A neural network approach to real-time electron density reconstruction from alkali beam emission spectroscopy data at W7-X

B.Molnar[1], G.I. Pokol[1,2], A. Jalalvand[3], O. Asztalos [1,2], M. Karacsonyi [1,2], L. Armbruster[4], M. Vecsei[4], A.H.Nielsen[5] and W7-X Team[6]

[1] *HUN-REN Centre for Energy Research, Institute for Atomic Energy Research, Budapest, Hungary*

[2] *Budapest University of Technology and Economics, Budapest, Hungary*

[3] *Princeton University, Princeton, NJ, USA*

[4]*Max-Planck-Institut für Plasmaphysik, Greifswald, Germany*

[5]*Technical University of Denmark, Copenhagen, Denmark*

[6] *See "O. Grulke et al 2024 Nucl. Fusion 64 112002" for the W7-X Team*

Understanding edge plasma behavior in stellarators is essential for optimizing confinement and stability. Electron density profiles provide key information for studying turbulent transport, detachment and impurity dynamics. The alkali beam emission spectroscopy (ABES) is a powerful diagnostic for measuring density fluctuations in the edge plasma, offering a temporal resolution below 10 μs and a spatial resolution under 0.5 cm. The standard approach of reconstructing absolute electron density from ABES is the Bayesian inversion, that maximizes the a posteriori probability of a given density profile producing the observed light profile [1]. While accurate, this method is computationally expensive and unsuitable for real-time plasma control. To accelerate the reconstruction, a direct linearized reconstruction technique has been developed at W7-X [2], relying on a database lookup to first find an approximate density profile. However, this approach is limited by database completeness and is prone to noisy input. Machine learning methods can provide fast surrogate models for density reconstruction. Multilayer perceptrons (MLPs) have demonstrated promising results in predicting electron density profiles from ABES data, as shown in JET tokamak experiments [3]. Our previous studies [4], focusing on the forward problem (predicting light profiles from density data), confirmed that extreme learning machines (ELMs) can capture the underlying physics of the forward problem but exhibit high sensitivity to noise.

Here, we propose solving the inverse problem using convolutional neural networks (CNNs). The advantage of CNNs over MLPs and ELMs is that CNNs can exploit spatial dependencies in the data and are sufficiently complex to capture the inverse mapping from light to density profiles. A vanilla Convolutional Neural Network (CNN) was built for a regression task by stacking 3 convolutional layers with ReLu activations and using 20% dropout. The CNN architecture is shown in Fig. 1.: it takes a light profile as input, and outputs a density profile.

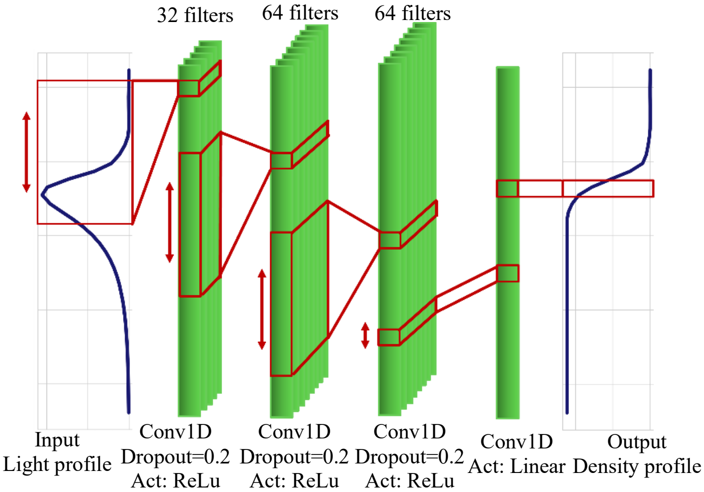

Figure 1. Architecture of the CNN used in this study. The network has a total of ~92000 trainable parameters.

For training and testing the network, a dataset of corresponding light and density profiles was assembled containing ~30.000 samples, generated from 211 density profiles via data augmentation. These profiles were collected from past reconstructions of W7-X experiments using the linearized inversion technique [2], representing 17 magnetic configurations of campaigns OP.2.2 and OP2.3. The collected dataset was augmented by adding an artificial bump at various locations with various widths and amplitudes. The resulting ~30.000 density profiles were then taken as inputs of forward simulations, and the corresponding light profiles were generated. The neural network was trained on the first 80% of shots from the campaigns and the rest was set aside for evaluation. The evaluation assessed the difference between predicted and true profiles and considered two metrics: per-profile and per-channel error. The Mean Absolute Scaled Error (MASE) describes the error per profile (and sums over channels):

$$\mathrm{MASE} = \frac{\frac{1}{M}\sum_{j=1}^{M}\left|y_j - \hat{y}_j\right|}{\frac{1}{M}\sum_{j=1}^{M}\left|y_j - \langle y_{train}\rangle_j\right|}$$

where M is the number of channels along the profile, y is the true profile, $\hat{y}$ is the predicted profile, $\langle y_{train}\rangle$ is the naive prediction – the average profile seen during training. The Mean Relative Error (MRE) describes the error per channel (and sums over profiles):

$$\mathrm{MRE} = \frac{1}{N}\sum_{i=1}^{N}\left|\frac{y_i - \hat{y}_i}{y_i}\right|$$

where N is the number of samples in the test data. The test set of ~6000 profiles was subject to evaluation, where the MASE and MRE metrics were calculated. The distribution of prediction errors over profiles was analyzed by plotting the histogram of MASE. This is shown in Fig. 2.: lower MASE means better prediction, while MASE around 1 would mean a prediction no better than a naïve guess (the average profile). We observed that MASE is well below 1, but a minor part of predictions is far worse than most predictions (note the logarithmic y-scale in Fig. 2.). These outlier profiles turned out to be originating from experiments where densities were in the order of $10^{20}m^{-3}$, while the network has been trained exclusively on densities below $10^{20}m^{-3}$, due to the temporal train/test split of the data. This means that the CNN model cannot generalize on out-of-sample data, it fails to reconstruct densities an order of magnitude higher than any density seen during training. This calls for a compilation of a comprehensive training data set to cover all density magnitudes. Another solution would be to alter the vanilla network architecture so it can generalize well on all absolute scales. Physics-informed neural networks [5] can be promising tools to tackle out-of-sample errors.

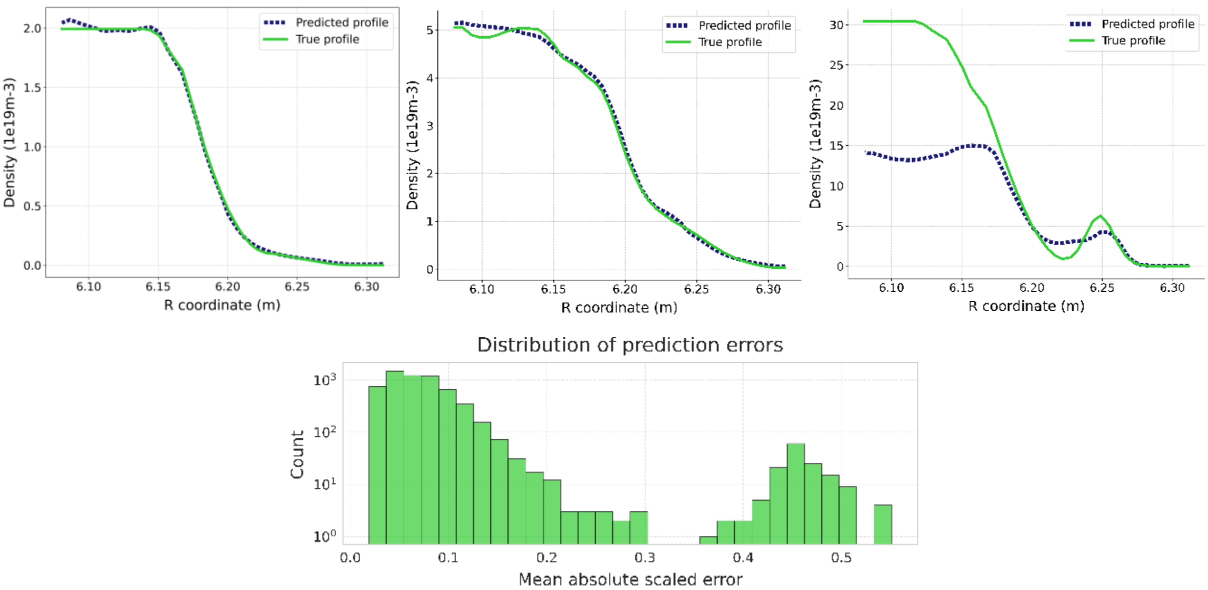

Figure 2. The histogram of per-profile errors (bottom) and predicted vs. true profiles of best to worst errors (top)

For further analysis, the out-of-sample data was excluded from the test set as to investigate the in-sample performance. The MRE, which indicates the error per channel, can be plotted as an average error profile and is shown in Fig. 3.

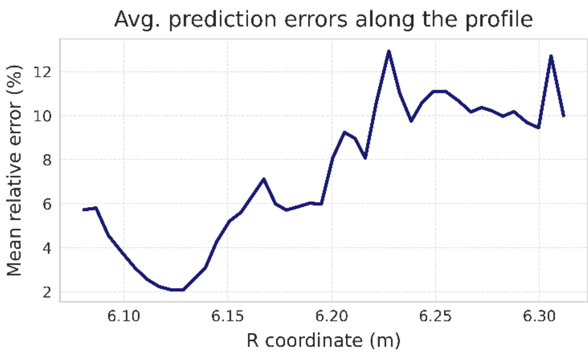

Figure 3. The average error profile of CNN predictions

The figure shows that the relative errors are in the range of 2-12% and are larger in the outer region and smaller in the inner region. That is due to the low absolute value of density in the region observed by the first few ABES channels. The errors are in the magnitude of those reported by a previous study on JET experiments [3] but are significantly better in the inner region. A comparison was made to MLPs used in the mentioned study [3], also ELMs, which were found to be successful for the forward problem [4] were considered for the comparison. The CNN was found to outperform both by a large margin, as seen by their per-profile (MASE) statistics on the test set in Table 1.

	ELM	MLP	CNN
Avg. MASE	0.228	0.277	**0.07**
Std. of MASE	0.203	0.115	**0.032**
Worst MASE	3.008	1.277	**0.295**
Best MASE	**0.014**	0.099	0.019

Table 1. Comparison of per-profile error (MASE) statistics of different architectures

Based on the promising results, the application of the CNN model at W7-X was considered. The current linearized Bayesian density reconstruction approach relies on a database lookup to find an approximate, initial density profile. This step could be replaced by the CNN solution if the predicted profiles were closer to the final solution than the profiles found by the lookup algorithm in the database. To investigate this, the error of the database lookup was calculated for the test set in terms of MASE, which gave the average and standard deviation of per-profile errors: 0.58 ± 0.78. From Table 1. we see that the CNN predicted profiles are much closer to the final solution (0.07 ± 0.032), thus the lookup method can indeed be replaced. Future work can target training a surrogate neural network model for real-time inference, substituting the whole reconstruction. The inference time using the current network is ~0.2 ms/profile, thus real-time targets can easily be achieved, but the issue of out-of-sample errors should be eliminated, and the network must be tested on measured light profiles and assess its robustness against noise. Development efforts should primarily target a compilation of an extensive set of measured data and leveraging physics-informed neural networks to cope with these issues.

References

[1] Fischer, R. et al. Plasma Physics and Controlled Fusion, 2008, 50 (8), 085009.

[2] Vécsei, M. et al. Review of Scientific Instruments, 2021, 92 (11), 113501.

[3] Pavone, A. et al. Plasma Physics and Controlled Fusion, 2020, 62(4), 045019.

[4] A. Jalalvand et al. 2023 International Joint Conference on Neural Networks (IJCNN), Gold Coast, Australia, 2023, pp. 1-8

[5] Yang et al., Journal of Computational Physics, 2021, 425 109913