

Projected Multi-Agent Consensus Equilibrium (PMACE) With Application to Ptychography

Qiuchen Zhai , *Student Member, IEEE*, Gregory T. Buzzard , *Senior Member, IEEE*, Kevin Mertes ,
Brendt Wohlberg , *Fellow, IEEE*, and Charles A. Bouman , *Life Fellow, IEEE*

Abstract—Multi-Agent Consensus Equilibrium (MACE) formulates an inverse imaging problem as a balance among multiple update agents such as data-fitting terms and denoisers. However, each such agent operates on a separate copy of the full image, leading to redundant memory use and slow convergence when each agent affects only a small subset of the full image. In this article, we extend MACE to Projected Multi-Agent Consensus Equilibrium (PMACE), in which each agent updates only a projected component of the full image, thus greatly reducing memory use for some applications. We describe PMACE in terms of an equilibrium problem and an equivalent fixed point problem and show that in most cases the PMACE equilibrium is not the solution of an optimization problem. To demonstrate the value of PMACE, we apply it to the problem of ptychography, in which a sample is reconstructed from the diffraction patterns resulting from coherent X-ray illumination at multiple overlapping spots. In our PMACE formulation, each spot corresponds to a separate data-fitting agent, with the final solution found as an equilibrium among all the agents. Our results demonstrate that the PMACE reconstruction algorithm generates more accurate reconstructions at a lower computational cost than existing ptychography algorithms when the spots are sparsely sampled.

Index Terms—Ptychography, consensus equilibrium, inverse problem, phase retrieval, iterative reconstruction.

I. INTRODUCTION

METHODS for inverse imaging typically seek to balance a fit to noisy data with some form of regularization or prior information. While the fit to data is often well-understood from physical models and can be captured using a log-likelihood

or other penalty function, the prior information is much more difficult to capture in such a functional form. On the other hand, algorithmic denoisers such as BM3D [1] and convolutional neural network (CNN) denoisers [2] encapsulate a great deal of prior information about images but only in an algorithmic way that is not easily amenable to description as a penalty function.

The recently developed Multi-Agent Consensus Equilibrium (MACE) [3] and Plug-and-Play (PnP) [4], [5], [6] methods provide a formulation and algorithms for incorporating both algorithmic prior information and information from penalty functions into a single reconstruction framework. Because the algorithmic prior information is encoded in input-output format as in an image denoiser, there is no cost function associated with this prior information, hence no cost function and no minimization associated with the reconstruction. Instead, MACE formulates the problem as an equilibrium among multiple input-output maps called agents.

The theory underlying MACE implies that PnP methods are equivalent to Bayesian methods when the agents are proximal maps of convex functions (described below), but that PnP methods are more general in that there is no corresponding natural optimization problem when the PnP agents are not proximal maps, as in the case of neural networks. Moreover, MACE generalizes PnP to incorporate more than 2 agents, which allows the problem to be divided into smaller subproblems and solved separately with distributed implementation. However, the original MACE framework requires that each agent maintain a separate copy of the entire reconstruction, which is very inefficient when each agent operates on only a small portion of the reconstruction.

Ptychography is an important computational imaging technique [7], [8] that is widely used in applications such as the imaging of manufactured nanomaterials [9], [10], [11], [12]. The imaging technique works by moving a coherent X-ray probe across an object plane to scan a sample in overlapping patches. A detector then measures the magnitude of the far-field diffraction pattern at each probe location. The magnitude and phase of the transmittance image of the sample can be recovered by exploiting the redundancy caused by the overlaps in the intensity measurements. Notably, the reconstructed phase is significant in estimating the structure of the sample.

Since ptychography is a phase recovery problem, it can be solved using conventional phase recovery methods [13]. However, a series of methods have been developed specifically for reconstruction of ptychographic data. The original method, known as ptychographic iterative engine (PIE) [14], has been

Manuscript received 27 March 2023; revised 11 September 2023; accepted 13 October 2023. Date of publication 30 October 2023; date of current version 17 November 2023. The work of Qiuchen Zhai and Brendt Wohlberg was supported by the LDRD Program of LANL under Grant 20200061DR. The work of Gregory T. Buzzard and Qiuchen Zhai was supported by NSF under Grant CCF-1763896. The work of Kevin Mertes was supported by the LDRD Program of LANL under Grants 20230070DR and 20200061DR. The work of Charles A. Bouman was supported by the U.S. Department of Energy and the Showalter Trust. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Mauricio Delbracio. (*Corresponding author: Qiuchen Zhai.*)

Qiuchen Zhai and Charles A. Bouman are with the School of Electrical and Computer Engineering, Purdue University, West Lafayette, IN 47907 USA (e-mail: qzhai@purdue.edu; bouman@purdue.edu).

Gregory T. Buzzard is with the Department of Mathematics, Purdue University, West Lafayette, IN 47907 USA (e-mail: buzzard@purdue.edu).

Kevin Mertes is with Physical Chemistry and Applied Spectroscopy Group, Los Alamos National Laboratory, Los Alamos, NM 87545 USA (e-mail: kmertes@lanl.gov).

Brendt Wohlberg is with Theoretical Division, Los Alamos National Laboratory, Los Alamos, NM 87545 USA (e-mail: brendt@ieee.org).

This article has supplementary downloadable material available at <https://doi.org/10.1109/TCI.2023.3328288>, provided by the authors.

Digital Object Identifier 10.1109/TCI.2023.3328288

generalized to methods such as extended PIE (ePIE) [15], regularized PIE (rPIE) [16], and momentum-accelerated PIE (mPIE) [17]. All these methods tend to have rapid convergence and work by serially enforcing local intensity measurements. However, since the PIE-type algorithms are based on serial updates of the image at each probe location, the PIE methods cannot be easily parallelized, and the methods also tend to require a large overlap between scan positions to achieve a high-quality reconstruction [18].

In addition, a variety of gradient-based methods including Wirtinger Flow (WF) [19], accelerated WF [20], Maximum Likelihood method [21], as well as other approaches based on proximal algorithms [22], [23] have also been proposed for ptychographic reconstruction. While methods based on gradient calculations can be parallelized across probe locations, these methods tend to have slower convergence than serial methods and are limited by the requirement that the problem must be formulated as an optimization.

More recently, a method known as scalable heterogeneous adaptive real-time ptychography (SHARP) has been proposed [24] and implemented [25] for ptychographic reconstruction that is intrinsically parallel in nature. The SHARP method is based on a general method for phase-retrieval known as relaxed averaged alternating reflections (RAAR) [26]. However, SHARP and RAAR are designed specifically for phase retrieval and not more general inverse problems.

In this article, we introduce Projected Multi-agent Consensus Equilibrium (PMACE) an intrinsically parallel algorithm for solving distributed inverse problems. PMACE extends the MACE framework to allow the state of a problem to be updated in smaller components that can be processed in parallel by localized agents. For some applications such as ptychography, PMACE can dramatically reduce memory requirements relative to MACE because each agent operates in parallel while storing only a small portion of the entire image. We also show that PMACE is equivalent to an optimization formulation for some choices of update agents, but that in most cases there is no naturally corresponding optimization problem, with the solution corresponding instead to a weighted equilibrium of the updates of the individual agents.

To illustrate the utility of the PMACE framework, we apply it to the problem of ptychographic reconstruction. We use the data from each scan position to implement a computationally efficient, non-optimization-based data-fitting agent that updates a local patch to better fit the measured intensity data for a single diffraction spot. These local patches are combined into a full image with a weighted average that uses the probe intensity to capture the uncertainty in measurement at a given location.

In contrast to ePIE, the inherent parallelism of individual agents allows the PMACE algorithm for ptychography to make parallel updates to the image at different probe locations. Therefore, as with SHARP, the PMACE algorithm is more suitable than PIE-type algorithms for implementation on parallel or distributed computers. However, we leave the exploration of efficient parallel implementation for future work.

Our experimental results demonstrate that PMACE outperforms existing state-of-the-art ptychographic reconstruction

algorithms in both convergence speed and reconstruction quality as the ptychographic spots are spaced further apart. In practice, reconstruction from sparse samples is of practical importance because it can be used to reduce the amount of data required for accurate ptychographic reconstruction.

II. MACE AND PMACE

A. MACE

Multi-Agent Consensus Equilibrium (MACE) [3] is a problem formulation that reconciles multiple agents, such as data fitting updates or denoisers, each acting to improve a candidate reconstruction. In this formulation, each agent, F_j , maintains an individual copy, x_j , of the full image (or volume), and $F_j(x_j)$ is an improved reconstruction according to the j th agent. These x_j are either real or complex vectors with a dimension appropriate to the problem, and $F_j(x_j)$ has the same dimension. We leave the dimensions implicit when they are clear from the context. These copies are stacked to form a full MACE state, given by

$$\mathbf{x} = [x_0, \dots, x_{J-1}]$$

and the agents are stacked to form a single operator defined by

$$\mathbf{F}(\mathbf{x}) = [F_0(x_0), \dots, F_{J-1}(x_{J-1})]. \quad (1)$$

Note that \mathbf{x} denotes the full MACE state, while x_j denotes the j th component image. Since each component is a full copy of the image to be reconstructed, the MACE state contains multiple, potentially inconsistent reconstructions. In order to produce a single MACE reconstruction, we define the operator

$$\mathbf{G}(\mathbf{x}) = [\bar{x}(\mathbf{x}), \dots, \bar{x}(\mathbf{x})], \quad (2)$$

where $\bar{x}(\mathbf{x}) = (1/J) \sum_{j=0}^{J-1} x_j$. So $\mathbf{G}(\mathbf{x})$ computes the average of the components in \mathbf{x} , and then returns a state vector formed by replicating this average J times. A key algebraic property of the averaging operator is that $\mathbf{G}(\mathbf{G}(\mathbf{x})) = \mathbf{G}(\mathbf{x})$.

The MACE equation is then given by

$$\mathbf{F}(\mathbf{x}^*) = \mathbf{G}(\mathbf{x}^*), \quad (3)$$

where \mathbf{x}^* solves the equation, and the final reconstruction is then given by $x^* = \bar{x}(\mathbf{x}^*)$. As shown in [3], (3) has the interpretation of finding a consensus equilibrium among all the agents, in that the output of each agent is the common image $\bar{x}(\mathbf{x}^*)$, while the updates $F_j(x_j^*) - x_j^*$ sum to 0.

In the case that each F_j is the proximal map for a convex, differentiable¹ function f_j , then x^* satisfies

$$\sum_{j=0}^{J-1} \nabla f_j(x^*) = 0, \quad (4)$$

and x^* is a minimizer of the function $f = \sum_j f_j$ [27].

However, when the MACE agents are not proximal maps, there may be no cost function that is minimized by solving the MACE equation. We discuss this in greater generality in Section II-B.

¹We assume differentiability for simplicity. Similar results hold for general convex functions by using subgradients.

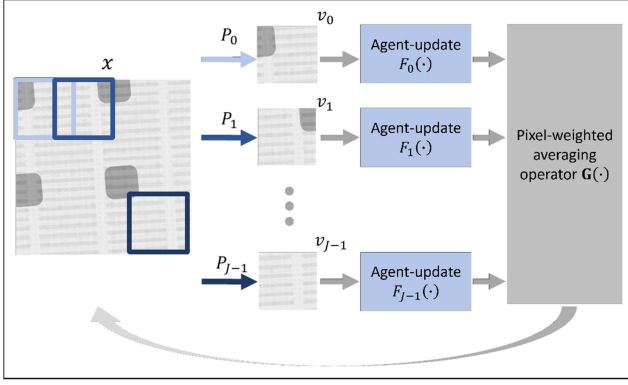


Fig. 1. Conceptual rendering of the PMACE pipeline. A reconstruction x is subdivided into possibly overlapping components, v_j , and distributed to multiple agents, F_j , each of which acts to improve this local reconstruction. These local reconstructions are combined using a pixel-weighted averaging operator, G , that reassembles the local components into a consistent global reconstruction.

We note that $\mathbf{F}(\mathbf{x})$ is an intrinsically parallel operation that can be efficiently distributed across many compute nodes, while the averaging operator $\mathbf{G}(\mathbf{x})$ requires communication across all nodes in order to first gather the individual components, then compute their average, and then broadcast the average back to individual nodes [27].

B. PMACE

In some problems, agents operate naturally on subsets or projections of a full image rather than on the image itself. Fig. 1 illustrates such a problem in which separate data is collected from each patch of an image, x . In this case, the agents can operate separately to make each component more consistent with its data. However, these individual components must also be reconciled into a single consistent reconstruction.

Mathematically, we express this extraction of a patch or component from an image as

$$v_j = P_j x$$

where P_j is a matrix with orthonormal rows, so that $P_j P_j^T = \mathbf{I}$. We use v_j to emphasize that these components lie in a different space than x . Typically, P_j simply selects a region of the image x , but more general transformations are allowed. To capture spatial information and relevant features without excessive computational cost, the selection of P_j should be based on the size of local measurements associated with individual components.

As in MACE, each component, v_j , is processed by a corresponding agent, F_j , whose role is to improve the component by better fitting to data, reducing noise, etc. Also as in MACE, PMACE stacks the component states and the agents to form the vector and operator given by

$$\mathbf{v} = [v_0, \dots, v_{J-1}], \quad (5)$$

$$\mathbf{F}(\mathbf{v}) = [F_0(v_0), \dots, F_{J-1}(v_{J-1})]. \quad (6)$$

Note that PMACE has a potentially huge advantage over MACE because it does not require J replicates of the entire

reconstruction. This can dramatically reduce memory and computation resources when implemented on a computer. For example, in ptychography as described below, if each reconstruction is a large image and there are 1000 agents corresponding to 1000 probe locations, then MACE becomes impractical to implement since it requires that the full image be replicated 1000 times in the computer's memory.

Since the PMACE components are not complete versions of the reconstruction, the averaging operator $\mathbf{G}(\mathbf{x})$ must be redefined. To do this, we project the components back to the full image, then use a pixel-wise weighted average to combine them and reproject them back to components. This leads to the new averaging operator

$$\mathbf{G}^P(\mathbf{v}) = [P_0 \bar{x}(\mathbf{v}), \dots, P_{J-1} \bar{x}(\mathbf{v})]. \quad (7)$$

where

$$\bar{x}(\mathbf{v}) = \Lambda^{-1} \sum_{j=0}^{J-1} P_j^T W v_j \quad (8)$$

$\Lambda = \sum_{j=0}^{J-1} P_j^T W P_j$, and W is a positive definite and (typically) diagonal weighting matrix. We assume that Λ is invertible, so if P_j selects subsets of the original image, then the set of J components must cover the entire image being reconstructed. The weight matrix, W , can be used to model the uncertainty in measurements for each component, with its selection depending on the nature of imaging modality and the noise characteristics in the data.

Using these operators, we look for a solution \mathbf{v}^* to the PMACE equation, defined as

$$\mathbf{F}(\mathbf{v}^*) = \mathbf{G}^P(\mathbf{v}^*) \quad (9)$$

in which case $x^* = \bar{x}(\mathbf{v}^*)$ is the final reconstruction. In words, (9) means that each agent generates a patch that is consistent with the entire reconstruction x^* .

C. Solving the PMACE Equations

To solve the PMACE (9), we follow [3] to show that it can be reformulated as a fixed point problem and apply Mann iterations to compute a solution.

By definition of \bar{x} , \mathbf{G}^P , and Λ , we have

$$\begin{aligned} \bar{x}(\mathbf{G}^P(\mathbf{v})) &= \Lambda^{-1} \sum_j P_j^T W P_j \bar{x}(\mathbf{v}) \\ &= \bar{x}(\mathbf{v}), \end{aligned}$$

which gives

$$\mathbf{G}^P(\mathbf{G}^P(\mathbf{v})) = \mathbf{G}^P(\mathbf{v}). \quad (10)$$

Thus \mathbf{G}^P , like the original averaging operator \mathbf{G} , is a linear projection. With this, a further calculation shows that

$$(2\mathbf{G}^P - \mathbf{I})^{-1} = (2\mathbf{G}^P - \mathbf{I}). \quad (11)$$

From this we obtain the equivalent fixed point formulation of (9) as

$$(2\mathbf{G}^P - \mathbf{I})(2\mathbf{F} - \mathbf{I})(\mathbf{v}^*) = \mathbf{v}^* \quad (12)$$

or $\mathbf{T}\mathbf{v}^* = \mathbf{v}^*$, where $\mathbf{T} = (2\mathbf{G}^P - \mathbf{I})(2\mathbf{F} - \mathbf{I})$.

If \mathbf{T} is non-expansive and has a fixed point, then Mann iterations with $\rho \in (0, 1)$, defined as

$$\mathbf{v} \leftarrow (1 - \rho)\mathbf{v} + \rho\mathbf{T}\mathbf{v} \quad (13)$$

are guaranteed to converge to a fixed point of \mathbf{T} and hence to a solution of (9). The parameter ρ is a kind of step size; its choice impacts the convergence rate and the stability of the iterations but not the final solution.

D. PMACE for Optimization and Beyond

A natural question is whether PMACE corresponds to the solution of an optimization problem. We show in Theorem II.2 that the answer is affirmative in certain special cases, but not in general. (Proofs are given in the Appendix.)

To understand the relationship between equilibrium and optimization, suppose that each F_j is the proximal map for a real-valued, differentiable, convex function $f_j(v_j)$:

$$F_j(v_j) = \arg \min_v \left\{ f_j(v) + \frac{1}{2\sigma^2} \|v - v_j\|^2 \right\}. \quad (14)$$

In this setting, a PMACE solution is equivalent to a zero of a naturally associated vector field.

Theorem II.1: Let $F_j, j = 0, \dots, J-1$ denote the proximal map function of a differentiable convex function f_j as specified in (14). Then \mathbf{v}^* is a solution to the PMACE equation in (9) if and only if $x^* = \bar{x}(\mathbf{v}^*)$ satisfies

$$\sum_{j=0}^{J-1} P_j^T W \nabla f_j(P_j x^*) = 0. \quad (15)$$

Note that (15) captures the idea of equilibrium in that the sum of the J vectors is 0. However, as described in more detail below, because of the weight matrix W , this vector field may not be a gradient field, hence may not correspond naturally to a minimization problem. To explain this more fully, we first recall some basic results on vector fields.

We say that $V(x)$ is a conservative vector field if it is the gradient of some potential function $f(x)$ [28, Section 9.2A]. In this case, the condition $V(x^*) = 0$ is a first-order necessary condition for the minimization of f . Moreover, a continuously differentiable vector field, $V(x)$, is conservative if and only if its Jacobian is self-adjoint (or symmetric) [28, Section 9.2C].

When each f_j is twice continuously differentiable, the Jacobian of the vector field in (15) (as a function of x in place of x^*) is

$$\sum_{j=0}^{J-1} P_j^T W H_{f_j}(P_j x) P_j \quad (16)$$

where $H_{f_j}(P_j x)$ is the symmetric Hessian matrix of 2nd order partial derivatives of f_j , evaluated at $P_j x$. If W is a multiple of the identity, then the matrix in (16) is symmetric, so the PMACE solution corresponds naturally to a minimization problem. However, this symmetry is not robust in that nonzero mixed partial derivatives in f_j together with distinct entries in the diagonal

matrix W will destroy this symmetry. This leads to the following theorem.

Theorem II.2: Let f_j and F_j be as in Theorem II.1.

- 1) If $W = rI$ for scalar r , then PMACE is equivalent to optimization in that \mathbf{v}^* is a solution to the PMACE equation in (9) if and only if $x^* = \bar{x}(\mathbf{v}^*)$ satisfies

$$x^* = \arg \min_x \left\{ \sum_{j=0}^{J-1} f_j(P_j x) \right\}. \quad (17)$$

- 2) For generic diagonal W and generic convex f_j , the PMACE formulation does not naturally arise as an optimization problem. That is, the vector field in (15) is not a conservative vector field, and hence is not the gradient field of a potential function.

In other words, even when the agents F_j are proximal maps, if the weight matrix W is not a multiple of the identity, then the natural vector field defining the PMACE solution is not a gradient field, so the PMACE solution is not naturally given by minimizing a cost function. Of course, it's possible to set up an optimization problem with the same solution by minimizing the norm of the vector field in (19), but this doesn't correspond to a consensus minimization problem as in (17).

When the f_j are convex but not everywhere differentiable, then the gradient ∇f_j in (15) is replaced with the possibly set-valued subdifferential ∂f_j , and the search for an x^* to give a zero of the left-hand side becomes the search for an x^* so that the left-hand side contains the 0 vector, which is known as a monotone inclusion problem. Results analogous to Theorem II.2 can be obtained in this setting but are not pursued here. For more about the relationship between convex optimization and monotone operators, see [29].

Non-proximal equilibrium: To extend to the case in which the agents F_j are not proximal maps, we note that the proximal map in (14) can be rewritten [30, Cor. 17.6] as $F_j(v_j) = v_j - \sigma^2 \nabla f_j(F_j(v_j))$. That is, a proximal map for a convex, differentiable function is an implicit gradient descent step in the sense that the gradient is evaluated at the endpoint rather than the starting point. With this and the fact that $F_j(v_j^*) = P_j x^*$, (15) can be rewritten as

$$\sum_{j=0}^{J-1} P_j^T W (F_j(v_j^*) - v_j^*) = 0. \quad (18)$$

The vector field in this case is defined not on an image x but on a stack of patches as in (5). In order to use this formulation to give a single reconstruction, x^* , we note that each v_j^* must satisfy the constraint $F_j(v_j^*) = P_j x^*$. We use this relationship to implicitly define a set-valued inverse $F_j^{-1}(P_j x)$ that maps x to all patches v_j consistent with x .

With this, we reformulate (18) and hence the PMACE equation in (9) as an inclusion problem.

Theorem II.3: Define the set-valued vector field

$$V(x) = \sum_{j=0}^{J-1} P_j^T W (P_j x - F_j^{-1}(P_j x)). \quad (19)$$

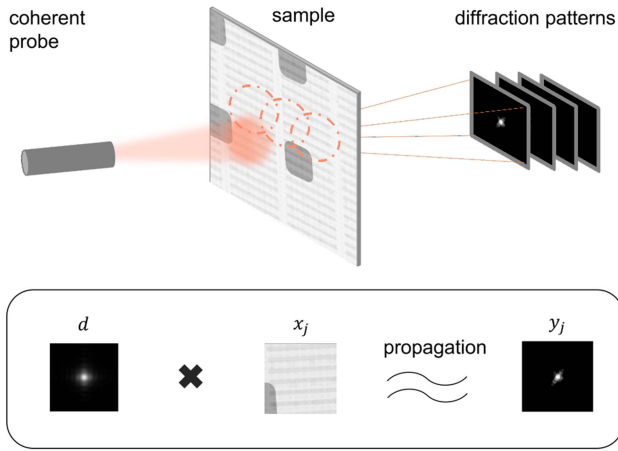


Fig. 2. In ptychographic imaging, a coherent probe is used to illuminate a sample at multiple, partially overlapping regions. The intensities of the resulting diffraction patterns are recorded by the detector.

Then the solution $x^* = \bar{x}(\mathbf{v}^*)$ of the PMACE equation in (9) is the solution of the inclusion problem

$$0 \in V(x^*). \quad (20)$$

We call this a consensus equilibrium problem because the common point x^* yields a set of vectors v_j^* with $F_j(v_j^*) = P_j x^*$ (consensus) such that the weighted average updates give the zero vector as in (18).

The verification that a given operator is monotone is often technically involved, so we do not address that here. Likewise, further conditions under which (19) corresponds to an optimization problem are beyond the scope of this article.

Finally, we note that, as with MACE or PnP [31], PMACE can incorporate regularization by either using a denoiser as one of the agents, or using a denoiser as part of the averaging operator, $\mathbf{G}(\mathbf{x})$. While this could potentially improve results for ptychography reconstruction, we leave this topic as an area for future research.

III. PMACE FOR PTYCHOGRAPHY

As noted in the introduction and illustrated in Fig. 2, a ptychographic reconstruction relies on far-field intensity measurements (also known as diffraction patterns) obtained by illuminating a sample with a coherent X-ray source at multiple locations, with significant overlap between adjacent illuminated regions. The goal of these measurements is to reconstruct the full complex transmittance image of the scanned sample. The phase of this complex transmittance image is particularly important since it gives precise information about the thickness of the sample at each location.

A potential point of confusion in this problem is that there are two sets of phases that must be estimated for this reconstruction. One set is the phase of the complex transmittance image, which is the primary quantity of interest because it gives very precise, quantitative information about the material properties of the sample. The other set contains the phases associated with the far-field intensity data. These far-field phases are essentially

the phases of the Fourier transform of the illuminated sample. Since the detector measures intensities only, the far-field phases must be estimated using the redundancies associated with overlapping intensity measurements – this is known as the phase-retrieval problem. However, for our purposes, the far-field phases are really just nuisance parameters that must be estimated in order to estimate the phase and magnitude of the complex transmittance.

To apply PMACE to the ptychography problem, we define each data-fitting agent to be a local update to better match the measurements of the far-field Fraunhofer diffraction pattern in a single patch. Our approach does not use a penalty function/proximal map formulation but rather defines an update agent for each patch update directly from data using a weighted average between a previous estimate and new estimate that fits the available data. This approach has the advantages of being computationally efficient and easily adjusted to account for the signal-to-noise ratio of the data.

While probe estimation is an important component of ptychography, we assume for clarity in this article that an accurate estimate of the probe function is known and consistent at all scanning positions. In a future paper we will apply PMACE to estimate both the probe and complex image simultaneously.

A. Forward Model

We let $x \in \mathbb{C}^{N_1 \times N_2}$ denote the complex material transmittance of the target object and seek to recover x from real-valued measurements $y_j \in \mathbb{R}^{N_p \times N_p}$ of the form

$$y_j = \sqrt{\text{Pois}(|\mathcal{F}DP_j x|^2)}. \quad (21)$$

Here \mathcal{F} denotes the 2D orthonormal discrete Fourier transform, $D = \text{Diag}(d)$ is a diagonal matrix representing the probe's complex illumination function d , and $P_j : \mathbb{C}^{N_1 \times N_2} \rightarrow \mathbb{C}^{N_p \times N_p}$ is the linear operator representing the extraction of an $N_p \times N_p$ patch corresponding to the j^{th} scan position, $j = 0, \dots, J-1$. Thus $\mathcal{F}DP_j x$ is the mean complex far-field diffraction pattern associated with the j^{th} scan position. We assume that the physical sensor measurements can be modeled by a vector of independent Poisson distributed random variables denoted by $\text{Pois}(\lambda)$ with mean λ given by the pixel-wise intensity of this diffraction pattern, and that these measurements are pre-processed by taking a square-root, which serves as a variance-stabilizing transform for the Poisson distribution [32].

B. PMACE Formulation

In order to use the PMACE algorithm, we will need to select a data-fitting agent, $F_j(v_j)$. The proximal map is not a good choice because its evaluation requires the solution of a non-convex optimization problem, which is computationally expensive and not guaranteed to achieve a global minimum. Fortunately, the PMACE method is not restricted to the use of proximal maps. So instead, our strategy will be to design a data-fitting agent that moves the solution in a direction that better fits the data.

Using this approach, we define the data-fitting agents $F_j(v_j)$ by taking a weighted average between a nonlinear projection

onto measured data and an input point. The nonlinear projection begins with a complex-valued image patch, $v_j = P_j x$, and isolates the phase in diffraction space by $\mathcal{F}Dv_j/|\mathcal{F}Dv_j|$ (the phase can be set arbitrarily if the magnitude is 0). This phase is multiplied point-wise by the amplitudes y_j , and then the result is inverted to return to image space.

Mathematically, we represent the full nonlinear projection from an input image patch to one that matches the given amplitudes in diffraction space is given by $D_\epsilon^{-1} \mathcal{F}^* (y_j \frac{\mathcal{F}Dv_j}{|\mathcal{F}Dv_j|})$ where $D_\epsilon^{-1} = \text{Diag}(d_\epsilon^{-1})$ is a numerically stable inverse of D .²

For the full data-fitting agent, we average this projection with the input point to obtain

$$F_j(v_j) = (1 - \alpha)v_j + \alpha D_\epsilon^{-1} \mathcal{F}^* \left(y_j \frac{\mathcal{F}Dv_j}{|\mathcal{F}Dv_j|} \right) \quad (22)$$

where α controls the strength of fit to data. In contrast with ρ in (13), where ρ is employed to optimize convergence speed, α regulates regularization, with the ratio $\alpha/(1 - \alpha)$ provides a rough interpretation of the signal-to-noise ratio of the data.

This interpolation between input and a data-fitting point is similar in spirit to a data-fitting proximal map. However, the Jacobian of (22) is not symmetric, so F_j is not a proximal map (where we convert complex images to real-valued vectors by stacking real and imaginary parts in order to take real derivatives). In fact, Theorem II.3 implies that the PMACE formulation using these F_j does not naturally correspond to a minimization problem. In the appendix, we show that F_j is invertible, which implies that $V(x)$ in (19) is a single-valued vector field, which is nonconservative since the Jacobian of (22) is not symmetric. Hence, as in Theorem II.2, this formulation is not naturally a minimization problem.

We use these F_j to define the vectorized operator \mathbf{F} as in (6). For \mathbf{G} in (7), we model the uncertainty arising from the point-to-point variation of probe illumination by using the weight matrix

$$W = |D|^\kappa \quad (23)$$

to define Λ and \bar{x} . Here $\kappa \in [1, 2]$ can be used to control the emphasis between brightly lit points near the center of the probe (large κ) versus dimly lit points near the edges of the probe (small κ). In our experiments, $\kappa \in [1.25, 1.5]$ gives the fastest convergence and best reconstruction quality; we display all results using $\kappa = 1.5$.

The PMACE solution is then determined by solving $\mathbf{F}(\mathbf{x}) = \mathbf{G}(\mathbf{x})$ or by iterating (13) to obtain a fixed point of \mathbf{T} . This has the effect of balancing the updates of the data fitting agents for each patch, subject to constructing an image that agrees with each individual patch.

Proper initialization of the PMACE algorithm is important for fast convergence to a good solution. This is particularly true since phase retrieval corresponds to a non-convex optimization problem, so that solutions can in general depend on the initial condition. We have found the following initialization to work

²We compute the numerically stable inverse of the diagonal entries as $d_\epsilon^{-1} = d^* / (|d|^2 + \epsilon)$ where $\epsilon = 10^{-6} \sqrt{\|d\|^2 / \dim(d)}$.

Algorithm 1: Mann Iteration for Computing PMACE Solution.

Input: Initialization: $x^{(0)} \in \mathbb{C}^{N_1 \times N_2}$; $\kappa \in [1.25, 1.5]$

Output: Final Reconstruction: $\hat{x} \in \mathbb{C}^{N_1 \times N_2}$

$\mathbf{w} = \mathbf{v} = [v_0^{(0)}, \dots, v_{J-1}^{(0)}]$, where $v_j^{(0)} = P_j x^{(0)}$

while not converged **do**

$\mathbf{w} \leftarrow \mathbf{F}(\mathbf{v})$

$\mathbf{z} \leftarrow \mathbf{G}(2\mathbf{w} - \mathbf{v})$

$\mathbf{v} \leftarrow \mathbf{v} + 2\rho(\mathbf{z} - \mathbf{w})$

end while

return $\hat{x} = \Lambda^{-1} \sum_{j=0}^{J-1} P_j^T |D|^\kappa v_j$

well in practice,

$$x^{(0)} \leftarrow \Lambda_0^{-1} \sum_{j=0}^{J-1} P_j^T \left(\frac{\|y_j\|}{\|d\|} \mathbf{1} \right) \quad (24)$$

where $\Lambda_0 = \sum_{j=0}^{J-1} P_j^T P_j$ and $\mathbf{1}$ is a column vector of 1 s. While simple, this creates an initial image with the correct scale at each probe location based on the relative magnitude of the probe and measured signal.

The resulting algorithm has similarities to SHARP [25] but also significant differences. First, (22) incorporates the data projector of [25], but (22) is in the image domain and averaged with the identity, as opposed to the SHARP data projector, which is in the so-called frame domain (probe times image) and which is not averaged. The averaging factor α replaces the stabilization factor β in the data projector of SHARP. More significantly, the averaging operator P_Q of SHARP uses a fixed weighting that is essentially equivalent to taking $\kappa = 2$ in (23), which tends to underweight the contributions of locations with small to moderate probe intensities.

The pseudo-code for computing the PMACE solution is shown in Algorithm 1. The algorithm starts with an initial guess of the complex transmittance of the unknown object $x^{(0)} \in \mathbb{C}^{N_1 \times N_2}$, and then constructs the stacked projections. Mann iterations are then used to revise the estimates of individual projections. The final reconstruction of the complex object is given by taking the weighted average of the estimates of individual projections.

IV. EXPERIMENTAL RESULTS

In this section, we present the results of our approach on both synthetic and measured data and compare with reconstruction results using state-of-the-art algorithms. Our results indicate that PMACE matches the best-competing algorithms in convergence speed and image quality for small probe spacing, and outperforms the competitors as the probe spacing between neighboring scan points increases.³

A. Synthetic Data Generation

Fig. 3 shows the complex-valued ground truth object image and probe used for our simulation experiments. A synthetic

³The code for PMACE is available at https://github.com/cabouman/ptycho_pmace.




Fig. 3. Ground truth images used in simulation experiments. The complex ground-truth object's (a) magnitude and (b) phase in radians. The complex probe's (c) magnitude and (d) phase in radians.

complex transmittance image of size 800×800 pixels was generated by modeling a 5-layer composite material.

The complex probe of size 256×256 pixels was simulated with a photon energy of 8.8 keV.

To simulate measurements for a single probe location, we multiply the projection of the complex transmittance image by the complex probe, then take the squared modulus of the Fourier transform to simulate forward propagation. We scale to simulate a peak photon rate, add a factor for dark current, then use this as a mean for a sample from a Poisson distribution. With r_p as the peak photon detection rate and λ as the mean dark current, we then have that

$$\hat{y}_j \leftarrow \sqrt{\text{Pois} \left(r_p \frac{|\mathcal{F}Dx_j|^2}{\max_k (\|\mathcal{F}Dx_k\|_\infty^2)} + \lambda \right)} \quad (25)$$

where $\|\cdot\|_\infty$ is the maximum magnitude of its argument, $\max_k(\cdot)$ denotes the maximum value over all k , and the operations are performed pointwise. Assuming a photon detector with 14-bit dynamic range and the presence of a half-bit of dark current, $r_p = 10^4$ and $\lambda = 0.5$ were used to simulate noisy data.

To simulate data from a full experiment, we first select a nominal distance (in the range of 20 to 76 pixels) between nearest-neighbor probe locations. We then use a randomly perturbed rectangular grid for the probe sampling pattern with random offsets selected uniformly in $[-5, 5]$ (in units of pixels). This has the effect of more realistically simulating actual ptychography experiments while also avoiding periodic reconstruction artifacts [15].

B. Experimental Methods for Synthetic Data Case

The key to estimating phase information from intensity measurements is the redundancy inherent in using multiple overlapping illuminations. An increase in probe spacing with a fixed

probe size will reduce image reconstruction quality [33], while using a larger probe will compensate for larger probe spacing [34]. Therefore, a measure of illumination overlap is more relevant for reconstruction quality than simple probe spacing.

To quantify this overlap, we define the overlap ratio between adjacent probes indexed by j and k as

$$r_{j,k} = \frac{\|P_j^T |D| \odot P_k^T |D|\|_1}{\||D| \odot |D|\|_1} \quad (26)$$

where \odot denotes point-wise multiplication and $\|\cdot\|_1$ denotes the l_1 -norm. We then define the overlap ratio r_{ovlp} of the full grid scan pattern as the average of the overlap ratios between paired adjacent probes

$$r_{\text{ovlp}} = \frac{1}{N} \sum_{\{(j,k), j \neq k\}} r_{j,k} \quad (27)$$

where N denotes the number of pairs of adjacent scan positions. The overlap ratio is near 1 when adjacent probes almost fully overlap and near 0 when they are almost disjoint.

To evaluate image reconstruction quality for the synthetic data case, we assume a known complex probe (we will consider joint image-probe estimation in a future paper). Given the complex probe profile in Fig. 3, we generated synthetic data using fixed (simulated) probe spacing of 68 pixels with a corresponding overlap ratio of $r_{\text{ovlp}} \approx 14.5\%$.

All reconstructions are performed using the ePIE [15], AWF [20], SHARP [25], and PMACE algorithms. To provide a fair comparison, we optimized one algorithmic parameter using grid-search for each method: the step size of ePIE, the relaxation parameter of SHARP, and the data-fitting parameter α for PMACE. We ran each method for 100 iterations. Each iteration for each method involves 2 FFTs per probe location, and these FFTs are the dominant computational cost for each method, so 100 iterations represent essentially the same amount of computation for each method.

In order to evaluate reconstruction quality, we used the Normalized Root-Mean-Square Error (NRMSE) between each reconstructed image \hat{x} and the ground truth image x . Since ptychography data is invariant to a constant phase shift in the image domain, we computed the NRMSE as

$$\text{NRMSE} = \min_c \frac{\|c\hat{x} - x\|}{\|x\|} \quad (28)$$

where c accounts for an unknown phase shift and possible gain.

C. Reconstruction Results on Synthetic Data

Figs. 4 and 5 show the results of the ePIE, AWF, SHARP, and PMACE algorithms. Fig. 4 shows both the phase of the reconstruction in radians and the error in the phase as compared to the ground truth, and it also lists the final NRMSE of the reconstruction. Fig. 5 shows the corresponding images for the magnitude of the reconstruction.

Since the probe spacing is relatively large and the probe overlap is small ($r_{\text{ovlp}} \approx 14.5\%$), ePIE cannot reconstruct the complex transmittance image accurately. AWF and SHARP do significantly better, but the difference images indicate that both

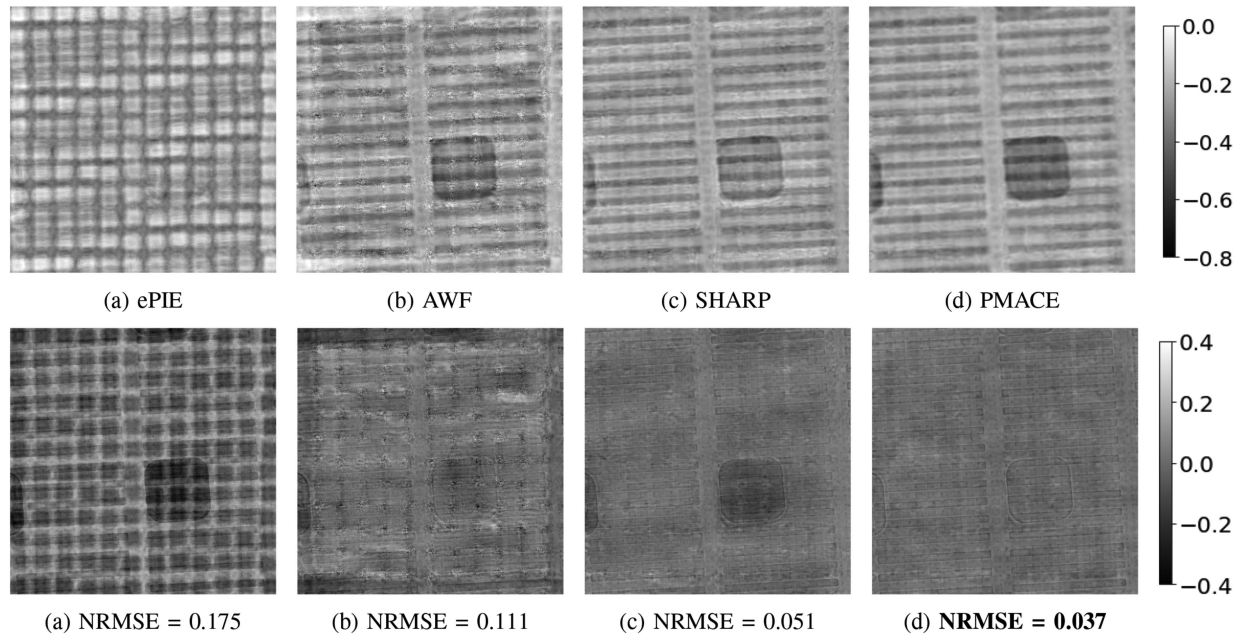


Fig. 4. Top: **Phase** (in radians) of the reconstructed complex transmittance images in Fig. 3 from synthetic data with a small overlap: $r_{\text{ovlp}} \approx 14.5\%$. Bottom: Difference between the reconstructed and ground truth phase, with associated NRMSEs in the subcaptions. Note that transmittance phase is typically more important than magnitude in applications.

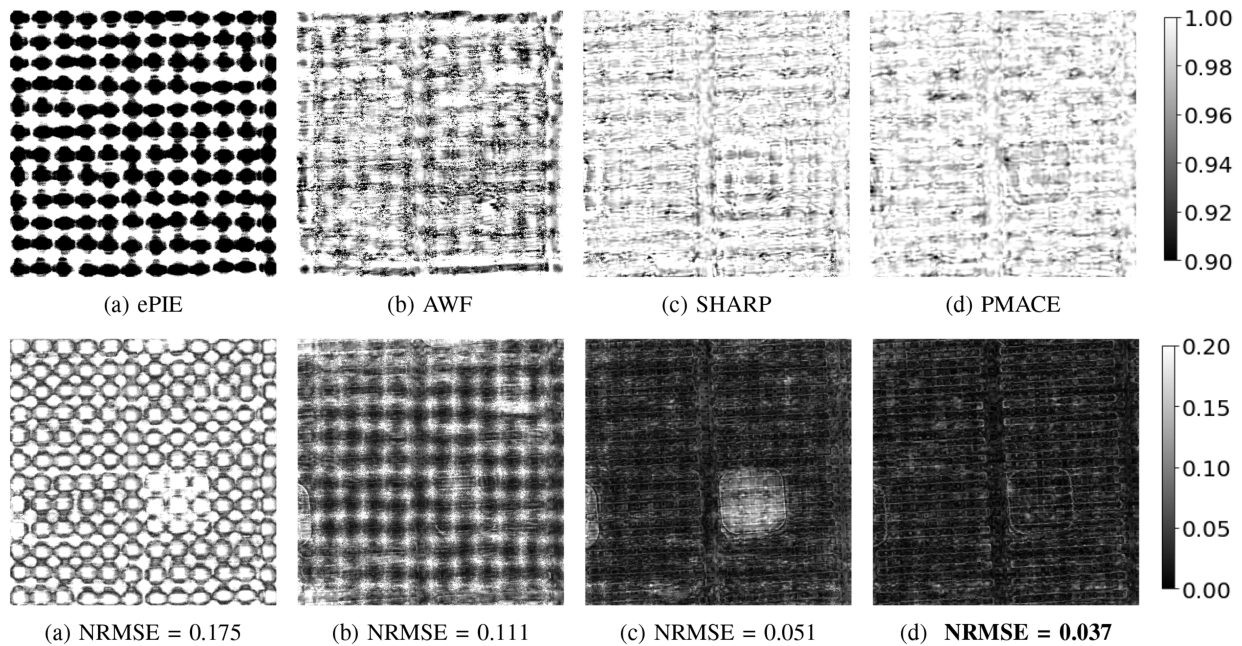


Fig. 5. Top: **Magnitude** of the reconstructed complex transmittance images in Fig. 3 from synthetic data with a small overlap: $r_{\text{ovlp}} \approx 14.5\%$. Bottom: Amplitudes of error between the complex reconstructions and ground truth, with NRMSEs in the subcaptions. Note that transmittance phases (Fig. 4) are typically more important in application than magnitudes.

of these methods introduce artifacts not found in the ground truth. Overall, PMACE reduces these artifacts and produces the best result both visually and based on the NRMSE. This is particularly relevant for the phase images, since transmittance phase is typically more important in applications than magnitude.

Fig. 6 shows the convergence of the NRMSE as a function of the number of iterations. Consistent with Figs. 4 and 5, the

NRMSE associated with ePIE actually increases from the initial condition and does not converge to a reasonable solution. We believe this is due to the fact that the problem is very ill-posed due to the small probe overlap. This is also consistent with other observations on the behavior of ePIE [14]. Similarly, AWF also performs poorly with a large initial error, but then converges slowly to a more reasonable but large final NRMSE. SHARP




Fig. 6. NRMSE between reconstruction and ground truth as a function of number of iterations for reconstructions on synthetic noisy data with a small overlap: $r_{\text{ovlp}} \approx 14.5\%$.

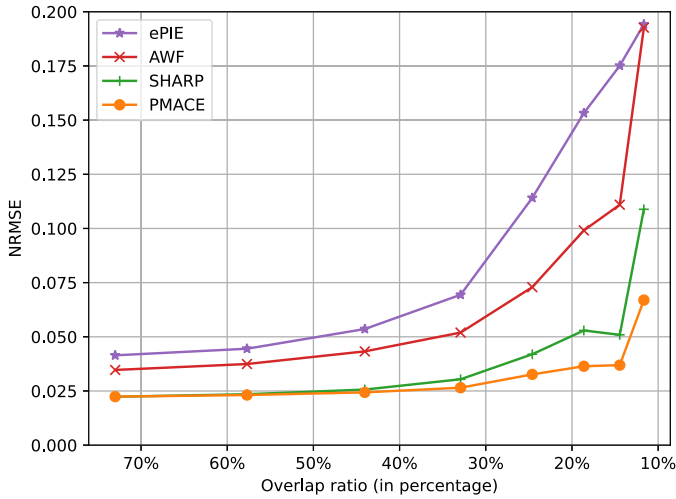


Fig. 7. NRMSE between reconstruction and ground truth for image of Fig. 3 as a function of the overlap ratio for each approach.

and PMACE have a large initial decrease in error followed by a slower decay, with SHARP leveling off fairly quickly and PMACE continuing to decrease to a good solution. Ultimately, PMACE achieves the lowest final error.

D. Reconstruction Quality Versus Probe Overlap

To determine the effect of probe overlap on reconstruction quality, we repeated the experiment above using multiple probe spacing values between 20 and 76 pixels, corresponding to overlap ratios from 73% to 12%. For each spacing value and each method, we optimized a single algorithmic parameter as described above, and we ran the corresponding method for 100 iterations.

Fig. 7 shows a plot of the final NRMSE of each method versus the overlap ratio. For large overlap ratios, all methods yield good image quality, with NRMSE values from about 2.5% to 4.5%.

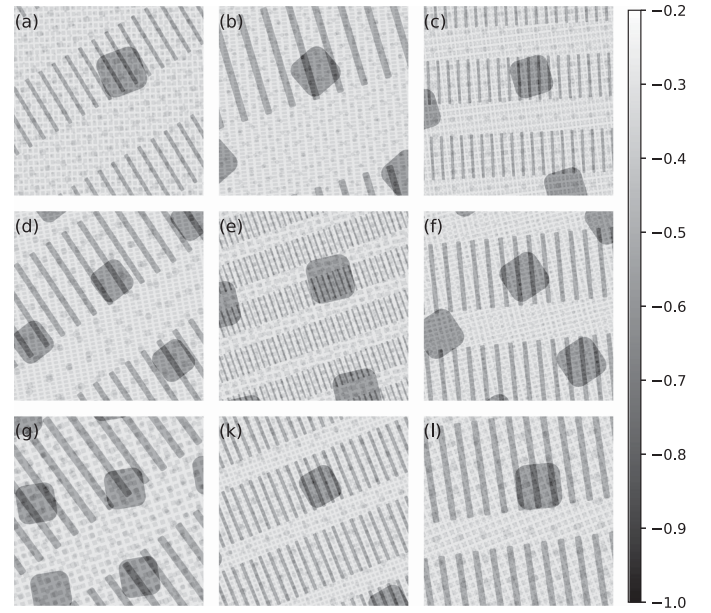


Fig. 8. Phase of the 9 additional complex transmittance images used as ground truth in simulation experiments of Fig. 9.

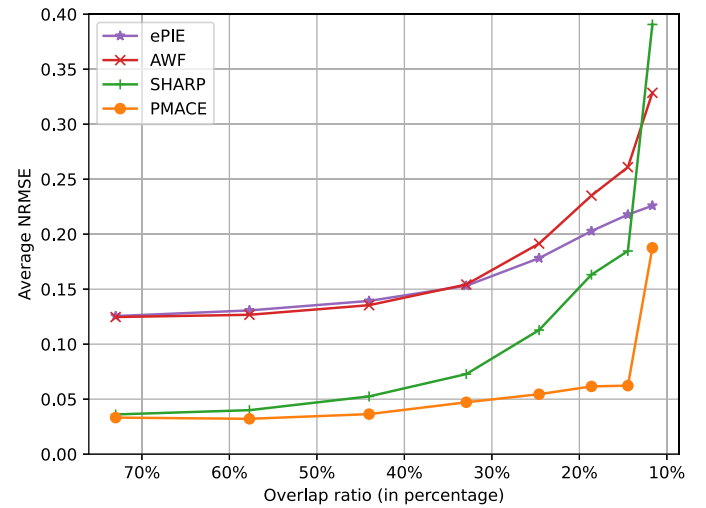


Fig. 9. Average NRMSE between reconstructions and ground truth as a function of the overlap ratio for each approach.




Fig. 10. The (a) magnitude and (b) phase of the probe used in all experiments on the measured data. The probe was estimated from the data using a variation of the ePIE algorithm.




Fig. 11. **Phases** (in radians) of complex reconstructions from measured data. The top row shows the full field of view along with outlines of insets displayed in rows 2–4. All methods reconstruct fine detail in the central region (row 3, outlined in yellow). ePIE and AWF yield artifacts near the edge of the field of view, while SHARP and PMACE produce more uniform results throughout the field of view.

However, ePIE and AWF underperform relative to SHARP and PMACE. As the overlap ratio decreases, NRMSE increases, with SHARP and PMACE maintaining good reconstruction quality of about 5% or less as the overlap ratio decreases to about 15%. PMACE achieves an NRMSE below 7% even with an overlap ratio r_{ovlp} below 12%, while the SHARP reconstruction has an NRMSE above 10% at this overlap ratio. ePIE and AWF are not competitive at small overlap ratios.

Fig. 9 shows a comparison of the algorithms for a set of 9 distinct 800×800 ground truth images shown in Fig. 8. Again, this result demonstrates that PMACE maintains the best reconstruction quality when the overlap percentage is low, and that SHARP and PMACE are significantly better than ePIE and AWF.

E. Reconstruction Results on Measured Data

In this subsection, we evaluate PMACE on measured data and report reconstruction results. We used the gold balls data set [35] consisting of 800 ptychographic measurements, which was generated using X-rays with a photon energy of 1 keV at beamline 5.3.2 of the Advanced Light Source. The scans were performed in a 20×40 grid with scan spacing of 30 nm and recorded by a detector placed 0.112 m downstream from the test sample of gold balls. This data set also includes 20 separately acquired dark images in addition to measured data.

For this experiment, the overlap ratio is approximately 84%, which is higher than any value used in our simulated data. This large overlap ratio and the results of Fig. 7 indicate that any of




Fig. 12. **Magnitudes** of complex reconstructions from measured data. PMACE captures fine detail in the largest field of view relative to the other methods.

the methods should be able to produce a reasonably accurate reconstruction.

We preprocessed the raw data by subtracting from each measurement the average of 20 dark scans. We also removed 6 abnormal measurements with high deviations in the data. Then we centered and cropped each measurement from 621×621 pixels to 512×512 pixels. To suppress noise, each of these diffraction measurements was multiplied by a 2D Tukey window that was generated by rotating a 1D Tukey window with shape parameter of 0.5. The resulting data set contains preprocessed 794 diffraction measurements.

Since the goal of this research is to compare the quality and speed of different ptychographic reconstruction algorithms, we used a fixed probe function for all algorithms in our comparisons. This fixed probe, shown in Fig. 10, was estimated from experimental data using a variation of the ePIE algorithm. We note that probe estimation is important in practice and has been incorporated into other algorithms, but we leave the problem of PMACE probe estimation as a topic for future research.

As in the case of simulated data, we varied one algorithmic parameter for each method and selected the parameter that produced the highest image contrast and fewest image artifacts. We ran each method for 100 iterations. As noted above, the fit to data is unchanged by a constant phase shift in the reconstruction. To provide a uniform visual comparison, we shift the phase of each reconstruction to achieve zero-mean phase in the central rectangle of 105×65 pixels, which was chosen since the centers of the probe illuminations are located in this region.

The reconstructed phase and magnitude images are shown in Figs. 11 and 12. We cropped the reconstructions to show the center 400×400 region of the image that was fully measured by the probes, and we used color outlines to highlight zoomed-in regions of the phase images. As can be observed, ePIE and AWF produce good image quality in the central region (yellow), but their image quality degrades in the red and blue regions. SHARP and PMACE produce good results in all three regions, with PMACE reconstructing the gold balls in the largest field of view. PMACE also appears to have slightly higher sharpness/contrast than SHARP, but since we do not have ground truth, it is difficult to know if the detail in the background is noise or valid detail.

Table I tabulates a measure of reconstruction quality for the ePIE, AWF, SHARP and PMACE reconstructions on the gold balls data set. Since there is no ground truth, we forward

TABLE I
FORWARD PROPAGATED NRMSE FOR EACH OF THE RECONSTRUCTION ALGORITHMS

Algorithm	ePIE	AWF	SHARP	PMACE
NRMSE	0.1102	0.1041	0.0914	0.0912

Notice that PMACE results in the lowest forward propagated error.

propagated the reconstructed complex images to the detector plane following (21) and calculate the NRMSE in the magnitude relative to the square-rooted data. In principle, a lower NRMSE at the detector plane implies a better fit to measurements; however, one must be careful since this metric does not account for overfitting of data. Notice that the PMACE reconstruction provides the best fit to the measured data. Hence these results support the claim that PMACE provides state-of-the-art quality under high-overlap conditions for measured data.

V. CONCLUSION

In this article, we introduce a theoretical foundation for the PMACE algorithm and show how it can be applied to the specific problem of ptychographic reconstruction. The PMACE algorithm is based on a set of agents each of which only acts on a small portion of the full reconstruction. This makes PMACE more suitable for large distributed inverse problems in which measured data only acts on a portion of the total image. As with MACE, the PMACE algorithm solves an equilibrium problem, which does not always correspond to an optimization problem. This makes PMACE more flexible for use in applications than traditional regularized inversion.

We applied PMACE to ptychography on both synthetic and measured data and found that it had improved image quality relative to competing algorithms, particularly for the case of large spacing between probe positions. This last point implies that our method requires less data for accurate reconstruction and can thus reduce acquisition time for a given level of reconstruction quality.

Finally, we note that the PMACE algorithm can be applied to other inverse problems in which each measurement only depends on a portion of the reconstruction. For example, this may be appropriate in tomography applications, such as laminography, in which each view corresponds to a portion of a larger planar object.

APPENDIX PROOFS

A. Proof of Theorem II.1

Proof: The assumptions on f_j and F_j imply [36, Section 6.1] that $F_j(v_j) = (I + \sigma^2 \nabla f_j)^{-1}(v_j)$ and hence $v_j = (I + \sigma^2 \nabla f_j)(F_j(v_j))$. With this in mind, suppose \mathbf{v}^* satisfies (9). By definition of \mathbf{F} , \mathbf{G}^P , and x^* , this means

$$F_j(v_j^*) = P_j \bar{x}(\mathbf{v}^*) = P_j x^* \quad (29)$$

for all j . Using $F_j = (I + \sigma^2 \nabla f_j)^{-1}$ and then applying $(I + \sigma^2 \nabla f_j)$ to both sides gives

$$v_j^* = P_j x^* + \sigma^2 \nabla f_j(P_j x^*). \quad (30)$$

Using this with the definition of x^* gives

$$x^* = \bar{x}(\mathbf{v}^*) = \Lambda^{-1} \sum_j P_j^T W(P_j x^* + \sigma^2 \nabla f_j(P_j x^*)), \quad (31)$$

so

$$\Lambda x^* = \sum_j P_j^T W P_j x^* + \sigma^2 \sum_j P_j^T W \nabla f_j(P_j x^*). \quad (32)$$

Since the first term on the right hand side is Λx^* , this simplifies to

$$\sigma^2 \sum_j P_j^T W \nabla f_j(P_j x^*) = 0, \quad (33)$$

which gives (15).

For the converse, if x^* satisfies (15), then by taking (30) as the definition of v_j^* , each step above is reversible to give (29) and hence (9). \square

B. Proof of Theorem II.2

Proof: 1. Note that \mathbf{G}^P is independent of $r > 0$ when $W = rI$, so we may assume $W = I$. A straightforward modification to the proof of [3, Theorem 2], shows that if \mathbf{F} is defined using the F_j , then a PMACE solution \mathbf{v}^* of (9) is equivalent using $x^* = \bar{x}(\mathbf{v}^*)$, $\lambda_j^* = v_j^* - P_j x^*$ to solutions (x^*, λ_j^*) that solve the following equations

$$F_j(P_j x^* + \lambda_j^*) = P_j x^* \quad (34)$$

$$\sum_j P_j^T \lambda_j^* = 0. \quad (35)$$

These conditions are equivalent to the first-order optimality conditions for the constrained optimization problem

$$x^* = \arg \min_{x, u_j} \left\{ \sum_{j=0}^{J-1} f_j(u_j) \right\} \text{ s.t. } u_j = P_j x \text{ for all } j, \quad (36)$$

which is equivalent to (17).

2. As noted above, the conservative property is equivalent to the symmetry of the Hessian matrix in (16), which is equivalent to

$$\sum_j P_j^T (W H_j(x) - H_j(x) W) P_j = 0, \quad (37)$$

for all x , where $H_j(x)$ is $H f_j$ evaluated at $P_j x$. After an arbitrarily small perturbation of the diagonal elements of W , we may assume that these diagonal entries are all distinct and positive. Then, if the left hand side of (37) is not identically 0, we are done. Otherwise, by an ϵ -small perturbation of f_0 (in the sup-norm on a ball large enough to include all images of interest), we can replace H_0 in (37) by $H_0 + \epsilon H$, with at least two off-diagonal entries of H nonzero. The assumption that $P_j P_j^T = I$ implies that P_0^T is injective, so with $H_0 + \epsilon H$ in place of H_0 , it is sufficient to show that $WH \neq HW$. Since W is diagonal with distinct, positive diagonal entries and H is symmetric, they commute if and only if H is also diagonal, which contradicts the choice of H . Hence after an arbitrarily small perturbation of W and/or f_0 , there is no potential function associated with the vector field (15). \square

C. Proof of Theorem II.3

Proof: Let \mathbf{v}^* denote the solution of (9) and $x^* = \bar{x}(\mathbf{v}^*)$. Since $\mathbf{G}^P(\mathbf{x}^*) = \mathbf{G}^P(\mathbf{G}^P(\mathbf{x}^*))$, \mathbf{v}^* satisfies (9) if and only if $F_j(v_j^*) = P_j x^*$ for all j and

$$\mathbf{G}^P(\mathbf{v}^*) = \mathbf{G}^P(\mathbf{F}(\mathbf{v}^*)). \quad (38)$$

This is equivalent to $v_j \in F_j^{-1}(P_j x^*)$ for all j and

$$P_i \Lambda^{-1} \sum_{j=0}^{J-1} P_j^T W v_j^* = P_i \Lambda^{-1} \sum_{j=0}^{J-1} P_j^T W F_j(v_j^*), \quad (39)$$

for all i . Since the matrix obtained by stacking the P_i is injective, this is equivalent to $v_j \in F_j^{-1}(P_j x^*)$ for all j and

$$\sum_{j=0}^{J-1} P_j^T W (v_j^* - F_j(v_j^*)) = 0. \quad (40)$$

Replacing both v_j^* with $F_j^{-1}(P_j x^*)$ gives the theorem. \square

D. Proof That the PMACE Forward Operator is Invertible

Proof: Recall that the PMACE forward operator from (22) is

$$F_j(v_j) = (1 - \alpha)v_j + \alpha D^{-1} \mathcal{F}^* \left(y_j \frac{\mathcal{F} D v_j}{|\mathcal{F} D v_j|} \right). \quad (41)$$

We first multiply both sides of (22) by $\mathcal{F} D$ to get

$$\mathcal{F} D F_j(v_j) = ((1 - \alpha) |\mathcal{F} D v_j| + \alpha y_j) \frac{\mathcal{F} D v_j}{|\mathcal{F} D v_j|}. \quad (42)$$

Taking absolute value of both sides and using $y_j \geq 0$ allows us to solve for $|\mathcal{F} D v_j|$ in terms of $F_j(v_j)$, y_j , and α . This leaves $\mathcal{F} D v_j$ as the only unknown in (42), so we solve for that and then multiply by $D^{-1} \mathcal{F}^*$ to get v_j . \square

REFERENCES

- [1] K. Dabov, A. Foi, V. Katkovnik, and K. Egiazarian, "Image denoising by sparse 3-D transform-domain collaborative filtering," *IEEE Trans. Image Process.*, vol. 16, no. 8, pp. 2080–2095, Aug. 2007.
- [2] K. Zhang, W. Zuo, Y. Chen, D. Meng, and L. Zhang, "Beyond a Gaussian denoiser: Residual learning of deep CNN for image denoising," *IEEE Trans. Image Process.*, vol. 26, no. 7, pp. 3142–3155, Jul. 2017.

- [3] G. T. Buzzard, S. H. Chan, S. Sreehari, and C. A. Bouman, "Plug-and-play unplugged: Optimization-free reconstruction using consensus equilibrium," *SIAM J. Imag. Sci.*, vol. 11, no. 3, pp. 2001–2020, 2018.
- [4] S. V. Venkatakrishnan, C. A. Bouman, and B. Wohlberg, "Plug-and-play priors for model based reconstruction," in *Proc. IEEE Glob. Conf. Signal Inf. Process.*, 2013, pp. 945–948.
- [5] S. Sreehari et al., "Plug-and-play priors for bright field electron tomography and sparse interpolation," *IEEE Trans. Comput. Imag.*, vol. 2, no. 4, pp. 408–423, Dec. 2016.
- [6] C. A. Bouman, *Foundations of Computational Imaging: A Model-Based Approach*. Philadelphia, PA, USA: SIAM, 2022, doi: [10.1137/1.9781611977134](https://doi.org/10.1137/1.9781611977134).
- [7] J. M. Rodenburg, "Ptychography and related diffractive imaging methods," *Adv. Imag. Electron Phys.*, vol. 150, pp. 87–184, 2008.
- [8] J. Rodenburg and A. Maiden, "Ptychography," in *Springer Handbook of Microscopy*. Berlin, Germany: Springer, 2019, pp. 819–904.
- [9] R. Wilke et al., "Hard X-ray imaging of bacterial cells: Nano-diffraction and ptychographic reconstruction," *Opt. Exp.*, vol. 20, no. 17, pp. 19232–19254, 2012.
- [10] P. Trtik, A. Diaz, M. Guizar-Sicairos, A. Menzel, and O. Bunk, "Density mapping of hardened cement paste using ptychographic X-ray computed tomography," *Cement Concrete Composites*, vol. 36, pp. 71–77, 2013.
- [11] M. Guizar-Sicairos et al., "High-throughput ptychography using eiger: Scanning X-ray nano-imaging of extended regions," *Opt. Exp.*, vol. 22, no. 12, pp. 14859–14870, 2014.
- [12] D. A. Shapiro et al., "Ptychographic imaging of nano-materials at the advanced light source with the nanosurveyor instrument," *J. Phys.: Conf. Ser.*, vol. 849, no. 1, 2017, Art. no. 012028.
- [13] J. R. Fienup, "Phase retrieval algorithms: A comparison," *Appl. Opt.*, vol. 21, no. 15, pp. 2758–2769, 1982.
- [14] J. M. Rodenburg and H. M. Faulkner, "A phase retrieval algorithm for shifting illumination," *Appl. Phys. Lett.*, vol. 85, no. 20, pp. 4795–4797, 2004.
- [15] A. M. Maiden and J. M. Rodenburg, "An improved ptychographical phase retrieval algorithm for diffractive imaging," *Ultramicroscopy*, vol. 109, no. 10, pp. 1256–1262, 2009.
- [16] P. Thibault and A. Menzel, "Reconstructing state mixtures from diffraction measurements," *Nature*, vol. 494, no. 7435, pp. 68–71, Feb. 2013, doi: [10.1038/nature11806](https://doi.org/10.1038/nature11806).
- [17] A. Maiden, D. Johnson, and P. Li, "Further improvements to the ptychographical iterative engine," *Optica*, vol. 4, no. 7, pp. 736–745, 2017.
- [18] O. Bunk, M. Dierolf, S. Kynde, I. Johnson, O. Marti, and F. Pfeiffer, "Influence of the overlap parameter on the convergence of the ptychographical iterative engine," *Ultramicroscopy*, vol. 108, no. 5, pp. 481–487, 2008.
- [19] E. J. Candes, X. Li, and M. Soltanolkotabi, "Phase retrieval via Wirtinger flow: Theory and algorithms," *IEEE Trans. Inf. Theory*, vol. 61, no. 4, pp. 1985–2007, Apr. 2015.
- [20] R. Xu et al., "Accelerated Wirtinger flow: A fast algorithm for ptychography," 2018, *arXiv:1806.05546*.
- [21] M. Odstrčil, A. Menzel, and M. Guizar-Sicairos, "Iterative least-squares solver for generalized maximum-likelihood ptychography," *Opt. Exp.*, vol. 26, no. 3, pp. 3108–3123, 2018.
- [22] F. Soulez, É. Thiébaud, A. Schutz, A. Ferrari, F. Courbin, and M. Unser, "Proximity operators for phase retrieval," *Appl. Opt.*, vol. 55, no. 26, pp. 7412–7421, 2016.
- [23] H. Yan, "Ptychographic phase retrieval by proximal algorithms," *New J. Phys.*, vol. 22, no. 2, 2020, Art. no. 023035.
- [24] S. Marchesini, Y.-C. Tu, and H.-T. Wu, "Alternating projection, ptychographic imaging and phase synchronization," *Appl. Comput. Harmon. Anal.*, vol. 41, no. 3, pp. 815–851, 2016.
- [25] S. Marchesini et al., "SHARP: A distributed GPU-based ptychographic solver," *J. Appl. Crystallogr.*, vol. 49, no. 4, pp. 1245–1252, 2016.
- [26] D. R. Luke, "Relaxed averaged alternating reflections for diffraction imaging," *Inverse Problems*, vol. 21, no. 1, p. 37, 2004.
- [27] V. Sridhar, X. Wang, G. T. Buzzard, and C. A. Bouman, "Distributed iterative CT reconstruction using multi-agent consensus equilibrium," *IEEE Trans. Comput. Imag.*, vol. 6, pp. 1153–1166, 2020.
- [28] R. E. Williamson and H. F. Trotter, "Multivariable mathematics: Linear algebra, calculus," in *Differential Equations*, vol. 2. Hoboken, NJ, USA: Pearson Prentice Hall, 2004.
- [29] P. L. Combettes, "Monotone operator theory in convex optimization," *Math. Program.*, vol. 170, pp. 177–206, 2018, doi: [10.1007/s10107-018-1303-3](https://doi.org/10.1007/s10107-018-1303-3).
- [30] H. H. Bauschke and P. L. Combettes, "Convex analysis and monotone operator theory in Hilbert spaces," in *CMS Books in Mathematics*. Berlin, Germany: Springer, 2011.
- [31] S. V. Venkatakrishnan, C. A. Bouman, and B. Wohlberg, "Plug-and-play priors for model based reconstruction," in *Proc. IEEE Glob. Conf. Signal Inf. Process.*, 2013, pp. 945–948.
- [32] P. Godard, M. Allain, V. Chamard, and J. Rodenburg, "Noise models for low counting rate coherent diffraction imaging," *Opt. Exp.*, vol. 20, no. 23, pp. 25914–25934, 2012.
- [33] X. Huang, H. Yan, R. Harder, Y. Hwu, I. K. Robinson, and Y. S. Chu, "Optimization of overlap uniformness for ptychography," *Opt. Exp.*, vol. 22, no. 10, pp. 12634–12644, 2014.
- [34] L. Zhou et al., "Low-dose phase retrieval of biological specimens using cryo-electron ptychography," *Nature Commun.*, vol. 11, no. 1, pp. 1–9, 2020.
- [35] S. Marchesini, "Ptychography gold ball example dataset," Lawrence Berkeley National Lab., Berkeley, CA, USA, Tech. Rep. CXIDB ID 65, 2017.
- [36] E. K. Ryu and S. Boyd, "A primer on monotone operator methods survey," *Appl. Comput. Math.*, vol. 15, no. 1, pp. 3–43, 2016.

Qiuchen Zhai (Student Member, IEEE) received the B.S. degree in electronic information engineering from the University of Electronic Science and Technology of China (UESTC), Chengdu, China, in 2018, and the M.S. degree in electrical and computer engineering from The Ohio State University, Columbus, OH, USA, in 2019. She is currently a Ph.D. candidate with the School of Electrical and Computer Engineering, Purdue University, West Lafayette, IN, USA. Her research interests include signal and image processing, computational imaging, inverse problems, and machine learning.

Gregory T. Buzzard (Senior Member, IEEE) is Professor of Mathematics at Purdue University, where he served as Department Head for 7 years. He has a Ph.D. from the University of Michigan and held postdoctoral appointments at Indiana University and Cornell University before joining Purdue. His research has led to theoretical advances in dynamical systems and experiment design and to new algorithms for image and volume reconstruction for a variety of sensing modalities. The unifying ideas in his recent work are iterative methods for image reconstruction and reduction of uncertainty through appropriate measurement schemes.

Kevin Mertes is a scientist at Los Alamos National Laboratory. He earned his Ph.D. from City University Of New York. He has conducted experimental research in quantum tunneling of magnetization in single molecule magnets, the metal-insulator transition in 2D electron gases and the mixing dynamics of spinor Bose-Einstein condensates. He has also conducted experimental and computational work in hyper-spectral remote sensing and x-ray ptychography.

Brendt Wohlberg (Fellow, IEEE) received the B.Sc. (Hons.) degree in applied mathematics, and the M.Sc. (Applied Science) and Ph.D. degrees in electrical engineering from the University of Cape Town, Cape Town, South Africa. He is a Staff Scientist with Theoretical Division at Los Alamos National Laboratory. He was the Chair of the Computational Imaging Special Interest Group (now the Computational Imaging Technical Committee) of the IEEE Signal Processing Society and was the Editor-in-Chief of the IEEE TRANSACTIONS ON COMPUTATIONAL IMAGING. He is currently Editor-in-Chief of the IEEE OPEN JOURNAL OF SIGNAL PROCESSING.

Charles A. Bouman (Fellow, IEEE) is the Showalter Professor of Electrical and Computer Engineering and Biomedical Engineering at Purdue University. He received his B.S.E.E. degree from the University of Pennsylvania, M.S. degree from the University of California at Berkeley, and Ph.D. from Princeton University. He is a member of the National Academy of Inventors as well as a Fellow of IEEE, AIMBE, IS&T, and SPIE. He has served as the IEEE Signal Processing Society's Vice President of Technical Directions as well as the Editor-in-Chief of the IEEE TRANSACTIONS ON IMAGE PROCESSING.