

# Equivariance-based self-supervised learning for audio signal recovery from clipped measurements

Victor Sechaud  
CNRS, ENS de Lyon  
Laboratoire de physique  
F-69007 Lyon, France  
victor.sechaud@ens-lyon.fr

Laurent Jacques  
UCLouvain  
ICTEAM  
Louvain-la-Neuve, Belgium  
laurent.jacques@uclouvain.be

Patrice Abry  
CNRS, ENS de Lyon  
Laboratoire de physique  
F-69007 Lyon, France  
patrice.abry@ens-lyon.fr

Julián Tachella  
CNRS, ENS de Lyon  
Laboratoire de physique  
F-69007 Lyon, France  
julian.tachella@ens-lyon.fr

**Abstract**—In numerous inverse problems, state-of-the-art solving strategies involve training neural networks from ground truth and associated measurement datasets that, however, may be expensive or impossible to collect. Recently, self-supervised learning techniques have emerged, with the major advantage of no longer requiring ground truth data. Most theoretical and experimental results on self-supervised learning focus on linear inverse problems. The present work aims to study self-supervised learning for the non-linear inverse problem of recovering audio signals from clipped measurements. An equivariance-based self-supervised loss is proposed and studied. Performance is assessed on simulated clipped measurements with controlled and varied levels of clipping, and further reported on standard real music signals. We show that the performance of the proposed equivariance-based self-supervised declipping strategy compares favorably to fully supervised learning while only requiring clipped measurements alone for training.

## I. INTRODUCTION

Inverse problems appear in various and numerous scientific and engineering applications (such as tomography, MRI [1] or phase retrieval [2]). They consist in recovering a signal  $\mathbf{x} \in \mathcal{X} \subset \mathbb{R}^n$  from measurements  $\mathbf{y} \in \mathcal{Y} \subset \mathbb{R}^m$ , through a forward operator  $A : \mathbb{R}^n \rightarrow \mathbb{R}^m$ , possibly corrupted with noise  $\epsilon \in \mathbb{R}^m$ :

$$\mathbf{y} = A(\mathbf{x}) + \epsilon. \quad (1)$$

Inverse problems are often ill-posed, because the forward operator may be incomplete, *i.e.*,  $m < n$ . Regularization strategies based on incorporating prior information have been widely studied and shown to have performance that strongly depends on the relevance of the chosen prior [3]. To overcome this limitation, most strategies are based on learning an inverse operator of  $A$  on  $\mathcal{X}$ ,  $f_\theta : \mathbf{y} \mapsto \mathbf{x}$ , from a training dataset containing ground truth signals and associated measurement  $\{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^N$ . Standard approaches minimize the supervised mean square error (MSE):

$$\min_{\theta} \sum_{i=1}^N \|f_\theta(\mathbf{y}_i) - \mathbf{x}_i\|^2 \quad (2)$$

where  $f_\theta$  is generally modeled by a neural network. Supervised learning however suffers from two main limitations: (i) training sets can be difficult or impossible to obtain, *e.g.*, in medical

VS's PhD is funded by ANR UNLIP. LJ's research is supported by the FRS-FNRS (QuadSense, T.0160.24).




Fig. 1: Learning to declip signals from measurement data alone. We propose a new self-supervised method that can learn to declip audio signals without ever seeing ground-truth reference signals by exploiting carefully chosen assumptions on the invariance of the reconstructed signal distribution.

or space imaging [4], and (ii) even when ground-truth datasets are available, there might be a significant distribution shift between training (often synthetic) and testing (real) datasets.

To get over these two issues, self-supervised learning based on equivariance principles was proposed in [5] which only requires a dataset of measurements  $\{\mathbf{y}_i\}_{i=1}^N$ . Their performance was theoretically assessed only for linear inverse problems [6].

The present work contributes to extending equivariance-based self-supervised learning to non-linear inverse problems by addressing audio signal recovery from clipped measurements, as sketched in Fig. 1. Declipping is a common non-linear distortion, typically occurring with analog-to-digital (ADC) converters when the dynamic range of the original (analog) signal is too high.

### A. Related works

1) *Audio declipping*: Several unsupervised techniques for audio declipping were proposed [7]–[9], with many employing variational methods, incorporating priors such as sparsity in either the Fourier or time domain. Both  $\ell_0$  and  $\ell_1$  minimization approaches are used, often performed via algorithms like Douglas-Rachford or Chambolle-Pock. Additionally, learning-based methods such as dictionary learning were proven effective, with a dictionary trained from time windows of clipped signal  $\mathbf{y}$ . Other prior information inspired on human perception [10], or leveraging the presence of data from multiple channels [11], [12], can be also used to improve the performance of these methods. The present work

aims to avoid the reliance on such hand-crafted assumptions by learning directly from measurements alone.

2) *Self-supervised learning for inverse problems*: In the realm of unsupervised learning, several methods like noise2noise [4], noise2void [13], have gained prominence for handling noise in measurements. These methods primarily focus on mitigating noise artifacts but struggle with addressing incomplete information stemming from non-invertible operators  $A$ . An unsupervised learning method, referred to as equivariant imaging (EI), was proposed in [5] to solve linear inverse problems. It differs from noise2x methods by exploiting invariance properties, which may exist in natural images, to learn where  $A$  does not provide information about the signal set  $\mathcal{X}$ . This is based on a natural intuition: the invariance of a signal set under a transformation implies that signals remain identifiable after the application of the transformation. Theoretical guarantees were provided in [6] for linear forward operators only. A first attempt to study EI self-supervised learning for non-linear operators was conducted in [14] in the one-bit compressed sensing framework, yet restricted to equivariance to translations and rotations. Here we show that equivariance to such geometrical transformations is not enough for learning from clipped measurements alone.

### B. Goals, contributions, and outline

The goal of the present work is to devise and assess the performance of equivariance-based self-supervised learning to solve the nonlinear inverse problem of signal declipping, where amplitude invariance must be used, rather than translation or rotation invariance. To that end, Sec. II recalls the formulation of declipping as an inverse problem and discusses the *naive* measurement consistency solutions and their limitations. Sec. III introduces scale invariance, discusses its use for equivariance learning, and designs the corresponding loss function. Sec. IV presents some numerical experiments. First, synthetic simulated clipped measurements with controlled and varied levels of clipping are used, thus permitting quantifying the impact of the choice of hyperparameters on reconstruction performance. Second, the application of the proposed equivariance-based self-supervised declipping strategy on a standard real sound dataset, shows that it performs comparably to fully supervised learning procedures.

## II. UNSUPERVISED LEARNING FOR DECLIPPING

### A. Clipping forward model

Clipping consists of applying a non-linear thresholding operator to the original data  $x$ :

$$\mathbf{y} = \eta(\mathbf{x}) \in \mathbb{R}^n, \quad (3)$$

with  $\eta$  the threshold operator of level  $\mu > 0$  (depending of the measurement device) defined for a scalar  $c$ :

$$\eta(c) = \begin{cases} \text{sign}(c)\mu & \text{if } |c| \geq \mu \\ c & \text{otherwise} \end{cases}$$

which is applied component-wise on vectors.

### B. Naive measurement consistency

In supervised learning, an MSE type loss, similar to (2) is generally minimized. In unsupervised learning, the dataset consists of measurements only. The training loss must thus be rewritten to avoid using the ground truth  $\mathbf{x}$ . The natural idea is to take the naive measurement consistency:

$$\mathcal{L}_{\text{NMC}}(\theta) = \sum_{i=1}^N \|\mathbf{y}_i - \eta(f_{\theta}(\mathbf{y}_i))\|^2 \quad (4)$$

which enforces the reconstructed signals to be consistent with the observed measurements. However, it is impossible to learn from measurements consistency only because any function  $f$  of the form

$$f(\mathbf{y})_j = \begin{cases} y_j & \text{if } |y_j| \leq \mu \\ v(y_j) & \text{if } |y_j| = \mu \end{cases} \quad (5)$$

with  $v$  being a function lying above the threshold level  $\mu$  satisfies

$$\eta(f(\mathbf{y})) = \mathbf{y} \quad (6)$$

and thus incurs zero training measurement loss.

## III. EQUIVARIANCE-BASED SELF-SUPERVISED LEARNING FOR DECLIPPING

### A. Exploiting scale invariance

To learn in the saturated part, we have to add another constraint that removes the limitation of measurement consistency alone. Here we assume the model  $\mathcal{X}$  is amplitude invariant, *i.e.*

$$\forall \mathbf{x} \in \mathcal{X}, \forall g \in \mathbb{R}_+, g\mathbf{x} \in \mathcal{X}. \quad (7)$$

The intuition behind using this invariance is that each measurement can be seen as a signal through another threshold-type operator  $\eta_g$ :

$$\eta(\mathbf{x}) = \eta(gg^{-1}\mathbf{x}) := \eta_g(\mathbf{x}'). \quad (8)$$

with  $\mathbf{x}' = g^{-1}\mathbf{x} \in \mathcal{X}$  and  $\eta_g(\cdot) = \eta(g\cdot)$  by assumptions. Thus, we have a family of different operators with different saturation levels and different types of missing information. This invariance assumption differs from previous work on equivariant imaging [5], [14] which relies on shift and/or rotations, as  $\eta$  is equivariant to these geometrical group actions. Indeed, noting  $T_g \in \mathbb{R}^{n \times n}$  as such transformations, we would have:

$$\eta(\mathbf{x}) = \eta(T_g T_{g^{-1}} \mathbf{x}) = T_g \eta(\mathbf{x}'), \quad (9)$$

since  $T_g$  are permutation matrices, and  $\eta$  is applied element-wise. It results in a family of operators with the same saturation level (and thus do not provide additional information). We can also generalize (7) for other transformations beyond amplitude invariance, as long as (8) is respected, such as affine group transformation  $\mathbf{x} \rightarrow a\mathbf{x} + \mathbf{b}$ , with  $g = (a, \mathbf{b}) \in \mathbb{R}_+ \times \mathbb{R}^n$ .

In practice, to exploit this invariance property we remark that since  $g\mathbf{x} \in \mathcal{X}$ ,  $f$  has to be able to reconstruct  $\mathbf{x}$  and  $g\mathbf{x}$  for all  $\mathbf{x}$  in  $\mathcal{X}$ , *i.e.*

$$f(\eta(g\mathbf{x})) = g\mathbf{x} \text{ and } f(\eta(\mathbf{x})) = \mathbf{x} \quad (10)$$




Fig. 2: Comparison of  $\mathcal{L}_{\text{NMC}}$  and  $\mathcal{L}_{\text{MC}}$  with  $\mu = 1$ . On the left: Curve of the  $\mathcal{L}_{\text{NMC}}$  in one dimension for two examples in both cases saturated and not. On the right: Curve of the new  $\mathcal{L}_{\text{MC}}$  for the same two examples.

thus

$$f(\eta(gx)) = g(f(\eta(x))). \quad (11)$$

In other words, the composition of  $f$  with  $\eta$  should be equivariant for the multiplicative group  $(\mathbb{R}_+, \cdot)$ .

### B. Proposed loss

Our approach first considers a loss function whose minimization enforces the equivariance of the reconstruction function to the multiplicative group. We define it as:

$$\mathcal{L}_{\text{EI}}(\theta) = \sum_{i=1}^N \mathbb{E}_{g \sim p_g} \left[ \|gf_{\theta}(\mathbf{y}_i) - f_{\theta}(\eta(gf_{\theta}(\mathbf{y}_i)))\|^2 \right] \quad (12)$$

where the expectation is computed with respect to a certain distribution  $p_g$  over the multiplicative group. The choice of this distribution is further discussed in Sec. IV. Second, our loss function should also include a second term that promotes measurement consistency. While the simple loss introduced in (4) fulfills this role, its gradient vanishes where the signal is not saturated ( $|y_j| < \mu$ ). Indeed, in (4), as  $\eta$  has zero gradients beyond the threshold levels, we cannot hope to have the direction that minimizes the loss (by gradient descent) whereas, if the loss function is adequate, its gradient should point towards  $y$ . We therefore choose to use the MSE where the signal is not saturated, and to alter it otherwise according to the following loss:

$$\mathcal{L}_{\text{MC}}(\theta) = \sum_{i=1}^N \|h(f_{\theta}(\mathbf{y}_i), \mathbf{y}_i)\|^2. \quad (13)$$

with  $h$  defined as

$$h(\mathbf{u}, \mathbf{v})_j = \begin{cases} v_j - u_j & \text{if } |v_j| < \mu \text{ or } \text{sign}(v_j)u_j < \mu \\ 0 & \text{if } |v_j| = \mu \text{ and } \text{sign}(v_j)u_j \geq \mu. \end{cases}$$

See Fig. 2 for an example with a single signal in one dimension.

Adding both losses together results in a new loss

$$\mathcal{L}(\theta) = \mathcal{L}_{\text{MC}}(\theta) + \mathcal{L}_{\text{EI}}(\theta) \quad (14)$$

which enforces both consistency with the first term and equivariance with the second.

### C. Bias free reconstruction

We choose a bias-free version of ReLU neural network, to respect the scale-invariant assumption as proposed in [15]. Indeed a feedforward neural network with ReLU activation functions and without bias is a homogeneous function:

$$f(\alpha y) = \alpha f(y) \quad \forall \alpha \in \mathbb{R}_+. \quad (15)$$

While this architectural choice does not render the system  $f \circ \eta$  amplitude equivariant, it can still improve the generalization properties of  $f$ , since this is a desirable property for the non-saturated parts of  $y$  as it is illustrated in Fig. 3.

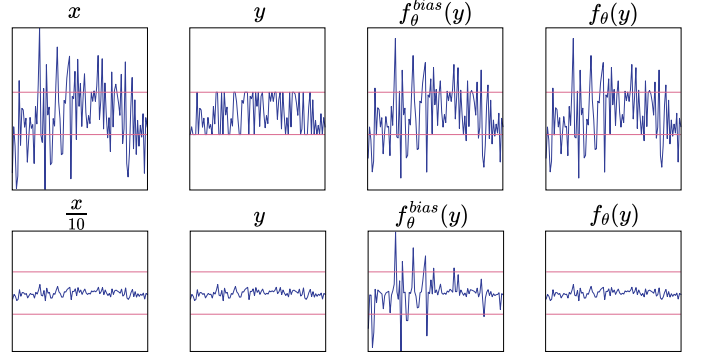


Fig. 3: A toy example where a network with learnable biases is not suitable for a scale-invariant signal set. From left to right: Top row: a signal, the associated measurement, and reconstruction; Bottom row: the same signal divided by 10 (still contained in  $\mathcal{X}$  by assumption) with the associated measurement and reconstruction.

### D. Masking

The learned reconstruction network can often slightly modify the unsaturated part of the signal. To remove this undesired effect, at test time, we compute a mask that modifies only saturated measurements, *i.e.*

$$\hat{x}_j = (1 - b_j)y_j + b_j f_{\theta}(y)_j. \quad (16)$$

$$\text{with } b_j = \frac{\max(0, |y_j| - \tau\mu)}{1 - \tau\mu}.$$

for all  $j = 1, \dots, n$ , where  $\tau$  is a hyper parameter set to 0.95 in all the experiments.

Another way to use information from unsaturated parts is to feed it into the neural network. For this purpose, we add a mask of the unsaturated part as a channel to the signal, before being fed into the neural network. In theory, we expect that the neural network itself learns to calculate this mask, but experimentally we have noticed that explicitly providing the mask as an input improves the results.

## IV. EXPERIMENTS

In this section, we present a series of experiments that evaluate the performance of the proposed method. Two datasets are used, one synthetic which allows us to handle the proportion of saturated samples and the size of the model, and the other




Fig. 4: Diagram illustrating the used model. A mask is computed from the measurement and both the mask and measurements are fed to the network.

corresponding to real sounds. The performance is evaluated using the Signal Distortion Ratio (SDR) between the ground-truth  $\mathbf{x}$  and the reconstruction  $\hat{\mathbf{x}}$  [7]:

$$\text{SDR}(\mathbf{x}, \hat{\mathbf{x}}) = 20 \log_{10} \left( \frac{\|\mathbf{x}\|_2}{\|\mathbf{x} - \hat{\mathbf{x}}\|_2} \right). \quad (17)$$

### A. Synthetic dataset

We start with experiments on a synthetic dataset constructed as follows: first, we generate a random universal subspace of  $\mathbb{R}^{100}$  with dimension  $d \in \{1, \dots, 100\}$  from basis vectors which each coordinate drawn from a standard normal distribution. We then generate  $n = 1000$  vectors  $\{\mathbf{x}\}_{i=1}^N$  on this subspace, which we rescale by a factor  $q_i \in \mathbb{R}_+$  chosen so that a proportion  $v \in (0, 1)$  of the resulting measurements are clipped (with the threshold level  $\mu$  set to 1). We evaluate the recovery performance depending on the parameters  $d$  and  $v$ . Fig. 5 demonstrates the network’s ability to reconstruct unsaturated signals, as a function of these two parameters. The test performance diminishes with a larger subspace dimension or an increased number of saturated samples. Indeed signals belonging to a higher dimensional subspace are harder to reconstruct like studying in the context of sparsity [16].

For this experiment, the distribution  $p_g$  referred in (12) is the uniform distribution on  $[0.5, 1.5]$ . We use a fully connected multilayer perceptron (MLP) with skip-connection for this experiment to remove any inductive biases associated with convolutional networks [17].

### B. Real dataset

Here we evaluate the method on a dataset of real sounds. We use the GTZAN dataset, which is composed of a collection of 10 genres with 100 audio files each, all having a length of 30 seconds with a sample rate of 22 050 Hz. We focus here on classical music but similar results could be obtained for each genre or by mixing them. Sounds are split into 30 1-second signals and they are then saturated with a threshold level  $\mu = 0.1$ . Finally, measurements  $\mathbf{y}$  that don’t have any saturated entry are removed. This results in a training dataset of 576 sounds and a testing dataset of size 63. For these experiments, we choose a standard UNet architecture which is well adapted for the time correlation of natural sound signals; and the uniform distribution on  $[0.1, 2]$  for  $p_g$ . Tab. I shows that the proposed self-supervised method obtains a performance on par with the fully supervised method. Tab. II shows that using




Fig. 5: Average reconstruction performance as a function of saturated part  $v$  and model dimension  $d$ . The value indicated corresponds to the mean SDR over the test dataset.

TABLE I: Reconstruction performance with blending for both supervised and unsupervised methods without using a mask in network input. SDR is averaged over the test dataset.

Methods	SDR
Identity	$10.39 \pm 3.77$
Supervised	$16.45 \pm 3.83$
Unsupervised	$15.28 \pm 3.81$

the mask can help to improve performance by approximately 0.5 dB.

To illustrate that it is possible to have a supervised training dataset that does not transfer well to the test set, we perform the following experiment: we create a supervised training dataset with music only and a test dataset with music and voices [18], [19] containing only measurements (no ground truth measurements). The supervised method (trained using (2)) learns from the training dataset and the self-supervised on the train and test datasets because it does not need ground truth, they are then both evaluated on the test dataset. The result presented in Tab. III shows that the self-supervised method is more robust when the training dataset differs from the test set.

TABLE II: Reconstruction performance for both mask used and not. SDR is averaged over all test dataset.

Methods	SDR
Identity	$10.39 \pm 3.77$
Supervised using mask	$16.43 \pm 3.81$
Unsupervised without mask	$15.28 \pm 3.81$
Unsupervised using mask	$15.8 \pm 4.26$

TABLE III: Comparison between supervised and self-supervised method, the first train on music dataset with ground-truth, the second music and voices dataset without ground-truth. both results are evaluated using the voices dataset.

Methods	SDR
Identity	$6.54 \pm 2.34$
Supervised	$10.94 \pm 2.00$
Unsupervised	$11.92 \pm 2.46$




Fig. 6: Average reconstruction performance as a function of  $g_{\max}$ . The dashed line corresponds to the mean SDR of the test synthetic dataset.

### C. Influence of parameter choices

The loss function (12) depends on a distribution  $p_g$  which in theory could be any unbounded distribution over  $\mathbb{R}_+$ , but in practice, signals have a bounded amplitude, and distribution over a bounded space  $[g_{\min}, g_{\max}]$  (not necessary uniform) may therefore suffice. Parameters have been chosen empirically: if  $g_{\max}$  is too small, the network will not learn; if it is too large, learning becomes unstable, making it difficult to converge. Fig 6 shows how the reconstruction behaves as a function of the value of  $g_{\max}$  for the synthetic dataset with  $v = 0.3$  and  $d = 10$ . The best performance in this case is obtained with  $g_{\max} = 1.5$ .

## V. CONCLUSION

In this study, we propose an unsupervised method that uses scale invariance to address the nonlinear saturation problem. Experimental results indicate that this method can perform on par with the standard supervised approach, but on average, it does not surpass it, since it uses strictly less information for training. The high variance observed in both methods suggests significant variability in the results. Looking ahead, there are several possible directions for future research. Firstly, providing theoretical guarantees for the proposed method could enhance its applicability and understanding. Additionally, integrating additional prior knowledge, such as multi-channel or psychoacoustics, may improve the method’s robustness and

performance in some scenarios. As the saturation problem is not specific to sounds, experiments on other databases such as images can be envisaged. Finally, we envisage extending our method to other non-linear problems beyond signal saturation, such as phase retrieval.

## REFERENCES

- [1] M. Lustig, D. Donoho, and J. M. Pauly, “Sparse mri: The application of compressed sensing for rapid mr imaging,” *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, vol. 58, no. 6, pp. 1182–1195, 2007.
- [2] Y. Shechtman, Y. C. Eldar, O. Cohen, H. N. Chapman, J. Miao, and M. Segev, “Phase retrieval with application to optical imaging: a contemporary overview,” *IEEE signal processing magazine*, vol. 32, no. 3, pp. 87–109, 2015.
- [3] L. I. Rudin, S. Osher, and E. Fatemi, “Nonlinear total variation based noise removal algorithms,” *Physica D: nonlinear phenomena*, vol. 60, no. 1-4, pp. 259–268, 1992.
- [4] J. Lehtinen, J. Munkberg, J. Hasselgren, S. Laine, T. Karras, M. Aittala, and T. Aila, “Noise2noise: Learning image restoration without clean data,” *arXiv preprint arXiv:1803.04189*, 2018.
- [5] D. Chen, J. Tachella, and M. E. Davies, “Equivariant imaging: Learning beyond the range space,” pp. 4379–4388, 2021.
- [6] J. Tachella, D. Chen, and M. E. Davies, “Sensing theorems for unsupervised learning in linear inverse problems,” *J. Mach. Learn. Res.*, vol. 24, pp. 39:1–39:45, 2022.
- [7] P. Závřiska, P. Rajmic, A. Ozerov, and L. Rencker, “A survey and an extensive evaluation of popular audio declipping methods,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 15, no. 1, pp. 5–24, 2020.
- [8] S. Kitić, L. Jacques, N. Madhu, M. P. Hopwood, A. Spriet, and C. De Vleeschouwer, “Consistent iterative hard thresholding for signal declipping,” in *Acoustics, Speech and Signal Processing (ICASSP), 2013 IEEE International Conference on*, IEEE, 2013.
- [9] C. Gaultier, S. Kitić, R. Gribonval, and N. Bertin, “Sparsity-based audio declipping methods: Selected overview, new algorithms, and large-scale evaluation,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 29, pp. 1174–1187, 2021.
- [10] B. Defraene, N. Mansour, S. De Hertogh, T. Van Waterschoot, M. Diehl, and M. Moonen, “Declipping of audio signals using perceptual compressed sensing,” *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 21, no. 12, pp. 2627–2637, 2013.
- [11] A. Ozerov, Ç. Bilen, and P. Pérez, “Multichannel audio declipping,” in *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 659–663, IEEE, 2016.
- [12] C. Gaultier, N. Bertin, and R. Gribonval, “Cascade: Channel-aware structured cospars audio declipper,” in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 571–575, IEEE, 2018.
- [13] A. Krull, T.-O. Buchholz, and F. Jug, “Noise2void-learning denoising from single noisy images,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 2129–2137, 2019.
- [14] J. Tachella and L. Jacques, “Learning to reconstruct signals from binary measurements alone,” *Transactions on Machine Learning Research*, 2023. Featured Certification.
- [15] S. Mohan, Z. Kadkhodaie, E. P. Simoncelli, and C. Fernandez-Granda, “Robust and interpretable blind image denoising via bias-free convolutional neural networks,” *arXiv preprint arXiv:1906.05478*, 2019.
- [16] S. Foucart and T. Needham, “Sparse recovery from saturated measurements,” *Information and Inference: A Journal of the IMA*, vol. 6, no. 2, pp. 196–212, 2017.
- [17] D. Ulyanov, A. Vedaldi, and V. Lempitsky, “Deep image prior,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 9446–9454, 2018.
- [18] N. Vryzas, R. Kotsakis, A. Liatsou, C. A. Dimoulas, and G. Kalliris, “Speech emotion recognition for performance interaction,” *Journal of the Audio Engineering Society*, vol. 66, no. 6, pp. 457–467, 2018.
- [19] N. Vryzas, M. Matsiola, R. Kotsakis, C. Dimoulas, and G. Kalliris, “Subjective evaluation of a speech emotion recognition interaction framework,” in *Proceedings of the Audio Mostly 2018 on Sound in Immersion and Emotion*, pp. 1–7, 2018.