Helsinki Tomography Challenge 2022 (HTC 2022)

Salla Latva-Äijö, Alexander Meaney, Siiri Rautio,
Samuli Siltanen, Fernando Silva de Moura, Tommi Heikkilä

Last updated 28th of October, 2022, v. 1.1

Contents

1 Introduction

The Finnish Inverse Problems Society (FIPS) proudly presents the Helsinki Tomography Challenge 2022 (HTC2022). We invite all scientists and research groups to test their limited-angle CT reconstruction algorithms on our real-world X-ray data.

The top participants of the challenge will be invited to a minisymposium at the Inverse Days Conference to be held in Kuopio, Finland, in December 2022.

1.1 What is limited-angle tomography?

Computed tomography means reconstructing the internal structure of a physical body using X-ray images of the body taken from different directions. Mathematically, the problem is to recover a non-negative function $f(x)$ from a collection of line integrals of f. Reconstruction of the original, full object requires that measurements are obtained continuously at least 180°around the object.

In limited-angle tomography reconstruction, the object is imaged using only limited angle interval of X-ray projections. The significance of this is that the reconstruction must be computed from an incomplete set of line integrals, a highly ill-posed and challenging task.

2 Challenge description

The purpose of the challenge is to recover the shapes of 2D targets imaged with limited-angle tomography, collected in the Industrial Mathematics Computed Tomography Laboratory at the University of Helsinki, Finland. The experimental setup, targets, and measurement protocol are described in the following sections.

2.1 Targets

The targets are homogenous acrylic disc phantoms of 70 mm in diameter, with holes of varying shapes made with a laser cutter. Each disk has a different number of irregular holes in random locations. Figure 1 show a few examples.

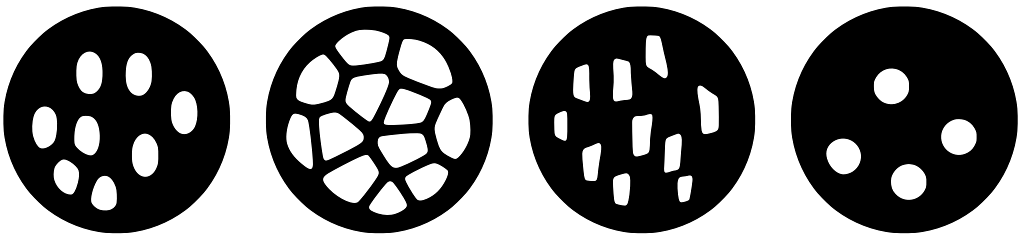

Figure 1: Target examples. Note that examples are provided to the competitors as the training set and, therefore, do not belong to the test set used to evaluate the submissions.

The dataset collected for the HTC2022 challenge consists of two separate sets, with identical experimental setup and settings. One set is provided to the competitors as training set for algorithm development, and the other will be used by the organizers to test the reconstruction algorithms. The test set will be made public after the end of the competition.

2.2 Challenge Outcomes

The outcome of the challenge should be an algorithm which takes in the X-ray data, *i.e.*, the sinogram and it's associated metadata about the measurement geometry, and produces a reconstruction which has been segmented into two components: air and plastic.

2.3 Training Dataset

The training set consists of five phantoms with full angular data. These are designed to facilitate algorithm development and benchmarking for the challenge itself. Four of the training phantoms contain holes. A fifth training phantom is a solid disc with no holes.

We encourage subsampling these datasets to create limited-data sinograms and comparing the reconstruction results to the ground truth obtainable from the full-data sinograms. Note that the phantoms are not all identically centered. Training data for each difficulty group can be created by subsampling these datasets to create limited-data sinograms that match the angular range of each group of the actual challenge data (See Table 1).

As the orientation of CT reconstructions can depend on the tools used, we have included example reconstructions for each of the training phantoms to demonstrate how the reconstructions obtained from the sinograms and the specified geometry should be oriented. The reconstructions have been computed using the filtered back-projection algorithm provided by the ASTRA Toolbox.

We have also included segmentation examples of the reconstructions to demonstrate the desired format for the final competition entries. The segmentation images were obtained by the following steps:

1. Set all negative pixel values in the reconstruction to zero
2. Determine a threshold level using Otsu's method
3. Globally threshold the image using the threshold level
4. Perform a morphological closing on the image using a disc with a radius of 3 pixels.

The competitors do not need to follow the above segmentation procedure, and are encouraged to explore various segmentation techniques for the limited-angle reconstructions.

The competitors are encouraged to generate extra training data using simulations. The organizing committee will not provide the code to generate new targets before the end of the competition.

The training dataset can be downloaded here.

2.4 Limited-angle tomography dataset for the challenge

The actual challenge data consists of 21 phantoms, arranged into seven groups of gradually increasing difficulty, with each level containing three different phantoms, labeled A, B, and C. As the difficulty level increases, the number of holes increases and their shapes become increasingly complex. Furthermore, the view-angle is reduced as the difficulty level increases, starting with a 90-degree field-of-view at level 1, and reducing by 10 degrees at each increasing level of difficulty. Each target is assigned to a single group, therefore, each target is used only once.

The limited data is then passed as input to the submitted algorithms for assessment of the reconstructions. See Section 3.2.1 for more details.

The test targets have been scanned using full-angle tomography, and have been appropriately subsampled to create the challenge data. This enables comparison of the limited-angle reconstruction to the ground truth obtainable from the full-angle data. The ground truth is obtained using the segmentation procedure described in Section 2.3.

The specifications of each difficulty group are defined in Table 1. In this table, *angular range* specifies the view-angle in the limited-angle data. The view-angles in the challenge data will not all begin from 0 degrees.

The complete test dataset will be made public by the end of the competition.

Table 1: Limited-angle tomography difficulty groups

Group	Angular range	Angular increment $\Delta\theta$ (^{o})	Number of projections
1	90°	0.5	181
2	80°	0.5	161
3	70°	0.5	141
4	60°	0.5	121
5	50°	0.5	101
6	40°	0.5	81
7	30°	0.5	61

2.5 Dataset File Format

The dataset is shared using the MATLAB .mat files (version 7.3). Each individual measurement dataset consists of data structure containing the sinogram and its associated metadata, including measurement geometry. In other words, each file contains the measurements for one tomographic image.

UPDATE 28.10.2022: For the actual test data in the challenge, the variable containing the data structure for the CT data will always be named `CtDataLimited`, while in the training data the corresponding variable name is `CtDataFull`.

Python users can load this type of file into their code using the mat73 module. Please refer to this link on how to install and use the module.

3 Rules of the competition

3.1 How to enter the competition

Deadline for registration: Register before 23:59 EET (Eastern European Time) on September 30, 2022, using this electronic form.

3.2 Deadlines and what needs to be submitted

UPDATE 28.10.2022:
Deadline for submission: Send your submission to htc2022("at")fips.fi before November 4, 2022 23:59 EET (Eastern European Time).

What needs to be submitted? Briefly, the codes must be in Matlab or Python 3.X and the algorithms must be shared with us as a private GitHub repository at latest on deadline. Check the following subsections for detailed instructions. Only submissions that fulfill the requirements listed below will be accepted.

The teams can submit more than one reconstruction algorithm to the challenge, however, each algorithm must be in a separate repository. The maximum number of algorithms is the number of members of the team. Your team do not need to register multiple times in case you decide to submit more than one algorithm to the challenge. The team can send a single email with the links to all the repositories.

After the deadline, there is a brief period during which we can troubleshoot the codes together with the competing teams. This is to ensure that we are able to run the codes. The troubleshoot communication is done mainly via 'Issues' section of the submitted repository, so pay attention to any activities in the repository after the deadline.

3.2.1 How we expect your code to be

1. Your code must contain a main routine that we can run to apply your algorithm automatically to every data file in the input directory. This is the file we will run to evaluate your code. Give it a name that is easy to identify, like main.m or main.py.

   Important: The input directory contains only the test dataset. No training dataset is provided to your code during the assessment. Therefore, any training procedures must be performed by your team before the submission.

2. Your main routine must require three input arguments, in this order:

   (a) (string) Folder where the input files are located (.mat files)

   (b) (string) Folder where the output files must be stored

   (c) (int) Group category number, following Table 1. (Values between 1 and 7)

   You can assume we are providing the full path to the directories. Below are the expected formats of the main routines in python and Matlab:

   Matlab users: Note the main function is a callable function:

```
function main(inputFolder,outputFolder,groupNbr)
    ...
    your code comes here
    ...
```

   This is how your main function will be called from the console:

```
>> main('/path/to/input/folder/', '/path/to/output/folder/', 3)
```

   Python users: The main function must be a callable function from the command line. To achieve this you can use either `sys.argv` or `argparse` module.

```
$ python3 main.py /path/to/input/folder/ /path/to/output/folder/ 3
```

3. The main routine must produce one image file of the reconstructed target in the output folder for each dataset file in the input folder, in other words, your code must go through all .mat files in the input folder, producing one image per file. The image with the reconstruction must have the same file name of the input data, apart from the extension. You can assume all files in the input folder belongs to the same difficulty level, provided as the third argument to your main function.

   Output image specification:

   - format: 24-bit PNG (RGB) or 32-bit PNG (RGBA) file. The alpha channel of 32-bit PNG will be discarded without any use.
   - dimensions: 512 x 512 pixels

- colors: The images must be black-and-white, where white represents disk material and black represents holes or background. Your image must contain only black and white pixels. Any other gray levels that eventually might be in the image will be truncated to black/white using a simple threshold of 50%.

4. The teams are allowed to use freely available Python modules or Matlab toolboxes. Toolboxes, libraries and modules with paid licenses can also be used if the organizing committee also has the license. For example, the most usual Matlab toolboxes for image processing and deconvolution can be used (Image processing toolbox, wavelet toolbox, PDE toolbox, computer vision toolbox, deep learning toolbox, optimization toolbox). The teams can contact us to check if other toolboxes are available.

   For getting started, we recommend taking a look at the HelTomo - Helsinki Tomography Toolbox, a Matlab toolbox which will allow you a quick and easy way to start working with the CT data. Note that full use of HelTomo also requires other additional Matlab toolboxes - please refer to the documentation for details.

   Note that using Heltomo for working with the data is by no means compulsory. The metadata for each dataset contains a full specification of the measurement geometry, and the competitors are free to use any and all computational tools they want to in computing the reconstructions and segmentations, as long as the organizing committee also has access to them.

3.2.2 How we expect your Github repository to be

1. Algorithms must be shared with us as a private GitHub repository at latest on deadline.

2. Competitors can update the contents of the repository as many times as needed **before the deadline**, even if they have already sent the link to the organizing committee. We will consider only the latest **RELEASE** of your repository on Github up to the deadline.

   Attention: Simple commits to the main branch will not be considered. You **MUST** create a releases with your final version for the competition. Please see Github's documentation on how to create releases. We will only consider the latest release before the dealine for the competition. If the latest release does not work, we will not try with older releases.

3. If your algorithm requires uploading large files to Github (e.g. with trained coefficients of a neural network), you can use Git Large File Storage (preferable way) or store them in another server and add the link to the Github installation instructions.

4. Your repository must contain a README.md file with at least the following sections:

   - Authors, institutions, addresses.
   - Brief description of your algorithm and a mention of the competition. Here you can also refer to published works of the team related to the algorithm you implemented.
   - Installation instructions, including any requirements.

     Matlab users: Please specify any toolboxes you use.

     Python users: Please specify the modules you use. If you use anaconda, please add to the repository an environment.yml file capable of creating an environment than can run your code (instructions). Otherwise, please add a requirements.txt file generated with `pip freeze` (instructions)

- Usage instructions.
- Present a few examples of the reconstructions from the training set.

5. Finally, the competitors must make their GitHub repositories public at latest on November 30, 2022. In the spirit of open science, only a public code can win HTC2022.

3.3 Scores and leaderboard

The scores and leaderboard are constructed step-wise as follows:

1. All teams start with difficulty level 1. The reconstructions of the three samples (A, B, and C) of this level will be assessed quantitatively following the criteria described in Section 3.3.1, and their scores will be summed, forming the total score of the first level S_1
$$S_1 = S_1^A + S_1^B + S_1^C.$$
The team with the highest score S_1 will be used as reference for the cut-off score of this level: any team with score S_1 at least 25% of the highest score will pass to the next level.

2. The same procedure is repeated for all difficulty levels, up to level 7, but considering only the teams that passed the cut-off score of the previous level:
$$S_N = S_N^A + S_N^B + S_N^C.$$
The cut-off of the levels stays fixed at 25% of the maximum at each level.

3. Denote by Nmax the hardest level that at least one team could enter. If there is only one team, they win. If there are several teams competing in level Nmax, they are ordered in the leaderboard according to their scores at that level.

4. In case of a tie in level Nmax, the previous levels, starting from Nmax-1, will be compared until one of the competitors win. If the tie persists, the organizing committee will make the final decision on the winner.

5. If one team submits more than one algorithm to the competition, then each submission will be temporarily assumed to belong to different 'virtual' teams when computing the scores and cut-offs. However, this team cannot be in more than one position in the leaderboard. In this situation, the organizing committee will consider only the highest performance algorithm when ranking the winners. For example, say team X ends with one algorithm at first place, and another at third place. In this case, the team wins the first place but the third place will be given to another team.

Special situations: The spirit of the competition is that the algorithm is a general-purpose algorithm, capable of reconstructing limited-angle tomography images of the targets. The organizing committee has the right to disqualify an algorithm trying to violate that spirit.

Conflict of interest: researchers affiliated with the Department of Mathematics and Statistics of University of Helsinki are welcome to participate in the competition but will not be added to the leaderboard and cannot win the competition.

3.3.1 Reconstruction assessment method

The reconstructions will be assessed quantitatively, comparing the reconstructed binary image I_r with the ground truth binary image I_t, assigning a numeric score. I_r is assumed to have a dimension of 512 x 512 pixels, otherwise a score 0 will be given to the reconstruction I_r.

The score is based on the confusion matrix of the classification of the pixels between empty (0) or material (1). The confusion matrix is composed by

$$TP = \sum_{i,j} (I_t \cap I_r)_{ij}$$

$$FP = \sum_{i,j} (\bar{I}_t \cap I_r)_{ij}$$

$$FN = \sum_{i,j} (I_t \cap \bar{I}_r)_{ij}$$

$$TN = \sum_{i,j} (\bar{I}_t \cap \bar{I}_r)_{ij}$$

$$\mathbf{M} = \begin{bmatrix} TP & FN \\ FP & TN \end{bmatrix}$$

The score of the reconstruction is given by the Matthews correlation coefficient (MCC)

$$S = \frac{TP \times TN - FP \times FN}{\sqrt{(TP+FP)(TP+FN)(TN+FP)(TN+FN)}},$$

where $S \in [-1, 1]$. A score of +1 (best) represents a perfect reconstruction, 0 no better than random reconstruction, and −1 (worst) indicates total disagreement between reconstruction and ground truth. A python code that implements the scoring will be provided to the competitors. The same code will be used to assess the algorithms.

4 Data collection

The challenge data was measured at the Industrial Mathematics Computed Tomography Laboratory at the University of Helsinki. The measurement device is a cone-beam computed tomography scanner designed and constructed in-house. The scanner consists of an X-ray source, a rotating sample holder, and an X-ray detector (Figure 2).

The data has already been pre-processed with background and flat-field corrections, and compensated for a slightly misaligned center of rotation in the cone-beam computed tomography scanner. The log-transforms from intensity measurements to attenuation data have also been already computed.

We highlight the following geometric definitions, needed to properly specify the X-ray projection operator for the measurement setup:

D_{sd} : Distance from source to detector

D_{so} : Distance from source to origin

D_{sd} : Distance from origin to detector

The X-ray detector data was binned by a factor of four after the measurements, giving a pixel size of 0.2 millimeters.

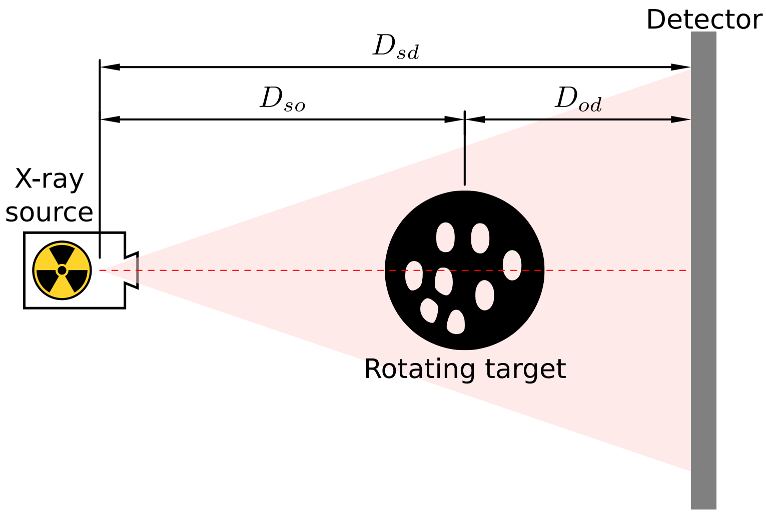

Figure 2: Experimental setup

5 The Grand Prize

The top participants of the challenge will be invited to a minisymposium at the Inverse Days Conference organized by the Finnish Inverse Problems Society (FIPS) to be held in Kuopio, Finland, in December 2022.

The winner also receives a vintage-looking tool for everyday use where determining the viewing angle is necessary.

6 Contact

To contact the HDC2021 organizers send an e-mail to htc2022(”at”)fips.fi.