

Joint Super-Resolved Compressive Sensing and Encryption for Earth Observation Applications

Original

Joint Super-Resolved Compressive Sensing and Encryption for Earth Observation Applications / Migliorati, Andrea; Valsesia, Diego; Bianchi, Tiziano; Magli, Enrico. - ELETTRONICO. - (2022), pp. 1-7. (Intervento presentato al convegno OBPDC2022 - 8th International Workshop on OnBoard Payload Data Compression tenutosi a Athens, Greece nel 28-30 September 2022) [10.5281/zenodo.7245135].

Availability:

This version is available at: 11583/2979932 since: 2023-07-07T10:29:36Z

Publisher:

ESA

Published

DOI:10.5281/zenodo.7245135

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)

28-30 SEPTEMBER 2022

JOINT SUPER-RESOLVED COMPRESSIVE SENSING AND ENCRYPTION FOR EARTH OBSERVATION APPLICATIONS

Andrea Migliorati ⁽¹⁾, Diego Valsesia ⁽¹⁾, Tiziano Bianchi ⁽¹⁾, Enrico Magli ⁽¹⁾

⁽¹⁾ *Politecnico di Torino – Department of Electronics and Telecommunications (DET)
Corso Duca degli Abruzzi 24, 10129 Torino, Italy
E-mail: NAME.SURNAME@polito.it*

INTRODUCTION

Over the years, we have witnessed a continuous increase in data generated by imaging sensors employed for Earth Observation (EO) applications, which has been caused by the evolution of spatial image acquisition technologies and the need to obtain data at a higher resolution than before to foster scientific advancement and improve commercial applications. Dealing with an increased amount of data, however, implies increased requirements for the computational and memory resources available to the onboard processing units. These units must be able to efficiently deal at the same time with large data flows, limited energy consumption constraints, and the bottleneck of data transmission to the ground. Researchers have been trying to tackle these challenging problems by working on improved ways of acquiring signals (i.e., images) from the imaging sensors. The Compressive Sensing (CS) paradigm enables sensing by optical light modulation followed by computational image reconstruction at the ground segment. Deep learning methods have proven to be very effective for the reconstruction process.

We present work done within the SUPRISE (SUPER-Resolved compRESSive InSTRument) EU-funded project for EO applications in the visible and medium infrared with the goal of implementing a super-spectral instrument with enhanced capability in onboard data processing, spatial resolution, and encryption functionalities. Unlike other traditional sensing techniques, SURPRISE employs data acquisition by the projection of the images in random spaces described by a given sensing matrix. Hence, assuming that the sensing matrix is kept secret for each projection space and that it is only known to the reconstruction instrument/hardware that reconstructs the image, this framework can be seen as a symmetric key encryption system with specific security properties. In this paper, we investigate the security features of the SURPRISE instrument model, performing a thorough theoretical and experimental evaluation that analyses the generation of the random matrices for each acquisition, the robustness against malicious attacks under different scenarios, and the computational overhead of the security scheme. We demonstrate that the SUPRISE instrument exhibits strong security properties which comes at the expense of very little additional complexity.

BACKGROUND ON COMPRESSIVE SENSING

The Compressive (or Compressed) Sensing (CS) framework, introduced in [1], is today one of the most important and employed techniques in the field of signal processing as it enables applications to efficiently deal with the increasingly larger amounts of information acquired by sensors. CS describes techniques to find sparse solutions to an underdetermined system, and specifically to set of linear equations, also enabling compressed signal sensing. In particular, when dealing with image acquisition and processing, CS attempts to overcome the Shannon-Nyquist sampling theorem by using a number of data samples as close as possible to the degrees of freedom of the image. This ensures that the signal representation can get significantly easier to manipulate and transmit over a data channel.

$$\mathbf{y} = \Phi \mathbf{x} \quad (1)$$

CS can be generally described by Eq.1, which indicates how an image \mathbf{x} can directly be represented by means of a narrow set of weighted linear combinations of samples called *measurements*, \mathbf{y} , via the projection on a subspace defined by the equation matrix ϕ called *sensing matrix* or *projection matrix*. Typically, \mathbf{y} has a greatly reduced dimensionality with respect to \mathbf{x} , while ϕ is chosen as a random matrix with independent and identically distributed random entries (i.e., iid). Authors in [2] demonstrated how, under certain conditions, CS can be effectively used as a secure encryption system where the sensing matrix ϕ acts as the encryption key. For these reasons, CS can be a great fit for EO applications as it can potentially lead to the development of innovative computational imaging instruments perfectly apt to deal with *big data from space*. CS can indeed provide native joint compression and encryption, hence enabling on board image processing and analysis while at the same time reducing computational requirements on board.




Fig. 1. Optical layout of the SURPRISE instrument.

In the recent years, the signal processing community has been developing new and more efficient CS algorithms based on Deep Learning (DL) techniques, which allowed for great performance improvement over traditional algorithms. DL methods can reconstruct an original image \mathbf{x} from its measurements \mathbf{y} in a very short time, at the expense of an initial so-called *training phase* on a specific dataset, as typical of the realm of Machine Learning (ML) algorithms to which DL belongs to.

DESIGN OF THE SURPRISE (SUPER-RESOLVED COMPRESSIVE) INSTRUMENT

In this section we describe the proposed SURPRISE instrument together with the corresponding image acquisition model that is needed to successfully reconstruct the original images from the measurements. The main components of the considered system are as reported in Fig.1:

- A scanning system enabling whiskbroom operations.
- A spatial light modulator (SLM) which modulates the scene according to a binary pattern, referred to as sensing mask, implemented with a digital micromirror device (DMD).
- A set of dichroic mirrors used to first split VIS-NIR light from IR light, and then to split IR light into two spectral bands of interest for the two dedicated MWIR channels.
- A set of optical condensers that perform the signal spatial integration and focus it on the channel's detector.
- A spectrometer and detection system used to implement the VIS-NIR channels.
- A single element MWIR detector and relevant bandpass filters used to implement the two MWIR channels.

As illustrated, the SURPRISE instrument is a super-resolved whiskbroom imager that acquires the scene in a serial fashion, one *macro-pixel* at a time, where *macro-pixel* indicates the image progressively captured in the instantaneous field of view (IFOV) of the device. Due to the coded aperture design exploiting a spatial light modulator, each macro-pixel can be reconstructed into a block of $n \times n$ pixels (or *micro-pixels*), where the value of n is a tunable parameter called *super-resolution factor*, as it specifies how many micro-pixels can be reconstructed from the measurements of a macro-pixel. For a single spatial location \mathbf{P} in the whiskbroom acquisition procedure, we can model the described system as follows by referencing Eq.1:

$$\mathbf{y}_i^P = \Phi_i^P \mathbf{x}_i^P, \quad (2)$$

where \mathbf{y}_i^P indicates the column vector of m measurements acquired by the detector in the spectral band i , \mathbf{x}_i^P represents the column vector of n image micro-pixels, and Φ_i^P is the binary sensing matrix which can be modelled as an $m \times n$ matrix with rows corresponding to the serially acquired measurements. The ratio between the number of measurements and the number of micro-pixels, m/n , is denoted as *compression ratio* (CR). By considering Eq.2 on a global image model in order to effectively exploit spatial correlation in the image reconstruction algorithm for a band i , we reach the following result:

$$\mathbf{y}_i = \Phi_i \mathbf{x}_i. \quad (3)$$

Hence, the SURPRISE instrument performs a whiskbroom scanning of a scene where the measurements in the spectral band i are acquired by means of sparse block-diagonal sensing matrix Φ_i . As anticipated, utmost importance is reserved to the choice of the reconstruction algorithm, on which depends accordingly the definition of the sensing matrix and measurements computation. In our case, the choice falls on the DL-based method called ISTA-Net+ [5] that relies on a Convolutional Neural Network (CNN) to implement the very fast, traditional CS method known as Iterative Shrinkage

Thresholding Algorithm (ISTA) [4]. While ISTA is known for reconstructing the original signals from the measurements by means of multiple thresholded iterations, its DL counterpart ISTA-Net+ adapts very well to natural images thanks to the employment of CNNs. Specifically, ISTA-Net+ consists of a certain number of basic blocks (called *phases*) which maps the iterative threshold updates of the original ISTA algorithm directly on the convolutional layers arranged in simple, repeatable units. From a macroscopical point of view, ISTA-Net+ employs a non-linear transformation via convolutional operators and ReLU [3] activation units to sparsify the input images. As shown in [5], ISTA-Net+ greatly improves performance on image reconstruction with respect to a wide range of traditional CS algorithms as well as more recent DL-based ones, for multiple considered compression ratios. Further, it reduced the computational burden with respect to typical CS methods. Also, thanks to the fact that no specific structural constraints are imposed on the sensing matrix, but rather the DL-based algorithm iteratively updates the CNN's parameters to solve a local minimum problem dependent on the chosen Φ , ISTA-Net+ could easily be extended to other CS domains than the one of natural images. These considerations allow us to conclude that ISTA-Net+ offers a great tool for CS thanks to its combination of scalability, flexibility, and high performance. For this reason, the reconstruction algorithm has been based on it; the final algorithm has employed significant extensions of ISTA-Net+, some of which have been reported in [14].

DESIGN OF THE ENCRYPTION ALGORITHM

As anticipated, the adopted DL-based approach to image acquisition within the SURPRISE device can also be seen as a symmetric key cryptographic system, where the sensing matrix Φ represents the cryptographic key, and the measurements \mathbf{y} can be thought as the encrypted payload obtained from the original image \mathbf{x} . Indeed, as known from the literature [9], it has been demonstrated that measurements computed from signals with the same energy are undistinguishable asymptotically with the dimensionality of the signals. In our scenario, we work under the assumption that the sensing matrix is successfully kept secret to the attacker and only known to the reconstruction hardware. A perfectly secure cryptographic system ensures the statistical independency between the original signal \mathbf{x} and the correspondent encrypted signal \mathbf{y} , such that $p(\mathbf{x}|\mathbf{y}) = p(\mathbf{x})$ (i.e., $I(\mathbf{x}, \mathbf{y}) = 0$). As known from the literature [2], perfect secrecy can be obtained only when two specific conditions are met: **(1)** the sensing matrix has entries sampled from independent Gaussian iid distributions and changed at every CS acquisition; **(2)** the energy of the measurements is constant. We can proceed to describe the adopted solutions for the SURPRISE instrument as to address the Conditions **(1)** and **(2)**, while a thorough security evaluation of the implemented framework is instead presented in the next section.

If Condition **(1)** is not met, such as in the case of real-life EO applications, then theoretical secrecy cannot be reached. However, it can be demonstrated that, for specific types of sensing matrices such as discrete-valued ones, the advantage an attacker might have at recovering the original signal becomes negligible asymptotically [7]. In the SURPRISE device, the implementation of the SLM sensor forces us to consider only sensing matrices which assume binary values, corresponding to the two possible positions the DMD mirrors can take (Fig.1). Hence, we must choose a suitable design based on binary sensing matrices whose security properties has already been investigated in the literature [8][9]. For example, it is known that a binary matrix can well approximate the security properties of a perfectly random Gaussian matrix as the dimension of the acquired signal grows [2]. With specific reference to Condition **(1)**, we can approximately reach computational security for our system by generating a sequence of pseudo-casual matrices, one for each acquisition (i.e., encryption), via cryptographically secure pseudo-casual bit generators [11]. In particular, we use the low-complexity stream cyphers known as Salsa-20/ChaCha-20 [43][44] that are extremely efficient as they are only based on three basic hardware operations: the sum of values from fixed-size registers, exclusive OR (XOR), and 1-bit data shift. As a matter of example, an optimized Salsa-20/ChaCha-20 implementation can produce 1 pseudo-random byte on a standard x86 architecture using 4 clock cycles. Also, the considered stream cyphers work in constant time regime, and this offers the advantage of greatly reducing the vulnerability to the so-called *side channel* attacks where the attacker tries to infer information by separately measuring the time for encryption and decryption. Further, Salsa-20/ChaCha-20 can be freely used in any desired system as they are not protected by patent and therefore publicly available. Note that the parameter 20 indicates the number of iterative rounds used by the stream cypher to generate the pseudo-casual bit sequence, and it can be reduced in order to meet a suitable trade-off between security and the available resources.

Secondarily, it can be observed that, due to the linearity of the acquisition process of the SURPRISE instrument, it is in theory always possible for an attacker to at least infer the energy e_x of the original signal \mathbf{x} by observing the measurements \mathbf{y} . For this reason, to overcome this security concern, we employ energy normalization of the measurements. Consequently, we encrypt the normalization factor with the same stream cypher used for generating the encryption key (i.e., the pseudo-random binary matrix). The encrypted normalization factor $E[e_y]$ can then be considered as an additional measurement, so that the encrypted payload y_n is defined as in the following:

$$y_n = \left\{ y / \sqrt{e_y}, E[e_y] \right\}. \quad (4)$$




Fig. 2. θ -distinguishability as a function of the macro-pixel size n ; the considered signals have a fixed sparsity equal to 20% of the number of micro-pixels in the macro-pixel; the curves have been computed assuming an independent reconstruction is carried out for contiguous macro-pixels, which is a worst-case scenario; the curves show that, even at very small macro-pixels ($n = 4$), θ takes very small values ($\theta \approx 0,01$)

In such fashion, by obtaining measurements with unitary energy, we ensure Condition (2) is respected while also greatly reducing the chances of inferring the energy of the original signal. On the downside, the adopted solution entails the transmission of an additional measurement to the original sets of measurements, therefore slightly increasing the complexity of the design. However, in the following sections we show that the increase in the computational overhead is not as such as to burden the system design in terms of resource consumption.

SECURITY ANALYSIS

Since its early stages, the CS framework allowed investigations into the possibility of providing at least some notion of confidentiality for the reconstructed signals. As anticipated in the previous section, our SURPRISE instrument design can indeed be considered a computationally secure system [7] even if it does not enjoy theoretical perfect secrecy from an Information Theory perspective [6], and even if we employ binary sensing matrices [10]. In the following, we proceed in measuring the security of the SURPRISE device in terms of *distinguishability*.

θ -distinguishability evaluation

To evaluate the security properties of the proposed system we use the concept of *distinguishability* [2], which measures the advantage a hypothetical attacker might have in attacking the considered system with respect to an ideal (i.e., perfectly secure) one. In other words, given two signals x_1 and x_2 , a cyphered signal y and a decoder $D(y)$, we measure the capacity of the decoder to understand if the cyphered signal has been obtained either from x_1 or x_2 . Defining P_d and P_f as the probability of making a correct decision and the probability of false alarm, respectively, two signals x_1 and x_2 are θ -distinguishable when, for every possible decoder $D(y)$, we have:

$$P_d - P_f < \theta \quad (5)$$

Accordingly, *distinguishability* can also be referred to as θ -distinguishability. The parameter θ measures the advantage for the attacker to choose the correct signal from the encrypted payload y with respect to a random guess. In particular, the decoder can pick the right signal (i.e., correctly decipher the encrypted signal) with a probability $P_d \leq 1/2 + \theta$, from which we can conclude that a system enjoying perfect secrecy would have $P_d = P_f = 1/2$ and hence $\theta = 0$. For this reason, the smaller the value of θ is, the closest the implemented system is to having theoretical security. We compute the θ -distinguishability values for our system as a function of the size of the macro-pixel ($n \times n$) in terms of micro-pixels, for a fixed signal sparsity level k . Specifically, we simulate encrypted signals (i.e., measurements in our framework) which are sparse in the DCT-2D domain by sampling first a uniform distribution and then a Bernoulli one to obtain each value of the signal [2].




Fig. 3. Results when recovering a full image with a key that is different than the one used for the acquisition in the case of a macro-pixel of size 32×32 micro-pixels; from left to right: original image, reconstruction with the SURPRISE instrument, and reconstruction with an independently generated sensing matrix for each macro-pixel but different to the correct key used during acquisition.

The considered signals are 1-dimensional, representing the column vector of the measurements obtained via the sensing matrix from the corresponding image. We consider a worst-case scenario where the attacker only attacks one macro-pixel at the time. Fig.2 reports our results going from macro-pixels of size $4 \times 4 = 16$ micro-pixels all the way to macro-pixels of size $32 \times 32 = 1024$ micro-pixels for both uniform and Bernoulli entries. By looking at the curves, it can be observed that an attacker knowing only two signals of size $32 \times 32 = 1024$ micro-pixels have been transmitted would only have an advantage of $\theta \approx 10^{-5}$ with respect to a random guess (i.e., $P_d \leq 1/2 + 10^{-5}$) when trying to infer the original one by observing measurements.

This advantage increases to $\theta \approx 0.01$ when the macro-pixel size goes down to $4 \times 4 = 16$ macro-pixels (i.e., $P_d \leq 1/2 + 10^{-1}$). However, practical attacks are carried out against entire images and not single macro-pixels. Hence, even if the attacker could more easily attack a macro-pixel in our secure framework than a traditional cryptographic system, the complexity of the attack on the whole image level is comparable to the one the attacker would face in the case of brute-force attacks on the encryption key (i.e., the sensing matrix). For example, if we assume the attacker could only be choose between two values for each macro-pixel, the probability of inferring an original image of size 128×128 (composed of $32 \times 32 = 1024$ macro-pixels, each of size 4×4 micro-pixels) from the measurements would be equal to $2^{-51} = 3.5634e^{-300}$. For reference, the probability of guessing a 256-bit cryptographic key would be equal to $2^{-256} = 8.6362e^{-78}$, which means that, in this situation, it would be easier for the attacker to carry out a brute-force search for the cryptographic key than trying to exploit the structural and security properties of the binary sensing matrix. Also, Figure 3 shows the results obtained when recovering a full EO image with a cryptographic key that is different than the one used for the acquisition, in the case of a macro-pixel of size 32×32 micro-pixels. The image on the left shows the original image, the image in the center shows the reconstruction with the SURPRISE instrument, and the image on the right refers to reconstruction with an independently generated sensing matrix for each macro-pixel, but different to the correct key used during acquisition (i.e., simulating an attack without knowledge of each macro-pixel's cryptographic key). As easily observed shown in Fig.3, it is impossible to infer any meaningful information from the image that has been reconstructed without knowing the encryption key, as the resulting image is completely scrambled with clear artifacts at the borders. This result is obtained from a specific EO sample image, but it can be extended to any acquired image.

Our analysis so far showed that casual binary matrices can offer computational security even when the macro-pixel has small sizes, down to 4×4 micro-pixels. However, all our conclusions about practical security would not stand if the procedure by which we generate the binary casual sensing matrices were not secure. For this reason, we proceed to assess the level of security ensured by the chosen pseudo-random bit generators Salsa-20/Chacha-20. To tackle this task, we refer to known results from the literature [12][13] which indeed show that the considered stream cyphers can be considered equivalent to pseudo-casual generator functions. Hence, the recovery of a pseudo-casual bit sequence generated with Salsa-20/Chacha-20 would require a hypothetical attacker the same number of resources required for a brute-force attack on the encryption key, which corresponds to 2^{255} operations, on average. For reference, it has been demonstrated that, even with 8 rounds (i.e., Salsa-8/Chacha-8), the number of operations required to attack the stream cypher is 2^{255} , showing how our solution provides security that is comparable to traditional cryptographic methods.

Finally, we analyze practical aspects related to the implementation of the casual binary sensing matrices. While ideally sensing matrices should be binary $0/1$ value, in practice, due to the non-ideality of the SLM, they have binary entries with a low value equal to α and a high value equal to β . The exact α, β values depend on the SLM calibration and change as a function of wavelength, so the same matrix cannot be reused to model the acquisition of all detectors in all bands. In our approach, we consider the following approximation: $\alpha = 0, \beta = 1$, and assume an unknown rescaling factor to be computed at instrument calibration time to adjust the sensing matrices' entries accordingly, before energy normalization of the measurements.




Fig. 4. Reconstruction example by exploiting the estimated energy of a single macro-pixel composed of 32×32 micro-pixels; left: original image; right: reconstruction from the energy of the measurements via nearest-neighbor interpolation.

We conclude the security evaluation by providing an example of reconstruction of the original image using the SURPRISE instrument by exploiting the estimated energy of a single macro-pixel composed of 32×32 micro-pixels at the time. Specifically, given the measurements \mathbf{y} , with l measurements for each macro-pixel, and their energy of the measurements e_y , we have that the estimated energy of the original macro-pixel \bar{x} is equal to $\bar{x} = \frac{2^m}{n} e_y$, where n denotes the macro-pixel size (i.e., the super-resolution factor). Hence, an attacker can estimate the micro-pixels inside the macro-pixel by applying the nearest-neighbor interpolation on each macro-pixel from the corresponding estimated energy \bar{x} . An example of such a reconstruction is provided in Figure 4. The figure on the left shows the original image, while the one on the right refers to the reconstruction inferred from the energy of the measurements. We observe that, even though the reconstruction from the measurements' energy hints at the underlying macrostructures that are present in the original image, the micro-details are heavily distorted, and no useful or sensitive information can be inferred from it.

COMPLEXITY EVALUATION

In this section, we estimate the additional computational overhead determined by the encryption system employed in the SURPRISE device as described previously. First, we evaluate the complexity disregarding energy normalization of the measurements. Secondly, we introduce the same measure and provide the final complexity assessment. We assume 20 rounds of the Salsa/Chacha stream cypher which correspond to a system with maximized security.

By construction, the considered stream cypher Salsa-20/Chacha-20 maps a 256-bit key, a 64-bit nonce and a 64-bit counter onto a 512-bit sequence that represents the final output of the stream (i.e., the desired pseudo-random bit sequence). Specifically, given two 32-bit strings a, b , Salsa-20/Chacha-20 is composed of a 20-iteration chain of three operations:

- 32 sums, required for computing $a + b \bmod 2^{32}$;
- 32 XOR operations, required for computing $a \oplus b$;
- 32 bit shifts, required for obtaining $a \ll b$.

As reported in [43], the estimated complexity of the stream cypher on the standard family of x86 processor goes from 4 to 14 cycles per byte, according to the specific architecture. Hence, in the example of a macro-pixel composed of $4 \times 4 = 16$ micro-pixels, the estimated complexity of the stream cypher sets around $8 \div 28$ cycles for each measurement. Now, we can compute the additional number of operations required by the energy normalization of the measurements, as described in Eq. (4), for each single measurement. For simplicity, we assume as negligible the effort to transmit the extra required measurement represented by the encrypted energy of the measurements $E[e_y]$. Considering the i -th measurement y_i composed of l elements, the additional number of operations is the amount required to compute the energy of the signal. In more detail, these amount to:

- $2 * l$ multiplications, required for computing the component-wise inner products $y_i * y_i$ and the divisions by $\sqrt{e_y}$, under the assumption the division operator is implemented as a multiplication by a pre-stored scaling factor;
- $l - 1$ sums, required for computing the sum of the component-wise inner products $(y_1 * y_1 + \dots + y_i * y_i + \dots + y_l * y_l)$;
- 1 square root, required for computing $\sqrt{e_y}$ which is then stored and used for every division.

In the case of the SURPRISE instrument, where a significant number of measurements is considered, we can assume that the cost of computing the square root of the energy of the measurement is negligible. Hence, the additional overhead due

to the energy normalization can be quantified in 2 multiplications and 1 sum for measurement. To finally conclude on the total overhead of the encryption system, we add the overhead caused by the energy normalization to the complexity of the stream cypher previously estimated. Assuming each measure is coded by a 16-bit representation within a generic x86 architecture, the energy normalization cost approximately takes value in the range of 3 to 4 additional cycles per measure.

CONCLUSIONS

In this work, we introduced the super-resolved compressive instrument called SURPRISE for EO applications in the visible and medium infrared. The device implements a super-spectral instrument with enhanced capability in onboard data processing, spatial resolution, and encryption functionalities. We employ data acquisition by the projection of the images in random spaces described by a given binary random sensing matrix. We rely on the Deep Learning technique ISTA-Net+ for reconstruction which greatly improves over traditional CS methods. Assuming the sensing matrix is kept secret for each projection space and changed at every acquisition, this framework can be seen as a symmetric key encryption system.

In our experimental evaluation, we demonstrated that the SUPRISE instrument exhibits strong security properties even in the case of limited size macro-pixels via both theoretical (*θ -distinguishability*) and practical considerations. We use the cryptographic primitives from the Salsa-20/ChaCha-20 family to implement a computationally secure stream cypher which we use to obtain the necessary pseudo-random bit sequences necessary for the random binary sensing matrix. Finally, we show that the developed encryption system comes at the expense of very little additional complexity to the unsecure design.

ACKNOWLEDGEMENTS

The SURPRISE project has received funding from the European Union's Horizon 2020 research and innovation programme under Grant Agreement No 870390.

REFERENCES

- [1] E. J. Candes e M. B. Wakin, «An introduction to compressive sampling» IEEE signal processing magazine, vol. 25, n. 2, p. 21–30, 2008.
- [2] T. Bianchi, V. Bioglio e E. Magli, «Analysis of one-time random projections for privacy preserving compressed sensing» IEEE Transactions on Information Forensics and Security, vol. 11, n. 2, pp. 313-327, 2016.
- [3] Fukushima, Kunihiko. "Cognitron: A self-organizing multi-layered neural network." Biological cybernetics 20.3 (1975): 121-136.
- [4] A. Chambolle, R. A. DeVore, N. Lee e B. J. Lucier, «Nonlinear wavelet image processing: Variational problems, compression, and noise removal through wavelet shrinkage,» IEEE Transaction on Image Processing, vol. 7, n. 3, pp. 319-35, 1998.
- [5] J. Zhang e B. Ghanem, «ISTA-Net: Interpretable Optimization-Inspired Deep Network for Image Compressive Sensing,» in 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2018.
- [6] C. E. Shannon, "Communication theory of secrecy systems," Bell. Syst.Tech. J., vol. 28, no. 4, pp. 656–715, Oct. 1949.
- [7] Y. Rachlin and D. Baron, "The secrecy of compressed sensing measurements," in Proc. Annu. Allerton Conf. Commun., Control, Comput., Sep. 2008, pp. 813–817.
- [8] M. Testa, T. Bianchi and E. Magli, "Secrecy Analysis of Finite-Precision Compressive Cryptosystems," in IEEE Transactions on Information Forensics and Security, vol. 15, pp. 1-13, 2020.
- [9] V. Cambareri, M. Mangia, F. Pareschi, R. Rovatti, and G. Setti, "Low-complexity multiclass encryption by compressed sensing," IEEE Trans. Signal Process., vol. 63, no. 9, pp. 2183–2195, May 2015.
- [10] M. Testa, D. Valsesia, T. Bianchi, and E. Magli, Compressed Sensing for Privacy-Preserving Data Processing. Singapore: Springer, 2018.
- [11] R. Fay, "Introducing the counter mode of operation to compressed sensing based encryption," Inf. Process. Lett., vol. 116, no. 4, pp. 279–283, 2016.
- [12] Paul Crowley, Truncated differential cryptanalysis of five rounds of Salsa-20, Cryptology ePrint Archive, Report 2005/375.
- [13] Shi Z., Zhang B., Feng D., Wu W. (2013) Improved Key Recovery Attacks on Reduced-Round Salsa-20 and ChaCha. In: Kwon T., Lee MK., Kwon D. (eds) Information Security and Cryptology – ICISC 2012. ICISC 2012. Lecture Notes in Computer Science, vol 7839. Springer, Berlin, Heidelberg.
- [14] M. Cilia, N. Prette, D. Valsesia, T. Bianchi, E. Magli, "Deep learning methods for satellite compressive imaging", Proc. OBPDC - Onboard Payload Data Compression workshop, 2022.