



PAPER

OPEN ACCESS

RECEIVED
5 November 2025REVISED
26 February 2026ACCEPTED FOR PUBLICATION
30 April 2026PUBLISHED
27 May 2026

Original content from
this work may be used
under the terms of the
Creative Commons
Attribution 4.0 licence.

Any further distribution
of this work must
maintain attribution to
the author(s) and the title
of the work, journal
citation and DOI.



Real-time CBCT reconstructions using Krylov solvers in repeated scanning procedures

Fred Hastings^{1,6}, S M Ragib Shahriar Islam^{2,3,4,6} , Malena Sabaté Landman^{5,6} , Sepideh Hatamikia^{2,3,4}, Carola-Bibiane Schönlieb¹ and Ander Biguri^{1,*}

¹ Department of Applied Mathematics and Theoretical Physics (DAMTP), University of Cambridge, Cambridge, United Kingdom

² Austrian Center for Medical Innovation and Technology (ACMIT), Wiener Neustadt, Austria

³ Center for Medical Physics and Biomedical Engineering, Medical University of Vienna, Vienna, Austria

⁴ Danube Private University (DPU), Krems, Austria

⁵ Mathematical Institute, University of Oxford, Oxford, United Kingdom

⁶ These authors contributed equally to this work.

* Author to whom any correspondence should be addressed.

E-mail: ab2860@cam.ac.uk

Keywords: cone beam computed tomography, inverse problems, Krylov subspace methods, PICCS regularization

Abstract

This work introduces a new efficient iterative solver for the reconstruction of real-time cone-beam computed tomography, which is based on the prior image constrained compressed sensing (PICCS) regularization and leverages the efficiency of Krylov subspace methods. In particular, we focus on the setting where a sequence of under-sampled computed tomography (CT) scans are taken on the same object with only local changes (e.g. changes in a tumour size or the introduction of a surgical tool). This is very common, for example, in image-guided surgery, where the amount of measurements is limited to ensure the safety of the patient. In this case, we can also typically assume that a (good) initial reconstruction for the solution exists, coming from a previously over-sampled scan, so we can use this information to aid the subsequent reconstructions. The effectiveness of this method is demonstrated in both a synthetic scan and using real CT data, where it can be observed that the PICCS framework is very effective for the reduction of artefacts, and that the new method is faster than other common alternatives used in the same setting.

1. Introduction

Advancements in fast iterative reconstruction techniques have significantly enhanced the utility of cone beam computed tomography (CBCT) in medical imaging, particularly for interventional procedures that demand rapid image acquisition and reconstruction. These developments are enabling CBCT to transition from primarily diagnostic use to a central role in real-time, image-guided interventions. Several clinical applications can benefit greatly from such advancements. For example, in image-guided needle-based interventions, such as percutaneous biopsies or robotic-assisted ablations, accurate instrument placement is critical (Kickuth *et al* 2015, Gulias-Soidan *et al* 2020, Hatamikia *et al* 2023). These procedures often require real-time CBCT imaging to guide and verify needle trajectories during the intervention. Since adjustments and repositioning may demand multiple scans within a short time frame, both rapid image reconstruction and low-dose capabilities are essential to maintain procedural efficiency and patient safety. Similarly, spinal surgeries involving pedicle screw placement rely heavily on intraoperative CBCT to verify correct screw positioning (Burström *et al* 2021, Fikuart *et al* 2025). Misplacement in these contexts can result in irreversible damage, including paralysis. Therefore, high-resolution CBCT imaging is often performed iteratively throughout the procedure. A fast and dose-efficient reconstruction method can significantly reduce surgery time while minimizing the patient's radiation exposure. Neuro-interventional procedures, such as cerebral angiography, also frequently involve repetitive CBCT scans. These may be needed to visualize parenchymal structures, monitor hemorrhage progression, or assess the

impact of ischemic events (Irie *et al* 2008, Srinivasan *et al* 2016, Kawauchi *et al* 2021). As these interventions demand both speed and accuracy, a reconstruction algorithm that delivers high-quality images with fast reconstruction—without compromising dose efficiency—is highly desirable. Furthermore, in image-guided radiation therapy, CBCT is routinely used for patient setup and alignment—sometimes daily or multiple times per week where an immediate CBCT scan prior to radiotherapy session is needed for treatment planning. Summarizing, two critical challenges emerge across all these applications; first, the need for rapid image acquisition and reconstruction, often in near real-time, to integrate imaging into procedural workflows, second, the need to minimize cumulative radiation dose, especially when multiple scans are required over short time intervals. Notably, these procedures often involve imaging the same anatomical region repeatedly within a session. This presents a unique opportunity to leverage prior scans as input to the reconstruction algorithm, thereby enhancing image quality even with fewer projections.

Several algorithms have been developed to address the challenges associated with using CBCT in image-guided interventions. One approach involves utilizing reconstruction techniques that inherently reduce radiation exposure compared to the conventional Feldkamp–Davis–Kress (FDK) algorithm (Feldkamp *et al* 1984), which is still prevalent in clinical settings. To this end, there has been a recent expansion of iterative solvers that sequentially update the reconstructed image to fit with the measured data for clinical use (Scherzer *et al* 2009, Liu 2014). These methods have been proven to produce equal quality images under much lower radiation exposure: either by lowering the x-ray dose, or lowering the amount of measurements (Baumueller *et al* 2012, Noël *et al* 2018, Mileto *et al* 2019). Moreover, these algorithms can easily be adapted to incorporate user-given prior information in the form of regularization. This is often related to assumed properties of the solution, e.g. Tikhonov if the images are smooth or total variation (TV) if they are piece-wise constant. More recently, prior image constrained compressed sensing (PICCS) (Chen *et al* 2008) was proposed to include structural similarities to previous reconstructions. This type of regularization is very suited for image guided interventions, where we assume only a small number of pixels change between consecutive images. Note that in this work we focus mainly on the appearance of new structures, such as tumours or needles, but mathematically the same exact principle applies to previously existing objects (e.g. spinal fixation pins appearing in clinical practice) being removed between scans. In fact, assuming a good initial reconstruction, it was shown in Lauzier *et al* (2012) that one can reduce the number of measurements by a very large amount using PICCS regularization in combination with an iterative solver (up to just using 20 projections). However, each iteration of any iterative algorithm takes more time than a single FDK reconstruction, and one generally needs to do tens or often hundreds of iterations to obtain sufficiently a good reconstruction. This has been a roadblock in the clinical adoption of such methods, and is particularly true for the PICCS framework when using the solvers that were proposed in the original and subsequent articles (Chen *et al* 2008, 2017, Lauzier *et al* 2012, Lee *et al* 2012). To overcome this drawback, we propose using a fast converging class of projection algorithms, namely Krylov subspace methods. These methods converge quadratically (as opposed to the general linear convergence of most standard iterative algorithms) and update the image within the iteration in one step (as opposed to the best performing standard algorithms that require multiple steps per iteration of the whole dataset). In fact, one can reconstruct a clinically sized image within a handful of seconds using these methods, as they require very few iterations to converge (Kulvait and Rose 2021, Sabaté Landman *et al* 2023). Traditionally, Krylov methods are used to solve linear systems or least-squares problems, which might include Tikhonov regularization, but have also been used in combination with more advanced regularization terms involving ℓ_p norms. One very successful approach to do this consists on approximating the ℓ_p norms (possibly in combination with a regularization matrix, as in the case of TV regularization), by a sequence of quadratic functionals, giving rise to a chain of (least-squares) minimization problems that need to be solved sequentially (Wohlberg and Rodriguez 2008, Daubechies *et al* 2010). For this reason, this technique, known as iteratively reweighted norm (IRN) scheme, involves two nested loops of iterations where the inner loop requires solving a possibly very large-scale least-squares problem. However, this can be done efficiently and with low memory requirements using Krylov methods based on short-recurrences. We emphasize, however, that unlike approaches based on flexible Krylov-subspaces (Chung and Gazzola 2019, Gazzola and Sabaté Landman 2019, Sabaté Landman *et al* 2023), or Krylov subspace recycling (Soodhalter *et al* 2020), the proposed method does not reuse Krylov subspace information across subproblems or scans, since this would require significantly more memory usage. Instead, we leverage the previous reconstructions only through an update of the regularization term.

In this work, we propose a new algorithm aimed at image-guided surgery which combines the benefits of PICCS regularization with the speed of Krylov methods, implemented in an IRN framework. The paper is organized as follows. Section 2 concerns the methods of this work. In particular, in section 2.1 we describe the mathematical model for CBCT and the PICCS regularization framework. In section 2.2.1

we recall necessary background, and in section 2.2.2 we describe the new algorithm in detail. Finally, in section 3, we evaluate the performance of the new method in simulated data, and subsequently we showcase its use in a lung needle surgery phantom experiment using a Phillips Allura FD20 Xper C-arm CBCT scanner.

2. Methods

This section outlines the methods used in this work. We begin by introducing the problem by formulating the mathematical model for CBCT along with the PICCS regularization framework. Following this, we review the necessary theoretical background for the proposed solution: the IRN scheme; and we finally describe the new algorithm, providing a detailed explanation of its implementation.

2.1. PICCS in Krylov methods

Computed tomography (CT) reconstruction can be modelled as a large scale linear system

$$\mathbf{Ax} + \mathbf{e} = \mathbf{b}, \quad (1)$$

where $\mathbf{x} \in \mathbb{R}^M$ is the vectorized 3D image with M voxels, $\mathbf{b} \in \mathbb{R}^N$ is the vector of measurements with added noise, where N is the number of detector pixels multiplied by the number of projection angles, $\mathbf{A} \in \mathbb{R}^{N \times M}$ is the sensing matrix and $\mathbf{e} \in \mathbb{R}^N$ is additive noise (Kak and Slaney 2001). Even though the noise is not white and Gaussian, this is usually assumed to simplify the computations in experiments that are far from the low photon count limit, see e.g. Mueller and Siltanen (2012), section 2.3.2. When the data is under-sampled or in fast acquisition regimes, recovering the solution \mathbf{x} from \mathbf{b} in equation (1) can be very sensitive to small perturbations in the measurements. This means that the noise or model errors can easily lead to very noisy reconstructions or artefacts. A good strategy to mitigate this effect is the use of additional information about the reconstruction. In particular, when performing a sequence of CT reconstructions on the same object with only minor differences, we can assume that a previously obtained (and reasonably accurate) reconstruction, denoted as \mathbf{x}_p , already exists. The key idea of this paper is to use information from \mathbf{x}_p to produce accurate reconstructions \mathbf{x} at later times while reducing the x-ray dose. Moreover, to be useful in a clinical setting, this has to be done efficiently in almost real-time.

The PICCS algorithm was proposed in Chen *et al* (2008) to incorporate information from \mathbf{x}_p in the solution to equation (1) in later times by solving the neighbouring variational regularization problem

$$\min_{\mathbf{x}} \left\{ \|\mathbf{Ax} - \mathbf{b}\|_2^2 + \alpha^2 \|\Phi_1(\mathbf{x})\|_1 + \lambda^2 \|\Phi_2(\mathbf{x} - \mathbf{x}_p)\|_1 \right\}, \quad (2)$$

where Φ_1 and Φ_2 are appropriately chosen sparsifying transforms, and where the regularization parameters λ, α are assumed to be known *a priori*. In particular, they propose choosing $\Phi_i = \mathbf{D}$, for $i = 1, 2$, where \mathbf{D} corresponds to a discrete approximation of a differential operator. The resulting regularization terms are also called TV functionals and the resulting minimization problem becomes:

$$\min_{\mathbf{x}} \left\{ \|\mathbf{Ax} - \mathbf{b}\|_2^2 + \alpha^2 \text{TV}(\mathbf{x}) + \lambda^2 \text{TV}(\mathbf{x} - \mathbf{x}_p) \right\}. \quad (3)$$

The first regularization term, $\text{TV}(\mathbf{x})$, is known to promote edge reconstructions and therefore reduces the noisy appearance of the recovered solution. Further, practitioners are often most interested in changes between consecutive images; for example observing how a tumour develops over time, or tracking a surgical tool moving in the body.

Therefore, one might use a $\text{TV}(\mathbf{x} - \mathbf{x}_p)$ regularization term to promote the salient differences between the two images to be clear. This is because piece-wise constant differences between images have a small TV norm, while small oscillatory differences (that are most likely to correspond to noise) have a large TV norm value. In contrast, penalizing ℓ_2 differences between the current and prior reconstructions, would penalize all differences equally.

Several variations on the PICCS framework exist in the literature. One such case, which we consider in this paper, is prior-image-registered penalized-likelihood estimation (PIPLE) (Stayman *et al* 2013), which involves the following minimization

$$\min_{\mathbf{x}} \left\{ \|\mathbf{Ax} - \mathbf{b}\|_2^2 + \alpha^2 \text{TV}(\mathbf{x}) + \lambda^2 \|\mathbf{x} - \mathbf{x}_p\|_2^2 \right\}, \quad (4)$$

where the regularization term including the prior reconstruction is relaxed to be an ℓ_2 . This is known to have more of an averaging effect between the prior reconstruction and the least-squares solution of the fit-to-data term. Moreover, in practice, a registration step might be necessary to align the prior image \mathbf{x}_p with the new reconstructions; this can be seamlessly implemented by adding a suitable deformation to \mathbf{x}_p in (4). This is done, for example, in PIRPLE (Stayman *et al* 2013), where a (linear) transformation operator parameterized by a vector of parameters is added to the framework (e.g. modelling 3D rigid body transformation). The framework defined in this paper can be easily adapted to this case, but it is not in the scope of this paper to explore this avenue.

2.2. The Method

The strategy used in this paper to solve the minimization problems in (3) and (4) is as follows. First, the functional in (3) (or (4)) is approximated by a smooth not quadratic functional and, second, this is solved by constructing a sequence of recursive quadratic approximations (or sub-problems), that are minimized sequentially. This method, which is part of a wider class of algorithms following a majorization–minimization scheme, requires solving a nested loop of iterations (Lanza *et al* 2017). The particular approximation strategy that we adopt in this paper, also known as IRN (Wohlberg and Rodriguez 2008), has been used for general TV regularization, and it is known to converge to the minimizer of the smoothed functional. In the following, the mathematical foundations of the presented algorithm are reviewed in detail, including the appropriate approximations of the TV, PIPLE, and PICCS functionals, as well as the efficient conjugate gradient least squares (CGLS) solver that we employ for the subproblems in the sequence, which exploits the use of fast and parallelizable GPU applications of the forward model.

2.2.1. Background on IRN for TV and extension to PIPLE

Let $\mathbf{D}_1 \in \mathbb{R}^{(M^{1/3}-1) \times M^{1/3}}$ be a scaled finite difference approximation of the first derivative operator in one dimension:

$$\mathbf{D}_1 = \begin{pmatrix} 1 & -1 & & \\ & \ddots & \ddots & \\ & & 1 & -1 \end{pmatrix},$$

so that, for 3-dimensional objects, we consider the following discrete approximation of the first derivative operator:

$$\mathbf{D} = \begin{pmatrix} \mathbf{D}_x \\ \mathbf{D}_y \\ \mathbf{D}_z \end{pmatrix} = \begin{pmatrix} \mathbf{D}_1 \otimes \mathbf{I} \otimes \mathbf{I} \\ \mathbf{I} \otimes \mathbf{D}_1 \otimes \mathbf{I} \\ \mathbf{I} \otimes \mathbf{I} \otimes \mathbf{D}_1 \end{pmatrix},$$

where \mathbf{D}_x , \mathbf{D}_y and \mathbf{D}_z correspond to directional derivatives and \mathbf{I} is the identity matrix of dimension $M^{1/3}$. We are interested in the following discrete isotropic TV regularization term, which can be defined in 3-dimensional objects as

$$\text{TV}(\mathbf{x}) = \left\| \sqrt{(\mathbf{D}_x \mathbf{x})^2 + (\mathbf{D}_y \mathbf{x})^2 + (\mathbf{D}_z \mathbf{x})^2} \right\|_1, \quad (5)$$

where the squaring and square root operations are applied entry-wise. Note that the functional in (5) is not differentiable at any vector with a 0 component, so we consider instead a smoothed version defined as

$$\text{TV}(\mathbf{x}) \approx \left\| \tilde{\mathbf{W}}(\mathbf{D}\mathbf{x}) \mathbf{D}\mathbf{x} \right\|_2^2, \quad (6)$$

with weights defined as

$$\tilde{\mathbf{W}}(\mathbf{D}\mathbf{x}) = \begin{pmatrix} \tilde{\mathbf{W}}_* & & \\ & \tilde{\mathbf{W}}_* & \\ & & \tilde{\mathbf{W}}_* \end{pmatrix} \in \mathbb{R}^{3M^{2/3}(M^{1/3}-1) \times 3M^{2/3}(M^{1/3}-1)} \quad (7)$$

with

$$\left[\tilde{\mathbf{W}}_*(\mathbf{D}\mathbf{x}) \right]_{ii} = \frac{1}{\sqrt{\sqrt{[\mathbf{D}_x \mathbf{x}]_i^2 + [\mathbf{D}_y \mathbf{x}]_i^2 + [\mathbf{D}_z \mathbf{x}]_i^2} + \tau^2}}, \quad (8)$$

and where τ is a smoothing parameter fixed ahead of the iterations. In this work, this framework is used to find approximate solutions of minimization problems involving TV regularization norms. Assuming that we have an initial guess for the solution \mathbf{x}_0 , then we can define a recursive sequence of reweighted least squares problems of the form

$$\mathbf{x}_k = \arg \min_{\mathbf{x}} \left\{ \|\mathbf{Ax} - \mathbf{b}\|_2^2 + \alpha^2 \|\tilde{\mathbf{W}}_k \mathbf{D}\mathbf{x}\|_2^2 \right\} \quad \text{where} \quad \mathbf{W}_k = \tilde{\mathbf{W}}(\mathbf{D}\mathbf{x}_{k-1}). \quad (9)$$

The regularization term in (9), ignoring additive constants that do not affect the minimization and multiplicative constants that are absorbed by the regularization parameter, is a quadratic tangent majorant of the smoothed version of the original regularization term, defined in (6), at $\mathbf{D}\mathbf{x}_{k-1}$. By definition, up to constant factors, this means that the regularization term in (9) is quadratic and is an upper bound for (6), and that at the point $\mathbf{D}\mathbf{x}_{k-1}$, the regularization term in (9) and (6) take the same value of both the functional and its gradient. Note that (9) has a unique solution if $\mathcal{N}(\mathbf{A}) \cap \mathcal{N}(\mathbf{W}_k \mathbf{D}) = \mathbf{0}$, where $\mathcal{N}(\cdot)$ denotes the null-space of a matrix. This is true in CT problems, where the only element in $\mathcal{N}(\tilde{\mathbf{W}}_k \mathbf{D})$ is the constant vector, which is not in the null space of \mathbf{A} . This class of methods is known to converge to the solution of the least-squares problem with smoothed TV regularization norm, and it is known as iteratively reweighted least squares or IRN scheme.

Moreover, if one wants to incorporate prior information from a solution \mathbf{x}_p , it is also possible to incorporate another term in the minimization problem. For the simplified PIPLE (Stayman *et al* 2013) with a TV and an ℓ_2 terms in (4), this corresponds to:

$$\mathbf{x}_k = \arg \min_{\mathbf{x}} \left\{ \|\mathbf{Ax} - \mathbf{b}\|_2^2 + \alpha^2 \|\tilde{\mathbf{W}}_k \mathbf{D}\mathbf{x}\|_2^2 + \lambda^2 \|\mathbf{x} - \mathbf{x}_p\|_2^2 \right\}, \quad (10)$$

where $\mathbf{W}_k = \tilde{\mathbf{W}}(\mathbf{D}\mathbf{x}_{k-1})$.

2.2.2. Using IRN for PICCS

Analogously to section 2.2.1, we approximate each TV functional in the PICCS problem (3) by a weighted 2-norm (including a mild smoothing):

$$\min_{\mathbf{x}} \left\{ \|\mathbf{Ax} - \mathbf{b}\|_2^2 + \alpha^2 \left\| \tilde{\mathbf{W}}(\mathbf{D}\mathbf{x}) \mathbf{D}\mathbf{x} \right\|_2^2 + \lambda^2 \left\| \tilde{\mathbf{W}}(\mathbf{D}\mathbf{x} - \mathbf{D}\mathbf{x}_p) \mathbf{D}(\mathbf{x} - \mathbf{x}_p) \right\|_2^2 \right\}, \quad (11)$$

with weights that now depend both on the current solution \mathbf{x} and on the given reference solution \mathbf{x}_p . Assuming we have an initial guess for the solution \mathbf{x}_0 , we define the corresponding recursive sequence of reweighted least squares problems as

$$\mathbf{x}_k = \min_{\mathbf{x}} \left\{ \|\mathbf{Ax} - \mathbf{b}\|_2^2 + \alpha^2 \left\| \mathbf{W}_k^{(1)} \mathbf{D}\mathbf{x} \right\|_2^2 + \lambda^2 \left\| \mathbf{W}_k^{(2)} \mathbf{D}(\mathbf{x} - \mathbf{x}_p) \right\|_2^2 \right\}, \quad (12)$$

where

$$\mathbf{W}_k^{(1)} = \tilde{\mathbf{W}}(\mathbf{D}\mathbf{x}_{k-1}) \quad \text{and} \quad \mathbf{W}_k^{(2)} = \tilde{\mathbf{W}}(\mathbf{D}\mathbf{x}_{k-1} - \mathbf{D}\mathbf{x}_p). \quad (13)$$

Note that, in practice, problem (12) is solved by considering the equivalent augmented least-squares problems:

$$\mathbf{x}_k = \min_{\mathbf{x}} \left\| \begin{bmatrix} \mathbf{A} \\ \alpha \mathbf{W}_k^{(1)} \mathbf{D} \\ \lambda \mathbf{W}_k^{(2)} \mathbf{D} \end{bmatrix} \mathbf{x} - \begin{bmatrix} \mathbf{b} \\ \mathbf{0} \\ \lambda \mathbf{W}_k^{(2)} \mathbf{D}\mathbf{x}_p \end{bmatrix} \right\|_2^2 = \min_{\mathbf{x}} \left\| \tilde{\mathbf{A}}\mathbf{x} - \tilde{\mathbf{b}} \right\|_2^2, \quad (14)$$

and approximately solving them using a few inner iterations of a Krylov method (LSQR, or the mathematically equivalent CGLS) on $\tilde{\mathbf{A}}\mathbf{x} = \tilde{\mathbf{b}}$. Note that this is very efficient, since Krylov methods only require matrix-vector products with the system matrix and its transpose, and do not suppose a high memory over cost, since they are based on short recurrences for the solution and residual updates.

Summarizing, we use the IRN process to solve the TV, PIPLE and PICCS regularization problems so we can leverage the potential of Krylov solvers on large-scale 3D CT problems in an effective way.

3. Results

In this section, we present two examples representing medical imaging settings where there exists a (reasonably high quality) initial reconstruction, and this is accurate up to local changes, i.e. the introduction of a foreign body in the image, which can represent a tumour or a surgical tool. To test the validity of the proposed algorithm in a controlled testing environment, the first example corresponds to a simulated head phantom, and associated digitally rendered radiographs (DRR), with an artificial tumour (available in the TIGRE toolbox (Biguri *et al* 2016, 2025b)). However, it is often true that simplified assumptions in algorithm validation based on toy examples do not often translate to the more nuanced and complex noise distributions of a medical scanner in the clinic. Thus, to further validate our method, the second example corresponds to real scan data, corresponding to an anthropomorphic Thorax phantom with an inserted metallic surgical needle, acquired in a Phillips Allura scanner at the Universitätsklinikum Wiener Neustadt (formerly known as Landesklinikum). Note that the corresponding reconstructions for this example can be subject to strong metal artefacts.

In both examples, we assume that a good initial reconstruction is given, which in practice would most likely correspond to the reconstruction obtained using high-dose measurement data, but that the subsequent scan is heavily undersampled.

In the two examples, we aim to highlight the effect of using different regularization choices, as well as the performance of different optimization algorithms associated to them. Therefore, the reconstructions obtained using the following algorithms are presented, along with other quantitative metrics about their performance:

- Direct solver: FDK algorithm.
- Iterative algorithms without explicit regularization: simultaneous iterative reconstruction technique (SIRT) (Kak and Slaney 1988) and CGLS (Hestenes 1952). Note that most iterative algorithms have implicit regularization properties if equipped with early stopping, see, e.g. (Hansen 2010, Sabaté Landman *et al* 2023).
- Algorithms with only TV regularization on the solution: iterative reweighted norm based TV regularized CGLS (IRN-TV) (Sabaté Landman *et al* 2023) and adaptive steepest descent projection onto convex sets (ASD-POCS-TV) (Sidky and Pan 2008).
- Algorithm with TV regularization on the solution and ℓ_2 norm regularization on the prior term: the proposed IRN based PIPLE algorithm in combination with CGLS (IRN-PIPLE).
- Algorithms with PICCS regularization: using ASD-POCS (ASD-POCS-PICCS) (Hatamikia *et al* 2023) and the proposed algorithm, using an IRN framework in combination with CGLS (IRN-PICCS).

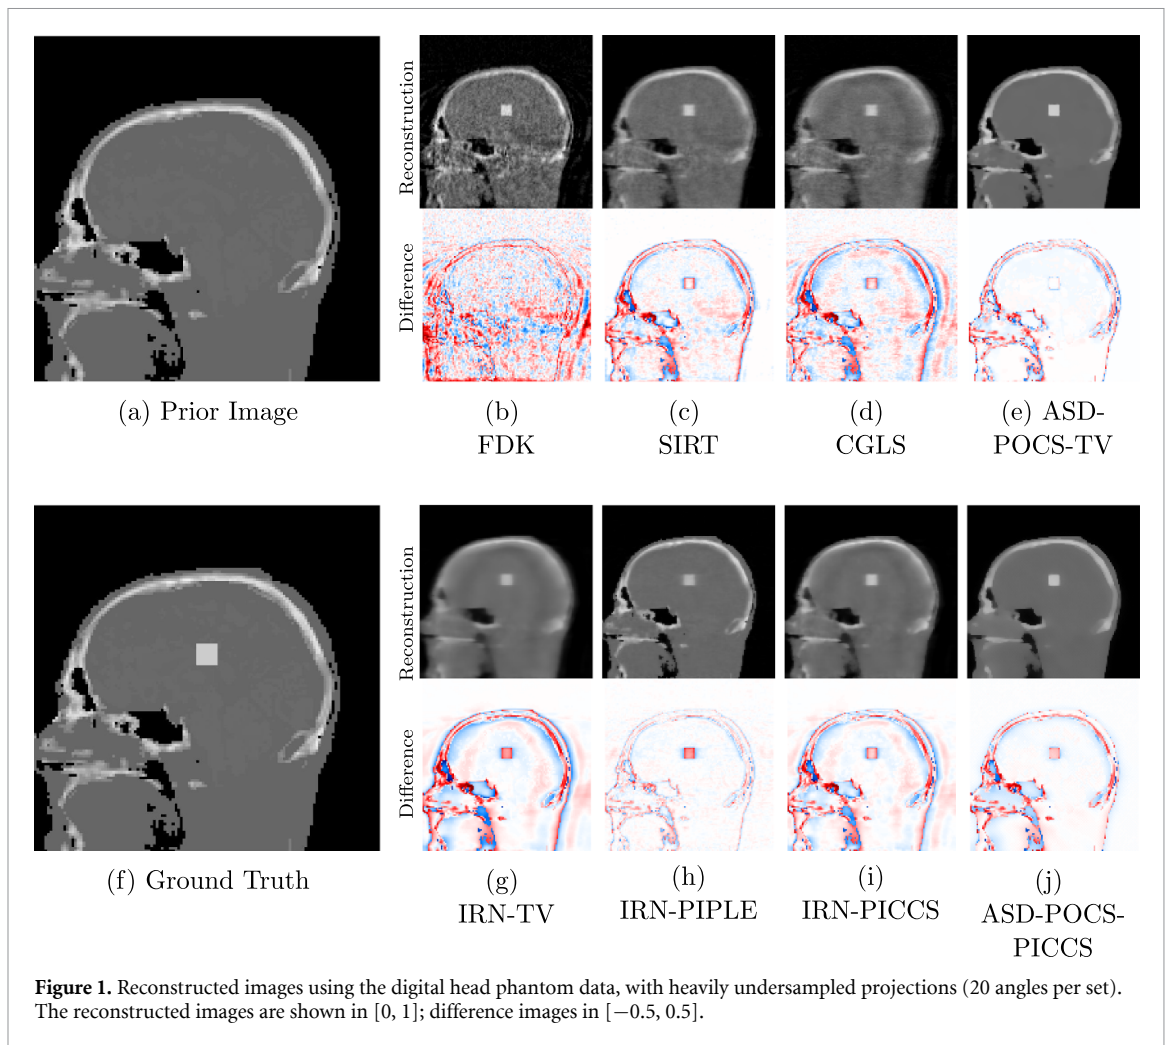
The reasoning to compare these algorithms is as follows. First, FDK is the clinical standard. Second, SIRT a standard iterative algorithm, but it displays slow convergence in comparison to the Krylov-based algorithm CGLS. Note that both produce implicitly regularized solutions. Third, ASD-POCS-TV is the most commonly used TV-regularized algorithm in the CT literature, and IRN-TV is a faster Krylov-based algorithm for the same problem. Finally, ASD-POCS-PICCS is the slow converging standard PICCS method, and IRN-PIPLE and IRN-PICCS are the proposed Krylov-based prior image regularized methods.

For all the setups, the corresponding geometry of the scanning devices can be replicated in the TIGRE toolbox when applying the reconstruction algorithms (Biguri *et al* 2016).

The PSNR, SSIM and HaarPSI were used for the quantitative analysis of the reconstructed images from all the performed experiments. In particular, HaarPSI is used due to its high correlation with qualitative clinicians' opinion of image quality (Breger 2025). Note that, when we show the reconstructions, only a slice of the 3D volumes is displayed. However, the quantitative error metrics are computed over the whole 3D images. The experiments were run in a standard NVIDIA GeForce RTX 4090 GPU in a workstation with AMD Ryzen Threadripper 1900X 16-Core Processor CPU and 135Gb of RAM.

3.1. Digital head phantom

In this synthetic example we consider cases where we want to assess if potential tumours have appeared (or changed in size or shape), using a very low number of projections, and we report reconstruction accuracy and computational time for different choices of regularization and optimization algorithm. In particular, we consider a digital head phantom, of size $128 \times 128 \times 128$ voxels, with a (synthetically) added cubic tumour of $8 \times 8 \times 8$ voxels. The DRRs projections were computed in TIGRE with resolution



$128 \times 128 \times 20$ projections. Note that this is a heavily undersampled dataset: we only consider 20 projections in each imaging step, while a standard clinical scan would require on the order of 600 projections. Even if this is a simple example, i.e. it is highly unlikely to observe a cubic tumour which such a strong contrast in practice, we believe this to be a highly illustrative test scenario, in which we can both assess the reconstruction quality of the fine structure in the digital head phantom, as well as the edges of the tumour.

For all the algorithms including prior image regularization, the head phantom image without the tumour was used as prior image. This is a reasonable choice, since we know we can do a high-dose scan once and obtain a very accurate reconstruction. Moreover, since this is a synthetic example, we can use the true solution as the ground truth image in our quantitative analysis.

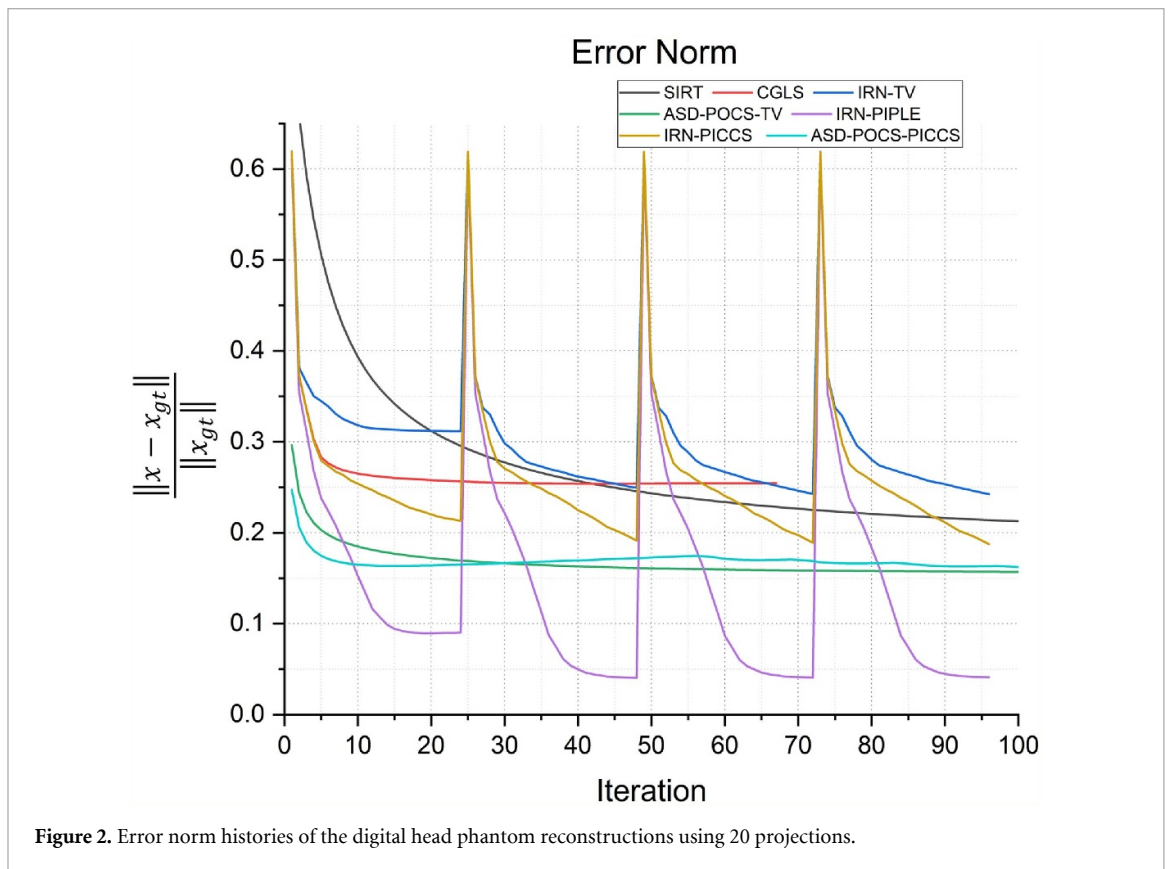
Figure 1 exhibits a slice of the reconstructed images using all the aforementioned algorithms. Here, (a) is the prior image x_p , (f) is the ground truth, and the rests are the reconstructed images (top) with the corresponding reconstructions errors (bottom), all displayed using the same colourbar. Note that the error is just the difference between the reconstruction with respect to the ground truth image, so that high pixel intensity indicates inaccurate reconstructions and a fully black image would indicate a perfect reconstruction. This local error information is more informative than only displaying the error norms, or global error, since one can easily observe there is a high local difference between reconstructions.

In terms of global information, table 1 contains the required number of iterations, reconstruction time, and the different image quality metric scores for the reconstructed images. Moreover, figure 2 shows the relative error norm histories corresponding to the different reconstruction algorithms.

Since this is a very undersampled dataset, FDK performs qualitatively poorly. Also for SIRT, and IRN-TV the quality drops significantly. The ASD-POCS-TV and ASD-POCS-PICCS are performing well, although a lot of texture information from the image is lost due to the smoothing process. However, the reconstruction time is higher than the other algorithms. The proposed algorithms perform better considering the available texture detail and the reconstruction time. Especially, the image in figure 1(h),

Table 1. Quantitative results for the digital head phantom example using 20 projections. Best value is bolded.

Algorithm	Iteration	Time	PSNR	SSIM	HaarPSI
FDK	N/A	00:00.2	20.6098	0.4329	0.3827
SIRT	100	00:30	22.2915	0.6594	0.5152
CGLS	20	00:04	21.5970	0.5898	0.4448
ASD-POCS	100	06:57	24.5381	0.8505	0.7091
IRN-TV	25	00:13	20.9140	0.6436	0.4429
IRN-TV	100(4 Outer)	00:45	21.7822	0.6895	0.4857
IRN-PIPLE	25	00:18	29.4520	0.9274	0.8524
IRN-PIPLE	100(4 Outer)	01:09	37.5399	0.9532	0.9394
IRN-PICCS	25	00:19	22.3996	0.7263	0.5298
IRN-PICCS	100(4 Outer)	01:16	24.1822	0.8070	0.5710
ASD-POCS-PICCS	20	02:46	23.8756	0.8012	0.6567

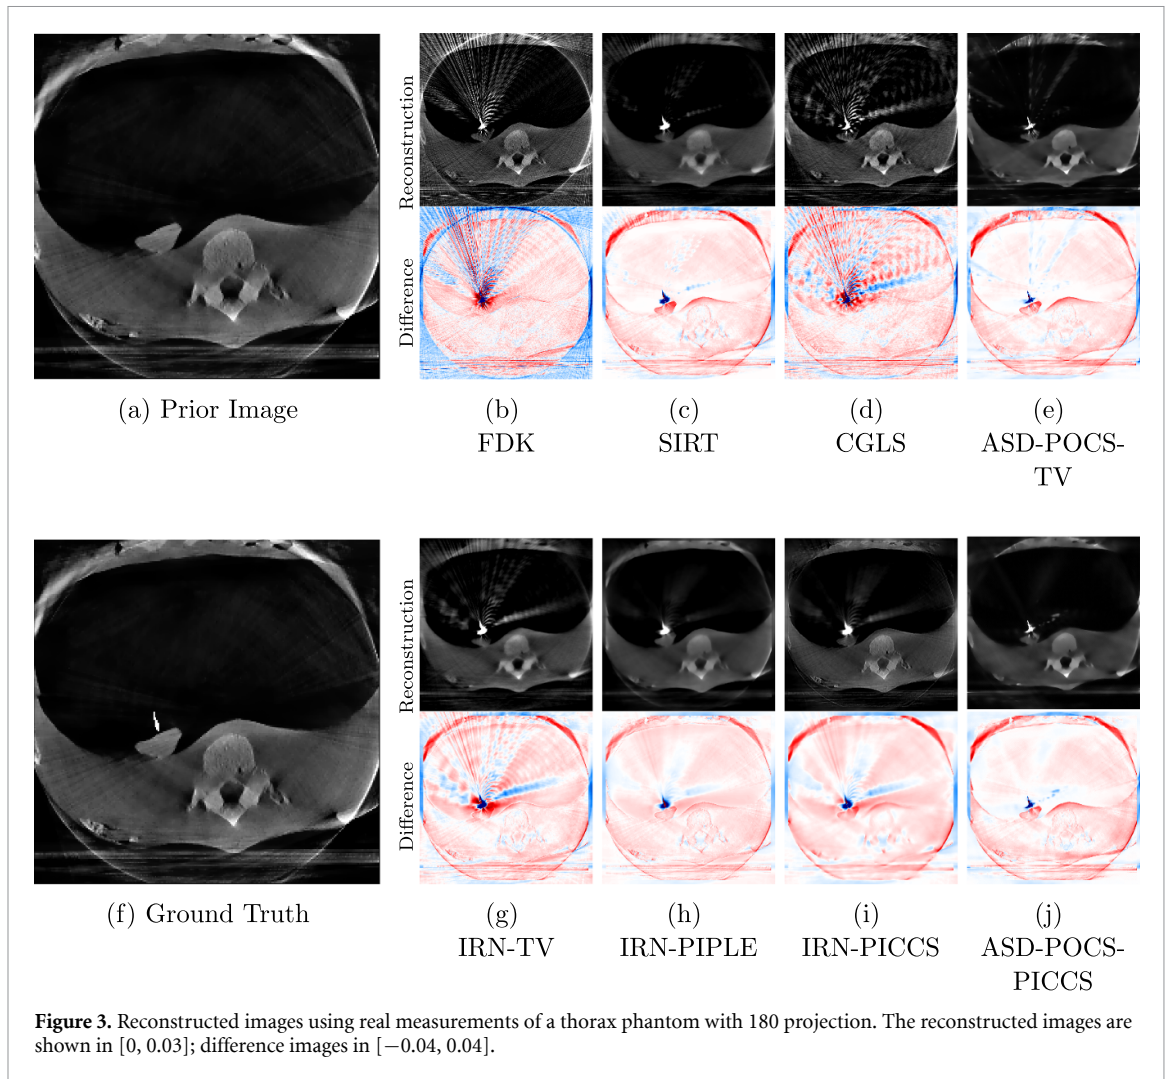
**Figure 2.** Error norm histories of the digital head phantom reconstructions using 20 projections.

IRN-PIPLE, is reconstructing the tumour better by preserving the texture information. Moreover, in figure 1(i), the tumour is more prominently visible due to the added TV regularization in the IRN-PICCS algorithm, however, with compromised original texture information of the image. Note that, for both proposed algorithms, the reconstruction time is in the range of 13–76 s only.

The error norm histories are displayed in figure 2. Note that the methods based on the IRN scheme have inherent restarts (in which the weights for the optimization are updated), which lead to the peaks in the error norms. This is an expected behaviour in these types of methods. However, the proposed IRN-PIPLE method obtains a reconstruction of higher quality in fewer iterations compared to other algorithms.

3.2. Lung needle surgery scan

In this example, we reconstruct a Thorax phantom scanned using the Phillips Allura FD20 Xper C-arm CBCT scanner. This device has source-to-axis and source-to-detector distances of 810 mm and 1195 mm respectively, and an installed detector of size $30 \times 40 \text{ cm}^2$ with 0.776 mm pixel pitch. Moreover, the data has a resolution of $364 \times 512 \times \text{number of projections}$ and the image is reconstructed with a resolution of $256 \times 256 \times 256$ voxels. Projections were acquired in a step-and-shoot protocol at 350 mAs of exposure and beam energy of 80 kV.

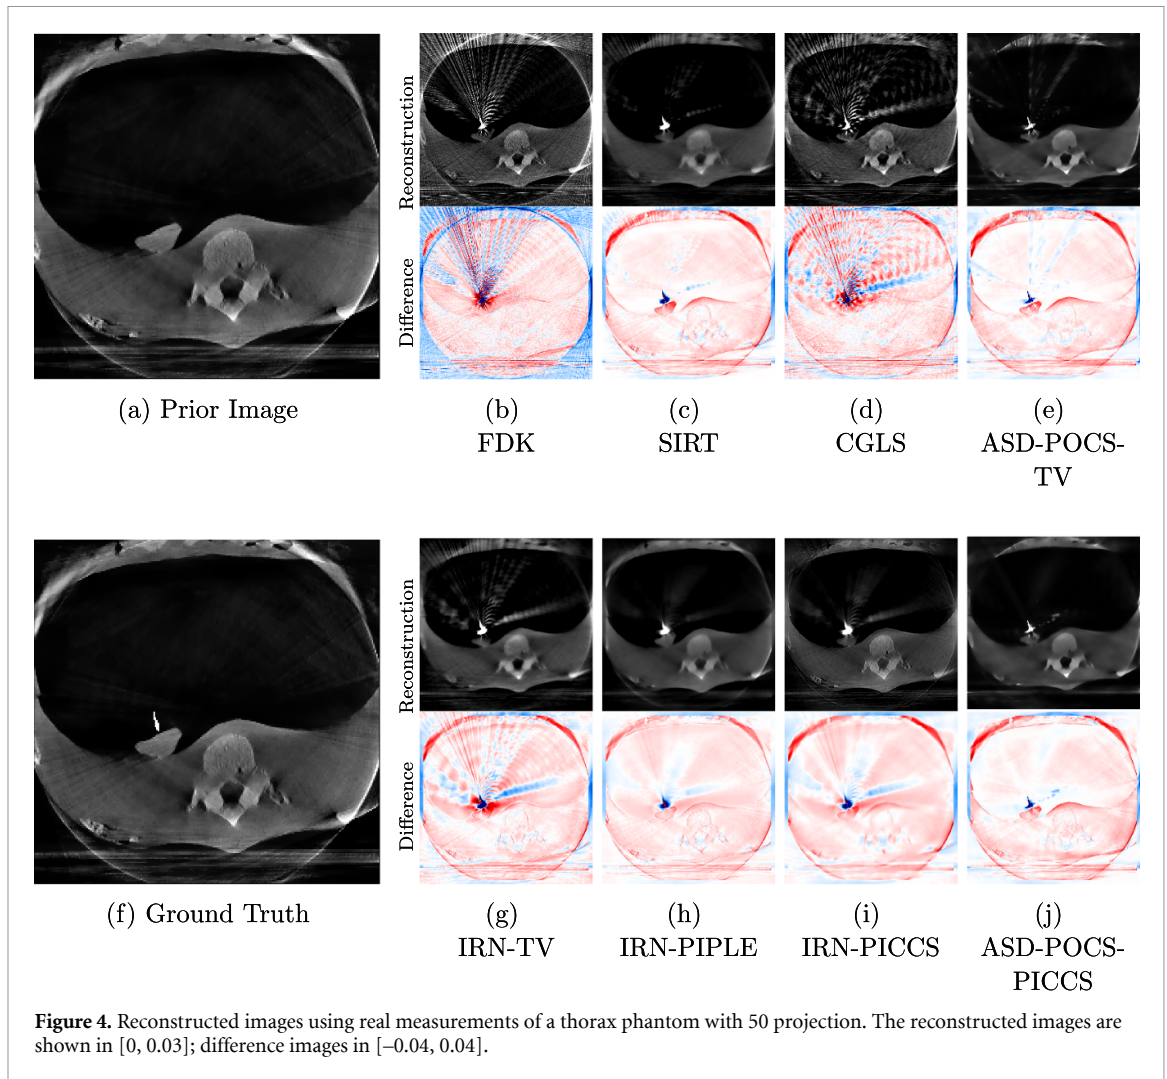


Since the data for this experiment corresponds to real measurements, we consider the ground truth to be a modelled approximation of the true solution. For this, a ground truth of the Thorax phantom is created using 3D Slicer, and the needle data is superimposed artificially in the appropriate location.

The reconstructions corresponding to the real-scan data with different number of projections (180, 50, and 20, respectively), can be observed in figures 3–5. In all three figures (a) is the prior image x_p used for prior image regularization, (f) is the ground truth image; and the rest of images correspond to different reconstructions (top) with their corresponding errors (bottom). Analogously to the previous experiment, the difference images (or error) should be completely black in the case of a perfect reconstruction and high intensity indicates more inaccurate reconstructions.

Note that the real data contains a lot of measurement noise, and due to the presence of the needle, the reconstructions can suffer from strong metal artefacts. Therefore, the reconstructed images using different algorithms, particularly without explicit regularization, tend to show strong errors even with a high number of projections. Accordingly, all the algorithms with explicit regularization perform better in this scenario, and particularly those including PIPLE or PICCS regularization. Moreover, the proposed IRN-PIPLE algorithm significantly outperforms other methods, and IRN-PICCS is not far behind. Using only TV regularization struggles to separate artefacts from features, oversmoothing images and thus losing important image features. Notably, the proposed methods reconstruct high quality images in less than 2 min in an implementation that is not optimized to solve this specific geometry.

Some evaluation metrics for this experiment can be seen in tables 2–4 (considering 180, 50, and 20 projections respectively) containing the required number of iterations, reconstruction time, and the different image quality metrics scores for the reconstructed images. Note that the image quality measurement metrics are performed over the whole 3D volume with respect to the (approximated) ground truth. Krylov algorithms have been run for 20 iterations, but given the slower convergence of ASD-POCS, 100 iterations are set as maximum. Albeit sometimes these last algorithms will stop earlier due



to their stopping criteria, which evaluates the direction of alternating minimization vectors and stops when they are opposite.

The proposed algorithms are the best performing in almost all cases, except for HaarPSI at high number of projections, where the prior image regularization is less meaningful. This can also be qualitatively observed from the reconstructed images.

4. Discussion

This study discusses different methods to obtain fast and accurate reconstructions in the context of repeated low-dosage CBCT sequence scanning of the same ROI, given a high quality initial scan. To do so, we focus on both the choice of appropriate regularization to ensure image quality, and the appropriate choice of the algorithm to ensure speed.

In terms of regularization, we compare implicit regularization given by early stopping of particular iterative methods, TV regularization, and methods that combine TV regularization on the image and prior based regularization on the L2 norm (PIPLE) and the TV norm (PICCS). First, as noted before in the literature, we confirm that methods based on image prior information (PIPLE and PICCS) very effectively reduce the number of projections required to obtain a meaningful reconstruction. For example, a standard clinical scan for needle surgery requires in the order of 600 projections, or 300 projections in low dose protocols, in each imaging step. This (and other works) suggest instead that using only 3% of the standard dose can be sufficient, given that a good regularization method is chosen. This reduction is crucial when we consider that, for image-guided needle surgery, several scans are taken during the entire procedure.

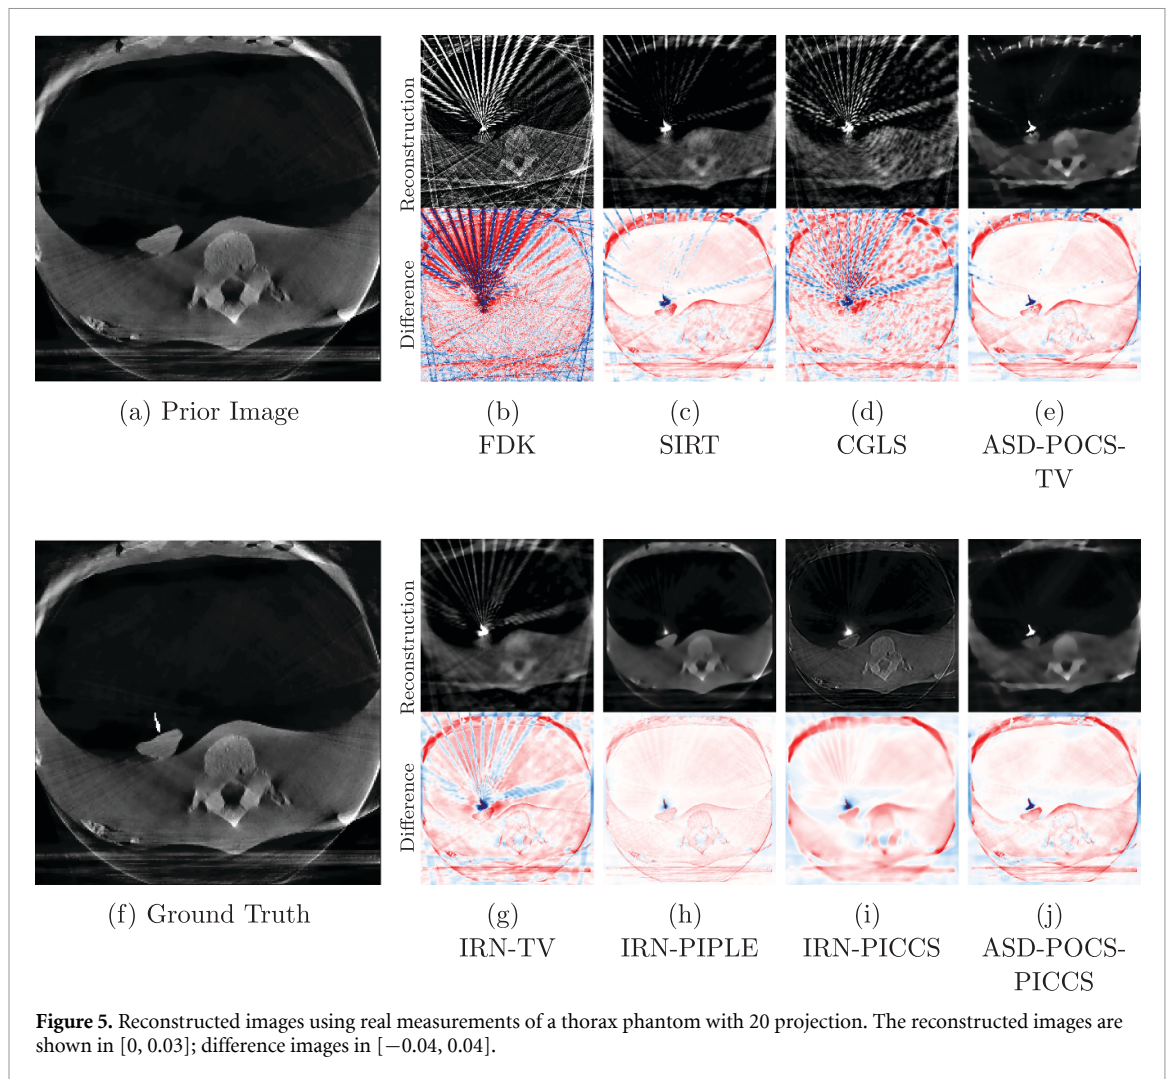


Figure 5. Reconstructed images using real measurements of a thorax phantom with 20 projection. The reconstructed images are shown in $[0, 0.03]$; difference images in $[-0.04, 0.04]$.

Table 2. Quantitative results for the thorax phantom example with real measurements involving 180 projections.

Algorithm	Iteration	Time	PSNR	SSIM	HaarPSI
FDK	N/A	00:00:01.33	33.86	0.8202	0.5870
SIRT	100	01:19	40.65	0.9464	0.6809
CGLS	20	00:15	34.05	0.8501	0.5822
ASD-POCS-TV	100	01:15:56	36.74	0.9463	0.6798
IRN-TV	20	01:05	37.51	0.8935	0.6079
IRN-PIPLE	20	01:24	43.85	0.9473	0.7789
IRN-PICCS	20	01:46	41.70	0.9181	0.6970
ASD-POCS-PICCS	100	01:42:22	36.74	0.8957	0.8434

One of the main disadvantages of using advanced regularization is that the most commonly used reconstruction algorithms become significantly more computationally expensive, which can be a bottleneck in their applicability. This is particularly true in image-guided surgery, where very fast reconstructions are needed. In this work, we propose using Krylov methods, a family of iterative methods that show significantly faster convergence than most other commonly used iterative methods.

This is the first time that Krylov methods are used in repeated low-dosage CBCT sequence scanning with image-prior regularization (PIPLE and PICCS). In particular, we combine a technique called IRN approximation, which requires solving a sequence of quadratic subproblems, and we use CGLS as the inner-solver. In this work, we observe that these methods significantly outperform other standard algorithms not just in quality but also in computational time for both simulated and real data.

More specifically, we show an example using heavily undersampled data from a digital head phantom which has no or minimum noise component in the projection space. For this synthetic data, we observe

Table 3. Quantitative results for the thorax phantom example with real measurements involving 50 projections.

Algorithm	Iteration	Time	PSNR	SSIM	HaarPSI
FDK	N/A	00:00.62	31.02	0.6415	0.4403
SIRT	100	00:55	40.68	0.9444	0.6754
CGLS	20	00:09	35.93	0.8302	0.5748
ASD-POCS-TV	100	21:16	36.67	0.9407	0.6626
IRN-TV	20	00:56	38.23	0.8893	0.6152
IRN-PIPLE	20	01:20	44.36	0.9709	0.8405
IRN-PICCS	20	01:43	43.76	0.9450	0.7771
ASD-POCS-PICCS	22 ^a	07:33	36.92	0.9476	0.6755

^a Set up for 100 iterations but stopped early due to stopping criteria.

Table 4. Quantitative results for the thorax phantom example with real measurements involving 20 projections.

Algorithm	Iteration	Time	PSNR	SSIM	HaarPSI
FDK	N/A	00:00.28	27.26	0.4932	0.2971
SIRT	100	00:40	41.20	0.9359	0.6795
CGLS	20	00:07	37.63	0.8301	0.6024
ASD-POCS-TV	100	08:53	37.37	0.9385	0.6528
IRN-TV	20	00:56	40.12	0.8935	0.6572
IRN-PIPLE	20	01:19	43.70	0.9752	0.8006
IRN-PICCS	20	01:41	43.73	0.9480	0.8158
ASD-POCS-PICCS	19 ^a	02:43	38.09	0.9462	0.6596

^a Set up for 100 iterations but stopped early due to stopping criteria.

that both IRN-PIPLE and IRN-PICCS highly perform with only 20 projections, and they have low reconstruction times with respect to the other compared algorithms. Similarly, for an example of a needle surgery dataset with real measurements, IRN-PIPLE and IRN-PICCS provided the best results in the case of having few projections (20 and 50).

One of the objectives of this work was to provide a PICCS regularized algorithm with a computational speed that would make it clinically viable. Results show that, for clinical data, we can obtain reconstruction times of around 100 seconds, significantly faster than the previously best performing algorithm, ASD-POCS-PICCS. The algorithms reported here are timed end-to-end, starting from data loaded in RAM. They are implemented in TIGRE, where the regularization steps for the ASD-POCS-based algorithms and the forward and backprojection operations are performed on the GPU, while the remaining computations in all algorithms are carried out on the CPU. With bespoke implementations in clinical hardware, this could be accelerated further (but not by an order of magnitude), providing a viable algorithm for on-line image guided interventions. Note that the final computational time also critically depends on scanner geometry, image size, and projections acquired, albeit this method is completely flexible to those parameters.

It is however important to note that, strictly speaking, ASD-POCS-PICCS and IRN-PICCS (similarly, IRN-TV and ASD-POCS) are solving the exact same mathematical problem, yet this work show their performance differing in practice. This is likely caused by the faster convergence of Krylov methods, reaching a better solution in significantly fewer iterations than ASD-POCS. However, hyperparameter tuning is also a likely factor in this difference. While Krylov methods only have one hyperparameter to set aside from the number of iterations, ASD-POCS-TV has eight, making finding the best performing set a highly complex task (Lohvithee *et al* 2017). Therefore, even if it is theoretically possible to find a set of parameters and iteration numbers such that both methods perform to the same level of quality, this is not easily done in practice. This is another advantage of Krylov methods: a much simpler set of parameters to set experimentally.

A remaining limitation of the proposed approach is the requirement for accurate alignment between the prior image and the reconstructed volume, which is common in all prior image-guided reconstruction. This is generally handled through a pre-processing step. A reconstruction is first computed without a prior, and this reconstruction using is then aligned to the prior image using deformable registration. Repeating this procedure can further improve reconstruction quality, but it requires solving many sub-problems, which can be efficiently handled using our method. In a clinical setting, this is less important

for rigid-body imaging, such as the head or limbs, but would require more attention in non-rigid areas of the body like the thorax.

Last, a relevant discussion to have at the time of publication of this work is the value of novel numerical algebra reconstruction algorithms in an environment of rapid deep learning (DL) development, and the potential for hybrid approaches that leverage the strengths of both frameworks. It is the authors' opinion that, at the time of writing, there are no published DL methods that produce good results for the size of images presented in this work, and the level of noise and under-sampling we propose. Only very recently CBCT has started to be tackled in a learned reconstruction manner (Biguri *et al* 2025a), and efficient applications of the most sophisticated methods have only started to be proposed in the computational mathematics literature (Vo and Tachella 2026). This is even more challenging if we consider hallucinations, training data bias, and other DL-specific problems that do not apply to Krylov methods. Thus, so far, we do not think there is a comparable DL method in the literature to the work proposed here. However, DL could be used to find a good prior in combination to our method. Namely, in this work we use a zero image as an initial estimate, but instead one could use a denoised (with DL) version of, e.g. the FDK reconstruction as input, starting the algorithm from a much better estimate. This might lead to requiring fewer iterations to converge, accelerating the process, and making the proposed methods DL-compatible. However, the discussion still stands.

5. Conclusions

In this work, we propose two fast and easy to tune algorithms for CBCT imaging when prior images of the same sample/patient are available, using the PIPEL and the PICCS regularization frameworks within a Krylov subspace solver. The resulting algorithms show very good performance, both in image quality and computational time, which makes them very strong algorithmic options for image guided interventions, such as image guided surgery or radiation therapy.


Acknowledgments

AB acknowledges the support of EPSRC Grant EP/W004445/1 and the Accelerate Programme for Scientific Discovery. MSL acknowledges partial support by the U S National Science Foundation, Grant DMS-02208294, and the support by the Hong Kong Innovation and Technology Commission (InnoHK Project CIMDA). CBS acknowledges support from the Philip Leverhulme Prize, the Royal Society Wolfson Fellowship, the EPSRC advanced career fellowship EP/V029428/1, the EPSRC programme Grant EP/V026259/1, and the EPSRC Grants EP/S026045/1 and EP/T003553/1, EP/N014588/1, EP/T017961/1, the Wellcome Innovator Awards 215733/Z/19/Z and 221633/Z/20/Z, the European Union Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie Grant Agreement No. 777826 NoMADS, the Cantab Capital Institute for the Mathematics of Information and the Alan Turing Institute. This research was supported by the NIHR Cambridge Biomedical Research Centre (NIHR203312). The views expressed are those of the author(s) and not necessarily those of the NIHR or the Department of Health and Social Care.

Data availability statement

No new data were created or analysed in this study.

ORCID iDs

S M Ragib Shahriar Islam  0009-0005-8620-3179

Malena Sabaté Landman  0000-0003-3215-2773

Carola-Bibiane Schönlieb  0000-0003-0099-6306

Ander Biguri  0000-0002-2636-3032

References

- BaumueLLer S, Winklehner A, Karlo C, Goetti R, Flohr T, Russi E W, Frauenfelder T and Alkadhi H 2012 Low-dose CT of the lung: potential value of iterative reconstructions *Eur. Radiol.* **22** 2597–606
- Biguri A et al 2025 Advancing the frontiers of deep learning for low-dose 3d cone-beam ct reconstruction *IEEE Open J. Signal Process.* **6** 942–9
- Biguri A et al 2025 TIGRE v3: efficient and easy to use iterative computed tomographic reconstruction toolbox for real datasets *Eng. Res. Express* **7** 015011
- Biguri A, Dosaanjh M, Hancock S and Soleimani M 2016 TIGRE: a MATLAB-GPU toolbox for CBCT image reconstruction *Biomed. Phys. Eng. Express* **2** 055010
- Breger A et al 2025 *A Study on the Adequacy of Common IQA Measures for Medical Images* (Springer) pp 451–62
- Burström G, Cewe P, Charalampidis A, Nachabe R, Söderman M, Gerdhem P, Elmi-Terander A and Edström E 2021 Intraoperative cone beam computed tomography is as reliable as conventional computed tomography for identification of pedicle screw breach in thoracolumbar spine surgery *Eur. Radiol.* **31** 2349–56
- Chen G-H, Tang J and Leng S 2008 Prior image constrained compressed sensing (PICCS): a method to accurately reconstruct dynamic CT images from highly undersampled projection data sets *Med. Phys.* **35** 660–3
- Chen Y, Liu J, Xie L, Hu Y, Shu H, Luo L, Zhang L, Gui Z and Coatrieux G 2017 Discriminative prior-prior image constrained compressed sensing reconstruction for low-dose CT imaging *Sci. Rep.* **7** 13868
- Chung J and Gazzola S 2019 Flexible krylov methods for ell_p regularization *SIAM J. Sci. Comput.* **41** S149–71
- Daubechies I, DeVore R, Fornasier M and Güntürk C 2010 Iteratively reweighted least squares minimization for sparse recovery *Commun. Pure Appl. Math.* **63** 38
- Feldkamp L A, Davis L C and Kress J W 1984 Practical cone-beam algorithm *J. Opt. Soc. Am. A* **1** 612–9
- Fikuart M, Bullert B, Vetter S Y, Franke J, Gruetzner P A and Swartman B 2025 Effect of the cone-beam CT acquisition trajectory on image quality in spine surgery: experimental cadaver study *Spine J.* **25** 154–64
- Gazzola S and Sabaté Landman M 2019 Flexible GMRES for total variation regularization *BIT Numer. Math.* **59** 721–46
- Gulias-Soidan D, Crus-Sanchez N M, Fraga-Manteiga D, Cao-González J I, Balboa-Barreiro V and González-Martín C 2020 Cone-beam CT-guided lung biopsies: results in 94 patients *Diagnostics* **10** 1068
- Hansen P C 2010 *Discrete Inverse Problems* (SIAM)
- Hatamikia S, Biguri A, Kronreif G, Russ T, Kettenbach J and Birkfellner W 2023 Source-detector trajectory optimization for CBCT metal artifact reduction based on PICCS reconstruction *Z. Med. Phys.* **34** 565–79
- Hestenes M R et al 1952 Methods of conjugate gradients for solving linear systems *J. Res. Natl Bur. Stand.* **49** 409–36
- Irie K, Murayama Y, Saguchi T, Ishibashi T, Ebara M, Takao H and Abe T 2008 Dynact soft-tissue visualization using an angiographic C-arm system: initial clinical experience in the operating room *Oper. Neurosurg.* **62** 266–72
- Kak A C and Slaney M 1988 IEEE engineering in medicine and biology society *Principles of Computerized Tomographic Imaging* (IEEE Press) Published under the sponsorship of the IEEE Engineering in Medicine and Biology Society
- Kak A C and Slaney M 2001 *Principles of Computerized Tomographic Imaging* (SIAM)
- Kawauchi S, Chida K, Moritake T, Hamada Y and Tsuruta W 2021 Radioprotection of eye lens using protective material in neuro cone-beam computed tomography: estimation of dose reduction rate and image quality *Phys. Medica* **82** 192–9
- Kickuth R, Reichling C, Bley T, Hahn D and Ritter C 2015 C-arm cone-beam CT combined with a new electromagnetic navigation system for guidance of percutaneous needle biopsies: initial clinical experience *RöFo-Fortschritte auf dem Gebiet der Röntgenstrahlen und der Bildgebenden Verfahren* vol 187 (© Georg Thieme Verlag KG) pp 569–76
- Kulvait V and Rose G 2021 Software implementation of the krylov methods based reconstruction for the 3d cone beam CT operator
- Lanza A, Morigi S, Selesnick I and Sgallari F 2017 Nonconvex nonsmooth optimization via convex–nonconvex majorization–minimization *Numer. Math.* **136** 06
- Lauzier P T, Tang J and Chen G-H 2012 Prior image constrained compressed sensing: implementation and performance evaluation *Med. Phys.* **39** 66–80
- Lee H, Xing L, Davidi R, Li R, Qian J and Lee R 2012 Improved compressed sensing-based cone-beam CT reconstruction using adaptive prior image constraints *Phys. Med. Biol.* **57** 2287
- Liu L 2014 Model-based iterative reconstruction: a promising algorithm for today's computed tomography imaging *J. Med. Imaging Radiat. Sci.* **45** 131–6
- Lohvithee M, Biguri A and Soleimani M 2017 Parameter selection in limited data cone-beam CT reconstruction using edge-preserving total variation algorithms *Phys. Med. Biol.* **62** 9295
- Mileto A, Guimaraes L S, McCollough C H, Fletcher J G and Yu L 2019 State of the art in abdominal CT: the limits of iterative reconstruction algorithms *Radiology* **293** 491–503
- Mueller J L and Siltanen S 2012 *Linear and Nonlinear Inverse Problems With Practical Applications* (Society for Industrial and Applied Mathematics)
- Noël P B, Engels S, Köhler T, Muenzel D, Franz D, Rasper M, Rummeny E J, Dobritz M and Fingerle A A 2018 Evaluation of an iterative model-based CT reconstruction algorithm by intra-patient comparison of standard and ultra-low-dose examinations *Acta Radiol.* **59** 1225–31
- Sabaté Landman M, Biguri A, Hatamikia S, Boardman R, Aston J and Schönlieb C-B 2023 On Krylov methods for large-scale CBCT reconstruction *Phys. Med. Biol.* **68** 155008
- Scherzer O, Grasmair M, Grossauer H, Haltmeier M and Lenzen F 2009 *Variational Methods in Imaging* (Springer)
- Sidky E Y and Pan X 2008 Image reconstruction in circular cone-beam computed tomography by constrained, total-variation minimization *Phys. Med. Biol.* **53** 4777
- Soodhalter K M, de Sturler E and Kilmer M E 2020 A survey of subspace recycling iterative methods *GAMM-Mitteilungen* **43** e202000016
- Srinivasan V M, Schafer S, Ghali M G Z, Arthur A and Duckworth E A M 2016 Cone-beam CT angiography (Dyna CT) for intraoperative localization of cerebral arteriovenous malformations *J. Neurointerv. Surg.* **8** 69–74
- Stayman J W, Dang H, Ding Y and Siewerdsen J H 2013 PIRPLE: a penalized-likelihood framework for incorporation of prior images in CT reconstruction *Phys. Med. Biol.* **58** 7563
- Vo R and Tachella J 2026 Efficient unrolled networks for large-scale 3d inverse problems
- Wohlberg B and Rodriguez P 2008 An iteratively reweighted norm algorithm for minimization of total variation functionals *IEEE Signal Process. Lett.* **14** 948–51