

Learning Maximally Monotone Operators for Image Recovery

Jean-Christophe Pesquet, Audrey Repetti, Matthieu Terris, Yves Wiaux

► To cite this version:

Jean-Christophe Pesquet, Audrey Repetti, Matthieu Terris, Yves Wiaux. Learning Maximally Monotone Operators for Image Recovery. 2020. hal-03087515v1

HAL Id: hal-03087515 https://hal.science/hal-03087515v1

Preprint submitted on 23 Dec 2020 (v1), last revised 20 Dec 2021 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers. L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

LEARNING MAXIMALLY MONOTONE OPERATORS FOR IMAGE RECOVERY*

JEAN-CHRISTOPHE PESQUET[†], AUDREY REPETTI[‡], MATTHIEU TERRIS[§], and YVES WIAUX[¶]

Abstract. We introduce a new paradigm for solving regularized variational problems. These are typically formulated to address ill-posed inverse problems encountered in signal and image processing. The objective function is traditionally defined by adding a regularization function to a data fit term, which is subsequently minimized by using iterative optimization algorithms. Recently, several works have proposed to replace the operator related to the regularization by a more sophisticated denoiser. These approaches, known as plug-and-play (PnP) methods, have shown excellent performance. Although it has been noticed that, under nonexpansiveness assumptions on the denoisers, the convergence of the resulting algorithm is guaranteed, little is known about characterizing the asymptotically delivered solution. In the current article, we propose to address this limitation. More specifically, instead of employing a functional regularization, we perform an operator regularization, where a maximally monotone operator (MMO) is learned in a supervised manner. This formulation is flexible as it allows the solution to be characterized through a broad range of variational inequalities, and it includes convex regularizations as special cases. From an algorithmic standpoint, the proposed approach consists in replacing the resolvent of the MMO by a neural network (NN). We provide a universal approximation theorem proving that nonexpansive NNs provide suitable models for the resolvent of a wide class of MMOs. The proposed approach thus provides a sound theoretical framework for analyzing the asymptotic behavior of first-order PnP algorithms. In addition, we propose a numerical strategy to train NNs corresponding to resolvents of MMOs. We apply our approach to image restoration problems and demonstrate its validity in terms of both convergence and quality.

Key words. Monotone operators, neural networks, convex optimization, plug-and-play methods, inverse problems, computational imaging, nonlinear approximation

AMS subject classifications. 47H05, 90C25, 90C59, 65K10, 49M27, 68T07, 68U10, 94A08.

1. Introduction. In many problems in data science, in particular when dealing with inverse problems, a variational approach is adopted which amounts to

(1.1)
$$\min_{x \in \mathcal{H}} f(x) + g(x)$$

where \mathcal{H} is the underlying data space, here assumed to be a real Hilbert space, $f: \mathcal{H} \to]-\infty, +\infty]$ is a data fit (or data fidelity) term related to some available data z (observations), and $g: \mathcal{H} \to]-\infty, +\infty]$ is some regularization function. The data fit term is often derived from statistical considerations on the observation model through the maximum likelihood principle. For many standard noise distributions,

^{*}Submitted 12/23/2020

Funding: M. Terris would like to thank Heriot-Watt University for the PhD funding. This work was supported by the UK Engineering and Physical Sciences Research Council (EPSRC) under grant EP/T028270/1. The work of J.-C. Pesquet was supported by Institut Universitaire de France and the ANR Chair in AI BRIGEABLE.

[†] Université Paris-Saclay, Inria, Center for Visual Computing, Gif sur Yvette, France (jeanchristophe.pesquet@centralesupelec.fr).

[‡]Second corresponding author. School of Mathematics and Computer Sciences and School of Engineering and Physical Sciences, Heriot-Watt University, Edinburgh, UK. Maxwell Institute for Mathematical Sciences, Bayes Centre, Edinburgh, UK (a.repetti@hw.ac.uk).

[§]First corresponding author. School of Engineering and Physical Sciences, Heriot-Watt University, Edinburgh, UK (mt114@hw.ac.uk).

[¶]School of Engineering and Physical Sciences, Heriot-Watt University, Edinburgh, UK (y.wiaux@hw.ac.uk).

the negative log-likelihood corresponds to a smooth function (e.g. Gaussian, Poisson-Gauss, or logistic distributions). The regularization term is often necessary to avoid overfitting or to overcome ill-posedness problems. A vast literature has been developed on the choice of this term. It often tends to promote the smoothness of the solution or to enforce its sparsity by adopting a functional analysis viewpoint. Good examples of such regularization functions are the total variation semi-norm [51] and its various extensions [12, 25], and penalizations based on wavelet (or "x-let") frame representations [26]. Alternatively, a Bayesian approach can be followed where this regularization is viewed as the negative-log of some prior distribution, in which case the minimizer of the objective function in (1.1) can be understood as a Maximum A Posteriori (MAP) estimator. In any case, the choice of this regularization introduces two main roadblocks. First, the function q has to be chosen so that the minimization problem in (1.1) be tractable, which limits its choice to relatively simple forms. Secondly, the definition of this function involves some parameters which need to be set. The simplest case consists of a single scaling parameter usually called the regularization factor, the choice of which is often very sensitive on the quality of the results. Note that, in some works, this regularization function is the indicator function of some set encoding some smoothness or sparsity constraint. For example, it can model some upper bound on some functional of the discrete gradient of the sought signal, this bound playing then a role equivalent to a regularization parameter [19]. Using an indicator function can also model standard constraints in some image restoration problems, where the image values are bounded [1, 10].

By denoting by $\Gamma_0(\mathcal{H})$ the class of lower-semicontinuous convex functions from \mathcal{H} to $]-\infty, +\infty]$ with a nonempty domain, let us now assume that both f and g belong to $\Gamma_0(\mathcal{H})$. The Moreau subdifferentials of these functions will be denoted by ∂f and ∂g , respectively. Under these convexity assumptions, if

(1.2)
$$0 \in \partial f(x) + \partial g(x),$$

then x is a solution to the minimization problem (1.1). Actually, under mild qualification conditions the sets of solutions to (1.1) and (1.2) coincide [7]. By reformulating the original optimization problem under the latter form, we have moved to the field of variational inequalities. Interestingly, it is a well-established fact that the subdifferential of a function in $\Gamma_0(\mathcal{H})$ is a maximally monotone operator (MMO), which means that (1.2) is a special case of the following monotone inclusion problem:

(1.3) Find
$$x \in \mathcal{H}$$
 such that $0 \in \partial f(x) + A(x)$.

where A is an MMO. We recall that a multivalued operator A defined on \mathcal{H} is maximally monotone if and only if, for every $(x_1, u_1) \in \mathcal{H}^2$,

(1.4)
$$u_1 \in Ax_1 \quad \Leftrightarrow \quad (\forall x_2 \in \mathcal{H})(\forall u_2 \in Ax_2)\langle x_1 - x_2 \mid u_1 - u_2 \rangle \ge 0.$$

Actually the class of monotone inclusion problems is much wider than the class of convex optimization problems and, in particular, includes saddle point problems and game theory equilibria [18]. What is also worth noting is that many existing algorithms for solving convex optimization have their equivalent for solving monotone inclusion problems. This suggests that it is more flexible and probably more efficient, to substitute (1.3) for (1.1) in problems encountered in data science. In other words, instead of performing a functional regularization, we can introduce an operator regularization through the maximally monotone mapping A. Although this extension of

(1.1) may appear both natural and elegant, it induces a high degree of freedom in the choice of the regularization strategy. However, if we except the standard case when $A = \partial g$, it is hard to have a good intuition about how to make a relevant choice for A. To circumvent this difficulty, our proposed approach will consist in learning A in a supervised manner by using some available dataset in the targeted application.

Since a MMO is fully characterized by its resolvent, our approach enters into the family of so-called plug-and-play (PnP) methods [57], where one replaces the proximity operator of an optimization algorithm with a denoiser, e.g. a denoising neural network (NN) [67]. It is worth mentioning that by doing so, any algorithm whose proof is based on MMO theory can be turned into a PnP algorithm, e.g., Forward-Backward (FB), Douglas-Rachford, Peaceman-Rachford, primal-dual approaches, and more [7, 20, 35]. To ensure the convergence of such PnP algorithms, it is known from fixed point theory that (under mild conditions) it is sufficient for the denoiser to be firmly nonexpansive. Unfortunately, most pre-defined denoisers do not satisfy this assumption, and learning a firmly nonexpansive denoiser remains challenging [52, 56]. The main bottleneck is the ability to tightly constrain the Lipschitz constant of a NN. During the last years, several works proposed to control the Lipschitz constant (see e.g. [5, 9, 16, 30, 45, 52, 54, 56, 63]). Nevertheless, only few of them are accurate enough to ensure the convergence of the associated PnP algorithm and often come at the price of strong computational and architectural restrictions (e.g., absence of residual skip connections) [9, 30, 52, 56]. The method proposed in [9] allows a tight control of convolutional layers, but in order to ensure the nonexpansiveness of the resulting architecture, one cannot use residual skip connections, despite their wide use in NNs for denoising applications. In [30], the authors propose to train an averaged NN by projecting the full convolutional layers on the Stiefel manifold and showcase the usage of their network in a PnP algorithm. Yet, the architecture proposed by the authors remains constrained by proximal calculus rules. The assumption [52, Assumption A] introduced by Ryu et al. allowed the authors to propose the first convergent NNbased PnP algorithm in a more general framework, but this assumption is rather non standard and applies only to FB and ADMM. In our previous work [56], we proposed a method to build firmly nonexpansive convolutional NNs; to the best of our knowledge, this was the first method ensuring the firm nonexpansiveness of a denoising NN. However, the resulting architecture was strongly constrained and did not improve over the state-of-the-art. Since building firmly nonexpansive denoisers is difficult, many works on PnP methods leverage ADMM algorithm which may appear easier to handle in practice [52]. At this point, it is worth mentioning that the convergence of ADMM requires restrictive conditions on the involved linear operators [35].

Another drawback of PnP algorithms is that, even if some results exist concerning their convergence to a limit point, little is known about the characterization of this limit point - given that it exists. The regularization by denoising (RED) approach [3, 17], provides a partial answer to this question. By considering a minimum mean square error (MMSE) denoiser, one can link the PnP algorithms based on FB or ADMM to a minimization problem [3, 62]. However, as underlined by the authors, the denoising NN is only an approximation to the MMSE regressor. Eventually, [17] proposes a comprehensive theoretical study of the RED framework under a demicontractivity assumption. This assumption remains however less convenient to check than the standard firm nonexpansiveness condition which allows the convergence of the resulting PnP algorithm to be ensured in a quite versatile context.

Our main contribution is to show that one can train a neural network (NN) so that it corresponds to the resolvent of some MMO. We first explore the theoretical side of the question by stating a universal approximation theorem. Then, we put emphasis on the algorithmic side of the problem. To do so, we propose to regularize the training loss with the spectral norm of the Jacobian of a suitable nonlinear operator. Although the resulting NN could be plugged into a variety of iterative algorithms, our work is focused on the standard FB algorithm. We illustrate the convergence of the corresponding PnP scheme in an image restoration problems. We show that our method compares positively in terms of quality to both state-of-the-art PnP methods and regularized optimization approaches.

This article is organized as follows. In section 2, we recall how MMOs can be mathematically characterized and explain how their resolvent can be modeled by an averaged residual neural network. We also establish that NNs are generic models for a wide class of MMOs. In section 3, we show the usefulness of learning MMOs in the context of plug-and-play (PnP) first-order algorithms employed for solving inverse problems. We also describe the training approach which has been adopted. In section 4, we provide illustrative results for the restoration of monochromatic and color images. Finally, some concluding remarks are made in section 5.

Notation: Throughout the article, we will denote by $\|\cdot\|$ the norm endowing any real Hilbert space \mathcal{H} . The same notation (being clear from the context) will be used to denote the norm of a bounded linear operator L from \mathcal{H} to some real Hilbert space \mathcal{G} , that is $\|L\| = \sup_{x \in \mathcal{H} \setminus \{0\}} \|Lx\| / \|x\|$. The inner product of \mathcal{H} associated to $\|\cdot\|$ will be denoted by $\langle \cdot | \cdot \rangle$, here again without making explicit the associated space. Let D be a subset of \mathcal{H} and $T: D \to \mathcal{H}$. The operator T is μ -Lipschitzian for $\mu > 0$ if, for every $(x, y) \in D^2$, $\|Tx - Ty\| \leq \mu \|x - y\|$. If T is 1-Lipschitzian, its is said to be nonexpansive. The operator T is firmly nonexpansive if, for every $(x, y) \in D^2$, $\|Tx - Ty\|^2 \leq \langle x - y | Tx - Ty \rangle$. Let $A: \mathcal{H} \Rightarrow \mathcal{H}$ be a multivariate operator, i.e., for every $x \in \mathcal{H}$, A(x) is a subset of \mathcal{H} . The graph of A is defined as gra $A = \{(x, u) \in \mathcal{H}^2 | u \in Ax\}$. The operator $A: \mathcal{H} \to 2^{\mathcal{H}}$ is monotone if, for every $(x, u) \in \operatorname{gra} A$ and $(y, v) \in \operatorname{gra} A, \langle x - y | u - v \rangle \geq 0$, and maximally-monotone if (1.4) holds, for every $(x_1, u_1) \in \mathcal{H}^2$. The resolvent of A is $J_A = (\operatorname{Id} + A)^{-1}$, where the inverse is here defined in the sense of the inversion of the graph of the operator. For further details on monotone operator theory, we refer the reader to [7].

2. Neural network models for maximally monotone operators.

2.1. A property of maximally monotone operators. Any multivalued operator operating on \mathcal{H} is fully characterized by its resolvent. A main property for our purpose is the following:

PROPOSITION 2.1. Let $A: \mathcal{H} \rightrightarrows \mathcal{H}$. A is a maximally monotone operator (MMO) if and only if there exists a nonexpansive (i.e. 1-Lipschitzian) operator $Q: \mathcal{H} \rightarrow \mathcal{H}$ such that

(2.1)
$$J_A \colon \mathcal{H} \to \mathcal{H} \colon x \mapsto \frac{x + Q(x)}{2},$$

that is

(2.2)
$$A = 2(\mathrm{Id} + Q)^{-1} - \mathrm{Id}.$$

Proof. This result is a direct consequence of Minty's theorem and the fact that any firmly nonexpansive operator can be expressed as the arithmetic mean of the identity operator and some nonexpansive operator Q (see [7]). (2.2) is deduced by inverting (2.1).

5

The above result means that the class of MMOs can be derived from the class of nonexpansive mappings. The focus should therefore turn on how to model operators in the latter class with neural networks.

2.2. Nonexpansive neural networks. Our objective will be next to derive a parametric model for the nonexpansive operator Q in (2.1). Due to their oustanding approximation capabilities, neural networks appear as good choices for building such models. We will restrict our attention to feedforward NNs.

MODEL 2.2. Let $(\mathcal{H}_m)_{0 \leq m \leq M}$ be real Hilbert spaces such that $\mathcal{H}_0 = \mathcal{H}_M = \mathcal{H}$. A feedforward NN having m layer and both input and ouput in \mathcal{H} can be seen as a composition of operators:

$$(2.3) Q = T_M \cdots T_1,$$

where

$$(2.4) \qquad (\forall m \in \{1, \dots, M\}) \quad T_m \colon \mathcal{H}_{m-1} \to \mathcal{H}_m \colon x \mapsto R_m(W_m x + b_m).$$

At each layer $m \in \{1, ..., M\}$, $R_m: \mathcal{H}_m \to \mathcal{H}_m$ is a nonlinear activation operator, $W_m: \mathcal{H}_{m-1} \to \mathcal{H}_m$ is a bounded linear operator corresponding to the weights of the network, and $b_m \in \mathcal{H}_m$ is a bias parameter vector.

In the remainder, we will use the following notation:

NOTATION 2.3. Let V and V' be nonempty subsets of some Euclidean space and let $\mathcal{N}_{\mathcal{F}}(V, V')$ denote the class of nonexpansive feedforward NNs with inputs in V and outputs in V' built from a given dictionary \mathcal{F} of allowable activation operators.

Also, we will make the following assumption:

ASSUMPTION 2.4. The identity operator as well as the sorting operator performed on blocks of size 2 belong to dictionary \mathcal{F} .

In other words, a network in $\mathcal{N}_{\mathcal{F}}(V, V')$ can be linear, or it can be built by using maxpooling with blocksize 2 and any other kind of activation function, say some given function $\rho \colon \mathbb{R} \to \mathbb{R}$, operating componentwise in some of its layers, provided that the resulting structure is 1-Lipschitzian.

The main difficulty is to design such a feedforward NN so that Q in (2.3) has a Lipschitz constant smaller or equal to 1. An extensive literature has been devoted to the estimation of Lipschitz constants of NNs [5, 53, 55], but the main goal was different from ours since these works were motivated by robustness issues in the presence of adversarial perturbations [28, 36, 48, 55]. Based on the results in [23], useful sufficient conditions for a NN to be nonexpansive are given below:

PROPOSITION 2.5. Let Q be a feedforward NN as defined in Model 2.2. Assume that, for every $m \in \{1, \ldots, M\}$, R_m is α_m -averaged with $\alpha_m \in [0, 1]$. Then Q is nonexpansive if one of the following conditions holds:

- (i) $||W_1|| \cdots ||W_M|| \leq 1;$
- (ii) for every $m \in \{1, ..., M-1\}$, $\mathcal{H}_m = \mathbb{R}^{K_m}$ with $K_m \in \mathbb{N} \setminus \{0\}$, R_m is a separable activation operator, in the sense that there exist real-valued one-variable functions $(\rho_{m,k})_{1 \leq k \leq K_m}$ such that, for every $x = (\xi_k)_{1 \leq k \leq K_m} \in \mathcal{H}_m$, $R_m(x) = (\rho_{m,k}(\xi_k))_{1 \leq k \leq K_m}$, and

(2.5)
$$(\forall \Lambda_1 \in \mathcal{D}_{1,\{1-2\alpha_1,1\}}) \dots (\forall \Lambda_{M-1} \in \mathcal{D}_{M-1,\{1-2\alpha_{M-1},1\}})$$

 $\|W_M \Lambda_{M-1} \cdots \Lambda_1 W_1\| \leq 1,$



Fig. 1: Neural network modelling the resolvent of a maximally monotone operator. The weight operators $(W_m)_{1 \leq m \leq M}$ have to be set according to the conditions provided in Proposition 2.5.

where, for every $m \in \{1, ..., M-1\}$, $\mathcal{D}_{m,\{1-2\alpha_m,1\}}$ denotes the set of diagonal $K_m \times K_m$ matrices with diagonal terms equal either to $1-2\alpha_m$ or 1;

(iii) for every $m \in \{1, ..., M\}$, $\mathcal{H}_m = \mathbb{R}^{K_m}$ with $K_m \in \mathbb{N} \setminus \{0\}$ and W_m is a matrix with nonnegative elements, $(R_m)_{1 \leq m \leq M-1}$ are separable activation operators, and

$$(2.6) ||W_M \cdots W_1|| \leqslant 1$$

Note that the α -averageness assumption on $(R_m)_{1 \leq m \leq M-1}$ means that, for every $m \in \{1, \ldots, M-1\}$, there exists a nonexpansive operator $\widetilde{R}_m \colon \mathcal{H}_m \to \mathcal{H}_m$ such that $R_m = (1 - \alpha_m) \operatorname{Id} + \alpha_m \widetilde{R}_m$. Actually, most of the activation operators employed in neural networks (ReLU, leaky ReLU, sigmoid, softmax,...) satisfy this assumption with $\alpha_m = 1/2$ [21]. A few others like the sorting operator used in max-pooling correspond to a value of the constant α_m larger than 1/2 [23]. It is also worth mentioning that, although Condition (i) in Proposition 2.5 is obviously the simplest one, it is usually quite restrictive, the weakest condition being given by (2.6) which requires yet the network weights to be nonnegative.

By summarizing the results of the previous section, Figure 1 shows a feedforward NN architecture for MMOs, for which Proposition 2.5 can be applied. It can be noticed that (2.1) induces the presence of a skip connection in the global structure.

2.3. Stationary maximally monotone operators. In the remainder, we will focus our attention on a particular subclass of operators.

DEFINITION 2.6. Let $(\mathcal{H}_k)_{1 \leq k \leq K}$ be real Hilbert spaces. An operator A defined on the product space space $\mathcal{H} = \mathcal{H}_1 \times \cdots \times \mathcal{H}_K$ will be said to be a stationary MMO if its resolvent J_A is an operator from \mathcal{H} to \mathcal{H} such that, for every $k \in \{1, \ldots, K\}$, there exists a bounded linear operator $\Pi_k \colon \mathcal{H} \to \mathcal{H}_k$ and a self-adjoint nonnegative operator $\Omega_k \colon \mathcal{H} \to \mathcal{H}$ such that

(2.7)
$$(\forall (x,y) \in \mathcal{H}^2) \quad \|\Pi_k (2J_A(x) - x - 2J_A(y) + y)\|^2 \leq \langle x - y \mid \Omega_k(x - y) \rangle$$

with

(2.8)
$$\sum_{k=1}^{K} \Pi_k^* \Pi_k = \mathrm{Id}$$

(2.9)
$$\left\|\sum_{k=1}^{K} \Omega_k\right\| \leqslant 1.$$

Immediate consequences of this definition are given below. In particular, we will see that stationary MMOs define a subclass of the set of MMOs. PROPOSITION 2.7. Let $(\mathcal{H}_k)_{1 \leq k \leq K}$ be real Hilbert spaces and let $\mathcal{H} = \mathcal{H}_1 \times \cdots \times \mathcal{H}_K$. Let $A: \mathcal{H} \rightrightarrows \mathcal{H}$.

- (i) If A is a stationary MMO on \mathcal{H} , then it is maximally monotone.
- (ii) Assume that (2.8) is satisfied where, for every $k \in \{1, ..., K\}$, $\Pi_k : \mathcal{H} \to \mathcal{H}_k$ is a bounded linear operator. If $\operatorname{ran}(A + \operatorname{Id}) = \mathcal{H}$ and (2.10)

$$(\forall (p,q) \in \mathcal{H}^2)(\forall p' \in A(p))(\forall q' \in A(q)) \ \langle \Pi_k(p-q) \mid \Pi_k(p'-q') \rangle \ge 0,$$

then A is a stationary MMO.

Proof. (i): Let A be a stationary MMO defined on \mathcal{H} . Summing over k in (2.7) yields, for every $(x, y) \in \mathcal{H}^2$,

$$(2.11) \quad \left\langle 2J_A(x) - x - 2J_A(y) + y \mid \left(\sum_{k=1}^K \Pi_k^* \Pi_k\right) (2J_A(x) - x - 2J_A(y) + y) \right\rangle$$
$$\leqslant \left\langle x - y \mid \sum_{k=1}^K \Omega_k(x - y) \right\rangle.$$

It thus follows from (2.8), (2.9), and the nonnegativity of $(\Omega_k)_{1 \le k \le K}$ that

(2.12)
$$\|2J_A(x) - x - 2J_A(y) + y\|^2 \leq \left\|\sum_{k=1}^K \Omega_k\right\| \|x - y\|^2 \leq \|x - y\|^2.$$

This shows that $2J_A$ – Id is a nonexpansive operator. Hence, based on Proposition 2.5, A is an MMO.

(ii): Let k be an arbitrary integer in $\{1, \ldots, K\}$. (2.10) can be reexpressed as

$$(2.13) \quad (\forall (p,q) \in \mathcal{H}^2)(\forall p' \in A(p))(\forall q' \in A(q))) \\ \langle \Pi_k^* \Pi_k(p-q) \mid p'-q'+p-q \rangle \ge \langle \Pi_k^* \Pi_k(p-q) \mid p-q \rangle.$$

In particular, this inequality holds if $p \in J_A(x)$ and $q \in J_A(y)$ where x and y are arbitrary elements of \mathcal{H} . Then, by definition of J_A , we have $x - p \in A(p), y - q \in A(q)$, and (2.13) yields

(2.14)
$$\langle \Pi_k^* \Pi_k(p-q) \mid x-y \rangle \ge \langle \Pi_k^* \Pi_k(p-q) \mid p-q \rangle$$

By summing over k and using (2.8), it follows that J_A is firmly nonexpansive and it is thus single valued. (2.14) is then equivalent to

(2.15)
$$\|\Pi_k (2J_A(x) - x - 2J_A(y) + y)\|^2 \leq \langle x - y \mid \Pi_k^* \Pi_k (x - y) \rangle.$$

This shows that Inequality (2.7) holds with $\Omega_k = \prod_k^* \prod_k$. Since (2.9) is then obviously satisfied, A is a stationary MMO.

A natural question at this point is: how generic are stationary MMOs? To provide a partial answer to this question, we feature a few examples of such operators.

EXAMPLE 2.8. For every $k \in \{1, ..., K\}$, let B_k be an MMO defined on a real Hilbert space \mathcal{H}_k and let B be the operator defined as

$$(2.16) \quad (\forall x = (x^{(k)})_{1 \leq k \leq K} \in \mathcal{H} = \mathcal{H}_1 \times \dots \times \mathcal{H}_K) \quad B(x) = B_1(x^{(1)}) \times \dots \times B_K(x^{(K)}).$$

Let $U: \mathcal{H} \to \mathcal{H}$ be a unitary linear operator. Then $A = U^*BU$ is a stationary MMO.

Proof. As B is an MMO and U is surjective, U^*BU is an MMO [7, Corollary 25.6]. We are thus guaranteed that ran $(\mathrm{Id} + A) = \mathcal{H}$ [7, Theorem 21.1]. For every $k \in \{1, \ldots, K\}$, let

(2.17)
$$D_k \colon \mathcal{H} \to \mathcal{H}_k \colon (x^{(\ell)})_{1 \leq \ell \leq K} \mapsto x^{(k)}$$

(2.18)
$$\Pi_k = D_k U.$$

It can be noticed that

(2.19)
$$\sum_{k=1}^{K} \Pi_{k}^{*} \Pi_{k} = U^{*} U = \text{Id}.$$

Let $(p,q) \in \mathcal{H}^2$. Every $(p',q') \in A(p) \times A(q)$ is such

$$(2.20) p' = U^* r$$

(2.21)
$$q' = U^* s,$$

where $r \in B(Up)$ and $s \in B(Uq)$. Using (2.18), (2.20), and (2.21) yield, for every $k \in \{1, \ldots, K\},\$

(2.22)
$$\langle \Pi_k(p-q) \mid \Pi_k(p'-q') \rangle = \langle D_k U p - D_k U q \mid D_k r - D_k s \rangle.$$

Because of the separable form of $B, D_k r \in B_k(D_k U p)$ and $D_k s \in B_k(D_k U q)$. It then follows from (2.22) and the monotonicity of B_k that

(2.23)
$$\langle \Pi_k(p-q) \mid \Pi_k(p'-q') \rangle \ge 0.$$

By invoking Proposition 2.7(ii), we conclude that A is a stationary MMO.

EXAMPLE 2.9. For every $k \in \{1, \ldots, K\}$, let $\varphi_k \in \Gamma_0(\mathbb{R})$, and let the function g be defined as

(2.24)
$$(\forall x = (x^{(k)})_{1 \le k \le K} \in \mathbb{R}^K) \quad g(x) = \sum_{k=1}^K \varphi_k(x^{(k)}).$$

Let $U \in \mathbb{R}^{K \times K}$ be an orthogonal matrix. Then the subdifferential of $g \circ U$ is a stationary MMO.

Proof. This corresponds to the special case of Example 2.8 when, for every $k \in$ $\{1, \ldots, K\}, \mathcal{H}_k = \mathbb{R} \text{ (see } [7, \text{ Theorem } 16.47, \text{Corollary } 22.23]).$

EXAMPLE 2.10. Let $(\mathcal{H}_k)_{1 \leq k \leq K}$ be real Hilbert spaces and let B be a bounded linear operator from $\mathcal{H} = \mathcal{H}_1 \times \cdots \times \mathcal{H}_K$ to \mathcal{H} such that one of the following conditions holds:

(i) $B + B^*$ is nonnegative

(ii) B is skewed

(iii) B is cocoercive.

Let $c \in \mathcal{H}$. Then the affine operator $A: \mathcal{H} \to \mathcal{H}: x \mapsto Bx + c$ is a stationary MMO.

Proof. If $B + B^*$ is nonnegative, B, hence A, are maximally monotone and $J_A =$ $J_B(\cdot - c)$ is firmly nonexpansive. As a consequence, the reflected resolvent of B

(2.25)
$$Q = 2(\mathrm{Id} + B)^{-1} - \mathrm{Id}$$

(2.18)

is nonexpansive. For every $k \in \{1, ..., K\}$, let D_k be the decimation operator defined in (2.17) and let

$$(2.26) \Pi_k = D_k$$

(2.27)
$$\Omega_k = Q^* D_k^* D_k Q$$

 Π_k satisfies (2.8) and, since

(2.28)
$$\left\|\sum_{k=1}^{K} \Omega_{k}\right\| = \|Q^{*}Q\| = \|Q\|^{2} \leq 1,$$

(2.9) is also satisfied. In addition, for every $(x, y) \in \mathcal{H}^2$ and, for every $k \in \{1, \ldots, K\}$, we have

(2.29)
$$\begin{aligned} \|\Pi_k (2J_A(x) - x - 2J_A(y) + y)\|^2 \\ &= \|\Pi_k (2J_B(x - c) - x + c - 2J_B(y - c) + y - c)\|^2 \\ &= \langle x - y \mid \Omega_k(x - y) \rangle, \end{aligned}$$

which shows that A is a stationary MMO.

Note finally that, if B is skewed or coccercive linear operator, then $B + B^*$ is non-negative.

EXAMPLE 2.11. Let $(\mathcal{H}_k)_{1 \leq k \leq K}$ be real Hilbert spaces, let $\mathcal{H} = \mathcal{H}_1 \times \cdots \times \mathcal{H}_K$, and let $A: \mathcal{H} \rightrightarrows \mathcal{H}$ be a stationary MMO. Then its inverse A^{-1} is a stationary MMO.

Proof. The resolvent of A^{-1} is given by $J_{A^{-1}} = \operatorname{Id} - J_A$. In addition, since A is stationary, there exist bounded linear operators $(\Pi_k)_{1 \leq k \leq K}$ and self-adjoint operators $(\Omega_k)_{1 \leq k \leq K}$ satisfying (2.7)-(2.9). For every $k \in \{1, \ldots, K\}$, we have then, for every $(x, y) \in \mathcal{H}^2$,

$$\|\Pi_k (2J_{A^{-1}}(x) - x - 2J_{A^{-1}}(y) + y)\|^2 = \|\Pi_k (2J_A(y) - y - 2J_A(x) + x)\|^2$$
(2.30) $\leq \langle y - x \mid \Omega_k(y - x) \rangle.$

EXAMPLE 2.12. Let $(\mathcal{H}_k)_{1 \leq k \leq K}$ be real Hilbert spaces, let $\mathcal{H} = \mathcal{H}_1 \times \cdots \times \mathcal{H}_K$, and let $A: \mathcal{H} \Rightarrow \mathcal{H}$ be a stationary MMO. Then, for every $\rho \in \mathbb{R} \setminus \{0\}$, $\rho A(\cdot/\rho)$ is a stationary MMO.

Proof. $B = \rho A(\cdot/\rho)$ is maximally monotone and its resolvent reads $J_B = \rho J_A(\cdot/\rho)$ [7, Corollary 23.26]. Using the same notation as previously, for every $k \in \{1, \ldots, K\}$ and for every $(x, y) \in \mathcal{H}^2$,

$$\|\Pi_k \left(2J_B(x) - x - 2J_B(y) + y\right)\|^2 = \rho^2 \left\|\Pi_k \left(2J_A \left(\frac{x}{\rho}\right) - \frac{x}{\rho} - 2J_A \left(\frac{y}{\rho}\right) + \frac{y}{\rho}\right)\right\|^2$$

$$(2.31) \qquad \leqslant \langle y - x \mid \Omega_k(y - x) \rangle. \qquad \Box$$

2.4. Universal approximation theorem. In this section we provide one of the main contributions of this article, consisting in a universal approximation theorem for MMOs defined on $\mathcal{H} = \mathbb{R}^{K}$. To this aim, we first need to introduce useful results, starting by recalling the definition of a lattice.

DEFINITION 2.13. A set \mathcal{L}_E of functions from a set E to \mathbb{R} is said to be a lattice if, for every $(h^{(1)}, h^{(2)}) \in \mathcal{L}_E^2$, $\min\{h^{(1)}, h^{(2)}\}$ and $\max\{h^{(1)}, h^{(2)}\}$ belong to \mathcal{L}_E . A sub-lattice of \mathcal{L}_E is a lattice included in \mathcal{L}_E . This notion of lattice is essential in the variant of the Stone-Weierstrass theorem provided below.

PROPOSITION 2.14. [5] Let (E, d) be a compact metric space with at least two distinct points. Let \mathcal{L}_E be a sub-lattice of $\operatorname{Lip}_1(E, \mathbb{R})$, the class of 1-Lipschtzian (i.e. nonexpansive) functions from E to \mathbb{R} . Assume that, for every $(u, v) \in E^2$ with $u \neq v$ and, for every $(\zeta, \eta) \in \mathbb{R}^2$ such that $|\zeta - \eta| \leq d(u, v)$, there exists a function $h \in \mathcal{L}_E$ such that $h(u) = \zeta$ and $h(v) = \eta$. Then \mathcal{L}_E is dense in $\operatorname{Lip}_1(E, \mathbb{R})$ for the uniform norm.

This allows us to derive the following approximation result that will be instrumental to prove our main result.

COROLLARY 2.15. Let V be a subspace of \mathbb{R}^K and let $h \in \text{Lip}_1(V, \mathbb{R})$. Let E be a compact subset of V. Then, for every $\varepsilon \in [0, +\infty[$, there exists $h_{\varepsilon} \in \mathcal{N}_{\mathcal{F}}(V, \mathbb{R})$, where \mathcal{F} is any dictionary of activation function satisfying Assumption 2.4, such that

(2.32)
$$(\forall x \in E) \quad |h(x) - h_{\varepsilon}(x)| \leq \varepsilon.$$

Proof. First note that $\mathcal{N}_{\mathcal{F}}(V, \mathbb{R})$ is a lattice. Indeed, if $h^{(1)}: V \to \mathbb{R}$ and $h^{(2)}: V \to \mathbb{R}$ are 1-Lipschitzian, then $\min\{h^{(1)}, h^{(2)}\}$ and $\max\{h^{(1)}, h^{(2)}\}$ are 1-Lipschitzian. In addition, if $h^{(1)}$ and $h^{(2)}$ are elements in $\mathcal{N}_{\mathcal{F}}(V, \mathbb{R})$, then by applying sorting operations on the two outputs of these two networks, $\min\{h^{(1)}, h^{(2)}\}$ and $\max\{h^{(1)}, h^{(2)}\}$ are generated. Each of these outputs can be further selected by applying weight matrices either equal to $[1 \ 0]$ or $[0 \ 1]$ as a last operation, so leading to a NN in $\mathcal{N}_{\mathcal{F}}(V, \mathbb{R})$.

Let E be a compact subset of V. Assume that E has at least two distinct points. Since $\mathcal{N}_{\mathcal{F}}(V,\mathbb{R})$ is a lattice, the set of restrictions to E of elements in $\mathcal{N}_{\mathcal{F}}(V,\mathbb{R})$ is a sub-lattice \mathcal{L}_E of $\operatorname{Lip}_1(E,\mathbb{R})$. In addition, let $(u,v) \in E^2$ with $u \neq v$ and let $(\zeta,\eta) \in \mathbb{R}^2$ be such that $|\zeta - \eta| \leq ||u - v||$. Set $h: V \to \mathbb{R}: x \mapsto w^{\top}(x - v) + \eta$ where $w = (\zeta - \eta)(u - v)/||u - v||^2$. Since $||w|| = |\zeta - \eta|/||u - v|| \leq 1$, h is a linear network in $\mathcal{N}_{\mathcal{F}}(V,\mathbb{R})$ and we have $h(u) = \zeta$ and $h(v) = \eta$. This shows that the restriction of hto E is an element of \mathcal{L}_E satisfying the assumptions of Proposition 2.14. It can thus be deduced from this proposition that (2.32) holds.

The inequality also trivially holds if E reduces to a single point x since it is always possible to find a linear network in $\mathcal{N}_{\mathcal{F}}(V, \mathbb{R})$ whose output equals h(x).

REMARK 2.16. This result is valid whatever the norm used on V.

We are now able to state a universal approximation theorem for MMOs defined on $\mathcal{H} = \mathbb{R}^{K}$ (i.e., for every $k \in \{1, \ldots, K\}$, $\mathcal{H}_{k} = \mathbb{R}$ in Definition 2.6).

THEOREM 2.17. Let $\mathcal{H} = \mathbb{R}^{K}$. Let $A: \mathcal{H} \Rightarrow \mathcal{H}$ be a stationary MMO. For every compact set $S \subset \mathcal{H}$ and every $\epsilon \in [0, +\infty[$, there exists a NN $Q_{\epsilon} \in \mathcal{N}_{\mathcal{F}}(\mathcal{H}, \mathcal{H})$, where \mathcal{F} is any dictionary of activation function satisfying Assumption 2.4, such that $A_{\epsilon} = 2(\mathrm{Id} + Q_{\epsilon})^{-1} - \mathrm{Id}$ satisfies the following properties.

- (i) For every $x \in S$, $||J_A(x) J_{A_{\epsilon}}(x)|| \leq \epsilon$.
- (ii) Let $x \in \mathcal{H}$ and let $y \in A(x)$ be such that $x + y \in S$. Then, there exists $x_{\epsilon} \in \mathcal{H}$ and $y_{\epsilon} \in A_{\epsilon}(x_{\epsilon})$ such that $||x - x_{\epsilon}|| \leq \epsilon$ and $||y - y_{\epsilon}|| \leq \epsilon$.

Proof. (i): If $A: \mathbb{R}^K \Rightarrow \mathbb{R}^K$ is a stationary MMO then it follows from Propositions 2.1 and 2.7(i) and that there exists a nonexpansive operator $Q: \mathbb{R}^K \to \mathbb{R}^K$ such that $J_A = (\mathrm{Id} + Q)/2$. In addition, according to Definition 2.6, there exist vectors $(p_k)_{1 \leq k \leq K}$ in \mathbb{R}^K such that, for every $k \in \{1, \ldots, K\}$,

(2.33)
$$(\forall (x,y) \in \mathcal{H}^2) \quad |\langle p_k \mid Q(x) - Q(y) \rangle|^2 \leq \langle x - y \mid \Omega_k(x - y) \rangle$$

where

(2.34)
$$\sum_{k=1}^{K} p_k p_k^{\top} = \mathrm{Id}$$

and $(\Omega_k)_{1 \leq k \leq K}$ are positive semidefinite matrices in $\mathbb{R}^{K \times K}$ satisfying (2.9).

Set $k \in \{1, ..., K\}$ and define $h_k : x \mapsto \langle p_k \mid Q(x) \rangle$. Let V_k be the nullspace of Ω_k and let V_k^{\perp} be its orthogonal space. We distinguish the cases when $V_k^{\perp} \neq \{0\}$ and when $V_k^{\perp} = \{0\}$. Assume that $V_k^{\perp} \neq \{0\}$. It follows from (2.33) that, for every $x \in V_k^{\perp}$ and $(y, z) \in V_k^2$,

(2.35)
$$h_k(x+y) = h_k(x+z) = \tilde{h}_k(x)$$

where $\widetilde{h}_k \colon V_k^{\perp} \to \mathbb{R}$ is such that

(2.36)
$$(\forall (x, x') \in (V_k^{\perp})^2) \quad |\widetilde{h}_k(x) - \widetilde{h}_k(x')| \leq ||x - x'||_{\Omega_k}$$

and $(\forall x \in \mathbb{R}^K) \|x\|_{\Omega_k} = \langle x \mid \Omega_k x \rangle^{1/2}$. $\|\cdot\|_{\Omega_k}$ defines a norm on V_k^{\perp} . Inequality (2.36) shows that \tilde{h}_k is 1-Lipschitzian on V_k^{\perp} equipped with this norm. Let S be a compact subset of \mathbb{R}^K and let $\operatorname{proj}_{V_k^{\perp}}$ be the orthogonal projection onto V_k^{\perp} . $E_k = \operatorname{proj}_{V_k^{\perp}}(S)$ is a compact set and, in view of Corollary 2.15, for every $\epsilon \in \mathbb{R}$, there exists $\tilde{h}_{k,\epsilon} \in \mathcal{N}_{\mathcal{F}}(V_k^{\perp},\mathbb{R})$ such that

(2.37)
$$(\forall x \in E_k) \quad |\tilde{h}_k(x) - \tilde{h}_{k,\epsilon}(x)| \leq \frac{2\epsilon}{\sqrt{K}}.$$

Set now $h_{k,\epsilon} = \tilde{h}_{k,\epsilon} \circ \operatorname{proj}_{V_k^{\perp}}$. According to (2.35) and (2.37), we have

$$(\forall x \in S) \quad |h_k(x) - h_{k,\epsilon}(x)| = |h_k(\operatorname{proj}_{V_k}(x) + \operatorname{proj}_{V_k^{\perp}}(x)) - h_{k,\epsilon}(\operatorname{proj}_{V_k}(x) + \operatorname{proj}_{V_k^{\perp}}(x))| = |\tilde{h}_k(\operatorname{proj}_{V_k^{\perp}}(x)) - \tilde{h}_{k,\epsilon}(\operatorname{proj}_{V_k^{\perp}}(x))|$$

$$(2.38) \qquad \leqslant \frac{2\epsilon}{\sqrt{K}}.$$

In addition, by using the Lipschitz property of $\tilde{h}_{k,\epsilon}$ with respect to norm $\|\cdot\|_{\Omega_k}$, for every $(x, x') \in \mathbb{R}^K$,

$$\begin{aligned} \left(h_{k,\epsilon}(x) - h_{k,\epsilon}(x')\right)^2 \\ &= \left(\tilde{h}_{k,\epsilon}(\operatorname{proj}_{V_k^{\perp}}(x)) - \tilde{h}_{k,\epsilon}(\operatorname{proj}_{V_k^{\perp}}(x'))\right)^2 \\ &\leqslant \|\operatorname{proj}_{V_k^{\perp}}(x) - \operatorname{proj}_{V_k^{\perp}}(x')\|_{\Omega_k}^2 \\ &= \left\langle \operatorname{proj}_{V_k^{\perp}}(x - x') \mid \Omega_k \operatorname{proj}_{V_k^{\perp}}(x - x') \right\rangle \\ &= \left\langle \Omega_k^{1/2} \operatorname{proj}_{V_k^{\perp}}(x - x') \mid \Omega_k^{1/2} \operatorname{proj}_{V_k^{\perp}}(x - x') \right\rangle \\ &= \left\langle x - x' \mid \Omega_k(x - x') \right\rangle. \end{aligned}$$

(2.39)

If $V_k^{\perp} = \{0\}$, then it follows from (2.33) that $h_k = 0$. Therefore, (2.38) holds with $h_{k,\epsilon} = 0$, which belongs to $\mathcal{N}_{\mathcal{F}}(V_k^{\perp}, \mathbb{R})$ and obviously satisfies (2.39).

Condition (2.34) means that $(p_k)_{1 \leq k \leq K}$ is an orthornormal basis of \mathbb{R}^K in the standard Euclidean metric. This implies that

(2.40)
$$(\forall x \in \mathbb{R}^K) \quad Q(x) = \sum_{k=1}^K h_k(x) \, p_k.$$

 Set

(2.41)
$$(\forall x \in \mathbb{R}^K) \quad Q_{\epsilon}(x) = \sum_{k=1}^K h_{k,\epsilon}(x) \, p_k.$$

It follows from (2.39) and (2.9) that, for every $(x, x') \in (\mathbb{R}^K)^2$,

$$||Q_{\epsilon}(x) - Q_{\epsilon}(x')||^{2}$$

$$= \sum_{k=1}^{K} \left(h_{k,\epsilon}(x) - h_{k,\epsilon}(x')\right)^{2}$$

$$\leqslant \sum_{k=1}^{K} \left\langle x - x' \mid \Omega_{k}(x - x') \right\rangle$$

$$\leqslant ||x - x'||^{2},$$
(2.42)

which shows that $Q_{\epsilon} \in \text{Lip}_1(\mathbb{R}^K, \mathbb{R}^K)$. In addition since, for every $x \in \mathbb{R}^K$,

(2.43)
$$Q_{\epsilon}(x) = W[h_{1,\epsilon}(x), \dots, h_{K,\epsilon}(x)]^{\top}$$

with $W = [p_1, \ldots, p_K]$ and, for every $k \in \mathbb{N}$, $h_{k,\epsilon} \in \mathcal{N}_{\mathcal{F}}(\mathbb{R}^K, \mathbb{R})$, Q_{ϵ} belongs to $\mathcal{N}_{\mathcal{F}}(\mathbb{R}^K, \mathbb{R}^K)$. Let $A_{\epsilon} = 2(\mathrm{Id} + Q_{\epsilon})^{-1} - \mathrm{Id}$. We finally deduce from (2.38) that, for every $x \in S$,

(2.44)
$$\|J_{A}(x) - J_{A_{\epsilon}}(x)\|^{2} = \left\|\frac{x + Q(x)}{2} - \frac{x + Q_{\epsilon}(x)}{2}\right\|^{2} = \frac{1}{4} \sum_{k=1}^{K} \left(h_{k}(x) - h_{k,\epsilon}(x)\right)^{2} \leqslant \epsilon^{2}.$$

(ii): Let $(x, y) \in (\mathbb{R}^K)^2$. We have

(2.45)
$$y \in A(x) \quad \Leftrightarrow \quad x = J_A(x+y).$$

Assume that $x + y \in S$. It follows from (i) that there exists $x_{\epsilon} \in \mathbb{R}^{K}$ such that $x_{\epsilon} = J_{A_{\epsilon}}(x + y)$ and $||x - x_{\epsilon}|| \leq \epsilon$. Let $y_{\epsilon} = x - x_{\epsilon} + y$. We have $x_{\epsilon} = J_{A_{\epsilon}}(x_{\epsilon} + y_{\epsilon})$, that is $y_{\epsilon} \in A_{\epsilon}(x_{\epsilon})$. In addition, $||y - y_{\epsilon}|| = ||x - x_{\epsilon}|| \leq \epsilon$.

We will next show that Theorem 2.17 extends to a wider class of MMOs.

COROLLARY 2.18. Let $\mathcal{H} = \mathbb{R}^K$. Let $(\omega_i)_{1 \leq i \leq I} \in [0,1]^I$ be such that $\sum_{i=1}^I \omega_i = 1$. For every $i \in \{1, \ldots, I\}$, let $A_i \colon \mathcal{H} \rightrightarrows \mathcal{H}$ be a stationary MMO. Then the same properties as in Theorem 2.17 hold if $A \colon \mathcal{H} \rightrightarrows \mathcal{H}$ is the MMO with resolvent $J_A = \sum_{i=1}^I \omega_i J_{A_i}$. *Proof.* First note that $J_A: \mathcal{H} \to \mathcal{H}$ is firmly nonexpansive [7, Proposition 4.6]), hence A is indeed an MMO. As a consequence of Theorem 2.17, for every compact set $S \subset \mathcal{H}$ and every $\epsilon \in]0, +\infty[$, there exist NNs $(Q_{i,\epsilon})_{1 \leq i \leq I}$ in $\mathcal{N}_{\mathcal{F}}(\mathcal{H}, \mathcal{H})$ such that $(A_{i,\epsilon})_{1 \leq i \leq I} = (2(\mathrm{Id} + Q_{i,\epsilon})^{-1} - \mathrm{Id})_{1 \leq i \leq I}$ satisfy:

(2.46)
$$(\forall i \in \{1, \dots, Q\}) (\forall x \in S) \quad ||J_{A_i}(x) - J_{A_{i,\epsilon}}(x)|| \leq \epsilon.$$

Let $Q_{\epsilon} = \sum_{i=1}^{I} \omega_i Q_{i,\epsilon}$. Then $Q_{\epsilon} \in \operatorname{Lip}_1(\mathbb{R}^K, \mathbb{R}^K)$ and, since it is built from a linear combination of the outputs of I NNs in $\mathcal{N}_{\mathcal{F}}(\mathcal{H}, \mathcal{H})$ driven with the same input, it belongs to $\mathcal{N}_{\mathcal{F}}(\mathcal{H}, \mathcal{H})$. In addition, $A_{\epsilon} = 2(\operatorname{Id} + Q_{\epsilon})^{-1} - \operatorname{Id}$ is such that

(2.47)
$$J_{A_{\epsilon}} = \frac{1}{2} \Big(\sum_{i=1}^{I} \omega_i Q_{i,\epsilon} + \mathrm{Id} \Big) = \sum_{i=1}^{I} \omega_i J_{A_{i,\epsilon}},$$

which allows us to deduce from (2.46) that

(2.48)
$$(\forall x \in S) \quad \|J_A(x) - J_{A_{\epsilon}}(x)\| \leq \sum_{i=1}^{I} \omega_i \|J_{A_i}(x) - J_{A_{i,\epsilon}}(x)\| \leq \epsilon.$$

The rest of the proof follows the same line as for Theorem 2.17.

REMARK 2.19. The above results are less accurate than standard universal approximations ones which, for example, guarantee an arbitrary close approximation to any continuous function with a network having only one hidden layer [32, 38]. Indeed, the requirement that the resolvent of a MMO must be firmly nonexpansive induces some significant increase of the difficulty of the mathematical problem. None-theless, the firm nonexpansiveness will enable us to build convergent PnP algorithms described in the next sections.

3. Proposed algorithm.

3.1. Forward-backward algorithm. Let us now come back to problems of the form (1.3). Such monotone inclusion problems can be tackled by a number of algorithms [18, 22], which are all grounded on the use of the resolvent of A (or a scaled version of this operator). For simplicity, let us assume that f is a smooth function. In this case, a famous algorithm for solving (1.3) is the forward-backward (FB) algorithm [14, 24], which is expressed as

(3.1)
$$(\forall n \in \mathbb{N}) \quad x_{n+1} = J_{\gamma A} (x_n - \gamma \nabla f(x_n))$$

where $\gamma > 0$. If a neural network \widetilde{J} is used to approximate $J_{\gamma A}$, then a natural substitute for (3.1) is

(3.2)
$$(\forall n \in \mathbb{N}) \quad x_{n+1} = J(x_n - \gamma \nabla f(x_n)).$$

The following convergence result then straightforwardly follows from standard asymptotic properties of the FB algorithm [24].

PROPOSITION 3.1. Let $\mu \in [0, +\infty[$ and let $\gamma \in [0, 2/\mu[$. Let $f: \mathcal{H} \to \mathbb{R}$ be a convex differentiable function with μ -Lipschitzian gradient. Let \widetilde{J} be a neural network such that \widetilde{J} is 1/2-averaged as in (2.1). Let \widetilde{A} be the maximally monotone operator equal to $(\widetilde{J}^{-1} - \mathrm{Id})$. Assume that the set S_{γ} of zeros of $\nabla f + \gamma^{-1}\widetilde{A}$ is nonempty. Then,



Fig. 2: Unfolded FB algorithm over N iterations.

the sequence $(x_n)_{n\in\mathbb{N}}$ generated by iteration (3.2) converges (weakly) to $\hat{x} \in S_{\gamma}$, i.e., \hat{x} satisfies

(3.3)
$$0 \in \nabla f(\widehat{x}) + \gamma^{-1} \widetilde{A}(\widehat{x}).$$

Remark 3.2.

- (i) The technical assumption S_γ ≠ Ø can be waived if ∇f + γ⁻¹à is strongly monotone [7, Corollary 23.37]). Then there exists a unique solution to (3.3). This is achieved if f is strongly convex or if J̃ = (Id+Q)/2 where Q = Q̃/(1+δ) with δ ∈]0,+∞[and Q̃: H → H nonexpansive. In the latter case, à = B + δ(2 + δ)⁻¹ Id where B: H ⇒ H is maximally monotone. An example of such a strongly monotone operator à is encountered in elastic net regularization.
- (ii) A classical result by Rockafellar states that A is the subdifferential of some convex lower-semincontinuous function if and only if A is maximally cyclically monotone [18, 50]. As we only enforce the maximal monotonicity, (3.3) does not necessarily correspond to a minimization problem in general.

If a finite number N of iterations of Algorithm (3.1) are performed, unfolding the FB algorithm results in the NN architecture given in Figure 2. If $\gamma < 2/\mu$, the gradient operator (Id $-\gamma\nabla f$) is a $\gamma\mu/2$ -averaged operator. It can thus be interpreted as an activation operator [23]. This activation operator is however non standard both because of its form and its dependence on the observed data z. A special case arises when f corresponds to a least squares data fit term, i.e.,

(3.4)
$$(\forall x \in \mathcal{H}) \quad f(x) = \frac{1}{2} \|Hx - z\|^2,$$

where z belongs to some real Hilbert space \mathcal{G} and H is a bounded operator from \mathcal{H} to \mathcal{G} modelling some underlying linear observation process (e.g. a degradation operator in image recovery). Then, $\nabla f \colon x \mapsto H^*(Hx - z)$ where H^* denotes the adjoint of H and $\mu = ||H||^2$. Hence, $\mathrm{Id} - \gamma \nabla f$ is an affine operator involving a selfadjoint weight operator Id $-\gamma H^*H$ and a bias γH^*z . The unfolded network has thus a structure similar to a residual network where groups of layers are identically repeated and the bias introduced in the gradient operator depends on z. A parallel could also be drawn with a recurrent neural network driven with a stationary input, which would here correspond to z. It is worth pointing out that, under the assumptions of Proposition 3.1, the unfolded network in Figure 2 is robust to adversarial input perturbations, since it is globally nonexpansive. Note finally that, in the case when f is given by (3.4), allowing the parameter γ and the operator J to be dependent on $n \in \{1, \ldots, N\}$ in Figure 2 would yield an extension of ISTA-net [64]. However, as shown in [15], convergence of such a scheme requires specific assumptions on the target signal model. Other works have also proposed NN architectures inspired from primal dual algorithms [2, 6, 33].

3.2. Training. A standard way of training a NN operating on $\mathcal{H} = \mathbb{R}^K$ for PnP algorithms is to train a denoiser for data corrupted with Gaussian noise [68]. Let $\overline{\boldsymbol{x}} = (\overline{x}_\ell)_{1 \leq \ell \leq L}$ be training set of L images of \mathcal{H} and let

(3.5)
$$(\forall \ell \in \{1, \dots, L\}) \qquad y_{\ell} = \overline{x}_{\ell} + \sigma_{\ell} w_{\ell}$$

be a noisy observation of \overline{x}_{ℓ} , where $\sigma_{\ell} \in [0, +\infty[$. In practice, either $\sigma_{\ell} \equiv \sigma > 0$ is chosen to be constant during training [67], or σ_{ℓ} is chosen to be a realization of a random variable with uniform distribution in $[0, \sigma]$, for $\sigma \in [0, +\infty[$ $(w_{\ell})_{1 \leq \ell \leq L}$ are assumed to be realizations of standard normal i.i.d. random variables. [69].

The NN \widetilde{J} described in the previous section will be optimally chosen within a family $\{\widetilde{J}_{\theta} \mid \theta \in \mathbb{R}^{P}\}$ of NNs. For example, the parameter vector θ will account for the convolutional kernels and biases of a given network architecture. An optimal value $\widehat{\theta}$ of the parameter vector is thus a solution to the following problem:

(3.6) minimize
$$\sum_{\ell=1}^{L} \|\widetilde{J}_{\theta}(y_{\ell}) - \overline{x}_{\ell}\|^2$$
 s.t. $Q_{\theta} = 2\widetilde{J}_{\theta} - \mathrm{Id}$ is nonexpansive.

(The squared ℓ_2 norm in (3.6) can be replaced by another cost function, e.g., an ℓ_1 norm [65].) The main difficulty with respect to a standard training procedure is the nonexpansiveness constraint stemming from Proposition 2.1 which is crucial to ensure the convergence of the overall PnP algorithm. In this context, the tight sufficient conditions described in Proposition 2.5 for building the associated nonexpansive operator Q_{θ} are however difficult to enforce. For example, the maximum value of the left-hand side in inequality (2.5) is NP-hard to compute [58] and estimating an accurate estimate of the Lipschitz constant of a NN requires some additional assumptions [48] or some techniques which do not scale well to high-dimensional data [28]. In turn, by assuming that, for every $\theta \in \mathbb{R}^P Q_{\theta}$ is differentiable, we leverage on the fact that Q_{θ} is nonexpansive if and only if its Jacobian ∇Q_{θ} satisfies

(3.7)
$$(\forall x \in \mathcal{H}) \qquad \|\nabla Q_{\theta}(x)\| \leq 1.$$

In practice, one cannot enforce the constraint in (3.7) for all $x \in \mathcal{H}$. We therefore propose to impose this constraint on every segment $[\overline{x}_{\ell}, \widetilde{J}_{\theta}(y_{\ell})]$ with $\ell \in \{1, \ldots, L\}$, or more precisely at points

(3.8)
$$\widetilde{x}_{\ell} = \varrho_{\ell} \overline{x}_{\ell} + (1 - \varrho_{\ell}) J_{\theta}(y_{\ell}),$$

where ϱ_{ℓ} is a realization of a random variable with uniform distribution on [0,1]. To cope with the resulting constraints, instead of using projection techniques which might be slow [56] and raise convergence issues when embedded in existing training algorithms [4], we propose to employ an exterior penalty approach. The final optimization problem thus reads

(3.9)
$$\min_{\theta} \sum_{\ell=1}^{L} \Phi_{\ell}(\theta),$$

where, for every $\ell \in \{1, \ldots, L\}$,

(3.10)
$$\Phi_{\ell}(\theta) = \|\widetilde{J}_{\theta}(y_{\ell}) - \overline{x}_{\ell}\|^{2} + \lambda \max\left\{\|\nabla Q_{\theta}(\tilde{x}_{\ell})\|^{2}, 1 - \varepsilon\right\},$$

 $\lambda \in [0, +\infty[$ is a penalization parameter, and $\varepsilon \in [0, 1[$ is a parameter allowing us to control the constraints. Standard results concerning penalization methods [40, Section 13.1], guarantee that, if $\hat{\theta}_{\lambda}$ is a solution to (3.9) for $\lambda \in [0, +\infty[$, then $(\forall \ell \in \{1, \ldots, L\})$ $\lim_{\lambda \to +\infty} \|\nabla Q_{\hat{\theta}_{\lambda}}(\tilde{x}_{\ell})\|^2 \leq 1 - \varepsilon$. Then, there exists $\overline{\lambda} \in [0, +\infty[$ such that, for every $\lambda \in [\overline{\lambda}, +\infty[$ and every $\ell \in \{1, \ldots, L\}$, $\|\nabla Q_{\hat{\theta}_{\lambda}}(\tilde{x}_{\ell})\| \leq 1$.

Remark 3.3.

- (i) Hereabove, we have made the assumptions that the network is differential. Automatic differentiation tools however are applicable to networks which contain nonsmooth linearities such as ReLU (see [11] for a theoretical justification for this fact).
- (ii) Note that this regularization strategy has the same flavour as the one in [31], where the loss is regularized with the Froebenius norm of the Jacobian. However, the latter is not enough to ensure convergence of the PnP method (3.2) which requires to constrain the spectral norm || ⋅ || of the Jacobian. Other works in the GAN literature have investigated similar regularizations [29, 49, 60].

To solve (3.9) numerically, we resort to the Adam optimizer [69] as described in Algorithm 3.1. This algorithm uses a fixed number of iterations $N \in \mathbb{N}^*$ and relies on approximations to the gradient of $\sum_{\ell} \Phi_{\ell}$ computed on randomly sampled batches of size D, selected from the training set of images $(\bar{x}_{\ell})_{1 \leq \ell \leq L}$. More precisely, at each iteration $t \in \{1, \ldots, N\}$, we build the approximated gradient $\frac{1}{D} \sum_{d=1}^{D} g_d$ (see lines 3-9), followed by an Adam update (line 10) consisting in a gradient step on θ_d with adaptive moment [34]. Then the approximated gradient is computed as follows. For every $d \in \{1, \ldots, D\}$, we select randomly an image from the training set (line 4), we draw at random a realization of a normal i.i.d. noise that we use to build a noisy observation y_d (line 5-6). We then build \tilde{x}_d as in (3.8) (lines 5-7) and compute the gradient g_d of the loss Φ_d w.r.t. to the parameter vector at its current estimate θ_n (line 8). Note that any other gradient-based algorithm, such as SGD or RMSprop [47] could be used to solve (3.9).

Algorithm 3.1 Adam algorithm to solve (3.9)

1: Let $D \in \mathbb{N}^*$ be the batch size, and $N \in \mathbb{N}^*$ be the number of training iterations. 2: for n = 1, ..., N do for d = 1, ..., D do 3: 4: Select randomly $\ell \in \{1, \ldots, L\};$ Draw at random $w_d \sim \mathcal{N}(0, 1)$ and $\varrho_d \sim \mathcal{U}([0, 1])$; 5: 6: $y_d = \overline{x}_\ell + \sigma w_d;$ $\widetilde{x}_d = \varrho_d \overline{x}_\ell + (1 - \varrho_d) \widetilde{J}_{\theta_n}(y_d);$ 7: $g_d = \nabla_\theta \Phi_d(\theta_n);$ 8: end for 9: $\theta_{n+1} = \operatorname{Adam}(\frac{1}{D}\sum_{d=1}^{D} g_d, \theta_n);$ 10: 11: end for 12: return \widetilde{J}_{θ_N}

REMARK 3.4. To compute the spectral norm $\|\nabla Q_{\theta}(x)\|$ for a given image $x \in \mathcal{H}$, we use the power iterative method where the Jacobian is computed by backpropagation.

- 4. Simulations and results.
- 4.1. Experimental setting.



Fig. 3: Blur kernels used in our simulations. (a)-(h) are kernels 1-8 from [39] respectively while (i) is the kernel from the GaussianA setup and (j) from the Square setup in [8].

Inverse Problem. We focus on inverse deblurring imaging problems, where the objective is to find an estimate $\hat{x} \in \mathbb{R}^K$ of an original unknown image $\overline{x} \in \mathbb{R}^K$, from degraded measurements $z \in \mathbb{R}^K$ given by

$$(4.1) z = H\overline{x} + e,$$

where $H: \mathbb{R}^K \to \mathbb{R}^K$ is a blur operator and $e \in \mathbb{R}^K$ is a realization of an additive white Gaussian random noise with zero-mean and standard deviation $\nu \in [0, +\infty[$. In this context, a standard choice for the data-fidelity term is given by (3.4) In our simulations, H models a blurring operator implemented as a circular convolution with impulse response h. We will consider different kernels h taken from [39] and [8], see Figure 3 for an illustration. The considered kernels are normalized such that the Lipschitz constant μ of the gradient of f is equal to 1.

Datasets. Our training dataset consists of 50000 test images from the ImageNet dataset [27] that we randomly split in 98% for training and 2% for validation. In the case of grayscale images, we investigate the behaviour of our method either on the full BSD68 dataset [43] or on a subset of 10 images, which we refer to as the BSD10 set. For color images, we consider both the BSD500 test set [43] and the Flickr30 test set [61].¹ Eventually, when some fine-tuning is required, we employ the Set12 and Set18 datasets [67] for grayscale and color images, respectively.

Network architecture and pretraining. In existing PnP algorithms involving NNs (see e.g. [37, 66, 67, 69]), the NN architecture \tilde{J} often relies on residual skip connections. This is equivalent, in (2.3), to set $Q = \text{Id} + \tilde{T}_M \dots \tilde{T}_1$ where, for every $m \in \{1, \dots, M\}$, \tilde{T}_m is standard neural network layer (affine operator followed by activation operator). More specifically, the architecture we consider for \tilde{J} is such that M = 20. It is derived from DnCNN-B architecture [66] from which we have removed batch normalization layers and where we have replaced ReLUs with LeakyReLUs (see Figure 4).

We first pretrain the model \tilde{J} in order to perform a blind denoising task without any Jacobian regularization. For each training batch, we generate randomly sampled patches of size 50 × 50 from images that are randomly rescaled and flipped. More precisely, we consider Problem (3.9)-(3.10) with $\lambda = 0$, and $(\sigma_{\ell})_{1 \leq \ell \leq L}$ chosen to be realizations of i.i.d. random variable with uniform distribution in [0, 0.1] for each patch. We use the Adam optimizer [34] to pretrain the network with learning rate 10^{-4} , clipped gradient norms at 10^{-2} , and considering 150 epochs, each consisting of 490 iterations of the optimizer. The learning rate is divided by 10 after 100 epochs. This pretrained network will serve as a basis for our subsequent studies. The details regarding the training of our networks will be given on a case-by-case basis in the following sections.

¹We consider normalised images, where the coefficient values are in [0,1] (resp. $[0,1]^3$) for grayscale (resp. color) images.



Fig. 4: Proposed DnCNN architecture of \tilde{J} , with a total of 20 convolutional layers. It corresponds to a modified version of the DnCNN-B architecture [66]. The number of channels C is indicated above arrows (C = 1 for grayscale images and C = 3 for color ones).

All models are trained on 2 Nvidia Tesla 32 Gb V100 GPUs and experiments are performed in $PyTorch^2$.

Goal. We aim to study the PnP-FB algorithm (3.2) where \tilde{J} , chosen according to the architecture given in Figure 4, has been trained in order to solve (3.10). We will first study the impact of the choice of the different parameters appearing in the training loss (3.10) on the convergence of the PnP-FB algorithm and on the reconstruction quality. Then, we will compare the proposed method to state-of-theart iterative algorithms either based on purely variational or PnP methods.

We evaluate the reconstruction quality with Peak Signal to Noise Ratio (PSNR) and Structural Similarity Index Measure (SSIM) metrics [59]. The PSNR between an image $x \in \mathbb{R}^{K}$ and the ground truth $\overline{x} \in \mathbb{R}^{K}$ is defined as

(4.2)
$$\operatorname{PSNR}(x,\overline{x}) = 20 \log_{10} \left(\frac{\sqrt{K} \max_{1 \le \ell \le L} \overline{x}_{\ell}}{\|x - \overline{x}_{\ell}\|} \right),$$

where, in our case, we have $\max_{1 \le \ell \le L} \overline{x}_{\ell} = 1$. The SSIM is given by

(4.3)
$$\operatorname{SSIM}(x,\overline{x}) = \frac{(2\mu_x\mu_{\overline{x}} + \vartheta_1)(2\sigma_{x\overline{x}} + \vartheta_2)}{(\mu_x^2 + \mu_{\overline{x}}^2 + \vartheta_1)(\sigma_x^2 + \sigma_{\overline{x}}^2 + \vartheta_2)},$$

where (μ_x, σ_x) and $(\mu_{\overline{x}}, \sigma_{\overline{x}})$ are the mean and the variance of x and \overline{x} respectively, $\sigma_{x\overline{x}}$ is the cross-covariance between x and \overline{x} , and $(\vartheta_1, \vartheta_2) = (10^{-4}, 9 \times 10^{-4})$.

4.2. Choice of the parameters. In this section, we study the influence of the parameters $(\lambda, \sigma, \gamma)$ on the results of the PnP-FB algorithm 3.2 applied to the NN in Figure 4. We recall that λ is the parameter acting on the Jacobian regularization, σ is the noise level for which the denoiser is trained, and γ is the stepsize in the PnP-FB algorithm (3.2).

Simulation settings. We consider problem (4.1) with H associated with the kernels shown in Figure 3(a)-(h), and $\nu = 0.01$. In this section, we consider the grayscale images from the BSD68 dataset.

To investigate the convergence behaviour of the PnP-FB algorithm, we consider the quantity defined at iteration $n \in \mathbb{N} \setminus \{0\}$ as

(4.4)
$$c_n = \|x_n - x_{n-1}\| / \|x_0\|,$$

where $(x_n)_{n \in \mathbb{N}}$ is the sequence generated by the PnP-FB algorithm (3.2). Note that the quantity $(c_n)_{n \in \mathbb{N}}$ is known to be monotonically decreasing if the network \widetilde{J} is firmly nonexpansive [7].

²Code publicly available at https://github.com/basp-group/xxx (upon acceptance of the paper)

Influence of the Jacobian penalization. First we study the influence of λ on the convergence behaviour of the PnP-FB algorithm (3.2). In particular we consider $\lambda \in \{5 \times 10^{-7}, 10^{-6}, 2 \times 10^{-6}, 5 \times 10^{-6}, 10^{-5}, 2 \times 10^{-5}, 4 \times 10^{-5}, 1.6 \times 10^{-4}, 3.2 \times 10^{-4}, 6.4 \times 10^{-4}\}.$

After pretraining, we train our DnCNN by considering the loss given in (3.10), in which we set $\varepsilon = 5 \times 10^{-2}$ and $\sigma = 0.01$. The batches are built as in the pretraining setting. The network is trained for 100 epochs and the learning rate is divided by 10 at epoch 80. The training is performed with Algorithm 3.1 where D = 100 and $N = 4.9 \times 10^4$. For Adam's parameters, we set the learning rate to 10^{-4} and the remaining parameters to the default values provided in [34].

To verify that our training loss enables the firm nonexpansiveness of our NN J, we evaluate the norm of the Jacobian $\|\nabla Q(y_\ell)\|$ on a set of noisy images $(y_\ell)_{1 \leq \ell \leq 68}$, obtained from the BSD68 test set considering the denoising problem (3.5). The maximum of these values is given in Table 1 for the different considered values of λ . We observe that the norm of the Jacobian decreases as λ increases and is smaller than 1 for $\lambda \geq 10^{-5}$.

We now investigate the convergence behaviour of the PnP-FB algorithm, depending on λ , considering BSD10 (a subset of BSD68). In our simulations, we set $\gamma = 1/\mu = 1$. In Figure 5 we show the values $(c_n)_{1 \leq n \leq 1000}$ for 1000 iterations, considering kernel (a) from Figure 3 for the different values of λ . The case $\lambda = 0$ corresponds to training a DnCNN without the Jacobian regularization. We observe that the stability of the PnP-FB algorithm greatly improves as λ increases: for $\lambda \geq 10^{-5}$, all curves are monotonically decreasing. These observations are in line with the metrics from Table 1 showing that $\|\nabla Q(y_\ell)\| \leq 1$ for $\lambda \geq 10^{-5}$.

These results confirm that by choosing an appropriate value of λ , one can ensure Q to be 1-Lipschitz, i.e. \tilde{J} to be firmly nonexpansive, and consequently we secure the convergence of the PnP-FB algorithm (3.2).

λ	0	5×10^{-7}	1×10^{-6}	2×10^{-6}	5×10^{-6}	1×10^{-5}	4×10^{-5}	1.6×10^{-4}	3.2×10^{-4}
$\ \nabla Q\ ^2$	31.36	1.65	1.349	1.156	1.028	0.9799	0.9449	0.9440	0.9401

Table 1: Numerical evaluation of the firm nonexpansiveness \tilde{J} on a denoising problem on the BSD68 test set for different values of λ .

Influence of the stepsize and training noise level. Second, we investigate the influence (σ, γ) on the reconstruction quality of the images restored with the PnP-FB algorithm. We train the NN \tilde{J} given in Figure 4 for $\sigma \in \{0.005, 0.006, 0.007, 0.008, 0.009, 0.01\}$. As per the procedure followed in the study of the parameter λ , after pretraining, we train \tilde{J} by considering the loss given in (3.10), in which we set $\varepsilon = 5 \times 10^{-2}$. The value of λ was fine-tuned around 10^{-5} . The batches are built as in the pretraining setting. The network is trained for 100 epochs and the learning rate is divided by 10 at epoch 80. The training is performed with Algorithm 3.1 where D = 100 and $N = 4.9 \times 10^4$. For Adam's parameters, we set the learning rate to 10^{-4} and the remaining parameters to the default values provided in [34]. We subsequently plug the trained DnCNN \tilde{J} in the PnP-FB algorithm (3.2), considering different values for $\gamma \in [0, 2[$. In these simulations, we focus on the case when the blur kernel in Problem (4.1) corresponds to the one shown in Figure 3(a).

Before discussing the simulation results, we present a heuristic argument suggesting that (i) σ should scale linearly with γ , and (ii) the appropriate scaling coefficient



Fig. 5: Influence of $\lambda \in \{0, 5 \times 10^{-7}, 10^{-6}, 2 \times 10^{-6}, 5 \times 10^{-6}, 10^{-5}, 4 \times 10^{-5}, 1.6 \times 10^{-4}\}$ on the stability of the PnP-FB algorithm for the deblurring problem with kernel in Figure 3(a). (a)-(h): On each graph, evolution of the quantity c_n defined in (4.4) for each image of the BSD10 test set, for a value of λ , along the iterations of the PnP-FB algorithm (3.2).

is given as $2\nu \|h\|$. Given the choice of the data-fidelity term (3.4), the PnP-FB algorithm (3.2) reads

(4.5)
$$(\forall n \in \mathbb{N}) \quad x_{n+1} = \widetilde{J}(x_n - \gamma H^*(Hx_n - H\overline{x} - e)).$$

We know that, under suitable conditions, the sequence $(x_n)_{n\in\mathbb{N}}$ generated by (4.5) converges to a fixed point \hat{x} , solution to the variational inclusion problem (3.3). We assume that \hat{x} lies close to \overline{x} up to a random residual $e' = H(\hat{x} - \overline{x})$, whose components are uncorrelated and with equal standard deviation, typically expected to be bounded from above by the standard deviation ν of the components of the original noise e. Around convergence, (4.5) therefore reads as

(4.6)
$$\widehat{x} = \widetilde{J} \left(\widehat{x} - \gamma H^* \left(e' - e \right) \right),$$

suggesting that, J is acting as a denoiser of \hat{x} for an effective noise $-\gamma H^*(e'-e)$. If the components of e'-e are uncorrelated, the standard deviation of this noise is bounded by $\gamma \nu_{\text{eff}}$, with $\nu_{\text{eff}} = 2\nu ||h||$, a value reached when e' = -e. This linear function of γ with scaling coefficient ν_{eff} thus provides a strong heuristic for the choice of the standard deviation σ of the training noise. For the considered kernel (shown in Figure 3(a)), we have $\nu_{\text{eff}} = 0.0045$, so the interval $\sigma \in [0.005, 0.01]$ also reads $\sigma \in [1.1 \nu_{\text{eff}}, 2.2 \nu_{\text{eff}}]$.

In Figure 6 we provide the average PSNR (left) and SSIM (right) values associated with the solutions to the deblurring problem for the considered simulations as a function of $\sigma/\gamma\nu_{\text{eff}}$. For each sub-figure, the different curves correspond to different values of γ . We observe that, whichever the values of γ , the reconstruction quality is sharply peaked around values of $\sigma/\gamma\nu_{\text{eff}}$ consistently around 1, thus supporting our heuristic argument. We also observe that the peak value increases with γ . We recall that, according to the conditions imposed on γ in Proposition 3.1 to guarantee theoretically the convergence of the sequence generated by PnP-FB algorithm, one has $\gamma < 2$. The values $\gamma = 1.99$ and $\sigma/\gamma\nu_{\text{eff}} = 1$ (resp. $\gamma = 1.99$ and $\sigma/\gamma\nu_{\text{eff}} = 0.9$) gives the best results for the PSNR (resp. SSIM).



Fig. 6: Influence of $\gamma \in]0, 1.99]$ and $\sigma \in [0.005, 0.01]$ on the reconstruction quality for the deblurring problem with kernel from Figure 3(a) on the BSD10 test set. For this experiment $\nu_{\text{eff}} = 0.0045$. Left: average PSNR, right: average SSIM.

In Figure 7 we provide visual results for an image from the BSD10 test set, to the deblurring problem for different values of γ and σ . The original unknown image \overline{x} and the observed blurred noisy image are displayed in Figure 7(a) and (g), respectively. On the top row, we set $\sigma = 2\nu_{\text{eff}}$, while the value of γ varies from 1 to 1.99. We observe that the reconstruction quality improves when γ increases, bringing the ratio $\sigma/\gamma \nu_{\rm eff}$ closer to unity. Precisely, in addition to the PSNR and SSIM values increasing with γ , we can see that the reconstructed image progressively loses its oversmoothed aspect, showing more details. The best reconstruction for this row is given in Figure 7(f), for $\gamma = 1.99$. On the bottom row, we set $\gamma = 1$ and vary σ from $1.3 \nu_{\text{eff}}$ to $2.2 \nu_{\text{eff}}$. We see that sharper details appear in the reconstructed image when σ decreases, again bringing the ratio $\sigma/\gamma\nu_{\rm eff}$ closer to unity. The best reconstructions for this row are given in Figure 7(h) and (i), corresponding to the cases $\sigma = 1.3 \nu_{\rm eff}$ and $\sigma = 1.6 \nu_{\text{eff}}$, respectively. Overall, as we have already noticed, the best reconstruction is obtained for $\gamma = 1.99$ and $\sigma / \gamma \nu_{\text{eff}} = 1$, for which the associated image is displayed in Figure 7(f). These results further support both our analysis of Figure 6 and our heuristic argument for a linear scaling of σ with γ , with scaling coefficient closely driven by the value ν_{eff} .

4.3. Comparison with other PnP methods. In this section we investigate the behaviour of the PnP-FB algorithm (3.2) with \tilde{J} corresponding either to the proposed DnCNN provided in Figure 4, or to other denoisers. In this section, we aim to solve problem (4.1), considering either grayscale or color images.

Grayscale images. We consider the deblurring problem (4.1) with H associated with the kernels from Figure 3(a)-(h), $\nu = 0.01$, evaluated on the BSD10 test set.

We choose the parameters of our method to be the ones leading to the best PSNR values in Figure 6, i.e. $\sigma = 0.009$ and $\gamma = 1.99$ corresponding to $\sigma/\gamma\nu_{\text{eff}} = 1$ for the kernel (a) of Figure 3, and we set $\lambda = 4 \times 10^{-6}$.

We compare our method with other PnP-FB algorithms, where the denoiser corresponds either to RealSN [52], BM3D [41], DnCNN [66], or standard proximity operators [24, 46]. In our simulations, we consider the proximal operators of the two following functions: (i) the ℓ_1 -norm composed with a sparsifying operator consisting in the concatenation of the first eight Daubechies (db) wavelet bases [13, 42], and (ii) the total variation (TV) norm [51]. In both cases, the regularization parameters are fine-tuned on the Set12 dataset [66] to maximize the reconstruction quality. Note



Fig. 7: Reconstructions of an image from the BSD10 test set obtained with the PnP-FB algorithm (3.2) for the deblurring problem with kernel from Figure 3(a) for which $\nu_{\text{eff}} = 0.0045$. Top row: results for $\gamma \in [1, 1.99]$ in algorithm (3.2) and $\sigma/\nu_{\text{eff}} = 2$ (i.e. $\sigma = 0.009$). Bottom row: results for $\sigma/\nu_{\text{eff}} \in [1.3, 2.2]$ during training in (3.10) and $\gamma = 1$ in algorithm (3.2).



Fig. 8: Convergence profile of the PnP-FB algorithm (3.2) for different denoisers plugged in as \tilde{J} , namely BM3D (a), RealSN (b) and the proposed firmly nonexpansive DnCNN (c). Results are shown for the deblurring problem (4.1) with kernel from Figure 3(a). Each graph shows the evolution of c_n defined in (4.4) for each image of the BSD10 test set.

that the training process for RealSN has been adapted for the problem of interest.

We first check the convergence of the PnP-FB algorithm considering the abovementioned different denoisers. We study the quantity $(c_n)_{n\in\mathbb{N}}$ defined in (4.4), considering the inverse problem (4.1) with kernel in Figure 3(a). Figure 8 shows the c_n values with respect to the iterations $n \in \{1, \ldots, 1000\}$ of the PnP-FB algorithm for various denoisers \tilde{J} : BM3D (Figure 8(a)), RealSN (Figure 8(b)), and the proposed firmly nonexpansive DnCNN (Figure 8(c)). On the one hand, we notice that the PnP-FB algorithm with BM3D or RealSN does not converge since $(c_n)_{n\in\mathbb{N}}$ does not tend to zero, which confirms that neither BM3D nor RealSN are firmly nonexpansive.

	LEARNING MAXIMALLY MONOTONE OPERATORS FOR IMAGE RECOVERY $=23$
--	--

denoiser	kernel (see Figure 3)								convergence
denoiser	(a)	(b)	(c)	(d)	(e)	(f)	(g)	(h)	convergence
Observation	23.36	22.93	23.43	19.49	23.84	19.85	20.75	20.67	
RealSN [52]	26.24	26.25	26.34	25.89	25.08	25.84	24.81	23.92	\checkmark
$\operatorname{prox}_{\mu_{\ell_1}} \ \Psi^{\dagger} \cdot \ _1$	29.44	29.20	29.31	28.87	30.90	30.81	29.40	29.06	\checkmark
$\operatorname{prox}_{\mu_{\mathrm{TV}} \ \cdot \ _{\mathrm{TV}}}$	29.70	29.35	29.43	29.15	30.67	30.62	29.61	29.23	\checkmark
DnCNN [66]	29.82	29.24	29.26	28.88	30.84	30.95	29.54	29.17	×
BM3D [41]	30.05	29.53	29.93	29.10	31.08	30.78	29.56	29.41	×
Proposed	30.91	30.47	30.46	30.24	31.72	31.75	30.60	30.23	✓

Table 2: Average PSNR values obtained by different denoisers plugged in the PnP-FB algorithm (3.2), to solve the deblurring problem (4.1) with kernels of Figure 3(a)-(h) considering the BSD10 test set. The last row provides the average SSIM values for the observed blurred image y in each experimental setting. Each algorithm is stopped after a fixed number of iterations equal to 1000. The best PSNR values are indicated in bold.

On the other hand, as expected, PnP-FB with our network, which has been trained to be firmly nonexpansive, shows a convergent behaviour with monotonic decrease of c_n .

In Table 2 we provide a quantitative analysis of the restoration quality obtained on the BSD10 dataset with the different denoisers. Although DnCNN and BM3D do not benefit from any convergence guarantees, we report the SNR values obtained after 1000 iterations. For all the eight considered kernels, the best PSNR values are delivered by the proposed firmly nonexpansive DnCNN.

In Figure 9 we show visual results and associated PSNR and SSIM values obtained with the different methods on the deblurring problem (4.1) with kernel from Figure 3(a). We notice that despite good PSNR and SSIM values, the proximal methods yield reconstructions with strong visual artifacts (wavelet artifacts in Figure 9(c) and cartoon effects in Figure 9(d)). PnP-FB with BM3D provides a smoother image with more appealing visual results, yet some grid-like artifacts appear in some places (see e.g. red boxed zoom in Figure 9(e)). RealSN introduces ripple and dotted artifacts, while DnCNN introduces geometrical artifacts, neither of those corresponding to features in the target image. For this image, we can observe that our method provides better visual results as well as higher PSNR and SSIM values than other methods.

The results presented in this section show that the introduction of the Jacobian regularizer in the training loss (3.10) not only allows to build convergent PnP-FB methods, but also improves the reconstruction quality over both FB algorithms involving standard proximity operators, and existing PnP-FB approaches.

Color images. We now apply our strategy to a color image deblurring problem of the form (4.1), where the noise level and blurring operator are chosen to reproduce the experimental settings of [8], focusing on the four following experiments: First, the Motion A (M. A) setup with blur kernel (h) from Figure 3 and $\nu = 0.01$; second, the Motion B (M. B) setup with blur kernel (c) from Figure 3 and $\nu = 0.01$; third, the Gaussian A (G. A) setup with kernel (i) from Figure 3 and $\nu = 0.008$; finally, the Square (S.) setup with kernel (j) from Figure 3 $\nu = 0.01$. The experiments in this section are run on the Flickr30 dataset and on the test set from BSD500³. We compare our method on these problems with the variational method VAR from [8], and three

³As in [8], we consider 256×256 centered-crop versions of the datasets.



Fig. 9: Reconstructions of an image from the BSD10 test set obtained with the PnP-FB algorithm (3.2), considering different denoisers as \tilde{J} , for the deblurring problem with kernel from Figure 3(a) and $\nu = 0.01$. Associated (PSNR, SSIM) values are indicated below each image, best values are highlighted in bold. Each algorithm is stopped after a fixed number of iteration equal to 1000.

PnP algorithms, namely PDHG [44], and the PnP-FB algorithm combined with the BM3D or DnCNN denoisers. It is worth mentioning that, among the above mentioned methods, only the proposed approach and VAR have convergence guaranties. The results for PDHG and VAR are borrowed from [8].

For the proposed method, we choose $\gamma = 1.99$ in the PnP-FB algorithm (3.2), and we keep the same DnCNN architecture for \tilde{J} given in Figure 4, only changing the number of input/output channels to C = 3. We first pretrain our network as described in subsection 4.1. We then keep on training it considering the loss given in (3.10), in which we set $\varepsilon = 5 \times 10^{-2}$, $\lambda = 10^{-5}$, and $\sigma = 0.007$.

The average PSNR and SSIM values obtained with the different considered reconstruction methods, and for the different experimental settings, are reported in Figure 10. This figure shows that our method significantly improves reconstruction quality over the other considered PnP methods.

Visual comparisons are provided in Figure 11 for the different approaches. These results show that our method also yields better visual results. The reconstructed images contain finer details and do not show the oversmoothed appearance of PnP-FB with DnCNN or slightly blurred aspect of PnP-FB with BM3D. Note, in particular, that thanks to its convergence, the proposed method shows an homogeneous performance over all images, unlike PnP-FB with DnCNN that may show some divergence effects (see the boat picture for Motion A, row (f)). One can observe that the improvement obtained with our approach are more noticeable on settings M. A and M. B than on G. A and S.



Fig. 10: Average PSNR and SSIM values obtained on the Flickr30 (top) and BSD500 (bottom) test sets using the experimental setups of [8]: G. A, M. A, M. B, and S., for different methods.

5. Conclusion. In this paper, we investigated the interplay between PnP algorithms and monotone operator theory, in order to propose a sound mathematical framework yielding both convergence guarantees and a good reconstruction quality in the context of computational imaging.

First, we established a universal approximation theorem for a wide range of MMOs, in particular the new class of stationary MMOs we have introduced. This theorem constitutes the theoretical backbone of our work by proving that the resolvents of these MMOs can be approximated by building nonexpansive NNs. Leveraging this result, we proposed to learn MMOs in a supervised manner for PnP algorithms. A main advantage of this approach is that it allows us to characterize their limit as a solution to a variational inclusion problem.

Second, we proposed a novel training loss to learn the resolvent of an MMO for high dimensional data, by imposing mild conditions on the underlying NN architecture. This loss uses information of the Jacobian of the NN, and can be optimized efficiently using existing training strategies. Finally, we demonstrated that the resulting PnP algorithms grounded on the FB scheme have good convergence properties. We showcased our method on an image deblurring problem and showed that the proposed PnP-FB algorithm outperforms both standard variational methods and state-of-the-art PnP algorithms.

Note that the ability of approximating resolvents as we did would be applicable to a much wider class of iterative algorithms than the forward-backward splitting [22]. In addition, we could consider a wider scope of applications than the restoration problems addressed in this work.

REFERENCES

- A. ABDULAZIZ, A. DABBECH, AND Y. WIAUX, Wideband super-resolution imaging in radio interferometry via low rankness and joint average sparsity models (hypersara), Monthly Notices of the Royal Astronomical Society, 489 (2019), pp. 1230–1248.
- J. ADLER AND O. ÖKTEM, Learned primal-dual reconstruction, IEEE Trans. on Medical Imaging, 37 (2018), pp. 1322–1332.

- [3] R. AHMAD, C. A. BOUMAN, G. T. BUZZARD, S. CHAN, S. LIU, E. T. REEHORST, AND P. SCHNITER, Plug-and-play methods for magnetic resonance imaging: Using denoisers for image recovery, IEEE Signal Process. Mag., 37 (2020), pp. 105–116.
- [4] A. ALACAOGLU, Y. MALITSKY, AND V. CEVHER, Convergence of adaptive algorithms for weakly convex constrained optimization, arXiv preprint arXiv:2006.06650, (2020).
- [5] C. ANIL, J. LUCAS, AND R. GROSSE, Sorting out lipschitz function approximation, in International Conference on Machine Learning, PMLR, 2019, pp. 291–301.
- [6] S. BANERT, A. RINGH, J. ADLER, J. KARLSSON, AND O. OKTEM, Data-driven nonsmooth optimization, SIAM J. on Optimization, 30 (2020), pp. 102–131.
- [7] H. H. BAUSCHKE AND P. L. COMBETTES, Convex analysis and monotone operator theory in Hilbert spaces, Springer, 2017.
- [8] C. BERTOCCHI, E. CHOUZENOUX, M.-C. CORBINEAU, J.-C. PESQUET, AND M. PRATO, Deep unfolding of a proximal interior point method for image restoration, Inverse Problems, 36 (2020), p. 034005.
- [9] A. BIBI, B. GHANEM, V. KOLTUN, AND R. RANFTL, Deep layers as stochastic solvers, in International Conference on Learning Representations, 2019.
- [10] J. BIRDI, A. REPETTI, AND Y. WIAUX, Sparse interferometric stokes imaging under the polarization constraint (polarized sara), Monthly Notices of the Royal Astronomical Society, 478 (2018), pp. 4442–4463.
- [11] J. BOLTE AND E. PAUWELS, Conservative set valued fields, automatic differentiation, stochastic gradient methods and deep learning, Mathematical Programming, (2020), pp. 1–33.
- [12] K. BREDIES, K. KUNISCH, AND T. POCK, Total generalized variation, SIAM J. on Imaging Sciences, 3 (2010), pp. 492–526.
- [13] R. E. CARRILLO, J. D. MCEWEN, D. VAN DE VILLE, J.-P. THIRAN, AND Y. WIAUX, Sparsity averaging for compressive imaging, IEEE Signal Process. Lett., 20 (2013), pp. 591–594.
- [14] G. H.-G. CHEN AND R. T. ROCKAFELLAR, Convergence rates in forward-backward splitting, SIAM J. Optim, 7 (1997), pp. 421–444.
- [15] X. CHEN, J. LIU, Z. WANG, AND W. YIN, Theoretical linear convergence of unfolded ista and its practical weights and thresholds, in Advances in Neural Information Processing Systems, 2018, pp. 9061–9071.
- [16] M. CISSE, P. BOJANOWSKI, E. GRAVE, Y. DAUPHIN, AND N. USUNIER, Parseval networks: Improving robustness to adversarial examples, in Proceedings of the 34th International Conference on Machine Learning, JMLR, 2017, pp. 854–863.
- [17] R. COHEN, M. ELAD, AND P. MILANFAR, Regularization by denoising via fixed-point projection (red-pro), arXiv preprint arXiv:2008.00226, (2020).
- [18] P. L. COMBETTES, Monotone operator theory in convex optimization, Math. Program., 170 (2018), pp. 177–206.
- [19] P. L. COMBETTES AND J.-C. PESQUET, Image restoration subject to a total variation constraint, IEEE Trans. Image Process., 13 (2004), pp. 1213–1222.
- [20] P. L. COMBETTES AND J.-C. PESQUET, Proximal splitting methods in signal processing, in Fixed-point algorithms for inverse problems in science and engineering, Springer, 2011, pp. 185–212.
- [21] P. L. COMBETTES AND J.-C. PESQUET, Deep neural network structures solving variational inequalities, Set-Valued and Variational Analysis, (2020), pp. 1–28.
- [22] P. L. COMBETTES AND J.-C. PESQUET, Fixed point strategies in data science, arXiv preprint arXiv:2008.02260, (2020).
- [23] P. L. COMBETTES AND J.-C. PESQUET, Lipschitz certificates for layered network structures driven by averaged activation operators, SIAM J. on Mathematics of Data Science, 2 (2020), pp. 529–557.
- [24] P. L. COMBETTES AND V. R. WAJS, Signal recovery by proximal forward-backward splitting;, Multiscale Model. Simul, 4 (2005), pp. 1168–1200.
- [25] L. CONDAT, Semi-local total variation for regularization of inverse problems, in Proceedings of the IEEE European Signal Processing Conference, 2014, pp. 1806–1810.
- [26] I. DAUBECHIES, M. DEFRIEE, AND C. DEMOL, An iterative thresholding algorithm for linear inverse problems with a sparsity constraint, Comm. Pure Appl. Math, 57 (2004), pp. 1413– 1457.
- [27] J. DENG, W. DONG, R. SOCHER, L.-J. LI, K. LI, AND L. FEI-FEI, *Imagenet: A large-scale hierarchical image database*, in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2009, pp. 248–255.
- [28] M. FAZLYAB, A. ROBEY, H. HASSANI, M. MORARI, AND G. PAPPAS, Efficient and accurate estimation of Lipschitz constants for deep neural networks, in Advances in Neural Information Processing Systems, 2019, pp. 11427–11438.

- [29] I. GULRAJANI, F. AHMED, M. ARJOVSKY, V. DUMOULIN, AND A. C. COURVILLE, *Improved training of wasserstein gans*, in Advances in Neural Information Processing Systems, 2017, pp. 5767–5777.
- [30] J. HERTRICH, S. NEUMAYER, AND G. STEIDL, Convolutional proximal neural networks and plugand-play algorithms, arXiv preprint arXiv:2011.02281, (2020).
- [31] J. HOFFMAN, D. A. ROBERTS, AND S. YAIDA, Robust learning with jacobian regularization, arXiv preprint arXiv:1908.02729, (2019).
- [32] K. HORNIK, M. STINCHCOMBE, AND H. WHITE, Multilayer feedforward networks are universal approximators, p, (1989), pp. 359–366.
- [33] M. JIU AND N. PUSTELNIK, A deep primal-dual proximal network for image restoration, arXiv preprint arXiv:2007.00959, (2020).
- [34] D. P. KINGMA AND J. BA, Adam: A method for stochastic optimization, in International Conference on Learning Representations, 2015.
- [35] N. KOMODAKIS AND J.-C. PESQUET, Playing with duality: An overview of recent primal-dual approaches for solving large-scale optimization problems, IEEE Signal Process. Mag., 32 (2015), pp. 31–54.
- [36] F. LATORRE, P. ROLLAND, AND V. CEVHER, Lipschitz constant estimation of neural networks via sparse polynomial optimization, in International Conference on Learning Representations, 2020.
- [37] C. LEDIG, L. THEIS, F. HUSZÁR, J. CABALLERO, A. CUNNINGHAM, A. ACOSTA, A. AITKEN, A. TEJANI, J. TOTZ, Z. WANG, ET AL., *Photo-realistic single image super-resolution using* a generative adversarial network, in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2017, pp. 4681–4690.
- [38] M. LESHNO, V. Y. LIN, A. PINKUS, AND S. SCHOCKEN, Multilayer feedforward networks with a nonpolynomial activation function can approximate any function, p, (1993), pp. 861–867.
- [39] A. LEVIN, Y. WEISS, F. DURAND, AND W. FREEMAN, Understanding and evaluating blind deconvolution algorithms, in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2009.
- [40] D. G. LUENBERGER, Y. YE, ET AL., Linear and nonlinear programming, Fourth Edition, vol. 228, Springer, 2016.
- [41] Y. MÄKINEN, L. AZZARI, AND A. FOI, Collaborative filtering of correlated noise: Exact transform-domain variance for improved shrinkage and patch matching, IEEE Trans. Image Process., 29 (2020), pp. 8339–8354.
- [42] S. MALLAT, A wavelet tour of signal processing: the sparse way, Academic press, 2008.
- [43] D. MARTIN, C. FOWLKES, D. TAL, AND J. MALIK, A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics, in Proceedings of the IEEE International Conference on Computer Vision, vol. 2, 2001, pp. 416–423.
- [44] T. MEINHARDT, M. MOLLER, C. HAZIRBAS, AND D. CREMERS, Learning proximal operators: Using denoising networks for regularizing inverse imaging problems, in Proceedings of the IEEE International Conference on Computer Vision, 2017, pp. 1781–1790.
- [45] T. MIYATO, T. KATAOKA, M. KOYAMA, AND Y. YOSHIDA, Spectral normalization for generative adversarial networks, in International Conference on Learning Representations, 2018.
- [46] J.-J. MOREAU, Proximité et dualité dans un espace hilbertien, Bulletin de la Société mathématique de France, 93 (1965), pp. 273–299.
- [47] M. C. MUKKAMALA AND M. HEIN, Variants of rmsprop and adagrad with logarithmic regret bounds, in International Conference on Machine Learning, JMLR, 2017, pp. 2545–2553.
- [48] A. NEACSU, J.-C. PESQUET, AND C. BURILEANU, Accuracy-robustness trade-off for positively weighted neural networks, in Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing, 2020, pp. 8389–8393.
- [49] H. PETZKA, A. FISCHER, AND D. LUKOVNIKOV, On the regularization of wasserstein gans, in International Conference on Learning Representations, 2018.
- [50] R. ROCKAFELLAR, Characterization of the subdifferentials of convex functions, Pacific J. Math., 17 (1966), pp. 497–510.
- [51] L. I. RUDIN, S. OSHER, AND E. FATEMI, Nonlinear total variation based noise removal algorithms, Physica D: nonlinear phenomena, 60 (1992), pp. 259–268.
- [52] E. RYU, J. LIU, S. WANG, X. CHEN, Z. WANG, AND W. YIN, Plug-and-play methods provably converge with properly trained denoisers, in International Conference on Machine Learning, PMLR, 2019, pp. 5546–5557.
- [53] K. SCAMAN AND A. VIRMAUX, Lipschitz regularity of deep neural networks: Analysis and efficient estimation, Adv. Neural Inform. Process. Syst, 31 (2018), pp. 3839–3848.
- [54] H. SEDGHI, V. GUPTA, AND P. M. LONG, The singular values of convolutional layers, arXiv

preprint arXiv:1805.10408, (2018).

- [55] C. SZEGEDY, W. ZAREMBA, I. SUTSKEVER, J. BRUNA, D. ERHAN, I. J. GOODFELLOW, AND R. FERGUS, *Intriguing properties of neural networks*, in International Conference on Learning Representations, 2014.
- [56] M. TERRIS, A. REPETTI, J.-C. PESQUET, AND Y. WIAUX, Building firmly nonexpansive convolutional neural networks, in Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing, 2020, pp. 8658–8662.
- [57] S. V. VENKATAKRISHNAN, C. A. BOUMAN, AND B. WOHLBERG, Plug-and-play priors for model based reconstruction, in 2013 IEEE Global Conference on Signal and Information Processing, 2013, pp. 945–948.
- [58] A. VIRMAUX AND K. SCAMAN, Lipschitz regularity of deep neural networks: analysis and efficient estimation, in Advances in Neural Information Processing Systems, 2018, pp. 3835– 3844.
- [59] Z. WANG, A. C. BOVIK, H. R. SHEIKH, AND E. P. SIMONCELLI, Image quality assessment: from error visibility to structural similarity, IEEE Trans. Image Process., 13 (2004), pp. 600–612.
- [60] X. WEI, B. GONG, Z. LIU, W. LU, AND L. WANG, Improving the improved training of wasserstein gans: A consistency term and its dual effect, in International Conference on Learning Representations, 2018.
- [61] L. XU, J. S. REN, C. LIU, AND J. JIA, Deep convolutional neural network for image deconvolution, in Advances in Neural Information Processing Systems, 2014, pp. 1790–1798.
- [62] X. XU, Y. SUN, J. LIU, B. WOHLBERG, AND U. S. KAMILOV, Provable convergence of plug-andplay priors with mmse denoisers, arXiv preprint arXiv:2005.07685, (2020).
- [63] Y. YOSHIDA AND T. MIYATO, Spectral norm regularization for improving the generalizability of deep learning, arXiv preprint arXiv:1705.10941, (2017).
- [64] J. ZHANG AND B. GHANEM, Ista-net: Interpretable optimization-inspired deep network for image compressive sensing, in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2018, pp. 1828–1837.
- [65] K. ZHANG, Y. LI, W. ZUO, L. ZHANG, L. VAN GOOL, AND R. TIMOFTE, Plug-and-play image restoration with deep denoiser prior, arXiv preprint, (2020).
- [66] K. ZHANG, W. ZUO, Y. CHEN, D. MENG, AND L. ZHANG, Beyond a gaussian denoiser: Residual learning of deep cnn for image denoising, IEEE Trans. Image Process., 26 (2017), pp. 3142– 3155.
- [67] K. ZHANG, W. ZUO, S. GU, AND L. ZHANG, Learning deep cnn denoiser prior for image restoration, in Proceedings of the IEEE conference on Computer Vision and Pattern Recognition, 2017, pp. 3929–3938.
- [68] K. ZHANG, W. ZUO, AND L. ZHANG, Learning a single convolutional super-resolution network for multiple degradations, in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2018, pp. 3262–3271.
- [69] K. ZHANG, W. ZUO, AND L. ZHANG, Deep plug-and-play super-resolution for arbitrary blur kernels, in IEEE Conference on Computer Vision and Pattern Recognition, 2019, pp. 1671– 1681.



LEARNING MAXIMALLY MONOTONE OPERATORS FOR IMAGE RECOVERY 29

Fig. 11: Visual results on Flickr30 (first two columns) and BSD500 (last three columns) datasets for different methods (VAR, PDHG, PnP-FB with BM3D, PnP-FB with DnCNN and ours) for different deblurring problems with $\nu = 0.01$. Associated (PSNR, SSIM) values are indicated below each image, best values are highlighted in bold. Except PDHG, each algorithm is stopped after 1000 iterations.