G_{p,P}(\xi) = \frac{1}{\sqrt{(2\pi)^n \det P}} e^{-\frac{1}{2}(\xi-p)^T P^{-1}(\xi-p)}$$

of the vector ξ parameterized by $p \in \mathbb{R}^n$ and by $P \in \mathbb{R}^{n \times n}$, that is the PDF of a random vector $z \sim \mathcal{G}(p, P)$.

The Gaussian nature of $x[t]$ allows a well-known derivation [12, Chapter 13] of the rate-distortion function $\rho(\delta)$.

The basic building block comes from the application of Property 1 to a single Gaussian source, [12, Lemma 13.3.2]

Property 2. *If $x \sim \mathcal{G}(0, \sigma^2)$ then the minimum rate achievable given a chosen distortion is*

$$\rho(\delta) = \begin{cases} \frac{1}{2} \log_2 \frac{\sigma^2}{\delta} & \text{if } 0 \leq \delta \leq \sigma^2 \\ 0 & \text{otherwise} \end{cases}$$

When we deal with a vector of independent Gaussians, Property 2 must be paired with the fact that the total distortion is the sum of the distortions imposed to individual components. This leads to a water-filling result [12, Theorem 13.3.3] depending on a parameter $\theta \in [0, \lambda_0^{\text{ok}}]$

$$\rho = \frac{1}{2} \sum_{j=0}^{n-1} \log_2 \left(\frac{\lambda_j^{\text{ok}}}{\min\{\theta, \lambda_j^{\text{ok}}\}} \right) \quad (5)$$

$$\delta = \sum_{j=0}^{n-1} \min\{\theta, \lambda_j^{\text{ok}}\} \quad (6)$$

To track the effect of water filling we define I_n as the n -dimensional identity matrix and use the matrices:

- $T_\theta = \min\{I_n, \theta(\Sigma^{\text{ok}})^{-1}\}$, whose diagonal elements account for the fraction of energy cancelled by distortion along each component;
- $S_\theta = I_n - T_\theta = \max\{0, I_n - \theta(\Sigma^{\text{ok}})^{-1}\}$, whose diagonal elements account for the fraction of energy that survives distortion along each component;
- $R_\theta = S_\theta T_\theta$.

The trade-off between compression and distortion that is optimized along the rate-distortion curve is due to a proper distortion $f_{\hat{x}|x}$ that produces a signal with a PDF $f_{\hat{x}}^{\text{ok}}$. Both $f_{\hat{x}|x}$ and $f_{\hat{x}}^{\text{ok}}$ are given by the following Property, whose derivation is in the appendix.

Property 3. *The optimal distortion is given by*

$$f_{\hat{x}|x}(\alpha, \beta) = G_{\beta S_\theta, \Sigma^{\text{ok}} R_\theta}(\alpha) \quad (7)$$

and the optimally distorted signal has the PDF

$$\hat{x}^{\text{ok}} \sim \mathcal{G}(0, \Sigma^{\text{ok}} S_\theta) \quad (8)$$

The compression mechanism $f_{\hat{x}|x}$ that optimally addresses the rate-distortion trade-off for the normal source is used also on the anomalous instances. This produces $\hat{x}^{\text{ko}}[t]$ whose PDF is given by the following Property, whose proof is in the appendix.

Property 4. *If an anomalous source $x^{\text{ko}} \sim \mathcal{G}(0, \Sigma^{\text{ko}})$ is encoded with the compression scheme $f_{\hat{x}|x}$ of Property 3, then*

$$\hat{x}^{\text{ko}} \sim \mathcal{G}(0, S_\theta \Sigma^{\text{ko}} S_\theta + \theta S_\theta) \quad (9)$$

Such a result has two noteworthy corner cases.

- If $\theta \rightarrow 0^+$ there is no distortion. In fact, since $S_\theta = I_n$, Property 4 gives $\hat{x}^{\text{ko}} \sim x^{\text{ko}}$.

- If $x^{\text{ko}} \sim x^{\text{ok}}$ there is no anomaly, $\Sigma^{\text{ok}} = \Sigma^{\text{ko}}$, and

$$S_\theta \Sigma^{\text{ko}} S_\theta + \theta S_\theta = [S_\theta + \theta(\Sigma^{\text{ok}})^{-1}] \Sigma^{\text{ok}} S_\theta = \Sigma^{\text{ok}} S_\theta$$

where the last equality holds since $S_\theta = \max\{0, I_n - \theta(\Sigma^{\text{ok}})^{-1}\}$, the possible disagreements between $S_\theta + \theta(\Sigma^{\text{ok}})^{-1}$ and I_n correspond to components multiplied by zero by the last S_θ factor. Hence, Property 4 can be compared with Property 3 to confirm that $\hat{x}^{\text{ko}} \sim \hat{x}^{\text{ok}}$.

IV. DISTINGUISHABILITY IN THE GAUSSIAN FRAMEWORK

A. Pointwise distinguishability

Assume that, after compression, $\hat{x}^{\text{ok}}$ and $\hat{x}^{\text{ko}}$ are both a Gaussian n -dimensional random vectors whose first n_θ components survive distortion and are non-null, while the other $n - n_\theta$ components are set to 0 and thus cannot be used to tell anomalous from normal cases.

Say that the first n_θ components of x^{ok} are distributed according to $\mathcal{G}(0, \hat{\Sigma}_\theta^{\text{ok}})$ while the corresponding components of x^{ko} are distributed according to $\mathcal{G}(0, \hat{\Sigma}_\theta^{\text{ko}})$.

Within this Gaussian assumption, we may obtain an expression for the distinguishability measures (3) and (4). The starting point is the specialization of the building block $L(f'; f'')$ when f' and f'' are Gaussian as given by the following Property whose derivation is in the appendix.

Property 5. *If $f'(\alpha) = G_{0, P'}(\alpha)$ and $f''(\alpha) = G_{0, P''}(\alpha)$ then*

$$L(f'; f'') = \frac{1}{2 \ln 2} \{ \ln [(2\pi)^n |P''|] + \text{tr} [(P'')^{-1} P'] \} \quad (10)$$

where $|\cdot|$ indicates the determinant of its matrix argument.

By properly combining the definitions in (3) and (4) with (10) we obtain

$$\zeta = \frac{1}{2 \ln 2} \text{tr} [(\hat{\Sigma}_\theta^{\text{ok}})^{-1} \hat{\Sigma}_\theta^{\text{ko}} - I_{n_\theta}] \quad (11)$$

$$\kappa = \frac{1}{2 \ln 2} \left\{ \text{tr} [(\hat{\Sigma}_\theta^{\text{ok}})^{-1} \hat{\Sigma}_\theta^{\text{ko}} - I_{n_\theta}] - \ln \left| \frac{\hat{\Sigma}_\theta^{\text{ko}}}{\hat{\Sigma}_\theta^{\text{ok}}} \right| \right\} \quad (12)$$

in which distinguishability measures are given in bits.

Clearly, both figures vanish when $\hat{\Sigma}_\theta^{\text{ok}} = \hat{\Sigma}_\theta^{\text{ko}}$, i.e., when along the directions surviving distortion, the normal and the anomalous signals coincide. Note also that, with respect to $\hat{\Sigma}_\theta^{\text{ko}}$, ζ is linear while κ is convex [13, Chapter 3].

Properties 3 and 4 imply that when the normal and anomalous signals are Gaussian before compression, performance of anomaly detectors depend on how much we are capable of distinguishing between the two distributions in (8) and (9).

In this case, we have $n_\theta = \arg \max_j \{\lambda_j^{\text{ok}} > \theta\}$, $\hat{\Sigma}_\theta^{\text{ok}}$ is the $n_\theta \times n_\theta$ upper-left submatrix of $\Sigma^{\text{ok}} S_\theta$ in (8), and $\hat{\Sigma}_\theta^{\text{ko}}$ is the $n_\theta \times n_\theta$ upper-left submatrix of $S_\theta \Sigma^{\text{ko}} S_\theta + \theta S_\theta$ in (9).

By straightforward substitution and exploiting the diagonality of both Σ^{ok} and S_θ , we obtain that $(\hat{\Sigma}_\theta^{\text{ok}})^{-1} \hat{\Sigma}_\theta^{\text{ko}}$ is the $n_\theta \times n_\theta$ upper-left submatrix of $(\Sigma^{\text{ok}})^{-1} \Sigma^{\text{ko}} S_\theta + T_\theta$.

As a noteworthy particular case, when the normal signal is white, i.e., when $\Sigma^{\text{ok}} = I_n$, we have that $\theta \in [0, 1]$ and that for any $\theta < 1$, $T_\theta = \theta I_n$, $S_\theta = (1 - \theta)I_n$ and $n_\theta = n$. Hence, $\hat{\Sigma}_\theta^{\text{ok}} = (1 - \theta)I_n$, $\hat{\Sigma}_\theta^{\text{ko}} = (1 - \theta)^2 \Sigma^{\text{ko}} + \theta(1 - \theta)I_n$, and $(\hat{\Sigma}_\theta^{\text{ok}})^{-1} \hat{\Sigma}_\theta^{\text{ko}} = (1 - \theta) \Sigma^{\text{ko}} + \theta I_n$.

Hence, the distinguishability measures in (11) and (12) become

$$\zeta = 0 \quad (13)$$

$$\kappa = -\frac{1}{2 \ln 2} \ln |(1 - \theta) \Sigma^{\text{ko}} + \theta I_n| \quad (14)$$

Having $\zeta = 0$ in case of normal signal distributed as white Gaussian noise is not surprising since the distinguishability modelled by ζ depends only on the statistics of x^{ok} that has no exploitable structure.

B. Average on the set of possible anomalies

Since anomalies are zero-mean Gaussian vectors, they are completely defined by their covariance matrix Σ^{ko} .

We decompose $\Sigma^{\text{ko}} = U^{\text{ko}} \Lambda^{\text{ko}} U^{\text{ko}^\top}$ with $\Lambda^{\text{ko}} = \text{diag}(\lambda_0^{\text{ko}}, \dots, \lambda_{n-1}^{\text{ko}})$ and U^{ko} orthonormal.

The set of all possible λ^{ko} is

$$\mathbb{S}^n = \left\{ \lambda \in \mathbb{R}^{+n} \mid \sum_{j=0}^{n-1} \lambda_j = n \right\}$$

while the set of all possible U^{ko} is that of orthonormal matrices

$$\mathbb{O}^n = \{U \in \mathbb{R}^{n \times n} \mid U^\top U = I_n\}$$

By indicating with $\mathcal{U}(\cdot)$ the uniform distribution in the argument domain, we will assume that when λ^{ko} is not known then $\lambda^{\text{ko}} \sim \mathcal{U}(\mathbb{S}^n)$ and than when U^{ko} is not known then $U^{\text{ko}} \sim \mathcal{U}(\mathbb{O}^n)$, independently of λ^{ko} .

Note now that $\mathbb{S}^n$ is invariant with respect to any permutation of the λ_j . Since $\lambda^{\text{ko}} \sim \mathcal{U}(\mathbb{S}^n)$, also $\mathbf{E}[\lambda^{\text{ko}}]$ must be invariant with respect to the same permutations so that $\mathbf{E}[\lambda_j^{\text{ko}}] = \mathbf{E}[\lambda_k^{\text{ko}}]$ for any j, k . Since λ^{ok} has a constrained sum and is the diagonal of Λ^{ko} we have $\mathbf{E}[\Lambda^{\text{ko}}] = I_n$. This implies

$$\begin{aligned} \mathbf{E}[\Sigma^{\text{ko}}] &= \mathbf{E}[U^{\text{ko}} \Lambda^{\text{ko}} U^{\text{ko}^\top}] = \\ &= \mathbf{E}[U^{\text{ko}} \mathbf{E}[\Lambda^{\text{ko}}] U^{\text{ko}^\top}] = \mathbf{E}[U^{\text{ko}} U^{\text{ko}^\top}] = I_n \end{aligned} \quad (15)$$

Hence, in our setting, the average anomaly is white and we may compute the corresponding distinguishability measures $\bar{\zeta}$ and $\bar{\kappa}$. Note that, in this case, $\hat{\Sigma}_\theta$ is the $n_\theta \times n_\theta$ upper-left submatrix of $(\Sigma^{\text{ok}})^{-1} S_\theta + T_\theta$, i.e., the diagonal matrix in which the j -th diagonal entry is the sum of the j -th diagonal entry of $(\Sigma^{\text{ok}})^{-1} S_\theta$ and of the j -th diagonal entry of T_θ , i.e.,

$$u_{\theta,j} = \lambda_j^{\text{ok}-1} (1 - \theta \lambda_j^{\text{ok}-1}) + \theta \lambda_j^{\text{ok}-1}$$

With these quantities, the expressions of the distinguishability measures become

$$\bar{\zeta} = \frac{1}{2 \ln 2} \sum_{j=0}^{n_\theta-1} (u_{\theta,j} - 1) \quad (16)$$

$$\bar{\kappa} = \frac{1}{2 \ln 2} \sum_{j=0}^{n_\theta-1} (u_{\theta,j} - 1 - \ln u_{\theta,j}) \quad (17)$$

Note that due to the linearity of ζ and the convexity of κ we have $\bar{\zeta} = \mathbf{E}[\zeta]$ and $\bar{\kappa} \leq \mathbf{E}[\kappa]$.

Moreover, the very simple structure of $\bar{\zeta}$ allows the derivation of the following Property whose proof is in the appendix.

Property 6. *If $\bar{k} = \arg \max_k \{\lambda_k^{\text{ok}} \geq \lambda_k^{\text{ko}} = 1\}$, then $\bar{\zeta} = 0$ for at least one point $0 < \theta < \lambda_k^{\text{ok}}$*

Property 6 means that in case of x^{ko} distributed as Gaussian white noise, there exists at least a critical level of distortion that makes the detectors that do not use information of the anomaly ineffective.

C. Asymptotic distinguishability

White signals are not only the average anomalies but are also *typical* anomalies in the sense specified by the following Property whose proof is in the appendix.

Property 7. *If $\lambda^{\text{ko}} \sim \mathcal{U}(\mathbb{S}^n)$ and $U^{\text{ko}} \sim \mathcal{U}(\mathbb{O}^n)$ then, as $n \rightarrow \infty$, $\Sigma^{\text{ko}} = U^{\text{ko}} \text{diag}(\lambda_0^{\text{ko}}, \dots, \lambda_{n-1}^{\text{ko}}) U^{\text{ko}^\top}$ tends to I_n in probability.*

Hence, when n increases, most of the possible anomalies behave as white signals and $\bar{\zeta} \simeq \mathbf{E}[|\zeta|]$, that thus enjoys the properties shown in the previous subsection.

V. NUMERICAL EXAMPLES

In this section we match the theoretical derivations with the quantitative assessment of the performance of some practical anomaly detectors applied to compressed signals.

Normal signals are assumed to be $x^{\text{ok}} \sim \mathcal{G}(0, \Sigma^{\text{ok}})$ where $\Sigma_{j,k}^{\text{ok}} = \omega^{|j-k|}$ and ω is set to yield different degree of non-whiteness. Non-whiteness is measured with the so-called *localization* defined as

$$\mathcal{L}_{x^{\text{ok}}} = \frac{\text{tr}(\Sigma^{\text{ok}^2})}{\text{tr}^2(\Sigma^{\text{ok}})} - \frac{1}{n}$$

that goes from $\mathcal{L}_{x^{\text{ok}}} = 0$ when the signal is white to $\mathcal{L}_{x^{\text{ok}}} = 1 - \frac{1}{n}$ when all the energy is concentrated along a single direction of the signal space (see [14] for more details). To show the effect of realistic localization [15] we consider values of ω corresponding to $\mathcal{L}_{x^{\text{ok}}} \in \{0, 0.05, 0.2\}$.

Anomalous signals are generated as $x^{\text{ko}} \sim \mathcal{G}(0, \Sigma^{\text{ko}})$, where $\Sigma^{\text{ko}} = U^{\text{ko}} \Lambda^{\text{ko}} U^{\text{ko}^\top}$ is randomly picked according to the uniform distribution defined in IV-B.

To generate $\lambda^{\text{ko}} \sim \mathcal{U}(\mathbb{S}^n)$, we follow [16] to first draw $\xi_j \sim \mathcal{U}([0, 1])$ for $j = 0, \dots, n-1$ and then set

$$\lambda_j^{\text{ko}} = \frac{\log \xi_j}{\sum_{k=0}^{n-1} \log \xi_k}$$

To generate $U^{\text{ko}} \sim \mathcal{U}(\mathbb{O}^n)$, we follow [17] and start by generating a matrix A within the Ginibre ensemble [18], i.e.,

Fig. 3. Trend of $\Delta_2 = \frac{1}{n} \sqrt{\sum_{j,k=0}^{n-1} [\Sigma_{j,k}^{\text{ko}} - (I_n)_{j,k}]^2}$ and $\Delta_\infty = \max_{j,k} |\Sigma_{j,k}^{\text{ko}} - (I_n)_{j,k}|$ when n increases. Solid lines are median trends while shaded areas contain 98% of the population.

with independent entries $A_{j,k} \sim \mathcal{G}(0, 1)$ for $j, k = 0, \dots, n-1$. We then set U^{ko} to the orthonormal factor of the QR -decomposition of A .

A first use of this random sampling is the possibility of pairing Property 7 with some numerical evidence. Fig. 3 reports the vanishing trend of the average squared and uniform deviation from I_n of a population of uniformly distributed covariance matrices Σ^{ko} . Though not a theoretical result, note that empirical evidence supports a classical $1/\sqrt{n}$ convergence.

As far as detector assessment is concerned, we decide to set $n = 32$ and consider three compression techniques tuned to the normal signal and applied to both normal and anomalous instances. More specifically x is mapped to $\hat{x}$ by

- the minimum-rate-given-distortion compression in (7) (Rate-Distortion Compression RDC);
- projecting x along the subspace spanned by the eigenvectors of Σ^{ok} with the largest eigenvalues (Principal Component Compression PCC);
- a family of autoencoders [19, Chapter 14] with an increasingly deficient latent representation (Auto-Encoder Compression AEC). Assuming that ℓ is the dimensionality of the representation, the encoder is a neural network with fully connected layers of dimensions $n, 4n, 2n, \ell$, and the decoder is the dual network whose layers have dimensions $2n, 4n$, and n and the number of input is ℓ . The family of autoencoders is trained to minimize distortion computed as in (1). To smooth performance degradation we first train an autoencoder with $\ell = n - 1$. Then, the node of the latent representation along which we measure the least average energy is dropped to produce a smaller network with an $(\ell - 1)$ -dimensional latent space, that is re-trained using the previous weights as initialization. This is repeated decreasing ℓ and thus considering larger distortion values.

The three schemes we consider address in a different way the trade-off between compression and distortion. This can be shown by pairing each of them with a quantization stage ensuring that rate values are finite. In particular we encode each component of $\hat{x}$ with 16 bits and this yields rates of less than $16n = 512$ bits per time step. We assume that quantization is fine enough to substantially preserve the Gaussian distribution

Fig. 4. Rate distortion curves for the three compression schemes we consider and for different value of the localization of the original signal.

of $\hat{x}$ and thus evaluate the mutual information between x and $\hat{x}$ as if they were jointly gaussian with a covariance matrix that we estimate by Monte Carlo simulation [20]. Such estimation yields the rate-distortion curves in Fig. 4.

As expected, RDC yields the smallest rates while PCC gives the largest ones. Between the two we have AEC, whose performance depends on the effectiveness of the training.

The compressed version of the signal is then passed to a detector whose task is to compute a score such that high-score instances should be more likely to be anomalous. The final binary decision is taken by matching the score against a threshold.

We consider two detectors not relying on information of the anomaly

- a Likelihood Detector (LD) whose score is the inverse of the log-likelihood of the instance with respect to the normal signal distribution, so that to the instance x we associate the score $-\log f_x^{\text{ok}}(\hat{x})$;
- a One-Class Support-Vector Machine (OCSVM) [21] trained on a set of instances of normal signals contaminated by 1% of unlabelled white instances to help the algorithm finding the envelope of normal instances.

We also consider two detectors that are able to leverage information on the anomaly

- a Neyman-Pearson Detector (NPD) [22, Chapter 3], whose score is the difference between the log-likelihoods of the instance with respect to the normal and anomalous distributions, so that to instance x we associate the score $\log f_x^{\text{ko}}(\hat{x}) - \log f_x^{\text{ok}}(\hat{x})$;
- a Deep Neural Network (DNN) with 3 fully connected hidden layers with $n_\theta, 2n, n$ neurons with a ReLu activation and a final sigmoid neuron producing the score. The network is trained by backpropagation with a binary cross-entropy loss against a dataset containing labelled normal and anomalous instances.

LD and NPD detectors can be employed only on signals compressed by RDC or by PCC method since they rely on the statistical characterization of the signals that is not available after the nonlinear processing in AEC.

Table I shows how many different anomalies and how many signal instances are generated for the training (when needed)

 Fig. 5. Distinguishability measures ζ , κ and ψ against normalized distortion d in case of RDC.

 Fig. 6. Distinguishability measures ζ , κ and ψ against normalized distortion d in case of PCC.

detector	training			assessment		
	# Σ^{ko}	#instances $\times \Sigma^{\text{ko}}$ ok	ko	# Σ^{ko}	#instances $\times \Sigma^{\text{ko}}$ ok	ko
LD				10^3	10^3	10^3
OCSVM	1	99×10^3	10^3	10^3	10^3	10^3
NPD				10^3	10^3	10^3
DNN	50	10^5	10^5	50	10^5	10^5

TABLE I

NUMBER OF ANOMALIES (Σ^{ko}) AND, FOR EACH ANOMALY, THE NUMBER OF NORMAL (ok) AND ANOMALOUS (ko) SIGNAL INSTANCES USED IN THE TRAINING AND ASSESSMENT OF THE DETECTORS.

and for the assessment of the detectors. Note that in the DNN case we limited to 50 anomalies since the training process must be repeated for each of them.

To be independent of the choice of thresholds, detectors' performance is assessed by the Area-Under-the-Curve (AUC) methodology [23]. AUC estimates the probability that given a random normal instance and a random anomalous instance, the former has a lower score with respect to the latter, as it should be in an ideal setting. Hence, AUC is a positive performance index.

Clearly, detectors with $\text{AUC} = \frac{1}{2}$ are no better than coin tossing. Yet, if $\text{AUC} < \frac{1}{2}$, the score has some ability to distinguishing normal and anomalous signals if it is interpreted in a reverse way. Hence, it is convenient to set our *empirical distinguishability* measure to

$$\psi = \left| \text{AUC} - \frac{1}{2} \right|$$

to obtain a non-negative quantity that increases up to $\psi = \frac{1}{2}$ when the detector improves.

In the following the trends of ψ are reported and matched with the trends of $|\zeta|$ and κ to show how theoretical properties reflect on real cases. Comparisons must be partially qualitative as the range and the significance of ζ and κ are different from those of ψ and as the latter refers to practical detectors. All plots are made against a normalized distortion $d = D/n$ in the range $d \in [0, 0.64]$ as larger relative distortions are usually beyond operative ranges.

A. RDC

Fig. 5 summarizes the results we have in this case with two rows of 3 plots each. The upper row of plots correspond to detectors that do not exploit information on the anomaly, while the lower row of plots concerns detectors that may leverage information on the anomaly. Colors correspond to different $\mathcal{L}_{x^{\text{ok}}}$, dashed trends assume that the anomaly is the average one, i.e., white, and shaded areas show the span of 50% of the Monte Carlo population. The profiles of $|\zeta|$ and κ on the left shall be matched with the profiles on the right that correspond to the four detectors we consider. No $|\zeta|$ profile appears for $\mathcal{L}_x = 0$ as in that case $\zeta = 0$.

Theory anticipates that without any knowledge of the anomaly (upper row), a limited amount of distortion may cause distinguishability to vanish and thus detectors to fail. This actually happens for practical detectors such as LD and

OCSVM. The distortion level at which detectors fail is also anticipated by $|\zeta|$ and depends on $\mathcal{L}_{x^{\text{ok}}}$ as predicted by Property 6. Overall, theory anticipates that in the low-distortion region more localized signals are more distinguishable from anomaly though they cause detector failures at smaller distortions with respect to less localized signals.

Detectors leveraging the knowledge of the anomaly (lower row) fail completely only at 100% distortion as revealed by the abstract distinguishability measure κ . Also in this case, by comparing the trend of κ with the zoomed areas in the NPD and DNN plots we see how theory anticipates that in the low-distortion region more localized signals tend to be more distinguishable from anomalies but cause a more definite performance degradation of detectors when d increases.

B. PCC

From the point of view of the rate-distortion trade-off PCC is largely suboptimal. Yet, due to its linear nature x and $\hat{x}$ are still jointly Gaussian. This allows to compute the theoretical $|\zeta|$ and κ by means of (11) and (12).

Fig. 6 summarizes the results we have in this case with plots of the same kind of Fig. 5. The qualitative behaviours commented in the previous subsection appear in the new plots and are anticipated by the theoretical trends.

The distortion levels at which anomaly-agnostic detectors fail change with respect to the RDC case but are still anticipated by the theoretical curves and Property 6.

In this case, the values of $|\zeta|$ beyond breakdown distortion levels increase slightly more than in the optimal compression scenario. Hence, by adopting a compression strategy that is suboptimal in the rate-distortion sense one may obtain a better distinguishability of the compressed normal signal from the compressed anomalies. This is, in fact, what happens in practice as highlighted by the LD and OCSVM plots in the first row of Fig. 6.

C. AEC

In this case compression is non-linear so that x and $\hat{x}$ may not be jointly Gaussian. This prevents us from computing the theoretical curves $|\zeta|$ and κ and from applying LD and NPD that rely on the knowledge of the distribution of the signals.

For this reason, Fig. 7 reports only the performance of OCSVM and DNN detectors.

Notice how the qualitative trends of those performance still follow, though with a larger level of approximation, what is indicated by the theoretical curves for other compressors.

VI. CONCLUSION

Massive sensing systems may rely on lossy compression to reduce the bitrate needed to transmit acquisitions to the cloud while theoretically maintaining the important information. At some intermediate point along their path to centralized servers, compressed sensor readings may be processed for early detection of anomalies in the systems under observation. Such detection must be performed on compressed data.

Fig. 7. Distinguishability measure ψ against normalized distortion d in case of AEC.

We here analyze the trade-off between compression and performance of anomaly detectors modeling rate, distortion and distinguishability by means of information-theoretic quantities.

Such a trade-off is different from the well-known compression-distortion trade-off.

Some numerical examples in a toy case show that theoretical prediction are able to anticipate a number of features of the performance trends of real detectors.

APPENDIX

Proof of Property 3. Distortion is tuned to the normal case that entails a memoryless sources. Hence we may drop time indications and concentrate on a vector x^{ok} with independent components $x_j^{ok} \sim \mathcal{G}(0, \lambda_j^{ok})$ for $j = 0, \dots, n-1$.

We know from [24] that for a given value of the parameter θ , each component x_j^{ok} is transformed separately into $\hat{x}_j^{ok}$. In particular, if $\lambda_j^{ok} \leq \theta$ then x_j^{ok} is encoded into $\hat{x}_j^{ok} = 0$. If $\lambda_j^{ok} > \theta$ then x_j^{ok} is encoded into $\hat{x}_j^{ok} = x_j^{ok} + \Delta_j$ where, to achieve the Shannon lower bound, Δ_j must be an instance of a Gaussian random variable independent of $\hat{x}_j^{ok}$. Hence, the three quantities $\hat{x}_j^{ok}$, x_j^{ok} and Δ_j must be such that $(\hat{x}_j^{ok}, x_j^{ok}, \Delta_j)^\top \sim \mathcal{G}(0, \Sigma_{\hat{x}_j^{ok}, x_j^{ok}, \Delta_j}^{ok})$ with

$$\Sigma_{\hat{x}_j^{ok}, x_j^{ok}, \Delta_j}^{ok} = \begin{pmatrix} \lambda_j^{ok} - \theta & \lambda_j^{ok} - \theta & 0 \\ \lambda_j^{ok} - \theta & \lambda_j^{ok} & -\theta \\ 0 & -\theta & \theta \end{pmatrix} \quad (18)$$

That explains in which sense $\hat{x}_j^{ok}$ encodes x_j^{ok} . In fact, the non-diagonal elements $\lambda_j^{ok} - \theta$ are positive and thus $\hat{x}_j^{ok}$ and x_j^{ok} are positively correlated.

From (18), if we agree to identify a Gaussian with 0 variance with a Dirac's delta we infer that $\hat{x}_j^{ok} \sim \mathcal{G}(0, \max\{0, \lambda_j^{ok} - \theta\})$ and thus $\hat{x}_j^{ok} \sim \mathcal{G}(0, \Sigma_j^{ok} S_\theta)$.

Moreover, $(\hat{x}_j^{ok}, x_j^{ok})^\top \sim \mathcal{G}(0, \Sigma_{\hat{x}_j^{ok}, x_j^{ok}}^{ok})$ with $\Sigma_{\hat{x}_j^{ok}, x_j^{ok}}^{ok}$ the upper-left 2×2 submatrix of $\Sigma_{\hat{x}_j^{ok}, x_j^{ok}, \Delta_j}^{ok}$ in (18).

If we assume that $\theta < \lambda_j^{ok}$, from the joint probability of x_j^{ok} and $\hat{x}_j^{ok}$, we may compute the action of $f_{\hat{x}|x}$ on the j -th component of x_j^{ok} as the PDF of $\hat{x}_j^{ok}$ given x_j^{ok} , i.e.,

$$\begin{aligned} f_{\hat{x}_j|x_j}^{ok}(\alpha, \beta) &= \frac{f_{\hat{x}_j, x_j}^{ok}(\alpha, \beta)}{f_{x_j}^{ok}(\beta)} = \frac{G_{0, \Sigma_{\hat{x}_j, x_j}^{ok}}(\alpha)}{G_{0, \lambda_j^{ok}}(\beta)} \\ &= \frac{1}{\sqrt{2\pi \frac{\theta}{\lambda_j^{ok}} (\lambda_j^{ok} - \theta)}} e^{-\frac{1}{2} \frac{[\lambda_j^{ok} \alpha - (\lambda_j^{ok} - \theta) \beta]^2}{\lambda_j^{ok} \theta (\lambda_j^{ok} - \theta)}} \end{aligned}$$

If we define $\tau_j = \min\{1, \theta/\lambda_j^{ok}\}$ we have $0 \leq \tau_j \leq 1$ and

$$f_{\hat{x}_j|x_j}^{ok}(\alpha, \beta) = \frac{1}{\sqrt{2\pi \lambda_j^{ok} \tau_j (1 - \tau_j)}} e^{-\frac{1}{2} \frac{[\alpha - (1 - \tau_j) \beta]^2}{\lambda_j^{ok} \tau_j (1 - \tau_j)}}$$

that becomes $\delta(\alpha)$ for $\tau_j \rightarrow 1$ (maximum distortion of this component implies that the corresponding output is set to 0) and $\delta(\alpha - \beta)$ for $\tau_j \rightarrow 0$ (no distortion of this component, the output is equal to the input).

We may collect the component-wise PDFs into a vector PDF by using the matrix $T_\theta = \text{diag}(\tau_0, \dots, \tau_{n-1}) = \min\{I_n, \theta(\Sigma^{ok})^{-1}\}$, and the matrix $R_\theta = (I_n - T_\theta)T_\theta$ thus yielding the thesis. $\square$

Proof of Property 4. The PDF of $\hat{x}^{ko}$ distorted by means of $f_{\hat{x}|x}$ can be computed as

$$f_{\hat{x}}^{ko}(\alpha) = \int_{\mathbb{R}^n} f_{\hat{x}, x}^{ko}(\alpha, \beta) d\beta = \int_{\mathbb{R}^n} f_{\hat{x}|x}(\alpha, \beta) f_x^{ko}(\beta) d\beta$$

Assume first to be in the low-distortion condition $\theta < \lambda_{n-1}^{ok}$ that implies $T_\theta = \theta(\Sigma^{ok})^{-1}$, and write

$$\begin{aligned} f_{\hat{x}}^{ko}(\alpha) &= \int_{\mathbb{R}^n} G_{S_\theta \beta, \Sigma^{ok} R_\theta}(\alpha) G_{0, \Sigma^{ko}}(\beta) d\beta \\ &= G_{0, \Sigma^{ok} R_\theta}(\alpha) \times \\ &\quad \int_{\mathbb{R}^n} e^{-\frac{1}{2} [\beta^\top S_\theta (\Sigma^{ok} R_\theta)^{-1} S_\theta \beta - 2\alpha (\Sigma^{ok} R_\theta)^{-1} S_\theta \beta]} \times \\ &\quad G_{0, \Sigma^{ko}}(\beta) d\beta \\ &= G_{0, \Sigma^{ok} R_\theta}(\alpha) \times \\ &\quad \frac{1}{\sqrt{(2\pi)^n \det \Sigma^{ko}}} \underbrace{\int_{\mathbb{R}^n} e^{-\frac{1}{2} (\beta^\top Q \beta - 2q^\top \beta)} d\beta}_{g(\alpha)} \end{aligned}$$

with $Q = S_\theta (\Sigma^{ok} R_\theta)^{-1} S_\theta + (\Sigma^{ko})^{-1} = (\theta I_n)^{-1} - (\Sigma^{ok})^{-1} + (\Sigma^{ko})^{-1}$ and $q = (\Sigma^{ok} R_\theta)^{-1} S_\theta \alpha = \alpha/\theta$. To compute $g(\alpha)$ let $Q = UDU^\top$ with D diagonal and U orthonormal, and set $\beta' = D^{1/2} U^\top \beta$ so that $\beta = UD^{-1/2} \beta'$ and $d\beta = d\beta'/\sqrt{\det Q}$. With this write

$$g(\alpha) = \frac{1}{\sqrt{\det Q}} \int_{\mathbb{R}^n} e^{-\frac{1}{2}(\beta'^\top \beta' - 2q^\top U D^{-1/2} \beta')} d\beta'$$

at the exponent of which one may add and subtract $q^\top Q^{-1} q = q^\top U D^{-1/2} D^{-1/2} U^\top q$ to yield

$$\begin{aligned} g(\alpha) &= \frac{1}{\sqrt{\det Q}} \int_{\mathbb{R}^n} e^{-\frac{1}{2}(\|\beta' - D^{-1/2} U^\top q\|^2 - q^\top Q^{-1} q)} d\beta' \\ &= \sqrt{\frac{(2\pi)^n}{\det Q}} e^{\frac{1}{2} q^\top Q^{-1} q} \end{aligned}$$

Putting this back into $f_{\hat{x}}^{\text{ok}}$ we get

$$f_{\hat{x}}^{\text{ko}}(\alpha) = G_{0,[(\theta I_n)^{-1} - (\Sigma^{\text{ok}})^{-1} + (\Sigma^{\text{ko}})^{-1}] \Sigma^{\text{ko}} \Sigma^{\text{ok}} R_\theta}(\alpha)$$

A straightforward expansion of the definitions under the low-distortion assumption finally rearranges the covariance matrix into

$$\begin{aligned} &[(\theta I)^{-1} - (\Sigma^{\text{ok}})^{-1} + (\Sigma^{\text{ko}})^{-1}] \Sigma^{\text{ko}} \Sigma^{\text{ok}} R_\theta = \\ &= [(\theta I_n)^{-1} - (\Sigma^{\text{ok}})^{-1} + (\Sigma^{\text{ko}})^{-1}] \Sigma^{\text{ko}} \Sigma^{\text{ok}} \theta (\Sigma^{\text{ok}})^{-1} S_\theta \\ &= [\Sigma^{\text{ko}} - \theta (\Sigma^{\text{ok}})^{-1} \Sigma^{\text{ko}} + \theta I_n] S_\theta \\ &= [I_n - \theta (\Sigma^{\text{ok}})^{-1}] \Sigma^{\text{ko}} S_\theta + \theta S_\theta \\ &= S_\theta \Sigma^{\text{ko}} S_\theta + \theta S_\theta \end{aligned} \quad (19)$$

as in the statement of the Property.

To address the case in which θ exceeds $\lambda_{n-1}^{\text{ok}}$ note that for $\theta \rightarrow (\lambda_{n-1}^{\text{ok}})^-$, the last diagonal entry of S_θ tends to 0 and thus by (19) the covariance tends to have zeros in its last row and column. Since a Gaussian with vanishing-variance can be considered Dirac's delta, this model the fact that the last component of both x^{ok} and x^{ko} is fully distorted and set to 0. With this, (19), is valid also for $\lambda_{n-1}^{\text{ok}} < \theta < \lambda_{n-2}^{\text{ok}}$. Yet, analogous considerations can be carried out for $\theta \rightarrow (\lambda_j^{\text{ok}})^-$ and $j = n-2, n-3, \dots, 0$ so that (19) is valid for any value of θ . $\square$

Proof of Property 5.

$$\begin{aligned} L(f'; f'') &= - \int_{\mathbb{R}^n} G_{0,P'}(\alpha) \log_2 [G_{0,P''}(\alpha)] d\alpha \\ &= \frac{1}{2} \log_2 [(2\pi)^n |P''|] \int_{\mathbb{R}^n} G_{0,P'}(\alpha) d\alpha \\ &\quad + \frac{1}{2 \ln 2} \int_{\mathbb{R}^n} \alpha^\top (P'')^{-1} \alpha G_{0,P'}(\alpha) d\alpha \\ &= \frac{1}{2} \log_2 [(2\pi)^n |P''|] + \frac{1}{2 \ln 2} \text{tr} [(P'')^{-1} P'] \end{aligned}$$

Where the last summand has been computed as the expectation of a quadratic form in a Gaussian multivariate for which Corollary 3.2b.1 in [25, chapter 3] gives a formula. $\square$

Proof of Property 6. From (16) we have that for, n large, $\bar{\zeta} \simeq \frac{1}{2 \ln 2} |Z|$ with $Z = \sum_{j=0}^{n_\theta-1} \alpha_j(\theta)$ and

$$\alpha_j(\theta) = 1 - \frac{1}{\lambda_j^{\text{ok}}} \left(1 - \frac{\theta}{\lambda_j^{\text{ok}}} \right) - \frac{\theta}{\lambda_j^{\text{ok}}}$$

Note that $\alpha_j(\theta)$ is continuous and its derivative is $\frac{\partial}{\partial \theta} \alpha_j = (1/\lambda_j^{\text{ok}} - 1)/\lambda_j^{\text{ok}}$.

For simplicity's sake assume $\lambda_0^{\text{ok}} > \lambda_1^{\text{ok}} > \dots > \lambda_{n-1}^{\text{ok}} > 0$, set $\lambda_n^{\text{ok}} = 0$, and define $\Theta_j =]\lambda_{j+1}^{\text{ok}}, \lambda_j^{\text{ok}}[$ for $j = 0, \dots, n-1$ so that if $\theta \in \Theta_j$ then $n_\theta = j+1$.

As a function of θ , Z is continuous. In fact, it is trivially continuous in each Θ_j . Yet, it is continuous also at any chosen $\lambda_{\bar{j}}^{\text{ok}}$ with $\bar{j} = 0, \dots, n-1$. To see why, note that

$$\begin{aligned} \lim_{\theta \rightarrow \lambda_{\bar{j}}^{\text{ok}}-} Z &= \lim_{\theta \rightarrow \lambda_{\bar{j}}^{\text{ok}}-} \sum_{j=0}^{\bar{j}} \alpha_j(\theta) \\ &= \lim_{\theta \rightarrow \lambda_{\bar{j}}^{\text{ok}}-} \alpha_{\bar{j}}(\theta) + \sum_{j=0}^{\bar{j}-1} \alpha_j(\theta) \\ &= \lim_{\theta \rightarrow \lambda_{\bar{j}}^{\text{ok}}+} \sum_{j=0}^{\bar{j}-1} \alpha_j(\theta) = \lim_{\theta \rightarrow \lambda_{\bar{j}}^{\text{ok}}+} Z \end{aligned}$$

where we have exploited the fact that the $\alpha_j(\theta)$ are continuous and thus their left and right limits coincide, and that $\alpha_{\bar{j}}(\lambda_{\bar{j}}^{\text{ok}}) = 0$.

On the left-hand side of its domain, When $\theta = \lambda_n^{\text{ok}} = 0$ (no distortion), we have $n_\theta = n$ and thus

$$Z = \sum_{j=0}^{n-1} 1 - \frac{1}{\lambda_j^{\text{ok}}} \leq 0$$

where the last inequality follows from the fact that $\sum_{j=0}^n \lambda_j^{\text{ok}} = n$ and thus $\sum_{j=0}^n 1/\lambda_j^{\text{ok}} \geq n$.

On the right-hand side of its domain, when $\theta = \lambda_0^{\text{ok}}$ (maximum distortion), we have $n_\theta = 0$ and thus $Z = 0$. Yet, we also have that

$$\frac{\partial}{\partial \theta} Z = \sum_{j=0}^{n_\theta-1} \left(\frac{1}{\lambda_j^{\text{ok}}} - 1 \right) \frac{1}{\lambda_j^{\text{ok}}}$$

in which the summands are negative if $\lambda_j^{\text{ok}} > 1$. Hence, if $\bar{k} = \arg \max_k \{\lambda_k^{\text{ok}} \geq 1\}$, for $\theta \geq \lambda_{\bar{k}}^{\text{ok}}$, all the summands in the above expression are negative and thus $\frac{\partial}{\partial \theta} Z < 0$ for $\lambda_{\bar{k}}^{\text{ok}} < \theta \leq \lambda_0^{\text{ok}}$.

Since $Z = 0$ at the end of that interval, it must be positive in its interior.

Since we know that Z is negative for $\theta = \lambda_n^{\text{ok}}$ and it is continuous for $\theta \in]\lambda_n^{\text{ok}}, \lambda_0^{\text{ok}}[$, it must pass through zero at least once whenever it is not positive, i.e., for $0 < \theta < \lambda_{\bar{k}}^{\text{ok}}$. $\square$

Proof of Property 7. We will use the following Lemma whose proof follows this one.

Lemma 1. *If $\lambda^{\text{ko}} \sim \mathcal{U}(\mathbb{S}^n)$, then for any integrable function $f: \mathbb{R} \rightarrow \mathbb{R}$ and any $j = 0, \dots, n-1$*

$$\mathbf{E} [f(\lambda_j^{\text{ko}})] = \frac{n-1}{n^{n-1}} \int_0^n f(p)(n-p)^{n-2} dp$$

From [26] we know that if Q is uniformly distributed in $\mathbb{O}^n$ (i.e., if it is distributed according to the Haar measure

on the orthogonal group) then, for any sequence of integers $M_n < n$ increasing with n but such that $M_n = o\left(\frac{n}{\log n}\right)$, the entries of the first M_n columns of Q converge in probability to independent random variables such that $\sqrt{n}Q_{j,k} \sim \mathcal{G}(0, 1)$ for $j = 0, \dots, n-1$, $k = 0, \dots, M_n-1$.

Such a property can be extended to any subset of M_n columns. In fact, given any subset of m_n columns of Q , there is a permutation matrix P such that QP has such columns as the first ones.

Yet, since $P \in \mathbb{O}^n$ and Q is distributed according to the Haar measure in that group, also QP is distributed according to that measure and the entries in those columns tend to independent Gaussians.

Divide now n by M_n as in $n = N_n M_n + m_n$, where N_n is the quotient and $0 \leq m_n < M_n$ is the remainder. We can look at Q as the concatenation of N_n matrices Q_i , for $i = 0, \dots, N_n-1$, each $n \times M_n$, and of a last matrix Q_{N_n} that is $n \times m_n$. From $M_n = o\left(\frac{n}{\log n}\right)$ we have $N_n = O(\log n)$ but we can choose M_n such that $N_n = o(n^\alpha)$ for any $\alpha > 0$.

If we set $U^{\text{ko}} = Q$ then, $\Sigma^{\text{ko}} = U^{\text{ko}} \Lambda^{\text{ko}} U^{\text{ko}\top}$ can be written componentwise

$$\begin{aligned} \Sigma_{j,k}^{\text{ok}} &= \sum_{l=0}^{n-1} U_{j,l}^{\text{ko}} \lambda_l^{\text{ko}} U_{k,l}^{\text{ko}} = \\ &= \underbrace{\sum_{i=0}^{N_n-1} \sum_{l=iM_n}^{(i+1)M_n-1} \frac{1}{n} \nu_{j,l} \lambda_l^{\text{ko}} \nu_{k,l}}_{W_i} + \underbrace{\sum_{l=n-m_n}^{n-1} \frac{1}{n} \nu_{j,l} \lambda_l^{\text{ko}} \nu_{k,l}}_{W_{N_n}} \end{aligned}$$

$i = 0, \dots, N_n-1$ scans the first submatrices Q_i , the last summand accounts for the *remainder* matrix Q_{N_n} , and, thanks to the above considerations, $\nu_{j,l} \sim \nu_{k,l} \sim \mathcal{G}(0, 1)$ for all j, k, l .

We now have to address that case $j \neq k$ and the case $j = k$ separately.

A. Asymptotics of $\Sigma_{j,k}^{\text{ko}}$ for $j \neq k$

Let us now consider W_0 as the representative of all other W_i for $i = 0, \dots, N_n-1$, written as $W_0 = \sum_{l=0}^{M_n-1} X_{n,l}$ with $X_{n,l} = \frac{1}{n} \nu_{j,l} \lambda_l^{\text{ko}} \nu_{k,l}$. All the normal random variables involved in such a sum are asymptotically independent but this is not true for the λ_l^{ko} since the eigenvalues are constrained to sum to n .

Hence, W_0 is a triangular array of row-dependent random variables whose asymptotic behaviour can be analyzed by means of [27, Theorem 2.1] that is essentially a Lindeberg-Feller Central Limit Theorem with the row-wise independence relaxed to asymptotic row-wise incorrelation. To analyze the asymptotics of W_0 we note that

$$\mathbf{E}[X_{n,l}] = \mathbf{E}[\lambda_l^{\text{ko}}] \mathbf{E}[\nu_{j,l}] \mathbf{E}[\nu_{k,l}] = 0$$

Note also that though not independent, the covariance and the correlation between $X_{n,l'}$ and $X_{n,l''}$ is

$$\begin{aligned} \mathbf{E}[X_{n,l'} X_{n,l''}] &= \\ &= \frac{1}{n^2} \mathbf{E}[\lambda_{l'}^{\text{ko}} \lambda_{l''}^{\text{ko}}] \mathbf{E}[\nu_{j,l'} \nu_{j,l''}] \mathbf{E}[\nu_{k,l'} \nu_{k,l''}] \\ &= \frac{1}{n^2} \begin{cases} 0 & \text{if } l' \neq l'' \\ \mathbf{E}[(\lambda_l^{\text{ko}})^2] & \text{if } l' = l'' = l \end{cases} \end{aligned}$$

With this we also know that

$$\begin{aligned} \sigma_{W_0}^2 &= \mathbf{E}[W_0^2] = \sum_{l'=0}^{M_n-1} \sum_{l''=0}^{M_n-1} \mathbf{E}[X_{n,l'} X_{n,l''}] = \\ &= \sum_{l=0}^{M_n-1} \mathbf{E}[X_{n,l}^2] = \frac{1}{n^2} \sum_{l=0}^{M_n-1} \mathbf{E}[(\lambda_l^{\text{ko}})^2] \end{aligned}$$

To compute the last expectation we may resort to Lemma 1 that gives

$$\mathbf{E}[(\lambda_l^{\text{ko}})^2] = \frac{n-1}{n^{n-1}} \int_0^n p^2 (n-p)^{n-2} dp = \frac{2n}{n+1}$$

Hence we have $\sigma_{W_0}^2 = \frac{2M_n}{n(n+1)} \rightarrow 0$ for $n \rightarrow \infty$.

This helps satisfying the Lindeberg condition since, if for a given $\epsilon > 0$ we indicate with $\mathbf{E}[X_{n,l}^2 | |X_{n,l}| \geq \epsilon]$ the expectation of $X_{n,l}^2$ restricted to its values that are not less than ϵ in modulus, then

$$\sum_{l=0}^{n-1} \mathbf{E}[X_{n,l}^2 | |X_{n,l}| \geq \epsilon] \leq \sum_{l=0}^{n-1} \mathbf{E}[X_{n,l}^2] = \sigma_{W_0}^2 = \frac{2M_n}{n(n+1)}$$

that vanishes asymptotically.

Finally, let $L_{M_n}, R_{M_n} \subset \{0, \dots, M_n-1\}$ be two index subsets such that $L_{M_n} \cap R_{M_n} = \emptyset$. If $g_{L_{M_n}}$ is any function of the random variables $X_{n,l}$ with $l \in L_{M_n}$ and $h_{R_{M_n}} = \prod_{l \in R_n} X_{n,l}$, the covariance between $g_{L_{M_n}}$ and $h_{R_{M_n}}$ is

$$\begin{aligned} \mathbf{E}[g_{L_{M_n}} h_{R_{M_n}}] &= \mathbf{E}\left[g_{L_{M_n}} \prod_{l \in R_n} \frac{1}{n} \lambda_l^{\text{ko}} \nu_{j,l} \nu_{k,l}\right] \\ &= \mathbf{E}\left[g_{L_{M_n}} \prod_{l \in R_n} \lambda_l^{\text{ko}}\right] \prod_{l \in R_n} \frac{1}{n} \mathbf{E}[\nu_{j,l}] \mathbf{E}[\nu_{k,l}] = 0 \end{aligned}$$

that is enough to satisfy the assumptions in equation (2.3) and in equation (2.4) of [27, Theorem 2.1]. From that Theorem, we finally now that $W_0 \sim \mathcal{G}\left(0, \frac{2M_n}{n(n+1)}\right)$ when $n \rightarrow \infty$. Clearly, the same happens to any W_i for $i = 0, \dots, N_n-1$, while $W_{N_n} \sim \mathcal{G}\left(0, \frac{2m_n}{n(n+1)}\right)$.

Let us now consider

$$\begin{aligned}
& \Pr \{ |\Sigma_{j,k}^{\text{ko}}| \leq \epsilon \} = \\
& = \Pr \left\{ \left| \sum_{i=0}^{N_n-1} W_i + W_{N_n} \right| \leq \epsilon \right\} \\
& \geq \Pr \left\{ \sum_{i=0}^{N_n-1} |W_i| + |W_{N_n}| \leq \epsilon \right\} \\
& \geq \Pr \left\{ \max \{ |W_0|, \dots, |W_{N_n}| \} \leq \frac{\epsilon}{N_n} \right\} \\
& = 1 - \Pr \left\{ \max \{ |W_0|, \dots, |W_{N_n}| \} > \frac{\epsilon}{N_n} \right\}
\end{aligned}$$

Yet

$$\begin{aligned}
& \Pr \left\{ \max \{ |W_0|, \dots, |W_{N_n}| \} > \frac{\epsilon}{N_n} \right\} \\
& = \Pr \left\{ |W_0| > \frac{\epsilon}{N_n} \vee \dots \vee |W_{N_n}| > \frac{\epsilon}{N_n} \right\} \\
& \leq N_n \Pr \left\{ |W_0| > \frac{\epsilon}{N_n} \right\}
\end{aligned}$$

so that

$$\Pr \{ |\Sigma_{j,k}^{\text{ko}}| \leq \epsilon \} \geq 1 - N_n \Pr \left\{ |W_0| > \frac{\epsilon}{N_n} \right\}$$

From this and from the asymptotic normality of W_0 we may say

$$\Pr \{ |\Sigma_{j,k}^{\text{ko}}| \leq \epsilon \} \geq 1 - \frac{N_n}{2} \text{erfc} \left(\frac{\epsilon}{2\sqrt{\frac{N_n^2 M_n}{n(n+1)}}} \right) \quad (20)$$

in which the probability tends to 1 for $n \rightarrow \infty$, and $M_n = o\left(\frac{n}{\log n}\right)$ chosen to have $N_n = o(n^{1/2})$.

B. Asymptotics of $\Sigma_{j,k}^{\text{ko}}$ for $j = k$

In this case, exploiting the fact that $\sum_{j=0}^{n-1} \lambda_j^{\text{ko}} = n$, we may write

$$\begin{aligned}
\Sigma_{j,j}^{\text{ok}} &= 1 + \sum_{i=0}^{N_n-1} \underbrace{\sum_{l=iM_n}^{(i+1)M_n-1} \frac{1}{n} (\nu_{j,l}^2 - 1) \lambda_l^{\text{ko}}}_{W_i} + \\
&\quad \underbrace{\sum_{l=n-m_n}^{n-1} \frac{1}{n} (\nu_{j,l}^2 - 1) \lambda_l^{\text{ko}}}_{W_{N_n}}
\end{aligned}$$

in which the summands W_i are a triangular arrays of elements with features similar to the previous ones.

In fact, we may focus on $W_0 = \sum_{l=0}^{M_n-1} X_{n,l}$ with $X_{n,l} = \frac{1}{n} \lambda_l^{\text{ko}} (\nu_{j,l}^2 - 1)$ and note that $\mathbf{E}[X_{n,l}] = 0$ and

$$\begin{aligned}
\mathbf{E}[X_{n,l'} X_{n,l''}] &= \\
&= \frac{1}{n^2} \mathbf{E}[\lambda_{l'}^{\text{ko}} \lambda_{l''}^{\text{ko}}] \mathbf{E}[(\nu_{j,l'}^2 - 1)(\nu_{j,l''}^2 - 1)] \\
&= \frac{2}{n^2} \begin{cases} 0 & \text{if } l' \neq l'' \\ \mathbf{E}[(\lambda_l^{\text{ko}})^2] & \text{if } l' = l'' = l \end{cases}
\end{aligned}$$

where we have exploited the fact that $\mathbf{E}[(\nu_{j,l}^2 - 1)^2] = \mathbf{E}[\nu_{j,l}^4] - 2\mathbf{E}[\nu_{j,l}^2] + 1 = 2$. With this, $\sigma_{W_0}^2 = \frac{4M_n}{n(n+1)} \rightarrow 0$ for $n \rightarrow \infty$.

As before, this makes the Lindeberg condition automatically satisfied and is also enough to satisfy the covariance constraints in equations (2.3) and (2.4) of [27, Theorem 2.1] from which we get that $W_0 \sim \mathcal{G}\left(0, \frac{4M_n}{n(n+1)}\right)$ for $n \rightarrow \infty$. An analogous path leads to the asymptotic behaviour $W_{N_n} \sim \mathcal{G}\left(0, \frac{4m_n}{n(n+1)}\right)$.

The same inequalities as before lead to

$$\Pr \{ |\Sigma_{j,j}^{\text{ko}} - 1| \leq \epsilon \} \geq 1 - \frac{N_n}{2} \text{erfc} \left(\frac{\epsilon}{4\sqrt{\frac{N_n^2 M_n}{n(n+1)}}} \right) \quad (21)$$

in which the probability tends to 1 for $n \rightarrow \infty$, and $M_n = o\left(\frac{n}{\log n}\right)$ chosen to have $N_n = o(n^{1/2})$.

From (20) and (21) we finally get that U^{ko} tends to I_n in probability as $n \rightarrow \infty$. $\square$

Proof of Lemma 1. For any function $f: \mathbb{R} \mapsto \mathbb{R}$ we have

$$\begin{aligned}
\mathcal{I}[f(p)] &= \int_{\mathbb{S}^n} f(p_0) dp_0 \dots dp_{n-1} \\
&= \int_0^n f(p_0) \int_0^{n-p_0} \int_0^{n-p_0-p_1} \dots \int_0^{n-p_0-p_1-\dots-p_{n-3}} dp_0 \dots dp_{n-2} \\
&= \int_0^n f(p_0) \frac{(n-p_0)^{n-2}}{(n-2)!} dp_0
\end{aligned}$$

Since λ^{ko} is uniformly distributed over $\mathbb{S}^n$ the probability density is the constant $1/\mathcal{I}[1] = n^{-(n-1)}(n-1)!$ and the expectation of f is

$$\begin{aligned}
\mathbf{E}[f(\lambda_j^{\text{ok}})] &= n^{-(n-1)}(n-1)! \mathcal{I}[f(p)] \\
&= \frac{(n-1)!}{n^{n-1}} \int_0^n f(p) \frac{(n-p)^{n-2}}{(n-2)!} dp \\
&= \frac{n-1}{n^{n-1}} \int_0^n f(p) (n-p)^{n-2} dp
\end{aligned}$$

$\square$

REFERENCES

- [1] W. Shi, J. Cao, Q. Zhang, Y. Li, and L. Xu, "Edge computing: Vision and challenges," *IEEE Internet of Things Journal*, vol. 3, no. 5, pp. 637–646, Oct 2016.
- [2] N. Tishby, F. C. Pereira, and W. Bialek, "The information bottleneck method," *Proceedings of the 37-th Annual Allerton Conference on Communication, Control and Computing*, 1999, pp. 368–377.
- [3] N. Slonim, N. Tishby, "Agglomerative information bottleneck", In *Proceedings of the 12th International Conference on Neural Information Processing Systems (NIPS'99)*. MIT Press, Cambridge, MA, USA, 1999, pp. 617–623.

- [4] Y. Blau, and T. Michaeli, "The Perception-Distortion Tradeoff," 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2018, pp. 6228-6237, doi: 10.1109/CVPR.2018.00652.
- [5] Y. Blau, and T. Michaeli, "Rethinking Lossy Compression: The Rate-Distortion-Perception Tradeoff," Proceedings of the 36th International Conference on Machine Learning, (ICML), vol. 97, pp. 675-685, 2019.
- [6] M. R. D. Rodrigues, N. Deligiannis, L. Lai, and Y. C. Eldar, "Rate-distortion trade-offs in acquisition of signal parameters," 2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), 2017, pp. 6105-6109, doi: 10.1109/ICASSP.2017.7953329.
- [7] N. Shlezinger, Y. C. Eldar and M. R. D. Rodrigues, "Hardware-Limited Task-Based Quantization," IEEE Transactions on Signal Processing, vol. 67, no. 20, pp. 5223-5238, 15 Oct.15, 2019, doi: 10.1109/TSP.2019.2935864.
- [8] B. Foo, Y. Andreopoulos, M. van der Schaar, "Analytical Rate-Distortion-Complexity Modeling of Wavelet-Based Video Coders," IEEE Transactions on Signal Processing, vol. 56, no. 2, pp. 797-815, 2008.
- [9] V. Menon and S. Kalyani, "Structured and Unstructured Outlier Identification for Robust PCA: A Fast Parameter Free Algorithm," IEEE Transactions on Signal Processing, vol. 67, no. 9, pp. 2439-2452, May 2019.
- [10] L. Kong, C. Qi and H. Qi, "Classical Multidimensional Scaling: A Subspace Perspective, Over-Denoising, and Outlier Detection," IEEE Transactions on Signal Processing, vol. 67, no. 14, pp. 3842-3857, July, 2019.
- [11] M. Rahmani and P. Li, "Closed-Form, Provable, and Robust PCA via Leverage Statistics and Innovation Search," IEEE Transactions on Signal Processing, vol. 69, pp. 3132-3144, 2021.
- [12] T.M. Cover, J.A. Thomas, *Elements of Information Theory*, John Wiley & Sons, 1991
- [13] S. Boyd, L. Vandenberghe, *Convex Optimization* Cambridge University Press, 2006
- [14] M. Mangia, F. Pareschi, V. Cambareri, R. Rovatti, and G. Setti, "Rakeness-based design of low-complexity compressed sensing," IEEE Transactions on Circuits and Systems I: Regular Papers, vol. 64, no. 5, 2017.
- [15] V. Cambareri, M. Mangia, F. Pareschi, R. Rovatti and G. Setti, "A rakeness-based design flow for Analog-to-Information conversion by Compressive Sensing," 2013 IEEE International Symposium on Circuits and Systems (ISCAS), 2013, pp. 1360-1363, doi: 10.1109/IS-CAS.2013.6572107.
- [16] S. Onn, I. Weissman, "Generating uniform random vectors over a simplex with implications to the volume of a certain polytope and to multivariate extremes," *Annals of Operation Research*, vol. 189, n. 1, pp. 331-342, 2011
- [17] F. Mezzadri, "How to generate random matrices from the classical compact groups," Notices of the American Mathematical Society, vol. 54, n. 5, pp. 592-604, 2006
- [18] J. Ginibre, "Statistical ensembles for complex, quaternion, and real matrices", *Journal of Mathematical Physics*, vol. 6, pp. 440-449, 1965
- [19] I. Goodfellow, Y. Bengio, A. Courville, "Deep Learning", MIT Press, 2016
- [20] A.B. Reinaldo, C.E., Javier, G.G. Marc, "Shannon Entropy and Mutual Information for Multivariate Skew-Elliptical Distributions," *Scandinavian Journal of Statistics*, 2013, vol. 40, n. 1, pp. 42-62, doi: 10.1111/j.1467-9469.2011.00774.x.
- [21] B. Schölkopf, R. Williamson, A. Smola, J. Shawe-Taylor, J. Platt, "Support Vector Method for Novelty Detection" Neural Information Processing Systems (NIPS), 1999, pp. 582-588
- [22] S. M. Kay, *Fundamentals of Statistical Signal Processing, Volume 2 Detection Theory*, Prentice Hall PTR, 1998
- [23] T. Fawcett, "An introduction to roc analysis," *Pattern Recognition Letters*, vol. 27, no. 8, pp. 861 - 874, 2006
- [24] A.N. Kolmogorov, "On the Shannon Theory of Information Transmission in the Case of Continuous Signals," *IRE Transactions on Information Theory*, pp. 102-108, December 1956
- [25] A. M. Mathai, Serge B. Provost, *Quadratic Forms in Random Variables: Theory and Applications*, Marcel Dekker, Inc., New York, 1992
- [26] T. Jiang, "How many entries of a typical orthogonal matrix can be approximated by independent normals?", *The Annals of Probability*, vol. 34, n. 4, pp. 1497-1529, 2006
- [27] Michael H. Neumann, "A central limit theorem for triangular arrays of weakly dependent random variables, with applications in statistics," *ESAIM: Probability & Statistics*, vol. 17, pp. 120-134, 2013