PROCEEDINGS OF SPIE

Versatile regularisation toolkit for iterative image reconstruction with proximal splitting algorithms

Daniil Kazantsev, Edoardo Pasca, Mark Basham, Martin Turner, Matthias J. Ehrhardt, et al.

SPIE.

Versatile regularisation toolkit for iterative image reconstruction with proximal splitting algorithms

Daniil Kazantsev[a,*]**, Edoardo Pasca**[b]**, Mark Basham**[a]**, Martin Turner**[b]**, Matthias J. Ehrhardt**[c]**, Kris Thielemans**[d]**, Benjamin A. Thomas**[d]**, Evgueni Ovtchinnikov**[b]**, Philip J. Withers**[e]**, Alun W. Ashton**[a]

[a]Diamond Light Source Ltd., Scientific Software Department, Diamond House, Harwell Science & Innovation Campus, Didcot, UK, OX11 0DE
[b]Scientific Computing Department, STFC, Rutherford Appleton Laboratory, Didcot, UK, OX11 0QX
[c]Institute for Mathematical Innovation, University of Bath, Bath, UK, BA2 7JU
[d]Institute of Nuclear Medicine, UCL, London, UK, NW1 2BU
[e]The Henry Royce Institute, School of Materials, The University of Manchester, UK, M13 9PL

Abstract. Ill-posed image recovery requires regularisation to ensure stability. The presented open-source regularisation toolkit consists of state-of-the-art variational algorithms which can be embedded in a plug-and-play fashion into the general framework of proximal splitting methods. The packaged regularisers aim to satisfy various prior expectations of the investigated objects, e.g., their structural characteristics, smooth or non-smooth surface morphology. The flexibility of the toolkit helps with the design of more advanced model-based iterative reconstruction methods for different imaging modalities while operating with simpler building blocks. The toolkit is written for CPU and GPU architectures and wrapped for Python/MATLAB. We demonstrate the functionality of the toolkit in application to Positron Emission Tomography (PET) and X-ray synchrotron computed tomography (CT).

Keywords: X-ray CT, PET, iterative methods, model-based, proximal-dual, regularization.

* The corresponding author daniil.kazantsev@diamond.ac.uk

1 Introduction

In tomography,[1] various mathematical reconstruction algorithms are applied to the measured projection data to obtain a volume. When the measured projection data is corrupt, noisy and/or angularly undersampled, the inversion is ill-posed and regularised iterative methods (RIM) should be used. A suitable regulariser can be chosen according to some prior expectations about the reconstructed object. For instance, in material science the reconstructed object is frequently piecewise-constant while in medical imaging piecewise-smooth.

One of the main disadvantages of using RIM for tomographic reconstruction is the computational cost involved in optimising the objective function which consists of the data fidelity and regularisation terms (see Fig. 1). When the terms of the objective function are differentiable, gradient or Hessian-based optimisation methods can be employed. The smoothness constraint, however, might not always be a desirable feature of the objective and one needs to resort to non-smooth optimisation strategies.

The framework of proximal splitting operators (PSO)[2] can be applied to indifferentiable cost functions and allows the objective decoupling resulting in simpler, frequently parallelisable, optimisation steps. This enables fast prototyping and implementation of novel RIMs.[3]

In this paper, we introduce the CCPi-Regularisation Toolkit (`CCPi-RGL`)[4] for effective and efficient regularisation of iterative methods for application the reconstruction of large 3D tomographic datasets. We demonstrate how the `CCPi-RGL` toolkit can be used to reconstruct positron emission tomography (PET) data with the STIR-SIRF package[5,6] and also high-resolution synchrotron tomographic data using MPI-based Savu software.[7]

2 `CCPi-RGL` software contents and methodology

The `CCPi-RGL` software contains various state-of-the-art variational regularisation techniques which include a second and fourth-order diffusion-based methods as well as local and non-local approaches. We refer the reader to the dedicated paper[4] where all available methods we describe in detail and benchmarks provided to help with method selection.
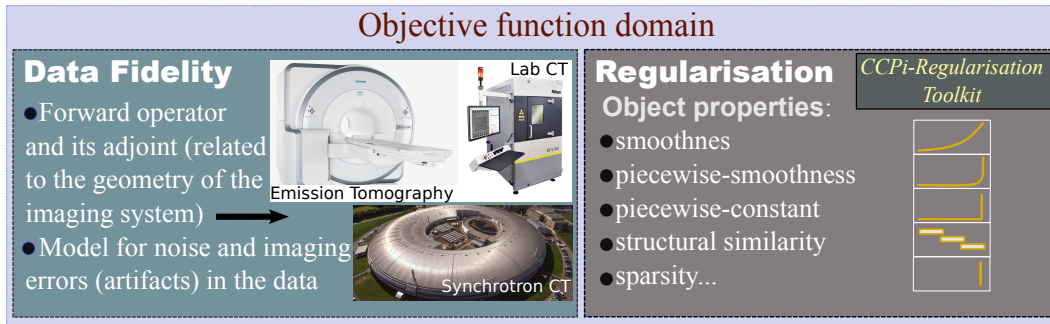

Fig 1 `CCPi-RGL` within the general optimisation framework.

The general optimisation form for tomographic image reconstruction can be formulated as:

$$\min_{\boldsymbol{x} \in \mathbb{R}^N} \mathcal{F}(\boldsymbol{x}) + g(\boldsymbol{x}) \equiv f(\boldsymbol{A}\boldsymbol{x}) + g(\boldsymbol{x}) \equiv \sum_{i=1}^{n} f_i(\boldsymbol{A}_i \boldsymbol{x}) + g(\boldsymbol{x}), \tag{1}$$

where $f_i : \mathbb{R}^{M_i} \to \mathbb{R}$, $f : \mathbb{R}^M \to \mathbb{R}$ is a continuously differentiable convex function with Lipschitz continuous gradient. Thus, $\mathcal{F}$ also has Lipschitz continuous gradient and we denote its constant by L. The functions f_i measure the *fidelity* of $\boldsymbol{A}\boldsymbol{x}$ to the projection data $\boldsymbol{b} \in \mathbb{R}^M$ where $\boldsymbol{A} = (\boldsymbol{A}_1; \ldots; \boldsymbol{A}_n) \in \mathbb{R}^{M \times N}$ is the linear forward operator and $\boldsymbol{x} \in \mathbb{R}^N$ is the unknown solution. The *regularisation* term $g : \mathbb{R}^N \to \mathbb{R}$ is a convex, possibly non-differentiable function expressing a prior knowledge of the unknown estimate $\boldsymbol{x}$ (see Fig. 1 (left)). In X-ray CT, the common choice for the data fidelity term is the Least-Squares (LS) model: $f(\boldsymbol{A}\boldsymbol{x}) = \|\boldsymbol{A}\boldsymbol{x} - \boldsymbol{b}\|_2^2$ or the Penalised Weighted Least Squares (PWLS): $f(\boldsymbol{A}\boldsymbol{x}) = \|\boldsymbol{A}\boldsymbol{x} - \boldsymbol{b}\|_W^2$, where $\boldsymbol{W} \in \mathbb{R}^{M \times M}$ is a diagonal matrix such as $\{W_{ii} = 1/\sigma_i^2\}_{i=1}^M$ and $\sigma_i^2 \approx y_i^2$ is the variance of the measurements. It is not uncommon, especially for emission tomography,[5] to use the more realistic non-linear Poisson model $f(\boldsymbol{A}\boldsymbol{x}) = \langle \boldsymbol{y}, \boldsymbol{A}\boldsymbol{x} \rangle + \langle \boldsymbol{z} \exp(-\boldsymbol{A}\boldsymbol{x}), \boldsymbol{1} \rangle$.

In order to solve problem (1) efficiently, we rely on the theory of the proximal methods[2] which helps to split the problem into parts which are easier to solve. Before presenting various splitting approaches, we introduce the notion of the *proximal operator*:

$$\text{prox}_{\tau g}(\boldsymbol{u}) = \min_{\boldsymbol{x} \in \mathbb{R}^N} g(\boldsymbol{x}) + \frac{1}{2\tau} \|\boldsymbol{x} - \boldsymbol{u}\|^2. \tag{2}$$

All regularisation algorithms of `CCPi-RGL` aim to solve (2) and therefore one needs to be concerned only with the f-related sub-problem which is specific to the imaging modality.

Here we present some PSO-based algorithms in which `CCPi-RGL` can be easily integrated.[4] A well-known FISTA algorithm[8] (see Alg. 1) can ensure $O(1/k^2)$ rate of convergence. FISTA requires f in (1) to be Lipschitz differentiable and include only one proximal step on each iteration

---

Algorithm 1 Fast Iterative Shrinkage-Thresholding Algorithm (FISTA)

Input: $\boldsymbol{x}^0 \in \mathbb{R}^N$, K;
  $\tau = 1/L, t_0 = 1$
  for $k = 0$ to $K - 1$ **do**
    1. $\boldsymbol{y}^k = \boldsymbol{x}^k + \left(\frac{t_k - 1}{t_{k+1}}\right)(\boldsymbol{x}^k - \boldsymbol{x}^{k-1})$
    2. $\boldsymbol{x}^{k+1} = \mathrm{prox}_{\tau g}\left(\boldsymbol{y}^k - \tau \nabla F(\boldsymbol{y}^k)\right)$
    3. $t_{k+1} = 1 + \sqrt{1 + 4t_k^2}/2$
  end for

(step 2). The group of *primal-dual* methods[9] relax the differentiability condition but normally rely on two proximal steps instead (see Alg. 2).

Algorithm 2 Primal Dual Hybrid Gradient (PDHG) algorithm

Input: $\boldsymbol{x}^0 \in \mathbb{R}^N$, $c > 0$, K;
  $\sigma = c/\|\boldsymbol{A}\|, \tau = 1/(c\|\boldsymbol{A}\|), \boldsymbol{y}^0 = 0 \in \mathbb{R}^M$
  for $k = 0$ to $K - 1$ **do**
    1. $\boldsymbol{x}^{k+1} = \mathrm{prox}_{\tau g}(\boldsymbol{x}^k - \tau \boldsymbol{A}^T \boldsymbol{y}^k)$
    2. $\boldsymbol{y}^{k+1} = \mathrm{prox}_{\sigma f^*}\left(\boldsymbol{y}^k + \sigma \boldsymbol{A}(2\boldsymbol{x}^{k+1} - \boldsymbol{x}^k)\right)$
  end for

Evaluation $\boldsymbol{A}$ or $\nabla F = \boldsymbol{A}^\top \circ \nabla f \circ \boldsymbol{A}$ in each iteration is cumbersome. An approach to overcome this drawback is *randomisation* which can be achieved within PDHG by selecting only a few dual variables $\boldsymbol{y}_i$ in each iteration, resulting in the stochastic PDHG[10,11] method (see Alg. 3).

Algorithm 3 Stochastic PDHG algorithm

Input: $\boldsymbol{x}^0 \in \mathbb{R}^N$, $c > 0$, K;
  $\sigma_i = c/\|\boldsymbol{A}_i\|, \tau = 1/(cn \max_j \|\boldsymbol{A}_j\|), \boldsymbol{y}^0 = 0 \in \mathbb{R}^M, \boldsymbol{z}^0 = 0 \in \mathbb{R}^N$
  for $k = 0$ to $K - 1$ **do**
    1. $\boldsymbol{x}^{k+1} = \mathrm{prox}_{\tau g}(\boldsymbol{x}^k - \tau \overline{\boldsymbol{z}}^k)$
    2. Select $j \in \{1, \ldots, n\}$ uniform at random.
    3. $\boldsymbol{y}_i^{k+1} = \begin{cases} \mathrm{prox}_{\sigma_i f_i^*}\left(\boldsymbol{y}_i^k + \sigma_i \boldsymbol{A}_i \boldsymbol{x}^{k+1}\right), & \text{if } i = j \\ \boldsymbol{y}_i^k, & \text{else} \end{cases}$
    4. $\Delta \boldsymbol{z} = \boldsymbol{A}_j^\top(\boldsymbol{y}_j^{k+1} - \boldsymbol{y}_j^k)$
    $\boldsymbol{z}^{k+1} = \boldsymbol{z}^k + \Delta \boldsymbol{z}, \overline{\boldsymbol{z}}^{k+1} = \boldsymbol{z}^k + n\Delta \boldsymbol{z}$
  end for

Under the linearisation conditions, the PDHG method becomes the well-known ADMM method[2] (see Alg. 4). Step 1. of the Alg. 4 is a quadratic optimisation problem when the data fidelity is chosen to be PWLS: $F(\boldsymbol{x}) = 1/2\|\boldsymbol{A}\boldsymbol{x} - \boldsymbol{b}\|_{\boldsymbol{W}}^2$. Therefore one needs to solve: $\boldsymbol{x}^{k+1} = (\boldsymbol{I} + \tau \boldsymbol{A}^\top \boldsymbol{W} \boldsymbol{A})^{-1}(\tau \boldsymbol{A}^\top \boldsymbol{W} \boldsymbol{b} + \boldsymbol{v}^k - \boldsymbol{u}^k)$ for which Krylov-type methods or Newton solvers can be used.

Algorithm 4 Alternating Directions of Multipliers (ADMM)

Input: $\boldsymbol{x}^0 \in \mathbb{R}^N$, step $\tau > 0$, K;
   $\boldsymbol{u}^0 = 0 \in \mathbb{R}^N$, $\boldsymbol{y}^0 = \boldsymbol{x}^0$
   for $k = 0$ to $K-1$ **do**
      1. $\boldsymbol{x}^{k+1} = \text{prox}_{\tau F}\left(\boldsymbol{v}^k - \boldsymbol{u}^k\right)$
      2. $\boldsymbol{v}^{k+1} = \text{prox}_{\tau g}\left(\boldsymbol{x}^{k+1} + \boldsymbol{u}^k\right)$
      3. $\boldsymbol{u}^{k+1} = \boldsymbol{u}^k + \boldsymbol{x}^{k+1} - \boldsymbol{v}^{k+1}$
   end for

---

3 Iterative reconstruction of tomographic data using the `CCPi-RGL` toolkit

Here we briefly demonstrate the usage of the `CCPi-RGL` toolkit applied to tomographic data of two different imaging modalities. Due to the limited space, the main aim is not to quantify or qualify the results, but rather to demonstrate the applicability of `CCPi-RGL` to different settings.

Our first application is to reconstruct Positron Emission Tomography (PET) data using the SPDHG[10] algorithm with the Kullback-Leibler (Poisson) data model and two different regularisers from `CCPi-RGL`: FGP-TV[8] and TGV.[12] We use STIR-SIRF[5,6] software to model realistic PET system blur and apply Poisson noise to the projection data generated from the thorax phantom (see Fig. 2). We run 50 iterations of the SPDHG algorithm and see that the reconstruction without
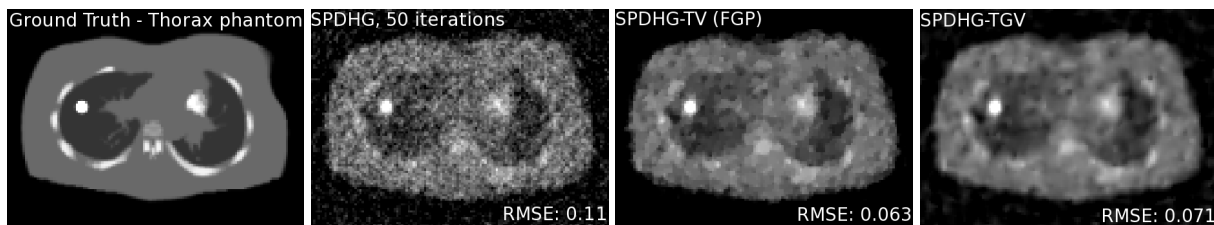

Fig 2 The SPDHG reconstructions of simulated PET data using no regularisation, FGP-TV, and TGV regularisers.

the regularisation is very noisy. Noise can be damped significantly by using regularisation. The higher-order priors, such as TGV might be more preferable than piecewise-constant TV models for emission tomography.

In Fig. 3, we apply `CCPi-RGL` to reconstruct time-resolved (4D) data of ice-cream under *in*
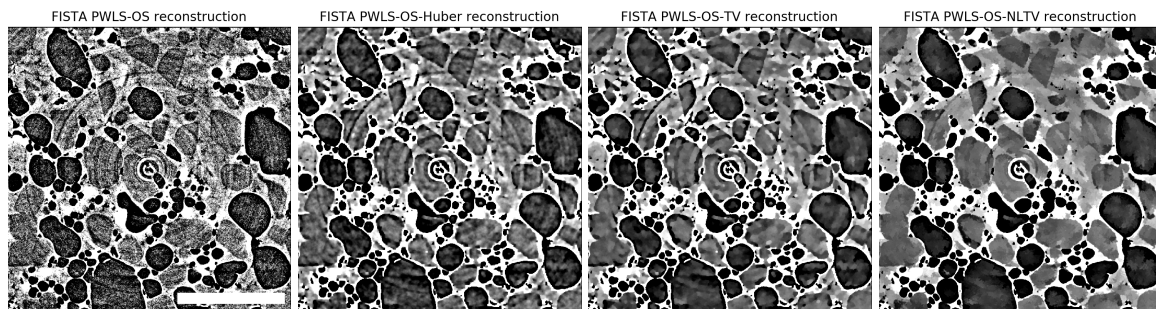

Fig 3 Cropped 1k^2 pixels axial slice of reconstructions of 2.5k^3 voxels volume using Huber (smoothed TV), FGP-TV[8] and NLTV[13] regularisers with FISTA PWLS-OS method (12 iterations). The scale bar corresponds to 100 μm.

situ thermal cycling conditions[14] collected at Diamond Light Source (DLS). We use the ordered-subsets accelerated FISTA reconstruction method (see Alg. 1) with the PWLS data model. Qualitatively the NLTV[13] regulariser outperforms its local rivals.

Additionally, we reconstructed dynamic alloy solidification data also obtained at DLS.[15] We use the ADMM algorithm (see Alg. 4) with the SB-TV[16] regulariser. It is evident that the regularised reconstruction can produce much improved images in terms of signal-to-noise ratio and resolution compare to direct techniques (Filtered Backprojection) or unregularised iterative methods.
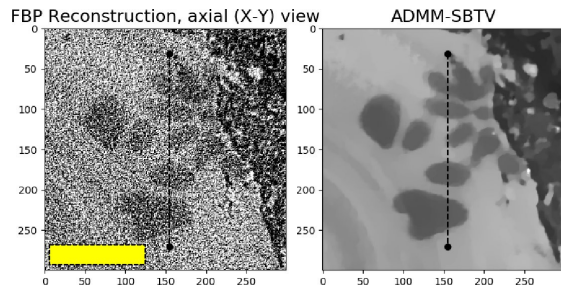

Fig 4 Magnified region of a reconstructed $1k^3$ voxels volume from the 3D projection data. Scale bar corresponds to 200μm.

4 Conclusions and acknowledgments

The presented open–source software `CCPi-RGL` can be used for tomographic image reconstruction across different imaging modalities. A selection of regularisers provides a desirable customisation and abstraction for fast prototyping of novel reconstruction algorithms. Due to its computationaly efficient implementation, the developed methods can be applied to big data problems across various disciplines.

References

1   T. M. Buzug, *Computed Tomography: From Photon Statistics to Modern Cone-Beam CT*, Springer, Berlin, Germany (2008).

2   A. Chambolle and T. Pock, "An introduction to continuous optimization for imaging," *Acta Numerica* **25**, 161–319 (2016).

3   S. V. Venkatakrishnan, C. A. Bouman, and B. Wohlberg, "Plug-and-play priors for model based reconstruction," in *IEEE In Global Conference on Signal and Information Processing*, 945–948, IEEE (2013).

4   D. Kazantsev, E. Pasca, M. Turner, *et al.*, "Ccpi-regularisation toolkit for computed tomographic image reconstruction with proximal splitting algorithms," *SoftwareX* **accepted**, – (2019).

5   K. Thielemans, C. Tsoumpas, S. Mustafovic, *et al.*, "Stir: software for tomographic image reconstruction release 2," *Physics in Medicine & Biology* **57(4)**, 867 (2012).

6   E. Ovtchinnikov, D. Atkinson, C. Kolbitsch, *et al.*, "Sirf: Synergistic image reconstruction framework," in *2017 IEEE Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC)*, 1–3, IEEE (2017).

7   N. Wadeson and M. Basham, "Savu: A python-based, mpi framework for simultaneous processing of multiple, n-dimensional, large tomography datasets," *arXiv preprint* **arXiv:1610.08015** (2016).

8   A. Beck and M. Teboulle, "A fast iterative shrinkage-thresholding algorithm for linear inverse problems," *SIAM journal on imaging sciences* **2(1)**, 183–202 (2009).

9   A. Chambolle and T. Pock, "A first-order primal-dual algorithm for convex problems with applications to imaging," *Journal of mathematical imaging and vision* **40**(1), 120–145 (2011).

10 A. Chambolle, M. J. Ehrhardt, P. Richtarik, *et al.*, "Stochastic primal-dual hybrid gradient algorithm with arbitrary sampling and imaging applications," *SIAM Journal on Optimization* **28(4)**, 2783–2808 (2018).

11 M. J. Ehrhardt, P. Markiewicz, and C.-B. Schönlieb, "Faster pet reconstruction with non-smooth priors by randomization and preconditioning," *arXiv preprint arXiv:1808.07150* (2018).

12 K. Bredies, K. Kunisch, and T. Pock, "Total generalized variation," *SIAM Journal on Imaging Sciences* **3(3)**, 492–526 (2010).

13 A. Elmoataz, O. Lezoray, and S. Bougleux, "Nonlocal discrete regularization on weighted graphs: a framework for image and manifold processing," *IEEE Trans. Image Processing* **17(7)**, 1047–1060 (2008).

14 E. Guo, D. Kazantsev, J. Mo, *et al.*, "Revealing the microstructural stability of a three-phase soft solid (ice cream) by 4d synchrotron x-ray tomography," *Journal of Food Engineering* **237**, 204–214 (2018).

15 D. Kazantsev, E. Guo, A. Phillion, *et al.*, "Model-based iterative reconstruction using higher-order regularization of dynamic synchrotron data," *Measurement Science and Technology* **28**(9), 094004 (2017).

16 T. Goldstein and S. Osher, "The split bregman method for l1-regularized problems," *SIAM journal on imaging sciences* **2(2)**, 323–343 (2009).