

Full-pulse Tomographic Reconstruction with Deep Neural Networks¹

Diogo R. Ferreira,* Pedro J. Carvalho,* Horácio Fernandes,* and JET Contributors²

EUROfusion Consortium, JET, Culham Science Centre, Abingdon, OX14 3DB, UK

**Instituto de Plasmas e Fusão Nuclear, Instituto Superior Técnico, Universidade de Lisboa, 1049-001 Lisboa, Portugal
diogo.ferreira@tecnico.ulisboa.pt*

ABSTRACT

Plasma tomography consists in reconstructing the 2D radiation profile in a poloidal cross-section of a fusion device, based on line-integrated measurements along several lines of sight. The reconstruction process is computationally intensive and, in practice, only a few reconstructions are usually computed per pulse. In this work, we trained a deep neural network based on a large collection of sample tomograms that have been produced at JET over several years. Once trained, the network is able to reproduce those results with high accuracy. More importantly, it can compute all the tomographic reconstructions for a given pulse in just a few seconds. This makes it possible to visualize several phenomena – such as plasma heating, disruptions and impurity transport – over the course of a discharge.

INTRODUCTION

One way to measure plasma radiation is through the use of bolometers, in particular foil bolometers [1]. These bolometers consist of a thin metal foil (about $10\ \mu\text{m}$) coupled with a temperature-sensitive resistance. As the metal foil absorbs radiation power, its temperature changes and there is a proportional change in resistance. This can be measured using a standard setup, such as a Wheatstone bridge. Overall, such system provides a linear response to the absorbed power, in the range from ultraviolet (UV) to soft X-ray [2].

At the Joint European Torus (JET) there is a multi-channel bolometer system comprising a horizontal camera and a vertical camera, with 24 bolometers each [3]. The horizontal camera has a pinhole structure that defines the lines of sight for each of its 24 bolometers. The vertical camera, on the other hand, uses a collimator block to achieve the same purpose [4]. In addition, the vertical camera has an extra 8 bolometers that can be used as reserve channels, so in total the system can provide 56 lines of sight over the plasma, as illustrated in Figure 1.

This two-camera bolometer system is the basis for many tomographic reconstructions that are routinely performed at JET during post-pulse analysis. The method that is used to compute such reconstructions has been developed by Ingesson et al. [5] and actually predates the current bolometer system, having been used with previous generations of soft X-ray diagnostics at JET. In essence, it is an iterative constrained optimization method that minimizes the error with respect to

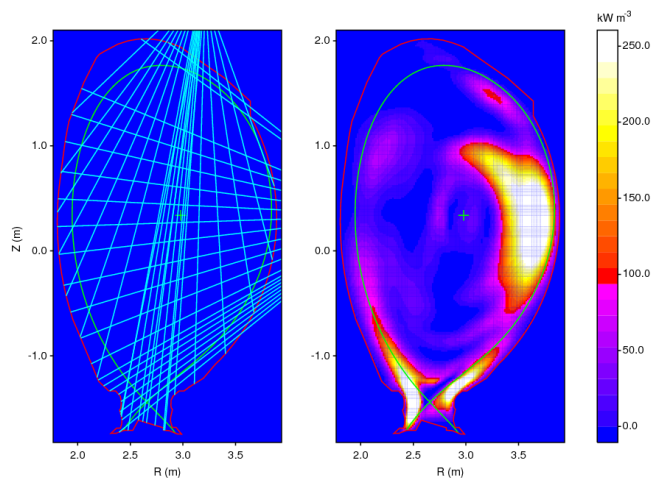


Fig. 1. Lines of sight for the vertical and horizontal cameras (left) and a sample tomographic reconstruction for pulse 89065 at $t=47.8\text{s}$ (right)

the observed measurements, while requiring the solution to be non-negative. To do so, it solves a generalized eigenvalue problem in order to find a solution as a function of Lagrange multipliers, and then adjusts these Lagrange multipliers iteratively until the non-negativity constraints are satisfied [6].

Both the solution to the generalized eigenvalue problem (which can be computed using a standard numerical library) and the iterative adjustments to the Lagrange multipliers take a significant amount of computation time. The total run-time depends on the actual data but, with the code available at JET, it can take more than 1h to produce a reconstruction. This makes it impractical to compute more than a few reconstructions per pulse. There is hardly an opportunity to see how the radiation profile develops across an entire pulse.

To appreciate the computational effort involved, consider the following: the bolometer system at JET has a sampling rate of 5 kHz; for the purpose of noise reduction, a window average of 5 ms is usually applied, which corresponds to 25 samples; subsampling by a factor of 25 yields an effective sampling rate of 200 Hz; so, in principle, it should be possible to have as much as 200 reconstructions per second of pulse time; for a pulse of about 30 seconds, this means a total of 6000 reconstructions; at an average of 1h per reconstruction, this would require 250 days.

Clearly, another way to compute the reconstructions for an entire pulse should be found. The following sections describe how a deep neural network was devised and trained on existing tomograms to produce the same results. A small amount of error is more than compensated by a large computational speedup (both to be quantified below). With this approach, it

¹This paper is an original manuscript describing the work presented at the 2nd IAEA Technical Meeting on Fusion Data Processing, Validation and Analysis that took place in Boston, USA, from May 30 to June 2, 2017.

²See the author list of “Overview of the JET results in support to ITER” by X. Litaudon et al. in *Nuclear Fusion*, **57**, 10, 102001 (2017).

becomes possible to analyze the time evolution of the radiation profile in great detail. The pulses discussed at the end of this paper, which illustrate disruptions and impurity transport, are only a few examples of what can be done.

DEEP NEURAL NETWORKS

Deep learning [7, 8] is having a tremendous impact in fields such as image processing and natural language processing. In particular, convolutional neural networks (CNNs) have been very successful at classifying input images into a set of output classes. This has been demonstrated in the recognition of hand-written digits [9], in the classification of Web images [10] and in the annotation of online videos [11], to cite only a few applications.

In general, CNNs have a common overall structure, which is depicted in Figure 2. This comprises, namely:

- An input layer, which receives an image, a multi-channel image (in case of color images), or a set of video frames.
- One or more convolutional layers, where multiple filters (in the form of a sliding window) are applied to the same input. Each filter produces a different *feature map*.
- A subsampling layer after each convolution or, alternatively, some form of subsampling applied during the convolution itself, namely by making the sliding window *stride* in larger steps across the input.
- One or more densely-connected layers at the end to perform classification based on the features extracted by the convolutional layers.
- An output layer with the same number of nodes as the number of output classes.

The idea is to have a first stage of convolutional layers to extract meaningful features from the input image, and a second stage of dense layers to perform the actual classification based

on those features. Every time a convolutional filter is applied, it produces a new feature map as output. The first convolutional layer operates directly on the input image; subsequent layers operate on the feature maps produced by the previous convolutional layer. The purpose of having subsampling (or a stride greater than 1) is to make the feature maps smaller and allow their number to progressively increase, while keeping the network under a manageable size.

In a typical CNN, the input is a 2D image and the output is a 1D vector of class probabilities. In tomography, however, the scenario is the opposite: the input is a 1D vector of bolometer measurements and the output is a 2D image of the plasma radiation profile. Therefore, for the purpose of tomographic reconstruction, it makes sense to think of the inverse of a CNN, i.e. a kind of “deconvolutional” network that is able to reconstruct a 2D image from its 1D projections.

In the literature, deconvolutional neural networks have been used for image generation [12, 13]. By specifying the class label and the camera position (1D data), the network is able to generate an object (2D image) of the specified class from the given camera view. In this case, the network architecture is basically the reverse of a CNN, with a couple of dense layers at the beginning, and a series of convolutional layers (with upsampling) to generate the output image.

In previous work [14], we have shown that such kind of network can produce tomographic reconstructions with high accuracy when trained on certain subsets of JET pulses. In the present work, we improved the network architecture by replacing the convolution+upsampling layers with proper deconvolutional layers (i.e. transposed convolutions [15]) in order to obtain the logical inverse of a CNN. We have also removed the requirement for any data preprocessing, so the measurement data coming from the bolometers can be fed directly to the network. These improvements allow the network to be trained much faster and with larger amounts of data.

The resulting network architecture is depicted in Figure 3. After the input layer for the 56 bolometer channels, there are two dense layers with 7500 nodes, which can be reshaped

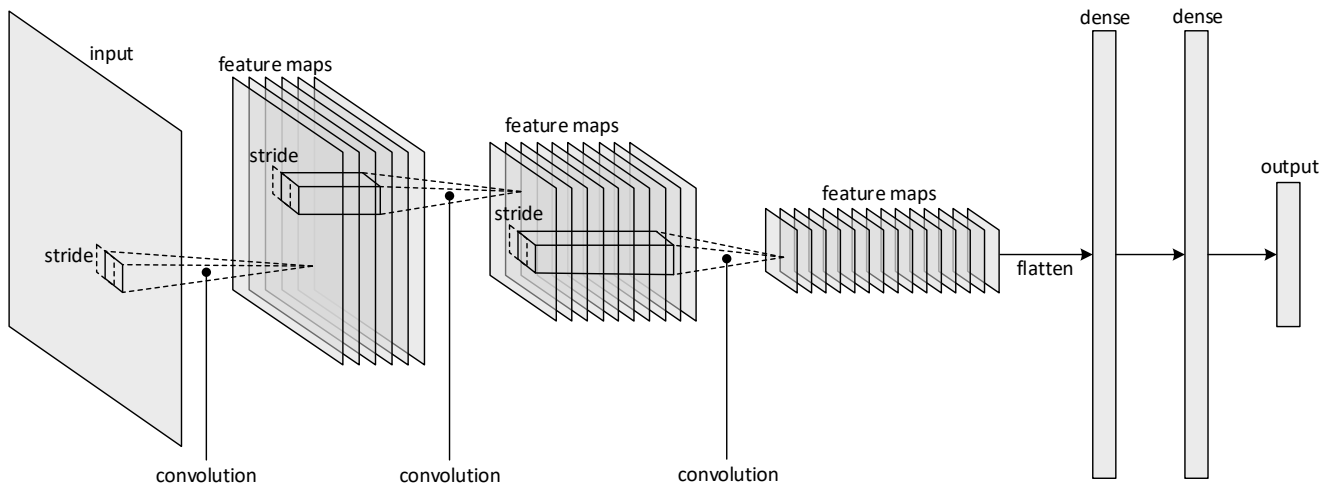


Fig. 2. Typical structure of a convolutional neural network (CNN)

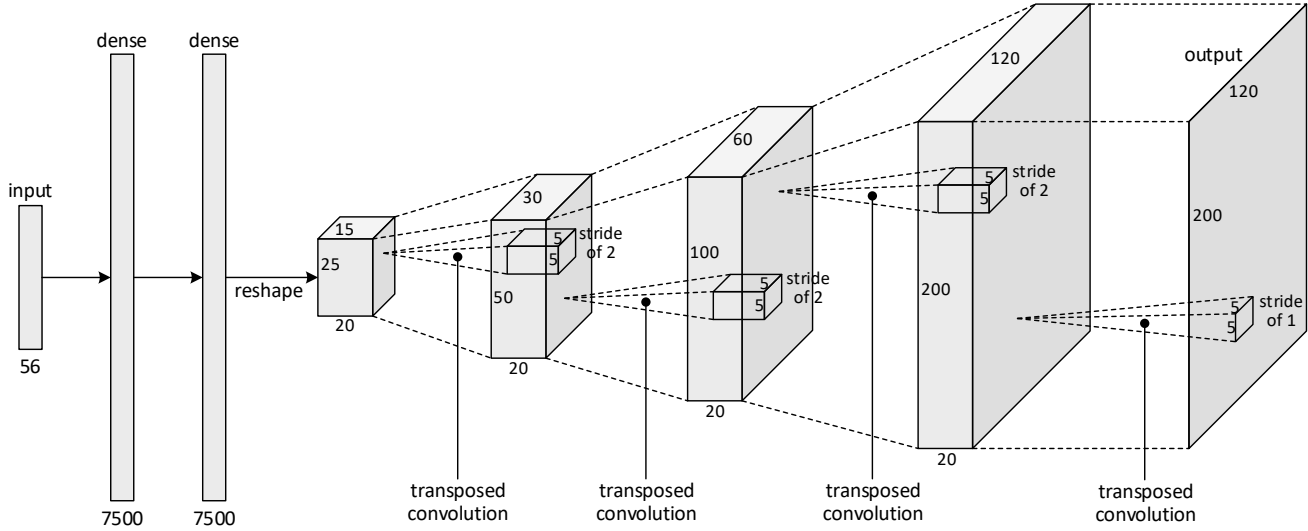


Fig. 3. Deconvolutional neural network for tomographic reconstruction

into a 3D structure of size $25 \times 15 \times 20$. This structure can be regarded as comprising 20 feature maps of size 25×15 . By applying a series of transposed convolutions, these feature maps can be brought up to a size of 200×120 , from which the output image is generated by one last convolution.

The transposed convolution is the inverse of the convolution in the sense that, if a sliding window would be applied to the output, the result would be the feature map given as input. It can be shown that learning a transposed convolution is equivalent to learning a weight matrix that is the transpose of a regular convolution [15], hence the name of this operation.

In Figure 3, the upsampling of the feature maps from 25×15 to 200×120 is achieved by having each transposed convolution operate with a stride of two pixels (i.e. one pixel is being skipped between each two consecutive positions of the sliding window). This means that the output is two times larger and taller than the input, except for the very last convolution which uses a stride of one to keep the same size.

TRAINING AND ACCURACY

In order to train the network in Figure 3, we gathered as many sample reconstructions as possible. For this purpose, we collected every single reconstruction that has been produced at JET since the installation of the ITER-like wall (ILW) in 2011 [16]. This yielded a total of 24203 sample tomograms, which have been separated into 90% (21783) for training and 10% (2420) for validation.

The network was trained using an adaptive gradient descent algorithm [17] with a relatively small learning rate (10^{-4}) and a large batch size (411). The batch size was chosen in order to have about 50 batches, and in order to avoid having any partially filled batch in the training set. Since 411 is a divisor of 21783, this gives $21783/411=53$ batches, and therefore there are 53 updates per epoch, where an epoch is one run through the whole training set.

The network was trained on an Nvidia Titan X graphics

processing unit (GPU) until the error in the validation set no longer improved. After 60 hours (12652 epochs) the minimum error on the validation set was achieved at epoch 6911, and there was no improvement after that. In fact, after that point the error on the validation set started to increase (this is marginally visible in Figure 4), which should be taken as a symptom of overfitting. Therefore, we kept the network weights from epoch 6911.

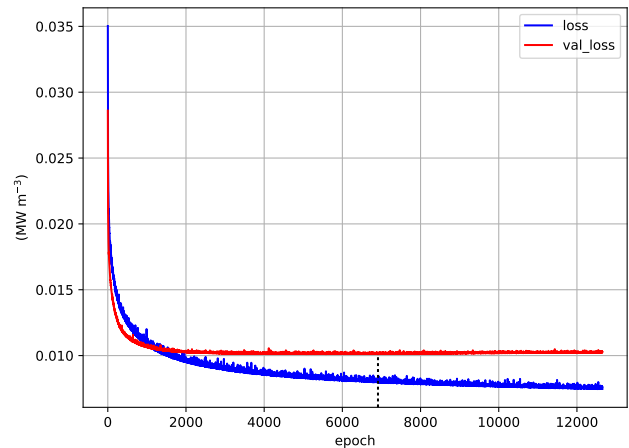


Fig. 4. Loss and validation loss during training

The minimum error on the validation set, measured as (pixel-wise) mean absolute error, was about 10 kW m^{-3} . The relatively small size of this error can be appreciated by comparing it to the dynamic range of the sample reconstruction in Figure 1.

As an example of the results that can be obtained with the trained network, Figure 5 shows three reconstructions from pulse 92213. On the left is the reconstruction produced by the iterative constrained optimization method, and on the right is the reconstruction produced by the neural network. It should

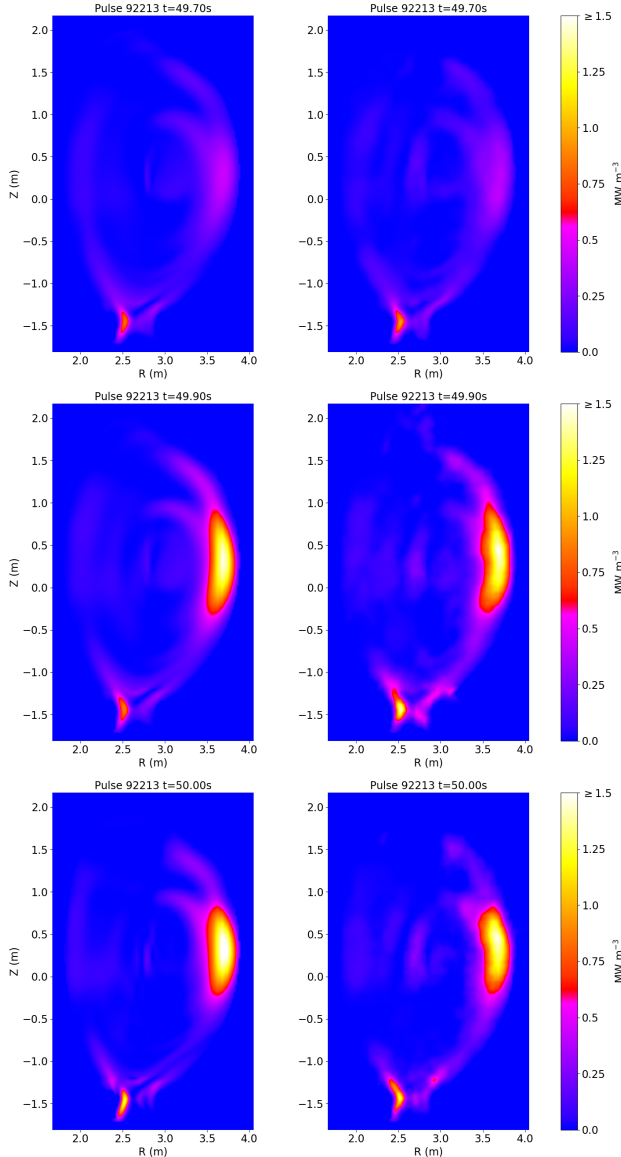


Fig. 5. Three reconstructions from pulse 92213 at $t=49.7s$ (top), $t=49.9s$ (middle) and $t=50.0s$ (bottom). Original reconstruction (left) vs. neural network (right).

be noted that these particular reconstructions were not part of the training set or the validation set.

Despite some perceptible differences, it is clear that the neural network knows exactly where the focus of radiation is (top and middle row) and can adapt even to small changes in shape (middle and bottom row). To quantify the differences between each pair of reconstructions, it is possible to use image quality metrics such as structural similarity (SSIM) [18], peak signal-to-noise ratio (PSNR) [19] and normalized root mean square error (NRMSE) [20]. Table I shows the kind of results that can be expected with these metrics.

The examples in Figure 5 have not been carefully selected to yield particularly good results and, in practice, it might be possible to obtain even better results. For example, Table II shows the results obtained on the validation set, where all

time	SSIM	PSNR	NRMSE
$t=49.7s$	0.9354	30.386	0.06925
$t=49.9s$	0.9239	28.161	0.08437
$t=50.0s$	0.9134	28.864	0.07750

TABLE I. Quality metrics on the reconstructions of Fig. 5

	SSIM	PSNR	NRMSE
mean	0.9573	33.791	0.0678
std. dev.	0.0609	6.064	0.0415

TABLE II. Quality metrics on the validation set

metrics are, on average, slightly better. Also, in previous work [14] we obtained even better results but it should be noted that, in that case, the network was trained separately on smaller and more uniform subsets of data.

In this work, the network was trained on a wide range of pulses, from 80000 (Aug. 2011, after the installation of the ILW) to 92504 (Nov. 2016, the last pulse at the time of this writing). It is possible that, during these 5 years, the bolometer system suffered some degradation (some channels might be broken or malfunctioning occasionally) and yet the neural network can make sense of the bolometer data to provide accurate reconstructions across this wide range of pulses.

FULL-PULSE RECONSTRUCTION

Once trained, the network can be used to generate the reconstruction for any given pulse and time. In fact, on the same GPU that has been used before, the trained network can produce about 3000 reconstructions per second (or 100 reconstructions per second on a standard quad-core CPU). This means that the computation time for 6000 reconstructions (which would take 250 days, as mentioned in the introduction) can be brought down to just 2s, making it possible to generate at once all the reconstructions for an entire pulse.

Figure 6 shows the reconstruction of pulse 92213 from $t=49.62s$ onwards, with a time increment of 0.1s (due only to space restrictions). The first row shows a focus of radiation developing on the outer wall, as we have seen in Figure 5. However, Figure 6 also shows what happens afterwards: the focus of radiation seems to slowly fade away (rows 1, 2), only to reappear later at the core with particularly strong intensity (rows 2, 3). (The dynamic range of these frames is the same as in Figure 5, i.e. 0 to 1.5 MW m^{-3} .)

The radiation peaking stays in the core for a relatively long time (at least from $t=52.32s$ to $t=53.42s$), while changing slightly in shape during that interval. Eventually, it also fades away as the heating systems are being turned off. However, just as it seemed that the plasma was about to soft land, there is a disruption at $t=53.72s$.

Clearly, there are phenomena of impurity sputtering, transport and accumulation in this pulse. Heavy ions (e.g. tungsten) coming from the wall and the divertor can lead to significant radiation losses, and these phenomena are especially noticeable in the radiation profile. It is therefore not surprising that bolometer systems are one of the key diagnostics used in

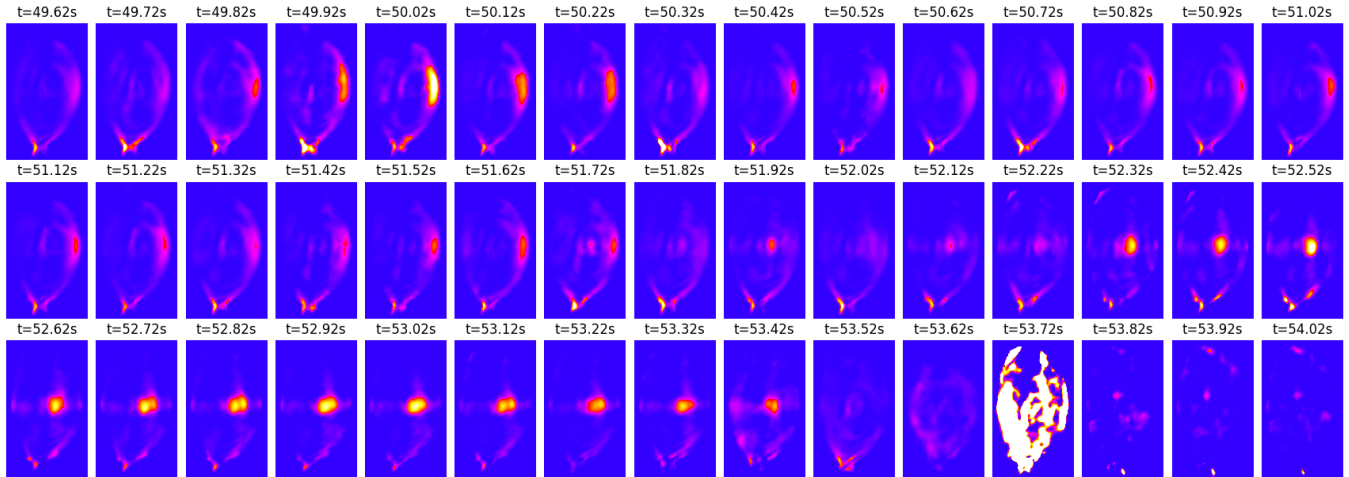


Fig. 6. Reconstruction of pulse 92213 from $t=49.62s$ to $t=54.02s$ with a time step of $0.1s$

disruption studies [21, 22, 23, 24], either by providing a measurement of total radiation across the pulse, or by providing the tomographic reconstruction at specific points of interest, e.g. during the thermal or current quench.

With full-pulse reconstruction, it is possible to carry out such analysis in more detail. For example, we can zoom into the final moments of pulse 92213 to observe what happens at the onset of (and even during) the disruption. Figure 7 shows the reconstruction from $t=53.688s$ to $t=53.732s$ with a time increment of just $1ms$. In these frames, there is an apparent swing from the outer wall to the inner wall, before the divertor lightens and a full-blown disruption occurs.

Although the dynamics of disruptions are not yet fully understood, one of the main causes is impurity accumulation in the core, which decreases the core temperature, creating a hollow temperature profile that eventually leads to core collapse [25, 26]. There are a number of experiments at JET where impurities are deliberately injected into the plasma in order to study impurity transport and how impurities get to the

core [27, 28, 29, 30]. In these experiments, the impurities are created by laser ablation of metal plate, and are subsequently injected into the plasma from the outer wall.

Figure 8 shows an excerpt of the full-pulse reconstruction for an impurity-injection experiment (92286). Initially, the pulse is relatively quiet, but in the second row there is a noticeable change, with the formation of a ring around the edge. This is the moment when impurities (in this case, tungsten) have been injected into the plasma (around $t=43.1s$). The ring turns into a cloud and, as the cloud dissipates, the impurities appear to have concentrated at the core.

This impurity transport can be analyzed in more detail by decreasing the time step, and also by decreasing the dynamic range of these plots in order to have a better visualization. Figure 9 shows a reconstruction with a time increment of $4ms$ and with a relatively low dynamic range of 0 to 30 kW m^{-3} . On the second frame in the top row ($t=43.110s$), the injection of tungsten is clearly visible at a midpoint on the outer wall.

Following the injection, the impurities flow in clockwise

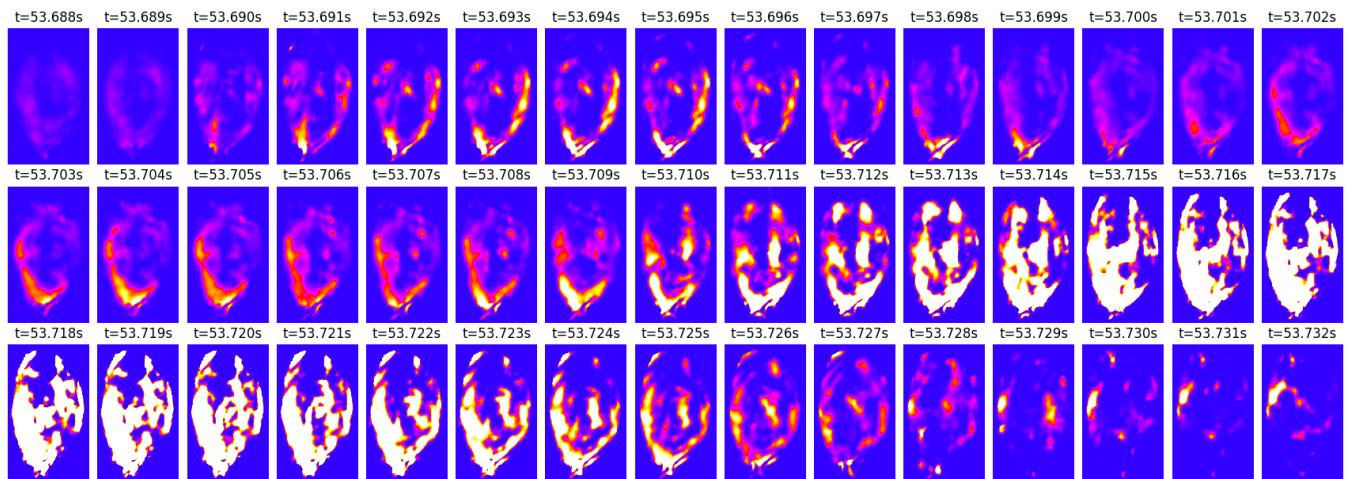


Fig. 7. Reconstruction of pulse 92213 around disruption time with a step of $0.001s$

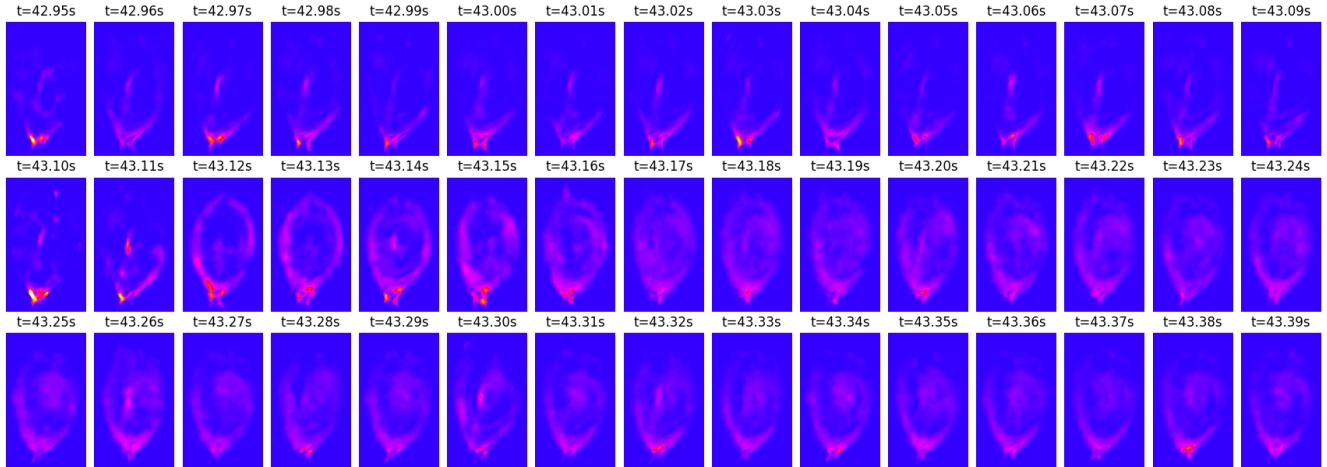


Fig. 8. Reconstruction of pulse 92286 from $t=42.95s$ to $t=43.39s$ with a time step of $0.01s$

direction towards the divertor and along the separatrix until forming a closed ring at $t=43.122s$. The ring thickens and the particles appear to follow a spiral movement towards the center. In the last row, a separate clump of radiation is already well established in the core, with another focus of activity going on in the divertor region.

CONCLUSION

A deep neural network, with an architecture that resembles the inverse of a traditional CNN, can be trained to produce tomographic reconstructions much faster than traditional methods. This makes it possible to reconstruct entire pulses, providing a detailed visualization of several phenomena, such as impurity transport and disruptions.

The same neural network approach can be applied to other diagnostics in the same machine, or to similar diagnostics in other machines, provided that there is sufficient training data. However, it should be noted that the neural network cannot produce better results than the method that was used to

generate the training data itself. In general, the neural network is just an approximation with arbitrarily good precision and, as such, its results may suffer from the same problems or artifacts that are present in the training data.

The tomographic routines at JET have been stable and have been widely used for a number of years. This provided a consistent a relatively large training dataset. In other machines, such as ITER, it will not be possible to use neural network methods from the start of operations, for lack of training data. However, it may be possible to train a neural network based on the part of the physics that is understood and can be simulated, such as the large-scale simulations provided by magnetohydrodynamics and impurity transport codes.

As future work, we intend to use the full-pulse reconstructions to analyze the precursors of disruptions in order to be able to mitigate, or even avoid those disruptions, if at all possible. As a testbed for ITER, JET is perhaps the last opportunity to tackle this problem before having to deal with it on a much larger scale.

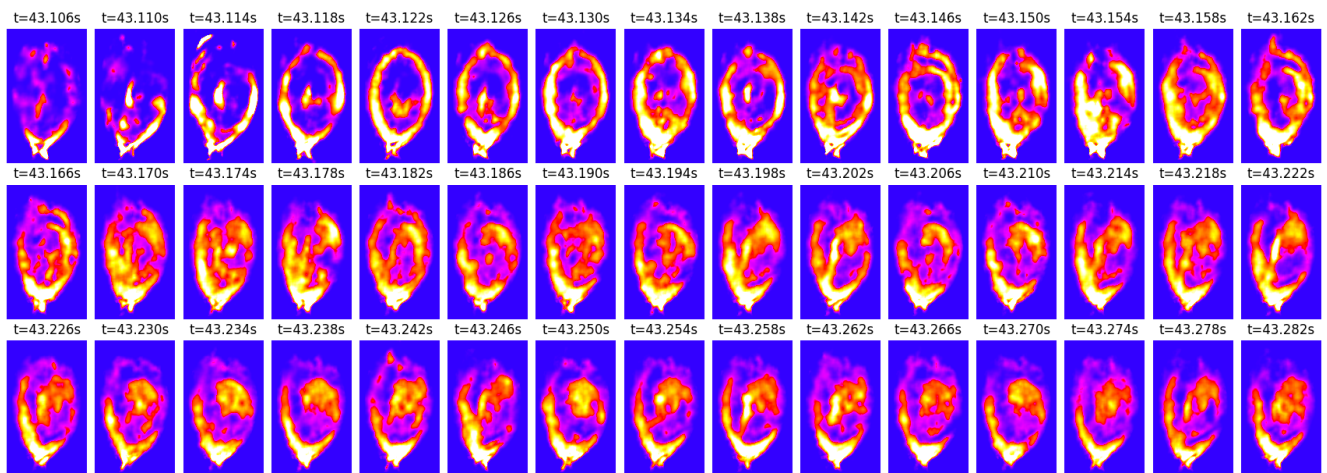


Fig. 9. Reconstruction of 92286 pulse from $t=43.106s$ to $t=43.282s$ with a time step of $0.004s$

ACKNOWLEDGMENTS

This work has been carried out within the framework of the EUROfusion Consortium and has received funding from the Euratom research and training programme 2014-2018 under grant agreement No 633053. IPFN activities also received financial support from *Fundação para a Ciência e a Tecnologia* (FCT) through project UID/FIS/50010/2013. The Titan X GPU used in this work was donated by NVIDIA Corporation.

We would also like to thank Ewa Pawelec at the University of Opole, in Poland, for insights into laser ablation experiments, and Peter Lomas at JET, UKAEA, for several helpful discussions regarding disruptions.

SUPPLEMENTARY MATERIALS

Videos of the reconstructed pulses presented in this paper (92213 and 92286) can be found online at: <http://web.tecnico.ulisboa.pt/diogo.ferreira/videos/>

REFERENCES

1. L. C. INGESSON, B. ALPER, B. J. PETERSON, and J.-C. VALLET, “Tomography Diagnostics: Bolometry and Soft-X-Ray Detection,” *Fusion Science and Technology*, **53**, 2, 528–576 (2008).
2. K. F. MAST, J. C. VALLET, C. ANDELFINGER, P. BETZLER, H. KRAUS, and G. SCHRAMM, “A low noise highly integrated bolometer array for absolute measurement of VUV and soft x radiation,” *Review of Scientific Instruments*, **62**, 3, 744–750 (1991).
3. A. HUBER, K. MCCORMICK, P. ANDREW, P. BEAUMONT, S. DALLEY, J. FINK, J. FUCHS, K. FULLARD, W. FUNDAMENSKI, L. INGESSON, F. MAST, S. JACHMICH, G. MATTHEWS, R. MERTENS, V. PHILIPPS, R. PITTS, S. SANDERS, and W. ZEIDNER, “Upgraded bolometer system on JET for improved radiation measurements,” *Fusion Engineering and Design*, **82**, 5, 1327–1334 (2007).
4. K. MCCORMICK, A. HUBER, C. INGESSON, F. MAST, J. FINK, W. ZEIDNER, A. GUIGON, and S. SANDERS, “New bolometry cameras for the JET Enhanced Performance Phase,” *Fusion Engineering and Design*, **74**, 1–4, 679–683 (2005).
5. L. INGESSON, B. ALPER, H. CHEN, A. EDWARDS, G. FEHMERS, J. FUCHS, R. GIANNELLA, R. GILL, L. LAURO-TARONI, and M. ROMANELLI, “Soft X ray tomography during ELMs and impurity injection in JET,” *Nuclear Fusion*, **38**, 11, 1675–1694 (1998).
6. G. C. FEHMERS, L. P. J. KAMP, and F. W. SLUIJTER, “An algorithm for quadratic optimization with one quadratic constraint and bounds on the variables,” *Inverse Problems*, **14**, 4, 893 (1998).
7. Y. LECUN, Y. BENGIO, and G. HINTON, “Deep learning,” *Nature*, **521**, 436–444 (2015).
8. I. GOODFELLOW, Y. BENGIO, and A. COURVILLE, *Deep Learning*, MIT Press (2016).
9. Y. LECUN, L. BOTTOU, Y. BENGIO, and P. HAFFNER, “Gradient-based learning applied to document recognition,” *Proceedings of the IEEE*, **86**, 11, 2278–2324 (1998).
10. A. KRIZHEVSKY, I. SUTSKEVER, and G. E. HINTON, “ImageNet Classification with Deep Convolutional Neural Networks,” *Advances in Neural Information Processing Systems*, **25**, 1097–1105 (2012).
11. A. KARPATY, G. TODERICI, S. SHETTY, T. LEUNG, R. SUKTHANKAR, and L. FEI-FEI, “Large-Scale Video Classification with Convolutional Neural Networks,” in “IEEE Conference on Computer Vision and Pattern Recognition,” (2014), pp. 1725–1732.
12. A. DOSOVITSKIY, J. T. SPRINGENBERG, and T. BROX, “Learning to generate chairs with convolutional neural networks,” in “IEEE Conference on Computer Vision and Pattern Recognition,” (June 2015), pp. 1538–1546.
13. A. DOSOVITSKIY, J. T. SPRINGENBERG, M. TATARCHENKO, and T. BROX, “Learning to Generate Chairs, Tables and Cars with Convolutional Networks,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **39**, 4, 692–705 (2017).
14. F. A. MATOS, D. R. FERREIRA, and P. J. CARVALHO, “Deep learning for plasma tomography using the bolometer system at JET,” *Fusion Engineering and Design*, **114**, 18–25 (January 2017).
15. V. DUMOULIN and F. VISIN, “A guide to convolution arithmetic for deep learning,” *arXiv:1603.07285* (2016).
16. G. F. MATTHEWS, M. BEURSKENS, S. BREZINSEK, M. GROTH, E. JOFFRIN, A. LOVING, M. KEAR, M.-L. MAYORAL, R. NEU, P. PRIOR, V. RICCARDO, F. RIMINI, M. RUBEL, G. SIPS, E. VILLEDIEU, P. DE VRIES, and M. L. WATKINS, “JET ITER-like wall – overview and experimental programme,” *Physica Scripta*, **2011**, T145, 014001 (2011).
17. D. P. KINGMA and J. BA, “Adam: A Method for Stochastic Optimization,” *arXiv preprint arXiv:1412.6980* (2014).
18. Z. WANG, A. C. BOVIK, H. R. SHEIKH, and E. P. SIMONCELLI, “Image quality assessment: from error visibility to structural similarity,” *IEEE Transactions on Image Processing*, **13**, 4, 600–612 (2004).
19. Q. HUYNH-THU and M. GHANBARI, “Scope of validity of PSNR in image/video quality assessment,” *Electronics Letters*, **44**, 13, 800–801 (2008).
20. J. R. FIENUP, “Invariant error metrics for image reconstruction,” *Applied Optics*, **36**, 32, 8352–8357 (1997).
21. G. ARNOUX, A. LOARTE, V. RICCARDO, W. FUNDAMENSKI, and A. HUBER, “Heat loads on plasma facing components during disruptions on JET,” *Nuclear Fusion*, **49**, 8, 085038 (2009).
22. V. RICCARDO, G. ARNOUX, P. CAHYNA, T. C. HENDER, A. HUBER, S. JACHMICH, V. KIPTILY, R. KOSLOWSKI, L. KRLIN, M. LEHNEN, A. LOARTE, E. NARDON, R. PAPROK, and D. TSKHAKAYA, “JET disruption studies in support of ITER,” *Plasma Physics and Controlled Fusion*, **52**, 12, 124018 (2010).
23. A. HUBER, G. ARNOUX, M. BEURSKENS, S. BOZHENKOV, S. BREZINSEK, T. EICH, C. FUCHS, W. FUNDAMENSKI, S. JACHMICH, U. KRUEZI, M. LEHNEN, A. LOARTE,

- G. MATTHEWS, P. MERTENS, P. MORGAN, V. PHILIPPS, R. PITTS, V. RICCARDO, U. SAMM, B. SCHWEER, G. SERGIENKO, and M. STAMP, "Radiation loads onto plasma-facing components of JET during transient events – Experimental results and implications for ITER," *Journal of Nuclear Materials*, **415**, 1, S821–S827 (2011).
24. M. LEHNEN, A. ALONSO, G. ARNOUX, N. BAUMGARTEN, S. A. BOZHENKOV, S. BREZINSEK, M. BRIX, T. EICH, S. N. GERASIMOV, A. HUBER, S. JACHMICH, U. KRUEZI, P. D. MORGAN, V. V. PLYUSNIN, C. REUX, V. RICCARDO, G. SERGIENKO, and M. F. STAMP, "Disruption mitigation by massive gas injection in JET," *Nuclear Fusion*, **51**, 12, 123010 (2011).
 25. P. C. DE VRIES, M. F. JOHNSON, B. ALPER, P. BURRATTI, T. C. HENDER, H. R. KOSLOWSKI, and V. RICCARDO, "Survey of disruption causes at JET," *Nuclear Fusion*, **51**, 5, 053018 (2011).
 26. P. C. DE VRIES, G. ARNOUX, A. HUBER, J. FLANAGAN, M. LEHNEN, V. RICCARDO, C. REUX, S. JACHMICH, C. LOWRY, G. CALABRO, D. FRIGIONE, M. TSALAS, N. HARTMANN, S. BREZINSEK, M. CLEVER, D. DOUAI, M. GROTH, T. C. HENDER, E. HODILLE, E. JOFFRIN, U. KRUEZI, G. F. MATTHEWS, J. MORRIS, R. NEU, V. PHILIPPS, G. SERGIENKO, and M. SERTOLI, "The impact of the ITER-like wall at JET on disruptions," *Plasma Physics and Controlled Fusion*, **54**, 12, 124032 (2012).
 27. D. PASINI, M. MATTIOLI, A. EDWARDS, R. GIANNELLA, R. GILL, N. HAWKES, G. MAGYAR, B. SAOUTIC, Z. WANG, and D. ZASCHE, "Impurity transport in JET using laser injected impurities in Ohmic and radiofrequency heated plasmas," *Nuclear Fusion*, **30**, 10, 2049 (1990).
 28. M. MATTIOLI, R. GIANNELLA, R. MYRNAS, C. DEMICHELIS, B. DENNE-HINNOV, T. D. D. WIT, and G. MAGYAR, "Laser blow-off injected impurity particle confinement times in JET and Tore Supra," *Nuclear Fusion*, **35**, 9, 1115 (1995).
 29. P. GALLI, A. CHERUBINI, R. D. ANGELIS, F. D. LUCA, M. ERBA, R. GIANNELLA, G. GORINI, A. JACCHIA, H. JÄDCKEL, P. MANTICA, V. PARAIL, L. PORTE, and A. TARONI, "Transient heat transport studies using laser ablated impurity injection in JET," *Nuclear Fusion*, **38**, 9, 1355 (1998).
 30. E. R. SOLANO, I. COFFEY, A. HUBER, S. S. HENDERSON, M. O'MULLANE, F. CASSON, H. SUMMERS, T. PÜTTERICH, M. E. PUIATTI, L. CARRARO, M. VALISA, S. MENMUIR, K. LAWSON, M. SERTOLI, B. LOMANOWSKI, M. REINKE, M. BERNERT, P. DREWELow, M. MASLOV, J. FLANAGAN, G. SERGIENKO, M. STAMP, S. BREZINSEK, L. MENESES, S. SCHMUCK, J. BOOM, C. GIROUD, C. GUILLEMAUT, E. DE LA LUNA, and A. BACIERO, "Characterising W radiation in JET-ILW plasmas," in "43rd EPS Conference on Plasma Physics," European Physical Society (2016), *Europhysics Conference Abstracts*, vol. 40A.