ISBN: 978-3-95450-238-7 ISSN: 2226-0358 doi:10.18429/JACoW-ICALEPCS2023-TUPDP114

MACHINE LEARNING BASED NOISE REDUCTION OF NEUTRON CAMERA IMAGES AT ORNL

I.V. Pogorelov*, J. Edelen†, M. Henderson, M. Kilpatrick, RadiaSoft LLC, Boulder, USA
S. Calder, B. Vacaliuc, ORNL RAD, Oak Ridge, USA
R.D. Gregory, G. Guyotte, C. Hoffmann, B. Krishna,
ORNL, Oak Ridge, USA

Abstract

Neutron cameras are utilized at the HB2A powder diffractometer to image the sample for alignment in the beam. Typically, neutron cameras are quite noisy as they are constantly being irradiated. Removal of this noise is challenging due to the irregular nature of the pixel intensity fluctuations and the tendency for it to change over time. RadiaSoft has developed a novel noise reduction method for neutron cameras that inscribes a lower envelope of the image signal. This process is then sped up using machine learning. Here we report on the results of our noise reduction method and describe our machine learning approach for speeding up the algorithm for use during operations.

INTRODUCTION

Neutron Cameras are a fairly ubiquitous piece of instrumentation specifically in neutron scattering experiments. Here instrument scientists utilize neutron cameras to visualize the sample location relative to the incident beam to ensure proper alignment in the beam. While these cameras are highly useful, due to the fact that they experience direct radiation the pixels degrade over time and noise develops that degrades the image quality. For many cases this degradation doesn't impact the ability to visualize the sample but for low absorbing samples it can be nearly indistinguishable from the beam. Over the years numerous efforts have been made to remedy noise in neutron camera images. The advent of machine learning has also led to a resurgence of these efforts as the problem is notoriously difficult to solve. Median filters are common approach to removing this noise [1] due to the nonuniform nature of the noise. Recent work on adaptive median filters [2] has shown some promise but median filters can be slow and are not always robust to changes in the noise characteristics over time. Another recent effort utilizes Generative Adversarial Networks to remove signal noise [3]. There is a strong machine learning upside to this technique however it relies on machine learning as the primary mechanism for removal of the noise as opposed to understanding the noise and using machine learning as a speed up tool. Other methods using principle component analysis have also shown promise in recent years [4]. Here we develop a denoising algorithm based on the fundamental characteristics of the noise in the nuetron camera data. We then use convolutional neural netwroks to speed up this calculation for use in real time during oprations. Our algorithm has been fully tested and deployed at the HB2A beamline at Oak Ridge National Laboratory.

DENOISING ALGORITHM

Speckle/"salt-and-pepper" noise is conspicuous in the raw images collected from the neutron camera at HB2A, as can be seen in the left panels of Fig. 1. Cleaning up the images improves our ability to generalize sample identification algorithms in images from neutron cameras and improves the ability for instrument scientists to visualize the sample in the beam. Our machine learning algorithm is trained using data generated using a first principles algorithm developed to remove noise by identifying its key characteristics.

We studied the noise in the images by examining the data on a slice-by-slice basis. We chose a horizontal slice as the aspect ratio of the beam is more favorable for understanding the noise characteristics in this direction. An inspection of the signal contained in a typical row of 640 pixels, such as that pictured in Fig. 2, shows that the statistical properties of the speckle noise are quite different from those of the Gaussian white noise that is predominantly encountered in practice and for which numerous denoising algorithms exist. Two salient features are, the large noise amplitude that is comparable in magnitude to the signal itself, and the "one-sidedness" of the noise, in the sense that the noise only assumes positive values. This special character of the noise makes it possible to attempt its removal by approximating the targeted denoised signal via inscribing an envelope from below, as shown in Fig. 2.

The inscribed envelope can be computed for each row of pixels, whereupon the denoised rows replace the original pixel rows in the image. Interestingly, not every way of inscribing an envelope is equally robust; the best-performing of the algorithms that we explored is based on a kind of min-pooling within a window moving along a line of pixels. It proved important to properly match the size of the min-pooling window to the size of the footprint of the spikes that make up the noise in the images. Figure 1 illustrates the results of application of this procedure to two of the more feature-rich sample images.

MACHINE LEARNING ADAPTATION

We developed a machine learning adaptation of our denoising algorithm for the purpose of speeding up execution during operations. While our baseline algorithm is not prohibitively slow, higher data rates demand faster image

* ilya@radiasoft.net
† jedelen@radiasoft.net

ISBN: 978-3-95450-238-7 ISSN: 2226-0358 doi:10.18429/JACoW-ICALEPCS2023-TUPDP114

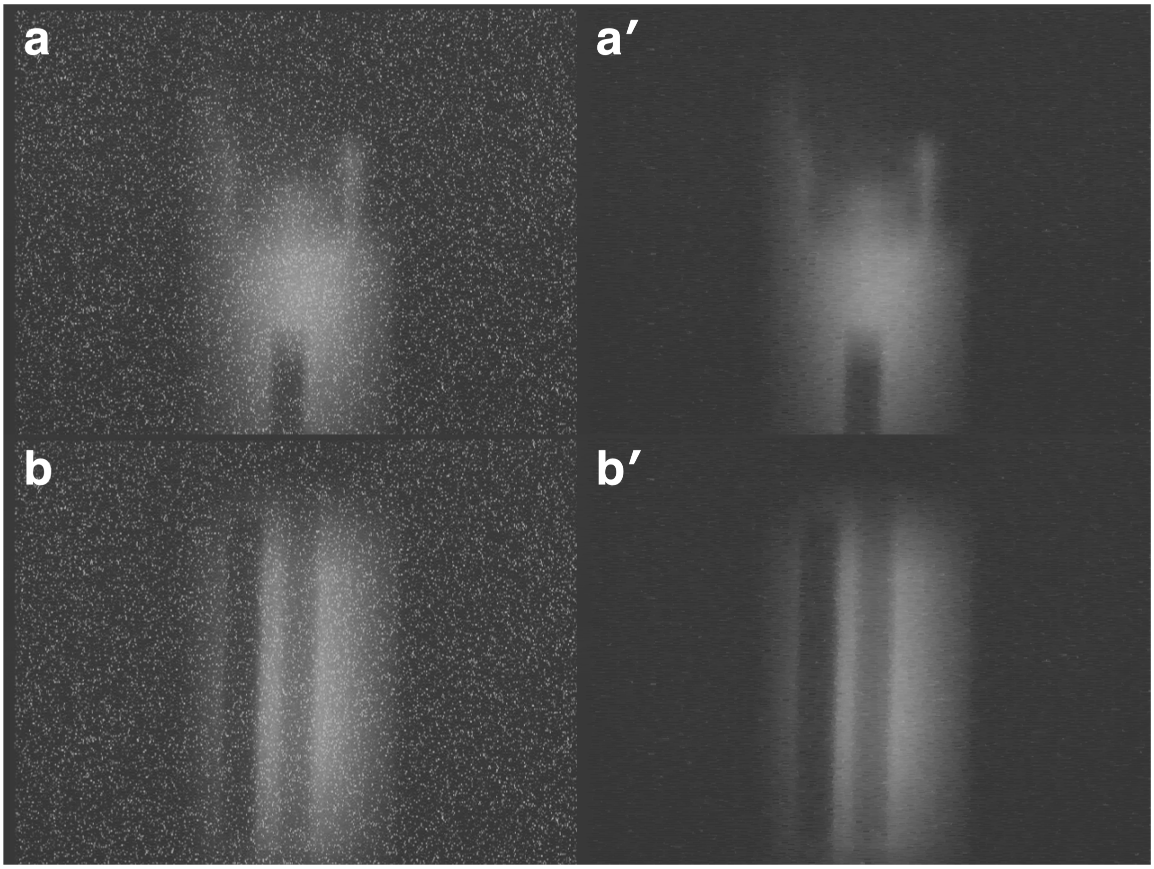

Figure 1: Raw neutron camera sample images at HB2A before (a, b) and after (a', b') “inscribed envelope” denoising.

processing. To this end, training a neural network on our algorithm allows for execution in real time.

We explored model architectures of two types, both belonging to the general family of convolutional neural networks for 2D image processing. The first type is similar to a conventional autoencoder architecture, that is, it consists of an encoder sequence of 2D convolution and pooling layers (with a concurrent increase in the number of feature maps), followed by a decoder sequence of 2D deconvolution layers that ends in an output layer of the same shape as the input image. Dropout layers can optionally be added at the training stage, as well. We varied the size and number of the layers, resulting in the number of trainable parameters ranging from approximately 20 thousand to about 2 million, and trained with different choices of optimizers and loss functions. The models were trained on 640-by-480-pixel HB2A images, with 820 images used for training and 50 each for validation and post-training testing. The target images were produced from the same set by applying our non-ML denoising procedure described above. We performed two kinds of training: First, using the raw images as input and denoised images as target; and second, via pre-training the net as an auto-encoder (*i.e.*, by using the denoised images as both input and target for the output), followed by a conventional supervised-learning training procedure.

The second kind of architectures that we explored for performing image denoising were of the U-Net type, which is similar to the convolution-deconvolution architecture described above, but involves in addition a coupling of the layers on the convolution path into the layers on the deconvolution path [5]. We studied training with and without dropout. The number of trainable parameters in the models we studied varied from about 200 thousand to about 1.5 million. As in the convolution-deconvolution case, we used the raw and denoised HB2A images for regression training, using either conventional training with raw and denoised images, or pre-training using denoised images as both input and targets. Pre-training did not have a noticeable effect on the efficacy of subsequent training.

The trained model that we chose for our timing studies and ultimately for deployment at the neutron imaging beamlines at ORNL is a convolutional neural network of the U-Net kind (as shown in Fig. 3). The contraction path consists of 5 pairs of Conv2D layers, the first 4 of which are each followed by a MaxPooling2D layer. The convolution kernels are 3×3, the number of filters is doubled progressively from 16 to 256, the pooling kernels are 2×2, the padding is ’same’. The expansion path consists of 4 pairs of the 3×3-kernel Conv2D layers with the number of filters progressively decreasing to 16, each pair of Conv2D layers preceded by a

ISBN: 978-3-95450-238-7 ISSN: 2226-0358 doi:10.18429/JACoW-ICALEPCS2023-TUPDP114

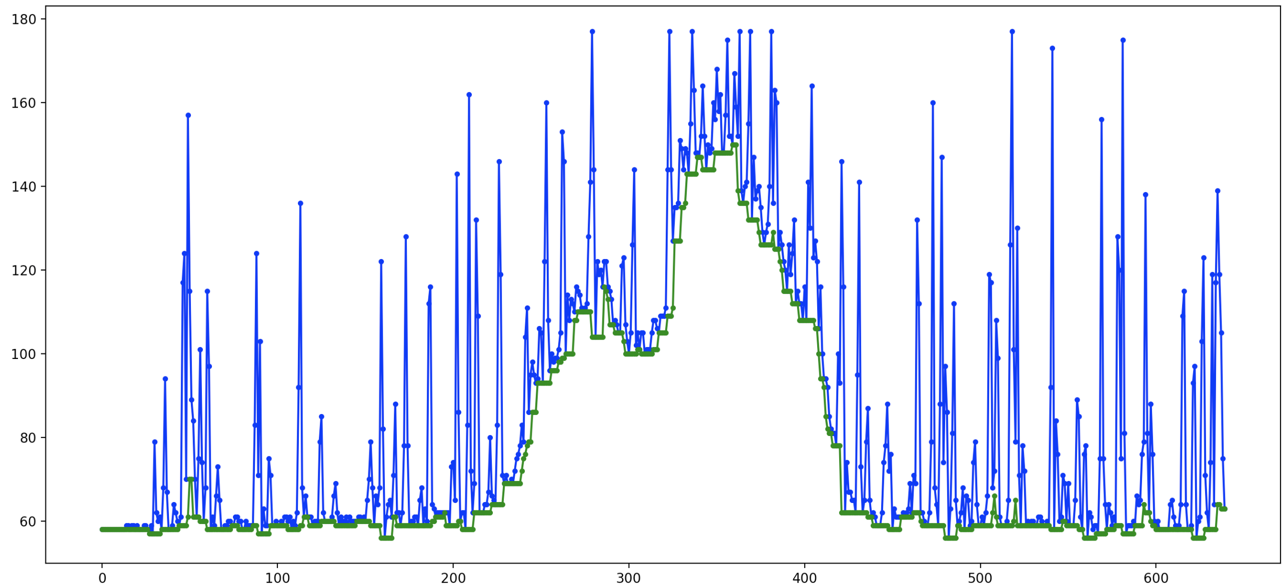

Figure 2: A "typical" 640-pixel-long row of pixels form an HB2A camera image. Denoised signal is approximated by an envelope (green) inscribed from below into the raw noisy data (blue).

Conv2DTranspose layer of kernel size 2 × 2 and stride 2 in each dimension, and a concatenation layer coupling into the layers on the contraction path, in a manner characteristic of the U-Net architecture. We use the tanh activation function in all layers except the output one, where the sigmoid activation function is used. The model has approximately 2M trainable parameters. The model was compiled with the tf.keras Adam optimizer and MSE loss function, and trained without dropout. The training image set consisted of 920 HB2A images pairs (raw and denoised with the non-ML denoiser), of which 820 were used for training, 50 for validation, and 50 for testing. The trained model generalizes quite well; an example of prediction made with the trained net can be seen in Fig. 4.

Figure 4 shows the result of our neural network model trained on data that was denoised using our first principles approach. Left is the original image, middle is processed with our baseline algorithm, and right is the model out put after training. The result is a neural network that does an excellent job reproducing our baseline algorithm with a significantly reduced execution time.

TIMING & PERFORMANCE

To evaluate the timing and performance we executed our baseline algorithm and the machine learning algorithm using the same computational resources. We then evaluated the execution time of the methods considering also the ancillary steps required for loading and scaling the images (necessary for the ML model). The result of our timing analysis is shown in Fig. 5. Here we executed both approaches on a large batch of test images and plotted the execution times as a histogram. The blue histogram shows the machine learning method while the red histogram is the baseline method.

The machine learning algorithm executed almost 8 times faster than the baseline algorithm down to below 100 ms. This would enable the use of of our algorithm with data acquisition rates of 10 Hz which is a typical speed for slow controls needed for sample alignment at these beamlines.

DEPLOYMENT AND TESTING

We deployed out algorithm on an EPICS IOC using the a flexible framework developed in Python specifically for fast deployment. The details of this framework are also presented in these proceedings. Because the IOC is only required to load the image and display a processed image the deployment details are quite simple. The raw image is made available as a process variable in epics. Our software loads the image, pre-processes it, applies the neural network, reprocesses it, and then serves the image to an epics PV that can be displayed in the user interface, Fig. 6. In collaboration with ORNL we installed our software on the beamline and integrated the denoised image into the instrument's GUI. There is now a switch that allows operators to view either the raw image or the denoised image.

CONCLUSIONS

We have successfully developed and implemented a novel denoising technique that utilizes an inscribed envelope to remove speckle from neutron camera images. We then trained a U-Net to replicate this process but with an 8x speed up allowing it to be seamlessly implemented into operations. We have tested and deployed our software at HB2A and are in the process of exploring the extension of this effort to other beamlines at ORNL.

ISBN: 978-3-95450-238-7 ISSN: 2226-0358 doi:10.18429/JACoW-ICALEPCS2023-TUPDP114

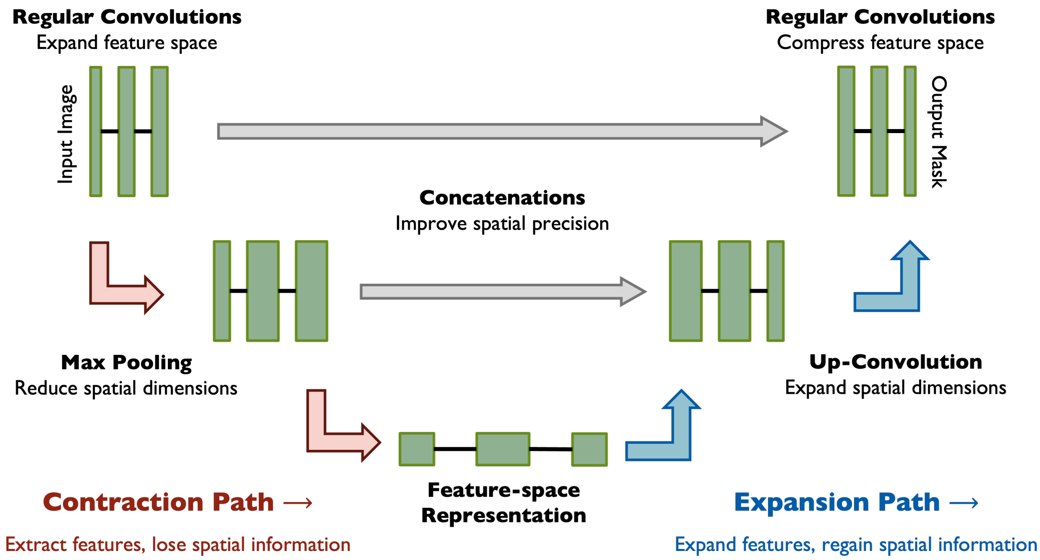

Figure 3: A depiction of the UNet architecture, with explanations of key component.

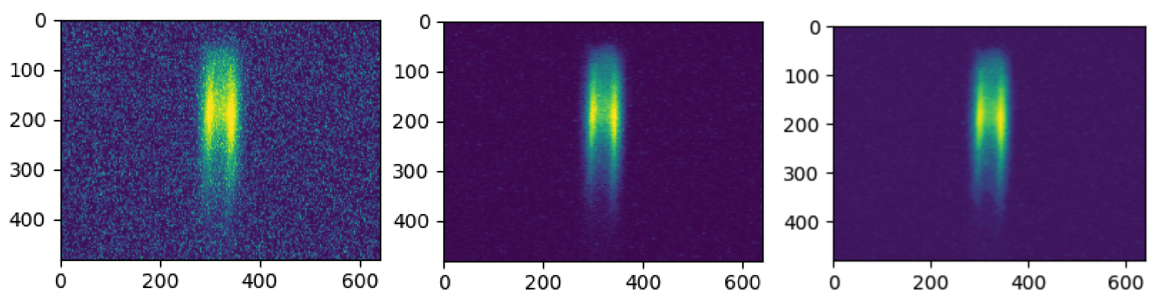

(a) Original image. (b) Denoised with the original algorithm. (c) Image denoised by a trained U-Net.

Figure 4: An HB2A camera image denoised with the original algorithm and a trained U-Net (see text).

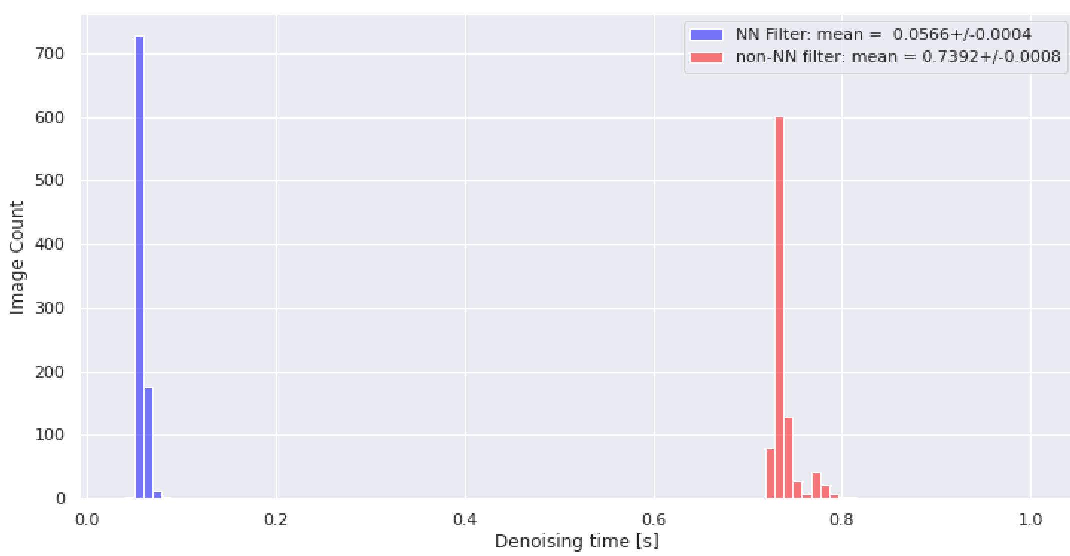

Figure 5: Timing results histogram for denoising with the original algorithm (red) versus the U-Net (blue) that was trained on the data denoised with the original algorithm, run on the same hardware.

ISBN: 978-3-95450-238-7 ISSN: 2226-0358 doi:10.18429/JACoW-ICALEPCS2023-TUPDP114

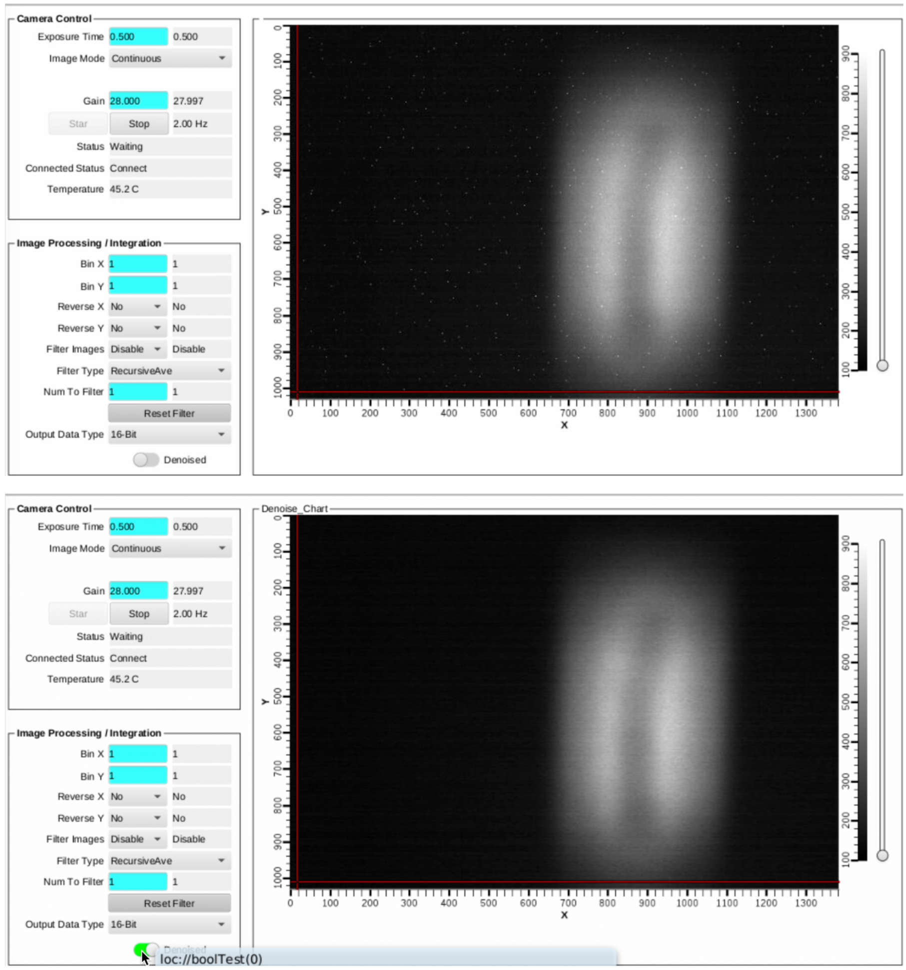

Figure 6: HB2A sample images in the GUI before (top) and after (bottom) "inscribed envelope" denoising during testing.

AKNOWLEDGEMENTS

This work is supported by the U.S. Department of Energy, Office of Science, Office of Basic Energy Sciences, SBIR and STTR Program under Award Number(s) DE-SC0021555.

REFERENCES

[1] X. Cao, Z. Zhang, S. Chen, and T. Li, "Application of improved self-adaptive weighted median filtering algorithm in neutron radiography", in *6th Int. Conf. Intell. Comput. Signal Process. (ICSP)*, 2021, pp. 50–55. doi:10.1109/ICSP51882.2021.9408883

[2] C. Zhao, Y. Yan, H. Li, T. Zhang, and S. Qiao, "An effective gamma white spots removal method for ccd-based neutron images denoising", *Fusion Eng. Des.*, vol. 150, p. 111 375, 2020. doi:10.1016/j.fusengdes.2019.111375

[3] C. Zhao, W. Yin, T. Zhang, X. Yao, and S. Qiao, "Neutron image denoising and deblurring based on generative adversarial networks", *Nucl. Instrum. Methods Phys. Res. A*, vol. 1055, p. 168 505, 2023. doi:10.1016/j.nima.2023.168505

[4] C. Zhao, Y. Yan, Y. Wang, and S. Qiao, "White spots noise removal of neutron images using improved robust principal component analysis", *Fusion Eng. Des.*, vol. 156, p. 111 739, 2020. doi:10.1016/j.fusengdes.2020.111739

[5] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation", in *Proc. MICCAI 2015*, Münich, Germany, 2015, pp. 234–241. doi:10.1007/978-3-319-24574-4_28