

A variational framework with composite sparse regularization for cryo-electron tomography reconstruction

Chenyun Yu^{1,2,†}, Ziheng Xu^{1,†}, Qiong Zeng³, Xiaohua Wan⁴, Haythem El-Messiry⁵, Fa Zhang^{4,*}, Renmin Han^{1,2,*}

¹Frontiers Science Center for Nonlinear Expectations (Ministry of Education), Research Center for Mathematics and Interdisciplinary Sciences, Qilu Hospital (Qingdao), Cheeloo College of Medicine, Shandong University, Qingdao, 266237, China

²College of Medical Information and Engineering, Ningxia Medical University, Yinchuan, 750004, China

³School of Computer Science and Technology, Shandong University, Qingdao, 266237, China

⁴School of Medical Technology, Beijing Institute of Technology, Beijing, 100081, China

⁵Canadian University Dubai, Dubai, 117781, United Arab Emirates

*Corresponding authors. Fa Zhang, School of Medical Technology, Beijing Institute of Technology, Beijing, 100081, China. E-mail: zhangfa@bit.edu.cn; Renmin Han, Frontiers Science Center for Nonlinear Expectations (Ministry of Education), Research Center for Mathematics and Interdisciplinary Sciences, Qilu Hospital (Qingdao), Cheeloo College of Medicine, Shandong University, Qingdao, 266237, China; College of Medical Information and Engineering, Ningxia Medical University, Yinchuan, 750004, China. E-mail: hanrenmin@sdu.edu.cn

[†]These authors contributed equally to this work.

Associate Editor: Pier Luigi Martelli

Abstract

Motivation: Cryogenic electron tomography (cryo-ET) enables *in situ* visualization of macromolecular and cellular structures from tilt-series projections. Reconstruction quality is often compromised by extremely low signal-to-noise ratio (SNR) and vignetting artifacts arising from detector truncation under constrained acquisition geometries. In practice, existing methods frequently struggle to balance noise robustness, computational efficiency, and stability under these conditions.

Results: We propose a robust, scalable, and parallelizable variational reconstruction framework that integrates a geometrically consistent data fidelity term with an implicit boundary-handling mechanism to mitigate truncation-induced artifacts without volume padding. A composite sparse regularizer integrating anisotropic total variation and curvelet-domain sparsity is employed to preserve structural boundaries and multiscale directional features. The resulting optimization problem is efficiently solved using the primal-dual hybrid gradient (PDHG) algorithm without nested inner iterations, for which we provide rigorous theoretical guarantees of stability and convergence. Experiments on simulated and experimental cryo-ET datasets demonstrate substantial noise suppression and contrast enhancement while preserving fine structural details under realistic, severely noise-limited and truncated acquisition conditions. These improvements lead to enhanced interpretability and facilitate downstream structural analysis, while achieving significantly reduced runtime compared to existing methods at comparable reconstruction quality.

Availability and implementation: Our code available at <https://github.com/icthrm/CSRT>. The real datasets used in this study are publicly available from EMPIAR and the Caltech Electron Tomography Database.

1 Introduction

Cryogenic electron tomography (cryo-ET) enables three-dimensional (3D) visualization of pleomorphic and heterogeneous biological specimens in a near-native state, providing direct access to *in situ* cellular architecture from a tilt series of two-dimensional (2D) transmission projections (Han and Xu 2024). Achieving high-fidelity reconstructions, however, is

severely impeded by fundamental physical and instrumental constraints. The electron dose must be kept extremely low to mitigate radiation damage, resulting in projection images with critically low signal-to-noise ratio (SNR). In addition, the accessible tilt range is restricted by microscope geometry, giving rise to the missing wedge problem; for subtomogram analysis, alignment and averaging can partially account for the missing wedge during similarity estimation and refinement, reducing the

Received: 14 February 2026. Revised: 2 April 2026. Accepted: 20 April 2026

© The Author(s) 2026. Published by Oxford University Press.

This is an Open Access article distributed under the terms of the Creative Commons Attribution License (<https://creativecommons.org/licenses/by/4.0/>), which permits unrestricted reuse, distribution, and reproduction in any medium, provided the original work is properly cited.

associated bias (Bartesaghi et al. 2008, Han et al. 2021). Beyond limited-angle sampling, cryo-ET reconstructions can also suffer from systematic bias due to the limited detector, often referred to as the long-object problem. Specifically, when the specimen extends beyond the reconstruction region of interest (ROI), the measured projections include contributions from outside the ROI that are absent from the forward model, leading to peripheral intensity fall-off and cupping or vignetting artifacts. Limited-detector compensation has been proposed to mitigate this mismatch and suppress vignetting (Xu et al. 2010).

A broad spectrum of reconstruction algorithms has been used in cryo-ET, with many microscopy pipelines relying on classical methods due to their simplicity and throughput. Direct approaches such as filtered back-projection (FBP) and weighted back-projection (WBP) are computationally efficient and therefore remain common in routine processing (Kak and Slaney (2001). Algebraic iterative methods, including the simultaneous iterative reconstruction technique (SIRT) (Gilbert 1972) and the simultaneous algebraic reconstruction technique (SART), improve data agreement through repeated forward-backward updates and are often favored for their robustness relative to direct back-projection under noisy and limited-angle conditions; under appropriate weighting, such iterations can be related to weighted least-squares (WLS) objectives (Van Der Sluis and Van Der Vorst 1990, Gregor and Fessler 2015). These algorithms are implemented and widely accessed through established software ecosystems for cryo-ET processing, including IMOD (Kremer et al. 1996), Tomo3D (Agulleiro and Fernandez 2011, 2015), EMAN2 (Tang et al. 2007), AuTom (Han et al. 2017, 2019), and acceleration-oriented toolkits such as TiltRec (Jiao et al. 2025) (see also (Pyle and Zanetti 2021)). Nevertheless, under low dose and missing-wedge sampling, classical reconstructions can exhibit substantial noise amplification, contrast loss, and iterative schemes may still rely largely on implicit regularization, for example through early stopping, reflecting the ill-posed nature of electron tomographic inversion (Fernandez 2012, Turoňová et al. 2015).

These limitations have motivated methods that incorporate stronger priors and better exploit the structure of cryo-ET data, most prominently through variational reconstruction frameworks and, more recently, deep learning. A statistically grounded model based iterative reconstruction (MBIR) framework was introduced in (Yan et al. 2019). It improves contrast by modeling the measurement process and incorporating object priors, but it is often computationally expensive because it relies on iterative MAP optimization and costly likelihood and prior terms (Sabne et al. 2017). A higher-order total variation (TV)-regularized reconstruction model was formulated in (Böhning et al. 2022) and solved using a primal-dual hybrid gradient (PDHG) algorithm (Chambolle and Pock 2011), while vignetting artifacts were reported that may require cropping in downstream subtomogram analysis. A GPU-accelerated proximal framework based on a linearized alternating direction method of multipliers (LADMM) was proposed in (Ramirez et al. 2024), but it incurs costly inner iterations, repeatedly invokes SART within the splitting updates, and pads the reconstruction volume to handle limited detector coverage, increasing memory and computation. Meanwhile, deep learning has become a complementary approach for cryo-ET denoising and missing-wedge compensation, spanning learning-based reconstruction and self-supervised fitting to measured data (Yang et al. 2021, Liu et al. 2022,

Wiedemann and Heckel 2024). However, such data-driven methods can raise concerns about scientific trustworthiness and reproducibility. When probing previously unknown *in situ* architectures, strict data fidelity is critical, since hallucinated features may mislead biological interpretation.

In this work, we propose a variational reconstruction framework to address these challenges via a regularized optimization model. First, we formulate a geometrically consistent WLS data fidelity term with an integrated limited-detector compensation scheme that suppresses vignetting artifacts without padding the reconstruction volume. Second, leveraging the modularity of proximal splitting, we implement a composite sparse regularization that combines anisotropic TV (ATV) with curvelet-domain sparsity, balancing boundary preservation with the recovery of multiscale directional features. Third, we solve the resulting non-smooth problem using an efficient PDHG algorithm, which decomposes the optimization into simple element-wise updates and avoids the costly nested inner iterations typical of LADMM-based methods, yielding improved scalability and higher computational efficiency among methods with comparable reconstruction quality. Moreover, we provide rigorous stability and convergence guarantees for the proposed algorithm, strengthening the reliability of the reconstructions for biological analysis. We validate the framework on simulated data and multiple experimental cryo-ET datasets using visual inspection, complementary quantitative metrics, and downstream task performance, demonstrating substantial noise suppression and contrast enhancement while preserving fine structural details.

2 Proposed method

This section presents the proposed variational framework for cryo-ET reconstruction, including its formulation, theoretical properties, and numerical solution. The reconstruction problem is introduced together with the associated notation, followed by the definition of an objective function that combines a truncation-aware data fidelity term and a composite sparse regularizer. Subsequent subsections detail the derivation of the weighted-norm fidelity model, the design of the regularization terms, and the PDHG-based optimization algorithm with stability and convergence guarantees.

2.1 Problem formulation

Let $\mathbf{u} \in \mathbb{R}^N$ denote the unknown 3D volume (vectorized) to be reconstructed from a tilt series, and let $\mathbf{b} \in \mathbb{R}^M$ denote the measured projection data after standard preprocessing (e.g., alignment and optional binning). The forward projection operator is written as $\mathbf{A} : \mathbb{R}^N \rightarrow \mathbb{R}^M$, with adjoint (backprojection) $\mathbf{A}^T$. We use a diagonal weighting matrix $\mathbf{R} \in \mathbb{R}^{M \times M}$ and the corresponding weighted norm.

$$\|\mathbf{x}\|_{\mathbf{R}}^2 \triangleq \mathbf{x}^T \mathbf{R} \mathbf{x}.$$

We estimate $\mathbf{u}$ in a variational manner by minimizing a sum of a data fidelity term and a regularization term:

$$\min_{\mathbf{u} \in \mathbb{R}^N} f(\mathbf{u}) + \Omega(\mathbf{u}).$$

The fidelity term $f(\mathbf{u})$ is constructed to be geometrically consistent under the limited-detector setting by incorporating




Figure 1 Illustration of the limited detector problem. (a) The gray area highlights regions of the specimen outside the defined reconstruction volume (ROI) that still contribute to the measured projection. The blue trapezoid represents the computed sum of weights within the ROI. (b) The blue rectangle represents the theoretical sum of weights corresponding to the full ray path through the specimen. The derived factor is applied for correction.

truncation compensation, aiming to suppress vignetting artifacts without padding the reconstruction volume. To regularize the reconstruction while retaining both boundary information and multiscale directional structures, we use a composite sparse regularizer that combines complementary priors rather than relying on a single penalty term.

$$\Omega(\mathbf{u}) = \lambda_1 \Omega_1(\mathbf{u}) + \lambda_2 \Omega_2(\mathbf{u}), \quad \lambda_1, \lambda_2 > 0,$$

Where in this work Ω_1 corresponds to ATV and Ω_2 promotes sparsity in the curvelet domain.

In the following, we first derive the weighted-norm data fidelity term together with the truncation compensation mechanism, and then present the composite regularization operators and the PDHG solver used to optimize the resulting non-smooth objective.

2.2 Weighted-norm data fidelity

The foundation of our reconstruction model is a WLS data fidelity term. In standard algebraic reconstruction, the iterative updates are typically expressed as

$$\mathbf{u}^{k+1} = \mathbf{u}^k - \mathbf{C}^T \mathbf{R}(\mathbf{A}\mathbf{u}^k - \mathbf{b}),$$

Where $\mathbf{R}$ and $\mathbf{C}$ are diagonal weighting matrices containing the inverse row and column sums of the system matrix $\mathbf{A}$, respectively,

$$r_{ii} = \frac{1}{\sum_{j=1}^N a_{ij}}, \quad c_{jj} = \frac{1}{\sum_{i=1}^M a_{ij}}.$$

Here, we assume that the system matrix entries are nonnegative, i.e., $a_{ij} \geq 0$, as they represent projection weights in tomographic imaging. Moreover, for the rows and columns involved in the reconstruction, the corresponding sums are strictly positive, that is, $\sum_{j=1}^N a_{ij} > 0$ and $\sum_{i=1}^M a_{ij} > 0$. This iteration implicitly minimizes the weighted objective function $f(\mathbf{u}) = \frac{1}{2} \|\mathbf{A}\mathbf{u} - \mathbf{b}\|_{\mathbf{R}}^2$ (Gregor and Fessler 2015).

However, a fundamental physical mismatch arises in cryo-ET due to the limited detector problem. As illustrated in Fig. 1, the acquired projection data $\mathbf{b}$ integrates information from the entire physical specimen, which often extends beyond the field of view. In contrast, the forward projection $\mathbf{A}\mathbf{u}$ accounts solely for

voxels within the defined reconstruction volume. Direct subtraction of these two terms leads to severe vignetting and cupping artifacts because $\mathbf{b}$ is physically scaled by the full object geometry, while $\mathbf{A}\mathbf{u}$ is scaled only by the ROI.

To resolve this inconsistency, we directly integrate the weight correction scheme in (Xu et al. 2010) into our variational framework. Let $\mathbf{W}$ denote the theoretical diagonal weighting matrix corresponding to the ray path lengths through the full specimen geometry. A physically consistent residual term, which compares the average density per unit length, is formulated as

$$\text{residual} = \mathbf{R}\mathbf{A}\mathbf{u} - \mathbf{W}\mathbf{b}. \quad (1)$$

Here, $\mathbf{R}\mathbf{A}\mathbf{u}$ represents the normalized projection of the reconstruction volume, and $\mathbf{W}\mathbf{b}$ represents the normalized acquired data. To incorporate this correction into our solver, we factor out the matrix $\mathbf{R}$:

$$\mathbf{R}\mathbf{A}\mathbf{u} - \mathbf{W}\mathbf{b} = \mathbf{R}(\mathbf{A}\mathbf{u} - \mathbf{R}^{-1}\mathbf{W}\mathbf{b}). \quad (2)$$

This derivation reveals that limited detector compensation is mathematically equivalent to replacing the raw projection data $\mathbf{b}$ with a corrected vector

$$\tilde{\mathbf{b}} = \mathbf{R}^{-1}\mathbf{W}\mathbf{b}. \quad (3)$$

Consequently, our final data fidelity term is defined as

$$f(\mathbf{u}) = \frac{1}{2} \|\mathbf{A}\mathbf{u} - \tilde{\mathbf{b}}\|_{\mathbf{R}}^2. \quad (4)$$

In practice, this correction is implemented as a pre-processing step. Since the exact geometry of the exterior object is unknown, the elements of $\mathbf{W}$ are approximated analytically by modeling the full specimen as a slab of uniform thickness d . The path length for a ray at tilt angle θ is estimated as $L = d / \cos(\theta)$, and the weights are set as the inverse of these lengths. By solving for $\tilde{\mathbf{b}}$ instead of $\mathbf{b}$, severe vignetting artifacts arising from the limited detector coverage are effectively eliminated.

2.3 Composite regularization

To obtain noise-robust yet structurally faithful reconstructions in cryo-ET, we use a composite regularization term that combines two complementary priors. We adopt ATV to preserve sharp boundaries and enforce local smoothness, and we further impose curvelet-domain sparsity to capture multiscale directional features. Together, these priors reduce noise while better maintaining fine biological structures than a single penalty term.

2.4 Anisotropic total variation

TV regularization has been widely adopted in electron tomogram reconstruction primarily for its robust denoising capabilities (Goris et al. 2012). Biologically, cryo-ET volumes typically consist of macromolecular complexes and organelles embedded in vitreous ice. The density transitions between these biological structures and the background solvent are ideally sharp, approximating a piecewise-constant intensity model (Deng et al. 2016).

Mathematically, TV minimization promotes sparsity in the gradient domain. In the context of cryo-ET, this property is

particularly effective at suppressing the severe high-frequency noise inherent to low-dose imaging while preserving the distinct boundaries of membranes and protein shells. By enforcing local smoothness within homogeneous regions, TV enhances the SNR and facilitates the visual delineation of structural features, which is critical for downstream segmentation tasks.

For a discrete 3D volume $\mathbf{u} \in \mathbb{R}^N$, the Isotropic TV (ITV) is often theoretically preferred for its rotational invariance:

$$\|\nabla \mathbf{u}\|_{2,1} = \sum_{i=1}^N \sqrt{(\mathbf{D}_x \mathbf{u})_i^2 + (\mathbf{D}_y \mathbf{u})_i^2 + (\mathbf{D}_z \mathbf{u})_i^2}. \quad (5)$$

However, the coupling of derivatives under the square root in ITV precludes a closed-form solution for its proximal operator, typically necessitating computationally expensive iterative sub-solvers. To ensure computational efficiency for large-scale tomograms, we adopt the ℓ_1 -based ATV, defined by:

$$\|\nabla \mathbf{u}\|_1 = \sum_{i=1}^N (|(\mathbf{D}_x \mathbf{u})_i| + |(\mathbf{D}_y \mathbf{u})_i| + |(\mathbf{D}_z \mathbf{u})_i|), \quad (6)$$

Where $\mathbf{D}_x, \mathbf{D}_y, \mathbf{D}_z$ represent the linear forward finite difference operators along the three spatial dimensions. The ATV formulation decouples the problem across spatial dimensions, yielding a fully separable proximal operator that allows the regularization update to be computed efficiently via simple component-wise soft-thresholding operations, thus avoiding the need for nested inner iterations. Consequently, we set $\Omega_1(\mathbf{u}) = \|\nabla \mathbf{u}\|_1$.

2.5 Curvelet-Domain sparsity

Standard TV regularization is widely used for preserving sharp boundaries (Rudin *et al.* 1992), but its axis-aligned gradient penalty can be insufficient for modeling curved and multiscale textures and may lead to “staircasing” artifacts (Böhning *et al.* 2022). Motivated by the directional and multiscale nature of cryo-ET structures, we introduce a curvelet sparsity prior as the second regularization term $\Omega_2(\mathbf{u})$. Curvelets provide an optimally sparse representation for images with C^2 -singularities such as smooth curves and wavefronts (Candès *et al.* 2006), and have been successfully applied in tomographic inverse problems, including sparse-view settings and hybrid formulations with TV (Yazdanpanah and Regentova 2016, Durand *et al.* 2017).

Let $\mathcal{C}$ denote the forward 3D fast discrete curvelet transform (FDCT) operator described in (Ying *et al.* 2005, Candès *et al.* 2006). We define $\Omega_2(\mathbf{u})$ as the ℓ_1 -norm of the curvelet coefficients:

$$\Omega_2(\mathbf{u}) = \|\mathcal{C}\mathbf{u}\|_1. \quad (7)$$

In our implementation, $\mathcal{C}$ is computed using the wrapping-based FDCT in the CurveLab toolbox, which offers favorable computational efficiency and near-isometry. The ℓ_1 penalty promotes sparsity in the curvelet domain by encouraging most coefficients to be small while retaining a limited number of sig-

nificant coefficients that encode dominant multiscale directional structures.

2.6 Numerical optimization via the PDHG algorithm

Substituting the weighted data-fidelity term and the composite sparsity regularizers into the generic variational model yields the unconstrained problem

$$\min_{\mathbf{u} \in \mathbb{R}^N} \left\{ \frac{1}{2} \|\mathbf{A}\mathbf{u} - \tilde{\mathbf{b}}\|_R^2 + \lambda_1 \|\nabla \mathbf{u}\|_1 + \lambda_2 \|\mathcal{C}\mathbf{u}\|_1 \right\}. \quad (8)$$

We solve this non-smooth problem using the primal-dual hybrid gradient (PDHG) algorithm (Chambolle and Pock 2011) by casting it into the saddle-point form

$$\min_{\mathbf{u}} \max_{\mathbf{v}} \langle \mathbf{K}\mathbf{u}, \mathbf{v} \rangle - F^*(\mathbf{v}) + G(\mathbf{u}), \quad G(\mathbf{u}) \equiv 0.$$

The stacked linear operator is

$$\mathbf{K} = \begin{bmatrix} \mathbf{R}^{1/2} \mathbf{A} \\ \nabla \\ \mathcal{C} \end{bmatrix}, \quad \mathbf{K}^* = [\mathbf{A}^T \mathbf{R}^{1/2} \quad \nabla^* \quad \mathcal{C}^*],$$

Where ∇^* and $\mathcal{C}^*$ denote the adjoints of the finite-difference and curvelet operators, respectively. Let $\mathbf{v} = (\mathbf{v}_A, \mathbf{v}_\nabla, \mathbf{v}_\mathcal{C})^T$, we define the block-separable function

$$F(\mathbf{v}) = \frac{1}{2} \|\mathbf{v}_A - \mathbf{R}^{1/2} \tilde{\mathbf{b}}\|_2^2 + \lambda_1 \|\mathbf{v}_\nabla\|_1 + \lambda_2 \|\mathbf{v}_\mathcal{C}\|_1.$$

PDHG alternates a proximal ascent step on the dual variable and a gradient descent step on the primal variable:

$$\mathbf{v}^{k+1} = \text{prox}_{F^*}(\mathbf{v}^k + \sigma \mathbf{K} \tilde{\mathbf{u}}^k), \quad (9)$$

$$\mathbf{u}^{k+1} = \mathbf{u}^k - \tau \mathbf{K}^* \mathbf{v}^{k+1}, \quad (10)$$

$$\tilde{\mathbf{u}}^{k+1} = \mathbf{u}^{k+1} + \theta(\mathbf{u}^{k+1} - \mathbf{u}^k). \quad (11)$$

Next, we derive closed-form updates for each dual block $(\mathbf{v}_A, \mathbf{v}_\nabla, \mathbf{v}_\mathcal{C})$.

2.7 Data fidelity update

The conjugate of the shifted quadratic data term F_A is $F_A^*(\mathbf{v}_A) = \frac{1}{2} \|\mathbf{v}_A\|_2^2 + \langle \mathbf{v}_A, \mathbf{R}^{1/2} \tilde{\mathbf{b}} \rangle$. Its proximal operator admits a closed-form linear solution:

$$\text{prox}_{F_A^*}(\mathbf{w}) = \frac{\mathbf{w} - \sigma \mathbf{R}^{1/2} \tilde{\mathbf{b}}}{1 + \sigma}. \quad (12)$$

Applying this to the general dual update step as shown in Eq. (9), the update for $\mathbf{v}_A$ becomes:

$$\mathbf{v}_A^{k+1} = \frac{\mathbf{v}_A^k + \sigma \mathbf{R}^{1/2} (\mathbf{A} \tilde{\mathbf{u}}^k - \tilde{\mathbf{b}})}{1 + \sigma}. \quad (13)$$

Direct evaluation of Eq. (13) requires the computationally expensive matrix square root $\mathbf{R}^{1/2}$. To circumvent this, we introduce the scaled dual variable $\tilde{\mathbf{v}}_A = \mathbf{R}^{1/2} \mathbf{v}_A$. Multiplying Eq. (13) by $\mathbf{R}^{1/2}$ yields an efficient update rule that relies solely on the original weight matrix $\mathbf{R}$:

$$\tilde{\mathbf{v}}_A^{k+1} = \frac{\tilde{\mathbf{v}}_A^k + \sigma \mathbf{R}(\mathbf{A}\tilde{\mathbf{u}}^k - \tilde{\mathbf{b}})}{1 + \sigma}. \quad (14)$$

Consequently, the contribution to the primal update step (10) simplifies to $\mathbf{A}^T \mathbf{R}^{1/2} \mathbf{v}_A^{k+1} = \mathbf{A}^T \tilde{\mathbf{v}}_A^{k+1}$.

2.8 Regularization update

For the regularization terms, the convex conjugates are indicator functions of ℓ_∞ -balls, i.e., $F_\nabla^* = \delta_{B_{\lambda_1}^\infty}$ and $F_C^* = \delta_{B_{\lambda_2}^\infty}$. Their proximal operators correspond to simple element-wise projections:

$$\text{prox}_{\sigma \delta_{B_{\lambda}^\infty}}(\mathbf{w}) = \Pi_\lambda(\mathbf{w}) = \text{sign}(\mathbf{w}) \odot \min(|\mathbf{w}|, \lambda). \quad (15)$$

Thus, the dual variables $\mathbf{v}_\nabla$ and $\mathbf{v}_C$ are updated by projecting their respective arguments onto the balls of radius λ_1 and λ_2 .

Substituting these specific operators into the iterative scheme yields the implementation detailed in Algorithm 1.

Algorithm 1 PDHG Algorithm for the Composite Regularized Model (8)

```

1: Operators  $\mathbf{A}, \mathbf{A}^T, \nabla, \nabla^*, C, C^*$ ; Weight matrix  $\mathbf{R}$ ; Corrected
   data  $\tilde{\mathbf{b}}$ ; Parameters  $\lambda_1, \lambda_2, \tau, \sigma, \theta$ .
2: Initialize primal variables  $\mathbf{u}^0 = \tilde{\mathbf{u}}^0 = \mathbf{0}$  and dual variables
    $\tilde{\mathbf{v}}_A^0, \mathbf{v}_\nabla^0, \mathbf{v}_C^0$  to  $\mathbf{0}$ .
3: for  $k = 0, 1, \dots, K_{\max} - 1$  do
4:   //Dual Update: Data Fidelity
5:   
$$\tilde{\mathbf{v}}_A^{k+1} \leftarrow \frac{\tilde{\mathbf{v}}_A^k + \sigma \mathbf{R}(\mathbf{A}\tilde{\mathbf{u}}^k - \tilde{\mathbf{b}})}{1 + \sigma}$$

6:   //Dual Update: Regularization ( $\ell_\infty$ -ball projections)
7:    $\mathbf{v}_\nabla^{k+1} \leftarrow \text{sign}(\mathbf{w}_\nabla) \odot \min(|\mathbf{w}_\nabla|, \lambda_1)$ , where
      $\mathbf{w}_\nabla = \mathbf{v}_\nabla^k + \sigma \nabla \tilde{\mathbf{u}}^k$ 
8:    $\mathbf{v}_C^{k+1} \leftarrow \text{sign}(\mathbf{w}_C) \odot \min(|\mathbf{w}_C|, \lambda_2)$ , where
      $\mathbf{w}_C = \mathbf{v}_C^k + \sigma C \tilde{\mathbf{u}}^k$ 
9:   //Primal Update
10:   $\mathbf{u}^{k+1} \leftarrow \mathbf{u}^k - \tau(\mathbf{A}^T \tilde{\mathbf{v}}_A^{k+1} + \nabla^* \mathbf{v}_\nabla^{k+1} + C^* \mathbf{v}_C^{k+1})$ 
11:  //Extrapolation
12:   $\tilde{\mathbf{u}}^{k+1} \leftarrow \mathbf{u}^{k+1} + \theta(\mathbf{u}^{k+1} - \mathbf{u}^k)$ 
13: end for
14: return  $\mathbf{u}^{k+1}$ 

```

2.9 Convergence analysis

We provide theoretical guarantees on the stability and convergence of the proposed PDHG solver under standard assumptions. Due to space limitations, the complete analysis, including supporting lemmas and proofs, is provided in Supplementary Section S1, available as supplementary data at Bioinformatics online.

3 Results

This section evaluates the proposed framework on simulated and experimental cryo-ET data. We report qualitative comparisons, leave-one-out FRC for measurement consistency, template matching for particle detection, and subtomogram averaging with FSC-based resolution on the simulated dataset.

3.1 Datasets and experimental setup

To evaluate the proposed method, we conducted experiments on one simulated dataset and four experimental cryo-ET datasets (EMPIAR-10643, EMPIAR-10499, EMPIAR-10110, and Nitrosop3). Dataset statistics and detailed acquisition parameters for each dataset are provided in Supplementary Section S2, available as supplementary data at Bioinformatics online. For all experimental datasets, the raw projections were spatially binned by a factor of 4 prior to reconstruction, and all tilt series were aligned using Markerauto2 Xu et al. (2024).

As comparison methods, we considered WBP, SIRT, MBIR, and LADMM. WBP and SIRT reconstructions were performed using Tomo3D, whose SIRT implementation also includes long object compensation. The parameters of the comparison methods were selected based on overall visual quality, noise suppression, and downstream task performance. In particular, the Hamming weighting filter frequency for WBP was set to 0.0, and SIRT was run for 10 iterations in all comparative experiments. Additional parameter sensitivity comparisons for WBP and SIRT are provided in Supplementary Section S3, available as supplementary data at Bioinformatics online. MBIR was configured following the parameter setting procedure described in Yan et al. (2019) and its supplementary material. For LADMM and our method, the parameters were determined by grid search, and the final setting was selected based on visual quality and downstream task performance. Moreover, once determined for a given dataset, the parameter settings of each method were used unchanged across all experiments involving that dataset. All experiments were conducted on a computing cluster node equipped with an Intel Xeon Platinum 8260 CPU and an NVIDIA A100 GPU with 40 GB memory.

3.2 Visual quality assessment

The XY slices in Fig. 2 highlight differences in structural visibility, contrast, and detail preservation among the reconstruction methods. WBP suffers from a low SNR, so only the approximate boundaries and the presence of particle like features can be recognized. SIRT makes boundaries and particle like features easier to distinguish than in WBP, but the contrast remains low and the noise level is still relatively high. MBIR performs well on EMPIAR-10643 and EMPIAR-10499, where discrete structures such as viral capsids and ribosomes are recovered with clearer interiors and more coherent edges; however, it is less favorable on the macroscopic Nitrosop3 dataset, where the reconstruction appears low in contrast and several cellular components remain poorly separated. LADMM suppresses background fluctuations, but it tends to




Figure 2 Visual comparison of XY and XZ slices from reconstruction results across real cryo-ET datasets. The columns correspond to different reconstruction methods (WBP, SIRT, MBIR, LADMM, and Ours), and the rows correspond to different biological specimens (EMPIAR-10643, EMPIAR-10499, and Nitrosop3).

over smooth the volume and weaken structural contrast, resulting in ambiguous cell envelope boundaries and a flattened intracellular appearance, particularly in Nitrosop3 where texture cues become less informative. In contrast, our method consistently enhances contrast and sharpness across all three datasets, improving the separation between structures and background while better preserving boundaries and fine structural details.

The XZ slices in Fig. 2 provide a more direct view of the limited detector problem, where the corresponding vignetting artifacts are more evident. Among the compared methods, SIRT in Tomo3D and our method are the two approaches that explicitly incorporate long object compensation. As shown in the XZ views, both methods effectively alleviate the vignetting artifacts caused by the limited detector problem, consistent with the behavior illustrated in Xu *et al.* (2010). In particular, the strong intensity imbalance near peripheral regions is substantially reduced, indicating that long object compensation is effective for regions whose tilted projections remain within the detector field of view. Some residual local intensity differences can still be observed in boundary areas with extremely limited projection support, but these are confined to regions where parts of the object extend beyond the recorded

projection area at certain tilt angles. Overall, the XZ results further support that our method handles the limited detector problem well.

Complementing the 2D slice analysis, Fig. 3 presents 3D iso-surface renderings of the HIV-1 dataset (EMPIAR-10643) visualized in ChimeraX Goddard *et al.* (2018) to assess structural fidelity. WBP is dominated by noise, producing a heavily fragmented volume in which the capsid particles are barely distinguishable. SIRT recovers the approximate particle shapes more clearly than WBP, but the reconstructed surfaces remain noisy and irregular, which still limits the clarity of the capsid shells. Regularized methods reduce these artifacts and recover overall particle shapes, but differences remain in structural completeness and surface texture. LADMM, although strongly denoised, shows structural loss, including incomplete shells and overly smooth surfaces that remove meaningful geometric variation. MBIR yields a clean visualization close to our result, yet Fig. 3f shows that it smooths intrinsic surface granularity and slightly reduces local detail on the capsid shell. By contrast, our method maintains a continuous and denoised surface while retaining more shell texture and local geometric variation, indicating a better balance between noise suppression and preservation of biologically relevant detail.




Figure 3 3D isosurface rendering comparison of the HIV-1 dataset (EMPIAR-10643) visualized using UCSF ChimeraX. (a) WBP. (b) SIRT. (c) MBIR. (d) LADMM. (e) Our proposed method. (f) Detailed close-up comparison between MBIR and our method. The dashed boxes in (c) and (e) indicate the zoomed-in regions displayed in (f), where the top crop corresponds to the MBIR result and the bottom crop corresponds to our method.

3.3 Fidelity assessment via leave-one-out FRC

We assessed reconstruction fidelity using leave-one-out Fourier ring correlation (FRC) [Cardone et al. \(2005\)](#), which evaluates how well a reconstructed tomogram predicts an unseen measurement. For each aligned projection X at a given tilt angle, we excluded it from the tilt series, reconstructed a tomogram from the remaining projections, and re-projected the tomogram at the omitted angle to obtain a prediction $\hat{X}$. We then computed the FRC between X and $\hat{X}$, where higher correlations over spatial frequencies indicate better measurement consistency and information restoration. LADMM was not included because the released implementation does not support leaving out individual projections.

[Figure 4A](#) compares the held-out raw projection at 0° with re-projections generated from tomograms reconstructed without that view. SIRT is visibly affected by artifacts, whereas WBP, MBIR, and our method appear more similar in overall structural recovery, although WBP exhibits a higher noise background. To quantify these differences, [Fig. 4B–C](#) report the leave-one-out FRC curves. WBP and SIRT show lower correlations with a clear frequency-dependent decay, while MBIR and our method consistently maintain higher correlations across a wide frequency range. Although minor dataset-specific variations exist, MBIR and our method closely track each other and provide a substantial improvement over the baselines. Overall, the results suggest that our variational framework achieves information restoration fidelity

comparable to the computationally intensive MBIR while remaining consistent with the original measurements.

3.4 Particle detection performance via template matching

To evaluate whether improved reconstruction quality translates to better downstream particle localization, we performed template matching using Dynamo [Castaño-Díez et al. \(2012\)](#). The dataset contains 400 ground-truth particles. Dynamo outputs a ranked list of candidate peaks; we matched each detected peak to the ground truth using nearest-neighbor association and counted it as a true positive if the Euclidean distance was below 10 pixels (approximately one-sixth of the particle diameter). Based on true positives (TP), false positives (FP), and false negatives (FN), we report Precision, Recall, and F1-score:

$$\text{Precision} = \frac{\text{TP}}{\text{TP} + \text{FP}}, \quad \text{Recall} = \frac{\text{TP}}{\text{TP} + \text{FN}},$$

$$\text{F1-score} = 2 \cdot \frac{\text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}}.$$

Precision quantifies how many selected candidates are correct, Recall measures the fraction of recovered ground-truth particles, and F1-score balances the two.

[Figure 5](#) and [Table 1](#) compare the methods in terms of both template matching accuracy and reconstruction efficiency. In terms of detection performance, MBIR achieves the best overall results, with the highest AP and an F1-score of 0.7900. Our method ranks second, reaching an F1-score of 0.7660, which is




Figure 4 Evaluation of information restoration fidelity on experimental cryo-ET data (Nitrosop3 and EMPIAR 10110). (A) Visual comparison between the original raw projection (Data, left column) excluded from reconstruction and the corresponding re-projections generated by WBP, SIRT, MBIR, and our method at 0° tilt angle. (B-C) Leave-one-out FRC curves calculated between the raw images and re-projections for the two datasets.




Figure 5 Template matching performance comparison. (A) Precision-Recall curves. The Average Precision (AP) for each method is noted in the legend. (B) F1-score as a function of the number of picked particles. The maximum F1-score and the corresponding number of particles n are indicated.

close to that of MBIR. In addition, our method improves both Precision and Recall compared with WBP, SIRT, and LADMM, indicating that it better preserves structural features relevant for localization. WBP provides the lowest baseline performance with an F1-score of 0.6140, reflecting its sensitivity to noise. SIRT substantially improves the result to 0.7420 by enhancing the visibility of major structures. LADMM achieves an F1-score of 0.6920, which is lower than that of SIRT, suggesting that its regularization may oversuppress or distort discriminative cues needed for template matching, even though the reconstruction appears visually smoother.

Table 1 Quantitative comparison of template matching performance and computational efficiency.

Method	TP ↑	FP ↓	FN ↓	Recall ↑	Precision ↑	F1-score ↑	Run time ↓
WBP	307	293	93	0.7675	0.5117	0.6140	4.3s
SIRT	371	229	29	0.9275	0.6183	0.7420	<u>17s</u>
MBIR	395	205	5	0.9875	0.6583	0.7900	7h42m
LADMM	346	254	54	0.8650	0.5767	0.6920	22m55s
Ours	<u>383</u>	<u>217</u>	<u>17</u>	<u>0.9575</u>	<u>0.6383</u>	<u>0.7660</u>	1h18m

Note. ↑ indicates higher is better, ↓ indicates lower is better. The run time denotes the total time required to reconstruct one tomogram. WBP, SIRT, MBIR, and our method were executed on CPU, whereas LADMM was executed on GPU.

Metrics were calculated based on the top 600 particles picked by Dynamo from a dataset containing 400 ground truth particles. The **best** results are highlighted in **bold**, and the second-best results are underlined.

We further assessed ranking reliability by analyzing the F1-score as a function of the number of picked particles in Fig. 5B. A reliable ranking should peak near the ground-truth count of 400, because the top candidates should contain mostly true particles before false positives accumulate. MBIR peaks at 0.9029 when selecting 380 candidates. Our method shows a similar behavior and peaks at 0.8322 when selecting 369 candidates,




Figure 6 Comparison of subtomogram averaging results. (A) Visual comparison between the reference 5lzf map (Ground Truth) and the averaged maps derived from tomograms reconstructed by different methods. (B) FSC curves calculated between the averaged maps and the ground truth reference. The dashed line indicates the FSC = 0.5 threshold.

which remains close to the ground truth and indicates that true positives are prioritized early in the ranked list. In contrast, LADMM and SIRT peak at substantially larger candidate pools, with 509 and 489 candidates, respectively. This shift implies that more false positives are mixed among higher-ranked detections, so a looser cutoff is needed to recover most true particles.

From the perspective of computational cost, the compared methods show clear differences in reconstruction time. WBP is the fastest method, requiring only 4.3 s to reconstruct one tomogram, followed by SIRT at 17 s. Their short runtimes are mainly due to their simpler reconstruction procedures and the efficient implementation in Tomo3D. WBP consists of projection filtering and a single backprojection step, whereas SIRT, although iterative, uses relatively simple forward- and backprojection-based updates. In contrast, LADMM, our method, and MBIR require 22 min 55 s, 1 h 18 min, and 7 h 42 min per tomogram, respectively, because they involve more complex optimization procedures and generally require more iterative updates, resulting in higher computational cost. These differences in reconstruction speed may be relevant in practical cryo-ET workflows, particularly for large-scale routine reconstruction, dataset screening, and parameter tuning. It should also be noted that LADMM was executed on GPU, whereas WBP, SIRT, MBIR, and our method were executed on CPU. Therefore, these runtimes should be interpreted as practical costs under our implementation settings rather than as a strictly hardware-normalized benchmark.

3.5 Resolution assessment via subtomogram averaging

To assess the preservation of high-frequency structural details and the reliability of the reconstructed volumes for downstream analysis, we performed subtomogram averaging (STA) on the simulated dataset. Specifically, 400 subtomograms (80^3 voxels, voxel size 4 Å) were extracted from the tomograms reconstructed by each method using the ground truth particle coordinates. Crucially, to decouple reconstruction quality from potential alignment inaccuracies, we bypassed the iterative alignment process and utilized the ground truth rotation and

translation parameters from the simulation to directly average the subtomograms.

Figure 6A displays the resulting averaged maps alongside the reference map. Visual inspection reveals that the averaged map from WBP contains more noticeable noise artifacts than those from the other methods. Conversely, while MBIR and LADMM effectively suppress noise, they tend to over-smooth the volume, resulting in the loss of fine surface textures. In contrast, the reconstructions obtained via SIRT and our method successfully retain sharper structural features.

The fidelity of the resulting density maps was quantified using the Fourier shell correlation (FSC) calculated against the ground truth reference map, as illustrated in Fig. 6B. Adopting the standard 0.5 threshold criterion [van Heel and Schatz \(2005\)](#), our method achieves a resolution of 19.06 Å. This represents a measurable improvement over SIRT (19.55 Å) and LADMM (19.52 Å), while marginally surpassing WBP (19.16 Å) and MBIR (19.22 Å). These quantitative results confirm that our proposed composite regularization strategy successfully balances robust noise suppression with the preservation of high-frequency structural information, effectively avoiding the resolution loss often associated with conventional regularization schemes.

4 Conclusion

We proposed a variational framework for 3D cryo-ET reconstruction under low SNR and geometrically incomplete data. The method combines a physically consistent WLS data-fidelity term with composite regularization to correct detector-truncation-induced vignetting without volume padding while preserving multi-scale biological structures. The resulting convex problem admits stability guarantees and is solved efficiently via a PDHG scheme with parallelizable updates. Experiments on simulated and experimental datasets demonstrate consistent improvements over standard baselines in noise suppression, contrast, particle detection, and resolution, while mitigating over-smoothing artifacts; moreover, it achieves quality comparable to computationally intensive model-based approaches at substantially lower computational cost.

Despite these results, several limitations warrant further investigation. The regularization hyperparameters are currently tuned manually across datasets, and future work will investigate automated selection strategies based on statistical noise estimation. In addition, the proximal framework naturally accommodates additional priors, and integrating learned priors such as plug-and-play or RED-style pre-trained denoisers is promising for extremely low-dose regimes, provided that stability and convergence can be maintained or appropriately characterized.

Author contributions

Chenyun Yu (Conceptualization [Lead], Data curation [Lead], Formal analysis [Lead], Investigation [Lead], Methodology [Lead], Software [Lead], Validation [Lead], Visualization [Lead], Writing—original draft [Lead], Writing—review & editing [Lead]), Zihé Xu (Data curation [Supporting], Writing—review & editing [Equal]), Qiong Zeng (Writing—review & editing [Supporting]), Xiaohua Wan (Writing—review & editing [Supporting]), Haythem El-Messiry (Writing—review & editing [Supporting]), Fa Zhang (Supervision [Supporting], Writing—review & editing [Supporting]), and Renmin Han (Funding acquisition [Lead], Project administration [Lead], Resources [Lead], Supervision [Lead], Writing—review & editing [Lead])

Supplementary material

[Supplementary material](#) is available at *Bioinformatics* online.

Conflicts of interest

None declared.

Funding

This research was supported by the National Key Research and Development Program of China [2021YFF0704300], the Fundamental Research Funds for the Central Universities, the Dubai Future Foundation [Award No. 2024CANAD-MES-061], the Natural Science Foundation of Ningxia Province [2026AAC040016], the National Natural Science Foundation of China [W2511070, 624B2089, 32241027, 62472034], and the Natural Science Foundation of Shandong Province [ZR2023YQ057].

References

- Agulleiro JI, Fernandez JJ. Fast tomographic reconstruction on multicore computers. *Bioinformatics* 2011;**27**:582–3.
- Agulleiro J-I, Fernandez J-J. Tomo3d 2.0 – exploitation of advanced vector extensions (avx) for 3d reconstruction. *J Struct Biol* 2015;**189**:147–52.
- Bartesaghi A, Sprechmann P, Liu J *et al*. Classification and 3d averaging with missing wedge correction in biological electron tomography. *J Struct Biol* 2008;**162**:436–50.
- Böhning J, Bharat TA, Collins SM. Compressed sensing for electron cryotomography and high-resolution subtomogram averaging of biological specimens. *Structure* 2022;**30**:408–17.e4.
- Candès E, Demanet L, Donoho D *et al*. Fast discrete curvelet transforms. *Multiscale Model Simul* 2006;**5**:861–99.
- Cardone G, Grünewald K, Steven AC. A resolution criterion for electron tomography based on cross-validation. *J Struct Biol* 2005;**151**:117–29.
- Castaño-Díez D, Kudryashev M, Arheit M *et al*. Dynamo: a flexible, user-friendly development tool for subtomogram averaging of cryo-em data in high-performance computing environments. *J Struct Biol* 2012;**178**:139–51.
- Chambolle A, Pock T. A first-order primal-dual algorithm for convex problems with applications to imaging. *J Math Imaging Vis* 2011;**40**:120–45.
- Deng Y, Chen Y, Zhang Y *et al*. Icon: 3d reconstruction with ‘missing-information’ restoration in biological electron tomography. *J Struct Biol* 2016;**195**:100–12.
- Durand S, Frapart Y-M, Kerebel M. Electron paramagnetic resonance image reconstruction with total variation and curvelets regularization. *Inverse Probl* 2017;**33**:114002.
- Fernandez J-J. Computational methods for electron tomography. *Micron* 2012;**43**:1010–30.
- Gilbert P. Iterative methods for the three-dimensional reconstruction of an object from projections. *J Theor Biol* 1972;**36**:105–17.
- Goddard TD, Huang CC, Meng EC *et al*. Ucsf chimeraX: meeting modern challenges in visualization and analysis. *Protein Sci* 2018;**27**:14–25.
- Goris B, Van den Broek W, Batenburg K *et al*. Electron tomography based on a total variation minimization reconstruction technique. *Ultramicroscopy* 2012;**113**:120–30.
- Gregor J, Fessler JA. Comparison of SIRT and SQS for regularized weighted least squares image reconstruction. *IEEE Trans Comput Imaging* 2015;**1**:44–55.
- Han R, Li L, Yang P *et al*. A novel constrained reconstruction model towards high-resolution subtomogram averaging. *Bioinformatics* 2021;**37**:1616–26.
- Han R, Wan X, Li L *et al*. Autom-dualx: a toolkit for fully automatic fiducial marker-based alignment of dual-axis tilt series with simultaneous reconstruction. *Bioinformatics* 2019;**35**:319–28.
- Han R, Wan X, Wang Z *et al*. Autom: a novel automatic platform for electron tomography reconstruction. *J Struct Biol* 2017;**199**:196–208.
- Han R, Xu M. *The Advance of Computational Methods in Cryo-Electron Tomography*. Singapore: World Scientific, 2024, 3–54.
- Jiao Y, Li H, Xue Y *et al*. Tiltrec: an ultra-fast and open-source toolkit for cryo-electron tomographic reconstruction. *Bioinformatics* 2025;**41**:btaf068.
- Kak AC, Slaney M. *Principles of Computerized Tomographic Imaging*. Philadelphia, PA, USA: Society for Industrial and Applied Mathematics, 2001.
- Kremer JR, Mastronarde DN, McIntosh J. Computer visualization of three-dimensional image data using IMOD. *J Struct Biol* 1996;**116**:71–6.
- Liu Y-T, Zhang H, Wang H *et al*. Isotropic reconstruction for electron tomography with deep learning. *Nat Commun* 2022;**13**:6482.

- Pyle E, Zanetti G. Current data processing strategies for cryo-electron tomography and subtomogram averaging. *Biochem J* 2021;**478**:1827–45.
- Ramirez JR, Rautek P, Bohak C *et al.* GPU accelerated 3D tomographic reconstruction and visualization from noisy electron microscopy tilt-series. *IEEE Trans Vis Comput Graph* 2024;**30**:3331–45.
- Rudin LI, Osher S, Fatemi E. Nonlinear total variation based noise removal algorithms. *Phys D Nonlinear Phenomena* 1992;**60**:259–68.
- Sabne A, Wang X, Kisner SJ *et al.* Model-based iterative ct image reconstruction on gpus. In *Proceedings of the 22nd ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming PPoPP '17*. New York, NY, USA: Association for Computing Machinery, 2017, 207–20.
- Tang G, Peng L, Baldwin PR *et al.* Eman2: an extensible image processing suite for electron microscopy. *J Struct Biol* 2007;**157**:38–46.
- Turoňová B, Marsalek L, Davidovič T *et al.* Progressive stochastic reconstruction technique (psrt) for cryo electron tomography. *J Struct Biol* 2015;**189**:195–206.
- Van Der Sluis A, Van Der Vorst H. SIRT- and CG-type methods for the iterative solution of sparse linear least-squares problems. *Linear Algebra Appl* 1990;**130**:257–303.
- van Heel M, Schatz M. Fourier shell correlation threshold criteria. *J Struct Biol* 2005;**151**:250–62.
- Wiedemann S, Heckel R. A deep learning method for simultaneous denoising and missing wedge reconstruction in cryogenic electron tomography. *Nat Commun* 2024;**15**:8255.
- Xu W, Xu F, Jones M *et al.* High-performance iterative electron tomography reconstruction with long-object compensation using graphics processing units (GPUS). *J Struct Biol* 2010;**171**:142–53.
- Xu Z, Li H, Wan X *et al.* Markerauto2: a fast and robust fully automatic fiducial marker-based tilt series alignment software for electron tomography. *Structure* 2024;**32**:1507–18.e5.
- Yan R, Venkatakrishnan SV, Liu J *et al.* MBIR: a cryo-et 3d reconstruction method that effectively minimizes missing wedge artifacts and restores missing information. *J Struct Biol* 2019;**206**:183–92.
- Yang Z, Zhang F, Han R. Self-supervised cryo-electron tomography volumetric image restoration from single noisy volume with sparsity constraint. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*. Piscataway, NJ, USA: IEEE, 2021, 4036–45.
- Yazdanpanah AP, Regentova EE. *Sparse-View CT Reconstruction Using Curvelet and TV-Based Regularization*. Springer International Publishing, 2016, 672–7.
- Ying L, Demanet L, Candes E. 3D discrete curvelet transform. In *Proceedings of SPIE*, Vol. 5914. Bellingham, WA, USA: SPIE, 2005, 591413.