

Optically-Coherent Image Formation and Denoising Using Plug and Play Inversion Framework

CASEY J. PELLIZZARI^{1,*}, RUSSELL TRAHAN III², HANYING ZHOU², SKIP WILLIAMS³, STACIE E. WILLIAMS⁴, BIJAN NEMATI², MICHAEL SHAO², AND CHARLES A. BOUMAN¹

¹School of Electrical and Computer Engineering, Purdue University, West Lafayette, IN, 47907

²Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA

³Air Force Research Laboratory, Maui HI

⁴Air Force Office of Scientific Research, Washington DC, 20036

* Corresponding author: cpellizz@purdue.edu

Compiled April 21, 2017

The performance of optically-coherent imaging systems can be limited by measurement and speckle noise. In this paper we develop an image formation framework for computing the MAP estimate of an object's reflectivity when imaged using coherent illumination and detection. The proposed approach allows for the use of Gaussian denoising algorithms (GDAs), without modification, to mitigate the exponentially distributed and signal-dependent noise that occurs in coherent imaging. Several GDAs are compared using both simulated and experimental data. The proposed framework is shown to be robust to noise and significantly reduce reconstruction error compared to the standard inversion technique. © 2017 Optical Society of America

OCIS codes: (100.3010) Image reconstruction techniques; (100.3190) Inverse problems; (110.1650) Coherence imaging.

<http://dx.doi.org/10.1364/ao.XX.XXXXXX>

1. INTRODUCTION

Optically-coherent imaging systems, such as synthetic aperture LADAR (SAL) [1–4] and digital holography (DH) [5–7], offer significant improvements in resolution and sensitivity compared to passive and non-coherent systems. Coherent detection allows apertures to be synthesized which are much larger than the receiver optics. In addition, detection involves measuring the modulation of a strong reference field by a potentially-weak signal field. This allows detection of signals with energies equivalent to a single photon [8].

Despite the benefits, coherent-imaging systems are typically based on simple image-formation techniques and suffer from high noise. For imaging modalities using pupil-plane detection, the data is inverted using Discrete Fourier Transforms (DFTs) which can amplify noise and produce artifacts known as sidelobes. In addition, these techniques produce an estimate of the complex-valued reflection coefficient, g , rather than the real-valued reflectance, r . The reflectance, given by $r = E[|g|^2]$, where $E[\cdot]$ indicates the expected value, is a smoother quantity which we are accustomed to seeing in conventional images and is of greater interest for many imaging applications [9]. Conversely, reconstructing the reflection coefficient leads to images with high-spatial-frequency variations known as speckle. The signal-dependent speckle noise obscures image detail while shot

noise from the reference field creates a noise floor which can overpower weak signals. The combination of the two noise sources is referred to as doubly stochastic and can make tasks such as object identification and characterization difficult.

There has been significant progress made in the area of image denoising for an additive white Gaussian noise (AWGN) model. Gaussian denoising algorithms (GDAs) like non-local means (NLM) [10] and block-matching and 3D filtering (BM3D) [11] have pushed the limits of image denoising by leveraging self similarity within images through non-local, patched-based filtering. For most coherent imaging systems, the shot noise is well modeled as AWGN in the data domain; however, the signal-dependent speckle noise is not. Therefore, GDAs are not effective at reducing speckle variation.

Modifications to GDAs have been proposed for multiplicative noise models which modify the image data using a variance-stabilizing transform [12–15] or which modify the GDA itself [16–18]. While both methods help mitigate speckle noise, they do not account for shot noise and therefore produce poor results at low SNRs. In addition, they are applied in post-processing rather than incorporated into an image-formation process.

More advanced image reconstruction techniques have been proposed for synthetic aperture radar (SAR) imaging which combine image formation and denoising using regularized inversion [19–29]. However, most of these advanced techniques

enforce sparsity in the image domain which is not typically the case for optically-coherent imaging. They also regularize the reflection coefficient magnitude, $|g|$, rather than the reflectance, r , which results in speckled images. In [9] a model-based iterative reconstruction (MBIR) algorithm designed specifically for SAL was proposed. The algorithm computes the maximum *a posteriori* (MAP) estimate of the reflectance, r , rather than the reflection coefficient, g , using a relatively-simple prior and is the basis for the work in this paper.

In [30] a new framework called Plug and Play (P&P) was proposed for computing MAP estimates using more-advanced prior models inherent in GDAs. Through variable splitting and the use of Alternating Direction Method of Multipliers (ADMM) [31], the forward and prior models can be decoupled. Optimization over terms relating to the prior model can be mathematically formulated as a Gaussian denoising operation, and a GDA can be used to reduce the cost function.

Gaussian denoising can be viewed as computing a MAP estimate with an identity operator as the forward model and with a prior model which is inherent to the algorithm. By using GDAs in the P&P framework, we inherit the associated prior model without having to explicitly define it. Therefore, we may leverage the advances in GDAs to compute the MAP estimate, even for non-Gaussian forward models.

In [30] and [32], P&P was shown to be effective for applications with AWGN such as tomography and super resolution. In [33] and [34] the framework was extended to Poisson inverse problems. Rond, Giryes, and Elad also developed a framework called multi-prior plug and play (MPP&P) which can leverage multiple GDAs. However, standard P&P cannot easily be used to compute the MAP estimate of the reflectance for optically-coherent imaging systems due to the non-tractable cost function associated with the forward model. Additionally, the algorithm has not been demonstrated for use with multiplicative speckle noise or for applications with multiple noise sources.

In this paper, we show how the P&P framework can be extended for use with optically-coherent imaging systems to compute the MAP estimate of the real-valued reflectance function. The major contributions are:

1. An extension of the P&P algorithm is proposed which uses the expectation maximization (EM) algorithm to overcome the non-tractable cost function associated with coherent imaging.
2. Simulated and laboratory data are used to demonstrate that P&P, and therefore GDAs, can be effectively used for denoising coherent images corrupted by doubly-stochastic, exponentially-distributed noise.
3. High-performance computing (HPC) resources are used to compare six GDAs and one despeckling algorithm within the P&P framework. We determine which algorithm works best for optically-coherent imaging applications at various SNRs and for different image types.
4. We investigate the benefits of MPP&P for a doubly-stochastic application where different priors can be chosen to target the effects of the different noise sources.

2. MAP ESTIMATION USING PLUG AND PLAY

For coherent imaging systems, we formulate the image reconstruction problem as the MAP estimate of the scene's real-valued

reflectance function, $r \in \mathbb{R}^N$, from complex-valued noisy data, $y \in \mathbb{C}^M$ given by

$$\begin{aligned}\hat{r} &= \underset{r \in \Omega}{\operatorname{argmin}} \{ -\log p(r|y) \}, \\ &= \underset{r \in \Omega}{\operatorname{argmin}} \{ l(r) + \beta s(r) \},\end{aligned}\quad (1)$$

where Ω is the set of feasible solutions, $l(r) = -\log p(y|r)$, and $\beta s(r) = -\log p(r)$ [30]¹. The parameter β can be used to tune the prior model.

The P&P algorithm decouples the forward and prior models by splitting the variable r into both r and v and converting the unconstrained optimization of Eq. (1) into a constrained problem given by

$$\hat{r} = \underset{\substack{r \in \Omega \\ r=v}}{\operatorname{argmin}} \{ l(r) + \beta s(v) \} . \quad (2)$$

The ADMM algorithm is then used to solve the constrained optimization problem according to Algorithm 1 [30, 32]. The algorithm parameters are $\sigma_\lambda^2 = 1/\lambda$ [30], where λ controls the gain of the penalty to enforce the constraint, and u which acts as the Lagrange multiplier [35]. The two primary operations of the P&P algorithm are the inversion operator given by

$$\mathbb{F}(\tilde{r}; \sigma_\lambda) = \underset{r}{\operatorname{argmin}} \left\{ l(r) + \frac{1}{2\sigma_\lambda^2} \|r - \tilde{r}\|_2^2 \right\}, \quad (3)$$

where $\tilde{r} = v - u$, and the denoising operator given by

$$\mathbb{H}(\tilde{v}; \sigma_n) = \underset{v}{\operatorname{argmin}} \left\{ \frac{1}{2\sigma_n^2} \|\tilde{v} - v\|_2^2 + s(v) \right\}, \quad (4)$$

where $\sigma_n^2 = \beta \sigma_\lambda^2$ and $\tilde{v} = r + u$. Using this notation, the P&P algorithm for computing the MAP estimate is shown as Algorithm 1 below.

Algorithm 1. Plug and Play Algorithm

```

Initialize:  $v^{(0)}, u^{(0)} = 0$ 
Repeat{
     $\tilde{r}^{(k+1)} = v^{(k)} - u^{(k)}$ 
     $r^{(k+1)} = \mathbb{F}(\tilde{r}^{(k+1)}; \sigma_\lambda)$ 
     $\tilde{v}^{(k+1)} = r^{(k+1)} + u^{(k)}$ 
     $v^{(k+1)} = \mathbb{H}(\tilde{v}^{(k+1)}; \sigma_n)$ 
     $u^{(k+1)} = u^{(k)} + (r^{(k+1)} - v^{(k+1)})$ 
}

```

The operator \mathbb{F} is a proximal mapping of $l(r)$ [35]. It is mathematically equivalent to the MAP estimate of r using a Gaussian prior with distribution $p(r) \sim N(\tilde{r}, \sigma_\lambda^2 I)$, where $N(\mu, R)$ indicates a Gaussian distribution with mean μ and covariance matrix R .

The operator \mathbb{H} is a proximal mapping of $s(v)$. It is mathematically equivalent to a Gaussian denoising operation. Specifically, it is equivalent to computing the MAP estimate of v from the noisy image \tilde{v} . The variable σ_n^2 represents the noise variance of an i.i.d. Gaussian forward model, and $s(v)$ is the prior [30, 32]. Therefore, the ADMM cost function can be reduced by replacing \mathbb{H} with a GDA. We thus inherit the prior model which is either explicitly or implicitly part of the GDA.

¹In practice, \hat{r} is proportional to the actual reflectance by an unknown multiplicative constant α

In general, multiple prior models can be combined into a single hybrid prior model by splitting the negative log-prior model into K terms given by [33]

$$-\log p(r) = \sum_{i=1}^K \beta_i s_i(r). \quad (5)$$

In this paper we investigate the MPP&P algorithm by combining two different priors, i.e. $K = 2$. One prior is chosen to help regularize the effects of measurement noise, and the other to help regularize the effects of speckle. In App. A, we provide the MPP&P algorithm.

3. COHERENT-IMAGING FORWARD MODEL

For imaging systems that use optically-coherent illumination and detection, the complex data can be represented using an additive noise model given by [9]

$$y = Ag + w. \quad (6)$$

In Eq. (6) the matrix $A \in \mathbb{C}^{M \times N}$ models the propagation geometry for the particular imaging modality, $g \in \mathbb{C}^N$ is a discrete representation of the complex-valued reflection coefficient for the scene being imaged, and $w \in \mathbb{C}^M$ models the measurement noise. While the number of measurements, M , is fixed, we may choose $N > M$ if we wish to reconstruct r with increased resolution. However, we use $N = M$ in this work.

Examples of the matrix A include a Discrete Space Fourier Transform (DSFT) for Digital Holography (DH), a non-uniform DSFT for Synthetic Aperture LADAR (SAL), and the identity matrix for simple image denoising. Field-modifying functions, such as phase errors and pupil functions, may be included in A when applicable, and jointly estimated as necessary [9].

The vector, g , can be modeled as a conditionally complex Gaussian random variable for surfaces which are rough relative to the illumination wavelength. Given the underlying reflectance of the scene, r , the conditional distribution of the reflection coefficient is given by

$$p(g|r) \sim CN(0, \mathcal{D}(r)), \quad (7)$$

where CN indicates a complex normal distribution and $\mathcal{D}(r)$ is a matrix with diagonal elements equal to the vector, r .

It is common for coherent detection systems to use a reference beam which is much stronger than the return signal. In such cases, shot noise driven by the power of the reference beam is the dominant source of measurement noise and can be modeled as additive, zero-mean, complex Gaussian white noise [1, 36]. The distribution of w is therefore given by

$$p(w) \sim CN(0, \sigma_w^2 I), \quad (8)$$

where σ_w^2 is the noise variance and I is the identity matrix.

Given Eqs. 6-8, it can be shown that the likelihood function of the data has a conditionally complex Gaussian distribution given by

$$p(y|r) \sim CN(0, A\mathcal{D}(r)A^H + \sigma_w^2 I), \quad (9)$$

where H indicates the Hermitian transpose [37]. The negative log-likelihood function can then be written as

$$\begin{aligned} l(r) &= -\log p(y|r), \\ &= \log |A\mathcal{D}(r)A^H + \sigma_w^2 I| + y^H [A\mathcal{D}(r)A^H + \sigma_w^2 I]^{-1} y + c, \end{aligned} \quad (10)$$

where c includes all terms which are constant with respect to r .

Unfortunately, it is often the case that evaluation of the determinate and inverse in Eq. (10) is a non-tractable problem. Therefore, using the P&P framework is not practical for optically-coherent-image formation without modification.

4. P&P USING EM ALGORITHM

In this section we describe how the P&P algorithm is extended for use with coherent imaging applications.

A. EM Algorithm Surrogate Function

For any function $q(x)$, we define a surrogate function, $\tilde{q}(x, x')$, to be an upper-bounding function such that $q(x) \leq \tilde{q}(x, x') + c$, where x' is the current value of x which determines the functional form of \tilde{q} , and c is a constant that ensures the two functions are equal at x' . Surrogate functions have the property that minimization of \tilde{q} implies the minimization of q . That is

$$\{\tilde{q}(x, x') < \tilde{q}(x', x')\} \Rightarrow \{q(x) < q(x')\}. \quad (11)$$

To overcome the non-tractable determinate and inverse in Eq. (10), we can use the EM algorithm to replace the log-likelihood function, $l(r)$, with a more tractable surrogate function. Thus, the P&P inversion operator, \mathbb{F} , will be replaced with an alternate function, $\tilde{\mathbb{F}}$. To start, we introduce the reflection coefficient, g , into the forward model as the missing data. The resulting surrogate function for $l(r)$, given by the EM algorithm, has the form

$$\tilde{l}(r; r') = -E [\log p(y, g | r) | Y = y, r'], \quad (12)$$

where r' is the current estimate of r , and the tilde indicates it's a surrogate function. The operator $E[\cdot | Y = y, r']$ is a conditional expectation over g , given y and r' .

Using Bayes' theorem and the fact that $p(y|g, r) = p(y|g)$, Eq. (12) becomes

$$\begin{aligned} \tilde{l}(r; r') &= -E [\log p(g|r) | Y = y, r'] + c, \\ &= \log |\mathcal{D}(r)| + \sum_{i=1}^N \frac{1}{r_i} E [|g_i|^2 | Y = y, r'] + c, \end{aligned} \quad (13)$$

where c is the sum of terms constant with respect to r . To evaluate the second moment in Eq. (13), we must specify the conditional posterior distribution of g . We can use Bayes' theorem and $p(y|g, r) = p(y|g)$ to write the posterior distribution as

$$\begin{aligned} p(g|y, r) &= \frac{p(y|g)p(g|r)}{p(y|r)}, \\ &= \frac{1}{z} \exp \left\{ -\frac{1}{\sigma_w^2} \|y - Ag\|^2 - g^H \mathcal{D}(r)^{-1} g \right\}, \end{aligned} \quad (14)$$

where z is the partition function which has absorbed any exponential terms that are constant with respect to g . By completing the square, it can be shown that the posterior distribution is a complex Gaussian with mean

$$\mu = C \frac{1}{\sigma_w^2} A^H y \quad (15)$$

and covariance

$$C = \left[\frac{1}{\sigma_w^2} A^H A + \mathcal{D}(r)^{-1} \right]^{-1}. \quad (16)$$

The surrogate for the log likelihood function then becomes

$$\tilde{l}(r; r') = \log |\mathcal{D}(r)| + \sum_{i=1}^N \frac{1}{r_i} \left(C_{i,i} + |\mu_i|^2 \right), \quad (17)$$

where μ_i is the i^{th} element of the posterior mean and $C_{i,i}$ is the i^{th} diagonal element of the posterior covariance. Both μ and C are computed using the current reflectance value r' .

The likelihood surrogate can then be substituted back into Eq. (3) to obtain a modified inversion operator given by

$$\tilde{\mathbb{F}}(\tilde{r}; r', \sigma_\lambda) = \underset{r}{\operatorname{argmin}} \left\{ \tilde{l}(r; r') + \frac{1}{2\sigma_\lambda^2} \|r - \tilde{r}\|_2^2 \right\}. \quad (18)$$

Evaluation of the expectation in Eq. (13) constitutes the E-step of the EM algorithm while the minimization in the modified inversion operation given by Eq. (18) represents the M-step of the EM algorithm. Putting these together, Algorithm 2 then shows the combined steps for the P&P algorithm with EM updates which uses the modified inversion operator.²

Algorithm 2. Plug and Play Algorithm with EM

```

Initialize:  $v^{(0)}, u^{(0)} = 0$ 
Repeat{
     $\tilde{r}^{(k+1)} = v^{(k)} - u^{(k)}$ 
     $r^{(k+1)} = \tilde{\mathbb{F}}(\tilde{r}^{(k+1)}; r^{(k)}, \sigma_\lambda)$ 
     $\tilde{v}^{(k+1)} = r^{(k+1)} + u^{(k)}$ 
     $v^{(k+1)} = \mathbb{H}(\tilde{v}^{(k+1)}; \sigma_n)$ 
     $u^{(k+1)} = u^{(k)} + (r^{(k+1)} - v^{(k+1)})$ 
}

```

Each iteration of the P&P algorithm requires an evaluation of the proximal mapping \mathbb{F} , which in turn requires the exact minimization of the cost function associated with the optimization in Eq. (3). However, in our case, we compute an approximation to the proximal mapping that results from the use of the surrogate function given by the EM algorithm's \tilde{l} function. More specifically, each update of Algorithm 2 uses the function $\tilde{\mathbb{F}}$ given in Eq (18), rather than the true \mathbb{F} of Eq. (3).

This raises the question as to whether the modified algorithm using $\tilde{\mathbb{F}}$ converges to the same result as the original formulation. Convergence of the P&P algorithm is itself a complex issue with most theoretical results only holding for the case of convex optimization [32, 38, 39].

However, we can show that a fixed point of the P&P algorithm will also be a fixed point of the modified Algorithm 2. This ensures that any solution to the P&P iterations will be a fixed point to the modified iterations of Algorithm 2; and in particular, any solution to the MAP optimization of Eq. (1) will also be a fixed-point solution to Algorithm 2.

In order to see this, note that any fixed point to the P&P algorithm must solve the equations

$$\begin{aligned} \mathbb{F}(r^* - u^*; \sigma_\lambda) &= r^* \\ \mathbb{H}(r^* + u^*; \sigma_n) &= r^*, \end{aligned} \quad (19)$$

where (r^*, u^*) are the values of (r, u) that achieve the fixed point to Algorithm 1. This relationship can be derived by substituting $v^{(k)} = r^{(k)} = r^*$ and $u^{(k)} = u^*$ in Algorithm 1. Assuming that

$l(r)$ and $s(v)$ are continuously differentiable, then this implies that

$$0 = \nabla_r \left\{ l(r) + \frac{1}{2\sigma_\lambda^2} \|r - \tilde{r}\|^2 \right\} \Big|_{r=r^*} \quad (20)$$

which implies that

$$\nabla_r l(r^*) = -\frac{1}{\sigma_\lambda^2} (r^* - \tilde{r}). \quad (21)$$

If $\tilde{l}(r; r')$ denotes the surrogate function to $l(r)$, then we know that

$$\nabla_r \tilde{l}(r; r^*) \Big|_{r=r^*} = \nabla_r l(r^*) = -\frac{1}{\sigma_\lambda^2} (r^* - \tilde{r}). \quad (22)$$

So therefore, we know that

$$0 = \nabla_r \left\{ \tilde{l}(r; r^*) + \frac{1}{2\sigma_\lambda^2} \|r - \tilde{r}\|^2 \right\} \Big|_{r=r^*}, \quad (23)$$

and that

$$\begin{aligned} \tilde{\mathbb{F}}(r^* - u^*; r^* - u^*, \sigma_\lambda) &= r^* \\ \mathbb{H}(r^* + u^*; \sigma_n) &= r^*. \end{aligned} \quad (24)$$

The relationship of (24) ensures that (r^*, u^*) is a fixed point of Algorithm 2.

B. Optimization

Algorithm 2 is implemented, in part, by executing the optimizations that make up the modified inversion and denoising operators. In this section we describe how that is accomplished.

B.1. Modified Inversion Operation

Since the elements of r are not coupled by the cost function associated with the modified inversion operator of Eq. (18), we use iterative coordinate descent (ICD) to optimize. The cost function for the i^{th} element is given by

$$c(r_i; \tilde{r}_i, \mu_i, C_{i,i}, \sigma_\lambda) = \log r_i + \frac{1}{r_i} \left(C_{i,i} + |\mu_i|^2 \right) + \frac{1}{2\sigma_\lambda^2} (r_i^2 - 2r_i \tilde{r}_i). \quad (25)$$

If Eq. (25) is differentiated with respect to r_i , set equal to zero, and both sides multiplied by r_i^2 , we see that the solution is contained within the set of roots for a 3rd order polynomial given by

$$\alpha_1 r_i^3 + \alpha_2 r_i^2 + \alpha_3 r_i + \alpha_4, \quad (26)$$

where

$$\alpha_1 = \frac{1}{\sigma_\lambda^2}, \quad \alpha_2 = \frac{-\tilde{r}_i}{\sigma_\lambda^2}, \quad \alpha_3 = 1, \quad \alpha_4 = -\left(C_{i,i} + |\mu_i|^2 \right). \quad (27)$$

Therefore, the inversion operation can be carried out by simply rooting this polynomial for each pixel. The general solution to a 3rd order polynomial can be found in closed form using the cubic equation. For convenience, we used MATLAB's *roots.m*. If multiple real-valued roots exist, we use the cost function to determine the best choice among them.

B.2. Denoising Operation

One of the core ideas behind P&P is that we use a GDA to reduce the cost function associated with minimization in Eq. (4). Since each GDA has a unique prior model associated with it, we must determine which works best for coherent imaging applications. Thus, we compare results using six well-known GDAs as well as a leading speckle-denoising algorithm [17]. In addition, two MPP&P algorithms are designed to specifically target the effects of speckle and shot noise. Each algorithm is described in App. B.

²To use this approach with MPP&P, we replace $l(r)$ with $\tilde{l}(r; r')$ in Eq. (34).

B.3. Stopping Criteria

The algorithm is stopped after either a set number of iterations, N_K , has been reached or

$$\frac{\|r^k - r^{k-1}\|}{\|r^{k-1}\|} \leq \epsilon_T, \quad (28)$$

where k is the iteration index and ϵ_T is a specified stopping threshold.

B.4. Initialization

Since the cost function associated with the minimization in Eq. (3) is non-convex, the initial conditions are important in avoiding local minima. To begin algorithm 2 we initialize according to

$$\begin{aligned} v^{(0)} &= |A^H y|^{\circ 2}, \\ u^{(0)} &= \mathbf{0}, \\ \sigma_\lambda &= \frac{1}{2} \sqrt{s^2(v^{(0)})}, \end{aligned} \quad (29)$$

where $|\cdot|^{\circ 2}$ indicates the element-wise magnitude square of a vector and $s^2(\cdot)$ computes the sample variance of a vector's elements [40]. These initial conditions were found to be effective for computing the MAP estimate of r in a conventional, non-P&P framework [9]. The noise power, σ_w^2 , is assumed to be known via system calibration.

5. RESULTS

In this section results are presented for both simulated and laboratory data. Although the proposed approach is applicable to several optically-coherent imaging modalities, we focused on inverse synthetic aperture LADAR (ISAL) data. High performance computing resources were used to test the proposed algorithms over a wide range of images, SNRs, and model parameters.

A. Simulation Results

To simulate mock data we started with the reflectance function, r , and generated a realization of the reflection coefficient, g , according to the distribution given in Eq. (7). We then used Eq. (6) to produce the measured data, y . The ISAL linear forward-model operator, A , described in App. B of [9] was used. The model assumes a band limited and Nyquist sampled signal, a square A matrix (i.e. $M = N$), and linear object rotation. Under these conditions, A represents a skewed DSFT given by Eq. (65) of [9]. Data was generated over a range of SNR values, where SNR is defined as

$$\text{SNR} = \frac{s^2(Ag)}{s^2(w)}. \quad (30)$$

Figure 1 shows the input images used for the reflectance functions. All but *test targ* are part of MATLAB's standard library of images. In some cases the contrast was increased to produce a more-discernible coherent image by raising each element to the power γ .

We found any unknown scaling constant, α , between r and \hat{r} by computing the least squares fit given by

$$\alpha^* = \underset{\alpha}{\operatorname{argmin}} \left\{ \|\alpha \hat{r} - r\|^2 \right\}. \quad (31)$$

We then measured the reconstruction distortion using normalized root mean square error (NRMSE) defined as

$$\text{NRMSE} = \sqrt{\frac{\|\alpha^* \hat{r} - r\|^2}{\|r\|^2}}. \quad (32)$$

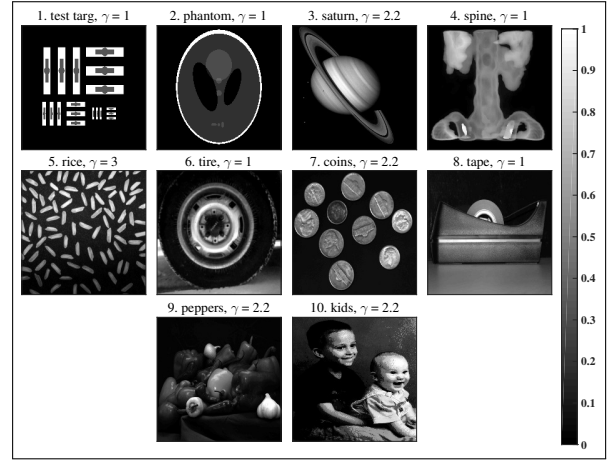


Fig. 1. Simulation input images along with gamma value used for increasing contrast.

NRMSE was used to find the optimal values of β for each algorithm (see App. B).

To compare the proposed algorithm against standard, Fourier-based reconstruction (FBR) techniques, we generated comparison images according to

$$\hat{r}_{\text{FBR}} = |W_{\text{FFT}} \mathcal{D}(T)y|^{\circ 2}, \quad (33)$$

where W_{FFT} is a two-dimensional FFT matrix and T is a $M \times 1$ vector of weights corresponding to a Taylor window.

Table 1 shows how NRMSE varied as a function of SNR for each algorithm. The results are ordered by their average over all 60 reconstructions (10 targets \times 6 SNRs). The remaining columns are the average over all 10 images for each SNR. The bottom five algorithms (P&P KSVD, MPP&P 1, P&P BM3D, P&P QGMMRF, and P&P TV) performed significantly better than the top four, making them the most suitable for reconstructing coherent images over a wide range of conditions. Not surprisingly, the technique using the despeckling algorithm BLR did poorly. This is largely due to poor performance at low SNRs where it is unable to mitigate the effects of measurement noise.

Table 2 shows how NRMSE varied as a function of image type. The results show that P&P BM3D performed best or near best for the simple, and more discrete-type images (i.e. images 1-4). Its performance tapered off for the more complex images (i.e. 5-10). Conversely, P&P KSVD performed best for the complex images, but did not do as well with the simple images.

Figure 2 shows a subset of the resulting reconstructions for three specific targets of differing complexity. The leftmost image in each row is the FBR reconstruction. To the right of each FBR image, are the top three reconstructions ordered by descending NRMSE values. It is interesting to note that although P&P BM3D didn't always have the lowest NRMSE, it produced images with sharper edges compared to P&P KSVD.

B. Laboratory Results

Experimental data was produced in the ISAL laboratory at JPL [41, 42]. The bench-top setup was a bi-static system consisting of a transmitter, heterodyne receiver, and rotating target. A 1310 nm tunable laser produced a 34 ms linear frequency modulated (LFM) chirped pulse with a chirp rate of 2 THz/s and a pulse repetition frequency (PRF) of 10 Hz. A self heterodyning system was employed with 10% of the transmitted signal used

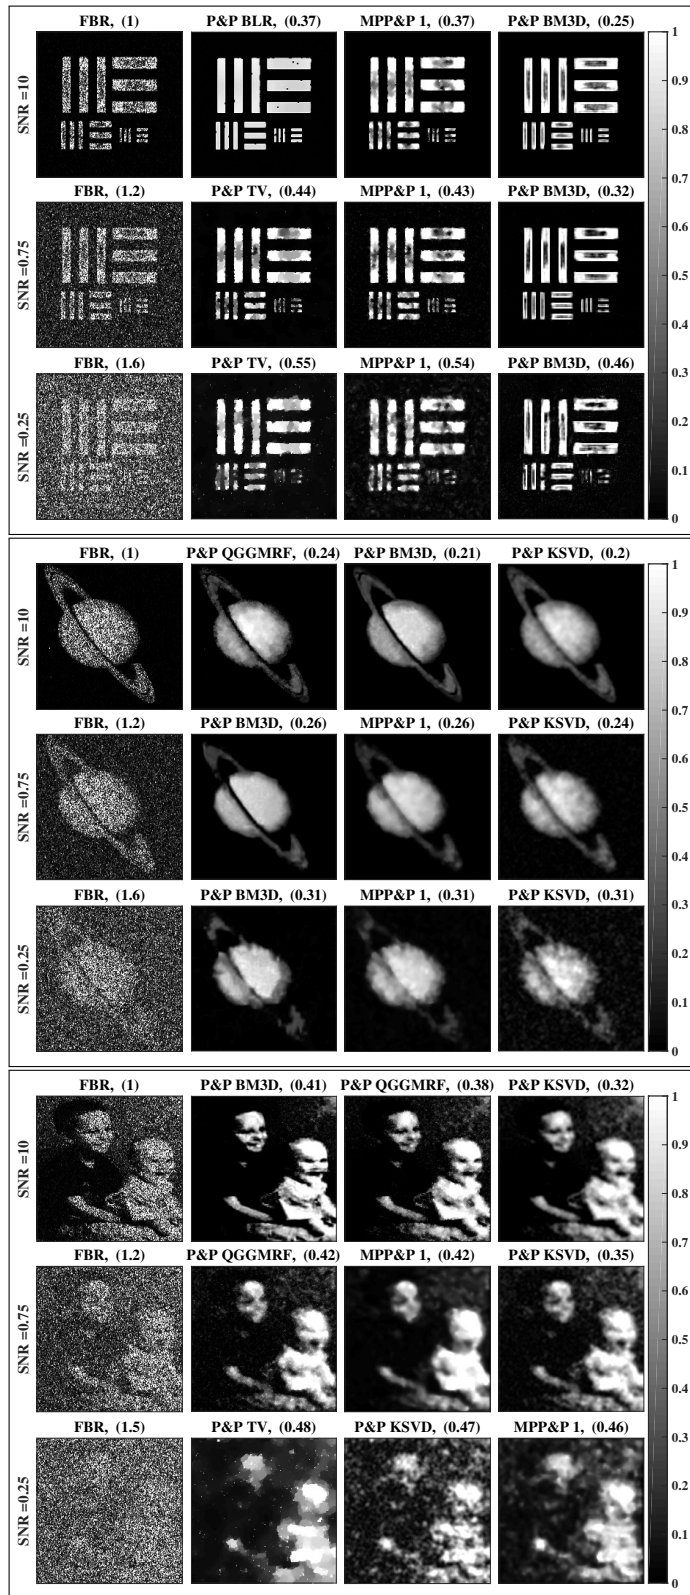


Fig. 2. Simulation results for three different targets at SNRs of 10, 0.75, and 0.25. The leftmost image in each row is the FBR reconstruction. To the right of each FBR image are the top three reconstructions ordered by descending NRMSE values. The NRMSE values are shown in parenthesis.

Table 1. NRMSE averaged over ten images vs. SNR

	SNR						Avg.
	10	3	1	0.75	0.5	0.25	
FBR	1.0	1.0	1.1	1.2	1.3	1.5	1.2
P&P NLM	0.91	0.88	0.98	1.0	1.1	1.3	1.0
MPP&P 2	0.42	0.45	0.5	0.52	0.64	1.1	0.61
P&P DR	0.53	0.55	0.62	0.61	0.61	0.65	0.59
P&P BLR	0.4	0.42	0.54	0.57	0.65	0.85	0.57
P&P TV	0.39	0.39	0.42	0.42	0.45	0.51	0.43
P&P QGGMRF	0.36	0.37	0.39	0.41	0.43	0.52	0.41
P&P BM3D	0.33	0.35	0.4	0.41	0.45	0.54	0.41
MPP&P 1	0.38	0.38	0.39	0.39	0.4	0.48	0.4
P&P KSVD	0.33	0.34	0.37	0.38	0.4	0.5	0.39

as the local oscillator. The detector sample rate was 1 MHz. A rotation stage was used to provide relative movement between the transmitter/receiver and target. The data was demodulated, low-pass filtered, and down sampled to isolate the narrow-band signal of interest. The resulting two-dimensional data array had 90 samples in range and 180 samples in cross-range. The approximate resolution of the FBR images is 2.3 mm in range and 0.88 mm in cross range. Phase errors were estimated and removed prior to image reconstruction using the technique described in [9].

Figure 3 shows the targets used in the experiment. The first is a stencil on a Lambertian plate and the second is a 3-D printed satellite made from a single material. The logo was rotated at $6.25 \mu\text{rad/s}$ for 120 s and the satellite was rotated at $12.5 \mu\text{rad/s}$ for 60 s.

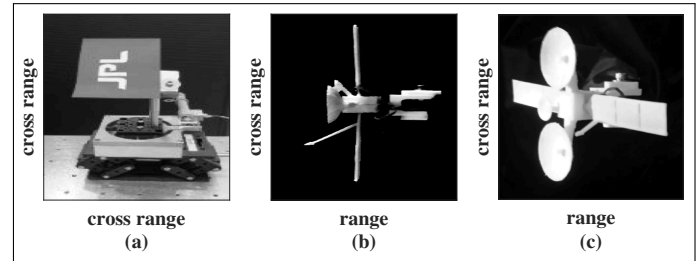


Fig. 3. Targets used for laboratory experiments. (a) JPL logo with range axis into the page. The stencil was placed on a Lambertian reflector angled at 45 degrees. (b) Mock satellite (Orientation 1). (c) Mock satellite (Orientation 2).

Figure 4 shows the reconstructions for each of the three data sets. The FBR images are shown in the top left of each block. For all algorithms, reconstructions were conducted over the same parameter range used in the simulated experiment. Since no truth image was available to compute the NRMSE, the best image had to be picked by eye. Upon visual inspection of the FBR images, we see that the data have high SNRs. Since speckle is the dominant noise source, we expect to observe a reduction in the speckle variation when using P&P, while a reduction in measurement noise may be less noticeable. The logo results show that, with such low resolution in the lab data, blurring from over regularization was a problem. P&P KSVD did reduce the speckle

Table 2. NRMSE averaged over six SNRs vs. image number

	Image Number										Avg.
	1	2	3	4	5	6	7	8	9	10	
FBR	1.2	1.2	1.2	1.2	1.2	1.2	1.2	1.1	1.2	1.2	1.2
P&P NLM	1.2	1.1	1.1	0.98	1.0	1.0	1.0	0.88	1.1	1.1	1.0
MPP&P 2	0.51	0.73	0.43	0.48	0.94	0.67	0.53	0.59	0.59	0.63	0.61
P&P DR	0.59	0.9	0.3	0.48	0.77	0.55	0.47	0.67	0.65	0.57	0.59
P&P BLR	0.55	0.71	0.43	0.43	0.76	0.65	0.48	0.51	0.64	0.56	0.57
P&P TV	0.45	0.48	0.3	0.32	0.67	0.46	0.39	0.39	0.42	0.45	0.43
P&P QGGMRF	0.45	0.5	0.27	0.35	0.5	0.41	0.38	0.46	0.38	0.43	0.41
P&P BM3D	0.33	0.38	0.25	0.34	0.59	0.44	0.4	0.47	0.44	0.48	0.41
MPP&P 1	0.44	0.49	0.27	0.32	0.54	0.41	0.36	0.38	0.39	0.42	0.4
P&P KSVD	0.52	0.6	0.24	0.29	0.45	0.33	0.35	0.38	0.34	0.36	0.39

variation compared to FBR, but the algorithm also blurred the letters J and P together. Conversely, P&P BM3 reduced the speckle and preserved the resolution. From the satellite data, we see that the algorithms that performed well with the simulated data were able to reduce the speckle variation. Even with the poor resolution and relatively small targets, the laboratory results appear consistent with the high-SNR simulated results in terms of speckle reduction and relative algorithm performance.

6. CONCLUSION

In this paper we have presented an extension of the P&P algorithm for use with optically-coherent imaging. The proposed algorithm reduces measurement noise and speckle variation by computing the MAP estimate of the reflectivity to produce a more natural looking image. To overcome the non-tractable MAP cost function, the EM algorithm was used to reduce the complexity of the problem by deriving a more tractable surrogate function. Additionally, the P&P framework allows us to use the denoising properties of GDAs as part of the MAP optimization process and modeling process.

The strength of the proposed algorithm was demonstrated using both simulated and laboratory data. HPC resources were used to test the framework over a wide range of GDAs, model parameters, input images, and SNRs. We also tested two MPP&P algorithms designed specifically for low-SNR coherent imaging. Simulated data was used to compare the various algorithms and quantitatively rank their performance. We also showed that the proposed approach can significantly reduce reconstruction error by reducing the effects of measurement noise and speckle variation. Finally, we validated our proposed algorithms using experimental data from the JPL ISAL laboratory. While the laboratory data was limited in terms of resolution and fixed SNR, the results were consistent with the simulated results.

APPENDIX A: MULTI-PRIOR PLUG AND PLAY

Algorithm 3 shows the update steps for MP&P. The inversion operator is given by

$$\mathbb{H}(\tilde{r}_1, \tilde{r}_2; \sigma_{\lambda_1}, \sigma_{\lambda_2}) = \underset{r}{\operatorname{argmin}} \left\{ l(r) + \sum_{i=1}^2 \frac{1}{2\sigma_{\lambda_i}^2} \|r - \tilde{r}_i\|_2^2 \right\}, \quad (34)$$

where $\tilde{r}_i = v_i - u_i$, and the two denoising operators are defined as

$$\mathbb{H}_i(\tilde{v}_i; \sigma_{n_i}) = \underset{v}{\operatorname{argmin}} \left\{ \frac{1}{2\sigma_{n_i}^2} \|\tilde{v}_i - v_i\|_2^2 + s_i(v_i) \right\}, \quad \text{for } i = 1, 2, \quad (35)$$

where $\tilde{v}_i = r + u_i$ and $\sigma_{n_i} = \beta_i \sigma_{\lambda_i}^2$.

Algorithm 3. Multi-Prior Plug and Play Algorithm

Initialize: $v_i^{(0)}, u_i^{(0)} = 0$ for $i = 1, 2$
Repeat{
 $\tilde{r}_i^{(k+1)} = v_i^{(k)} - u_i^{(k)}$ for $i = 1, 2$
 $r^{(k+1)} = \mathbb{F}(\tilde{r}_1^{(k+1)}, \tilde{r}_2^{(k+1)}; \sigma_{\lambda_1}, \sigma_{\lambda_2})$
 $\tilde{v}_i^{(k+1)} = r^{(k+1)} + u_i^{(k)}$ for $i = 1, 2$
 $v_i^{(k+1)} = \mathbb{H}_i(\tilde{v}_i^{(k+1)}; \sigma_{n_i})$ for $i = 1, 2$
 $u_i^{(k+1)} = u_i^{(k)} + (r^{(k+1)} - v_i^{(k+1)})$ for $i = 1, 2$
}

APPENDIX B: ALGORITHM DESCRIPTIONS

In this section we describe the denoising algorithms tested in the P&P and MPP&P frameworks. A brief description is provided for how each algorithm computes a denoised output, \hat{x} , from noisy input, y . Most denoising algorithms have several parameters which must be selected in addition to the P&P parameters, σ_λ and β . To reduce the degrees of freedom, we automate the choice of σ_λ according to Eq. (29) and fix all the denoising parameters except for one. Generally, this tuning parameter is the estimate of the noise variance, $\sigma_n^2 = \beta \sigma_\lambda^2$, which controls the strength of the denoising operation. For each of the 60 simulated reconstructions per algorithm (6 SNRs \times 10 images), we chose the parameter which minimized NRMSE. Table 3 provides the range of optimal parameters used for each algorithm.

Plug and Play with Non-Local Means (P&P NLM)

Non-local means (NLM) produces a point-wise estimate of each pixel which leverages the redundancy in natural images [10].

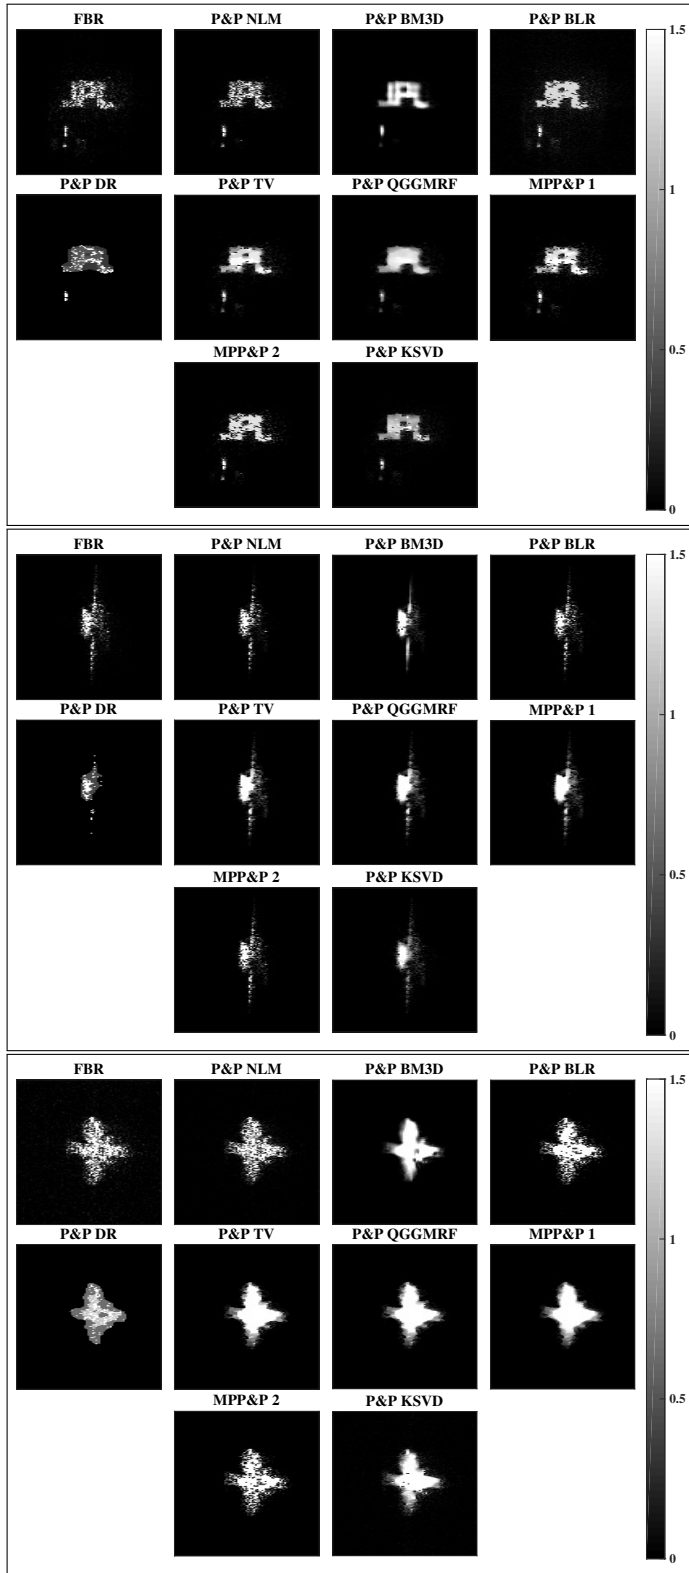


Fig. 4. Reconstructions from laboratory data. For each data set, the reconstructions for all 10 algorithms are shown.

Table 3. Summary of Optimal Tuning Parameters

Alg	Parameter	Range	Avg
P&P NLM	β	0.05 – 0.51	0.16
P&P BM3D	β	110 – 250	200
P&P TV	β	0.42 – 5.0	0.49
P&P DR	β	4.2 – 15	12
P&P QGGMRF	T	0.15 – 1.0	0.56
P&P KSVD	β	3 – 5	3.9
P&P BLR	h	0.22 – 1.0	0.69
MPP&P 1	β_2	0.83 – 9.7	3.4
MPP&P 2	β_2	0.2 – 0.3	0.27

Each pixel becomes a weighted average of other pixels in the image based on the similarity between their neighborhoods. The weights between pixels, y_i and y_j , are given by

$$w_{i,j} = \frac{1}{c} \exp \left\{ \frac{1}{\sigma_n^2} \|b_i - b_j\|_2^2 \right\}, \quad (36)$$

where c is a normalizing constant and b_* is the $l_n \times l_n$ neighborhood for the $*^{th}$ pixel. Typically the search for pixels with similar neighborhoods is limited to some finite search window of size $l_s \times l_s$. We fixed the variables $l_n = 3$, $l_s = 41$, and varied σ_n^2 using β .

Plug and Play w/ Block-Matching and 3D Filtering (P&P BM3D)

Block-matching and 3D filtering (BM3D) is a patched based denoising algorithm which stacks similar patches and collaboratively filters them using shrinkage techniques in a transform domain [11]. BM3D code was downloaded from [43]. The noise profile was set to “*vn*”, for very noisy, and the hard thresholding block size to 4×4 . All other parameters were left at their default values. The noise variance σ_n^2 was varied using β . We found that large values of σ_n^2 were required to sufficiently regularize the speckle variations. In many cases the optimal value of σ_n exceeded the range of values in y .

Plug and Play with Total Variation (P&P TV)

Total variation (TV) performs image denoising by solving the optimization problem given by

$$\hat{x} = \underset{x}{\operatorname{argmin}} \left\{ \|y - x\|_2^2 - \lambda \|Dx\| \right\}, \quad (37)$$

where D computes the magnitude of the directional derivative for each element in x . A TV algorithm was downloaded from [44]. Following the example script that was supplied with the code, we set $\lambda = \sigma_n$ and fixed the number of iterations at $n_{iter} = 10$ for each loop of P&P. All other parameters were set at their default values.

Plug and Play with Discrete Representation (P&P DR)

The discrete representation (DR) algorithm computes the MAP estimate of x using a discrete-state prior which biases the output towards one of K discrete states [45, 46]. The joint, non-convex

optimization problem to compute both the class labels, b , for each element of x , and the class values, is given by

$$(\hat{\mu}, \hat{b}) = \underset{\mu, b}{\operatorname{argmin}} \left\{ \frac{1}{2\sigma_n^2} \|y - \mu(b)\|_2^2 + \sum_{\{i,j\} \in \mathcal{P}} w_{i,j} \delta(b_i \neq b_j) \right\}. \quad (38)$$

The operator $\mu(b)$ converts each element of the class vector, b , to its corresponding value, \mathcal{P} is the set of all pair-wise cliques, $w_{i,j}$ is the weight between elements i and j within a neighborhood, and δ is an indicator function which is zero if the two classes are the same. The final output is given by $\hat{x} = \hat{\mu}(\hat{b})$. The neighborhood size was set to 3×3 with weights equal to $1/12$ for diagonal neighbors and $1/6$ for horizontal or vertical neighbors. Each time DR was called from within the P&P framework it was run for three iterations. The prior model was limited to $K = 8$ states. To initialize DR, we used the MATLAB functions *imquantize.m* and *multithresh.m* according to

$$x^{(0)} = \text{imquantize}(y_{LP}, \text{multithresh}(y_{LP}, K - 1)). \quad (39)$$

The variable y_{LP} is a low-pass filtered version of the input image using a 5×5 Gaussian kernel with $\sigma = 1$ pixel. The noise variance σ_n^2 was varied using β .

Plug and Play with Q-Generalized Gaussian Markov Random Field (P&P QGGMRF)

The Q-Generalized Gaussian Markov Random Field (QGGMRF) computes the MAP estimate of x using a Markov random field as the prior [35, 47]. For image denoising, the output is the solution to the optimization problem given by

$$\hat{x} = \underset{x}{\operatorname{argmin}} \left\{ \frac{1}{2\sigma_n^2} \|y - x\|_2^2 + \sum_{\{i,j\} \in \mathcal{P}} w_{i,j} \rho\left(\frac{\Delta}{\sigma_r}\right) \right\}, \quad (40)$$

where \mathcal{P} is the set of all pair-wise cliques, $w_{i,j}$ is the weight between elements i and j , $\Delta = x_i - x_j$, and σ_r controls the variation in r . The potential function, which determines how much influence pixels have on each other, is given by

$$\rho\left(\frac{\Delta}{\sigma_r}\right) = \frac{|\Delta|^p}{p\sigma_r^p} \left(\frac{|\frac{\Delta}{\sigma_r}|^{q-p}}{1 + |\frac{\Delta}{\sigma_r}|^{q-p}} \right). \quad (41)$$

The threshold value, T , controls the transition of the potential function from having the exponent q to having the exponent p . We limited the algorithm to 10 iterations and fixed the following parameters: $q = 2$, $p = 1.1$, and w was a 3×3 Gaussian kernel with $\sigma = 0.8$ pixels. The scale parameter was set to $\sigma_r = \frac{1}{2} \sqrt{s^2(y)}$. To tune P&P QGGMRF, we fixed the noise variance at $\sigma_n = \sigma_\lambda$ and varied the threshold, T .

Plug and Play with K-SVD (P&P K-SVD)

K-SVD is a denoising algorithm that uses redundant dictionaries for sparse representation of patches within an image [48]. The dictionary, D , and corresponding coefficients, α , can be learned from the noisy image using a joint optimization given by

$$(\hat{D}, \hat{\alpha}, \hat{x}) = \underset{D, \alpha, x}{\operatorname{argmin}} \left\{ \lambda \|x - y\|_2^2 + \sum_{p \in P} \mu \|\alpha_p\|_0 + \sum_{p \in P} \|(D\alpha)_p - x_p\|_2^2 \right\}. \quad (42)$$

The variable λ is the Lagrangian multiplier, μ controls the relative weight of the second term, P is the set of all patches in the image, and the subscript p indicates just the elements belonging to the p^{th} patch. The second term enforces sparsity in each patch and the third term ensures that the error of the sparse representation is bounded. K-SVD code was downloaded from [49]. Each time K-SVD was called from within the P&P framework it was run for a single iteration. The block size was set to 4×4 , the dictionary size was 100 atoms, and 40×10^3 training signals were used. The code chooses λ automatically based on the user-supplied noise variance σ_n^2 .

Plug and Play with Bayesian Likelihood Ratio (P&P BLR)

The image denoising algorithm in [17], which we will call the Bayesian likelihood ratio (BLR) algorithm, is a generalization of the NLM algorithm to account for non-Gaussian noise distributions. It has become recognized as one of the top performing denoising algorithms for speckle removal [16]. The denoising process is identical to NLM except that patch similarity is determined using a Bayesian likelihood ratio with weights given by

$$w_{i,j} = \frac{1}{c} \times \left[\sum_j \frac{\int p(y_i | \theta_{i,j} = t) p(y_j | \theta_{i,j} = t) p(\theta_{i,j} = t) dt}{\int p(y_i | \theta_i = t_i) p(\theta_i = t_i) dt_i \int p(y_j | \theta_j = t_j) p(\theta_j = t_j) dt_j} \right]^{\frac{1}{h}}. \quad (43)$$

The variable, c , is a normalizing constant, h controls the strength of the filter, θ_n is the underlying parameter of the random variable Y_n which represents the noiseless pixel value in most cases. We fixed the neighborhood size at 3×3 , the search window at 41×41 , and varied h to tune the algorithm.

Multi-Prior Plug and Play 1 (MPP&P 1)

MPP&P 1 uses the approach described in App. A to combine two prior models. BM3D was chosen as \mathcal{H}_1 since it was found to be most effective at reducing background shot noise, and QGGMRF was chosen as \mathcal{H}_2 to promote smoothness in the reconstruction. The ratio of the tuning parameters was fixed at $\beta_1 = 1.5 \times \beta_2$.

Multi-Prior Plug and Play 2 (MPP&P 2)

MPP&P 2 also uses the approach described in App. A to combine two priors. BM3D was again chosen as \mathcal{H}_1 , and BLR was chosen as \mathcal{H}_2 to help reduce speckle variations. The ratio of the tuning parameters was fixed at $\beta_1 = 5 \times \beta_2$.

REFERENCES

1. R. L. Lucke and L. J. Rickard, "Photon-limited synthetic aperture imaging for planet surface studies," *Applied Optics* **41**, 5,084–5,095 (2002).
2. S. M. Beck, J. R. Buck, W. F. Buell, R. P. Dickinson, D. A. Kozlowski, N. J. Marechal, and T. J. Wright, "Synthetic-aperture imaging laser radar: laboratory demonstration and signal processing," *Applied Optics* **44**, 7621–7629 (2005).
3. W. Buell, N. Marechal, J. Buck, R. Dickinson, D. Kozlowski, T. Wright, and S. Beck, "Demonstration of synthetic aperture imaging LADAR," in "Defense and Security," (International Society for Optics and Photonics, 2005), pp. 152–166.
4. B. Krause, J. Buck, C. Ryan, D. Hwang, P. Kondratko, A. Malm, A. Gleason, and S. Ashby, "Synthetic aperture LADAR flight demonstration," in "CLEO: Science and Innovations," (Optical Society of America, 2011), p. PDPB7.

5. J. C. Marron and K. S. Schroeder, "Holographic laser radar," *Optics Letters* **18**, 385–387 (1993).
6. B. D. Duncan and M. P. Dierking, "Holographic aperture LADAR," *Applied Optics* **48**, 1168–1177 (2009).
7. J. Marron, R. Kendrick, S. Thurman, N. Seldomridge, T. Grow, C. Embry, and A. Bratcher, "Extended-range digital holographic imaging," in "SPIE Defense, Security, and Sensing," (International Society for Optics and Photonics, 2010), pp. 76841J–76841J.
8. P. J. Winzer and W. R. Leeb, "Coherent LIDAR at low signal powers: basic considerations on optical heterodyning," *Journal of modern Optics* **45**, 1549–1555 (1998).
9. C. Pellizzari, R. Trahan III, H. Zhou, S. Williams, S. Williams, B. Nemati, M. Shao, and C. A. Bouman, "Synthetic aperture LADAR: A model based approach," Submitted to *IEEE Transactions on Computational Imaging* (TBD).
10. A. Buades, B. Coll, and J.-M. Morel, "A review of image denoising algorithms, with a new one," *Multiscale Modeling & Simulation* **4**, 490–530 (2005).
11. K. Dabov, A. Foi, V. Katkovnik, and K. Egiazarian, "Image denoising by sparse 3-D transform-domain collaborative filtering," *IEEE Transactions on image processing* **16**, 2080–2095 (2007).
12. H. Arsenault and G. April, "Properties of speckle integrated with a finite aperture and logarithmically transformed," *JOSA* **66**, 1160–1163 (1976).
13. H. H. Arsenault and M. Levesque, "Combined homomorphic and local-statistics processing for restoration of images degraded by signal-dependent noise," *Applied Optics* **23**, 845–850 (1984).
14. P. F. Yan and C. Chen, "An algorithm for filtering multiplicative noise in wide range," GRETSI, Saint Martin d'Hères, France (1986).
15. M. Mäkitalo, A. Foi, D. Fevrale, and V. Lukin, "Denoising of single-look SAR images based on variance stabilization and nonlocal filters," in "2010 International Conference on Mathematical Methods in Electromagnetic Theory," (IEEE, 2010), pp. 1–4.
16. C.-A. Deledalle, L. Denis, and F. Tupin, "Iterative weighted maximum likelihood denoising with probabilistic patch-based weights," *IEEE Transactions on Image Processing* **18**, 2661–2672 (2009).
17. C.-A. Deledalle, L. Denis, and F. Tupin, "How to compare noisy patches? Patch similarity beyond gaussian noise," *International journal of computer vision* **99**, 86–102 (2012).
18. S. Parrilli, M. Poderico, C. V. Angelino, and L. Verdoliva, "A nonlocal SAR image denoising algorithm based on LLMMSE wavelet shrinkage," *IEEE Transactions on Geoscience and Remote Sensing* **50**, 606–616 (2012).
19. M. Cetin and W. C. Karl, "Feature-enhanced synthetic aperture radar image formation based on nonquadratic regularization," *Image Processing, IEEE Transactions on* **10**, 623–631 (2001).
20. L. C. Potter, P. Schniter, and J. Ziniel, "Sparse reconstruction for radar," in "SPIE Defense and Security Symposium," (International Society for Optics and Photonics, 2008), pp. 697003–697003.
21. L. C. Potter, E. Ertin, J. T. Parker, and M. Cetin, "Sparsity and compressed sensing in radar imaging," *Proceedings of the IEEE* **98**, 1006–1020 (2010).
22. N. Ö. Onhon and M. Cetin, "Joint sparsity-driven inversion and model error correction for radar imaging," in "Acoustics Speech and Signal Processing (ICASSP), 2010 IEEE International Conference on," (IEEE, 2010), pp. 1206–1209.
23. M. Cetin, I. Stojanovic, N. O. Onhon, K. R. Varshney, S. Samadi, W. C. Karl, and A. S. Willsky, "Sparsity-driven synthetic aperture radar imaging: Reconstruction, autofocus, moving targets, and compressed sensing," *Signal Processing Magazine, IEEE* **31**, 27–40 (2014).
24. S. Ugur and O. Arikan, "SAR image reconstruction and autofocus by compressed sensing," *Digital Signal Processing* **22**, 923–932 (2012).
25. G. Xu, M. Xing, L. Zhang, Y. Liu, and Y. Li, "Bayesian inverse synthetic aperture radar imaging," *Geoscience and Remote Sensing Letters, IEEE* **8**, 1150–1154 (2011).
26. D. Vu, M. Xue, X. Tan, and J. Li, "A Bayesian approach to SAR imaging," *Digital Signal Processing* **23**, 852–858 (2013).
27. L. Zhao, L. Wang, G. Bi, and L. Yang, "An autofocus technique for high-resolution inverse synthetic aperture radar imagery," *Geoscience and Remote Sensing, IEEE Transactions on* **52**, 6392–6403 (2014).
28. W. Su, H. Wang, B. Deng, Y. Qin, and Y. Ling, "High resolution of ISAR imaging based on enhanced sparse Bayesian learning," in "Signal Processing (ICSP), 2014 12th International Conference on," (IEEE, 2014), pp. 2063–2067.
29. W. Su, H. Wang, B. Deng, R. Wang, and Y. Qin, "Sparse Bayesian learning in ISAR tomography imaging," *Journal of Central South University* **22**, 1790–1800 (2015).
30. S. V. Venkatakrishnan, C. A. Bouman, and B. Wohlberg, "Plug-and-play priors for model based reconstruction," in "Global Conference on Signal and Information Processing (GlobalSIP), 2013 IEEE," (IEEE, 2013), pp. 945–948.
31. S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, "Distributed optimization and statistical learning via the alternating direction method of multipliers," *Foundations and Trends® in Machine Learning* **3**, 1–122 (2011).
32. S. Sreehari, S. Venkatakrishnan, B. Wohlberg, L. F. Drummy, J. P. Simmons, and C. A. Bouman, "Plug-and-play priors for bright field electron tomography and sparse interpolation," *arXiv preprint arXiv:1512.07331* (2015).
33. A. Rond, R. Giryes, and M. Elad, "Poisson inverse problems by the plug-and-play scheme," *arXiv preprint arXiv:1511.02500* (2015).
34. S. H. Chan, X. Wang, and O. A. Elgendy, "Plug-and-Play ADMM for image restoration: Fixed point convergence and applications," *arXiv preprint arXiv:1605.01710* (2016).
35. C. A. Bouman, *Model Based Image Processing* (Unpublished, 2015).
36. V. V. Protopopov, *Laser heterodyning*, vol. 149 (Springer, 2009).
37. S. M. Kay, *Fundamentals of Statistical Signal Processing, Detection Theory*, vol. II (Prentice Hall Signal Processing Series, 1993).
38. J. Eckstein and D. P. Bertsekas, "On the douglas—rachford splitting method and the proximal point algorithm for maximal monotone operators," *Mathematical Programming* **55**, 293–318 (1992).
39. M. Hong, Z.-Q. Luo, and M. Razaviyayn, "Convergence analysis of alternating direction method of multipliers for a family of nonconvex problems," *SIAM Journal on Optimization* **26**, 337–364 (2016).
40. C. Forbes, M. Evans, N. Hastings, and B. Peacock, *Statistical distributions* (John Wiley & Sons, 2011).
41. R. Trahan, B. Nemati, H. Zhou, M. Shao, I. Hahn, and W. Schulze, "Low CNR inverse synthetic aperture LADAR imaging demonstration with atmospheric turbulence," in "Defense + Commercial Sensing, 2016 SPIE," (SPIE, 2016).
42. R. Trahan, H. Zhou, B. Nemati, M. Shao, and W. Schulze, "Low cost chirp linearization for long-range isal imaging application," in "Defense + Commercial Sensing, 2016 SPIE," (SPIE, 2016).
43. K. Dabov, "Image and video denoising by sparse 3D transform-domain collaborative filtering," .
44. G. Peyre, "Toolbox image - a toolbox for general purpose image processing," .
45. T. Frese, C. A. Bouman, and K. Sauer, "Multiscale Bayesian methods for discrete tomography," in "Discrete Tomography," (Springer, 1999), pp. 237–264.
46. S. V. Venkatakrishnan, "Model-based iterative reconstruction for micro-scale and nano-scale imaging," Ph.D. thesis, Purdue University (2014).
47. C. A. Bouman, "Digital image processing laboratory: MAP image restoration," <https://engineering.purdue.edu/bouman/grad-labs/MAP-Image-Restoration/> (2010).
48. M. Elad and M. Aharon, "Image denoising via sparse and redundant representations over learned dictionaries," *IEEE Transactions on Image processing* **15**, 3736–3745 (2006).
49. R. Rubinstein, "<http://www.cs.technion.ac.il/~ronrubin/software.html>," .