Artificial Intelligence and Augmented Intelligence for Automated Investigations for Scientific Discovery

Artificial intelligence for reconstruction and super-resolution of chemical tomography
Final Report
Project Dates: 01-05-2020 - 30-12-2020
Science and Technology Facilities Council

Keith Tobias Butler
Rutherford Appleton Laboratory
Simon Jacques
Finden Ltd.

Report Date: 21/04/2021

Artificial intelligence for reconstruction and super-resolution of chemical tomography
AI3SD-Project-Series:Report4_Butler_Final
Report Date: 21/04/2021
DOI: 10.5258/SOTON/P0037

Principal Investigator: *Professor Jeremy Frey*
Co-Investigator: *Professor Mahesan Niranjan*
Network+ Coordinator: *Dr Samantha Kanza*

Contents

1 Project Details

Title	Artificial intelligence for reconstruction and super-resolution of chemical tomography
Funding reference	AI3SD-FundingCall2_017
Lead Institution	STFC
Project Dates	01-05-2020 - 30-12-2020
Website	https://superres-tomo.readthedocs.io
Keywords	Chemical tomography, machine learning, CNN, diffraction

2 Project Team

2.1 Principal Investigator

Name and Title	Keith Tobias Butler
Employer name / University Department Name	STFC, Scientific Computing Department
Work Email	keith.butler@stfc.ac.uk
Website Link	http://keeeto.github.io

2.2 Co-Investigators

Name and Title	Dr Simon Jacques (Managing Director)
Employer name / University Department Name	Finden Ltd
Work Email	simon@finden.co.uk
Website Link	www.finden.co.uk

Name and Title	Dr Antony Vamvakeros (Research Scientist)
Employer name / University Department Name	Finden Ltd
Work Email	antony@finden.co.uk
Website Link	www.finden.co.uk

2.3 Researchers & Collaborators

Mr. Hongyang Dong, a PhD student at UCL chemistry joined the team during the project. Hongyang joined us in September 2020 and worked on generative adversarial networks for reconstruction. hongyang.dong.18@ucl.ac.uk

3 Publicity Summary

X-ray scatter-based tomography allows unprecedented insight into the chemical and physical state of functional materials and devices. Such tomographies can be used as research tools but also offer the prospect of routine scanning for security and inspection systems and potential for medical scanning.

In conventional X-ray tomography, the images that are obtained give maps of density within the object and the composing pixels contain single grey scale values. In scatter based tomography, each pixel instead contains spectrum or equivalent chemical signal i.e. a 1D array (or higher) of numbers. An X-ray scatter tomography slice becomes a data cube with the two conventional spatial dimensions and a third spectral dimension. Such image data is termed hyperspectral.

Whilst, hyperspectral tomography can match or even exceed the resolution offered by conventional X-ray absorption contrast tomography, the latter is more highly optimised and offers modalities that can generate images in a fraction of the time and dose to image the same volume. In practice it is often the case that hyperspectral tomography resolution is sacrificed to accelerate collection time. This project aims to exploit machine learning approaches to marry hyperspectral chemical tomography with conventional X-ray absorption tomography to achieve chemical images with the rich information of the former in combination with the resolution and speed of collection of the latter.

4 Executive Summary

X-ray based scatter tomography are extremely powerful non destructive analytical techniques that can provide insight into the chemical and physical states of functional materials and devices even under operating conditions. These approaches have potential for applications beyond their current use as research tools, but to date this has not been realised, in large part, due to the long collection times and high dose rates associated with measurement. The aim in this project was to facilitate shorter data collections that can still yield high resolution images by using machine learning approaches to obtain super-resolution. To achieve this, we used the information from the large number of data points within the hyperspectral X-ray scatter dataset and combined with the traditional conventional X-ray absorption signal which can be easily and quickly measured. This project has focused on (1) developing and applying novel AI-based methods for chemical image (volume) reconstruction and (2) enhancing the spatial resolution of the chemical images.

The approaches reported herein, have been applied to the X-ray diffraction computed tomography (XRD-CT) technique. The state-of-the-art with this method yields large 3D volumes, containing many hundreds of thousands or even millions of diffraction patterns. These are extremely challenging and time consuming to processes and analyse. The continuing development in instrumentation means that this big data problem is only increasing and indeed this problem becomes even worse when higher resolution chemical images are obtained.

The new machine learning based algorithms that have been developed will semi-automate the image reconstruction process while the image enhancement allows the collection of less data to achieve the same resolution (both cases falling under the big data handling umbrella). These new reconstruction and image enhancement approaches are immediately beneficial in terms of advancing these techniques and offering the prospect of translation beyond their use as research tools. Of great value should be the baseline model and a benchmark dataset that we have made publicly available, which can stimulate research across the community.

5 Aims and Objectives

We aim to develop generally applicable, freely available tools for using machine learning in X-ray computed tomography. Specifically we will develop:

- Open datasets for testing tomographic reconstruction and super-resolution approaches

- A benchmark baseline for tomographic reconstruction with neural networks

- A benchmark baseline for super-resolution enhancement of XRD-CT images using neural networks

- A documented, open-source repository for these tools and data

Our overarching aim was to accelerate machine learning application in XRD-CT by providing examples, baselines and test datasets.

6 Methodology

6.1 Scientific Methodology

The first part of the project focused on creating training libraries for the neural networks. To avoid bias, we have created libraries based on: (1) synthetic data containing random shapes using the well-known scikit-image python package, (2) the DIV2K dataset: DIVerse 2K resolution high quality images as used for the challenges @ NTIRE (CVPR 2017 and CVPR 2018) and @ PIRM (ECCV 2018) and (3) images reconstructed from previously acquired micro-CT and XRD-CT experimental datasets. For the experimental tomographic datasets, the sinograms were first centered, scaled (i.e. assuming equal summed intensity per tomographic angle), background was subtracted (e.g. air scattering for the XRD-CT data) and then the images were reconstructing with the filtered back projection algorithm setting all negative values to zero. Specifically for the XRD-CT datasets, appropriate filters (i.e. trimmed mean filter) were applied to the raw 2D diffraction images during radial integration to avoid the formation of hotspots in the sinograms. These processed images are considered to be the ground truth for these libraries. Where needed, these images were rescaled to lower resolution using bilinear interpolation and artificial sinograms were created using the astra toolbox in python. The performance of the reconstruction models is evaluated by comparing the reconstruction results of these sinograms with the ones obtained using the filtered back projection (i.e. reconstruction of synthetic sinograms). Similarly, for the super-resolution, the performance of the CNNs will be evaluated using the aforementioned three libraries.

6.2 AI Methodology

We used a mixed architecture for CNN reconstruction. It starts with four 2D convolutional layers whose strides are equal to 2, followed by four fully connected layers. Then the 1D output from the last fully connected layer is transformed back to a 2D image and then sent to the next three 2D convolutional layers whose strides are equal to 1. For fully connected layers, there are 1000 nodes inside each of the first three of them, and the last fully connected layer has the size equal to the number of pixels of the reconstructed image. Besides, there is also one dropout layer after each fully connected layer to avoid overfitting.

The accuracy and quality of the reconstructed images can be improved by increasing the number of nodes in the first three dense layers, but those numbers are significantly restricted by the computing resources. Increasing the number of nodes inside each layer can lead to a dramatic increase of the trainable weights, which makes the model harder to fit. Therefore, we added four more convolutional layers afterwards. The convolutional layers have significantly fewer weights than dense layers. They are used to take out the best-fitting features of the images and refine the reconstruction.

For the super-resolution, we are planning first to implement and evaluate the performance of the EDSR, WDSR and SRGAN neural networks using the aforementioned three libraries containing different types of image data. The next step will involve exploring their performance for upscaling micro-CT and for the first time, XRD-CT data. Finally, new architectures will be explored specifically for the XRD-CT data using a dual input of low resolution XRD-CT images and a high resolution micro-CT image acquired at the same position (i.e. XRD-CT and micro-CT corresponding at the same sample cross-section).

7 Results

a. Libraries
We have made a publicly available set of libraries for training and testing tomographic reconstruction and X-ray image super-resolution techniques. There is one library of 180,000 sinogram image pairs for reconstruction. As described above there are also several libraries of experimental micro-CT and XRD-CT data at different resolutions (for the same data), which can be used to train and test super-resolution methods. The links and details of the libraries can be found on the project documentation pages:
https://superres-tomo.readthedocs.io/en/latest/benchmark_data.html.

b. CNN for reconstruction
For the purpose of this project, we have decided to use the currently latest stable version of Tensorflow (v.2.2.0) in python for the development and testing of the neural networks. We designed, implemented and evaluated a large number of CNN architectures, exploring also the impact of various hyperparameters (e.g. learning rate) and loss functions. A representation of the CNN reconstruction architecture is shown in Figure 1.

Figure 1: A representation of the CNN reconstruction architecture. The green blocks represent convolutional layers, the orange blocks represent fully connected layers.

From the various performance tests, we found the cnn_reconstruct discussed in the 6b section as one of the most promising ones, especially due to its ability for upscaling (i.e. it can handle relatively large images while maintaining a number of parameters in the order of 107 - 108). The learning rate was set to 0.00025 and the root mean squared was used as the loss function. We created a library using a combination of experimental XRD-CT datasets using catalyst particles consisting of 8,000 pairs of sinograms-images. Some examples are shown in Figure 2 where the original image and the ones obtained from the reconstruction using the filtered back projection algorithm and the cnn respectively are presented.


Figure 2: Performance of the reconstruction CNN and its comparison with the results obtained with the filtered back projection algorithm using a library containing XRD-CT sinograms-images.

It can be seen here that the performance of the cnn_reconstruct is superior to the conventional filtered back projection algorithm as it can correctly reconstruct the shape and intensity of the catalyst particles while at the same time suppressing the background noise. However, when

the same CNN was tested using the library containing random shapes (55,000 pairs of sinograms-images), the performance was worse. These results are presented in Figure 3.

It can be seen that the filtered back projection algorithm can retain the sharp edges of the shapes and their overall shape while the reconstruction CNN fails to do so. The problem here arises from the training data as the various shape images can vary significantly in content while there is a high degree of correlation between the XRD-CT images (i.e. the XRD-CT images present in each XRD-CT dataset). These results are very important as they illustrate the strong dependence of the reconstruction CNN on the training data and the difficulty in creating a reconstruction CNN that can handle very different data (i.e. images that are not well correlated with the training data). This major issue is rarely discussed in literature and the current results from this project show that there should be more discussion on the impact and nature of training data used in supervised learning reconstruction CNNs. As an example, in literature one can often encounter CNNs used in medical imaging (reconstruction, denoising etc) claimed to exhibit superior performance compared to conventional methods but the CNNs have been trained using only training libraries containing medical CT for a body part/organ.


Figure 3: Performance of the reconstruction CNN and its comparison with the results obtained with the filtered back projection algorithm using a library containing synthetic sinograms-images of random shapes.

c. CNN transferability

We decided to explore further the transferability of the supervised-learning reconstruction CNNs by examining four X-ray diffraction computed tomography (XRD-CT) training datasets and performing t-SNE analysis of these datasets. Specifically, two experimental and two simulated XRD-CT datasets were used in this work. The cnn_reconstruct was trained using each dataset and was then used to reconstruct the tomographic images. As shown in Figure 4, the performance of the CNN strongly depends on the type of data it has experienced during training.
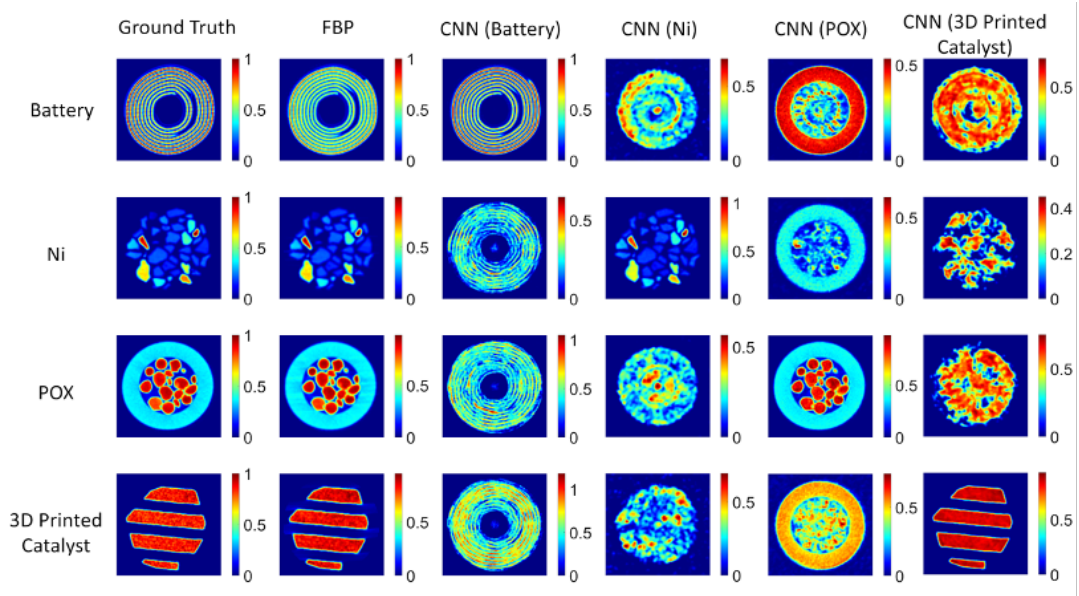
Figure 4: Representative examples of reconstructed images using CNNs trained on different data sets. The columns represent the different training sets used for the CNN, as well as ground truth and FBP results, and the rows represent the different test examples presented to the CNN.

We then combined the various libraries and trained the reconstruction CNN using this combined library. As shown in Figure 5, the CNN is now able to reconstruct accurately the images provided from different sources. We have performed a thorough investigation including t-SNE analysis and we are currently preparing a manuscript for submission to a peer reviewed journal.


Figure 5: Representative examples of reconstructed images using CNNs trained on mixed data sets. The columns represent the different size training sets used for the CNN, as well as ground truth and FBP results, and the rows represent the different test examples presented to the CNN.

d. Generative Adversarial Network for reconstruction

Initially in the proposed work plan we had intended to test the recently developed Mixed Scale Dense Networks (https://doi.org/10.1073/pnas.1715832114) for both image reconstruction and super-resolution. However when we further explored the architecture and its implementation, it has become apparent that the MSDN is restricted to inputs and outputs with exactly the same dimensions. This makes reconstruction impossible and limits its application for reconstruction and super-resolution. As a result we decided to concentrate our efforts on CNNs and Generative Adversarial Networks (GANs). We implemented an unsupervised machine learning approach for image reconstruction using a GAN inspired by the recently published GANrec which eliminates the bias caused by the nature of the training data in supervised learning approaches.

A schematic representation of the GAN is presented in Figure 6. The first step involves a coarse tomographic image reconstruction which is performed external to the CNN using for example the conventional filtered back-projection algorithm. The resultant image is then passed to the Generator CNN of the GAN and the resulting image is forward-projected using the radon transform to yield a sinogram. This simulated sinogram is then passed to the Discriminator CNN which compares it to the experimental sinogram. The Discriminator CNN is trained to understand whether the simulated sinogram is real or fake (i.e. matching the experimental sinogram or not). At the end of every iteration, the weights of both the Discriminator and Generator CNNs are updated. We have worked on optimizing the architecture of both the Generator and Discriminator CNNs, minimizing the number of parameters (i.e. width and depth of the neural networks). The optimised GAN architecture has high accuracy, is scalable (sinogram/image size) and as shown in Figure 7 does not suffer from transferability issues (i.e. can handle sinograms/images from different sources).
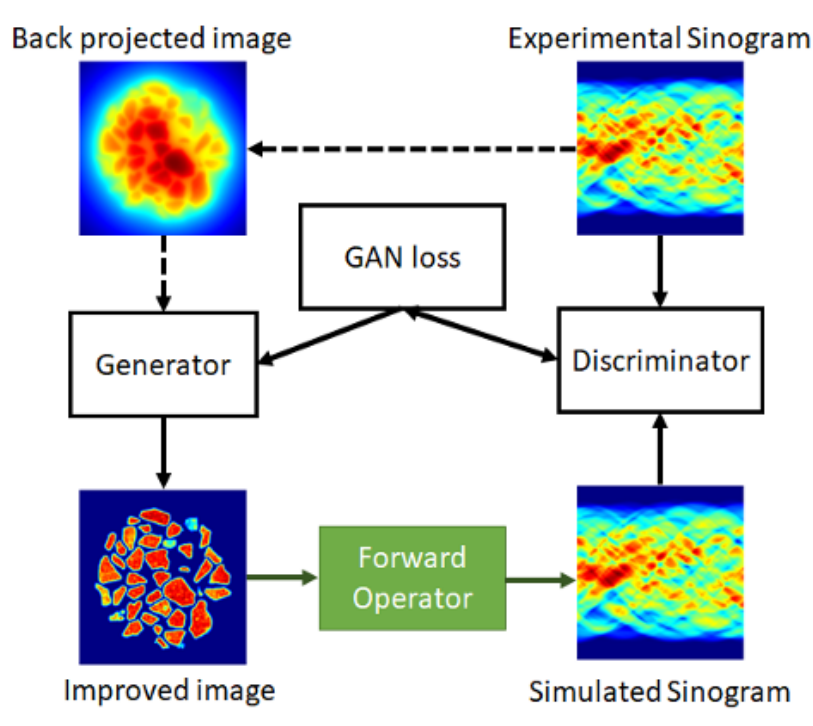

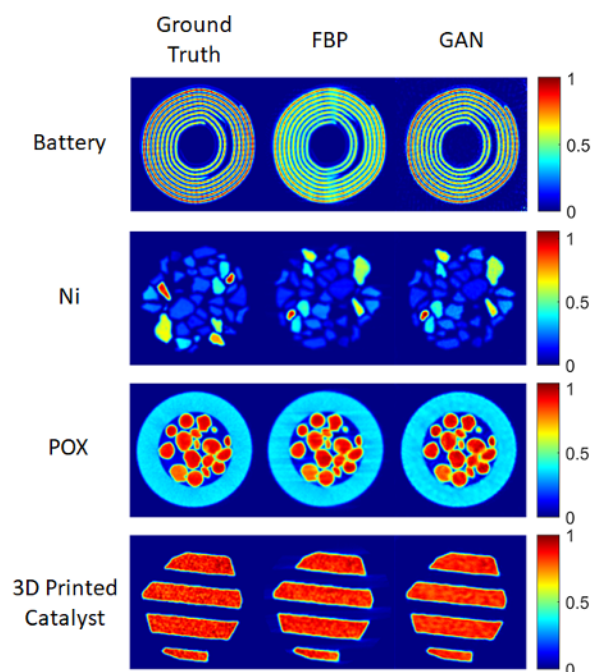Figure 6: Illustration of the GAN model used for tomographic image reconstruction.

Figure 7: Comparison between reconstructed images using the filtered back-projection algorithm (middle) and the GAN (right).

e. Resnet for Super-resolution

We have explored various models and approaches for a super-resolution neural network and have developed a new architecture employing multiple blocks of Residual Networks (RNs). This new architecture, that we call Chemical Image Super-resolution Network (ChemISR-Net), uses a dual input consisting of a low resolution chemical tomography image and a high resolution absorption-contrast image to yield a high resolution chemical tomography image. A schematic representation of the ChemISR-Net is presented in Figure 8.
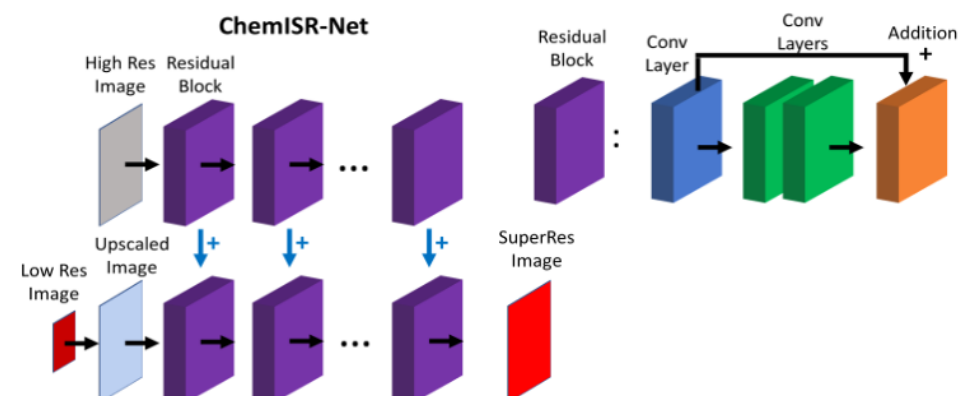

Figure 8: Schematic representation of the ChemISR-Net used for Super-resolution of chemical (hyperspectral) images using a dual input.

The ChemISR-Net contains only convolutional layers (i.e. no dense layers), which keeps the number of parameters low making it scalable and possible to train without the need of a ultra high spec workstation PC. Apart from simulated data, we evaluated the performance of the ChemISR-Net using experimental XRD-CT data we have collected using commercial cylindrical AAA Li-ion LCO batteries. We first performed Rietveld analysis of the XRD-CT data to yield maps corresponding to the concentration (scale factors) and weight fraction of each crystalline

component present in the cells. These maps were then used to calculate X-ray absorption-contrast tomographic images from the XRD-CT data. A comparison of a calculated image and a previously acquired experimental micro-CT of the same type AAA Li-ion LCO battery is shown in Figure 9. It can be clearly seen that our calculated image resembles very accurately the experimental micro-CT images.
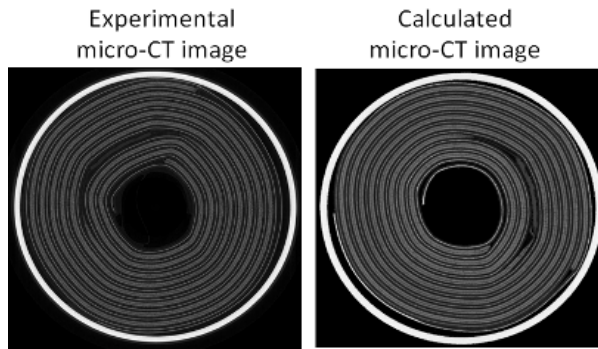

Figure 9: Left: Experimental micro-CT image of a commercial cylindrical AAA Li-ion LCO battery cell. Right: Calculated micro-CT image using experimental XRD-CT data collected from the same type of cell.

We downsampled the XRD-CT images by a factor of 4 (images of 520x520 to 130x130 pixels) and trained the ChemISR-Net using these images and the calculated high resolution micro-CT image as inputs. The trained network was then applied to a different experimental XRD-CT dataset from another battery cell. As shown in Figure 10, the ChemISR-Net yields very promising results, far superior to the conventional bilinear interpolation method. For comparison, we tested the ChemISR-Net against a single input RN architecture (i.e. equivalent to the bottom part of Figure) and it still outperforms it significantly (mean squared error of 0.0510 and 0.0395 for the single and dual input networks respectively).
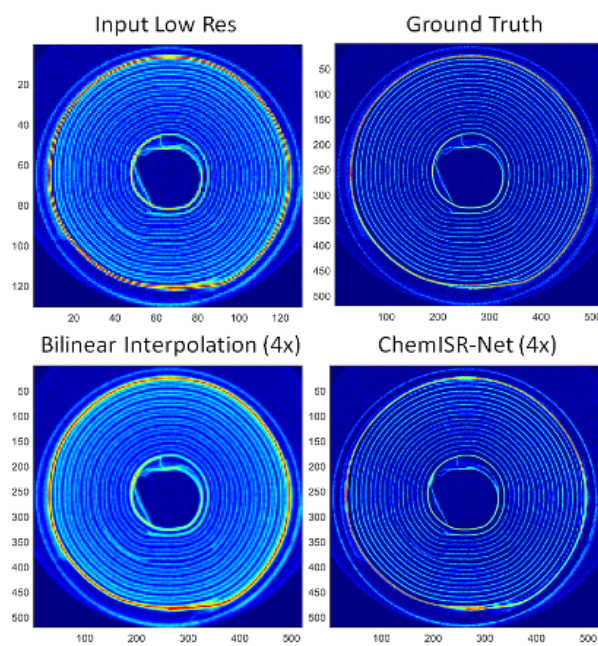

Figure 10: Comparison of the upscaled images obtained using bilinear interpolation (bottom left) and the ChemISR-Net from an XRD-CT image that corresponds to the separator of the Li-ion cell (top left: downsampled image, top right: ground truth image).

8 Outputs

We have collected the outputs from the project in a repository, hosted on GitHub. We have also been building documentation to describe the outputs of the project, as well as an API, with extensive tutorials to allow others to use the tools and the data resulting from the project. The documentation can be found at
https:/superres-tomo.readthedocs.io/en/latest/about.html and the code repository at
https://github.com/keeeto/super-resolution-ml

a. Libraries
The data libraries are stored in hdf5 format, to facilitate easy and efficient use in machine learning projects. The library locations and details are documented in the project docs at:
https://superres-tomo.readthedocs.io/en/latest/benchmark_data.html

b. Networks
We have developed networks for image reconstruction (CNN and GAN), segmentation, denoising and super-resolution. We have packaged the first three together in a consistent fashion in our GitHub repository (https://github.com/keeeto/super-resolution-ml) to ensure that the code can be used by others. We have also written extensive tutorials to demonstrate how to train and apply the networks for these different tasks:

- Reconstruction with a CNN: https://superres-tomo.readthedocs.io/en/latest/tutorials.html#reconstruction-with-a-cnn

- Reconstruction with a dense network: https://superres-tomo.readthedocs.io/en/latest/tutorials.html#reconstruction-with-a-dense-network

- Segmentation 'urlhttps://superres-tomo.readthedocs.io/en/latest/tutorials.html#segmentation-of-x-ray-images

- Denoising https://superres-tomo.readthedocs.io/en/latest/tutorials.html#denoising-of-x-ray-images

- The t-SNE will be added to the public repository after publication

- The ChemISR-Net will be added to the public repository after publication

c. Manuscripts
We have prepared a manuscript based on our analysis regarding the bias introduced by the type of training data used in supervised learning CNNs for tomographic image reconstruction. This has been submitted to the journal Machine Learning Science and Technology and is currently in second round of review.
We are also planning to prepare a manuscript for our super-resolution architecture
We will also hopefully contribute to a book chapter in machine learning approaches used in computed tomography

d. Presentations
The work has been or will be presented at:

- AI3SD Winter Series Seminars

- Invited talk at the MRS conference 2020

- Poster at the MC15 conference 2021

- Invited talk at Berkeley Lab workshop on autonomous materials discovery

- nvited talk at MRS Fall 2021

- Invited talk at MRS Fall 2021

9 Conclusions

In this project we explored and developed neural networks for reconstruction and enhancing X-ray diffraction computed tomography (XRD-CT) images. We found initially that naïve application of convolutional neural networks in a supervised learning approach can suffer seriously from training set bias – this is a finding seldom discussed in the many papers on this topic that are published. We explored and developed two routes to avoiding this bias (i) ensuring balanced datasets, which can be assessed using manifold learning to identify data clusters and ensure that new examples passed to the network are from within the distribution used for training the network; (ii) an unsupervised approach using generative adversarial networks. Both of our approaches provide reliable, scalable approaches for using deep learning to reconstruct XRD-CT data. We developed and employed a new network for enhancing resolution of XRD-CT images by fusing information from higher-resolution (but lower chemical information) micro-CT images. The super-resolution network greatly out-performs traditional image upscaling approaches and provides a route to significant advances in the spatial resolution of XRD-CT.

10 Future Plans

We are going to produce further public datasets for image reconstruction. Additionally, we are writing automated testing for the GitHub repository to ensure the stability of the code. We will continue to publish, document and test the new models and methods that we develop from the project. We are currently in discussions about incorporating our code with the popular existing Xlearn Tomography code (https://xlearn.readthedocs.io/en/latest/) - this will ensure long term support for the tools.

We have also secured funding from Innovate UK to work on projects arising from this work and we are applying for funding from STFC to continue the Super-resolution work to blend neutron and X-ray tomography experiments.

11 References

N/A.

12 Data & Software Links

We developed a new open-source software resource for using deep learning for tomographic reconstruction and image processing. The code is fully documented and automatically tested to ensure usability and reliability. We have also made all of our datasets available as a community resource for testing our models and other models.

- Documentation: https://superres-tomo.readthedocs.io/en/latest/about.html

- Code: https://github.com/keeeto/super_tomo_py

- Datasets: https://superres-tomo.readthedocs.io/en/latest/benchmark_data.html