Applying Deep Learning to GPI data to analyze plasma edge turbulence

Marion E. Smedberg[1], D. R. Smith[2], S. J. Zweben[3]

[1]Eindhoven U. of Technology, [2]U. of Wisconsin-Madison, [3]Princeton Plasma Physics Labs

Background

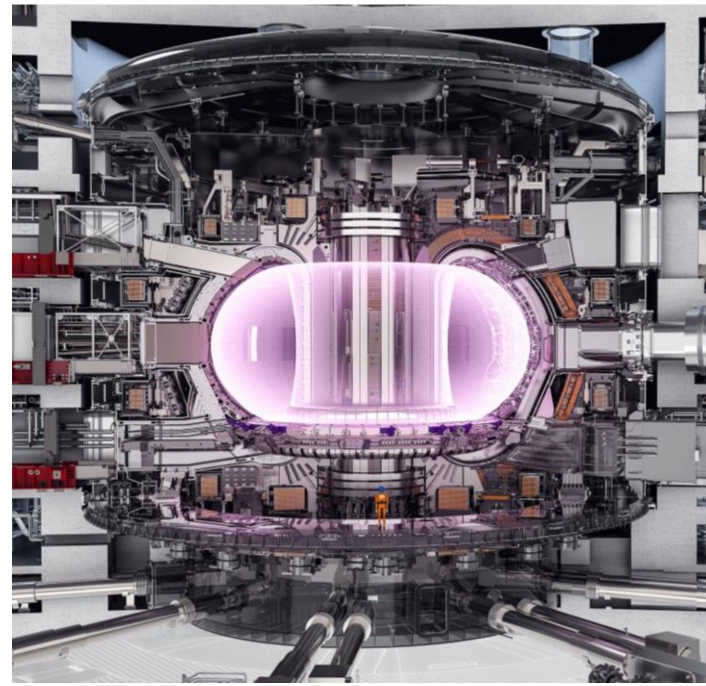

Understanding the edge dynamics of magnetically confined plasmas is an important aspect of increasing confinement and, eventually, putting fusion power on the grid, especially for large fusion machines such as ITER. Edge turbulence transports hot particles out of the plasma, both lowering the fusion power and heating up plasma-facing materials; therefore, **understanding and controlling edge turbulence is a critical part of creating smaller and/or cheaper machines.** This project used deep learning (DL) methods to analyze diagnostic data from the plasma edge of NSTX; **the goal was both to search for unidentified patterns and to lay the groundwork for future DL projects in plasma physics.**

The diagnostic: Gas Puff Imaging

In this project we analyzed images of the plasma edge of NSTX captured using a technique called <u>Gas Puff Imaging (GPI)</u> (Fig 1). In this method, neutral gas is puffed into a region of the plasma edge, where the neutral gas's excited electrons emit characteristic wavelengths of light. Those wavelengths are then filtered out and captured by a high-speed camera.

(a)

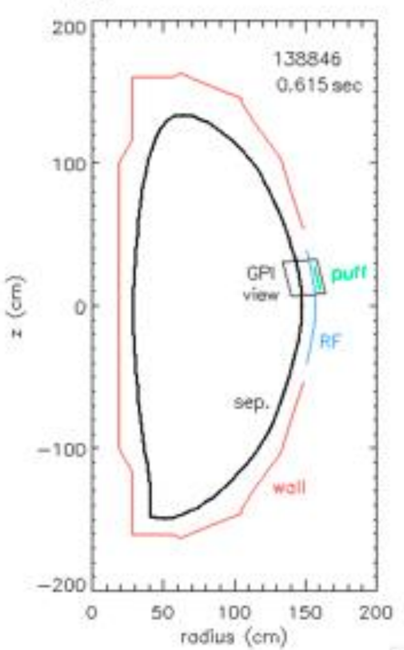

(b)

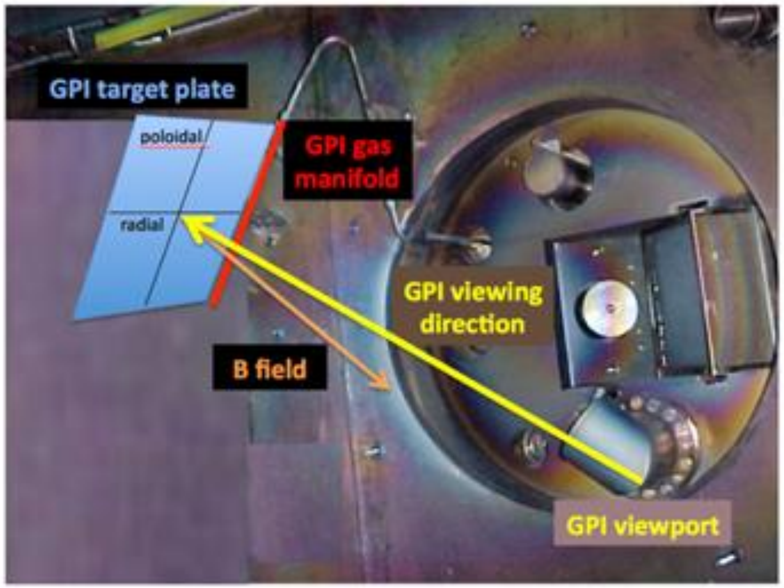

Figure 1: GPI hardware in NSTX. (a) is a typical radial vs. vertical cross-section of NSTX showing the GPI field of view. (b) shows the GPI hardware inside the vessel. Reproduced with permission from Zweben *et al.*, Phys Plasmas **24**, 102509 (2017).

Deep Learning (DL)

DL is a subset of machine learning (ML) usually used with images or data that is best represented as an array. Both ML and DL are being developed as techniques to efficiently analyze large amounts of data, which is an important capability to develop as scientific discovery becomes more and more data-heavy. In this project, we explored the capability of two deep learning frameworks (a Convolutional Neural Network or CNN, Fig 2, and an Autoencoder or AE, Fig 5) to classify or find patterns in GPI data.

Motivation

Compared to other STEM and physics fields, ML and DL are not widely used in plasma physics. This project aims to (1) explore utility of DL in fusion science, (2) efficiently identify patterns in large datasets, and (3) develop relationships between identified patterns and physics properties (e.g. L-mode or H-mode confinement). **Success of this project could encourage more use of these techniques in plasma computation and data analysis.**

Deep Learning (DL)

DL is a subset of machine learning (ML) usually used with images or data that is best represented as an array. Both ML and DL are being developed as techniques to efficiently analyze large amounts of data, which is an important capability to develop as scientific discovery becomes more and more data-heavy. In this project, we explored the capability of two deep learning frameworks (a Convolutional Neural Network or CNN, Fig 2, and an Autoencoder or AE, Fig 5) to classify or find patterns in GPI data.

Motivation

Compared to other STEM and physics fields, ML and DL are not widely used in plasma physics. This project aims to (1) explore utility of DL in fusion science, (2) efficiently identify patterns in large datasets, and (3) develop relationships between identified patterns and physics properties (e.g. L-mode or H-mode confinement). **Success of this project could encourage more use of these techniques in plasma computation and data analysis.**

Methods and Results: CNN

Figure 2: A graphic of the CNN's architecture

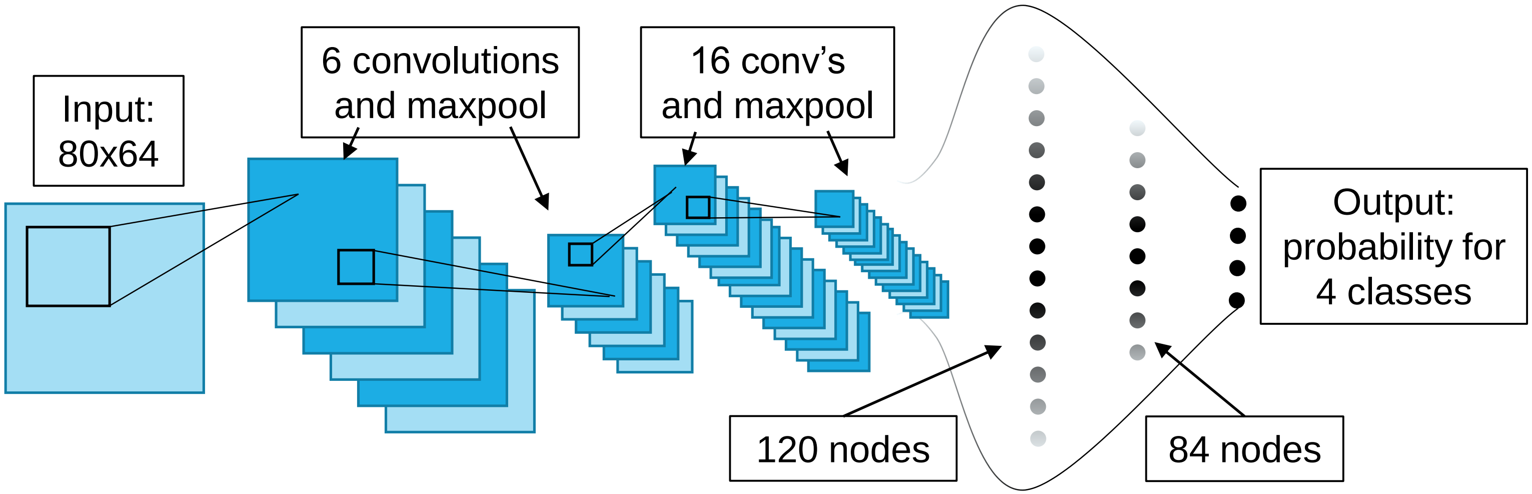

We used a CNN (Fig 2) to classify images into one of four classes (see pictures at right). Table 1 shows how changing functions or parameters in the network affected the classification accuracy. **Once we had a network that could accurately classify images, we analyzed how the network was doing this task,** i.e. what features or patterns the network was picking out that led to it assigning a particular class.

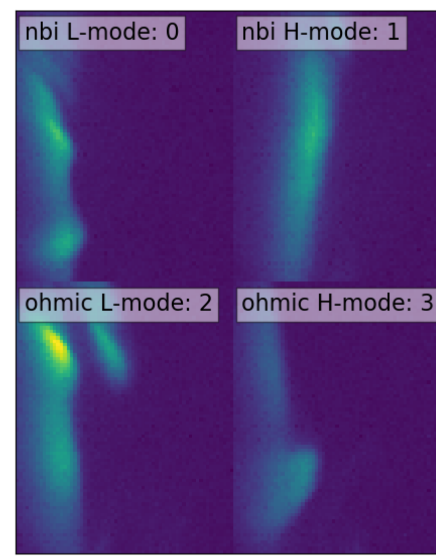

Training data size	Activation function	Optimizer	Batch size	% correct
20,000	Tanh	SGD	4	58
20,000	RELU	SGD	10	80
20,000	RELU	Adam	2	92
20,000	RELU	Adam	10	87
40,000	RELU	Adam	4	98

Table 1: Versions of the CNN, with classification accuracy. **We found that the choice of activation function, optimizer, batch size, and learning rate** (1e-4 for all examples shown here) **had a significant impact on the accuracy.**

Analyzing the CNN: Interpretability

After we trained the network to higher than 98% accuracy, we wanted to interpret the network: how was it classifying the images? We were particularly interested in what image features or combinations of features led to a particular class. **We found that**:

- the part of the image that was most important for classification varied. When we obscured (occluded) part of the image then ran it through the net, the class probabilities would change (Fig 3). **This technique seems to indicate that the network is sensitive to small scale/small amplitude features.**
- **certain nodes** in the "feature space" (last hidden layer, 84 nodes) **were only associated with one class and were very good at finding images with certain features** (Fig 4).

Figure 3 (right): Two examples of occlusion tests. For both (a) and (b), the top left is the original image and top right shows how much the correct label probability changed for a given occlusion location. Notably, the regions where the probability increased are very different; far outside the plasma in (a), but inside the plasma (especially covering the bright spot, bottom right image) for (b).

Figure 4 (below): Two examples of nodes that were only associated with one class. The images that maximally activated these nodes were all from one class and showed very similar visual characteristics.

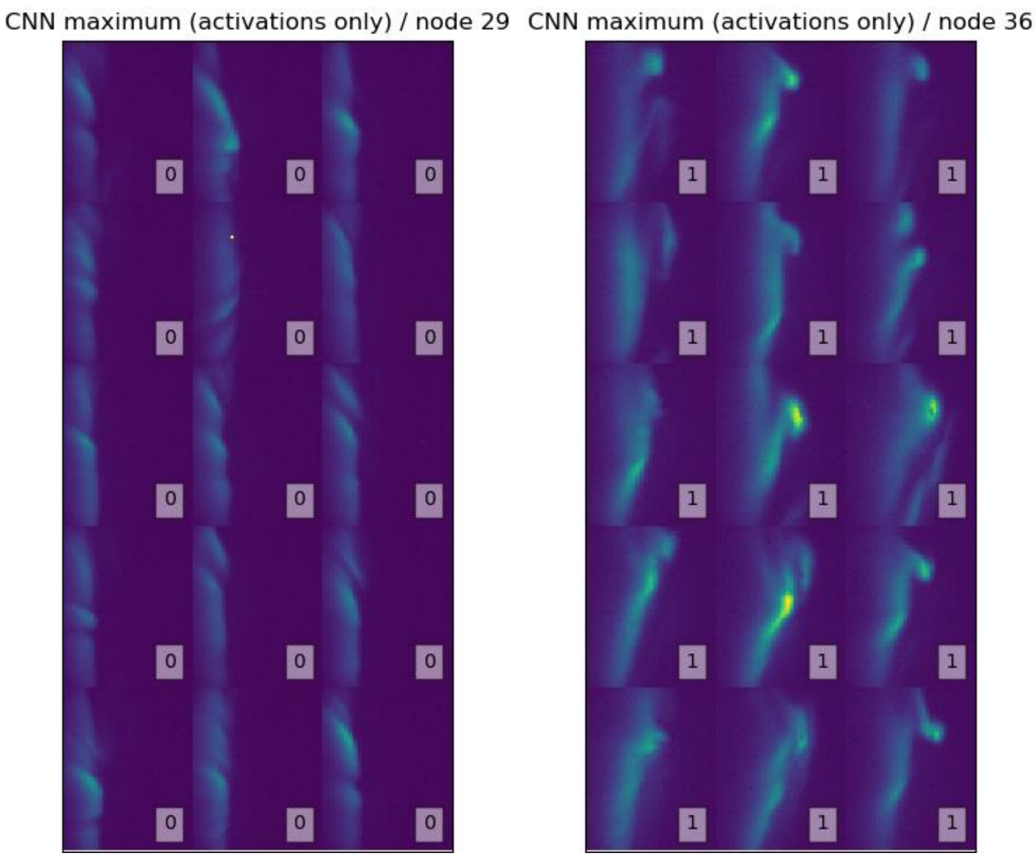

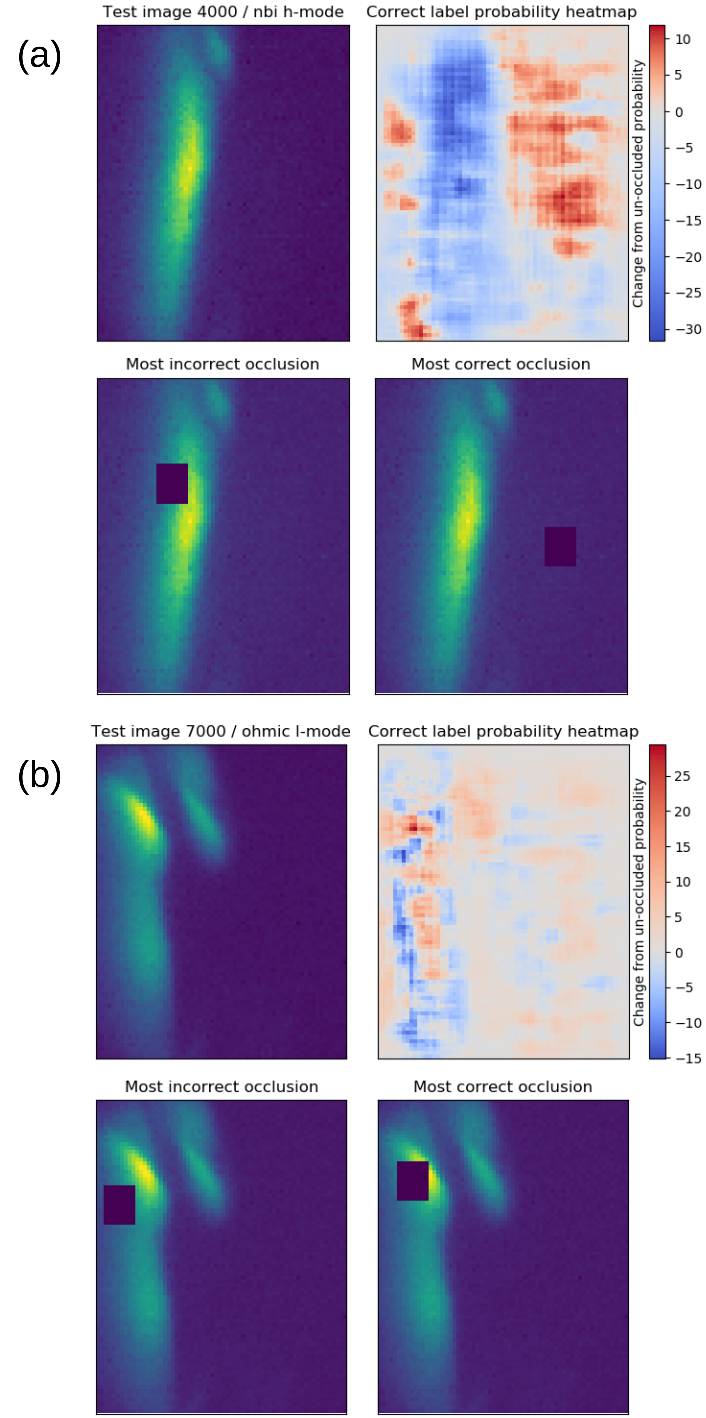

Potential Improvements

In the first set of 6 convolutions, 3 did not learn anything:

In the 84-node feature space, about 10 nodes never activated.

Both facts could be used to downsize the network, lowering training time with a minimal impact on accuracy.

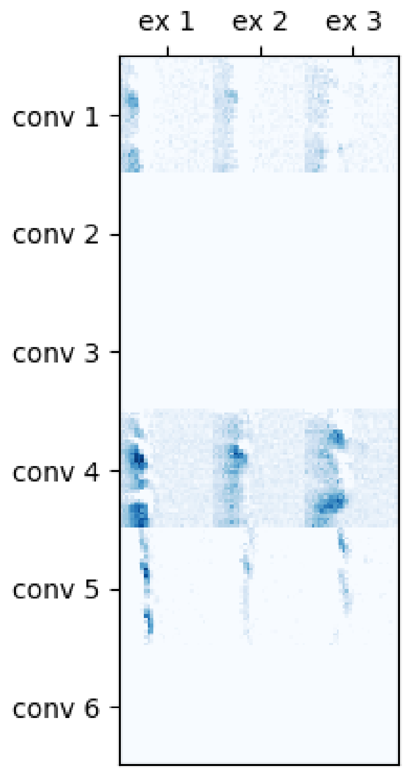

Future Questions

In the feature space, more nodes are weighted towards the h-mode classes than l-mode classes. What is the significance of this?

Would normalizing or transforming the data affect the training or accuracy?

How does the trained network do when given data from shots it was not trained on?

Methods and Results: Autoencoder

Figure 5: A graphic of the Autoencoder's architecture

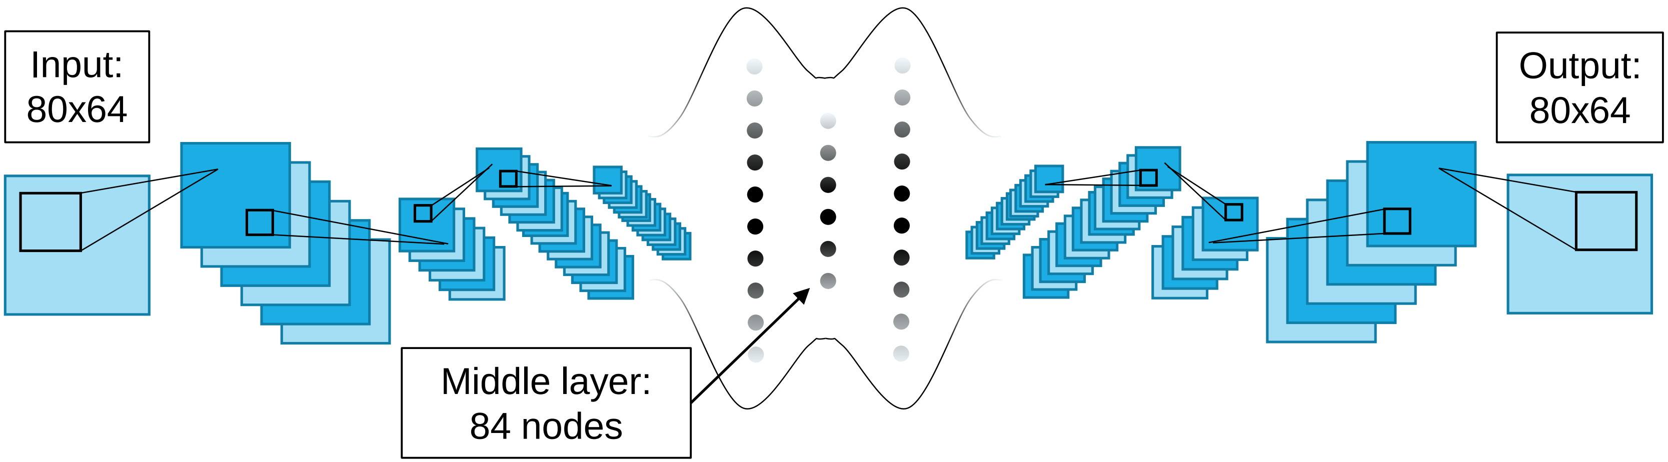

We used a different architecture, called an Autoencoder (AE) (Fig 5), to see whether the dimensionality of GPI images could be reduced. AEs are trained to reconstruct images from a lower dimensional space (the middle layer); we then analyze this layer to see what it learned. The idea is **if there are underlying patterns in the data, the AE will find them when reducing the dimensionality.**

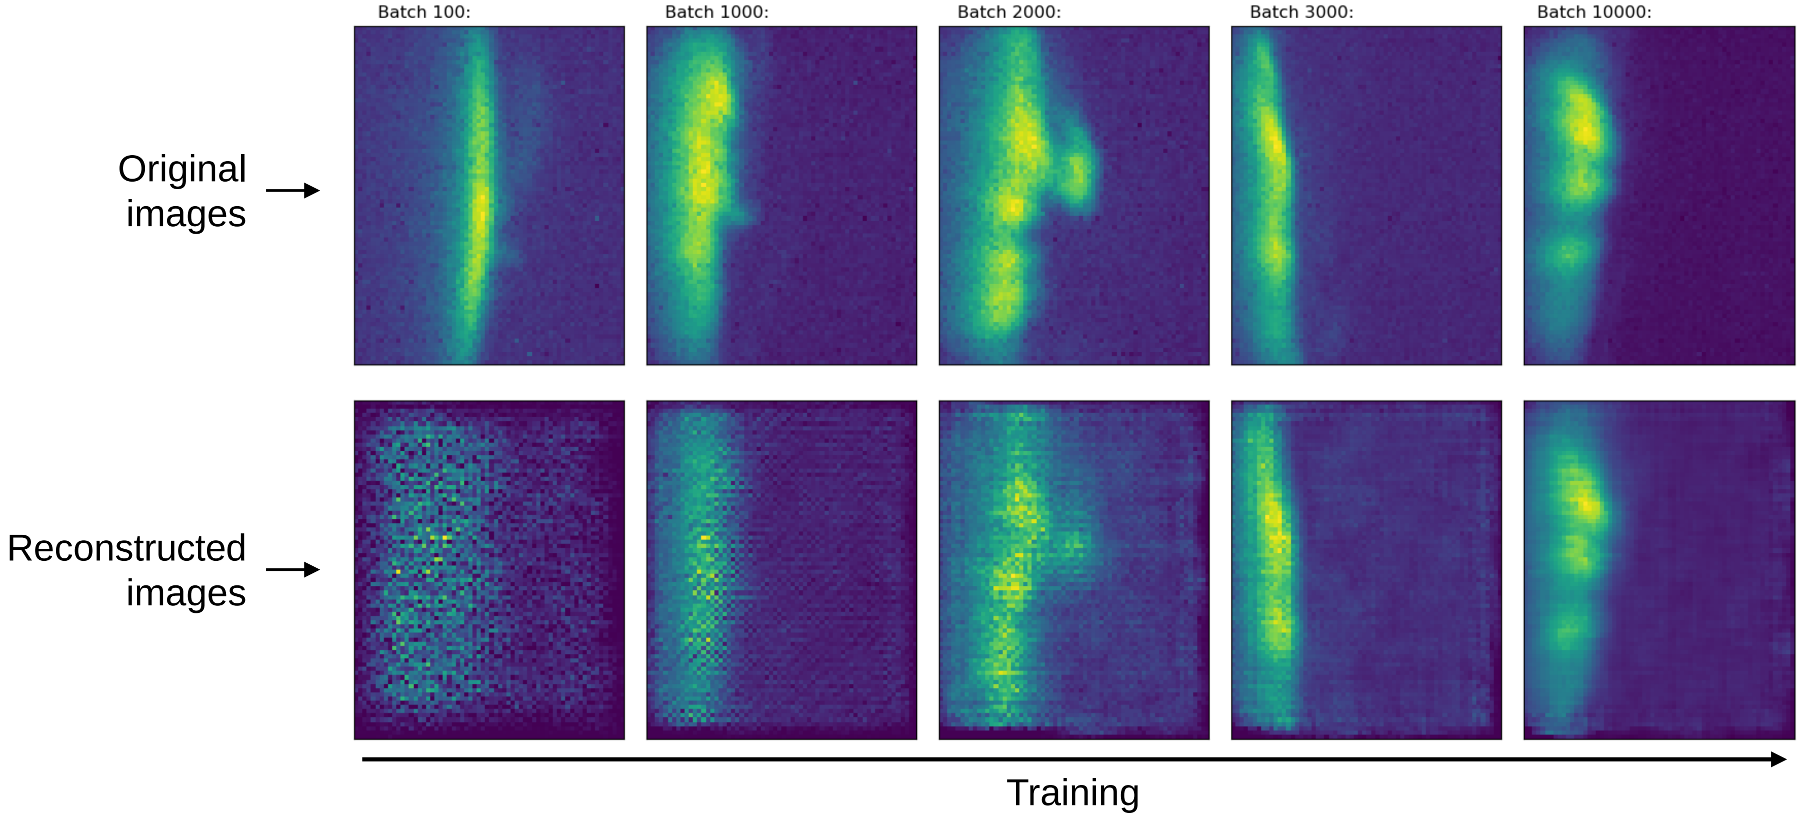

Figure 6: The AE's training progress, which proceeds left to right. The top row is original training images, while the bottom row is the AE's reconstructed versions, showing the network improving at its task of reconstructing images accurately.

By the final training batch (rightmost column in Fig 6), **the AE was able to reconstruct GPI images reasonably well**.

When we analyzed the middle layer, **a few preliminary outcomes stood out that contrasted with the CNN** (see Table 2).

Table 2: Contrasting the CNN and AE. **Significant differences probably arise because the two networks were trained differently**; the CNN was trained to classify images into categories given ground truth labels (supervised learning), whereas the AE was trained to reproduce images without any label information (unsupervised learning).

CNN	Autoencoder
Supervised learning	Unsupervised learning
10 inactive nodes in feature vector	20 inactive nodes in middle layer
Some feature vector nodes were associated with only one class	No middle layer nodes were associated with only one class
3 out of 6 convolutions were blank, convolved images grainy and picked out clear features	1 out of 6 convolutions were blank, convolved images smooth and did not pick out clear features

Potential Improvements

Similar to the CNN, blank convolutions and nodes in all layers that were always inactive could be downsized.

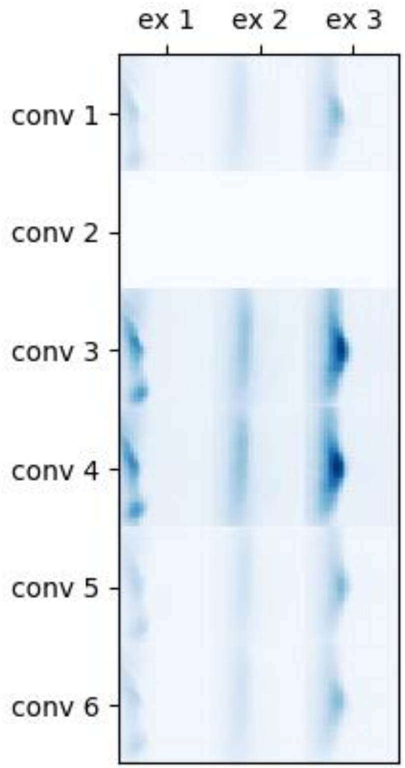

The middle layer especially should be downsized to see what affect that has on dimensionality and reconstruction.

Future Questions

Would normalizing or transforming the data affect the training or accuracy? What about using BCELoss?

What would happen if you took the reconstructed images from an AE and run them through a trained CNN?

What features would appear if we fed the AE an artificial middle layer with predetermined nodes activated?

Conclusion

This project was meant to be a demonstration of capability: can ML/DL techniques be useful in plasma physics as a data processing and analysis tool?

So far, we've been able to use DL to:

- accurately classify GPI image data into four relevant categories
- identify images with similar visual features
- pick up on small scale/small amplitude signals that would be difficult for a human to find
- reduce dimensionality of GPI images with low loss of information

Based on these results, in the future, DL might be used to:

- classify images into categories other than the ones used here (e.g. ELM, RF heating, etc.)
- suggest classifications for new or unlabeled data
- provide real-time information about plasma state
- inform theory or further analysis through what features/spatial locations a network learns is important

To match other STEM and physics fields, plasma physics can incorporate Machine Learning and Deep Learning into its computation and data analysis to automate certain processes, save time, or gain insight. With more research into interpretability, **ML/DL could be developed into a robust method in plasma physics for dealing with large datasets and problems that would be intractable without the aid of computers.**

Acknowledgements

I would like to thank the SULI program and PPPL for the chance to do this amazing internship; Deedee and Arturo for working so hard for us; Stewart for his information and support; and Dave for being an encouraging and attentive advisor.

About Me

In May, I graduated from Notre Dame of Maryland University in Baltimore with a B.A. in Physics. This fall I start a Masters degree in the “Science and Technology of Nuclear Fusion” at Eindhoven University of Technology in the Netherlands. I want to be a part of bringing clean energy from nuclear fusion to the grid.

Questions?

m.e.smedberg@student.tue.nl