

PAPER • OPEN ACCESS

TIGRE v3: Efficient and easy to use iterative computed tomographic reconstruction toolbox for real datasets

To cite this article: Ander Biguri *et al* 2025 *Eng. Res. Express* **7** 015011

View the [article online](#) for updates and enhancements.

You may also like

- [Experimental investigation on suppression of gearbox vibrations using particle damping technique](#)
Kuldip A Patil, Kaveri A Katake, Kailasnath B Sutar et al.
- [Rolling bearing fault diagnosis method based on IAGTO-VMD and kurtosis-MIC criterion](#)
Bailin Li, Song Tang, Siyuan Chen et al.
- [Assessing fingerprint image quality: an ensemble technique perspective](#)
Parul Choudhary, Pooja Pathak and Phalguni Gupta

Engineering Research Express



PAPER

OPEN ACCESS

RECEIVED

17 January 2025

REVISED

14 February 2025

ACCEPTED FOR PUBLICATION

27 February 2025

PUBLISHED

7 March 2025

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.



TIGRE v3: Efficient and easy to use iterative computed tomographic reconstruction toolbox for real datasets

Ander Biguri¹ , Tomoyuki Sadakane² , Reuben Lindroos¹, Yi Liu³, Malena Sabaté Landman⁴, Yi Du^{5,6}, Manasavee Lohvithee⁷ , Stefanie Kaser⁸ , Sepideh Hatamikia^{9,10}, Robert Bryll¹¹ , Emilien Valat¹², Sarinrat Wonglee¹³, Thomas Blumensath¹⁴ and Carola-Bibiane Schönlieb¹

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

² J. Morita Manufacturing Corp., Japan

³ X-ray Imaging, Sensing and Sorting, CSIRO Mineral Resources, Sydney, Australia

⁴ Mathematical Institute, University of Oxford, Oxford, United Kingdom

⁵ Key Laboratory of Carcinogenesis and Translational Research (Ministry of Education/Beijing), Department of Radiation Oncology, Peking University Cancer Hospital & Institute, People's Republic of China

⁶ Institute of Medical Technology, Peking University Health Science Center, People's Republic of China

⁷ Department of Nuclear Engineering, Chulalongkorn University, Bangkok, Thailand

⁸ Institute of High Energy Physics, Austrian Academy of Sciences, Vienna, Austria ¹⁵

⁹ Research Center for Clinical AI-Research in Omics and Medical Data Science (CAROM), Department of Medicine, Danube Private University (DPU), Krems, Austria

¹⁰ Austrian Center for Medical Innovation and Technology (ACMIT), Wiener Neustadt, Austria

¹¹ Mitutoyo Research & Development America Inc., Kirkland, United States of America

¹² Department of Mathematics, KTH Royal Institute of Technology, Stockholm, Sweden

¹³ Nuclear Research and Development Center, Thailand Institute of Nuclear Technology (TINT), Nakhon Nayok, Thailand

¹⁴ Institute of Sound and Vibration Research (ISVR), University of Southampton, Southampton, United Kingdom

¹⁵ Now at Institute for Biomedical Engineering, ETH Zürich and University of Zürich, Switzerland.

E-mail: ander.biguri@gmail.com

Keywords: tomography, software, CBCT, inverse problems, medical imaging

Abstract

Computed Tomography (CT) has been widely adopted in medicine and it is increasingly being used in scientific and industrial applications. Parallely, research in different mathematical areas concerning discrete inverse problems has led to the development of new sophisticated numerical solvers that can be applied in the context of CT. The Tomographic Iterative GPU-based Reconstruction (TIGRE) toolbox was born almost a decade ago precisely in the gap between mathematics and high performance computing for real CT data, providing user-friendly open-source software tools for image reconstruction. However, since its inception, the tools' features and codebase have had over a twenty-fold increase, and are now including greater geometric flexibility, a variety of modern algorithms for image reconstruction, high-performance computing features and support for other CT modalities, like proton CT. The purpose of this work is two-fold: first, it provides a structured overview of the current version of the TIGRE toolbox, providing appropriate descriptions and references, and serving as a comprehensive and peer-reviewed guide for the user; second, it is an opportunity to illustrate the performance of several of the available solvers showcasing real CT acquisitions, which are typically not be openly available to algorithm developers.

1. Introduction

Tomography, and particularly computed tomography (CT), is a non-destructive imaging technique that has recently become essential for visualizing the internal structure of very different objects, gaining widespread popularity among engineers, biologists, geophysicists, material scientists, and medical professionals. Since tomographic imaging relies on indirect measurements, it requires using reconstruction algorithms which can produce images with varying properties and qualities, see, e.g. [1, 2].

The mathematical model corresponding to CT problems has two main properties that make it challenging to solve, so decades after the inception of tomography it still remains of high research interest. First, in the limited measurements regime and in the presence of noise, the reconstructions are very sensitive to small perturbations in the measurements. This is called ill-posedness, see, e.g. [1, 3]. Moreover, these are scenarios of capital interest, i.e. when one reduces radiation doses in medical CT, reduces scanning time in industrial CT, or when physical limitations reduce the amount of available measuring angles. Second, modern applications tend to require bigger reconstructions due to an increase in resolution, in object size, or in the amount of measurements [4].

There is a constant need for efficient algorithms that can support a trend towards more ill-posed and higher dimensional CT problems, while maintaining image reconstruction quality. This can be achieved by leveraging the noisy measurements with prior information about the reconstruction. Many algorithms exist for this purpose: ranging from analytical methods, algebraic row-action methods, variational model-based methods or, more recently, machine learning based reconstruction methods.

However, algorithm developers often lack the expertise or resources to create efficient, lab-ready software, so users tend to optimize scanning procedures for proprietary software that typically uses direct reconstruction algorithms (e.g., filtered backprojection, FBP), requiring high-quality, high-dose, over-sampled data. As a result, many practitioners still rely on direct classical solvers, despite evidence showing that iterative algorithms can significantly improve reconstruction quality [5, 6].

The Tomographic Iterative GPU-based Reconstruction (TIGRE) toolbox [7] is an open source translational research tool designed to serve both algorithm developers (applied mathematicians and computer scientist without a software engineering background) and applied scientist working with real CT data. In this spirit, the TIGRE toolbox offers an easy to install platform where one can both reconstruct real data using advanced algorithms (with only a very basic level of mathematics and software knowledge) or write new mathematical methods for reconstruction and test them on high quality realistic test problems (with little knowledge of data managing pipelines or GPU-acceleration). Moreover, these new implemented and tested algorithms can be afterwards appended to the TIGRE toolbox to make them available to researchers wanting to try new algorithms on their data. Since its first release, the TIGRE toolbox has indeed become a community platform, used by researchers and scientists across the field of tomography. Indeed, since its publication the codebase of the toolbox has grown in over an order of magnitude, and, to the date, it has been cited in over 300 papers.

It is crucial to emphasize the significance of open source software (OSS) within academia. In order to create reproducible and reliable research, there is a need for free and transparent unified frameworks that offer a space for different research communities to meet and promote knowledge-exchange. The TIGRE toolbox contributes and enriches the increasingly wide ecosystem of OSS available for inverse problems in computed tomography, which differ in their approach to software development and distribution to applied scientists. For general x-ray tomography, other GPU-based OSS exist, such as the widely used ASTRA toolbox [8] and other software that uses it. This is for example the case of the PyTorch [9] compatible Tomosipo [10], or the CIL toolbox [11] which also strongly focuses on the μ CT community. The latter uses both ASTRA and TIGRE as GPU projectors, yet providing several user-friendly tools to easily craft new iterative algorithms suited to the relevant experiments. On the other spectrum of target users, RTK [12] provides reconstruction code for data associated to medical scanners. Lastly, the TomoPy toolbox [13] provides similar tools to the rest of the mentioned packages (including ASTRA wrappers), but was initially designed for synchrotron x-ray tomography. Recently, LEAP [14] was also released, with auto-differentiable operators. Software focused on other types of tomography is also available. For example EIDORS [15] or pyEIT [16] for Electrical Impedance Tomography, STIR [17] for emission tomography (such as Positron Emission Tomography and Single Particle Emission Computed Tomography), and SIRF [18] or CASToR [19] for multi-modal tomography. Finally, other OSS that are mostly used by the discrete inverse problems (mathematics) community can be used for tomography ODL [20], AIRtools [21] and IRtools [22], among others. However, these packages sometimes lack the required efficiency to support high dimensional realistic CT test problems. This is not an exhaustive list of the available software for tomography reconstruction. Each piece of software has nuanced differences in use cases, programming languages, licensing and support, which are out of the scope of this work to discuss.

This work focuses on the TIGRE toolbox most current version 3.0, offering a structured overview of the toolbox including appropriate descriptions and references, and highlighting some of the most important features added since its first release. Subsequently, this paper aims to serve as a comprehensive and peer-reviewed guide for the user. Moreover, this work also aims to illustrate the performance of several of the solvers that are available on the TIGRE toolbox on real CT acquisitions. It is also worth highlighting that TIGRE has functioning demos for all available features, together with explanations of the algorithms. These can be found in the Demos folder of the library, for each programming language.

The article is structured as follows. Section 2 briefly introduces the mathematical background of iterative reconstruction algorithms, focusing on the algorithms currently implemented in the TIGRE toolbox. In this section, a structured description of the design of the TIGRE toolbox is also included, as well as highlights of the

several new important features of the software such as novelties in the implementation, geometric flexibility and multi-GPU support. Next, section 3 presents five different examples based on real datasets, alongside the code needed to reproduce them. These correspond to a clinical scenario with patient data, a synchrotron parallel beam tomographic reconstruction, a Monte Carlo simulation of a proton CT problem, a neutron tomographic acquisition and an industrial μ CT dataset. Finally, some conclusions are given in section 4. A list of all the solvers provided by the TIGRE toolbox in its current state can be found in Appendix A.

The article contains code snippets, arbitrarily chosen to be in Python or MATLAB in order to have running examples in both languages. However, unless explicitly stated, the exact same code can be used in either language by simply replacing the basic syntax of the calls, i.e. the function names and traces are shared between languages.

2. Methods

This section provides both a context for the TIGRE toolbox and a description of the most important features that have been added since it was first published, almost a decade since this article. First, a brief description of the tomography reconstruction problem is given, along with a (non-exhaustive) review of iterative solvers. Afterwards, a brief description of the software architecture is included with the relevant references. Last, this section describes all the significant features that have been added to the toolbox, giving an in-depth description for those that have not been published elsewhere. These are of very different nature and have been added in the following order:

- newly implemented mathematical optimization algorithms,
- new Python codebase implementation,
- new features concerning the applicability of the codes to real data:
 - increased geometric flexibility of the operators,
 - multi-GPU features for high performance computing of large tomographs,
 - automatic data loaders,
 - pre-processing steps for proton CT;
- PyTorch wrapper implementation (given the growing interest of data-driven methods in reconstruction).

2.1. Tomographic reconstruction and iterative algorithms

The CT measuring (or forward) process can be modelled mathematically using the Radon transform, which is an integral operator that considers line integrals over the domain. The most common approach to CT reconstructions is to consider the large scale linear system of equations arising from the discretization of such measuring process, which is of the form:

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

where A is the (known) system matrix describing the forward operation; x is the (unknown) vector of image values in lexicographical order; b is the (known) vector of measured (noisy) linear attenuation values in lexicographical order; and \tilde{e} represent the noise affecting the measurements (both measurement noise and model errors introduced by the linearization of the real forward operator). The least-squares solution to equation (1) can be found using the Moore–Penrose pseudoinverse of the forward operator:

$$x = A^\dagger b, \quad (2)$$

however, this can be far from the true solution x due to noise amplification.

For CT problems, approximations of the analytical solution to the original Radon transform problem exist, the discretization of which is also known as the Filtered Backprojection (FBP) in the CT community, or the Feldkamp-Davis-Kress (FDK) algorithm for Cone Beam CT (CBCT). These methods produce useful reconstructions when the problem is not very ill-posed and there is a low level of noise in the measurements. This is also the most commonly used approach in practice, generally defaulting the scanning settings to maximize the quality of the results, assuming the FBP algorithm is used for the reconstruction, rather than considering alternatives to obtain meaningful CT reconstructions with lower dose or scanning time.

However, there are great advantages associated to reducing the x-ray dose and scanning times in terms of cost and safety. Moreover, it is sometimes the case where the physical limitations of the problem only allow a limited sampling space (e.g. limited angle tomography). In all these scenarios, the corresponding problem in

equation (1) is ill-posed, or, in other words, small perturbations on the measurement data can cause large perturbations in the reconstruction. This means that, in the presence of noise, the approximate solution obtained using equation (2) will be corrupted or even totally dominated by the noise amplification. To alleviate this, one needs to resort to regularization. One of the most well known approaches to regularization is the use of variational methods, which consist on solving the minimization problems of the form:

$$\min_x D(Ax, b) + \alpha R(x), \quad (3)$$

given a distance function $D(\cdot, \cdot)$ and a regularization function $R(\cdot)$ that penalizes constraints in the solution and can be related to the *prior* of the image if one sets the original problem (1) in the Bayesian framework [23]. Here, the *hyperparameter* α controls the trade-off between the image fitting the measured data and the image fitting the prior information. Usually, the choice of the distance function $D(\cdot, \cdot)$ is informed by the assumptions on the distribution of the unknown noise. In CT applications, the most common choice for this is to use the 2-norm, because the measured noise is well approximated by a Gaussian distribution after the standard log-transform of the data [4, Chapter 2.3.2]. However, some approaches consider the noise associated to the photon counts in the raw data, which has a Poisson distribution, so a different distance function $D(\cdot, \cdot)$, called Kullback–Leibler (KL) divergence, needs to be considered instead.

Different algorithms for given combinations and choices of D and R have been provided over the years by the optimization and linear algebra communities, with different behaviours and convergence properties that can lead to superior image quality. These methods, almost always iterative in nature, are covered by the generic term *iterative algorithms*. For example, a common method to solve any convex optimization method is the Gradient Descend (GD) method. Given the distance function of the 2-norm and no regularization, one can obtain a solution of equation (3) using the iterative update

$$x^k = x^{k-1} + \lambda A^T (Ax^{k-1} - b), \quad (4)$$

where k is the iteration number and λ a step size, desirably inversely proportional to the norm of A . Similarly, a common algorithm to solve the same problem given the KL-divergence is the Maximum Likelihood Expectation Maximization (MLEM) algorithm, with an update of the form

$$x^k = \frac{x^{k-1}}{A1} A^T \frac{b}{Ax^{k-1}}, \quad (5)$$

given 1 an all-ones vector.

There is an abundance of alternative methods to solve CT reconstruction problems, some accepting variants with explicit regularization and some with implicit regularization properties instead. These have significant differences in convergence behaviour and solution characteristics. TIGRE provides a platform with a vast array of the most common solvers, implemented in such a way that they can be directly used with real datasets. In particular, TIGRE has the following algorithms, categorized in different classes:

- **Direct methods.** These methods implement an analytic expression of the pseudoinverse, as in equation (2), or suitable approximations of it. Included are FBP and FDK[24], with various filter types and geometric correction methods, such as Wang [25] and Parker [26] weights for detector displacement and angular range corrections, respectively.
- **Kaczmarz-type algorithms.** These methods are generalizations of GD: the approximate images are updated similarly to (4) but with a reweighing of the rows and columns to normalize for length. They were rediscovered as the Algebraic Reconstruction Technique (ART) [27] in 1970 for CT with the inclusion of non-negativity constraints. Different modifications include the Simultaneous version (SART) [28], the Ordered Subsets (OS) version¹⁶ (OS-SART) [29] and the confusingly named Simultaneous Iterative Reconstruction algorithm (SIRT) [1]. The main difference between them is the update size.
- **Projection Onto Convex Sets (POCS) algorithms.** These algorithms allow solving two convex optimization problems jointly, in the TIGRE implementation, these correspond specifically to the ones related to the data constraints and the regularization. Given a 2-norm data constraint, and various flavours of Total Variation (TV) regularization (promoting pairwise smooth images), several POCS-like algorithms are available. These include the well-known Adaptive Steepest Descend version (ASD-POCS) [30] and the Projection-Controlled Steepest Descent (PCSD) [31], which reduces the required parameter selection. Versions of these two algorithms exists for the Adaptive weighed TV norm [32, 33]¹⁷, and OS versions exist for the first with both TV norms A version that uses a Bregman outer iteration is also available, B-ASD-POCS- β [34].

¹⁶ See section 2.2.2

¹⁷ See section 2.2.3

- **Krylov Subspace algorithms.** These are methods that project the original optimization problem into small relevant subspaces of increasing dimension, producing fast converging algorithms with intrinsic regularization properties. This can be used in combination with explicit 2-norm regularization. In the following, the word ‘hybrid’ is used to denote explicit 2-norm regularization on the projected problem [35]. In TIGRE, the following are implemented allowing, when possible, to compute the regularization parameter automatically: Conjugate Gradient Least Squares (CGLS) [36], LSQR [37], hybrid LSQR [37], AB/BA- Generalized Minimal Residual method (GMRES) [38], LSMR [39]. Moreover, two different methods that allow for approximations of the TV regularization are also implemented: using an Iterative Reweighed Norm (IRN) approach IRN-TV-CGLS [40] (this is an inner-outer scheme where quadratic approximations of the functional are solved using CGLS) and a hybrid-flexible approach: (f)LSQR-TV [41] (which avoids inner-outer loops by using a single flexible Krylov subspace).
- **Statistical minimization algorithms.** These methods are based on the understanding of the problem (3) as a Bayesian minimization problem tied with the KL-divergence distance function. MLEM has already been introduced in equation (5), and its OS-counterpart, widely used in emission tomography, OSEM[42], is also available.
- **Proximal algorithms.** More recently, optimization algorithms that combine *proximal* operators have been shown to be successful in image reconstruction. For the 2-norm and TV proximal operators, TIGRE provides the Fast Iterative Shrinkage-Thresholding algorithm (FISTA) [43] and further improvements for faster speed [44]. Moreover, TIGRE also includes the SART-TV algorithm [45], which was not originally proposed within this framework, but that can also be interpreted as a proximal method.

Some of the algorithms mentioned above were already available in the original TIGRE release, but many have been added over the years into its current version. New introductions to TIGRE are described in more detail in the following sections wherever appropriate, but the authors suggest reading the original articles for a more in depth description. A summary of all the available algorithms (as of this date), can be found in Appendix A, along with their corresponding citations, both in the original literature and in the corresponding TIGRE-specific implementation papers.

2.2. Improvements on the iterative reconstruction tools in TIGRE

This section describes the enhancement of iterative reconstruction algorithms available in TIGRE since the first release. This section is divided into three, describing first the addition of Krylov Subspace algorithms to TIGRE, then Ordered Subsets (OS) types of block-action algorithms, and finally a brief description of various minor yet useful features and regularization tools included since its first release.

2.2.1. Krylov subspace algorithms

Krylov subspace methods are a class of iterative solvers for linear equations of the form (1). Because they are naturally regularizing, they are especially suited for inverse problems. Moreover, these are projection methods which only require matrix-vector products with the system matrix (and its adjoint), so they are very suited to solve large-scale problems. The TIGRE toolbox features a collection of some of the most standard Krylov Subspace methods for non-square matrices (CGLS, LSQR, LSMR), possibly in combination with Tikhonov regularization, as well as recent developments including total variation (TV) regularization and mismatched backprojectors. An in detail explanation of the algorithms as well as its TIGRE implementation can be found in [46].

2.2.2. Ordered subsets algorithms

Ordered Subsets (OS) methods are named after the fact that only a subset of the data is used to update the solution of the tomographic reconstruction problem at each iteration. Particularly, at iteration k , the solution x^k is updated using a row-wise subsample of the matrix A , namely A_s , and therefore only the corresponding part of the measurements b , namely b^s . These algorithms are most commonly known as *block-iterative methods* in the mathematical community [47]. For example, similarly to equation (4), an approximate gradient step for the 2-norm distance for solving equation (3) with OS would be

$$x^k = x^{k-1} + \lambda A_s^T (A_s x^{k-1} - b_s), \quad (6)$$

where s defines a row-wise block of the matrix A which changes at each iteration.

Compared to algorithms that use the entire matrix in their update step (e.g. SIRT), OS algorithms require more (partial) updates to use all the measured data. However, they are well known for accelerating the rate of convergence, in the sense that many fewer iterations (defined as the number of times that the whole matrix A is seen) are required to produce a reconstruction of comparable quality. On the other extreme, algorithms that

update the solution projection by projection display faster converge in terms of the number times that one needs to access the whole system matrix, but they are often very slow in wall-time due to the large computational cost of having to update the solution a higher amount of times. Therefore, OS algorithms produce a reasonable middle ground algorithm between fast mathematical convergence and fast runtime [48]. This is why their use in CT is common, particularly in Kaczmarz-like methods (e.g., OS-SART [49]) and the most used algorithm in emission tomography like PET and SPECT is of this family (namely OS-EM [42]). Moreover, there are many algorithms that can be used in combination with an OS scheme. This is the case of methods designed for solving problems with explicit regularization that require update steps that are only related to the data-minimization term, without any particular prescriptions on the form of this update. In the original version of TIGRE, OS-SART was available. Since, ASD-POCS, AwASD-POCS, PCSD, AwPCSD and MLEM now have an OS counterpart.

When invoking these algorithms, a parameter that regulates the size of each block (in number of projections per block) can be selected by the user, as seen in Snippet 2.2.2 for ASD-POCS and its block-action counterpart. Updating can be done sequentially, randomly or via maximum information updates.

Code Snippet 1: OS-ASD-POCS algorithm in Python

```
1 import tigre. Algorithms as algs
2 n_iter=10
3 res_asd_pocs = algs.asd_pocs(projections, geo, angles, n_iter)
4 res_os_asd_pocs = algs.os_asd_pocs(projections, geo, angles, n_iter, blocksize=20)
```

2.2.3. Algorithm features

New features that are useful for most algorithms have been introduced to TIGRE since its original release. Here is a short summary:

- **Residual computing.** The residual norm history $\|Ax^k - b\|$, for each iteration number k , is an interesting value to track when comparing iterative methods, as it directly describes how well the current reconstruction of the image fits the data. This computation per iteration is costly (it requires an extra forward projection) but TIGRE now allows users to ask for this residual, if desired, for algorithm performance comparison. Note that Krylov methods include a way of estimating this quantity at no extra cost per iteration. Last, note that an exact fit is not desired, as that would also fit the noise present within the sinogram b .
- **Ground truth comparison.** The error norm histories $\|x_{gt} - x^k\|$, given a known ground truth image x_{gt} , are also of capital interest when studying algorithm performance. Note that, due to the ill-posedness of the problem, metrics solely based on the residual norm might be misleading. TIGRE now allows to input a ground truth to be compared against the current iterate during reconstruction.
- **Automatic regularization selection.** Krylov subspace algorithms have a single regularization hyperparameter when used in combination with variational regularization, i.e. α in equation (3), and there exist extensive theory on how to chose this parameter optimally [35]. Some algorithms in TIGRE include therefore automatic regularization parameter selection. If the noise-level value is either known or can be estimated, it can be inputted as a parameter of the algorithm to guide the choice of regularization parameter in the optimization problem, using the so-called Discrepancy Principe (DP) [50]. Otherwise, the Generalized Cross Validation (GCV) [51] is used instead to update the regularization parameter at each iteration. For more details on the implementation see [46].
- **Adaptive weighted total variation.** The well known Total Variation (TV) regularization promotes piece-wise constant reconstructions. However, this can sometimes have a significant effect in the shape of sharp edges, often the most important part of CT image analysis. Adaptive weighted TV (AwTV) [32], introduces a penalty term $w_{i,j}$ in order to only account for small variations in pixel values (mostly caused by noise), but not penalize large ones (which are more likely to represent real edges). Thus, the new TV term with adaptive weights (for the two-dimensional case) is:

$$\text{AwTV}(\mathbf{x}) = \sum_{i,j=1}^{n-1} \sqrt{w_j(x_{i,j} - x_{i,j-1})^2 + w_i(x_{i,j} - x_{i-1,j})^2}, \quad (7)$$

with weights defined as

$$w_i = e^{-\left(\frac{x_{i,j} - x_{i-1,j}}{\delta}\right)^2}, \quad w_j = e^{-\left(\frac{x_{i,j} - x_{i,j-1}}{\delta}\right)^2}, \quad (8)$$

with a parameter δ , that produces exactly the TV norm when $\delta \rightarrow \infty$. The AwTV norm can be now used with any POCS-type algorithm, and a gradient descend minimizer of the function is available for users.

- Data redundancy weighting on iterative algorithms. The data redundancy weighting applied in FDK for the cases where the detector is offsetted have been shown to improve reconstruction of SART and gradient descend-like algorithms [52, 53]. All algorithms that accept this type of weighing now have it by default.

2.3. Python implementation

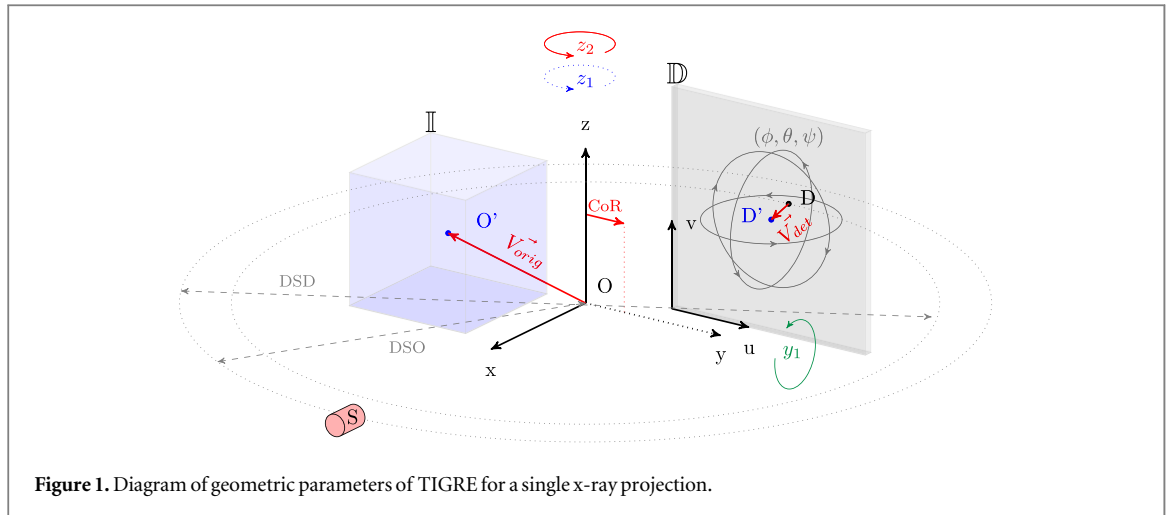
On its original release, TIGRE was only a MATLAB toolbox, which limited its definition as a free and open source toolbox, as MATLAB is neither free nor open source itself. Now, all TIGRE's MATLAB codebase has been duplicated into the Python language, and can be found in the same repository at github.com/CERN/TIGRE/tree/master/Python, including demos of all features. The structure and functionality of both implementations is the same, however, few technical differences exist: such as memory unrolling ordering for multidimensional arrays, the use of Cython rather than MATLABs MEX interface or the import hierarchy that does not exist in its counterpart.

Moreover, there is a notable difference between the MATLAB and the Python releases which goes past the technical aspects: the Python version uses an Object Oriented Programming (OOP) paradigm with a system of classes and inheritance, where all algorithm are children of a `IterativeReconAlg` class that sets some common methods and attributes and offers a framework for other algorithms to build upon. Examples of the Python interface are given across this document.

2.4. Geometric flexibility

In its original form, TIGRE allowed for most standard circular CBCT geometries, such as detector shifts, or reconstructions of images not centered in the axis of rotation. However, since the first publication [7], it has been evident that more complex geometries are both required for several non-standard but common scans, such as helical CT, tomosynthesis [54] or laminography [55], among others. Moreover, researchers on CT methods are showing an increased interest in exploring with complex scanning paths and setups, and having the right computational tools for both producing simulations and reconstructions can be of crucial importance. Thus, TIGRE's geometric flexibility has been updated with several important features:

- Detector in-place rotation. The detector is assumed to lie in the orthonormal plane to the source-to-detector line. However, this is not necessarily true either by design or by mechanical inaccuracies. For example, scanners settings like those used in laminography or tomosynthesis purposely have measurements where the detector is not orthonormal to the source. This can be easily expressed (equivalently) by a rotation of the detector instead of a circular CBCT trajectory. Moreover, many scanners have mechanical inaccuracies that arise after use, and calibration methods exists to correct such changes algorithmically [56]. TIGRE now accepts in-place (ϕ, θ, ψ) rotation angles for the detector.
- Center of Rotation (CoR) correction. In μ -CT scans, where the sample rotates, rather than the machine, it is standard that the axis of rotation is slightly shifted away from the source-to-detector line. While this could be described with a changing image combining a detector shift and a detector in-place rotation, it is significantly easier to describe as a shift in the axis of rotation itself, as TIGRE allows now.
- Arbitrary axis of rotation. To fully describe any geometry, one would also need to allow for any arbitrary axis of rotation, not only a vertical one, as a standard circular CBCT machine produces. This enables the use of TIGRE not only on experimental lab setups, but also in more standard C-arm CBCT machines. The implementation allows now to include Euler ZYZ (z_1, y_1, z_2) angles as gantry angle rotations. In the diagram presented in figure 1, a positive angle of rotation corresponds to an anti-clock wise turn, as represented by the arrow in the rotation angle symbols.
- Curved detectors. While most CBCT detectors are planar, 2D CT detectors are often curved (with equidistant source-detector distance in each pixel). Recently, some experimentation is being done on curved CBCT detectors. In that spirit, TIGRE now allows to consider curved detectors by applying a 'flattening' transformation. When projecting data from a curved detector onto a virtual flat detector, the pixel size of the flat detector is determined by the radial projection of an arc segment from the curved detector. The values are then interpolated to ensure that the pixels on the virtual flat detector are uniformly sampled. As a consequence, due to the geometry of the projection, the virtual flat detector ends up with more pixels than the curved detector because each pixel on the flat detector represents a smaller portion of the curved detector's surface.



- Per-projection geometry. In order to allow a fully flexible geometry for any arbitrary CBCT scan that exists and could be imagined, the defined geometry parameters should be able to vary for each projection measured. Thus, TIGRE now allows for all geometry defining parameters, except the image and detector size in both voxels and real world units, to be defined individually for each projection. This allows for arbitrary scanning geometries, as shown in [57].

All this considered, the geometric diagram of TIGRE shown in the original release (figure 1 in the original article [7]), should be updated to figure 1 presented here.

Figure 1 can be better understood if we provide a brief description of the default scanning geometry assumed by TIGRE. Here, the scanner is assumed to rotate around a fixed, static object. This volume is centered at the Oxyz axis, which does not rotate with the machine, but the origin O is allowed for offsets, so some other scanning modes, such as the helical scanning can be realized. The detector is assumed to be centered in the source-origin (S-O) line, at point D (with distance DSO from the origin). An auxiliary axis exists for the detector, uv, with an origin in the bottom left corner of the detector and always aligned with the detector itself. The rest of the parameters in the geometry definition of TIGRE, which allow for modifications of the default settings, are described in Appendix B.

2.5. Multi-GPU support

Since CT reconstruction is a very computationally heavy task, it benefits from access to any computational resources available. With the raise of multi-GPU machines, now a staple computational resource in most CT laboratories and hospitals alike, TIGRE has been extended to support multiple-GPU computing [58]. This allows faster computations in multiple GPU machines and also enables larger-than-memory computations.

The TIGRE implementation of this support is done via a novel memory management approach that seamlessly distributes any computationally heavy operation (forward and backprojector, and regularization optimizers) across multi-GPUs (either user selected, or all available), in such a way that the computation scales linearly with the number of GPUs. In addition to providing a faster execution time, the splitting algorithm that TIGRE implements can be used to reconstruct images that are larger than the GPU memory. This is, for example, a very common scenario in μ -CT applications, where both the projection and the image sizes are often bigger than most commonplace GPU DRAM sizes. This allows, as long as the CPU RAM is large enough, to reconstruct large images on arbitrarily small GPUs with a minuscule extra computational footprint added. This is achieved by optimized overlapping of memory copies and computational kernels in the GPUs, such that memory is only copied in or out of the GPU while computing is performed, and the projectors almost never need to wait for a memory transfer to finish. More information can be found in [58]. Moreover, in the cases where the CPU RAM is still too small for the problems in hand, TIGRE also allows the use of swap memory. However, this will be slower for two main reasons: the novel memory management system cannot be used in this case, and swap memory is slower than RAM. Lastly, the TIGRE toolbox includes a GPU-management system that allows the user to select which GPUs they want to use, either by name or GPU id. This is an essential tool for users computing in heterogeneous platforms or shared resources. Code Snippet 2.5 showcases how this can be selected, and passed into an algorithm.

Code Snippet 2: Optional GPU selection in Python

```
1 from tigre.utilities import gpu
2 gpuids = gpu.getGpuIds('1 GeForce RTX 2080 Ti') # Will only select these GPUs
3
4 gpuids = gpu.GpuIds() # Get all GPUs
5 gpuids.devices = int32(2:3) # Only use IDs 2 and 3
6
7 import tigre. Algorithms as algs
8 # Define projections, geo, angles, n_iter
9 res_cgls = algs.cgls(projections, geo, angles, n_iter, gpuids=gpuids)
```

It is worth noting that having this multi-GPU features is crucial for the μ -CT community to gain access to the iterative reconstruction field, allowing the testing of large acquisitions with modern reconstruction methods [59–61].

2.6. Data loaders

A key and often non-trivial feature of a user-driven toolbox is the capability of loading datasets directly from the output of the devices that acquired the measurements. This is particularly true for CT, because a standard file format for sinograms does not exist. On the contrary, each scanner manufacturer has a unique way of storing the measurements and scanning parameters, often changing between different versions of software of the same manufacturer. Correct data and system geometry loading is often key. In TIGRE, data loaders for 6 different major CT manufacturers are available, 2 from the medical community (Philips, Varian)¹⁸ and 4 from the μ CT community (Comet Yxlon, Nikon, Bruker, Diondo). Additionally, the Data Exchange (DXChange) format [62], proposed for synchrotron tomography sharing, is also supported. Among these, the Varian loader contains a much higher set of functionalities for the preprocessing of the data, and TIGRE emulates the Varian CBCT machines pipeline¹⁹. This is explained in more detail in the article by Yi *et al* [63].

The data loaders are included in TIGRE such that loading any complex dataset can be called in a single line of code containing image loading, geometry reading and any other preprocessing operations needed, such as applying the Beer-law to the projection. The following snippet highlights how it is used:

Code Snippet 3: Data loading script in Python

```
1 import tigre.utilities.io as tigreo
2 projections, geometry, angles = tigreo.NikonDataLoader('/path/to/dataset')
```

Few other features are available, such as the ability to load partial datasets, or collections of datasets, when available.

2.7. Proton computed tomography

Since the early 2000s, proton therapy centers have been playing an increasingly important role in cancer therapy. This is because the unique properties of these particles allow treating deep-seated tumors with high precision (sharp dose-deposition). In order to fully exploit the potential of proton therapy, the planning CT could be measured with protons as well. This would allow bypassing the inaccuracies in treatment planning that arise from the conversion of Hounsfield units (obtained from a standard planning CT) to relative stopping powers of protons, and could effectively improve therapy [64].

The properties of an imaging setup for proton computed tomography (pCT), which was first described in [65], differ from those of standard CT, and hence both the modelling and the subsequent reconstruction algorithms need to be modified accordingly. The main difference is that proton paths through matter cannot be accurately approximated by a straight line, due to multiple Coulomb scattering, and therefore information obtained from particle tracking has to be included in the reconstruction process. TIGRE includes a novel preprocessing step to do this [66], which allows generating optimized (rebinned) radiographs from the pCT measurements. This method is based in [67], where a maximum-likelihood approach combined with a cubic spline estimate for the proton path was used to obtain proton radiographs with high resolution. For improved

¹⁸ Philips and Varian data loaders are only available in MATLAB at the time of writing.

¹⁹ The authors do not have proprietary information on Varian software, the functions here are reverse engineered from domain knowledge, not private information.

pixel weighting in the radiographs, a cylindrical object hull can be used as input in the newly implemented binning step in TIGRE. While the binning itself is implemented in CUDA, the user can call the method, after defining the geometry, by a one-line MATLAB command:

Code Snippet 4: Binning of pCT data into optimized radiographs

```
1 geo = % Define geometry as usual
2 % New geometry parameters
3 geo.DSID = 300' % Distance between source and upstream detector.
4 geo.DSD = 700' % Distance between source and downstream detector.
5 geo.hull = [1 5 0' 15 0; 0' 4 0]' % Convex hull
6
7 pRad = pCTCubicSpline_mex ( data . posIn, data . posOut, data . dirIn, ...
8 data . dirOut, data . Wepl, eIn, geo )
```

In the code snippet, *pRad* refers to one optimized proton radiograph: for a full pCT, this step has to be repeated for each radiograph. The data input is a list of the proton positions and directions upstream and downstream the patient, the water-equivalent path length measured for each proton, the incident proton energy and the pCT imaging geometry (which now also contains the locations of the particle tracking detectors and optionally an object hull). A demo file has been added to TIGRE to facilitate easier usage. After all proton radiographs have been preprocessed, any implemented reconstruction algorithm in TIGRE can be used seamlessly without any further adaption for pCT. This code is only available in MATLAB at the time of writing.

2.8. PyTorch operator wrappers

State-of-the-art CT reconstruction includes data-driven methods, many of which use CT operators (forward and backprojection) within the model, e.g. Learned Primal Dual [68]. As for the time of writing, Torch, and in particular PyTorch [9], is the main machine learning library used in academia and industry. TIGRE now includes an optional PyTorch binding that allows TIGRE projectors to be included inside PyTorch models, and be treated as a linear differentiable operator by the automatic differentiation engine. Code Snippet 2.8 shows how TIGRE auto-differentiable operators can be created to use within PyTorch.

1 Code Snippet 5: Use of TIGRE's PyTorch bindings

```
2
3 geo = tigre.geometry(mode="fan")
4 angles = np.linspace(0, np.pi, 200)
5 ax, atb = create_PyTorch_operator(geo, angles)
6 # Now ax and atb can be used inside autograd-enabled pipelines
7 import torch
8 device = torch.device("cuda:1")
9 input_volume = torch.randn([2, geo.nVoxel[0], geo.nVoxel[1], geo.nVoxel[2]],
10 requires_grad=True).to(device)
11 sino_batch = ax(input_volume) # returns batched torch tensor
12 image_batch = atb(sino_batch)
```

While this allows for TIGRE to be used within PyTorch, it is important to mention that the embedding of the operators is not as sophisticated as it is with other tools, e.g. tomosipo [10]. Particularly, TIGRE's approach requires the memory to be transferred from a PyTorch tensor to a NumPy array, which more GPU ↔ CPU memory transfers are needed than in tomosipo.

3. Results

In order to showcase the usability of TIGRE in a variety of real datasets, we present several experiments alongside the code needed to reproduce them. The first example concerns a clinical scenario with patient data. The second example consists of a synchrotron parallel beam tomographic reconstruction. The third example corresponds to a Monte Carlo simulation of a proton CT problem (due to its early stages of development as a device). Next, a neutron CT scan is showcased. Finally, the last example involves an industrial μ CT dataset.

It is important to note that this document is not making scientific statements about which reconstruction techniques are best in any of the experiments presented here. This is because, often, the performance of CT

imaging reconstructions is heavily dependent on secondary tasks, e.g. the best image for segmentation is not necessarily the best image for diagnosis. Similarly, human eyes are biased towards natural looking images, rather than the maps of scalar fields outputted in CT reconstructions. Moreover, experts from different fields have sometimes trained their perception to filter out noise, e.g. a radiologist knows how to identify the differences between an artifact and a lesion in medical CT. Therefore, the only claim we make in this work is that reconstructions obtained using different algorithms are different, and that they may be of interest in different scenarios. For this reason, there is a real value on making more reconstruction algorithms easily available for both clinician and algorithm developers, as well as platforms to promote knowledge transfer between them.

It is also worth noting that all the experiments presented in this section have been run in a personal high-end desktop (Ryzen 5, 32 GB RAM, RTX 4070), and not in a large workstation, highlighting the power of TIGRE to produce computationally fast results in day-to-day use machines. Computational times for each of the methods are reported when introduced.

3.1. Medical dataset from a Varian CBCT machine

With the data loading and pre-processing tools added to TIGRE for .xim files, particularly for the Varian range of onboard CBCT machines, iterative algorithms can easily be utilized in clinical images. This experiment shows how to read and reconstruct such images, in particular from a Varian Edge CBCT situated at Beijing Cancer Hospital.

The first dataset is a head scan with projections of size 1024×768 taken on a limited arc range of 200° , uniformly sampled with 493 projections, where the size of the reconstruction is $512 \times 512 \times 362$ voxels. The projections are loaded and pre-processed with the data loader, and then reconstructed using three different algorithms: FDK (4.5 s), OS-SART with 50 iterations (118 s), and OS-ASD-POCS with 50 iterations (160 s). The latter includes TV regularization to reduce the noise in the reconstruction. The script that produces the reconstructions can be seen in Snippet 3.1 and two different slices of the head at different crano-caudal slice values can be seen in figure 2. Each reconstruction took less than 5 minutes (few seconds for FDK) in a personal desktop setup. In the reconstruction, one can see that OS-SART and OS-ASD-POCS produce images with less pixel-level noise than FDK, particularly noticeable in OS-ASD-POCS, as it has TV regularization, which promotes flatter images.

Code Snippet 6: MATLAB code to reconstruct Varian CBCT data 1 % Load TIGRE

2 InitTIGRE'

3 % Load

4 [projections, geo, angles] = VarianDataLoader ('./path/to/Head/scan');

5 % Reconstruct

6 recon_fdk = FDK(projections, geo, angles);

7 recon_os_sart = OS_SART(projections, geo, angles, 50);

8 recon_tv = OS_ASD_POCS(projections, geo, angles, 50);

3.2. Synchrotron dataset

Synchrotrons are large electron accelerators that have become an essential experimental facility for scientific discovery because they produce high quality x-rays of monochromatic energy and low noise statistics. This can be used to obtain extremely high quality CT images, and it also allows for also more complex experiments, such as x-ray crystallography, x-ray diffraction imaging and several different modalities. For this reason, they are routinely used for material science, biology and many other applications.

While synchrotron experiments often produce very high quality data, and thus iterative reconstruction is often unnecessary, this is not always the case. In this experiment, taken from the TomoBank data repository [69], a single noisy time-frame from a dynamic acquisition is reconstructed ($tomo_{ID} = 00031$). Specific information about the experiment can be found in [70, 71]: in short, often scientist want to perform in-situ experiments and image the evolution of certain processes. In these cases, the acquisition time has to be severely reduced with respect to the optimal time one would need to capture the dynamics of the experiment. Thus, imaging data is very low dose and therefore contains very high levels of noise, which produces severe artifacts in the commonly used FBP reconstruction, as seen in the top-left of figure 3. In this experiment, the projections are of size 960×900 and sampled over 900 angles, uniformly distributed around 180° range. The reconstructed images have $960 \times 960 \times 900$ voxels.

In this experiment, several algorithms are used to reconstruct the same data to showcase the different nature of the resulting images. In Snippet 3.2 we can see the code that produces the images in figure 3, where six different reconstruction can be seen. The leftmost column shows FBP with the standard ramp filter on top and FBP with a noise-rejecting Hamming filter on the bottom. The second column shows two iterative reconstruction algorithms with no explicit regularization function, LSQR on the top and SIRT on the bottom. The last column shows TV-regularized algorithms, namely FISTA and OS-ASD-POCS, with their higher noise-rejection behavior.




Figure 2. Reconstruction from a Varian Edge machine, 200° arc head scan at Beijing Cancer Hospital. Rows show different slices of the head, columns show the FDK, OS-SART (50 iterations) and OS-ASD-POCS (50 iterations) reconstructions, from left to right. Visualization is linear attenuation coefficients in range [0-0.05].

Code Snippet 7: MATLAB code to reconstruct synchrotron data 1% Load TIGRE

```
2 InitTIGRE ;
3 % Load
4 [projections, geo, angles] = DXChangeDataLoader("./data_file.h5");
5 % Reconstruct
6 recon_fbp = FBP(data, geo, angles);
7 recon_fbp_hamming = FBP(data, geo, angles,'filter','hamming');
8 recon_lsqr = LSQR(data,geo,angles,20);
9 recon_sirt = SIRT(data,geo,angles,150);
10 recon_fista = FISTA(data,geo,angles,50,'tviter',50);
11 recon_os_asd_pocs = OS_AS_D_POCS(data,geo,angles,20,'blocksize',80);
```

3.3. Proton computed tomography

Using the code demonstrated in Code Snippet 2.7, optimized proton radiographs were calculated in [66], using as input parameters the measured water-equivalent thickness, the initial proton energy, the upstream and downstream positions and directions for each proton and the pCT geometry. The pCT geometry closely follows the standard geometry definition in TIGRE, however, the locations of upstream and downstream detectors (as opposed to only one detector in conventional CT), as well as the convex hull (here implemented as a cylinder) have to be declared.

For this example, Catphan phantoms modules [72] (15 cm diameter) were simulated with the Monte Carlo toolkit Geant4 [73] and irradiated with 225 protons/mm². The synthetic measurements were then combined into optimized binned radiographs in TIGRE with the new pCT implementation, to create 90 radiographs (4 degree steps over a range of 360 degree) of the high resolution phantom (PMMA body with aluminium line pair insets). The experiment data was obtained from the the 'non-ideal' parallel simulation case (realistic energy resolution of $\Delta E/E = 1\%$, detector resolutions of $\sigma_s = 0.15$ mm and detector thicknesses of 300 μ m), that was described in full detail in [66]. These optimized radiographs can be used as CT-like data in TIGRE and hence inputted to the implemented algorithms without further adaptation.

Code Snippet 8: Reconstruction of optimized proton radiographs.

```
1 % Load optimized proton radiographs
2 recon_fdk = FDK(projections,geo,angles);
3 recon_tv = ASD_POCS(projections,geo,angles,20);
4 recon_osem = OSEM(projections,geo,angles,50);
```





Figure 3. Synchrotron dataset acquired with settings for a dynamic in-situ experiment. Reconstructions are, from left to right and top to bottom: FBP with linear filter (17 s), LSQR with 20 iterations (249 s), FISTA with 80 iterations (1330 s), FBP with Hamming filter (17 s), SIRT with 150 iterations (1330 s) and OS-ASD-POCS with 20 iterations (856 s).


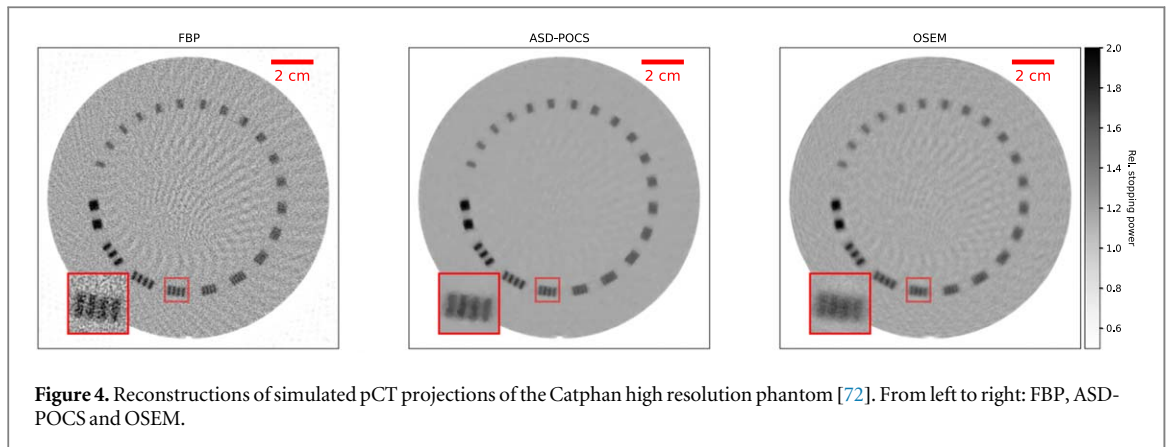
The resulting reconstructions are displayed in figure 4, for FBP (0.4 s), OSEM at 50 iterations (58 s) and ASD-POCS at 20 iterations (110 s). While the filtered-backprojection results in the noisiest image, OSEM results in a smoother but blurrier image. In all cases, despite the very low number of projections used, the algorithms still allow viewing details of the aluminium line pair insets of the phantom. The least amount of noise in the reconstruction was obtained with the ASD-POCS algorithm.

3.4. Neutron tomography dataset

Neutron imaging is a highly accurate non-invasive method which can be used to image objects containing light elements encased inside a heavier material [74]. This is because the properties of neutrons are complementary to those of x-rays, i.e. the transmissivity of neutrons with metallic elements is higher while their transmissivity with light elements is lower. However, this imaging technique sometimes leads to very slow measuring times, often making it impractical to obtain full angle measurements. For more information, we refer the reader to Appendix C.

The neutron image dataset in this experiment was obtained at the Thai Research Reactor-modification 1 (TRR-1/M1) located at the Thailand Institute of Nuclear Technology (TINT). In this facility, it generally takes at least 14 hours to acquire the 501 projection images, covering 180 degrees in intervals of 0.36 degrees required to obtain a crisp reconstruction using FBP. This long scanning time is currently limiting their potential applications, so they are exploring the use of iterative algorithms to produce reconstructions of similar quality with a limited number of projections.

For the presented experiment, the thermal-neutron flux at the imaging position was $2.5 \text{ \AA} \text{---} 10^5 \text{ cm}^{-2} \text{ s}^{-1}$. The main instruments of the neutron imaging system consist of a neutron-to-photon conversion plate made from 6LiF/ZnS, a 45-degree mirror, a sample rotation system operated by in-house developed software, a Nikon 50-mm/f1.2 lens, and a 2048×2048 -pixel CCD camera. The Field-of-View (FOV) is $20 \times 20 \text{ cm}^2$ and the L/D ratio is 78. In this experiment, four reconstruction algorithms (FDK (0.4 s), SART with 100 iterations (440 s), SART-TV with 100 iterations (794 s), and SIRT with 100 iterations (66 s)) were used to reconstruct the same neutron tomography dataset with an extremely limited set of neutron projections. The sizes of the projection data are 887×887 pixels, taken on an arc range of 180° in 18° increments, leading to a total of 10 projection images per set (as opposed to the standard 0.9° increments, leading to a total of 201 projection images per set). The size of the reconstructions is $512 \times 512 \times 720$ voxels. The filter Shepp-logan was used for the FDK reconstruction. One hundred iterations were executed for the SART, SART-TV, and SIRT algorithms. The codes that produce the reconstructed images in figure 5 can be seen in Snippet 3.4.



Code Snippet 9: MATLAB code to reconstruct neutron tomography data

```
1 % Load TIGRE
2 InitTIGRE ;
3 % Load data
4 [proj,geo,angles] = TINTDataLoader(data_folder) % This function is not in TIGRE
5 % Reconstruction
6 recon_fdk_shepp_logan = FBP(proj, geo, angles,'filter','shepp-logan');
7 recon_sart = SART(proj,geo,angles,100); %SART for 100 iterations
8 recon_sirt = SIRT(proj,geo,angles,100); %SIRT for 100 iterations
9 recon_sart_tv = SART_TV(proj,geo,angles,100); %SART_TV for 100 iterations
```

A representation of the reconstructed images is shown in figure 5: the left-most column shows the SIRT reconstruction; the second column shows the SART_TV algorithm, which is based on SART with additional TV regularization; the third column shows the standard SART reconstruction; and the right-most column shows the FDK reconstruction with a Shepp-logan filter. Since this experiment considered an extremely limited set of data, the resulting images lack definition around the edges of the padlock components in all reconstructions. However, it can be obviously seen that the results obtained using the FDK algorithm show more severe artifacts. If we compare the results obtained using the three iterative algorithms, we can observe in the top part of the padlock on column three of figure 5 that the reconstruction given by SART preserves edges better than the other algorithms.

3.5. Diondo μ CT dataset

When radiation dose is not an issue, micrometre-range pixel resolution μ CT scanners are often used, particularly increasingly being used in industry for non destructive testing. The high resolution, however, is linked to long scanning times; and while this is acceptable for some applications, it is a hindrance if the scanning is used for continuous inspection of samples. In μ CT, when there is no limit to the scanning time and if we have access to the object from all directions and if the object is approximately cylindrical with no




Figure 6. Scan of an Energizer AAAA battery acquired in a diondo d5 scanner at the μ Vis lab at the University of Southampton using a helical acquisition mode. Top shows FDK and bottom FISTA with TV regularization (20 iterations).

significant variations in x-ray path length through the object, the FDK algorithm produces very high quality reconstructions: with enough radiation exposure, the signal to noise ratio of the data is sufficiently high to minimally impact the reconstruction. This is not true when faster scans are needed, as either the projection data has to be acquired faster, thus with more noise, or less projection data are acquired. This example falls into the latter category where, to speed-up measuring times, an Energizer AAAA battery is measured in only a subset of 258 out of the 1258 projection angles prescribed in a conventional scanning regime. The scan was acquired at the University of Southampton μ Vis lab, which is part of the National x-ray Computed Tomography (NXCT) facilities of the UK, on a diondo d5 using a 300 kVp x-ray transmission source, with a detector size of 2000×2000 (reduced by selecting the 800×2000 pixel central region), in helical trajectory mode.

A representation of the reconstructions obtained using FDK (21 s) and the FISTA algorithm with TV regularization run over 20 iterations (425 s) can be seen in figure 6, following the code in Snippet 3.5. The images are $600 \times 600 \times 2000$ voxels. The reconstruction obtained using FISTA seems to be much less noisy, albeit a bit blurry in some areas. A better choice of hyperparameters may be able to improve the reconstruction, however it is also important to note that FDK is sometimes sufficiently performing, and this may be one of such cases. It is likely that FISTA would be a better algorithm to use for segmentation, but if visual inspection is only needed, FDK does a perfectly fine job. This is ultimately true for any reconstruction algorithm and application: image quality is dependent on what the image is needed for, it is not itself a target.

Code Snippet 10: MATLAB code to reconstruct Diondo Helical data

```
1 % Load TIGRE
2 InitTIGRE ;
3 % Load
4 [projections, geo, angles] = DiondoDataLoader('./path/to/battery/scan');
5 % Reconstruct
6 recon_fdk = FDK(projections,geo,angles);
7 recon_fista = FISTA(proj,geo,angles,20,'hyper',2e2,'init','FDK','tviter',50);
```

4. Discussion and conclusions

In this work, we present nearly a decade's worth of scientific advancements in the field of CT reconstruction, encapsulated in the open-source code contributions to the TIGRE toolbox. These contributions correspond to a twenty-fold increase in the code-base of the toolbox, and have allowed an effective and practical transfer of

knowledge from the mathematical fields of optimization and inverse problems into the experimentalists' world, as the various examples in this article, and hundreds of users of TIGRE show.

One of the main goals of the TIGRE toolbox is to provide tested and robust implementations of iterative algorithms that can be used on real data acquisitions in lab-environments. Moreover, whenever possible, the released implementations have been standardized (in terms of parameters choices, etc) to ease comparison and applicability. Currently, it already contains 23 standardized algorithms, and this number keeps increasing.

As highlighted in this article, stand-alone reconstruction algorithms are not sufficient for reconstructing real data. On the one hand, effective CT reconstruction requires operators (forward and backward) that are not only geometrically adaptable but also computationally efficient, and that can be incorporated easily in any structured algorithm implementation. Since its first release, TIGRE has kept on improving the computational speed and flexibility of its operators, facilitating the seamless reconstruction of experimental data. On the other hand, handling data from various scanners is a complex task involving thousands of lines of code, which is crucial both to simplify the testing of different algorithms by practitioners, as it is to successfully evaluate the performance of a new algorithm for different problems. To address this, TIGRE now includes data-loading support for seven different manufacturers, all of them verified on real scanners.

While great efforts to simplify the use of iterative algorithms for non-experts has been done, the authors acknowledge that it is sometimes complex to know which algorithm to use under specific scanning conditions or how to fine-tune them to the appropriate parameters. Suggestions and guidelines exist in the TIGRE repository, but future work will rely on improving automatic parameter selection and attempting to have a more systematic way to help users select algorithms.

On the topic of algorithm selection, it is important to mention that defining what makes an algorithm good is not possible. It is not uncommon to promote algorithm performance assuming that the final image quality compared to a reference is the most important feature (even if the most commonly used metrics fail at this task [75, 76]), however image quality itself is a relative metric. For example, an image for diagnosis in medical CT or an image for radiation treatment would focus on different things, detectability in the first case, but quantitative accuracy in the second. The best image reconstruction for each of these tasks is not the same. Similarly, a μ CT image for metrology or material characterization would be evaluated very differently on what constitutes quality. This is why it is so relevant that many iterative algorithms exist, and why it is important that these can be readily used.

In this work, we also want to highlight the broad range of applicability of iterative algorithms, and of the toolbox itself, with respect to different imaging types. We highlighted examples of all the main 3D CT modalities, namely synchrotron CT, medical CBCT and μ CT, and also neutron CT and proton-CT. The (forward) mathematical model behind the first four is based on the same operator, the Radon transform, but the way the data is handled and loaded is very different, as well as the nuisances of each application. Moreover, proton-CT, has also been adapted to fit seamlessly into the TIGRE framework. In this way, the experiments in this paper aim to showcase how TIGRE can be easily used in very different kinds of tomographic problems.

TIGRE has grown significantly since its inception, mostly driven by user feedback and needs, and it is the intention of the authors for this momentum to continue. Its original publication shows a snapshot of what the software was when it was released, and here we offer another snapshot almost a decade from then. Now, the purpose of this work is to provide a structured overview of the current version of the toolbox, providing appropriate descriptions and references; and to serve as a comprehensive and peer-reviewed guide for the users. It is, however, not the intention of this paper to provide a complete version of the toolbox, which we hope may become incomplete very soon with the addition of new capabilities. We thus urge mathematicians, CT experimentalist and anyone in-between to join into the TIGRE software community and contribute, from code, to issues, to feature requests at github.com/CERN/TIGRE/.

Acknowledgments

A. Biguri wants to thank all named and anonymous contributors of TIGRE toolbox over the past years, for keeping TIGRE alive and updated with the latest technological trends. The software is alive thanks to all these contributions. A. Biguri would also like to acknowledge CT users that have come with interesting problems to solve and many bugs to fix, as without the requests of these users, TIGRE would not be able to accommodate so many real use-cases. A. Biguri will finally also like to acknowledge CERN, EPSRC EP/T003553/1, EP/N014588/1 and the Accelerate Programme for Scientific Discovery for funding that allowed to develop and maintain such tool. The authors acknowledge the μ -VIS x-ray Imaging Centre (muvis.org), part of the National Facility for laboratory-based x-ray CT (nxct.ac.uk—EPSRC: EP/T02593X/1), at the University of Southampton for provision of tomographic imaging facilities. C.B.S additionally 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, 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). Please note that the content of this publication reflects the authors' views and that neither NIHR, the Department of Health and Social Care, IMI, the European Union, EFPIA, nor the DRAGON consortium are responsible for any use that may be made of the information contained therein. This work was partially supported by the National Science Foundation program under grant DMS-2208294. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

Data availability statement

The data cannot be made publicly available upon publication because the cost of preparing, depositing and hosting the data would be prohibitive within the terms of this research project. The data that support the findings of this study are available upon reasonable request from the authors.

Author contribution

A. Biguri is the main maintainer and developer of TIGRE and has been involved in every part of the work presented here. T. Sadakane and Y. Liu are active contributors to the tool and have significantly added to the general TIGRE development and issue fixing, particularly in the Python side. R. Lindroos developed the first iteration of the Python-TIGRE library. M. Sabaté Landman's contributions are related to the implementation of Krylov subspace methods, algorithm features and paper structure. M. Lohvithee developed total variation regularization algorithms and the experiments with neutron CT. S. Kaser created the proton CT methods. S. Hatamikia developed arbitrary rotation code and invaluable algorithmic testing expertise. Y Du developed, tested and validated the TIGRE module (TIGRE-Varian) for Varian linac onboard CBCT in clinical settings [63]. R. Bryll significantly improved the speed of the GPU-driven operators. E. Valat implemented the PyTorch bindings. S. Wonglee set the neutron tomography experiments and acquisition. C.B. Schönlieb and T. Blumensath provided knowledge and support on reconstruction methods.

Appendix A: Table of algorithms

Table A1. Here we provide a list of all the algorithms provided in the TIGRE toolbox (as of November 2024), along with the appropriate citations to the original papers, and the references that are specific to the TIGRE toolbox. Note that [7] refers to the first publication of the toolbox.

Algorithm name	Original description	TIGRE description
Direct methods		
FBP and FDK	[25], [26]	[7]
Kaczmarz-type algorithms		
Simultaneous ART (SART)	[28]	[7]
Ordered Subsets ART (OS-SART)	[29]	section 2.2.2, [7]
Simultaneous Iterative Reconstruction algorithm (SIRT)	[1]	[7]
Projection Onto Convex Sets (POCS)		
Adaptative Steepest Descent version (ASD-POCS)	[30]	[7], [33]
Projection-Controlled Steepest Descent (PCSD)	[31]	[33]
Adaptive weighed TV norm OS version with both TV norms	[32, 33]	section 2.2.3, [33] section 2.2.2
Bregman outer iterations + ASD-POCS (B-ASD-POCS- β)	[34]	[7]
Krylov Subspace algorithms		
Conjugate Gradient Least Squares (CGLS)	[36]	[7], [46]
LSQR	[37]	[46]
hybrid LSQR	[37]	[46]
AB/BA-GMRES	[38]	[46]
LSMR	[39]	[46]
IRN-TV-CGLS	[40]	[46]
hybrid-fLSQR-TV	[41]	[46]
Statistical minimization algorithms		
OSEM	[42]	section 2.1
Proximal algorithms		
Fast Iterative Shrinkage-Thresholding algorithm (FISTA)	[43]	section 2.1
Fast variant of FISTA	[44]	section 2.1
SART-TV	[45]	[7]

Appendix B: Geometric parameters in TIGRE

Table B1. Geometry parameters of TIGRE and description.

Geometric Parameter	Description	Reference
DSD	Distance between the source (S) and the detector (D)	—
DSO	Distance between the source (S) and origin (O)	—
sVoxel	Size of the volume (\mathbb{I}) in real world units	—
nVoxel	Number of voxels of the discretization of the \mathbb{I}	—
sVoxel	Size of each voxel in the discretization of \mathbb{I}	—
offOrigin (\vec{V}_{orig})	Shift of location of desired center of the volume w.r.t. the default	Oxyz
COR	Center of Rotation shift. Moves the rotation axis	Oy
angles (z_1, y_1, z_2)	Euler ZYZ rotation angles	Oxyz
sDetector	Size of the detector (\mathbb{D}) in real world units	—
nDetector	Number of pixel of the discretization of the \mathbb{D}	—
sDetector	Size of each pixel in the discretization of \mathbb{D}	—
offDetector (\vec{V}_{det})	Shift of location of desired center of the detector w.r.t. the default	uv
rotDetector (ϕ, θ, ψ)	In place rotations of the detector	D

Appendix C: Neutron tomography

Neutron imaging is a highly accurate and reliable non-invasive measurement method that has been utilized for various applications such as cultural heritage, thermal-hydraulics, nuclear engineering studies, and botanical sciences [74]. The properties of neutrons are complementary to those of x-rays and gamma-rays i.e. the transmissivity of neutrons with metallic elements is higher while their transmissivity with light elements is lower. Therefore, neutron imaging is excellent for inspecting specimens containing light elements inside the object to be scanned, especially if they are covered or enveloped by heavy elements on the outside.

Despite the advantages of neutron tomography, this modality has not yet been widely employed for routine inspection in industrial applications. One of the main reasons for this is the requirement of having a high neutron flux to produce high quality neutron tomography imaging, which has to be generated in a nuclear reactor or in a spallation source and is therefore not always easily accessible. Moreover, some common challenges can affect the power at which the reactors operate, as well as some potential benefits in terms of safety and money saving, causing a low neutron flux to arrive at the neutron imaging facility and resulting in very slow measuring times. This is important as neutrons also lead to the objects being ‘activated’ i.e. they become radioactive. In practice, the scanning time required for a full angle measurement can easily become impractical and compromise the imaging of many dynamic processes. Iterative algorithms have been shown to produce reconstructions with superior quality using a limited number of projections with respect to the commonly used FBP algorithm. Therefore, it is a promising avenue to explore the use of these methods in this context, with the aim of reducing the number of required projections and shorten acquisition times at neutron imaging facilities.

An example of neutron imaging facility is the Thai Research Reactor-modification 1 (TRR-1/M1), located at the Thailand Institute of Nuclear Technology (TINT). The TRR-1/M1 research reactor is an open pool type of TRIGA-Mark III, whose main purpose is public research and non-destructive investigation using radiography and tomography techniques. This reactor has a limited resource of fuel rods, which employ old, discontinued technology. Hence, the reactor has to operate at 1 MW to save fuel. In this facility, the acquisition time required to obtain a good-quality projection for one particular angle takes approximately 100 seconds. This long scanning time is currently limiting their potential applications.

ORCID iDs

Ander Biguri  <https://orcid.org/0000-0002-2636-3032>
Tomoyuki Sadakane  <https://orcid.org/0009-0002-4705-4164>
Manasavee Lohvithee  <https://orcid.org/0009-0004-4105-511X>
Stefanie Kaser  <https://orcid.org/0000-0001-7356-1835>
Robert Bryll  <https://orcid.org/0000-0002-2139-8279>
Thomas Blumensath  <https://orcid.org/0000-0002-7489-265X>
Carola-Bibiane Schönlieb  <https://orcid.org/0000-0003-0099-6306>

References

- [1] 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
- [2] Hansen P C, Jørgensen J and Lionheart W R B 2021 *Computed Tomography: Algorithms, Insight, and Just Enough Theory* (Society for Industrial and Applied Mathematics) PA
- [3] Hansen P C 2010 *Discrete Inverse Problems* (SIAM)
- [4] Mueller J L and Siltanen S 2012 *Linear and Nonlinear Inverse Problems with Practical Applications* (Society for Industrial and Applied Mathematics) PA
- [5] Padole A, Ali Khawaja R D, Kalra M K and Singh S 2015 CT radiation dose and iterative reconstruction techniques *American Journal of Roentgenology* **204** W384–92
- [6] Stiller W 2018 Basics of iterative reconstruction methods in computed tomography: a vendor-independent overview *Eur. J. Radiol.* **109** 147–54
- [7] Biguri A, Dosanjh M, Hancock S and Soleimani M 2016 TIGRE: a MATLAB-GPU toolbox for CBCT image reconstruction *Biomed. Phys. Eng. Express* **2** 055010
- [8] Van Aarle W, Palenstijn W J, De Beenhouwer J, Altantzis T, Bals S, Batenburg K J and Sijbers J 2015 The ASTRA toolbox: a platform for advanced algorithm development in electron tomography *Ultramicroscopy* **157** 35–47
- [9] Paszke A et al 2019 PyTorch: an imperative style, high-performance deep learning library ed H Wallach et al *Advances in Neural Information Processing Systems* vol 32 (Curran Associates, Inc)
- [10] Hendriksen A A, Schut D, Palenstijn W J, Viganó N, Kim J, Pelt D M, Van Leeuwen T and Batenburg K J 2021 Tomosipo: fast, flexible, and convenient 3d tomography for complex scanning geometries in python *Opt. Express* **29** 40494–513
- [11] Jørgensen J S et al 2021 Core imaging library-part i: a versatile python framework for tomographic imaging *Phil. Trans. R. Soc. A* **379** 20200192
- [12] Rit S, Vila Oliva M, Brousmiche S, Labarbe R, Sarrut D and Sharp G C 2014 The reconstruction toolkit (RTK), an open-source cone-beam CT reconstruction toolkit based on the insight toolkit (ITK) *In Journal of Physics: Conference Series* vol 489 (IOP Publishing) 012079
- [13] Gürsoy D, De Carlo F, Xiao X and Jacobsen C 2014 TomoPy: a framework for the analysis of synchrotron tomographic data *J. Synchrotron Radiat.* **21** 1188–93
- [14] Kim H and Champley K 2023 Differentiable forward projector for x-ray computed tomography arXiv:2307.05801
- [15] Adler A and Lionheart W R B 2006 Uses and abuses of EIDORS: an extensible software base for EIT *Physiol. Meas.* **27** S25
- [16] Liu B, Yang B, Xu C, Xia J, Dai M, Ji Z, You F, Dong X, Shi X and Fu F 2018 pyEIT: a python based framework for electrical impedance tomography *Software* **7** 304–8
- [17] Thielemans K, Tsoumpas C, Mustafovic S, Beisel T, Aguiar P, Dikaio N and Jacobson M W 2012 Stir: software for tomographic image reconstruction release 2 *Phys. Med. Biol.* **57** 867
- [18] Ovtchinnikov E et al 2020 SIRF: synergistic image reconstruction framework *Comput. Phys. Commun.* **249** 107087
- [19] Merlin T, Stute S, Benoit D, Bert J, Carlier T, Comtat C, Filipovic M, Lamare F and Lamare D 2018 Visvikis. CASToR: a generic data organization and processing code framework for multi-modal and multi-dimensional tomographic reconstruction *Phys. Med. Biol.* **63** 185005
- [20] Adler J, Kohr H and Ozan Ö 2017 ktem. Operator discretization library (ODL) (<https://odlgroup.github.io/odl/>)
- [21] Hansen P C and Jørgensen J S 2018 AIR Tools II: algebraic iterative reconstruction methods, improved implementation *Numer. Algorithms* **79** 107–37
- [22] Gazzola S, Hansen P C and Nagy J G 2019 IRtools: a MATLAB package of iterative regularization methods and large-scale test problems *Numer. Algorithms* **81** 773–811
- [23] Calvetti D and Somersalo E 2007 Introduction to Bayesian Scientific Computing *Ten Lectures on Subjective Computing* **2** 01
- [24] Feldkamp L A, Davis L C and Kress J W 1984 Practical cone-beam algorithm *J. Opt. Soc. Am. A* **1** 612–9
- [25] Wang G 2002 X-ray micro-CT with a displaced detector array *Med. Phys.* **29** 1634–6
- [26] Wesarg S, Ebert M and Bortfeld T 2002 Parker weights revisited *Med. Phys.* **29** 372–8
- [27] Gordon R, Bender R and Herman G T 1970 Algebraic reconstruction techniques (ART) for three-dimensional electron microscopy and x-ray photography *J. Theor. Biol.* **29** 471–81
- [28] Andersen A H and Kak A C 1984 Simultaneous algebraic reconstruction technique (SART): a superior implementation of the ART algorithm *Ultrason. Imaging* **6** 81–94
- [29] Censor Y and Elfving T 2002 Block-iterative algorithms with diagonally scaled oblique projections for the linear feasibility problem *SIAM J. Matrix Anal. Appl.* **24** 40–58
- [30] 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
- [31] Liu L, Lin W and Jin M 2015 Reconstruction of sparse-view X-ray computed tomography using adaptive iterative algorithms *Comput. Biol. Med.* **56** 97–106
- [32] Liu Y, Ma J, Fan Y and Liang Z 2012 Adaptive-weighted total variation minimization for sparse data toward low-dose X-ray computed tomography image reconstruction *Phys. Med. Biol.* **57** 7923
- [33] 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.* **42** 9295

- [34] Xue H, Zhang L, Cheng Z, Xing Y and Xiao Y 2010 An improved tv minimization algorithm for incomplete data problem in computer tomography *In Nuclear Science Symposium Conference Record (NSS/MIC)* (IEEE) 2621–4
- [35] Chung J and Gazzola S 2024 Computational methods for large-scale inverse problems: a survey on hybrid projection methods *SIAM Rev.* **66** 205–84
- [36] Hestenes M R *et al* 1952 Methods of conjugate gradients for solving linear systems *J. Res. Natl Bur. Stand.* **49** 409–36
- [37] Paige C C and Saunders M A 1982 Algorithm 583: LSQR: sparse linear equations and least squares problems *ACM Transactions on Mathematical Software (TOMS)* **8** 195–209
- [38] Hansen P C, Hayami K and Morikuni K 2022 GMRES methods for tomographic reconstruction with an unmatched back projector *J. Comput. Appl. Math.* **413** 114352
- [39] Fong D C-L and Saunders M 2011 LSMR: an iterative algorithm for sparse least-squares problems *SIAM J. Sci. Comput.* **33** 2950–71
- [40] Wohlberg B and Rodriguez P 2007 An iteratively reweighted norm algorithm for minimization of total variation functionals *IEEE Signal Process. Lett.* **14** 948–51
- [41] Gazzola S and Sabaté Landman M 2019 Flexible GMRES for total variation regularization *BIT Numerical Mathematics* **59** 721–46
- [42] Hudson H M and Larkin R S 1994 Accelerated image reconstruction using ordered subsets of projection data *IEEE Trans. Med. Imaging* **13** 601–9
- [43] Beck A and Teboulle M 2009 A fast iterative shrinkage-thresholding algorithm for linear inverse problems *SIAM J. Imag. Sci.* **2** 183–202
- [44] Liang J and Schönlieb C-B 2018 *Faster FISTA* (<https://doi.org/10.48550/arXiv.1807.04005>)
- [45] Biguri A 2017 Iterative reconstruction and motion compensation in computed tomography on GPUs *PhD Thesis* University of Bath 11
- [46] Landman M S, 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** 15
- [47] Censor Y and Zenios S A 1997 *Parallel Optimization: Theory, Algorithms, and Applications* (Oxford University Press)
- [48] Wang G and Jiang M 2004 Ordered-subset simultaneous algebraic reconstruction techniques (OS-SART) *J. X-Ray Sci. Technol.* **12** 169–77
- [49] Censor Y and Elfving T 2002 Block-iterative algorithms with diagonally scaled oblique projections for the linear feasibility problem *SIAM J. Matrix Anal. Appl.* **24** 40–58
- [50] Morozov V A 1966 On the solution of functional equations by the method of regularization *Soviet Math. Doklady* 414417
- [51] Golub G H, Heath M and Wahba G 1979 Generalized cross-validation as a method for choosing a good ridge parameter *Technometrics* **21** 215–23
- [52] Sanctorem J G, Wassenbergh S V, Nguyen V, Beenhouwer J D, Sijbers J and Dirckx J J J 2021 Extended imaging volume in cone-beam x-ray tomography using the weighted simultaneous iterative reconstruction technique *Phys. Med. Biol.* **66** 165008
- [53] Bian J, Wang J, Han X, Sidky E Y, Shao L and Pan X 2012 Optimization-based image reconstruction from sparse-view data in offset-detector cbct *Phys. Med. Biol.* **58** 205
- [54] Vedantham S, Karellas A, Vijayaraghavan G R and Kopans D B 2015 Digital breast tomosynthesis: state of the art *Radiology* **277** 663–84
- [55] O'Brien N S, Boardman R P, Sinclair I and Blumensath T 2016 Recent advances in X-ray cone-beam computed laminography *J. X-Ray Sci. Technol.* **24** 691–707
- [56] Yang K, Kwan A L C, Miller D W F and Boone J M 2006 A geometric calibration method for cone beam CT systems *Med. Phys.* **33** 1695–706
- [57] Hatamikia S, Biguri A, Kronreif G, Figl M, Russ T, Kettenbach J, Buschmann M and Birkfellner W 2021 Toward on-the-fly trajectory optimization for C-arm CBCT under strong kinematic constraints *PLoS One* **16** e0245508
- [58] Biguri A *et al* 2020 Arbitrarily large tomography with iterative algorithms on multiple GPUs using the TIGRE toolbox *J. Parallel Distrib. Comput.* **146** 52–63
- [59] Sun W, Chretien S, Biguri A, Soleimani M, Blumensath T and Talbott J 2023 The realisation of fast X-ray computed tomography using a limited number of projection images for dimensional metrology *NDT & E International* **137** 102852
- [60] Tekseth K R and Breiby D W 2024 4D imaging of two-phase flow in porous media using laboratory-based micro-computed tomography *Water Resour. Res.* **60** e2023WR036514
- [61] Du Plessis A, Yadroitsav I, Yadroitsava I and Roux S G L 2018 X-ray microcomputed tomography in additive manufacturing: a review of the current technology and applications *3D Printing and Additive Manufacturing* **5** 227–47
- [62] De Carlo F *et al* 2014 Scientific data exchange: a schema for HDF5-based storage of raw and analyzed data *J. Synchrotron Radiat.* **21** 12241230
- [63] Du Y, Wang R, Biguri A, Zhao X, Peng Y and Wu H 2022 TIGRE-VarianCBCT for on-board cone-beam computed tomography, an open-source toolkit for imaging, dosimetry and clinical research *Phys. Medica* **102** 33–45
- [64] Meyer S *et al* 2019 Dosimetric accuracy and radiobiological implications of ion computed tomography for proton therapy treatment planning *Phys. Med. Biol.* **64** 125008
- [65] Schulte R *et al* 2004 Conceptual design of a proton computed tomography system for applications in proton radiation therapy *IEEE Trans. Nucl. Sci.* **51** 866–72
- [66] Kaser S, Bergauer T, Biguri A, Birkfellner W, Hatamikia S, Hirtl A, Irmeler C, Kirchmayer B and Ulrich-Pur F 2022 Extension of the open-source TIGRE toolbox for proton imaging *Zeitschrift für Medizinische Physik* **33** 552–66
- [67] C-A, Brousmiche S, Portillo S K N, Beaulieu L and Seco J 2016 A maximum likelihood method for high resolution proton radiography/proton CT *Phys. Med. Biol.* **61** 8232
- [68] Adler J and Öktem O 2018 and Ozan Öktem. Learned primal-dual reconstruction *IEEE Trans. Med. Imaging* **37** 1322–32
- [69] De Carlo F *et al* 2018 Tomobank: a tomographic data repository for computational x-ray science *Meas. Sci. Technol.* **29** 034004
- [70] Gundlach-Graham A, Burger M, Allner S, Schwarz G, Wang H A O, Gyr L, Grolimund D, Hattendorf B and Gunther D 2015 High-speed, high-resolution, multielemental laser ablation-inductively coupled plasma-time-of-flight mass spectrometry imaging: part i. instrumentation and two-dimensional imaging of geological samples *Anal. Chem.* **87** 8250–8
- [71] Burger M, Gundlach-Graham A, Allner S, Schwarz G, Wang H A O, Gyr L, Burgener S, Hattendorf B, Grolimund D and Gunther D 2015 High-speed, high-resolution, multielemental la-icp-tofms imaging: Part ii. critical evaluation of quantitative three-dimensional imaging of major, minor, and trace elements in geological samples *Anal. Chem.* **87** 8259–67
- [72] The Phantom Laboratory Catphan 600a Manual (<https://phantomlab.com/catphan-600>) Accessed: 2024-10-26
- [73] Agostinelli S *et al* 2003 GEANT4—a simulation toolkit *Nucl. Instrum. Methods Phys. Res., Sect. A* **506** 250–303
- [74] Lohvithee M, Rassame S and Hibiki T 2022 Applications of neutron computed tomography to thermal-hydraulics research *Prog. Nucl. Energy* **149** 104262
- [75] Breger A *et al* 2024 A study of why we need to reassess full reference image quality assessment with medical images arXiv:2405.19097
- [76] Breger A *et al* 2024 A study on the adequacy of common iqa measures for medical images arXiv:2405.19224