



King's Research Portal

DOI:

[10.1016/j.cpc.2019.107087](https://doi.org/10.1016/j.cpc.2019.107087)

Document Version

Peer reviewed version

[Link to publication record in King's Research Portal](#)

Citation for published version (APA):

Ovtchinnikov, E., Brown, R., Kolbitsch, C., Pasca, E., da Costa-Luis, C. O., Gillman, A., Thomas, B., Efthimiou, N., Mayer, J., Wadhwa, P., Ehrhardt, M., Ellis, S., Jørgensen, J., Matthews, J., Prieto Vasquez, C., Reader, A. J., Tsoumpas, C., Turner, M., Atkinson, D., & Thielemans, K. (2020). SIRF: Synergistic Image Reconstruction Framework. *COMPUTER PHYSICS COMMUNICATIONS*, 249, Article 107087. <https://doi.org/10.1016/j.cpc.2019.107087>

Citing this paper

Please note that where the full-text provided on King's Research Portal is the Author Accepted Manuscript or Post-Print version this may differ from the final Published version. If citing, it is advised that you check and use the publisher's definitive version for pagination, volume/issue, and date of publication details. And where the final published version is provided on the Research Portal, if citing you are again advised to check the publisher's website for any subsequent corrections.

General rights

Copyright and moral rights for the publications made accessible in the Research Portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognize and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the Research Portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the Research Portal

Take down policy

If you believe that this document breaches copyright please contact librarypure@kcl.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

SIRF: Synergistic Image Reconstruction Framework

Evgueni Ovtchinnikov^a, Richard Brown^{b,*}, Christoph Kolbitsch^f,
 Edoardo Pasca^a, Casper da Costa-Luis^d, Ashley G. Gillman^k,
 Benjamin A. Thomas^b, Nikos Efthimiou^e, Johannes Mayer^f,
 Palak Wadhwa^{h,j}, Matthias J. Ehrhardt^l, Sam Ellis^d, Jakob S. Jørgensenⁱ,
 Julian Matthews^g, Claudia Prieto^d, Andrew J. Reader^d,
 Charalampos Tsoumpas^{h,j}, Martin Turnerⁱ, David Atkinson^c,
 Kris Thielemans^{b,**}

^a*Scientific Computing Department, Rutherford-Appleton Laboratory, UK Research and Innovation, Harwell Campus, Didcot OX11 0QX, UK*

^b*Institute of Nuclear Medicine, University College London, 235 Euston Rd, London NW1 2BU, UK*

^c*Centre for Medical Imaging, University College London, 43-45 Foley Street, London W1W 7TS, UK*

^d*School of Biomedical Engineering & Imaging Sciences, King's College London, Lambeth Wing, St. Thomas' Hospital, Westminster Bridge Road, London SE1 7EH, UK*

^e*PET Research Centre, School of Health Sciences, University of Hull, Hull HU6 7RX, UK*

^f*Physikalisch-Technische Bundesanstalt (PTB), Bundesallee 100, Braunschweig D-38116 and Abbestrasse 2-12, Berlin D-10587, Germany*

^g*Division of Informatics, Imaging and Data Sciences, University of Manchester, 46 Grafton Street, Manchester M13 9NT, UK*

^h*Biomedical Imaging Science Department, University of Leeds, Leeds LS2 9JT, UK*

ⁱ*School of Mathematics, University of Manchester, Oxford Road, Manchester M13 9PL, UK*

^j*Invicro London, Du Cane Rd, London W12 0NN, UK*

^k*Australian e-Health Research Centre, CSIRO, Brisbane, Queensland 4029, Australia*

^l*Institute for Mathematical Innovation, University of Bath, Convocation Ave, Bath BA2 7JU, UK*

Abstract

The combination of positron emission tomography (PET) with magnetic resonance (MR) imaging opens the way to more accurate diagnosis and improved patient management. At present, the data acquired by PET-MR scanners

*Joint first author

**Corresponding author. *E-mail address:* k.thielemans@ucl.ac.uk

1
2
3
4
5
6
7
8
9 are essentially processed separately, but the opportunity to improve accuracy
10 of the tomographic reconstruction via synergy of the two imaging techniques
11 is an active area of research.
12

13 In this paper, we present Release 2.1.0 of the CCP-PETMR Synergistic
14 Image Reconstruction Framework (SIRF) software suite, providing an open-
15 source software platform for efficient implementation and validation of novel
16 reconstruction algorithms. SIRF provides user-friendly Python and MAT-
17 LAB interfaces built on top of C++ libraries. SIRF uses advanced PET and
18 MR reconstruction software packages and tools. Currently, for PET this is
19 Software for Tomographic Image Reconstruction (STIR); for MR, Gadgetron
20 and ISMRMRD; and for image registration tools, NiftyReg. The software
21 aims to be capable of reconstructing images from acquired scanner data,
22 whilst being simple enough to be used for educational purposes.
23

24 The most recent version of the software can be downloaded from <http://www.ccppetmr.ac.uk/downloads> and <https://github.com/CCPPETMR/>.
25
26

27 *Keywords:* Image Reconstruction; PET-MR; Multi-modality; Medical
28 Imaging; Open-Source Software
29
30
31

32 **PROGRAM SUMMARY**

33 *Program Title:* Synergistic Image Reconstruction Framework (SIRF)
34

35 *Licensing provisions (please choose one):* GPL 3 and Apache-2.0
36

37 *Programming languages:* C++, C, Python, MATLAB
38

39 *Nature of problem:*

40 In current practice, data acquired by PET-MR scanners are processed separately.
41 Methods for improving the accuracy of the tomographic reconstruction using the
42 synergy of the two imaging techniques are actively being investigated by the PET-
43 MR research and development community, however, practical application is heav-
44 ily reliant on software. Open-source software available to the PET-MR commu-
45 nity – such as the PET package Software for Tomographic Image Reconstruction
46 (STIR) [1] and the MR package Gadgetron [2] – provide a basis for new syner-
47 gistic PET-MR software. However, these two software packages are independent
48 and have very different software architectures. They are mostly written in C++ but
49 many researchers in the PET-MR community are more familiar with script-style
50 languages, such as Python and MATLAB, which enable rapid prototyping of novel
51 reconstruction algorithms. In the current situation it is difficult for researchers to
52 exploit any synergy between PET and MR data. Furthermore, techniques from
53 one field cannot easily be applied in the other.
54

55 *Solution method:*
56
57
58

1
2
3
4
5
6
7
8
9 In SIRF, the bulk of computation is performed by available advanced open-source
10 reconstruction and registration software (currently STIR, Gadgetron and NiftyReg)
11 that can use multithreading and GPUs. The SIRF C++ code provides a thin layer
12 on top of these existing libraries. The SIRF layer has unified data-containers and
13 access mechanisms. This C++ layer provides the basis for a simple and intuitive
14 Python and MATLAB interface, enabling users to quickly develop and test their
15 reconstruction algorithms using these scripting languages only. At the same time,
16 advanced users proficient in C++ can directly utilise wider SIRF functionality via
17 the SIRF C++ libraries that we provide.
18
19
20
21
22

23 1. Introduction

24
25 Magnetic resonance (MR) images can provide high resolution anatomical
26 images with excellent soft tissue contrast, but provide only limited func-
27 tional and metabolic information. In contrast, positron emission tomogra-
28 phy (PET) is able to provide highly sensitive and specific functional imaging,
29 but at low spatial resolution (4-7 mm). The many PET tracers available can
30 provide information on a wide range of processes. Recognition of the comple-
31 mentary aspects of the two imaging modalities prompted the development
32 of integrated clinical PET-MR scanners. Hybrid PET-MR imaging [3, 4]
33 promises superior application in oncological, paediatric [5], cardiac [6] and
34 neurological [7, 8] imaging, both in clinical and research settings. However,
35 the technical hardware developments have not yet been paralleled by the
36 development of an integrated software platform for processing data acquired
37 by such scanners.
38
39
40

41 Image reconstruction is a challenging problem for each modality, albeit
42 for somewhat different reasons. PET image quality is limited by noise and
43 resolution thus it requires regularisation during image reconstruction [9]. On
44 the other hand, MR data are often undersampled for acquisition speed [10],
45 and model-based reconstruction is often needed [11, 12]. Both modalities
46 could benefit from additional information in the reconstruction. For PET,
47 this has most often been exploited by using a high quality MR image as “side-
48 information” [13]. Similar strategies can be used for dynamic MR data [14].
49 Recently, researchers have started exploring methods for joint reconstruction
50 of both data-sets, exploring the underlying structural similarity [15, 16, 17,
51 18, 19]. However, this research has been confined to single institutions using
52 their own internal software packages.
53
54
55
56
57
58
59
60
61
62
63
64
65

1
2
3
4
5
6
7
8
9 Many open-source packages exist for single-modality image reconstruction, for instance for PET [1, 20, 21, 22, 23] or MRI [24, 2] (an extensive list of Open Source imaging software is provided at <https://www.opensourceimaging.org/category/software>, and https://www.ismrm.org/mri_unbound/sequence.htm and <http://stir.sourceforge.net/links> contain numerous links to MR and PET software). However, none of these packages allow the reconstruction of data from both modalities, and are therefore not suitable for joint reconstruction of PET and MR data. The Collaborative Computational Project for PET-MR (CCP-PETMR, <http://www.ccppetmr.ac.uk>) is a network established in April 2015 with the aim of facilitating the investigation of novel synergistic PET-MR image reconstruction methods by providing the PET-MR research community with a software-development platform. The platform should be simple enough to use for research and educational purposes and, at the same time, powerful enough to be able to handle raw data acquired by PET, MR and PET-MR scanners within a reasonable processing time. Both of the available integrated clinical PET-MR systems, the Siemens 3T Biograph mMR (Siemens Healthineers, Erlangen, Germany) [25] and the General Electric Signa PET/MR [26, 27], are being targeted by this software platform, and a suitable data format exchange is being negotiated with the manufacturers. This paper describes Release 2.1.0 of the open-source Synergistic Image Reconstruction Framework (SIRF).

2. SIRF structure

2.1. Rationale

41 Image reconstruction requires complex software and support for different
42 types of scanners. Developing software that handles two substantially differ-
43 ent imaging modalities is therefore a challenging task. While it would have
44 been possible to extend one of the many open-source software projects that
45 aim to support a single modality, this would inevitably have led to duplicat-
46 ing the effort. Instead, SIRF is designed to act as a “framework” providing
47 a layer on top of existing packages for medical image reconstruction, referred
48 to as *engines* in this document. This layer aims to present a consistent in-
49 terface for data manipulation and image reconstruction, independent of the
50 modality and the underlying engine. Such an architecture has the additional
51 advantage that it allows SIRF to benefit from the active development in each
52 of the underlying engines. Note that these design choices are similar to those
53
54
55
56
57
58

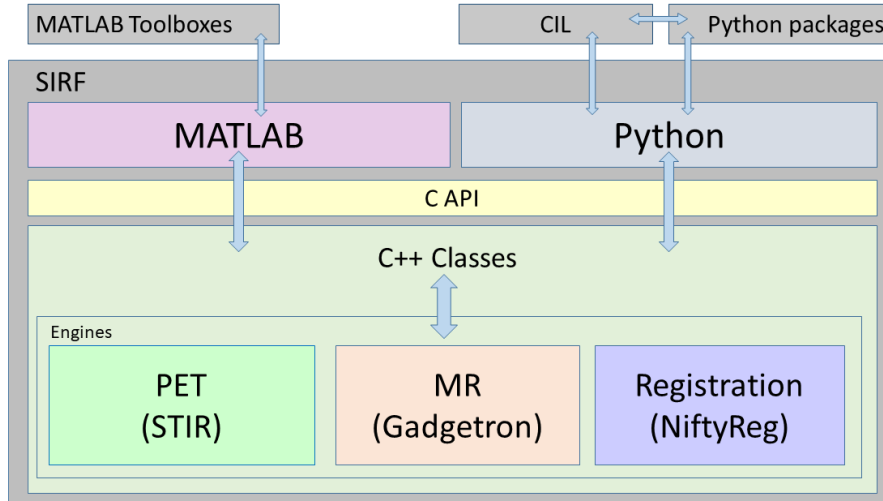


Figure 1: The architecture of SIRF: interfacing of underlying engines and most of the development is in C++. Interfacing MATLAB and Python with the underlying C++ code is done via a thin C API. CIL is a Python library described in §5.1.1.

of Operator Discretisation Library (ODL) [28] and Core Imaging Library (CIL), see §5.1.1. However, both are Python-only and do not support MRI.

An important choice is the programming language that the users of our platform will need to adopt for their own software development. Two languages and environments that are widely used by researchers are MATLAB and Python, in part because of the excellent prototyping abilities of these languages, while C++ is more commonly used for intensive computational applications, often benefiting from multi-threading via OpenMP. Most of the existing open-source software packages consist of a core C++-library, sometimes with interfaces to MATLAB and/or Python. In SIRF, we have chosen to do most of the development in C++, including the interfacing to the underlying engines, as described in Figure 1. The C++ code is then interfaced to simple Python and MATLAB building blocks. We aim to have close similarity between the C++, Python and MATLAB interfaces to SIRF, but expect that most of the users of our platform will use Python or MATLAB, and these interfaces are therefore currently most complete and well-documented.

SIRF uses an object-oriented approach to programming. A short description of this approach is presented in Appendix A and is illustrated by SIRF

1
2
3
4
5
6
7
8
9 usage examples.

10 11 *2.2. Underlying engines*

12
13 At present, we use STIR (Software for Tomographic Image Reconstruction) [1] as a PET engine, Gadgetron [2] as an MR engine and NiftyReg as
14 the engine for registration and resampling [29, 30]. We briefly describe these
15 engines here, but refer to their respective publications and web-sites for more
16 detail.
17
18

19 STIR [1] implements a library of C++ classes for performing PET recon-
20 struction and related tasks such as data input / output. Parameters of a
21 STIR reconstructor object are normally defined by the user in a text-based
22 parameter file, and a set of executables is provided that read the parameter
23 file specified in the command line and perform the required tasks, so that
24 the user does not need to be an expert in C++ or any other programming
25 language. For users familiar with Python or MATLAB, the current devel-
26 opment version of STIR offers an alternative usage via Python / MATLAB
27 scripts. This requires building Python and MATLAB interfaces to STIR by
28 SWIG [31]. However, these interfaces currently do not cover all functionality
29 of STIR. Moreover, they expose the complexity of underlying C++ library.
30
31

32 With Gadgetron [2], MR image reconstruction is performed by specifying
33 a chain of ‘gadgets’, i.e. pieces of code implementing specific tasks within
34 a streaming architecture. The chain of gadgets is executed on a server,
35 which can be launched in the background of the same computer, or it can
36 be another computer, or a Virtual Machine, or a cloud-based computational
37 infrastructure. A client can communicate with the server via TCP/IP. In
38 order to set up the gadget chain, the server needs to receive a text file in
39 the Extensible Markup Language (i.e. xml file format) describing the chain.
40 Then, the first gadget in the chain is activated and waits for the acquired data
41 to arrive from the client in chunks of a certain size. Having processed a chunk
42 of data, the first gadget passes the result to the second, starts processing the
43 next chunk and so on. The last gadget sends the reconstructed images back
44 to the client. As with STIR, a Gadgetron user is not required to have any
45 knowledge of C++. Instead, the reconstruction tasks can be configured via
46 parameter files in xml format.
47
48

49 NiftyReg [29, 30] is a C++ image registration library, restricted to images
50 stored in the NIfTI image format [32]. Rigid, affine and non-rigid image
51 registrations are possible, and these are available in both symmetric and
52 non-symmetric versions. From the registrations, it is possible to extract
53
54
55
56
57
58

1
2
3
4
5
6
7
8
9 transformation matrices (for rigid and affine registrations), and deformation
10 and displacement field images. With motion information extracted, it is then
11 possible to perform resampling as required.
12

13 *2.3. SIRF software layers*

14 *2.3.1. C++*

15 The C++ layer provides the basic class hierarchies, in particular data con-
16 tainer classes, and specific implementations where we wrap the functional-
17 ity provided by the engines as well as implementing any extra components
18 needed for the general usage of SIRF.
19
20
21
22

23 *2.3.2. Interfacing into Python and MATLAB*

24 At present, the SIRF user interface consists of the modules `sirf.Gadgetron`,
25 `sirf.STIR` and `sirf.Reg` wrapping the engines, together with some addi-
26 tional functionality in `sirf.Utilities`.
27

28 To avoid difficulties interfacing between the C++ code and the targeted
29 script languages, we include an additional layer by wrapping C++ code into
30 C. This extra layer simplifies interfacing and widens the scope of our targets
31 to practically any programming language in use.
32

33 Interfacing of C into MATLAB is facilitated by the availability of MAT-
34 LAB functions (`loadlibrary`, `calllib` etc.) that allow direct calls of func-
35 tions from C libraries (with a little ‘syntactic sugar’ needed under Windows).
36 For Python we use SWIG [31], which can also be used for a number of other
37 prospective languages.
38

39 The MATLAB and Python interfaces based on this C library are not
40 user-friendly and not intended to be accessed directly. Instead in addition
41 to these, we have Object-Oriented MATLAB and Python modules mirroring
42 and extending the C++ class structure.
43
44
45

46 **3. SIRF objects**

47 This section briefly describes the type of objects with which SIRF users
48 need to operate in order to perform image reconstruction and related tasks. A
49 more detailed description is provided in the SIRF User’s Guide, and examples
50 are given in §4.
51
52
53
54
55
56
57
58

1
2
3
4
5
6
7
8
9 *3.1. Image and acquisition data objects*

10 In SIRF, acquired data are handled by objects of class `AcquisitionData`
11 and images by objects of class `ImageData`. Both are derived from an abstract
12 data class `DataContainer` that implements data algebra: the user’s code can
13 compute norms, dot products and linear combinations of data objects of the
14 same class viewed as vectors.
15
16

17 We created a hierarchy of image data objects that allows us to handle
18 representations of images employed in different reconstruction engines in a
19 unified way. On the top of this hierarchy, we have an abstract base class
20 `ImageData`, from which the abstract classes `PETImageData` and `MRImageData`
21 are derived. Further down this hierarchy, we have various classes encapsulat-
22 ing different representations (e.g file or array) of image data used in different
23 reconstruction engines, such as `STIRImageData`. SIRF uses a patient-based
24 coordinate system to encode the physical position of a given volume element.
25 This allows spatial context to be integrated between images from different
26 engines.
27
28
29

30
31 *3.2. Acquisition models*

32 An ‘acquisition model’ is the terminology for a mathematical model of
33 the scanner acquisition process, which can be represented by
34

$$35 \qquad A(x) = b \qquad (1)$$

36
37 where x represents the scanned object, typically a vector of values at points
38 in three-dimensional space defining a discretised image, b approximates the
39 actual data collected by the scanner and A is an operator (generally non-
40 linear) that models the acquisition process.
41
42

43 The computation of $A(x)$ for a given image data representation x is re-
44 ferred to as the *forward projection*. The *backprojection* is the application of
45 the adjoint of the Fréchet derivative of A to acquisition data. For linear mod-
46 els, the backprojection corresponds to the Hermitian conjugate operation.
47
48

49 In SIRF, PET and MR acquisition models are implemented by their re-
50 spective `AcquisitionModel` classes.
51

52 *3.3. Reconstructors*

53 Reconstructor classes implement algorithms for solving equation (1). In
54 the current SIRF version, we expose Gadgetron MR reconstructors that im-
55 plement direct algorithms for Cartesian acquisitions using Fourier transforms.
56
57
58

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

Planned future versions will also expose iterative algorithms for Cartesian and non-Cartesian acquisitions. For PET, the SIRF Filtered Back Projection reconstructor implements a direct algorithm. Other reconstructors implement iterative algorithms maximising the (optionally penalised) Poisson log-likelihood (see §3.4), with the possibility of using ordered subsets [33].

With SIRF iterative reconstructors, the user has an option of performing one iteration at a time in a loop, which allows them to inspect the current state of the iterated image, apply their own filtering and so on.

3.4. Objective functions

SIRF iterative reconstructors maximise a certain objective function related to the acquisition model (1). The user is provided with facilities to compute the value of this function and its gradient, so that they can employ their own or third-party optimisation algorithms. Penalty terms can be added to the objective function, which opens a way to employing images produced by an MR reconstructor as anatomical priors for PET reconstruction [34, 35, 13], see also §5.1.2.

3.5. Registration and resampling

As previously mentioned, through the wrapping of NiftyReg SIRF provides registration and resampling functionality. For rigid and affine registrations, the user may use `NiftyAladinSym`; for non-rigid registration, the user may use `NiftyF3dSym`. Lastly, resampling can be carried out with the use of `NiftyResample`. Example usage of the registration and resampling functionality can be found in §4.3.

Once registrations have been performed, it is clearly important to be able to manipulate the motion information. Affine transformation matrices can therefore be handled with `AffineTransformation`, and displacement and deformation field images are handled with `NiftiImageData3DDisplacement` and `NiftiImageData3DDeformation`, respectively.

It is possible to convert between these different transformation types (although, naturally, it is not possible to convert from non-rigid transformations back to transformation matrices). Moreover, SIRF provides functionality to compose multiple transformations (of differing types if desired), as well as averaging affine transformation matrices via their quaternions.

1
2
3
4
5
6
7
8
9 *3.6. Other functionality*

10 In addition to the functionality described in §3.1-§3.5, Python and MAT-
11 LAB interfaces provide functionality for data input/output, object cloning,
12 inspecting data dimensions and so on.
13

14 Data-access functionality warrants a special note, as all the actual acquisi-
15 tion and image data are handled exclusively by the engines. For performance
16 reasons, the user does not have direct access to the underlying engine-specific
17 data structures and normally would not need to know whether it is stored
18 in the memory, in a file, in GPU memory, etc. Each data container has a
19 method `as_array` that returns a copy of its data as a Python or MATLAB
20 array, and a method `fill` that fills the container with the data from a user's
21 array. In addition, SIRF contains converters to directly go between data
22 structures of the different engines, if possible, albeit usually with some loss
23 of associated meta-data.
24
25
26
27

28
29 **4. Examples of usage**

30
31 *4.1. PET*

32 Although many other radiotracers are available, PET is most often used to
33 scan patients who have been injected with an ^{18}F -labelled Fluorodeoxyglu-
34 cose (FDG) solution that provides information about glucose metabolism.
35 FDG is commonly used to diagnose and stage tumours in oncology, but it
36 also has uses in various other clinical domains such as in cardiology, which
37 helps to assess the myocardial viability. PET data acquisition is usually
38 carried out for about 30 minutes, about 60 minutes after the tracer injection.
39

40 Here, we show how an FDG-PET scan acquired on a 3T Siemens Bio-
41 graph mMR [25] can be reconstructed using an ordered subsets expectation
42 maximisation (OSEM) algorithm [33]. The following pseudo-code uses the
43 Python interface of SIRF but can easily be adapted for MATLAB. Only parts
44 of the necessary code are provided here. For a detailed line-by-line tutorial
45 refer to the example file `reconstruct_from_listmode.py`.
46
47

48 The following files are required for this example reconstruction (the vari-
49 ables for the filenames for the following code snippets are given in brackets):
50
51

- 52 • Normalisation file in Siemens interfile format (`norm_file`)
- 53
- 54 • Attenuation map for the subject (usually obtained from a dedicated
55 MR scan) in interfile format (`human_attenuation_file`).
56
57

- Attenuation map for the table (and MR receive coils for head scans) in interfile format (`hardware_attenuation_file`).
- Listmode data in Siemens PETLINK format with Siemens interfile-like header (`listmode_file`).
- Template for the projection data calculation in interfile format. This has to fit to the scanner geometry. Examples for this can be found in the sub-module `data`, for instance `mMR_template_span11.hs` in the `data/examples/PET/mMR/` folder.

As a first step, the listmode data are transformed to sinogram space using:

```
import sirf.STIR as PET
lm2sino = PET.ListmodeToSinograms()
```

The `lm2sino` object is then provided with the listmode file (`listmode_file`), the template file (`template_file`) and information about the time interval which shall be transformed to sinogram space. The calculation of sinogram data is performed by the method `process`, and the calculated data are returned by the method `get_output`. After calculating the sinogram data, a sinogram of the mean of the random events can also be estimated [36] using the `lm2sino` object:

```
randoms = lm2sino.estimate_randoms()
```

The attenuation maps can be read in as `ImageData` objects and then combined:

```
attn_image_hum = PET.ImageData(human_attenuation_file)
attn_image_hw = PET.ImageData(hardware_attenuation_file)
attn_image = attn_image_hum + attn_image_hw
```

Figure 2 shows an example for the combined attenuation map.

After combining the attenuation data and the information from the normalisation file, a reconstruction object for OSEM is created:

```
recon = PET.OSMAPOSLReconstructor()
recon.set_num_subsets(21)
recon.set_num_subiterations(63)
```




Figure 2: Combined attenuation map showing both attenuation coefficients for the subject and the scanner table in three different orientations.

The OSMAPOSL reconstructor is an implementation using ordered subsets of Green’s One Step Late algorithm for Maximum A Posteriori reconstruction [37] which reduces to OSEM if no prior is set. Although this algorithm does not converge to the global maximum of the objective function but to an approximate maximum, it is still widely used in practice because it substantially speeds-up the calculations. See §5 for examples of other algorithms.

The objective function is defined as the logarithm of the Poisson likelihood, using a default (ray-tracing) linear acquisition model, with the data set to `acq_data`, the output returned by `lm2sino.get_output()`:

```
obj_fun = PET.make_Poisson_loglikelihood(acq_data)
```

The objective function requires an acquisition model. Here, we explicitly specify to use forward projection by a ray-tracing matrix multiplication:

```
acq_model = AcquisitionModelUsingRayTracingMatrix()
obj_fun.set_acquisition_model(acq_model)
```

but see the example script for details on how to include extra components in the acquisition model, such as attenuation and randoms. Results shown here also include Compton scatter in the acquisition model, in which the scatter is estimated using the Single Scatter Simulation method [38] as implemented

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65




Figure 3: Transversal, sagittal and coronal views of the reconstructed PET images with SIRF reconstructions on the left and the reconstructions by the scanner software on the right. The high uptake of FDG in the heart muscle can be clearly seen.

in STIR [39], although this functionality is currently still on a development branch of SIRF and STIR.

Following the set up of the reconstruction object, the iterative reconstruction can be carried out and the reconstructed PET data can be obtained:

```
recon.process()  
image_data = recon.get_current_estimate()
```

The final reconstructed PET image is shown in Figure 3. In addition, the standard PET reconstruction from the PET-MR scanner is also shown for comparison purposes. Post-filtering after image reconstruction was different for the images.

4.2. MR

Cardiac function – the contracting motion of the heart during the cardiac cycle – is an important clinical parameter when checking for potential problems in the heart. MR is often used for cardiac functional imaging (so-called cine imaging) due to its excellent soft tissue contrast. In order to minimise respiratory motion artefacts, 2D cine imaging is carried out during breath-hold. MR raw data are acquired in k-space over multiple cardiac cycles, and data from the same cardiac phase are retrospectively combined for image reconstruction. This yields dynamic 2D images showing the heart at different phases of the cardiac cycle. Parallel imaging techniques, such as GRAPPA [40], undersample data to speed up the acquisition and reduce breath-hold duration while maintaining high spatial and temporal resolution [10].

Here, we show how to reconstruct a 2D cine MR scan acquired with a 3T mMR Biograph during a 10 s breath-hold (repetition/echo time: 3.4 / 1.5 ms, bSSFP-sequence, flip angle: 50° , image resolution: 1.3 mm, slice thickness: 6 mm, undersampling factor: 2) using the Python interface of SIRF.

As a first step, the filename of the MR raw data file in ISMRMRD format [41] is passed to the MR AcquisitionData class to create the corresponding object:

```
import sirf.Gadgetron as MR
acq_data = MR.AcquisitionData(mr_file)
```


Preprocessing of the data, such as removal of readout oversampling and noise decorrelation for multi-coil data, is carried out:

```
preproc_data = MR.preprocess_acquisition_data(acq_data)
```

The undersampling of the acquired k-space leads to a violation of the Nyquist-Shannon sampling theorem and hence to “undersampling/aliasing” artefacts in the reconstructed image. GRAPPA utilises spatial information from the local receiver coils to minimise these artefacts [42]. A full GRAPPA reconstruction is implemented in Gadgetron and can be called in a straightforward manner from the Python or Matlab interface of SIRF by creating a CartesianGRAPPAReconstructor object:

```
recon = MR.CartesianGRAPPAReconstructor()
```

or by specifying the gadgets explicitly



25 Figure 4: Reconstructed time frames of a MR cine acquisition using SIRF compared to
26 the standard image reconstruction of the 3T Siemens mMR Biograph.
27


```
28
29 recon_gadgets = ['AcquisitionAccumulateTriggerGadget',
30                  'BucketToBufferGadget',
31                  'GenericReconCartesianReferencePrepGadget',
32                  'GenericReconCartesianGrappaGadget',
33                  'GenericReconFieldOfViewAdjustmentGadget',
34                  'GenericReconImageArrayScalingGadget',
35                  'ImageArraySplitGadget'
36                  ]
37
38 ]
```

```
39
40 recon = MR.Reconstructor(recon_gadgets)
41
```

42 Once the preprocessed data are passed to this object, the reconstruction
43 is carried out and the reconstructed image data is retrieved:
44

```
45
46 recon.set_input(preproc_data)
47 recon.process()
48 image_data = recon.get_output()
49
50
```

51 Figure 4 shows the cine data reconstructed with SIRF compared to the
52 standard reconstruction from the MR scanner. In contrast to the standard
53 image reconstruction performed on the PET-MR scanner, SIRF also pro-
54 vides further information, such as a g-factor map that describes the noise
55
56
57



18 Figure 5: G-factor maps for the cine time frames shown in Figure 4. These maps provide
19 a measure of the noise amplification related to the conditioning of the inverse problem, in
20 this case a GRAPPA reconstruction.
21

22
23 amplification that results from reconstructing undersampled data. The g-
24 factor depends upon the coil geometry, coil positioning, undersampling factor,
25 undersampling pattern and the object. (Figure 5). An inhomogeneous
26 distribution of values in the g-factor map can be, for example, an indicator
27 that the sequence parameters were not optimised for the used receiver coil
28 configuration. The calculation of the g-factor map can be enabled for the
29 GRAPPA gadget using:
30
31

```
32  
33 GenericReconCartesianGrappaGadget(send_out_gfactor=true)
```

34 and the g-factor maps are then retrieved using

```
35  
36  
37 gfactor_data = recon.get_output('gfactor')
```

38
39
40 The final number of dynamic 2D cine images depends on the subject's
41 heart rate. In order to make the images more easily comparable between
42 different subjects, the data are commonly interpolated to a predefined num-
43 ber of dynamic images. This can also be done using SIRF by adding the
44 following lines to the reconstruction gadget chain:
45
46

```
47  
48 'PhysioInterpolationGadget(phases=30, mode=0,  
49 first_beat_on_trigger=true, interp_method=BSpline)'
```

50
51 This yields 30 dynamic cine images which can be retrieved using:

```
52  
53  
54 image_data = recon.get_output('Image PhysioInterp')
```

4.3. PET-MR

Using SIRF, it is possible to perform motion-corrected static and dynamic PET reconstructions. This is highlighted in [43] (and briefly shown here), in which the example of head motion is given.

In this method, motion-free time frames are supplied from an external source (such as by using principal component analysis (PCA) on the PET listmode data [44, 45]). Using these time frames, the listmode data are then binned into corresponding sinograms. Non-attenuation corrected (NAC) reconstructions are performed on a frame-by-frame basis, and the resulting images are registered to a reference frame using rigid registration as follows:

```
import sirf.Reg as reg
algo = reg.NiftyAladinSym()
algo.set_parameter("SetPerformRigid","1")
algo.set_parameter("SetPerformAffine","0")
algo.set_floating_image(NAC_i)
algo.set_reference_image(NAC_ref)
algo.process()
TM = algo.get_transformation_matrix_forward()
```

Frame-by-frame attenuation-corrected (AC) reconstructions can then be performed by first transforming the MRAC into the space of the individual PET frames using the extracted transformation matrix, as such:

```
res = reg.NiftyResample()
res.set_floating_image(MRAC)
res.set_reference_image(NAC_i)
res.add_transformation(TM)
res.set_interpolation_type_to_linear()
res.process()
res.get_output().write("MRAC_def.nii")
```

Following the AC reconstruction of each of the frames, they can then be resampled back into the reference space, and indirect parametric estimation can be performed via STIR's Patlak analysis [46].

Figure 6 shows the movement during the scan by comparing the reference frame (left) to another frame earlier in the scan (middle). The earlier scan is then registered to the reference frame to account for motion (right).




Figure 6: NAC reconstructions of the reference frame (left) and an earlier frame (middle), where movement occurred between the two. Finally, the early frame registered to the reference frame (right).

A well-known problem with the above methodology is that the registration of early NAC frames is often not reliable as these images are usually too noisy and the activity distribution is changing rapidly over time. In such circumstances, PET-MR is particularly useful if MR acquisitions of high temporal resolution, such as arterial spin labelling (ASL), were acquired during these early PET frames.

Figure 7 then shows the movement throughout the whole scan, as detected by frame-by-frame registration of NAC and ASL images.

Rapid movement according to the ASL registrations (at 900 s, for example) corresponds to gaps in the NAC data, suggesting that PCA was successfully used to determine motion-free time frames for the PET reconstructions.

The NAC registrations appear to imply that rapid jittering motion is present at the beginning of the scan. This is likely a demonstration of the difficulty of registering early NAC images. This highlights the advantages of using complimentary MR information when available for PET-MR acquisitions.

Figure 8 shows indirect parametric reconstructions using Patlak analysis. On the left, motion correction was used (with motion information from a combination of NAC and ASL registrations), whereas no motion correction was used on the right-hand side. The superior section of the brain is underestimated when motion correction is not used as motion is in the superior-inferior direction (corresponding to the z-direction in Figure 7). This, again, shows the benefit of using complimentary PET-MR motion information.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

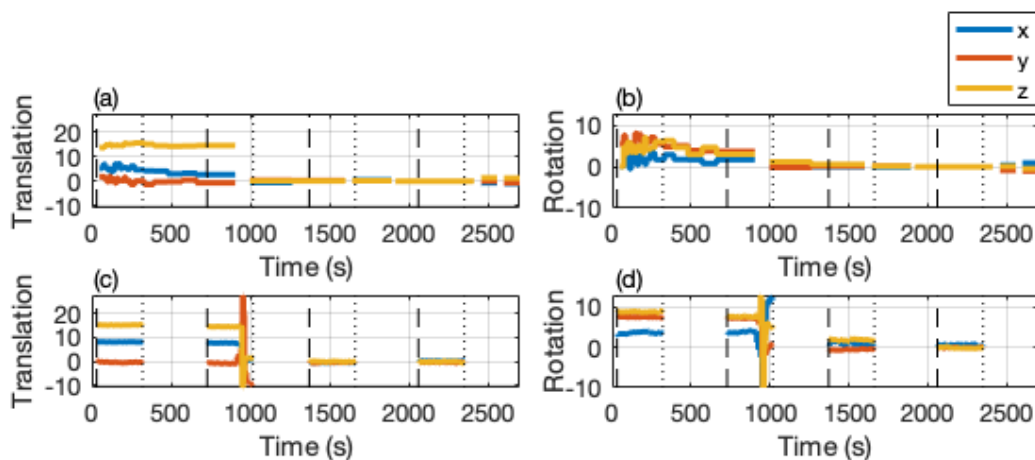


Figure 7: Motion based on frame-by-frame registrations. The top and bottom rows represent motion from registering NAC and ASL images, respectively. Further, the left-hand column shows translations (mm), whereas the right-hand column shows rotational motion via Euler angles (degrees). Gaps in the NAC data are present because motion-free frames have been determined (in this example with PCA) and periods of motion have been removed. Gaps in ASL data are present because ASL images were not acquired for the whole duration of the scan. Vertical lines show the start and end of ASL acquisitions (longer and shorter dashes for the start and end, respectively).




Figure 8: Indirect parametric reconstructions both (a) with and (b) without motion correction. Without motion correction, the top of the brain is underestimated due to motion in the inferior-superior direction.

5. Advanced usage

5.1. Implementing custom reconstruction algorithms

By using the functionality mentioned in this paper, it is possible to implement reconstruction algorithms directly in MATLAB or Python, for instance using `AcquisitionModel` operations, or in terms of the objective function and its `get_subset_gradient` method (see, for instance, the steepest ascent example).

In addition, it is often also convenient to use parts of an existing algorithm. Therefore, SIRF iterative reconstructors expose functionality such as:

- `update` and `update_current_estimate` to compute one update
- `get_current_estimate` and `set_current_estimate` to modify the estimate, for instance with filtering.

During a recent Hackathon [47], such functionality was used to implement new algorithms in Python:

- Total Variation regularised least squares with fast iterative shrinkage-thresholding algorithm (FISTA) [48], and
- De Pierro's maximum *a posteriori* expectation maximisation (MAP-EM) algorithm [49].

5.1.1. Use of CCPi optimisation and regularisation algorithms

The Collaborative Computational Project in Tomographic Imaging (CCPi, <https://www.ccp.i.ac.uk/>), provides a collection of software modules to tackle different aspects of the X-ray CT data analysis pipeline from pre-processing to reconstruction to visualisation; this collection is referred to as the Core Imaging Library (CIL), [50]. The relevant parts of CIL employed to integrate into SIRF are the following:

- CCPi-Framework: An object-oriented optimisation framework ¹

¹<https://github.com/vais-ral/CCPi-Framework/>

- CCPi-Regularisation-Toolkit: A set of CPU/GPU optimised regularisation modules for iterative image reconstruction and other image processing tasks [51] ²
- CCPi-FrameworkPlugins: A set of wrappers to allow usage of CCPi-Regularisation-Toolkit corresponding modules in the CCPi-Framework ³

The CCPi-Framework is developed in Python following an object-orientated approach and most of the classes follow the same naming convention as in SIRF, including `DataContainer`, `ImageData` and `AcquisitionData`. The CCPi-Framework was developed on the basis of the discrete imaging model $A(x) = b$, as discussed in §3.2. While in SIRF the operator A is called ‘AcquisitionModel’, in the CCPi-Framework it is called ‘Operator’ as it was decided to use the names of the underlying mathematical objects. The CCPi-Framework provides a range of generic implementations of optimisation algorithms to address a variety of smooth and non-smooth optimisation problems relevant for reconstruction purposes. One can seamlessly use SIRF objects in CCPi-Framework algorithms thanks to the dynamic/duck typing of Python and to the introduction of aliases of methods that are in CCPi-Framework objects to methods in SIRF objects. For example, the `forward` method of an `AcquisitionModel` is aliased as `direct`, and the `backward` method as `adjoint`.

The CCPi-Regularisation-Toolkit provides a range of different regularisation modules including Total Variation (TV) [52] and variants thereof, implemented in C/CUDA for accelerated computation on multi-core CPUs and GPUs. Through the CCPi-FrameworkPlugins wrappers, these tools are available in the CCPi-Framework and as a result can also be used from SIRF. In the remainder of the paper we will simply refer to CIL as meaning the collection of modules described above.

With the integration of CIL into SIRF, a whole plethora of optimisation algorithms and regularisers are now available. To demonstrate these, the Gradient Descent (GD) and FISTA algorithms [48] were used to solve the least squares and a regularised least-squared problem, respectively.

FISTA can be used to minimise objective functions which are the sum of a differentiable first term and a second term which is not necessarily differ-

²<https://github.com/vais-ral/CCPi-Regularisation-Toolkit>

³<https://github.com/vais-ral/CCPi-FrameworkPlugins>

1
2
3
4
5
6
7
8
9 entiable. The algorithm alternates between the two terms of the objective
10 function allowing for the use non-differentiable regularisers, such as TV, to
11 encourage preservation of edges in the reconstruction. The GD and FISTA
12 algorithms of CIL are provided by the CCPi-Framework, whereas the TV
13 regulariser and in particular its proximal operator is provided by the CCPi-
14 Regularisation-Toolkit by implementing the Fast Gradient Projection algo-
15 rithm [53].
16
17

18 Figure 9 shows the results of the reconstruction of a simulation of the
19 Shepp-Logan phantom using CIL algorithms. First, we used SIRF to re-
20 construct Shepp-Logan phantom x_{SL} from the acquisition data generated by
21 the ISMRMRD utility `ismrmrd.generate_cartesian_shepp_logan` and then
22 used SIRF AcquisitionModel class to generate simulated acquisition data
23 $b = A(x_{SL})$ (in the notation of (1)). Then, we applied the GD algorithm
24 to the least squares problem (minimisation of $\|A(x) - b\|_2^2$) and the FISTA
25 algorithm to the TV-regularised least-squares problem, with the TV term op-
26 erating only on the real part of the solution x . The results for Figure 9a and
27 9b were obtained for 100 and 50 iterations, respectively. The edge-preserving
28 effect of the TV-regulariser is clearly seen in Figure 9b. These examples are
29 provided as proof of concept of the versatile reconstruction options enabled
30 by the integration of CIL into SIRF. Future work will expand the set of al-
31 gorithms from CIL available in SIRF and test their application to real data.
32 The code for this example can be found in the SIRF-Contribs repository.
33
34
35
36
37

38 5.1.2. De Pierro's MAP-EM

39 The goal here was to incorporate new MAP-EM algorithms that could be
40 used both independently and as a part of synergistic PET-MR reconstruction
41 algorithms and to demonstrate the prototyping potential of SIRF.
42

43 The code was implemented using SIRF's pre-existing example Python
44 code for the one-step-late isotropically-weighted quadratically-penalised MAP-
45 EM algorithm as a starting point. The code was then tested for two cases:
46
47

- 48 1. Uniform (i.e., isotropic) weights; and
- 49 2. Bowsher weights [54], calculated using a prior class also introduced
50 during the Hackathon.
51

52 Figure 10 shows the use of the newly implemented algorithm. In Fig-
53 ure 10b a standard OSEM reconstructed image is displayed next to images re-
54 constructed using the uniform (Figure 10c) and Bowsher (Figure 10d) weights
55 of the De Pierro's algorithm. In Figure 10a, the MPRAGE that was used
56
57
58




Figure 9: Reconstructions of an MR simulation of the Shepp-Logan phantom using newly implemented algorithms that are now available in SIRF thanks to the integration of CIL. The results correspond to (a) the gradient descent (after 100 iterations) and (b) Total Variation regularised reconstruction using FISTA (after 50 iterations).

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65




Figure 10: Reconstructions using the recently-implemented De Pierro's MAP-EM algorithm. (a) The anatomical MR image (MPRAGE), (b) a standard OSEM reconstruction for comparison, (c) a reconstruction using De Pierro's MAP-EM with uniform weights, (d) a reconstruction using De Pierro's MAP-EM with Bowsher weights (using the MPRAGE).

1
2
3
4
5
6
7
8
9 to calculate the Bowsher weights is also shown. From these images, it can
10 be seen that the weight for both Figures 10b and 10c is substantially strong;
11 the objective of this task, however, was not to fine-tune the parameters, but
12 simply to demonstrate the images reconstructed by this algorithm. The code
13 for this implementation can be found in our SIRF-Contribs repository.
14
15

16 **6. SIRF software management and deployment routes**

17 *6.1. Software distribution via GitHub*

18
19 We use GitHub with its issue tracking, milestones and release manage-
20 ment facilities (visit <https://github.com/CCPPETMR/SIRF>). SIRF contains
21 a series of examples that also serve as manual tests. In addition, SIRF has
22 a test-suite for automatic testing using CTest. We also employ the Travis
23 Continuous Integration system in which every GitHub code change (commit)
24 is compiled and tested automatically. The build/test status is reported and
25 can be checked online at any time (see <https://travis-ci.org/CCPPETMR>).
26 Furthermore, we maintain a separate repository *SIRF-Exercises* which con-
27 tains Python examples using Jupyter notebooks, <https://jupyter.org>.
28 These examples have been used as the basis for hands-on training sessions.
29
30
31
32
33

34 *6.2. SuperBuild installation*

35
36 The functionality of SIRF relies upon a multitude of open-source software
37 packages. In addition to the reconstruction packages (currently STIR and
38 Gadgetron), SIRF requires a number of general-purpose C++ libraries such
39 as Boost, HDF5 etc. All these SIRF dependencies need to be built using
40 appropriate versions and configurations, and in turn have their own depen-
41 dencies. To reduce the burden on the user of providing all the associated
42 packages, *SIRF-SuperBuild* has been developed and it is the recommended
43 route if building SIRF from source.
44
45

46 The SIRF-SuperBuild is a CMake project which downloads, configures,
47 builds and installs all dependencies required to build SIRF. Once these pack-
48 ages are installed, the SuperBuild will build SIRF. This mechanism works
49 for Linux and macOS, with Windows support pending. By default, the lat-
50 est release of SIRF will be installed. The user can select whether to build
51 the Python and/or MATLAB bindings using CMake. For developers and ad-
52 vanced users, it is possible to configure the SuperBuild to reuse packages that
53 already exist on the system – skipping the building of particular dependencies
54 – and also to select development package versions.
55
56
57
58

1
2
3
4
5
6
7
8
9 Changes to the SuperBuild are also tested on Travis. This includes auto-
10 matic testing of dependencies.
11

12 6.3. Docker 13

14 Docker (available from <https://docs.docker.com/install>) is a container-
15 based low-overhead alternative to Virtual Machines (§6.4). The command
16 `docker pull ccppetmr/sirf` will download the latest ready-made SIRF
17 docker image. The SuperBuild repository (§6.2) contains the docker source
18 code and documentation. Docker images are automatically built on Travis.
19 Various different SIRF images (corresponding to different versions and soft-
20 ware) are available and listed at [https://hub.docker.com/r/ccppetmr/
21 sirf](https://hub.docker.com/r/ccppetmr/sirf).
22
23

24 The docker images can be used to run Gadgetron as a server (for in-
25 stance on Windows where building Gadgetron is difficult) and/or a Jupyter
26 notebook server such that a user can run SIRF from a web browser running
27 on their local machine without further configuration. It is also possible to
28 run Python scripts via docker without Jupyter, although in this case display
29 of any figures needs X forwarding or remote desktop set-up. Docker images
30 with GPU support are also provided for systems which support nvidia-docker
31 (<https://github.com/NVIDIA/nvidia-docker>).
32
33
34

35 6.4. Virtual machine 36

37 A binary installation of SIRF is provided by means of a VirtualBox vir-
38 tual machine (VM), running Ubuntu 18.04 LTS, which can be downloaded
39 from the project's website <http://www.ccppetmr.ac.uk>. The VM contains
40 a (minimal) Linux system complete with developer tools, X windowing sys-
41 tem and Jupyter notebook server.
42
43

44 The `CCPPETMR_VM` GitHub project contains configurations and scripts
45 to create a VM and build SIRF using the SIRF-SuperBuild project. A
46 new VM is created by means of the Vagrant application [https://www.
47 vagrantup.com](https://www.vagrantup.com), which takes as the input a text file that describes the type
48 of machine and how to configure and provision it. A simple bash script, run
49 by Vagrant, configures the SuperBuild and builds SIRF on the VM.
50

51 The VM can be used as a fully-fledged Linux machine. Alternatively,
52 it can be used to run the Gadgetron and/or Jupyter notebook server like
53 the Docker image (for instance on Windows versions, where Docker needs
54 VirtualBox).
55
56
57
58

6.5. Microsoft Azure cloud

SIRF can also be deployed on Azure cloud infrastructure. We use an application called Terraform, see <https://www.terraform.io>, to create a virtual machine running SIRF on Azure. Terraform handles the interaction with Azure to provision a virtual machine in the cloud. SIRF is then built using the scripting mechanism for virtual machine configuration as described in §6.4. In addition to SIRF, the Azure virtual machine automatically runs a Jupyter notebook server which can be accessed through any browser worldwide. Remote Desktop, to access the machine as if it was a local desktop, has also been enabled. Code to deploy SIRF virtual machines on cloud infrastructure besides Azure will be published on GitHub in the future.

7. Discussion

In this paper, we presented release 2.1.0 of SIRF, open-source software for PET-MR image reconstruction. SIRF is designed to provide a simple and object-oriented framework to enable implementation of novel algorithms exploring joint processing of dual-modality imaging data. Data access and manipulation is consistent between the two modalities. SIRF has already been used in several teaching courses as well as in hackathon events [47], showing its potential for rapid code development that can then be used for testing in patient data. SIRF relies on other open-source software packages (“engines”) for many of the underlying operations. This has the benefit of capitalising on the active development of specialist tools. On the other hand, it inherits the capabilities and limitations of these engines. For instance, while support for PET and MR data of the Siemens mMR Biograph [25] is mostly complete, support for the GE Signa PET/MR [27] has not been made available, yet. Recently, there has been good progress regarding the PET data of this scanner model in STIR [55, 56], and the convertor from the GE native MR data format to ISMRMRD is under active development [57]. The architecture of SIRF also allows adding other “engines” in the future with initial discussions started related to CASToR [20]. To cope with version changes in the employed engine packages, SIRF releases link to specific versions. At the same time, we provide advanced users with the opportunity to try the latest versions of all used software, in which case we, of course, cannot always guarantee that everything will work.

The current release of SIRF provides the basis for the implementation of existing and novel algorithms for exploiting the synergy in multi-modality

1
2
3
4
5
6
7
8
9 imaging data with initial work in combining high quality side information
10 [58, 59, 60, 61] and joint simultaneous reconstruction of PET and MRI images
11 [15, 19]. Making these algorithms freely available to the wider community
12 of researchers will be an important step towards assessing their benefits for
13 clinical application and translation.
14

15 In addition to PET-MR, there is active research in exploiting synergy
16 in other modalities, where multiple parameters of one underlying object can
17 be extracted from multiple acquired data-sets, such as multi-sequence MR
18 [62, 16, 63, 64], dual-energy/spectral CT [65, 66, 67, 68], and even CT-
19 MR [69, 70]. Many of the algorithms developed in PET-MR research are
20 transferable to those research areas and vice versa. The architecture of SIRF
21 would enable incorporating these modalities as well as including PET/CT,
22 SPECT/CT and SPECT/MR in a common framework in the future.
23

24 Another focus for the development of SIRF in the future will be around
25 deployment options. Ensuring SIRF is readily available and simple to use by
26 a non-technical researcher will allow SIRF to be used in higher-level medi-
27 cal research. The use of SIRF in such research potentially enables superior
28 quantitative accuracy in medical investigations, and the ability to demon-
29 strate the utility of the implemented algorithms in a clinical research context
30 before being demonstrated in routine clinical practice. Existing deployment
31 options were outlined in §6, but in the future may be augmented by deploy-
32 ing Gadgetron’s inline facilities, or as a hosted reconstruction service. A
33 production deployment of SIRF is a possibility, which is however not within
34 the grasp of our team, mainly because of regulatory issues. This is less of an
35 issue for pre-clinical systems.
36
37
38
39
40

41 42 **8. Conclusion**

43
44 The open-source Synergistic Image Reconstruction Framework (SIRF) is
45 a new software tool for research in image reconstruction and related manip-
46 ulations including motion estimation and compensation. It currently targets
47 joint image reconstruction of PET-MR data but with a view towards exten-
48 sion to other cases where synergistic reconstruction can be beneficial. SIRF
49 enables researchers to explore the benefits of synergistic reconstruction in a
50 framework that has been successfully used for educational purposes whilst
51 being powerful enough for processing data from state-of-the-art scanners.
52
53
54
55
56
57
58

Acknowledgements

CCP PET-MR is funded by the EPSRC (grant EP/M022587/1) and the associated Software Flagship project “A framework for efficient synergistic spatiotemporal reconstruction of PET-MR dynamic data” (grant EP/P022200/1). This work made use of computational support by CoSeC, the Computational Science Centre for Research Communities. We thank Siemens Healthineers and General Electric Healthcare for information and support with file formats. We also wish to thank Michael Hansen, David Hansen and Hui Xue for help with Gadgetron, and Camila Munoz, James Bland and Abolfazl Mehranian for helping with the coding during the hackathon leading to §5.1.2, and the wider CCP PET-MR community for feedback. We wish to acknowledge CCPi with the funding support by the EPSRC grants (EP/P02226X/1) ‘Reconstruction Toolkit for Multichannel CT’ and the network initiative (EP/M022498/1). Lastly, we wish to thank St Thomas’ PET Centre for the use of their patient data.

Appendix A. Object-oriented paradigm

This appendix is mainly intended for readers who are not familiar with Object-Oriented programming.

In Object-Oriented approach, instead of having data containers (arrays, files etc.) and functions that operate on them, one employs *objects*, which contain data and come along with sets of functions, called *methods*, which operate on data. In object-oriented languages, an object type is referred to as its *class*. Each object has a special method called constructor, which has the same name as the object class name and must be called to create that object. For example, to create an object of class `ImageData` that handles MR image data and fill it with data stored in the HDF5 file `my_image.h5` one needs to call for an assignment, such as:

```
image = ImageData('my_image.h5')
```

We note that an MR `ImageData` object contains not only the voxel values, but also a number of parameters specified by the ISMRMRD format of MR image data [41]. The object data is encapsulated, i.e. it is not directly accessible from the user’s code (being handled mostly by the underpinning C++ code) and is processed by the object methods. For example, to display the data encapsulated by `image`, one needs to call its method `show()`:

1
2
3
4
5
6
7
8
9 `image.show()`

10
11 and if anyone needs to copy the data into a MATLAB array, they would use
12 the method `as_array()`:
13

14
15 `image_data_array = image.as_array()`
16

17
18 Parameters of objects are modified/accessed via `set/get` methods (muta-
19 tors and accessors). For example, the value of an objective function handled
20 by the object named `obj_fun` on an image data object is computed by its
21 method `get_value()` as:
22

23
24 `obj_fun_value = obj_fun.get_value(image)`
25

26 The mutators are also responsible for basic error checking.
27

28 Some classes are derived from other classes, which means that they con-
29 tain all the methods of the classes they are derived from (i.e. *inherited*
30 methods). For example, the class `AcquisitionModelUsingMatrix` is derived
31 from the class `AcquisitionModel`, so an object of the former class inherits
32 all methods of the latter.
33
34

35 Bibliography

- 36
37 [1] K. Thielemans, C. Tsoumpas, S. Mustafovic, T. Beisel, P. Aguiar,
38 N. Dikaios, M. W. Jacobson, *Physics in Medicine and Biology* 57 (4)
39 (2012) 867–883.
40
41 [2] M. S. Hansen, T. S. Sørensen, *Magnetic Resonance in Medicine* 69 (6)
42 (2013) 1768–1776.
43
44 [3] H. F. Wehrl, A. W. Sauter, M. R. Divine, B. J. Pichler, *Journal of*
45 *Nuclear Medicine* 56 (2) (2015) 165–168.
46
47 [4] Z. Chen, S. D. Jamadar, S. Li, F. Sforazzini, J. Baran, N. Ferris, N. J.
48 Shah, G. F. Egan, *Human Brain Mapping* 39 (12) (2018) 5126–5144.
49
50 [5] F. Nensa, K. Beiderwellen, P. Heusch, A. Wetter, *Diagnostic and Inter-*
51 *ventional Radiology (Ankara, Turkey)* 20 (5) (2014) 438–447.
52
53
54
55
56
57
58

- 1
2
3
4
5
6
7
8
9 [6] P. M. Robson, D. Dey, D. E. Newby, D. Berman, D. Li, Z. A. Fayad,
10 M. R. Dweck, *JACC: Cardiovascular Imaging* 10 (10 Part A) (2017)
11 1165–1179.
12
13 [7] H. Barthel, O. Sabri, in: L. Umutlu, K. Herrmann (Eds.), *PET/MR*
14 *Imaging: Current and Emerging Applications*, Springer International
15 Publishing, Cham, 2018, pp. 99–106.
16
17 [8] X. Y. Zhang, Z. L. Yang, G. M. Lu, G. F. Yang, L. J. Zhang, *Frontiers*
18 *in Molecular Neuroscience* 10 (2017).
19
20 [9] J. Qi, R. M. Leahy, *Physics in Medicine and Biology* 51 (15) (2006)
21 R541–R578.
22
23 [10] K. P. Pruessmann, *NMR in biomedicine* 19 (3) (2006) 288–299.
24
25 [11] J. A. Fessler, *IEEE Signal Processing Magazine* 27 (4) (2010) 81–89.
26
27 [12] K. G. Hollingsworth, *Physics in Medicine and Biology* 60 (21) (2015)
28 R297–R322.
29
30 [13] B. Bai, Q. Li, R. M. Leahy, *Seminars in Nuclear Medicine* 43 (1) (2013)
31 30–44.
32
33 [14] A. Mehranian, E. De Vita, R. Neji, C. J. McGinnity, A. Hammers, A. J.
34 Reader, in: *PSMR 2018, Isola de Elba, Italy, 2018*, p. 3.
35
36 [15] M. J. Ehrhardt, K. Thielemans, L. Pizarro, D. Atkinson, S. Ourselin,
37 B. F. Hutton, S. R. Arridge, *Inverse Problems* 31 (2015) 015001.
38
39 [16] F. Knoll, M. Holler, T. Koesters, R. Otazo, K. Bredies, D. K. Sodickson,
40 *IEEE Transactions on Medical Imaging* 36 (1) (2017) 1–16.
41
42 [17] J. Rasch, E.-M. Brinkmann, M. Burger, *Inverse Problems* 34 (1) (2017)
43 014001.
44
45 [18] A. Mehranian, M. A. Belzunce, C. Prieto, A. Hammers, A. J. Reader,
46 *IEEE Transactions on Medical Imaging* 37 (1) (2018) 20–34.
47
48 [19] A. Mehranian, M. A. Belzunce, C. J. McGinnity, A. Bustin, C. Prieto,
49 A. Hammers, A. J. Reader, *Magnetic resonance in medicine* 81 (3) (2019)
50 2120–2134.
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

- 1
2
3
4
5
6
7
8
9 [20] T. Merlin, S. Stute, D. Benoit, J. Bert, T. Carlier, C. Comtat, M. Filipovic, F. Lamare, D. Visvikis, *Physics in Medicine & Biology* 63 (18) (2018) 185005.
10
11
12
13 [21] T. Kösters, K. P. Schäfers, F. Wübbeling, in: 2011 IEEE Nuclear Science Symposium Conference Record, 2011, pp. 4365–4368.
14
15
16 [22] S. Pedemonte, K. Van Leemput, N. Fuin, occiput.io.
17 <http://tomographylab.scienceontheweb.net>
18
19 [23] P. J. Markiewicz, M. J. Ehrhardt, K. Erlandsson, P. J. Noonan, A. Barnes, J. M. Schott, D. Atkinson, S. R. Arridge, B. F. Hutton, S. Ourselin, *Neuroinformatics* 16 (1) (2018) 95–115.
20
21
22 [24] M. Uecker, J. Tamir, F. Ong, C. Holme, M. Lustig, BART: version 0.4.01 (Jun. 2017).
23 <https://doi.org/10.5281/zenodo.817472>
24
25 [25] G. Delso, S. Fürst, B. Jakoby, R. Ladebeck, C. Ganter, S. G. Nekolla, M. Schwaiger, S. I. Ziegler, *Journal of Nuclear Medicine* 52 (12) (2011) 1914–1922.
26
27
28 [26] A. M. Grant, T. W. Deller, M. M. Khalighi, S. H. Maramraju, G. Delso, C. S. Levin, *Medical Physics* 43 (5) (2016) 2334–2343.
29
30
31 [27] C. S. Levin, S. H. Maramraju, M. M. Khalighi, T. W. Deller, G. Delso, F. Jansen, *IEEE Transactions on Medical Imaging* 35 (8) (2016) 1907–1914.
32
33
34 [28] Operator Discretization Library.
35 <https://odlgroup.github.io/odl/index.html>
36
37 [29] M. Modat, G. R. Ridgway, Z. A. Taylor, M. Lehmann, J. Barnes, D. J. Hawkes, N. C. Fox, S. Ourselin, *Computer Methods and Programs in Biomedicine* 98 (3) (2010) 278–284.
38
39
40 [30] M. Modat, D. M. Cash, P. Daga, G. P. Winston, J. S. Duncan, S. Ourselin, *Journal of Medical Imaging (Bellingham, Wash.)* 1 (2) (2014) 024003.
41
42
43 [31] Simplified Wrapper and Interface Generator.
44 <http://www.swig.org/index.php>
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

- 1
2
3
4
5
6
7
8
9 [32] NIfTI-2 Data Format – Neuroimaging Informatics Technology Initiative.
10 <https://nifti.nimh.nih.gov/nifti-2>
11
- 12 [33] H. Hudson, R. Larkin, IEEE Transactions on Medical Imaging 13 (4)
13 (1994) 601–609.
14
- 15 [34] J. E. Bowsher, V. E. Johnson, T. G. Turkington, R. J. Jaszczak, C. E.
16 Floyd, R. E. Coleman, Medical Imaging, IEEE Transactions on 15 (5)
17 (1996) 673–686.
18
- 19 [35] M. J. Ehrhardt, P. Markiewicz, M. Liljeroth, A. Barnes,
20 V. Kolehmainen, J. S. Duncan, L. Pizarro, D. Atkinson, B. F.
21 Hutton, S. Ourselin, K. Thielemans, S. R. Arridge, IEEE Transactions
22 on Medical Imaging 35 (9) (2016) 2189–2199.
23
- 24 [36] D. Hogg, K. Thielemans, S. Mustafovic, T. J. Spinks, in: 2002 IEEE
25 Nuclear Science Symposium Conference Record, Vol. 3, IEEE, 2002, pp.
26 1519–1523, event-place: Norfolk, VA, USA.
27
- 28 [37] P. Green, IEEE Transactions on Medical Imaging 9 (1) (1990) 84–93.
29
- 30 [38] C. C. Watson, D. Newport, M. E. Casey, in: P. Grangeat, J. L. Amans
31 (Eds.), Fully Three-Dimensional Image Reconstruction in Radiology and
32 Nuclear Medicine, Kluwer Academic, Amsterdam, Netherlands, 1996,
33 pp. 255–268.
34
- 35 [39] I. Polycarpou, K. Thielemans, R. Manjeshwar, P. Aguiar, P. Marsden,
36 C. Tsoumpas, Annals of Nuclear Medicine 25 (9) (2011) 643–649.
37
- 38 [40] M. A. Griswold, P. M. Jakob, R. M. Heidemann, M. Nittka, V. Jellus,
39 J. Wang, B. Kiefer, A. Haase, Magn Reson Med. 47(6) (2002) 1202 –
40 1210.
41
- 42 [41] S. J. Inati, J. D. Naegele, N. R. Zwart, V. Roopchansingh, M. J. Lizak,
43 D. C. Hansen, C.-Y. Liu, D. Atkinson, P. Kellman, S. Kozerke, H. Xue,
44 A. E. Campbell-Washburn, T. S. Sørensen, M. S. Hansen, Magn. Reson.
45 Med. 77 (1) (2017) 411–421.
46
- 47 [42] M. A. Griswold, P. M. Jakob, R. M. Heidemann, M. Nittka, V. Jellus,
48 J. Wang, B. Kiefer, A. Haase, Magnetic Resonance in Medicine 47 (6)
49 (2002) 1202–1210.
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

- 1
2
3
4
5
6
7
8
9 [43] R. Brown, B. A. Thomas, A. Rashidnasab, K. Erlandsson, E. Ovtchinnikov, E. Pasca, A. J. Reader, J. C. Matthews, C. Tsoumpas, K. Thielemans, in: 2018 IEEE Nuclear Science Symposium and Medical Imaging Conference (NSS/MIC), IEEE, Sydney, Australia, 2018.
- 10
11
12
13
14
15 [44] K. Thielemans, P. Schleyer, J. Dunn, P. K. Marsden, R. M. Manjeshwar, in: 2013 IEEE Nuclear Science Symposium and Medical Imaging Conference (2013 NSS/MIC), 2013, pp. 1–5.
- 16
17
18
19 [45] P. J. Schleyer, J. T. Dunn, S. Reeves, S. Brownings, P. K. Marsden, K. Thielemans, *Physics in Medicine and Biology* 60 (16) (2015) 6441–6458.
- 20
21
22
23
24 [46] C. Tsoumpas, K. Thielemans, *Nuclear Medicine Communications* 30 (7) (2009) 490–493.
- 25
26
27 [47] R. Brown, S. Ellis, E. Pasca, E. Ovtchinnikov, A. Gillman, C. Munoz, J. Bland, A. Mehranian, C. Prieto, A. J. Reader, K. Thielemans, in: Abstracts of PSMR2019, Muenich, Germany, 2019.
- 28
29
30
31 [48] A. Beck, M. Teboulle, *SIAM Journal on Imaging Sciences* 2 (1) (2009) 183–202.
- 32
33
34 [49] A. R. De Pierro, *IEEE Transactions on Medical Imaging* 14 (1) (1995) 132–137.
- 35
36
37 [50] CCPi, Ccpi core imaging library (june 2019).
<https://www.ccp.i.ac.uk/CIL>
- 38
39
40
41 [51] D. Kazantsev, E. Pasca, M. J. Turner, P. J. Withers, *SoftwareX* 9 (2019) 317 – 323.
- 42
43
44 [52] L. I. Rudin, S. Osher, E. Fatemi, *Physica D: Nonlinear Phenomena* 60 (1) (1992) 259 – 268.
- 45
46
47 [53] A. Beck, M. Teboulle, *IEEE Transactions on Image Processing* 18 (11) (2009) 2419–2434.
- 48
49
50
51 [54] J. E. Bowsher, Hong Yuan, L. W. Hedlund, T. G. Turkington, G. Akabani, A. Badea, W. C. Kurylo, C. T. Wheeler, G. P. Cofer, M. W. Dewhurst, G. A. Johnson, in: *IEEE Symposium Conference Record Nuclear Science 2004.*, Vol. 4, 2004, pp. 2488–2492 Vol. 4.
- 52
53
54
55
56
57
58

- 1
2
3
4
5
6
7
8
9 [55] N. Efthimiou, E. Emond, P. Wadhwa, C. Cawthorne, C. Tsoumpas,
10 K. Thielemans, *Physics in Medicine and Biology* 64 (3) (2019) 035004.
11
12 [56] P. Wadhwa, K. Thielemans, N. Efthimiou, K. Wangerin, D. Timothy,
13 G. Delso, K. Nicholas, D. Daniel, E. Emond, M. Tohme, F. Janssen,
14 R. N. Gunn, D. L. Buckley, W. A. Hallett, C. Tsoumpas, *Methods* (in-
15 press).
16
17 [57] V. Roopchansingh, J. A. Derbyshire, GE to ISMRMRD converter,
18 original-date: 2016-02-16T14:34:23Z (May 2019).
19 https://github.com/ismrmrd/ge_to_ismrmrd
20
21 [58] G. Wang, J. Qi, *IEEE Transactions on Medical Imaging* 34 (1) (2015)
22 61–71.
23
24 [59] J. Bland, M. A. Belzunce, S. Ellis, C. J. McGinnity, A. Hammers, A. J.
25 Reader, *IEEE Transactions on Radiation and Plasma Medical Sciences*
26 2 (5) (2018) 470–482.
27
28 [60] D. Deidda, N. A. Karakatsanis, P. M. Robson, Y.-J. Tsai, N. Efthimiou,
29 K. Thielemans, Z. A. Fayad, R. G. Aykroyd, C. Tsoumpas, *Inverse*
30 *Problems* 35 (4) (2019) 044001.
31
32 [61] D. Deidda, N. A. Karakatsanis, P. M. Robson, N. Efthimiou, Z. A.
33 Fayad, R. G. Aykroyd, C. Tsoumpas, *IEEE Transactions on Radiation*
34 *and Plasma Medical Sciences* 3 (4) (2019) 400–409.
35
36 [62] M. J. Ehrhardt, M. M. Betcke, *SIAM Journal on Imaging Sciences* 9 (3)
37 (2016) 1084–1106.
38
39 [63] B. Bilgic, T. H. Kim, C. Liao, M. K. Manhard, L. L. Wald, J. P. Haldar,
40 K. Setsompop, *Magnetic resonance in medicine* 80 (2) (2018) 619–632.
41
42 [64] A. Mehranian, C. Prieto, R. Neji, C. J. McGinnity, A. Hammers, A. J.
43 Reader, in: *ISMRM 2018, Paris, France, 2018*.
44
45 [65] C. H. McCollough, S. Leng, L. Yu, J. G. Fletcher, *Radiology* 276 (3)
46 (2015) 637–653.
47
48 [66] I. Danad, Z. A. Fayad, M. J. Willemink, J. K. Min, *JACC: Cardiovas-*
49 *cular Imaging* 8 (6) (2015) 710–723.
50
51
52
53
54
55
56
57
58

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65

[67] A. C. Silva, B. G. Morse, A. K. Hara, R. G. Paden, N. Hongo, W. Pavlicek, *RadioGraphics* 31 (4) (2011) 1031–1046.

[68] B. Chen, Z. Zhang, E. Y. Sidky, D. Xia, X. Pan, *Physics in Medicine and Biology* 62 (22) (2017) 8763–8793.

[69] G. Wang, J. Zhang, H. Gao, V. Weir, H. Yu, W. Cong, X. Xu, H. Shen, J. Bennett, M. Furth, Y. Wang, M. Vannier, *PLoS ONE* 7 (6) (2012) e39700.

[70] Y. Xi, J. Zhao, J. R. Bennett, M. R. Stacy, A. J. Sinusas, G. Wang, *IEEE Transactions on Biomedical Engineering* 63 (6) (2016) 1301–1309.