DOLPHIn - Dictionary Learning for Phase Retrieval

Andreas Tillmann, Yonina Eldar, Julien Mairal

► To cite this version:

Andreas Tillmann, Yonina Eldar, Julien Mairal. DOLPHIn - Dictionary Learning for Phase Retrieval. IEEE Transactions on Signal Processing, 2016, 64 (24), pp.6485-6500. 10.1109/TSP.2016.2607180 . hal-01387428

HAL Id: hal-01387428

<https://hal.inria.fr/hal-01387428>

Submitted on 25 Oct 2016

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.

DOLPHIn – Dictionary Learning for Phase Retrieval

Andreas M. Tillmann, Yonina C. Eldar, *Fellow, IEEE*, and Julien Mairal, *Senior Member, IEEE*

Abstract—We propose a new algorithm to learn a dictionary for reconstructing and sparsely encoding signals from measurements without phase. Specifically, we consider the task of estimating a two-dimensional image from squared-magnitude measurements of a complex-valued linear transformation of the original image. Several recent phase retrieval algorithms exploit underlying sparsity of the unknown signal in order to improve recovery performance. In this work, we consider such a sparse signal prior in the context of phase retrieval, when the sparsifying dictionary is not known in advance. Our algorithm jointly reconstructs the unknown signal—possibly corrupted by noise—and learns a dictionary such that each patch of the estimated image can be sparsely represented. Numerical experiments demonstrate that our approach can obtain significantly better reconstructions for phase retrieval problems with noise than methods that cannot exploit such “hidden” sparsity. Moreover, on the theoretical side, we provide a convergence result for our method.

Index Terms—(MLR-DICT, MLR-LEAR, OPT-NCVX, OPT-SOPT) Machine Learning, Signal Reconstruction, Image Reconstruction

I. INTRODUCTION

PHASE retrieval has been an active research topic for decades [1], [2]. The underlying goal is to estimate an unknown signal from the modulus of a complex-valued linear transformation of the signal. With such nonlinear measurements, the phase information is lost (hence the name “phase retrieval”), rendering the recovery task ill-posed and, perhaps not surprisingly, NP-hard [3]. Traditional approaches consider cases where the solution is unique up to a global phase shift, which can never be uniquely resolved, and devise signal reconstruction algorithms for such settings. Uniqueness properties and the empirical success of recovery algorithms usually hinge on oversampling the signal, i.e., taking more measurements than the number of signal components.

The most popular techniques for phase retrieval are based on alternating projections, see [4], [5], [6] for overviews.

Copyright (c) 2015 IEEE. Personal use of this material is permitted. However, permission to use this material for any other purposes must be obtained from the IEEE by sending a request to pubs-permissions@ieee.org.

A. M. Tillmann is with the TU Darmstadt, Research Group Optimization, Dolivostr. 15, 64293 Darmstadt, Germany (e-mail: tillmann@mathematik.tu-darmstadt.de). Y. C. Eldar is with the Department of Electrical Engineering, Technion – Israel Institute of Technology, Haifa 32000, Israel (e-mail: yonina@ee.technion.ac.il). J. Mairal is with Inria, Lear Team, Laboratoire Jean Kuntzmann, CNRS, Université Grenoble Alpes, 655, Avenue de l’Europe, 38330 Montbonnot, France (e-mail: julien.mairal@inria.fr).

The work of J. Mairal was funded by the French National Research Agency [Macaron project, ANR-14-CE23-0003-01]. The work of Y. Eldar was funded by the European Union’s Horizon 2020 research and innovation programme under grant agreement ERC-BNYQ, and by the Israel Science Foundation under Grant no. 335/14.

Manuscript received February 6, 2016, revised July 26, 2016, accepted August 25, 2016.

These methods usually require precise prior information about the signal (such as knowledge of the support set) and often converge to erroneous results. More recent approaches include semidefinite programming relaxations [7], [8], [9], [10], [11] and gradient-based methods such as Wirtinger Flow [12], [13].

In recent years, new phase retrieval techniques were developed for recovering *sparse* signals, which are linear combinations of only a few atoms from a known dictionary [8], [14], [13], [15]. With a sparsity assumption, these algorithms obtained better recovery performance than traditional non-sparse approaches. The main idea is akin to compressed sensing, where one works with fewer (linear) measurements than signal components [16], [17], [18]. An important motivation for developing sparse recovery techniques was that many classes of signals admit a sparse approximation in some basis or overcomplete dictionary [19], [20], [21]. While sometimes such dictionaries are known explicitly, better results have been achieved by adapting the dictionary to the data, e.g., for image denoising [20]. Numerous algorithms have been developed for this task, see, e.g., [19], [22], [23]. In this traditional setting, the signal measurements are *linear* and a large database of training signals is used to train the dictionary.

In this work, we propose a dictionary learning formulation for simultaneously solving the signal reconstruction and sparse representation problems given *nonlinear*, *phaseless* and *noisy* measurements. To optimize the resulting (nonconvex) objective function, our algorithm—referred to as DOLPHIn (DictiOnary Learning for PHase retrIeval)—alternates between several minimization steps, thus monotonically reducing the value of the objective until a stationary point is found (if step sizes are chosen appropriately). Specifically, we iterate between best fitting the data and sparsely representing the recovered signal. DOLPHIn combines projected gradient descent steps to update the signal, iterative shrinkage to obtain a sparse approximation [24], and block-coordinate descent for the dictionary update [23].

In various experiments on image reconstruction problems, we demonstrate the ability of DOLPHIn to achieve significantly improved results when the oversampling ratio is low and the noise level high, compared to the recent state-of-the-art Wirtinger Flow (WF) method [12], which cannot exploit sparsity if the dictionary is unknown. In this two-dimensional setting, we break an image down into small patches and train a dictionary such that each patch can be sparsely represented using this dictionary. The patch size as well as the amount of overlap between patches can be freely chosen, which allows us to control the trade-off between the amount of computation required to reconstruct the signal and the quality of the result.

The paper is organized as follows: In Sections II and III,

we introduce the DOLPHIn framework and algorithm. Then, in Section IV, we present numerical experiments and implementation details, along with discussions about (hyper-)parameter selection and variants of DOLPHIn. We conclude the paper in Section V. The appendix provides further details on the mathematical derivation of the DOLPHIn algorithm. A short preliminary version of this work appeared in the conference paper [25].

II. PHASE RETRIEVAL MEETS DICTIONARY LEARNING

In mathematical terms, the phase retrieval problem can be formulated as solving a nonlinear system of equations:

$$\text{Find } \mathbf{x} \in \mathcal{X} \subseteq \mathbb{C}^N \quad \text{s.t.} \quad |f_i(\mathbf{x})|^2 = y_i \quad \forall i = 1, \dots, M, \quad (1)$$

where the functions $f_i : \mathbb{C}^N \rightarrow \mathbb{C}$ are linear operators and the scalars y_i are nonlinear measurements of the unknown original signal $\hat{\mathbf{x}}$ in $\mathcal{X}$, obtained by removing the phase information. The set $\mathcal{X}$ represents constraints corresponding to additional prior information about $\hat{\mathbf{x}}$. For instance, when dealing with real-valued bounded signals, this may typically be a box constraint $\mathcal{X} = [0, 1]^N$. Other common constraints include information about the support set—that is, the set of nonzero coefficients of $\hat{\mathbf{x}}$. Classical phase retrieval concerns the recovery of $\hat{\mathbf{x}}$ given the (squared) modulus of the signal’s Fourier transform. Other commonly considered cases pertain to randomized measurements (f_i are random linear functions) or coded diffraction patterns, i.e., concatenations of random signal masks and Fourier transforms (see, e.g., [2], [12]).

A. Prior and Related Work

The most popular methods for classical phase retrieval—Fienup’s algorithm [5] and many related approaches [2], [4], [6], [26], [27]—are based on alternating projections onto the sets $\mathcal{Y} := \{\mathbf{x} \in \mathbb{C}^N \text{ s.t. } |f_i(\mathbf{x})| = y_i \forall i\}$ (or $\{\mathbf{x} \in \mathbb{C}^N \text{ s.t. } |f_i(\mathbf{x})|^2 = y_i \forall i\}$) and onto the set $\mathcal{X}$. However, the nonconvexity of $\mathcal{Y}$ makes the projection not uniquely defined and possibly hard to compute. The success of such projection-based methods hinges critically on precise prior knowledge (which, in general, will not be available in practice) and on the choice of a projection operator onto $\mathcal{Y}$. Ultimately, convergence to $\hat{\mathbf{x}}$ (up to global phase) is in general not guaranteed and these methods often fail in practice.

Further algorithmic techniques to tackle (1) include two different semidefinite relaxation approaches, PhaseLift [7] and PhaseCut [11]. PhaseLift “lifts” (1) into the space of (complex) positive semidefinite rank-1 matrices via the variable transformation $\mathbf{X} := \mathbf{x}\mathbf{x}^*$. Then, the nonlinear constraints $|f_i(\mathbf{x})|^2 = y_i$ are equivalent to *linear* constraints with respect to the matrix variable $\mathbf{X}$. By suitably relaxing the immediate but intractable rank-minimization objective, one obtains a convex semidefinite program (SDP). Similarly, PhaseCut introduces a separate variable $\mathbf{u}$ for the phase, allowing to eliminate $\mathbf{x}$, and then lifts $\mathbf{u}$ to obtain an equivalent problem with a rank-1-constraint, which can be dropped to obtain a different SDP relaxation of (1). Despite some guarantees on when these relaxations are tight, i.e., allow for correctly recovering the

solution to (1) (again up to a global phase factor), their practical applicability is limited due to the dimension of the SDP that grows quadratically with the problem dimension.

A recent method that works in the original variable space is the so-called *Wirtinger Flow* algorithm [12]. Here, (1) is recast as the optimization problem

$$\min_{\mathbf{x} \in \mathbb{C}^N} \frac{1}{4M} \sum_{i=1}^M (|f_i(\mathbf{x})|^2 - y_i)^2, \quad (2)$$

which is approximately solved by a gradient descent algorithm. Note that in the case of complex variables, the concept of a gradient is not well-defined, but as shown in [12], a strongly related expression termed the “Wirtinger derivative” can be used instead and indeed reduces to the actual gradient in the real case. For the case of i.i.d. Gaussian random measurements, local convergence with high probability can be proven for the method, and a certain spectral initialization provides sufficiently accurate signal estimates for these results to be applicable. Further variants of the Wirtinger Flow (WF) method that have been investigated are the Truncated WF [28], which involves improving search directions by a statistically motivated technique to filter out components that bear “too much” influence, and Thresholded WF [13], which allows for improved reconstruction of *sparse* signals (i.e., ones with only a few significant components or nonzero elements), in particular when the measurements are corrupted by noise.

The concept of sparsity has been successfully employed in the context of signal reconstruction from *linear* measurements, perhaps most prominently in the field of *compressed sensing* [16], [17], [18], [29] during the past decade. There, the task is to recover an unknown signal $\hat{\mathbf{x}} \in \mathbb{C}^N$ from $M < N$ linear measurements—that is, finding the desired solution among the infinitely many solutions of an underdetermined system of linear equations. For signals that are (exactly or approximately) sparse with respect to some basis or dictionary, i.e., when $\hat{\mathbf{x}} \approx \mathbf{D}\hat{\mathbf{a}}$ for a matrix $\mathbf{D}$ and a vector $\hat{\mathbf{a}}$ that has few nonzero entries, such recovery problems have been shown to be solvable in a very broad variety of settings and applications, and with a host of different algorithms. Dictionaries enabling sparse signal representations are sometimes, but not generally, known in advance. The goal of *dictionary learning* is to improve upon the sparsity achievable with a given (analytical) dictionary, or to find a suitable dictionary in the first place. Given a set of training signals, the task consists of finding a dictionary such that every training signal can be well-approximated by linear combinations of just a few atoms. Again, many methods have been developed for this purpose (see, e.g., [19], [20], [21], [22], [23]) and demonstrated to work well in different practical applications.

Signal sparsity (or compressability) can also be beneficially exploited in phase retrieval methods, cf. [8], [9], [13], [14], [15]. However, to the best of our knowledge, existing methods assume that the signal is sparse itself or sparse with respect to a fixed pre-defined dictionary. This motivates the development of new algorithms and formulations to *jointly* learn suitable dictionaries and reconstruct input signals from nonlinear measurements.

B. Dictionary Learning for Phase Retrieval

In this paper, we consider the problem of phase retrieval by focusing on image reconstruction applications. Therefore, we will work in a two-dimensional setting directly. However, it should be noted that all expressions and algorithms can also easily be formulated for one-dimensional signals like (1), as detailed in the appendix. We will also consider the case of noisy measurements, and will show that our approach based on dictionary learning is particularly robust to noise, which is an important feature in practice.

Concretely, we wish to recover an image $\hat{\mathbf{X}}$ in $[0, 1]^{N_1 \times N_2}$ from noise-corrupted phaseless nonlinear measurements

$$\mathbf{Y} := |\mathcal{F}(\hat{\mathbf{X}})|^2 + \mathbf{N}, \quad (3)$$

where $\mathcal{F} : \mathbb{C}^{N_1 \times N_2} \rightarrow \mathbb{C}^{M_1 \times M_2}$ is a linear operator, $\mathbf{N}$ is a real matrix whose entries represent noise, and the complex modulus and squares are taken component-wise. As mentioned earlier, signal sparsity is known to improve the performance of phase retrieval algorithms, but a sparsifying transform is not always known in advance, or a better choice than a predefined selection can sometimes be obtained by adapting the dictionary to the data. In the context of image reconstruction, this motivates learning a dictionary $\mathbf{D}$ in $\mathbb{R}^{s \times n}$ such that each $s_1 \times s_2$ patch $\hat{\mathbf{x}}^i$ of $\hat{\mathbf{X}}$, represented as a vector of size $s = s_1 s_2$, can be approximated by $\hat{\mathbf{x}}^i \approx \mathbf{D} \mathbf{a}^i$ with a sparse vector $\mathbf{a}^i$ in $\mathbb{R}^n$. Here, n is chosen a priori and the number of patches depends on whether the patches are overlapping or not. In general, $\mathbf{D}$ is chosen such that $n \geq s$. With linear measurements, the paradigm would be similar to the successful image denoising technique of [20], but the problem (3) is significantly more difficult to solve due to the modulus operator.

Before detailing our algorithm for solving (3), we introduce the following notation. Because our approach is patch-based (as most dictionary learning formulations), we consider the linear operator $\mathcal{E} : \mathbb{C}^{N_1 \times N_2} \rightarrow \mathbb{C}^{s \times p}$ that extracts the p patches $\mathbf{x}^i$ (which may overlap or not) from an image $\mathbf{X}$ and forms the matrix $\mathcal{E}(\mathbf{X}) = (\mathbf{x}^1, \dots, \mathbf{x}^p)$. Similarly, we define the linear operator $\mathcal{R} : \mathbb{C}^{s \times p} \rightarrow \mathbb{C}^{N_1 \times N_2}$ that reverses this process, i.e., builds an image from a matrix containing vectorized patches as its columns. Thus, in particular, we have $\mathcal{R}(\mathcal{E}(\mathbf{X})) = \mathbf{X}$. When the patches do not overlap, the operator $\mathcal{R}$ simply places every small patch at its appropriate location in a larger $N_1 \times N_2$ image. When they overlap, the operator $\mathcal{R}$ averages the pixel values from the patches that fall into the same location. Further, let $\mathbf{A} := (\mathbf{a}^1, \dots, \mathbf{a}^p)$ in $\mathbb{R}^{n \times p}$ be the matrix containing the patch representation coefficient vectors as columns. Then, our desired sparse-approximation relation “ $\mathbf{x}^i \approx \mathbf{D} \mathbf{a}^i$ for all i ” can be expressed as $\mathcal{E}(\mathbf{X}) \approx \mathbf{D} \mathbf{A}$.

With this notation in hand, we may now introduce our method, called DOLPHIn (DictIOnary Learning for PHase retrIeval). We consider an optimization problem which can be interpreted as a combination of an optimization-based approach to phase retrieval—minimizing the residual norm with respect to the set of nonlinear equations induced by the phaseless measurements, cf. (2)—and a (patch-based) dictionary learning model similar to that used for image denoising in [20]. The model contains three variables: The

image, or phase retrieval solution $\mathbf{X}$, the dictionary $\mathbf{D}$ and the matrix $\mathbf{A}$ containing as columns the coefficient vectors of the representation $\mathbf{X} \approx \mathcal{R}(\mathbf{D} \mathbf{A})$. The phase retrieval task consists of estimating $\mathbf{X}$ and the dictionary learning or sparse coding task consists of estimating $\mathbf{D}$ and $\mathbf{A}$; a common objective function provides feedback between the two objectives, with the goal of improving the phase retrieval reconstruction procedure by encouraging the patches of $\mathbf{X}$ to admit a sparse approximation.

Formally, the DOLPHIn formulation consists of minimizing

$$\begin{aligned} \min_{\mathbf{X}, \mathbf{D}, \mathbf{A}} \quad & \frac{1}{4} \|\mathbf{Y} - |\mathcal{F}(\mathbf{X})|^2\|_F^2 + \frac{\mu}{2} \|\mathcal{E}(\mathbf{X}) - \mathbf{D} \mathbf{A}\|_F^2 + \lambda \sum_{i=1}^p \|\mathbf{a}^i\|_1 \\ \text{s.t.} \quad & \mathbf{X} \in [0, 1]^{N_1 \times N_2}, \quad \mathbf{D} \in \mathcal{D}. \end{aligned} \quad (4)$$

Here, $\|\mathbf{X}\|_F$ denotes the Frobenius matrix-norm, which generalizes the Euclidean norm to matrices. The parameters $\mu, \lambda > 0$ in the objective (4) provide a way to control the trade-off between the data fidelity term from the phase retrieval problem and the approximation sparsity of the image patches¹. To that effect, we use the ℓ_1 -norm, which is well-known to have a sparsity-inducing effect [30]. In order to avoid scaling ambiguities, we also restrict $\mathbf{D}$ to be in the subset $\mathcal{D} := \{\mathbf{D} \in \mathbb{R}^{s \times n} : \|\mathbf{d}^j\|_2 \leq 1 \ \forall j = 1, \dots, n\}$ of matrices with column ℓ_2 -norms at most 1, and assume $n < p$ (otherwise, each patch is trivially representable by a 1-sparse vector $\mathbf{a}^i$ by including $\mathbf{x}^i / \|\mathbf{x}^i\|_2$ as a column of $\mathbf{D}$).

The model (4) could also easily be modified to include further side constraints, a different type of nonlinear measurements, or multiple images or measurements, respectively; we omit these extensions for simplicity.

III. ALGORITHMIC FRAMEWORK

Similar to classical dictionary learning [19], [22], [21], [31] and phase retrieval, problem (4) is nonconvex and difficult to solve. Therefore, we adopt an algorithm that provides monotonic decrease of the objective while converging to a stationary point (see Section III-D below).

The algorithmic framework we employ is that of *alternating minimization*: For each variable $\mathbf{A}$, $\mathbf{X}$ and $\mathbf{D}$ in turn, we take one step towards solving (4) with respect to this variable alone, keeping the other ones fixed. Each of these subproblems is convex in the remaining unfixed optimization variable, and well-known efficient algorithms can be employed accordingly. We summarize our method in Algorithm 1, where the superscript $*$ denotes the adjoint operator (for a matrix $\mathbf{Z}$, $\mathbf{Z}^*$ is thus the conjugate transpose), $\Re(\cdot)$ extracts the real part of a complex-valued argument, and $\odot$ denotes the Hadamard (element-wise) product of two matrices. The algorithm also involves the classical soft-thresholding operator $\mathcal{S}_\tau(\mathbf{Z}) := \max\{0, |\mathbf{Z}| - \tau\} \odot \text{sign}(\mathbf{Z})$ and the Euclidean projection $\mathcal{P}_{\mathcal{X}}(\mathbf{Z}) := \max\{0, \min\{1, |\mathbf{Z}|\}\}$ onto $\mathcal{X} := [0, 1]^{N_1 \times N_2}$; here, all operations are meant component-wise.

¹We discuss suitable choices and sensitivity of the model to these parameters in detail in Section IV-D.

To avoid training the dictionary on potentially useless early estimates, the algorithm performs two phases—while the iteration counter ℓ is smaller than K_1 , the dictionary is not updated. Below, we explain the algorithmic steps in more detail.

Note that DOLPHIn actually produces two distinct reconstructions of the desired image, namely $\mathbf{X}$ (the per se “image variable”) and $\mathcal{R}(\mathbf{DA})$ (the image assembled from the sparsely coded patches)². Our numerical experiments in Section IV show that in many cases, $\mathcal{R}(\mathbf{DA})$ is in fact slightly or even significantly better than $\mathbf{X}$ with respect to at least one quantitative quality measure and is therefore also considered a possible reconstruction output of Algorithm 1 (at least in the noisy setups we consider in this paper). Nevertheless, $\mathbf{X}$ is sometimes more visually appealing and can be used, for instance, to refine parameter settings (if it differs strongly from $\mathcal{R}(\mathbf{DA})$) or to assess the influence of the patch-based “regularization” on the pure non-sparse Wirtinger Flow method corresponding to the formulation where λ and μ are set to zero.

A. Updating the Patch Representation Vectors

Updating $\mathbf{A}$ (i.e., considering (4) with $\mathbf{D}$ and $\mathbf{X}$ fixed at their current values) consists of decreasing the objective

$$\sum_{i=1}^p \left(\frac{1}{2} \|\mathbf{D}_{(\ell)} \mathbf{a}^i - \mathbf{x}_{(\ell)}^i\|_2^2 + \frac{\lambda}{\mu} \|\mathbf{a}^i\|_1 \right), \quad (5)$$

which is separable in the patches $i = 1 \dots p$. Therefore, we can update all vectors $\mathbf{a}^i$ independently and/or in parallel. To do so, we choose to perform one step of the well-known algorithm ISTA (see, e.g., [24]), which is a gradient-based method that is able to take into account a non-smooth regularizer such as the ℓ_1 -norm. Concretely, the following update is performed for each $i = 1, \dots, p$:

$$\mathbf{a}_{(\ell+1)}^i = \mathcal{S}_{\gamma_\ell^A \lambda / \mu} \left(\mathbf{a}_{(\ell)}^i - \gamma_\ell^A \mathbf{D}_{(\ell)}^\top (\mathbf{D}_{(\ell)} \mathbf{a}_{(\ell)}^i - \mathbf{x}_{(\ell)}^i) \right). \quad (6)$$

This update involves a gradient descent step (the gradient with respect to $\mathbf{a}^i$ of the smooth term in each summand of (5) is $\mathbf{D}_{(\ell)}^\top (\mathbf{D}_{(\ell)} \mathbf{a}_{(\ell)}^i - \mathbf{x}_{(\ell)}^i)$, respectively) followed by soft-thresholding. Constructing $\mathbf{A}_{(\ell+1)}$ from the $\mathbf{a}_{(\ell+1)}^i$ as specified above is equivalent to Step 2 of Algorithm 1.

The step size parameter γ_ℓ^A can be chosen in $(0, 1/L_A)$, where L_A is an upper bound on the Lipschitz constant of the gradient; here, $L_A = \|\mathbf{D}_{(\ell)}^\top \mathbf{D}_{(\ell)}\|_2 = \|\mathbf{D}_{(\ell)}\|_2^2$ would be appropriate, but a less computationally demanding strategy is to use a backtracking scheme to automatically update L_A [24].

A technical subtlety is noteworthy in this regard: We can either find one γ_ℓ^A that works for the whole matrix-variable update problem—this is what is stated implicitly in Step 2— or we could find different values, say $\gamma_\ell^{a,i}$, for each column $\mathbf{a}^i$, $i = 1, \dots, p$, of $\mathbf{A}$ separately. Our implementation does the latter, since it employs a backtracking strategy for each column update independently.

²Technically, $\mathcal{R}(\mathbf{DA})$ might contain entries not in $\mathcal{X}$, so one should project once more. Throughout, we often omit this step for simplicity; differences (if any) between $\mathcal{R}(\mathbf{DA})$ and $\mathcal{P}_{\mathcal{X}}(\mathcal{R}(\mathbf{DA}))$ were insignificant in all our tests.

B. Updating the Image Estimate

With $\mathbf{D} = \mathbf{D}_{(\ell)}$ and $\mathbf{A} = \mathbf{A}_{(\ell+1)}$ fixed, updating $\mathbf{X}$ consists of decreasing the objective

$$\frac{1}{4} \|\mathbf{Y} - |\mathcal{F}(\mathbf{X})|^2\|_F^2 + \frac{\mu}{2} \|\mathcal{E}(\mathbf{X}) - \mathbf{DA}\|_F^2 \quad (7)$$

with $\mathbf{X} \in \mathcal{X} = [0, 1]^{N_1 \times N_2}$.

This problem can be seen as a regularized version of the phase retrieval problem (with regularization parameter μ) that encourages the patches of $\mathbf{X}$ to be close to the sparse approximation $\mathbf{DA}$ obtained during the previous (inner) iterations.

Our approach to decrease the value of the objective (7) is by a projected gradient descent step. In fact, for $\mu = 0$, this step reduces to the Wirtinger flow method [12], but with necessary modifications to take into account the constraints on $\mathbf{X}$ (real-valuedness and variable bounds $[0, 1]$).

The gradient of $\varphi(\mathbf{X}) := \frac{1}{4} \|\mathbf{Y} - |\mathcal{F}(\mathbf{X})|^2\|_F^2$ with respect to $\mathbf{X}$ can be computed as

$$\nabla \varphi(\mathbf{X}) = \Re \left(\mathcal{F}^* (\mathcal{F}(\mathbf{X}) \odot (|\mathcal{F}(\mathbf{X})|^2 - \mathbf{Y})) \right),$$

by using the chain rule. For $\psi(\mathbf{X}) := \frac{\mu}{2} \|\mathcal{E}(\mathbf{X}) - \mathbf{DA}\|_F^2$, the gradient is given by

$$\nabla \psi(\mathbf{X}) = \mu \mathcal{E}^* (\mathcal{E}(\mathbf{X}) - \mathbf{DA}) = \mu \mathbf{R} \odot \mathcal{R}(\mathcal{E}(\mathbf{X}) - \mathbf{DA}),$$

where $\mathbf{R}$ is an $N_1 \times N_2$ matrix whose entries r_{ij} equal the number of patches the respective pixel x_{ij} is contained in. Note that if the whole image is divided into a complete set of nonoverlapping patches, $\mathbf{R}$ will just be the all-ones matrix; otherwise, the element-wise multiplication with $\mathbf{R}$ undoes the averaging of pixel values performed by $\mathcal{R}$ when assembling an image from overlapping patches.

Finally, the gradient w.r.t. $\mathbf{X}$ of the objective in (7) is $\nabla \varphi(\mathbf{X}) + \nabla \psi(\mathbf{X}) \in \mathbb{R}^{N_1 \times N_2}$, and the update in Step 3 of Algorithm 1 is indeed shown to be a projected gradient descent step. Typically, a backtracking (line search) strategy is used for choosing the step size γ_ℓ^X ; see Theorem 2 in Section III-D for a selection rule that gives theoretical convergence, and also Section IV-B for a heuristic alternative.

C. Updating the Dictionary

To update the dictionary, i.e., to approximately solve (4) w.r.t. $\mathbf{D}$ alone, keeping $\mathbf{X}$ and $\mathbf{A}$ fixed at their current values, we employ one pass of a block-coordinate descent (BCD) algorithm on the columns of the dictionary [23]. The objective to decrease may be written as

$$\frac{1}{2} \sum_{i=1}^p \|\mathbf{D} \mathbf{a}_{(\ell+1)}^i - \mathbf{x}_{(\ell+1)}^i\|_2^2 \quad \text{s.t.} \quad \mathbf{D} \in \mathcal{D}, \quad (8)$$

and the update rule given by Steps 4–13 corresponds³ to one iteration of [21, Algorithm 11] applied to (8).

³In [21, Algo. 11], and in our implementation, we simply normalize the columns of $\mathbf{D}$; it is easily seen that any solution with $\|\mathbf{d}^j\|_2 < 1$ for some j is suboptimal (w.r.t. (4)) since raising it to 1 allows to reduce coefficients in $\mathbf{A}$ and thus to improve the ℓ_1 -term of the DOLPHIn objective (4). However, using the projection is more convenient for proving the convergence results without adding more technical subtleties w.r.t. this aspect.

Algorithm 1 Dictionary learning for phase retrieval (DOLPHIn)

Input: Initial image estimate $\mathbf{X}_{(0)} \in [0, 1]^{N_1 \times N_2}$, initial dictionary $\mathbf{D}_{(0)} \in \mathcal{D} \subset \mathbb{R}^{s \times n}$, parameters $\mu, \lambda > 0$, maximum number of iterations K_1, K_2

Output: Learned dictionary $\mathbf{D} = \mathbf{D}_{(K)}$, patch representations $\mathbf{A} = \mathbf{A}_{(K)}$, image reconstructions $\mathbf{X} = \mathbf{X}_{(K)}$ and $\mathcal{R}(\mathbf{D}\mathbf{A})$

1: **for** $\ell = 0, 1, 2, \dots, K_1 + K_2 =: K$ **do**

2: choose step size γ_ℓ^A as explained in Section III-A and update

$$\mathbf{A}_{(\ell+1)} \leftarrow \mathcal{S}_{\gamma_\ell^A \lambda / \mu} \left(\mathbf{A}_\ell - \gamma_\ell^A \mathbf{D}_{(\ell)}^\top (\mathbf{D}_{(\ell)} \mathbf{A}_\ell - \mathcal{E}(\mathbf{X}_{(\ell)})) \right)$$

3: choose step size γ_ℓ^X as explained in Section III-D or IV-B and update

$$\mathbf{X}_{(\ell+1)} \leftarrow \mathcal{P}_X \left(\mathbf{X}_{(\ell)} - \gamma_\ell^X \left(\Re \left(\mathcal{F}^* (\mathcal{F}(\mathbf{X}) \odot (|\mathcal{F}(\mathbf{X})|^2 - \mathbf{Y})) \right) + \mu \mathbf{R} \odot \mathcal{R}(\mathcal{E}(\mathbf{X}) - \mathbf{D}\mathbf{A}) \right) \right),$$

where $\mathbf{R}$ is defined in Section III-B

4: **if** $\ell < K_1$ **then**

5: do not update the dictionary: $\mathbf{D}_{(\ell+1)} \leftarrow \mathbf{D}_{(\ell)}$

6: **else**

7: set $\mathbf{B} \leftarrow \mathcal{E}(\mathbf{X}_{(\ell)}) \mathbf{A}_{(\ell)}^\top$ and $\mathbf{C} \leftarrow \mathbf{A}_{(\ell)} \mathbf{A}_{(\ell)}^\top$

8: **for** $j = 1, \dots, n$ **do**

9: **if** $C_{jj} > 0$ **then**

10: update j -th column: $(\mathbf{D}_{(\ell+1)})_{\cdot j} \leftarrow \frac{1}{C_{jj}} (\mathbf{B}_{\cdot j} - \mathbf{D}_{(\ell)} \mathbf{C}_{\cdot j}) + (\mathbf{D}_{(\ell)})_{\cdot j}$

11: **else**

12: reset j -th column: e.g., $(\mathbf{D}_{(\ell+1)})_{\cdot j} \leftarrow$ random $\mathcal{N}(0, 1)$ vector (in $\mathbb{R}^s$)

13: project $(\mathbf{D}_{(\ell+1)})_{\cdot j} \leftarrow \frac{1}{\max\{1, \|(\mathbf{D}_{(\ell+1)})_{\cdot j}\|_2\}} (\mathbf{D}_{(\ell+1)})_{\cdot j}$

To see this, note that each column update problem has a closed-form solution:

$$\begin{aligned} (\mathbf{d}^j)_{(\ell+1)} &= \mathcal{P}_{\|\cdot\|_2 \leq 1} \left(\frac{1}{\sum_{i=1}^p (a_j^i)^2} \sum_{i=1}^p a_j^i (\mathbf{x}^i - \sum_{\substack{k=1 \\ k \neq j}}^n a_k^i \mathbf{d}^k) \right) \\ &= \frac{1}{\max\{1, \|\frac{1}{w_j} \mathbf{q}^j\|_2\}} \left(\frac{1}{w_j} \mathbf{q}^j \right) \end{aligned}$$

with $w_j := \sum_i (a_j^i)^2$ and $\mathbf{q}^j := \sum_i a_j^i (\mathbf{x}^i - \sum_{k \neq j} a_k^i \mathbf{d}^k)$; here, we abbreviated $\mathbf{a}^i := \mathbf{a}_{(\ell+1)}^i$, $\mathbf{x}^i := \mathbf{x}_{(\ell+1)}^i$. If $w_j = 0$, and thus $a_j^i = 0$ for all $i = 1, \dots, p$, then column $\mathbf{d}^j$ is not used in any of the current patch representations; in that case, the column's update problem has a constant objective and is therefore solved by any $\mathbf{d}$ with $\|\mathbf{d}\|_2 \leq 1$, e.g., a normalized random vector as in Step 12 of Algorithm 1. The computations performed in Steps 8–13 of Algorithm 1 are equivalent to these solutions, expressed differently using the matrices $\mathbf{B}$ and $\mathbf{C}$ defined there. Note that the operations could be parallelized to speed up computation.

D. Convergence of the Algorithm

As mentioned at the beginning of this section, with appropriate step size choices, DOLPHIn (Algorithm 1) exhibits the property of monotonically decreasing the objective function value (4) at each iteration. In particular, many line-search type step size selection mechanisms aim precisely at reducing the objective; for simplicity, we will simply refer to such sub-routines as “suitable backtracking schemes” below. Concrete examples are the ISTA backtracking from [24, Section 3] we can employ in the update of $\mathbf{A}$, or the rule given in

Theorem 2 for the projected gradient descent update of $\mathbf{X}$ (a different choice is described in Section IV-B); further variants are discussed, e.g., in [32].

Proposition 1: Let $(\mathbf{A}_{(\ell)}, \mathbf{X}_{(\ell)}, \mathbf{D}_{(\ell)})$ be the current iterates (after the ℓ -th iteration) of Algorithm 1 with step sizes γ_ℓ^X and γ_ℓ^A determined by suitable backtracking schemes (or arbitrary $0 < \gamma_\ell^A < 1/\|\mathbf{D}_{(\ell)}^\top \mathbf{D}_{(\ell)}\|_2$, resp.) and let $f_{i,j,k}$ denote the objective function value of the DOLPHIn model (4) at $(\mathbf{A}_{(i)}, \mathbf{X}_{(j)}, \mathbf{D}_{(k)})$. Then, DOLPHIn either terminates in the $(\ell + 1)$ -th iteration, or it holds that $f_{\ell+1, \ell+1, \ell+1} \leq f_{\ell, \ell, \ell}$.

Proof: Since we use ISTA to update $\mathbf{A}$, it follows from [24] that $f_{\ell+1, \ell, \ell} \leq f_{\ell, \ell, \ell}$. Similarly, a suitable backtracking strategy is known to enforce descent in the projected gradient method when the gradient is locally Lipschitz-continuous, whence $f_{\ell+1, \ell+1, \ell} \leq f_{\ell+1, \ell, \ell}$. Finally, $f_{\ell+1, \ell+1, \ell+1} \leq f_{\ell+1, \ell+1, \ell}$ follows from standard results for BCD methods applied to convex problems, see, e.g., [33]. Combining these inequalities proves the claim. ■

The case of termination in Proposition 1 can occur when the backtracking scheme is combined with a maximal number of trial steps, which are often used as a safeguard against numerical stability problems or as a heuristic stopping condition to terminate the algorithm if no (sufficient) improvement can be reached even with tiny step sizes. Note also that the assertion of Proposition 1 trivially holds true if all step sizes are 0; naturally, a true descent of the objective requires a strictly positive step size in at least one update step. In our algorithm, step size positivity can always be guaranteed since all these updates involve objective functions whose gradient is Lipschitz continuous, and backtracking essentially finds step sizes inversely proportional to (local) Lipschitz constants. (Due

to non-expansiveness, the projection in the $\mathbf{X}$ -update poses no problem either).

Adopting a specific Armijo-type step size selection rule for the $\mathbf{X}$ -update allows us to infer a convergence result, stated in Theorem 2 below. To simplify the presentation, let $f_\ell^X(\mathbf{X}) := \frac{1}{4} \|\mathbf{Y} - |\mathcal{F}(\mathbf{X})|^2\|_F^2 + \frac{\mu}{2} \|\mathcal{E}(\mathbf{X}) - \mathbf{D}_{(\ell)} \mathbf{A}_{(\ell+1)}\|_F^2$ and $\mathbf{\Gamma}_\ell^X := \nabla f_\ell^X(\mathbf{X}_{(\ell)})$ (cf. $\nabla \varphi(\mathbf{X}) + \nabla \psi(\mathbf{X})$ in Section III-B).

Theorem 2: Let $\eta \in (0, 1)$ and $\bar{\eta} > 0$. Consider the DOLPHIn variant consisting of Algorithm 1 with $K_2 = \infty$ and the following Armijo rule to be used in Step 3:

Determine γ_ℓ^X as the largest number in $\{\bar{\eta}\eta^k\}_{k=0,1,2,\dots}$ such that $f_\ell^X(\mathcal{P}_X(\mathbf{X}_{(\ell)} - \gamma_\ell^X \mathbf{\Gamma}_\ell^X)) - f_\ell^X(\mathbf{X}_{(\ell)}) \leq -\frac{1}{2\gamma_\ell^X} \|\mathcal{P}_X(\mathbf{X}_{(\ell)} - \gamma_\ell^X \mathbf{\Gamma}_\ell^X) - \mathbf{X}_{(\ell)}\|_F^2$ and set $\mathbf{X}_{(\ell+1)} := \mathcal{P}_X(\mathbf{X}_{(\ell)} - \gamma_\ell^X \mathbf{\Gamma}_\ell^X)$.

If $\mu \geq 1$ and, for some $0 < \underline{\nu} \leq \bar{\nu}$, it holds that $\underline{\nu} \leq \gamma_\ell^X, \gamma_\ell^A$ (or $\gamma_\ell^{a,i}$), $\sum_{i=1}^p (\mathbf{a}_{(\ell)}^i)_j^2 \leq \bar{\nu}$ for all ℓ and j (and i), then every accumulation point of the sequence $\{(\mathbf{A}_{(\ell)}, \mathbf{X}_{(\ell)}, \mathbf{D}_{(\ell)})\}_{\ell=0,1,2,\dots}$ of DOLPHIn iterates is a stationary point of problem (4).

Proof: The proof works by expressing Algorithm 1 as a specific instantiation of the coordinate gradient descent (CGD) method from [34] and analyzing the objective descent achievable in each update of $\mathbf{A}$, $\mathbf{X}$ and $\mathbf{D}$, respectively. The technical details make the rigorous formal proof somewhat lengthy (it can be found in the appendix of the earlier preprint version [35] of this paper); due to space limitations, we only sketch it here.

The CGD method works by solving subproblems to obtain directions of improvement for blocks of variables at a time—in our case, (the columns of) $\mathbf{A}$, the matrix $\mathbf{X}$, and the columns of $\mathbf{D}$ correspond to such blocks—and then taking steps along these directions. More specifically, the directions are generated using a (suitably parameterized) strictly convex quadratic approximation of the objective (built using the gradient). Essentially due to the strict convexity, it is then always possible to make a *positive* step along such a direction that decreases the (original) objective, unless stationarity already holds. Using a certain Armijo line-search rule designed to find such positive step sizes which achieve a sufficient objective reduction, [34, Theorem 1] ensures (under mild further assumptions, which in our case essentially translate to the stated boundedness requirement of the step sizes) that every accumulation point of the iterate sequence is indeed a stationary point of the addressed (block-separable) problem.

To embed DOLPHIn into the CGD framework, we can interpret the difference between one iterate and the next (w.r.t. the variable “block” under consideration) as the improvement direction, and proceed to show that we can always choose a step size equal to 1 in the Armijo-criterion from [34] (cf. (9) and (46) therein). For this to work out, we need to impose slightly stricter conditions on other parameters used to define that rule than what is needed in [34]; these conditions are derived directly from known descent properties of the $\mathbf{D}$ - and $\mathbf{A}$ -updates of our method (essentially, ISTA descent properties as in [24]). That way, the $\mathbf{D}$ - and $\mathbf{A}$ -updates automatically satisfy the specific CGD Armijo rule, and the actual backtracking scheme for the $\mathbf{X}$ -update given in the present theorem can be shown to assert that our $\mathbf{X}$ -

update does so as well. (The step sizes used in DOLPHIn could also be reinterpreted in the CGD framework as scaling factors of diagonal Hessian approximations of the combined objective to be used in the direction-finding subproblems. With such simple Hessian approximations, the obtained directions are then indeed equivalent to the iterate-differences resulting from the DOLPHIn update schemes.) The claim then follows directly from [34, Theorem 1(e)] (and its extensions discussed in Section 8). ■

A more formal explanation for why the step sizes can be chosen positive in each step can be found on page 392 of [34]; the boundedness of approximate Hessians is stated in [34, Assumption 1]. Arguably, assuming step sizes are bounded away from zero by a constant may become problematic in theory (imagine an Armijo-generated step size sequence converging to zero), but will not pose a problem in practice where one always faces the limitations of numerical machine precision. (Note also that, in practice, the number of line-search trials can be effectively reduced by choosing $\bar{\eta}$ based on the previous step size [34].)

Our implementation uses a different backtracking scheme for the $\mathbf{X}$ -update (see Section IV-B) that can be viewed as a cheaper heuristic alternative to the stated Armijo-rule which still ensures monotonic objective descent (and hence is “suitable” in the context of Proposition 1), also enables strictly positive steps, and empirically performs equally well. Finally, we remark that the condition $\mu \geq 1$ in Theorem 2 can be dropped if the relevant objective parts of problem (4) are not rescaled for the $\mathbf{A}$ - and $\mathbf{D}$ -updates, respectively.

To conclude the discussion of convergence, we point out that one can obtain a linear rate of convergence for DOLPHIn with the Armijo rule from Theorem 2, by extending the results of [34, Theorem 2 (cf. Section 8)].

IV. NUMERICAL RESULTS

In this section, we discuss various numerical experiments to study the effectiveness of the DOLPHIn algorithm. To that end, we consider several types of linear operators $\mathcal{F}$ within our model (4) (namely, different types of Gaussian random operators and coded diffraction models). Details on the experimental setup and our implementation are given in the first two subsections, before presenting the main numerical results in Subsection IV-C. Our experiments demonstrate that with noisy measurements, DOLPHIn gives significantly better image reconstructions than the Wirtinger Flow method [12], one recent state-of-the-art phase retrieval algorithm, thereby showing that introducing sparsity via a (simultaneously) learned dictionary is indeed a promising new approach for signal reconstruction from noisy, phaseless, nonlinear measurements. Furthermore, we discuss sensitivity of DOLPHIn with regard to various (hyper-)parameter choices (Subsections. IV-D, IV-E and IV-F) and a variant in which the ℓ_1 -regularization term in the objective is replaced by explicit constraints on the sparsity of the patch-representation coefficient vectors $\mathbf{a}^i$ (Subsection. IV-G).

A. Experimental Setup

We consider several linear operators $\mathcal{F}$ corresponding to different types of measurements that are classical in the phase

retrieval literature. We denote by $\mathbf{F}$ the (normalized) 2D-Fourier operator (implemented using fast Fourier transforms), and introduce two complex Gaussian matrices $\mathbf{G} \in \mathbb{C}^{M_1 \times N_1}$, $\mathbf{H} \in \mathbb{C}^{M_2 \times N_2}$, whose entries are i.i.d. samples from the distribution $\mathcal{N}(0, \mathbf{I}/2) + i\mathcal{N}(0, \mathbf{I}/2)$. Then, we experiment with the operators $\mathcal{F}(\mathbf{X}) = \mathbf{G}\mathbf{X}$, $\mathcal{F}(\mathbf{X}) = \mathbf{G}\mathbf{X}\mathbf{G}^*$, $\mathcal{F}(\mathbf{X}) = \mathbf{G}\mathbf{X}\mathbf{H}^*$, and the coded diffraction pattern model

$$\mathcal{F}(\mathbf{X}) = \begin{pmatrix} \mathbf{F}(\overline{\mathbf{M}}_1 \odot \mathbf{X}) \\ \vdots \\ \mathbf{F}(\overline{\mathbf{M}}_m \odot \mathbf{X}) \end{pmatrix}, \quad \mathcal{F}^*(\mathbf{Z}) = \sum_{j=1}^m (\mathbf{M}_j \odot \mathbf{F}^*(\mathbf{Z}_j)), \quad (9)$$

where $\mathbf{Z}_j := \mathbf{Z}_{\{(j-1)N_1+1, \dots, jN_1\}}$, (i.e., $\mathbf{Z}^\top = (\mathbf{Z}_1^\top, \dots, \mathbf{Z}_m^\top)$) and the $\mathbf{M}_j$'s are admissible coded diffraction patterns (CDPs), see for instance [12, Section 4.1]; in our experiments we used ternary CDPs, such that each $\mathbf{M}_j$ is in $\{0, \pm 1\}^{N_1 \times N_2}$. (Later, we will also consider octanary CDPs with $\mathbf{M}_j \in \{\pm\sqrt{2}/2, \pm i\sqrt{2}/2, \pm\sqrt{3}, \pm i\sqrt{3}\} \in \mathbb{C}^{N_1 \times N_2}$.)

To reconstruct $\hat{\mathbf{X}}$, we choose an oversampling setting where $M_1 = 4N_1$, $M_2 = 4N_2$ and/or $m = 2$, respectively. Moreover, we corrupt our measurements with additive white Gaussian noise $\mathbf{N}$ such that $\text{SNR}(\mathbf{Y}, |\mathcal{F}(\hat{\mathbf{X}})|^2 + \mathbf{N}) = 10$ dB for the Gaussian-type, and 20 dB for CDP measurements, respectively. Note that these settings yield, in particular, a relatively heavy noise level for the Gaussian cases and a relatively low oversampling ratio for the CDPs.

B. Implementation Details

We choose to initialize our algorithm with a simple random image $\mathbf{X}_{(0)}$ in $\mathcal{X}$ to demonstrate the robustness of our approach with respect to its initialization. Nevertheless, other choices are possible. For instance, one may also initialize $\mathbf{X}_{(0)}$ with a power-method scheme similar to that proposed in [12], modified to account for the real-valuedness and box-constraints. The dictionary is initialized as $\mathbf{D}_{(0)} = (\mathbf{I}, \mathbf{F}_D)$ in $\mathbb{R}^{s \times 2s}$, where $\mathbf{F}_D$ corresponds to the two-dimensional discrete cosine transform (see, e.g., [20]).

To update $\mathbf{A}$, we use the ISTA implementation from the SPAMS package⁴ [23] with its integrated backtracking line search (for L_A). Regarding the step sizes γ_ℓ^X for the update of $\mathbf{X}$ (Step 3 of Algorithm 1), we adopt the following simple strategy, which is similar to that from [24] and may be viewed as a heuristic to the Armijo rule from Theorem 2: Whenever the gradient step leads to a reduction in the objective function value, we accept it. Otherwise, we recompute the step with γ_ℓ^X halved until a reduction is achieved; here, as a safeguard against numerical issues, we implemented a limit of 100 trials (forcing termination in case all were unsuccessful), but this was never reached in any of our computational experiments. Regardless of whether γ_ℓ^X was reduced or not, we reset its value to $1.68\gamma_\ell^X$ for the next round; the initial step size is $\gamma_0^X = 10^4/f_{(0)}$, where $f_{(0)}$ is the objective function of the DOLPHIn model (4), evaluated at $\mathbf{X}_{(0)}$, $\mathbf{D}_{(0)}$ and least-squares patch representations $\arg \min_{\mathbf{A}} \|\mathcal{E}(\mathbf{X}_{(0)}) - \mathbf{D}_{(0)}\mathbf{A}\|_{\mathbf{F}}^2$. (Note that, while this rule deviates from the theoretical convergence

Theorem 2, Proposition 1 and the remarks following it remain applicable.)

Finally, we consider nonoverlapping 8×8 patches and run DOLPHIn (Algorithm 1) with $K_1 = 25$ and $K_2 = 50$; the regularization/penalty parameter values can be read from Table I (there, m_Y is the number of elements of $\mathbf{Y}$). We remark that these parameter values were empirically benchmarked to work well for the measurement setups and instances considered here; a discussion about the stability of our approach with respect to these parameter choices is presented below in Section IV-D. Further experiments with a sparsity-constrained DOLPHIn variant and using overlapping patches are discussed in Section IV-G.

Our DOLPHIn code is available online on the first author's webpage⁵.

C. Computational Experiments

We test our method on a collection of typical (grayscale) test images used in the literature, see Figure 1. All experiments were carried out on a Linux 64-bit quad-core machine (2.8 GHz, 8 GB RAM) running Matlab R2016a (single-thread).

We evaluate our approach with the following question in mind: *Can we improve upon the quality of reconstruction compared to standard phase retrieval algorithms?* Standard methods cannot exploit sparsity if the underlying basis or dictionary is unknown; as we will see, the introduced (patch-) sparsity indeed allows for better recovery results (at least in the oversampling and noise regimes considered here).

To evaluate the achievable sparsity, we look at the average number of nonzeros in the columns of $\mathbf{A}$ after running our algorithm. Generally, smaller values indicate an improved suitability of the learned dictionary for sparse patch coding (high values often occur if the regularization parameter λ is too small and the dictionary is learning the noise, which is something we would like to avoid). To assess the quality of the image reconstructions, we consider two standard measures, namely the peak signal-to-noise ratio (PSNR) of a reconstruction as well as its structural similarity index (SSIM) [36]. For PSNR, larger values are better, and SSIM-values closer to 1 (always ranging between 0 and 1) indicate better visual quality.

Table I displays the CPU times, PSNR- and SSIM-values and mean patch representation vector sparsity levels obtained for the various measurement types, averaged over the instance groups of the same size. The concrete examples in Figures 2 and 3 show the results from DOLPHIn and plain Wirtinger Flow (WF; the real-valued, $[0, 1]$ -box constrained variant, which corresponds to running Algorithm 1 with $\mu = 0$ and omitting the updates of $\mathbf{A}$ and $\mathbf{D}$). In all tests, we let the Wirtinger Flow method run for the same number of iterations (75) and use the same starting points as for the DOLPHIn runs. Note that instead of random $\mathbf{X}_{(0)}$, we could also use a spectral initialization similar to the one proposed for the (unconstrained) Wirtinger Flow algorithm, see [12]. Such initialization can improve WF reconstruction (at least in the noiseless case), and may also provide better initial estimates for DOLPHIn. We have experimented with such a

⁴<http://spams-devel.gforge.inria.fr/>

⁵<http://www.mathematik.tu-darmstadt.de/~tillmann/>

Fig. 1. Test images. (a)–(c): cameraman, house and peppers (size 256×256). (d)–(h): lena, barbara, boat, fingerprint and mandrill (size 512×512).

TABLE I

TEST RESULTS FOR m_Y GAUSSIAN-TYPE AND CODED DIFFRACTION PATTERN (CDP) MEASUREMENTS. WE REPORT MEAN VALUES (GEOMETRIC MEAN FOR CPU TIMES) PER MEASUREMENT TYPE, OBTAINED FROM THREE INSTANCES WITH RANDOM $\mathbf{X}_{(0)}$ AND RANDOM NOISE FOR EACH OF THE THREE 256×256 AND FIVE 512×512 IMAGES, W.R.T. THE RECONSTRUCTIONS FROM DOLPHIN ($\mathbf{X}_{\text{DOLPHIN}}$ AND $\mathcal{P}_{\mathcal{X}}(\mathcal{R}(\mathbf{DA}))$) WITH PARAMETERS (μ, λ) AND (REAL-VALUED, $[0, 1]$ -CONSTRAINED) WIRTINGER FLOW ($\mathbf{X}_{\text{WF}}$), RESPECTIVELY. (CPU TIMES IN SECONDS, PSNR IN DECIBELS).

$\mathcal{F}$ type	reconstruction	256 $\times$ 256 instances					512 $\times$ 512 instances				
		$(\mu, \lambda)/m_Y$	time	PSNR	SSIM	$\varnothing \ \mathbf{a}^i\ _0$	$(\mu, \lambda)/m_Y$	time	PSNR	SSIM	$\varnothing \ \mathbf{a}^i\ _0$
$\mathbf{G}\hat{\mathbf{X}}$	$\mathbf{X}_{\text{DOLPHIN}}$	(0.5,0.105)	12.69	24.69	0.5747	3.77	(0.5,0.105)	68.62	24.42	0.6547	6.30
	$\mathcal{P}_{\mathcal{X}}(\mathcal{R}(\mathbf{DA}))$			23.08	0.6644				22.66	0.6807	
	$\mathbf{X}_{\text{WF}}$		7.49	19.00	0.2898			49.23	18.83	0.3777	
$\mathbf{G}\hat{\mathbf{X}}\mathbf{G}^*$	$\mathbf{X}_{\text{DOLPHIN}}$	(0.5,0.210)	51.78	22.67	0.4135	7.45	(0.5,0.210)	357.49	22.59	0.5296	11.37
	$\mathcal{P}_{\mathcal{X}}(\mathcal{R}(\mathbf{DA}))$			23.70	0.7309				23.43	0.7685	
	$\mathbf{X}_{\text{WF}}$		47.76	22.66	0.4131			349.28	22.58	0.5290	
$\mathbf{G}\hat{\mathbf{X}}\mathbf{H}^*$	$\mathbf{X}_{\text{DOLPHIN}}$	(0.5,0.210)	52.18	22.67	0.4132	7.50	(0.5,0.210)	357.66	22.57	0.5286	11.38
	$\mathcal{P}_{\mathcal{X}}(\mathcal{R}(\mathbf{DA}))$			23.68	0.7315				23.43	0.7667	
	$\mathbf{X}_{\text{WF}}$		48.24	22.65	0.4127			348.54	22.55	0.5282	
CDP (cf. (9))	$\mathbf{X}_{\text{DOLPHIN}}$	(0.05,0.003)	8.56	27.15	0.7416	7.85	(0.05,0.003)	36.72	27.33	0.7819	11.48
	$\mathcal{P}_{\mathcal{X}}(\mathcal{R}(\mathbf{DA}))$			26.58	0.7654				26.33	0.7664	
	$\mathbf{X}_{\text{WF}}$		2.83	13.10	0.1170			14.79	12.70	0.1447	

spectral approach and found the results comparable to what is achievable with random $\mathbf{X}_{(0)}$, both for WF and DOLPHIN. Therefore, we do not report these experiments in the paper.

The DOLPHIN method consistently provides better image reconstructions than WF, which clearly shows that our approach successfully introduces sparsity into the phase retrieval problem and exploits it for estimating the solution. As can be seen from Table I, the obtained dictionaries allow for rather sparse representation vectors, with the effect of making better use of the information provided by the measurements, and also denoising the image along the way. The latter fact can be seen in the examples (Fig. 2 and 3, see also Fig. 4) and also inferred from the significantly higher PSNR and SSIM values

for the estimates $\mathbf{X}_{\text{DOLPHIN}}$ and $\mathcal{R}(\mathbf{DA})$ (or $\mathcal{P}_{\mathcal{X}}(\mathcal{R}(\mathbf{DA}))$, resp.) obtained from DOLPHIN compared to the reconstruction $\mathbf{X}_{\text{WF}}$ of the WF algorithm (which cannot make use of hidden sparsity). The gain in reconstruction quality is more visible in the example of Fig. 3 (cf. Fig. 4) than for that in Fig. 2, though both cases assert higher quantitative measures. Furthermore, note that DOLPHIN naturally has higher running times than WF, since it performs more work per iteration (also, different types of measurement operators require different amounts of time to evaluate). Note also that storing $\mathbf{A}$ and $\mathbf{D}$ instead of an actual image $\mathbf{X}$ (such as the WF reconstruction) requires saving only about half as many numbers (including integer index pairs for the nonzero entries in $\mathbf{A}$).

Fig. 2. DOLPHIn example: Image original is the 512×512 “fingerprint” picture, measurements are noisy Gaussian $\mathbf{G}\hat{\mathbf{X}}$ ($M_1 = 4N_1$, noise-SNR 10 dB), $(\mu, \lambda) = (0.5, 0.105)m_Y$. (a) final dictionary (excerpt, magnified), (b) image reconstruction $\mathbf{X}_{\text{DOLPHIn}}$, (c) image reconstruction $\mathcal{R}(\mathbf{DA})$ from sparsely coded patches, (d) reconstruction $\mathbf{X}_{\text{WF}}$ after 75 WF iterations. Final PSNR values: 22.37 dB for $\mathcal{R}(\mathbf{DA})$, 23.74 dB for $\mathbf{X}_{\text{DOLPHIn}}$, 18.19 dB for $\mathbf{X}_{\text{WF}}$; final SSIM values: 0.7903 for $\mathcal{R}(\mathbf{DA})$, 0.8152 for $\mathbf{X}_{\text{DOLPHIn}}$, 0.5924 for $\mathbf{X}_{\text{WF}}$; average $\|\mathbf{a}^i\|_0$ is 10.06.

Fig. 3. DOLPHIn example: Image original is a 2816×2112 photo of the “Waldspirale” building in Darmstadt, Germany; measurements are noisy CDPs (obtained using two ternary masks), noise-SNR 20 dB, $(\mu, \lambda) = (0.05, 0.007)m_Y$. (a) image reconstruction $\mathbf{X}_{\text{DOLPHIn}}$, (b) image reconstruction $\mathcal{R}(\mathbf{DA})$ from sparsely coded patches, (c) reconstruction $\mathbf{X}_{\text{WF}}$ after 75 WF iterations. Final PSNR values: 23.40 dB for $\mathcal{R}(\mathbf{DA})$, 24.72 dB for $\mathbf{X}_{\text{DOLPHIn}}$, 12.63 dB for $\mathbf{X}_{\text{WF}}$; final SSIM values: 0.6675 for $\mathcal{R}(\mathbf{DA})$, 0.6071 for $\mathbf{X}_{\text{DOLPHIn}}$, 0.0986 for $\mathbf{X}_{\text{WF}}$; average $\|\mathbf{a}^i\|_0$ is 12.82. (Total reconstruction time roughly 30 min (DOLPHIn) and 20 min (WF), resp.) Original image taken from Wikimedia Commons, under Creative Commons Attribution-Share Alike 3.0 Unported license.

Fig. 4. DOLPHIn example “Waldspirale” image, zoomed-in 100×100 pixel parts (magnified). (a) original image, (b) image reconstruction $\mathbf{X}_{\text{DOLPHIn}}$, (b) reconstruction $\mathcal{R}(\mathbf{DA})$ from sparsely coded patches, (c) reconstruction $\mathbf{X}_{\text{WF}}$ after 75 WF iterations. The slight block artefacts visible in (b) and (c) are due to the nonoverlapping patch approach in experiments and could easily be mitigated by introducing some patch overlap (cf., e.g., Fig. 6).

As indicated earlier, the reconstruction $\mathcal{R}(\mathbf{DA})$ is quite often better than $\mathbf{X}_{\text{DOLPHIn}}$ w.r.t. at least one of either PSNR or SSIM value. Nonetheless, $\mathbf{X}_{\text{DOLPHIn}}$ may be visually more appealing than $\mathcal{R}(\mathbf{DA})$ even if the latter exhibits a higher quantitative quality measure (as is the case, for instance, in the example of Figures 3 and 4); Furthermore, occasionally $\mathbf{X}_{\text{DOLPHIn}}$ achieves notably better (quantitative) measures than $\mathcal{R}(\mathbf{DA})$; an intuitive explanation may be that if, while the sparse coding of patches served well to eliminate the noise and—by means of the patch-fit objective term—to successfully “push” the $\mathbf{X}$ -update steps toward a solution of good quality, that solution eventually becomes “so good”, then the fact

that $\mathcal{R}(\mathbf{DA})$ is designed to be only an *approximation* (of $\mathbf{X}$) predominates.

On the other hand, $\mathbf{X}_{\text{DOLPHIn}}$ is sometimes very close to $\mathbf{X}_{\text{WF}}$, which indicates a suboptimal setting of the parameters μ and λ that control how much “feedback” the patch-fitting objective term introduces into the Wirtinger-Flow-like $\mathbf{X}$ -update in the DOLPHIn algorithm. We discuss parameter choices in more detail in the following subsection.

D. Hyperparameter Choices and Sensitivity

The DOLPHIn algorithm requires several parameters to be specified a priori. Most can be referred to as *design*

Fig. 5. Influence of parameter λ on best achievable SSIM values for reconstructed images and sensitivity w.r.t. sampling ratios and noise levels, for different measurement types. Fixed other parameters: $\mu = 0.1m_Y$, $K_1 = 25$, $K_2 = 50$, $s_1 = s_2 = 8$ (nonoverlapping patches), $\mathbf{D}_{(0)} = (\mathbf{I}, \mathbf{F}_D)$, $\mathbf{X}_{(0)} \in \mathcal{X}$ random. (a)–(c): Averages over reconstructions from ternary CDP measurements of the three 256×256 images. The plots show (a) the best achievable SSIM for $\lambda/m_Y \in \{0.0001, 0.0002, \dots, 0.01\}$ (left vertical axis, solid lines) and average patch-sparsity of corresponding solutions (right vertical axis, dashed lines) for noise level $\text{SNR}(\mathbf{Y}, |\mathcal{F}(\mathbf{X})|^2) = 20$ dB and (b) choice of λ yielding best SSIM for different noise levels, for number of masks 2 (black), 5 (blue), 10 (red) or 20 (green), respectively; (c) choice of λ to achieve best SSIM for different number of masks, for noise-SNRs 10 (black), 15 (blue), 20 (red), 30 (yellow), 50 (light blue) and ∞ (green), respectively. (d)–(f): Averages over reconstructions from Gaussian measurements ($\mathbf{Y} = |\mathbf{G}\mathbf{X}|^2$) of the five 512×512 images. The plots display the same kind of information as (a)–(c), but in (d) with $\lambda/m_Y \in \{0.0001, 0.0051, 0.0101, \dots, 0.0951\}$ for noise-SNR 15 dB and in (e) for different noise levels, for sampling ratios $M_1/N_1 = 2$ (black), 4 (blue) and 8 (red), respectively; and in (f) with $M_1/N_1 = 2$ for noise-SNRs 10 (black), 15 (blue), 20 (red), 30 (green) and ∞ (yellow).

parameters; the most prominent ones are the size of image patches ($s_1 \times s_2$), whether patches should overlap or not (not given a name here), and the number n of dictionary atoms to learn. Furthermore, there are certain *algorithmic parameters* (in a broad sense) that need to be fixed, e.g., the iteration limits K_1 and K_2 or the initial dictionary $\mathbf{D}_{(0)}$ and image estimate $\mathbf{X}_{(0)}$. The arguably most important parameters, however, are the model or *regularization parameters* μ and λ . For any fixed combination of design and algorithmic parameters in a certain measurement setup (fixed measurement type/model and (assumed) noise level), it is conceivable that one can find some values for μ and λ that work well for most instances, while the converse—choosing, say, iteration limits for fixed μ , λ and other parameters—is clearly not a very practical approach.

As is common for regularization parameters, a good “general-purpose” way to choose μ and λ a priori is unfortunately not known. To obtain the specific choices used in our experiments, we fixed all the other parameters (including noise SNR and oversampling ratios), then (for each measurement model) ran preliminary tests to identify values μ for which good results could be produced with some λ , and finally fixed μ at such values and ran extensive benchmark tests to find λ

values that give the best results.

For DOLPHIn, μ offers some control over how much “feedback” from the current sparse approximation of the current image estimate is taken into account in the update step to produce the next image iterate—overly large values hinder the progress made by the Wirtinger-Flow-like part of the $\mathbf{X}$ -update, while too small values marginalize the influence of the approximation $\mathcal{R}(\mathbf{D}\mathbf{A})$, with one consequence being that the automatic denoising feature is essentially lost. Nevertheless, in our experiments we found that DOLPHIn is not strongly sensitive to the specific choice of μ once a certain regime has been identified in which one is able to produce meaningful results (for some choice of λ). Hence, λ may be considered the most important parameter; note that this intuitively makes sense, as λ controls how strongly sparsity of the patch representations is actually promoted, the exploitation of which to obtain improved reconstructions being the driving purpose behind our development of DOLPHIn.

Figure 5 illustrates the sensitivity of DOLPHIn with respect to λ , in terms of reconstruction quality and achievable patch-sparsity, for different noise levels, and exemplary measurement types and problem sizes. (In this figure, image quality is

measured by SSIM values alone; the plots using PSNR instead look very similar and were thus omitted. Note, however, that the parameters λ yielding the best SSIM and PSNR values, respectively, need not be the same.) As shown by (a) and (d), there is a clear correlation between the best reconstruction quality that can be achieved (in noisy settings) and the average sparsity of the patch representation vectors $\mathbf{a}^i$. For larger noise, clearly a larger λ is needed to achieve good results—see (b) and (e)—which shows that a stronger promotion of patch-sparsity is an adequate response to increased noise, as is known for linear sparsity-based denoising as well. Similarly, increasing the number of measurements allows to pick a smaller λ whatever the noise level actually is, as can be seen in (c) and (f), respectively. The dependence of the best λ on the noise level appears to follow an exponential curve (w.r.t. the reciprocal SNR) which is “dampened” by the sampling ratio, i.e., becoming less steep and pronounced the more measurements are available, cf. (b) and (e). Indeed, again referring to the subplots (c) and (f), at a fixed noise level the best λ values seem to decrease exponentially with growing number of measurements. It remains subject of future research to investigate these dependencies in more detail, e.g., to come up with more or less general (functional) rules for choosing λ .

E. Impact of Increased Inner Iteration Counts

It is worth considering whether more inner iterations—i.e., consecutive update steps for the different variable blocks—lead to further improvements of the results and/or faster convergence. In general, this is an open question for block-coordinate descent algorithms, so the choices are typically made empirically. Our default choices of $a = 1$ ISTA iterations for the $\mathbf{A}$ -update (Step 2 in Algorithm 1), $x = 1$ projected gradient descent steps for the $\mathbf{X}$ -update (Step 3) and $d = 1$ iterations of the BCD scheme for the $\mathbf{D}$ -update (Steps 4–13) primarily reflect the desire to keep the overall iteration cost low. To assess whether another choice might yield significant improvements, we evaluated the DOLPHIn performance for all combinations of $a \in \{1, 3, 5\}$ and $d \in \{1, 3, 5\}$, keeping all other parameters equal to the settings from the experiments reported on above. (We also tried these combinations together with an increased iteration count for the $\mathbf{X}$ -update, but already for $x = 2$ or $x = 3$ the results were far worse than with just 1 projected gradient descent step; the reason can likely be found in the fact that without adapting $\mathbf{A}$ and $\mathbf{D}$ to a modified $\mathbf{X}$ -iterate, the patch-fitting term of the objective tries to keep $\mathbf{X}$ close to a then-unsuitable estimate $\mathcal{R}(\mathbf{DA})$ based on older $\mathbf{X}$ -iterates, which apparently has a quite notable negative influence on the achievable progress in the $\mathbf{X}$ -update loop.)

The results are summarized in condensed format in Table II, from which we can read off the spread of the best and worst results (among the best ones achievable with either $\mathbf{X}$ or $\mathcal{P}_{\mathcal{X}}(\mathcal{R}(\mathbf{DA}))$) for each measurement-instance combination among all combinations $(a, d) \in \{1, 3, 5\}^2$. (Full tables for each test run can be found alongside our DOLPHIn code on the first author’s webpage.) As the table shows, the results are all quite close; while some settings lead to sparser patch

representations, the overall quality of the best reconstructions for the various combinations usually differ only slightly, and no particular combination stands out clearly as being better than all others. In particular, comparing the results with those in Table I, we find that our default settings provide consistently good results; they may be improved upon with some other combination of a and d , but at least with the same total iteration horizon ($K_1 + K_2 = 75$), the improvements are often only marginal. Since the overall runtime reductions (if any) obtained with other choices for (a, d) are also very small, there does not seem to be a clear advantage to using more than a single iteration for either update problem.

F. Influence of the First DOLPHIn Phase

Our algorithm keeps the dictionary fixed at its initialization for the first K_1 iterations in order to prevent the dictionary from training on relatively useless first reconstruction iterates. Indeed, if all variables *including* $\mathbf{D}$ are updated right from the beginning (i.e., $K_1 = 0$, $K_2 = 75$), then we end up with inferior results (keeping all other parameters unchanged): The obtained patch representations are much less sparse, the quality of the image estimate $\mathcal{R}(\mathbf{DA})$ decreases drastically, and also the quality of the reconstruction $\mathbf{X}$ becomes notably worse. This demonstrates that the dictionary apparently “learns too much noise” when updated from the beginning, and the positive effect of filtering out quite a lot of noise in the first phase by regularizing with sparsely coded patches using a fixed initial dictionary is almost completely lost. To give an example, for the CDP testset on 256×256 images, the average values obtained by DOLPHIn when updating also the dictionary from the first iteration onward are: 9.24 seconds runtime (versus 8.56 for default DOLPHIn, cf. Table I), mean patch sparsity $\mathbb{E}\|\mathbf{a}^i\|_0 \approx 20.09$ (vs. 7.85), PSNR 26.80 dB and 8.88 dB (vs. 27.15 and 26.58) and SSIM-values 0.6931 and 0.0098 (vs. 0.7416 and 0.7654) for the reconstructions $\mathbf{X}$ and $\mathcal{P}_{\mathcal{X}}(\mathcal{R}(\mathbf{DA}))$, respectively.

On the other hand, one might argue that if the first iterates are relatively worthless, the effort of updating $\mathbf{A}$ in the first DOLPHIn phase (i.e., the first K_1 iterations) could be saved as well. However, the influence of the objective terms involving $\mathbf{D}$ and $\mathbf{A}$ should then be completely removed from the algorithm for the first K_1 iterations; otherwise, the patch-fitting term will certainly hinder progress made by the $\mathbf{X}$ -update because it then amounts to trying to keep $\mathbf{X}$ close to the initial estimate $\mathcal{R}(\mathbf{DA})$, which obviously needs not bear any resemblance to the sought solution at all. Thus, if both $\mathbf{D}$ and $\mathbf{A}$ are to be unused in the first phase, then one should temporarily set $\mu = 0$, with the consequence that the first phase reduces to pure projected gradient descent for $\mathbf{X}$ with respect to the phase retrieval objective—i.e., essentially, Wirtinger Flow. Therefore, proceeding like this simply amounts to a different initialization of $\mathbf{X}$. Experiments with this DOLPHIn variant ($K_1 = 25$ initial WF iterations followed by $K_2 = 50$ full DOLPHIn iterations including $\mathbf{A}$ and $\mathbf{D}$; all other parameters again left unchanged) showed that the achievable patch-sparsities remain about the same for the Gaussian measurement types but become much worse for

TABLE II

TEST RESULTS FOR DOLPHIN VARIANTS WITH DIFFERENT COMBINATIONS OF INNER ITERATION NUMBERS a AND d FOR THE $\mathbf{A}$ - AND $\mathbf{D}$ -UPDATES, RESP. REPORTED ARE THE BEST MEAN VALUES ACHIEVABLE (VIA EITHER $\mathbf{X}_{\text{DOLPHIN}}$ OR $\mathcal{P}_{\mathcal{X}}(\mathcal{R}(\mathbf{DA}))$) WITH ANY OF THE CONSIDERED COMBINATIONS $(a, d) \in \{1, 3, 5\} \times \{1, 3, 5\}$ (FIRST ROWS FOR EACH MEASUREMENT TYPE) AND THE WORST AMONG THE BEST VALUES FOR EACH COMBINATION (SECOND ROWS), ALONG WITH THE RESPECTIVE COMBINATIONS YIELDING THE STATED VALUES. ALL OTHER PARAMETERS ARE IDENTICAL TO THOSE USED IN THE EXPERIMENTS FOR DEFAULT DOLPHIN ($a = d = 1$), CF. TABLE I.

$\mathcal{F}$ type	256 $\times$ 256 instances				512 $\times$ 512 instances			
	time	PSNR	SSIM	$\emptyset \ \mathbf{a}^i\ _0$	time	PSNR	SSIM	$\emptyset \ \mathbf{a}^i\ _0$
$\mathbf{GX}$	12.61 (1, 3) 15.30 (5, 5)	24.72 (1, 5) 24.48 (3, 3)	0.6680 (1, 5) 0.6485 (3, 3)	3.69 (1, 5) 4.94 (3, 1)	68.18 (1, 3) 79.28 (5, 5)	24.43 (1, 3) 24.30 (5, 3)	0.6903 (3, 3) 0.6803 (1, 5)	6.25 (1, 5) 8.55 (3, 5)
$\mathbf{GXG}^*$	51.02 (1, 3) 54.98 (3, 1)	23.71 (1, 5) 23.57 (5, 5)	0.7330 (1, 5) 0.7188 (1, 3)	6.71 (5, 5) 7.55 (1, 3)	357.49 (1, 1) 371.20 (5, 3)	23.44 (1, 5) 23.39 (5, 1)	0.7685 (1, 1) 0.7633 (1, 3)	9.54 (5, 5) 11.61 (1, 3)
$\mathbf{GXH}^*$	50.95 (1, 3) 56.05 (3, 1)	23.68 (1, 1) 23.56 (5, 3)	0.7315 (1, 1) 0.7143 (5, 3)	6.64 (5, {1, 5}) 7.59 (1, 3)	357.66 (1, 1) 373.43 (5, 5)	23.44 (1, {3, 5}) 23.40 (5, {3, 5})	0.7693 (1, 3) 0.7650 (5, 3)	9.46 (5, 5) 11.57 (1, 5)
CDP (cf. (9))	8.56 (1, 1) 11.40 (5, 5)	27.19 (3, 1) 27.04 (3, 5)	0.7692 (3, 1) 0.7654 (1, 1)	7.71 (1, 3) 9.49 (3, 1)	35.66 (1, 3) 47.07 (5, 3)	27.38 (5, 1) 27.33 (1, {1, 3})	0.7837 (3, 3) 0.7818 (1, 5)	11.40 (1, 3) 13.43 (3, 1)

the CDP setup, and that the average PSNR and SSIM values become (often clearly) worse in virtually all cases. In the example of measurements $\mathbf{GX}$ of the 512 $\times$ 512 test images, the above-described DOLPHIn variant runs 62.34 seconds on average (as less effort is spent in the first phase, this is naturally lower than the 68.62 seconds default DOLPHIn takes), produces slightly lower average patch-sparsity (5.88 vs. 6.30 for default DOLPHIn), but for both $\mathbf{X}$ and $\mathcal{P}_{\mathcal{X}}(\mathcal{R}(\mathbf{DA}))$, the PSNR and SSIM values are notably worse (22.32 dB and 20.55 dB vs. 24.42 dB and 22.66 dB, and 0.6165 and 0.5921 vs. 0.6547 and 0.6807, resp.). The reason for the observed behavior can be found in the inferior performance of WF without exploiting patch-sparsity (cf. also Table I); note that the results also further demonstrate DOLPHIn's robustness w.r.t. the initial point—apparently, the initial point obtained from some WF iterations is not more helpful for DOLPHIn than a random first guess.

Finally, it is also worth considering what happens if the dictionary updates are turned off completely, i.e., $K_2 = 0$. Then, DOLPHIn reduces to patch-sparsity regularized Wirtinger Flow, a WF variant that apparently was not considered previously. Additional experiments with $K_1 = 75$, $K_2 = 0$ and $\mathbf{D} = \mathbf{D}_{(0)} = (\mathbf{I}, \mathbf{F}_D)$ (other parameters left unchanged) showed that this variant consistently produces higher sparsity (i.e., smaller average number of nonzero entries) of the patch representation coefficient vectors, but that the best reconstruction ($\mathbf{X}$ or $\mathcal{P}_{\mathcal{X}}(\mathcal{R}(\mathbf{DA}))$) is always significantly inferior to the best one produced by our default version of DOLPHIn. The first observation is explained by the fact that with $\mathbf{D}$ fixed throughout, the patch coding ($\mathbf{A}$ -update) only needs to adapt w.r.t. new $\mathbf{X}$ -iterates but not a new $\mathbf{D}$; at least if the new $\mathbf{X}$ is not too different from the previous one, the former representation coefficient vectors still yield an acceptable approximation, which no longer holds true if the dictionary was also modified. While it should also be mentioned that the patch-sparsity regularized version performs much better than plain WF already, the second observation clearly indicates the additional benefit of working with trained dictionaries, i.e., superiority of DOLPHIn also over the regularized WF variant.

Full tables containing results for all testruns reported on

in this subsection are again available online along with our DOLPHIn code on the first author's website.

G. Sparsity-Constrained DOLPHIn Variant

From Table I and Figure 5, (a) and (d), it becomes apparent that a sparsity level of 8 ± 4 accompanies the good reconstructions by DOLPHIn. This suggests use in a DOLPHIn variant we already briefly hinted at: Instead of using ℓ_1 -norm regularization, we could incorporate explicit sparsity constraints on the $\mathbf{a}^i$. The corresponding DOLPHIn model then reads

$$\begin{aligned} \min_{\mathbf{X}, \mathbf{D}, \mathbf{A}} \quad & \frac{1}{4} \|\mathbf{Y} - |\mathcal{F}(\mathbf{X})|^2\|_{\mathbf{F}}^2 + \frac{\mu}{2} \|\mathcal{E}(\mathbf{X}) - \mathbf{DA}\|_{\mathbf{F}}^2 \\ \text{s.t.} \quad & \mathbf{X} \in [0, 1]^{N_1 \times N_2}, \quad \mathbf{D} \in \mathcal{D}, \quad \|\mathbf{a}^i\|_0 \leq k \quad \forall i = 1, \dots, p, \end{aligned} \quad (10)$$

where k is the target sparsity level. We no longer need to tune the λ parameter, and can let our previous experimental results guide the choice of k . Note that the only modification to Algorithm 1 concerns the update of $\mathbf{A}$ (Step 2), which now requires solving or approximating

$$\min \frac{1}{2} \|\mathcal{E}(\mathbf{X}) - \mathbf{DA}\|_{\mathbf{F}}^2 \quad \text{s.t.} \quad \|\mathbf{a}^i\|_0 \leq k \quad \forall i = 1, \dots, p.$$

In our implementation, we do so by running Orthogonal Matching Pursuit (OMP) [37] for each column $\mathbf{a}^i$ of $\mathbf{A}$ separately until either the sparsity bound is reached or $\|\mathbf{x}^i - \mathbf{D}\mathbf{a}^i\|_2 \leq \varepsilon$, where we set a default of $\varepsilon := 0.1$. (The value of ε is again a parameter that might need thorough benchmarking; obviously, it is related to the amount of noise one wants to filter out by sparsely representing the patches—higher noise levels will require larger ε values. The 0.1 default worked quite well in our experiments, but could probably be improved by benchmarking as well.) The OMP code we used is also part of the SPAMS package.

The effect of the parameter μ is more pronounced in the sparsity-constrained DOLPHIn than in Algorithm 1; however, it appears its choice is less dependent on the measurement model used. By just a few experimental runs, we found that good results in all our test cases can be achieved using $K_1 = K_2 = 25$ iterations, where in the first K_1 (with fixed

TABLE III

TEST RESULTS FOR SPARSITY-CONSTRAINED DOLPHIN, USING OVERLAPPING PATCHES, FOR m_Y GAUSSIAN-TYPE AND CODED DIFFRACTION PATTERN (CDP) MEASUREMENTS. REPORTED ARE MEAN VALUES (GEOMETRIC MEAN FOR CPU TIMES) PER MEASUREMENT TYPE, OBTAINED FROM THREE INSTANCES WITH RANDOM $\mathbf{X}_{(0)}$ AND RANDOM NOISE FOR EACH OF THE THREE 256×256 AND FIVE 512×512 IMAGES, W.R.T. THE RECONSTRUCTIONS FROM DOLPHIN ($\mathbf{X}_{\text{DOLPHIN}}$ AND $\mathcal{P}_{\mathcal{X}}(\mathcal{R}(\mathbf{DA}))$) WITH PARAMETERS (μ_1, μ_2) AND (REAL-VALUED, $[0, 1]$ -CONSTRAINED) WIRTINGER FLOW ($\mathbf{X}_{\text{WF}}$), RESPECTIVELY. (CPU TIMES IN SECONDS, PSNR IN DECIBELS).

$\mathcal{F}$ type	reconstruction	256 $\times$ 256 instances						512 $\times$ 512 instances					
		$(\mu_1, \mu_2)/m_Y$	time	PSNR	SSIM	$\emptyset \ \mathbf{a}^i\ _0$		$(\mu_1, \mu_2)/m_Y$	time	PSNR	SSIM	$\emptyset \ \mathbf{a}^i\ _0$	
$\mathbf{GX}$	$\mathbf{X}_{\text{DOLPHIN}}$	(0.005,0.0084)	23.58	26.79	0.6245			(0.005,0.0084)	137.15	26.73	0.7090		
	$\mathcal{P}_{\mathcal{X}}(\mathcal{R}(\mathbf{DA}))$			27.98	0.7721	2.71				27.60	0.7786	3.28	
	$\mathbf{X}_{\text{WF}}$		5.04	19.00	0.2889	–			32.63	18.94	0.3811	–	
$\mathbf{GXG}^*$	$\mathbf{X}_{\text{DOLPHIN}}$	(0.005,0.0084)	57.52	22.71	0.4143			(0.005,0.0084)	358.29	22.49	0.5234		
	$\mathcal{P}_{\mathcal{X}}(\mathcal{R}(\mathbf{DA}))$			27.29	0.6129	7.96				27.47	0.7202	8.00	
	$\mathbf{X}_{\text{WF}}$		32.44	22.71	0.4145	–			232.26	22.49	0.5239	–	
$\mathbf{GXH}^*$	$\mathbf{X}_{\text{DOLPHIN}}$	(0.005,0.0084)	57.67	22.54	0.4082			(0.005,0.0084)	356.27	22.56	0.5272		
	$\mathcal{P}_{\mathcal{X}}(\mathcal{R}(\mathbf{DA}))$			27.14	0.6059	7.97				27.55	0.7233	8.00	
	$\mathbf{X}_{\text{WF}}$		32.32	22.56	0.4088	–			232.62	22.56	0.5276	–	
CDP (cf. (9))	$\mathbf{X}_{\text{DOLPHIN}}$	(0.005,0.0084)	22.10	28.00	0.8041			(0.005,0.0084)	112.33	26.30	0.7400		
	$\mathcal{P}_{\mathcal{X}}(\mathcal{R}(\mathbf{DA}))$			26.87	0.7789	1.52				24.95	0.6557	1.51	
	$\mathbf{X}_{\text{WF}}$		2.70	21.95	0.3880	–			13.68	21.73	0.4935	–	

dictionary), we use a value of $\mu = \mu_1 = 0.005m_Y$ along with a sparsity bound $k = k_1 = 4$, and in the final K_2 iterations (in which the dictionary is updated), $\mu = \mu_2 = 1.68\mu_1 = 0.0084m_Y$ and $k = k_2 = 8$. Results on the same instances considered before (using the same $\mathbf{D}_{(0)}$) are presented in Table III, with the exception that for the CDP case, we used 2 complex-valued octanary masks (cf. [12]) here; the initial image estimates and measurement noise were again chosen randomly. Note also that for these test, we used complete sets of *overlapping* patches. This greatly increases p and hence the number of subproblems to be solved in the $\mathbf{A}$ -update step, which is the main reason for the increased running times compared to Table I for the standard DOLPHIN method. (It should however also be mentioned that OMP requires up to k iterations per subproblem, while in Algorithm 1 we explicitly restricted the number of ISTA iterations in the $\mathbf{A}$ -update to just a single one.)

A concrete example is given in Figure 6; here, we consider the color “mandrill” image, for which the reconstruction algorithms (given just two quite heavily noise-corrupted octanary CDP measurements) were run on each of the three RGB channels separately. The sparsity-constrained DOLPHIN reconstructions appear superior to the plain WF solution both visually and in terms of the quality measures PSNR and SSIM.

V. DISCUSSION AND CONCLUSION

In this paper, we introduced a new method, called DOLPHIN, for dictionary learning from noisy nonlinear measurements without phase information. In the context of image reconstruction, the algorithm fuses a variant of the recent Wirtinger Flow method for phase retrieval with a patch-based dictionary learning model to obtain sparse representations of image patches, and yields monotonic objective decrease or (with appropriate step size selection) convergence to a stationary point for the nonconvex combined DOLPHIN model.

Our experiments demonstrate that dictionary learning for phase retrieval with a patch-based sparsity is a promising

direction, especially for cases in which the original Wirtinger Flow approach fails (due to high noise levels and/or relatively low sampling rates).

Several aspects remain open for future research. For instance, regarding the generally difficult task of parameter tuning, additional benchmarking for to-be-identified instance settings of special interest could give further insights on how to choose, e.g., the regularization parameters in relation to varying noise levels.

It may also be worth developing further variants of our algorithm; the successful use of ℓ_0 -constraints instead of the ℓ_1 -penalty, combining OMP with our framework, is just one example. Perhaps most importantly, future research will be directed towards the “classic” phase retrieval problem in which one is given the (squared) magnitudes of Fourier measurements, see, e.g., [2], [1], [5]. Here, the WF method fails, and existing other (projection-based) methods are not always reliable either. The hope is that introducing sparsity via a learned dictionary will also enable improved reconstructions in the Fourier setting.

To evaluate the quality of the learned dictionary, one might also ask how DOLPHIN compares to the straightforward approach to first run (standard) phase retrieval and then learn dictionary and sparse patch representations from the result. Some preliminary experiments (see also those in Section IV-F pertaining to keeping both $\mathbf{A}$ and $\mathbf{D}$ fixed in the first DOLPHIN phase) indicate that both approaches produce comparable results in the noise-free setting, while our numerical results demonstrate a denoising feature of our algorithm that the simple approach obviously lacks.

Similarly, it will be of interest to see if the dictionaries learned by DOLPHIN can be used successfully within sparsity-aware methods (e.g., the Thresholded WF proposed in [13], if that were modified to handle local (patch-based) sparsity instead of global priors). In particular, learning dictionaries for patch representations of images from a whole *class* of images would then be an interesting point to consider. To that

Fig. 6. Sparsity-Constrained DOLPHIn example: Image original is the 512×512 RGB “mandrill” picture, measurements are noisy CDPs (obtained using two complex-valued octanary masks, noise-SNR 10 dB, per color channel), $\mu_1 = 0.003 m_Y$, $\mu_2 = 0.00504 m_Y$ (other parameters as described in Section IV-G). $\mathbf{D}_{(0)} = (\mathbf{I}, \mathbf{F}_D)$ for R-channel, final dictionary then served as initial dictionary for G-channel, whose final dictionary in turn was initial dictionary for B-channel; $\mathbf{X}_{(0)} \in \mathcal{X}$ random for each channel. (a) final dictionary (excerpt) for B-channel (b) image reconstruction $\mathbf{X}_{\text{DOLPHIn}}$, (c) image reconstruction $\mathcal{R}(\mathbf{DA})$ from sparsely coded patches, (d) reconstruction $\mathbf{X}_{\text{WF}}$ after 50 WF iterations. Final PSNR values: 20.53 dB for $\mathcal{R}(\mathbf{DA})$, 20.79 dB for $\mathbf{X}_{\text{DOLPHIn}}$, 14.47 dB for $\mathbf{X}_{\text{WF}}$; final SSIM values: 0.4780 for $\mathcal{R}(\mathbf{DA})$, 0.5242 for $\mathbf{X}_{\text{DOLPHIn}}$, 0.2961 for $\mathbf{X}_{\text{WF}}$; average $\|\mathbf{a}^i\|_0$ is 5.30. (Means over all color channels.) (e)–(h): zoomed-in 100×100 pixel parts (magnified) of (e) original image, (f) $\mathbf{X}_{\text{DOLPHIn}}$, (g) $\mathcal{R}(\mathbf{DA})$ and (h) $\mathbf{X}_{\text{WF}}$.

end, note that the DOLPHIn model and algorithm can easily be extended to multiple input images whose patches are all to be represented using a single dictionary by summing up the objectives for each separate image, but with the same $\mathbf{D}$ throughout.

Another interesting aspect to evaluate is by how much reconstruction quality and achievable sparsity degrade due to the loss of phase (or, more generally, measurement nonlinearity), compared to the linear measurement case.

APPENDIX

In the following, we derive the DOLPHIn algorithm for the one-dimensional setting (1). In particular, the (gradient) formulas for the 2D-case can be obtained by applying the ones given below to the vectorized image $\mathbf{x} = \text{vec}(\mathbf{X})$ (stacking columns of $\mathbf{X}$ on top of each other to form the vector $\mathbf{x}$), the vectorized matrix $\mathbf{a} = \text{vec}(\mathbf{A})$, and interpreting the matrix $\mathbf{F} \in \mathbb{C}^{M \times N}$ as describing the linear transformation corresponding to $\mathcal{F}$ in terms of the vectorized variables.

We now have a patch-extraction matrix $\mathbf{P}_e \in \mathbb{R}^{ps \times N}$ which gives us $\mathbf{P}_e \mathbf{x} = ((\mathbf{x}^1)^\top, \dots, (\mathbf{x}^p)^\top)^\top$ (in the vectorized 2D-case, $\mathbf{x}^i$ then is the vectorized i -th patch of $\mathbf{X}$, i.e., $\mathbf{P}_e$ corresponds to $\mathcal{E}$). Similarly, we have a patch-reassembly matrix $\mathbf{P}_a \in \mathbb{R}^{N \times ps}$; then, the reassembled signal vector will be $\mathbf{P}_a((\mathbf{a}^1)^\top \mathbf{D}^\top, \dots, (\mathbf{a}^p)^\top \mathbf{D}^\top)^\top$ (so $\mathbf{P}_a$ corresponds to $\mathcal{R}$). Note that $\mathbf{P}_a = \mathbf{P}_e^\dagger = (\mathbf{P}_e^\top \mathbf{P}_e)^{-1} \mathbf{P}_e^\top$; in particular, $\mathbf{x} = \mathbf{P}_a \mathbf{P}_e \mathbf{x}$, and $\mathbf{P}_e^\top \mathbf{P}_e$ is a diagonal matrix for which each diagonal entry is associated to a specific vector component and gives the number of patches this component is contained in. (Thus, if $\mathbf{x} = \text{vec}(\mathbf{X})$ is a vectorized 2D-image, $\mathbf{P}_e^\top \mathbf{P}_e \mathbf{x} =$

$\text{vec}(\mathbf{R} \odot \mathbf{X})$ with $\mathbf{R}$ as defined in Section III-B.) Note that for nonoverlapping patches, $\mathbf{P}_e$ is simply a permutation matrix, and $\mathbf{P}_a = \mathbf{P}_e^\top$ (so $\mathbf{P}_a \mathbf{P}_e = \mathbf{I}$). Also, applying just $\mathbf{P}_e^\top$ actually reassembles a signal from patches by simply adding the component’s values *without* averaging.

We wish to represent each patch as $\mathbf{x}^i \approx \mathbf{D} \mathbf{a}^i$ with sparse coefficient vectors $\mathbf{a}^i$; with $\mathbf{a} := ((\mathbf{a}^1)^\top, \dots, (\mathbf{a}^p)^\top)^\top$ and $\hat{\mathbf{D}} := \mathbf{I}_p \otimes \mathbf{D}$, this sparse-approximation relation can be written as $\mathbf{P}_e \mathbf{x} \approx \hat{\mathbf{D}} \mathbf{a}$. Our model to tackle the 1D-problem (1) reads

$$\begin{aligned} \min_{\mathbf{x}, \mathbf{D}, \alpha} \quad & \frac{1}{4} \|\mathbf{y} - |\mathbf{F} \mathbf{x}|^2\|_2^2 + \frac{\mu}{2} \|\mathbf{P}_e \mathbf{x} - \hat{\mathbf{D}} \mathbf{a}\|_2^2 + \lambda \|\mathbf{a}\|_1 \\ \text{s.t.} \quad & \mathbf{x} \in \mathcal{X} := [0, 1]^N, \quad \mathbf{D} \in \mathcal{D}; \end{aligned} \quad (11)$$

here, $\mathbf{y} := |\mathbf{F} \hat{\mathbf{x}}|^2 + \mathbf{n}$, with $\hat{\mathbf{x}}$ the original signal we wish to recover and $\mathbf{n}$ a noise vector.

The update formulas for $\mathbf{a}$ (separately for $\mathbf{a}^1, \dots, \mathbf{a}^p$) and $\mathbf{D}$ remain the same as described before, see Sections III-A and III-C, respectively. However, the update problem for the phase retrieval solution—now derived from (11), with $\mathbf{D}$ and $\mathbf{a}$ fixed at their current values—becomes decreasing the objective

$$\frac{1}{4} \|\mathbf{y} - |\mathbf{F} \mathbf{x}|^2\|_2^2 + \frac{\mu}{2} \|\mathbf{P}_e \mathbf{x} - \hat{\mathbf{D}} \mathbf{a}\|_2^2 \quad \text{with } \mathbf{x} \in \mathcal{X}. \quad (12)$$

We approach this by means of a projected gradient descent step; since we consider real-valued $\mathbf{x}$ -variables, this essentially amounts to (one iteration of) the Wirtinger Flow method [12], accommodating the $[0, 1]$ -box constraints via projection onto them after the (Wirtinger) gradient step. The step size will be chosen to achieve a reduction of the objective w.r.t. its value at the previous $\mathbf{x}$.

The objective function (12) can be rewritten as

$$\begin{aligned}\xi(\mathbf{x}) &:= \varphi(\mathbf{x}) + \psi(\mathbf{x}) \\ &:= \frac{1}{4} \sum_{j=1}^N (y_j - \mathbf{x}^\top \mathbf{F}_j^* \mathbf{F}_j \mathbf{x})^2 + \frac{\mu}{2} \|\mathbf{P}_e \mathbf{x} - \hat{\mathbf{D}} \mathbf{a}\|_2^2.\end{aligned}$$

The gradient of $\psi(\mathbf{x})$ is straightforwardly found to be

$$\nabla \psi(\mathbf{x}) = \mu \mathbf{P}_e^\top (\mathbf{P}_e \mathbf{x} - \hat{\mathbf{D}} \mathbf{a}).$$

Regarding $\nabla \varphi(\mathbf{x}) = \frac{1}{2} \sum_{j=1}^N (y_j - \mathbf{x}^\top \mathbf{M}^j \mathbf{x}) \cdot \nabla (y_j - \mathbf{x}^\top \mathbf{M}^j \mathbf{x})$, where $\mathbf{M}^j := \mathbf{F}_j^* \mathbf{F}_j$. ($\mathbf{F}_j$ is the j -th row of $\mathbf{F}$), we first note that $(\mathbf{M}^j)^* = \mathbf{M}^j$ and hence, in particular, $M_{ik}^j = M_{ki}^j$ (i.e., $\Re(M_{ik}^j) = \Re(M_{ki}^j)$ and $\Im(M_{ik}^j) = -\Im(M_{ki}^j)$). Thus, it is easily seen that for each $i = 1, \dots, N$, the terms in the double-sum $\mathbf{x}^\top \mathbf{M}^j \mathbf{x} = \sum_{\ell=1}^N \sum_{k=1}^N M_{\ell k}^j x_\ell x_k$ that contain x_i are precisely

$$\begin{aligned}M_{ii}^j x_i^2 + x_i \sum_{k=1, k \neq i}^N M_{ik}^j x_k + x_i \sum_{\ell=1, \ell \neq i}^N M_{\ell i}^j x_\ell \\ = M_{ii}^j x_i^2 + \left(2 \sum_{k=1, k \neq i}^N \Re(M_{ik}^j) x_k\right) x_i.\end{aligned}$$

Hence, $\frac{\partial}{\partial x_i} (y_j - \mathbf{x}^\top \mathbf{M}^j \mathbf{x}) = -2 \sum_{k=1}^N \Re(M_{ik}^j) x_k = -2 \Re(\mathbf{M}_i^j) \mathbf{x}$, and therefore $\nabla \varphi(\mathbf{x})$ is equal to

$$\begin{aligned}-\frac{1}{2} \sum_{j=1}^N (y_j - \mathbf{x}^\top \mathbf{M}^j \mathbf{x}) \cdot \left(2 \Re(\mathbf{M}_1^j) \mathbf{x}, \dots, 2 \Re(\mathbf{M}_N^j) \mathbf{x}\right)^\top \\ = -\sum_{j=1}^N (y_j - \mathbf{x}^\top \mathbf{F}_j^* \mathbf{F}_j \mathbf{x}) \Re(\mathbf{F}_j^* \mathbf{F}_j) \mathbf{x}.\end{aligned}$$

Consequently,

$$\begin{aligned}\nabla \xi(\mathbf{x}) &= \nabla \varphi(\mathbf{x}) + \nabla \psi(\mathbf{x}) \\ &= \mu \mathbf{P}_e^\top (\mathbf{P}_e \mathbf{x} - \hat{\mathbf{D}} \mathbf{a}) - \sum_{j=1}^N (y_j - \mathbf{x}^\top \mathbf{F}_j^* \mathbf{F}_j \mathbf{x}) \Re(\mathbf{F}_j^* \mathbf{F}_j) \mathbf{x}.\end{aligned}\tag{13}$$

Thus, with the projection $\mathcal{P}_{\mathcal{X}}(\mathbf{x}) = \max\{0, \min\{1, \mathbf{x}\}\}$ (component-wise) and a step size $\gamma_{(\ell)}^X > 0$, the update of the phase retrieval solution estimate in the ℓ -th DOLPHIn iteration for the 1D-case sets $\mathbf{x}^{(l+1)}$ to the value

$$\begin{aligned}\mathcal{P}_{\mathcal{X}}\left(\mathbf{x}^{(\ell)} - \gamma_{(\ell)}^X \left(\sum_{j=1}^N ((\mathbf{x}^{(\ell)})^\top \mathbf{F}_j^* \mathbf{F}_j \mathbf{x}^{(\ell)} - y_j) \Re(\mathbf{F}_j^* \mathbf{F}_j) \mathbf{x}^{(\ell)} \right. \right. \\ \left. \left. + \mu \mathbf{P}_e^\top (\mathbf{P}_e \mathbf{x}^{(\ell)} - \hat{\mathbf{D}}^{(\ell)} \mathbf{a}^{(\ell)}) \right) \right).\end{aligned}$$

The expression (13) can be further simplified by rewriting $\nabla \varphi(\mathbf{x})$: Since the only complex-valued part within $\nabla \varphi(\mathbf{x})$ is $\mathbf{F}_j^* \mathbf{F}_j$, we can take the real part of the whole expression instead of just this product, i.e.,

$$\nabla \varphi(\mathbf{x}) = -\sum_{j=1}^N (y_j - \mathbf{x}^\top \mathbf{F}_j^* \mathbf{F}_j \mathbf{x}) \Re(\mathbf{F}_j^* \mathbf{F}_j) \mathbf{x}$$

$$= \Re\left(-\sum_{j=1}^N (y_j - \mathbf{x}^\top \mathbf{F}_j^* \mathbf{F}_j \mathbf{x}) \mathbf{F}_j^* \mathbf{F}_j \mathbf{x}\right).$$

Further, using $\mathbf{F}_j^* = (\mathbf{F}_j)^\top = (\mathbf{F}^*)_{\cdot j}$ and rearranging terms, this becomes

$$\begin{aligned}\nabla \varphi(\mathbf{x}) &= \Re\left(\sum_{j=1}^N (\mathbf{F}^*)_{\cdot j} (|\mathbf{F} \mathbf{x}|^2 - y_j) \mathbf{F}_j \mathbf{x}\right) \\ &= \Re\left(\mathbf{F}^* ((|\mathbf{F} \mathbf{x}|^2 - \mathbf{y}) \odot \mathbf{F} \mathbf{x})\right).\end{aligned}$$

From this last form, it is particularly easy to obtain the gradient matrix $\nabla \varphi(\mathbf{X})$ in the 2D-case, which corresponds precisely to the gradient $\nabla \varphi(\mathbf{x})$ with $\mathbf{x}$ the vectorized matrix variable $\mathbf{X}$ and $\mathbf{F}$ representing the linear operator $\mathcal{F}(\mathbf{X})$ (in terms of $\mathbf{x}$), reshaped to match the size of $\mathbf{X}$ (i.e., reversing the vectorization process afterwards). Similarly, the expression for $\nabla \psi(\mathbf{X})$ can be derived from $\nabla \psi(\mathbf{x})$ w.r.t. the vectorized variable; here, the product with $\mathbf{P}_e^\top$ then needs to be replaced by an application of the adjoint $\mathcal{E}^*$, which can be recognized as $\mathcal{E}^*(\cdot) = \mathbf{R} \odot \mathcal{R}(\cdot)$ by translating the effect of multiplying by $\mathbf{P}_e^\top$ to the vectorized variable back to matrix form. (Analogously, one obtains $\mathcal{R}^*(\mathbf{Z}) = \mathcal{E}((1/\mathbf{R}) \odot \mathbf{Z})$, where $1/\mathbf{R}$ has entries $1/r_{ij}$, i.e., is the entry-wise reciprocal of $\mathbf{R}$.)

ACKNOWLEDGEMENTS

We thank the anonymous referees for their comments, which helped to improve the paper.

REFERENCES

- [1] S. Marchesini, "A Unified Evaluation of Iterative Projection Algorithms for Phase Retrieval," *Rev. Sci. Instrum.*, vol. 78, pp. 011301–1–10, 2007.
- [2] Y. Shechtman, Y. C. Eldar, O. Cohen, H. N. Chapman, J. Miao, and M. Segev, "Phase Retrieval with Application to Optical Imaging," *IEEE Signal Process. Mag.*, vol. 32, no. 3, pp. 87–109, 2015.
- [3] H. Sahinoglou and S. Cabrera, "On phase retrieval of finite-length sequences using the initial time sample," *IEEE Trans. Circuits Syst.*, vol. 38, no. 5, pp. 954–958, 1991.
- [4] R. W. Gerchberg and W. O. Saxton, "A practical algorithm for the determination of the phase from image and diffraction plane pictures," *Optik*, vol. 35, no. 2, pp. 237–246, 1972.
- [5] J. R. Fienup, "Phase Retrieval Algorithms: A Comparison," *Appl. Opt.*, vol. 21, no. 15, pp. 2758–2769, 1982.
- [6] H. H. Bauschke, P. L. Combettes, and D. R. Luke, "Phase retrieval, error reduction algorithm, and Fienup variants: A view from convex optimization," *J. Opt. Soc. Amer. A*, vol. 19, no. 7, pp. 1334–1345, 2002.
- [7] E. J. Candès, Y. C. Eldar, T. Strohmer, and V. Voroninski, "Phase retrieval via matrix completion," *SIAM J. on Imaging Sci.*, vol. 6, no. 1, pp. 199–225, 2013.
- [8] Y. Shechtman, Y. C. Eldar, A. Szameit, and M. Segev, "Sparsity based sub-wavelength imaging with partially incoherent light via quadratic compressed sensing," *Opt. Express*, vol. 19, no. 16, pp. 14807–14822, 2011.
- [9] K. Jaganathan, S. Oymak, and B. Hassibi, "Recovery of sparse 1-D signals from the magnitudes of their Fourier transform," arXiv:1206.1405[cs.IT], 2012.
- [10] H. Ohlsson, A. Y. Yang, R. Dong, and S. S. Sastry, "Compressive Phase Retrieval From Squared Output Measurements Via Semidefinite Programming," arXiv:1111.6323[math.ST], 2012.
- [11] I. Waldspurger, A. d'Aspremont, and S. Mallat, "Phase recovery, Max-Cut and complex semidefinite programming," *Math. Prog. A*, vol. 149, no. 1, pp. 47–81, 2015.
- [12] E. J. Candès, X. Li, and M. Soltanolkotabi, "Phase Retrieval via Wirtinger Flow: Theory and Algorithms," arXiv:1407.1065[cs.IT], 2014, to appear in IEEE Trans. Inf. Theory.
- [13] T. T. Cai, X. Li, and Z. Ma, "Optimal Rates of Convergence for Noisy Sparse Phase Retrieval via Thresholded Wirtinger Flow," arXiv:1506.03382[math.ST], 2015.

- [14] Y. Shechtman, A. Beck, and Y. C. Eldar, "GESPAR: Efficient Phase Retrieval of Sparse Signals," *IEEE Trans. Signal Process.*, vol. 62, no. 4, pp. 928–938, 2014.
- [15] M. L. Moravec, J. K. Romberg, and R. G. Baraniuk, "Compressive Phase Retrieval," in *Proc. SPIE 6701: Wavelets XII*, D. V. D. Ville, V. K. Goyal, and M. Papadakis, Eds. International Society for Optical Engineering, 2007, pp. 670 120–1–11.
- [16] S. Foucart and H. Rauhut, *A Mathematical Introduction to Compressive Sensing*, ser. Applied and Numerical Harmonic Analysis. Birkhäuser, 2013.
- [17] Y. C. Eldar and G. Kutyniok, Eds., *Compressed Sensing: Theory and Applications*. Cambridge University Press, 2012.
- [18] Y. C. Eldar, *Sampling Theory – Beyond Bandlimited Systems*. Cambridge University Press, 2015.
- [19] B. A. Olshausen and D. J. Field, "Emergence of Simple-Cell Receptive Field Properties by Learning a Sparse Code for Natural Images," *Nature*, vol. 381, pp. 607–609, 1996.
- [20] M. Elad and M. Aharon, "Image Denoising Via Sparse and Redundant Representations Over Learned Dictionaries," *IEEE Trans. Image Process.*, vol. 15, no. 12, pp. 3736–3745, 2006.
- [21] J. Mairal, F. Bach, and J. Ponce, "Sparse Modeling for Image and Vision Processing," *Found. Trends Comp. Graphics and Vision*, vol. 8, no. 2-3, pp. 85–283, 2014.
- [22] M. Aharon, M. Elad, and A. Bruckstein, "K-SVD: An Algorithm for Designing Overcomplete Dictionaries for Sparse Representation," *IEEE Trans. Signal Process.*, vol. 54, no. 11, pp. 4311–4322, 2006.
- [23] J. Mairal, F. Bach, J. Ponce, and G. Sapiro, "Online learning for matrix factorization and sparse coding," *J. Mach. Learn. Res.*, vol. 11, pp. 19–60, 2010.
- [24] A. Beck and M. Teboulle, "A Fast Iterative Shrinkage-Thresholding Algorithm for Linear Inverse Problems," *SIAM J. Imaging Sci.*, vol. 2, no. 1, pp. 183–202, 2009.
- [25] A. M. Tillmann, Y. C. Eldar, and J. Mairal, "Dictionary Learning from Phaseless Measurements," in *Proc. 41st IEEE Internat. Conf. on Acoustics, Speech, and Signal Processing (ICASSP)*, 2016, pp. 4702–4706.
- [26] H. H. Bauschke, P. L. Combettes, and R. D. Luke, "A Hybrid Projection Reflection Method for Phase Retrieval," *J. Opt. Soc. Am. A*, vol. 20, no. 6, pp. 1025–1034, 2003.
- [27] R. D. Luke, "Relaxed averaged alternating reflections for diffraction imaging," *Inverse Probl.*, vol. 21, pp. 37–50, 2005.
- [28] Y. Chen and E. J. Candès, "Solving Random Quadratic Systems of Equations Is Nearly as Easy as Solving Linear Systems," arXiv:1505.05114 [cs.IT], 2015.
- [29] M. Elad, *Sparse and Redundant Representations: From Theory to Applications in Signal and Image Processing*. Springer, 2010.
- [30] S. S. Chen, D. L. Donoho, and M. A. Saunders, "Atomic Decomposition by Basis Pursuit," *SIAM J. Sci. Comput.*, vol. 20, no. 1, pp. 33–61, 1998.
- [31] A. M. Tillmann, "On the Computational Intractability of Exact and Approximate Dictionary Learning," *IEEE Signal Process. Lett.*, vol. 22, no. 1, pp. 45–49, 2015.
- [32] A. N. Iusem, "On the convergence properties of the projected gradient method for convex optimization," *Comput. Appl. Math.*, vol. 22, no. 1, pp. 37–52, 2003.
- [33] D. P. Bertsekas, *Nonlinear Programming*, 2nd ed. Athena Scientific, 1999.
- [34] P. Tseng and S. Yun, "A coordinate gradient descent method for nonsmooth separable minimization," *Math. Prog. B*, vol. 117, no. 1, pp. 387–423, 2009.
- [35] A. M. Tillmann, Y. C. Eldar, and J. Mairal, "DOLPHIn – Dictionary Learning for Phase Retrieval," arXiv:1602.02263 [math.OC], 2016, preprint, available online.
- [36] 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.*, vol. 13, no. 4, pp. 600–612, 2004.
- [37] Y. C. Pati, R. Rezaifar, and P. S. Krishnaprasad, "Orthogonal Matching Pursuit: Recursive Function Approximation with Applications to Wavelet Decomposition," in *Proc. 27th Annual Asilomar Conf. on Signals, Systems and Computers*, vol. 1. IEEE Computer Society Press, 1993, pp. 40–44.

Andreas M. Tillmann graduated in financial and business mathematics (Dipl.-Math. Oec. degree) from TU Braunschweig, Germany, in 2009, and obtained the Dr. rer. nat. degree from TU Darmstadt, Germany, in 2013.

He was with the Institute for Mathematical Optimization at TU Braunschweig from 03/2009 to 06/2012; in 07/2012, he joined the Research Group Optimization at TU Darmstadt, Germany. From 10/2014 to 03/2015, he was an interim professor for mathematical optimization at TU Braunschweig. His research interests are in discrete and continuous optimization, particularly in sparse recovery and compressed sensing, as well as computational complexity.

Dr. Tillmann won the Best Student Paper Award at the SPARS conference 2013 in Lausanne, Switzerland. He is a member of the Mathematical Optimization Society (MOS), the International Association of Applied Mathematics and Mechanics (GAMM), and an IEEE Signal Processing Society Affiliate. He serves as a reviewer for IEEE Transactions on Information Theory, IEEE Signal Processing Letters and several other journals.

Yonina C. Eldar (S'98–M'02–SM'07–F'12) received the B.Sc. degree in Physics in 1995 and the B.Sc. degree in Electrical Engineering in 1996 both from Tel-Aviv University (TAU), Tel-Aviv, Israel, and the Ph.D. degree in Electrical Engineering and Computer Science in 2002 from the Massachusetts Institute of Technology (MIT), Cambridge, USA.

From January to July, 2002, she was a Postdoctoral Fellow at the Digital Signal Processing Group at MIT. She is currently a Professor in the Department of Electrical Engineering at the Technion –

Israel Institute of Technology, Haifa, Israel, where she holds the Edwards Chair in Engineering. She is also a Research Affiliate with the Research Laboratory of Electronics at MIT and was a Visiting Professor at Stanford University, Stanford, CA. Her research interests are in the broad areas of statistical signal processing, sampling theory and compressed sensing, optimization methods, and their applications to biology and optics.

Dr. Eldar has received many awards for excellence in research and teaching, including the IEEE Signal Processing Society Technical Achievement Award (2013), the IEEE/AESS Fred Nathanson Memorial Radar Award (2014), and the IEEE Kiyo Tomiyasu Award (2016). She also received several best paper awards and best demo awards together with her research students and colleagues, including the SIAM Outstanding Paper Prize and the Signal Processing Society Best Paper Award, and was selected as one of the 50 most influential women in Israel. She is a member of the Young Israel Academy of Science and Humanities and the Israel Committee for Higher Education. Currently, she is the Editor in Chief of Foundations and Trends in Signal Processing, a member of the IEEE Sensor Array and Multichannel Technical Committee and serves on several other IEEE committees. She is author and co-author of several books and many published articles.

Julien Mairal (SM'16) is a research scientist at Inria. He received a graduate degree from Ecole Polytechnique, France, in 2005, and a PhD degree from Ecole Normale Supérieure, Cachan, France, in 2010. Then, he was a postdoctoral researcher at the statistics department of UC Berkeley, before joining Inria in 2012. His research interests include machine learning, computer vision, mathematical optimization, and statistical image and signal processing.

Dr. Mairal received the Cor Baayen prize, awarded every year by ERCIM to a promising young researcher in computer science and applied mathematics, in 2013. In 2016, he received a Starting Grant from the European Research Council (ERC). Since 2015, he is a senior associate editor of IEEE Signal Processing Letters, an associate editor for the International Journal of Computer Vision (IJCV) and the Journal of Mathematical Imaging and Vision. He also served as area chair for major machine learning and computer vision conferences such as ICML and ICCV in 2015, and CVPR, ECCV, ICLR and NIPS in 2016.