

Seeking Higher-Order Modes in Free Electron Lasers

Peihao Sun* and Yanwen Sun†

*Department of Physics, Stanford University
SLAC National Accelerator Laboratory*

compiled: March 23, 2018

In free-electron lasers (FELs), the X-ray pulses are generated from spontaneous undulator radiation of electrons, which causes shot-to-shot fluctuations in the intensity, pointing, and spatial profile of the X-ray beam. In this project, we use deep neural networks to classify large amounts of X-ray images, which enables us to obtain statistical information of this intrinsically stochastic process. Both a supervised model and an unsupervised model are built, and the results are compared.

GitHub repo link

<https://github.com/phsunrise/higher-order-fel-modes>

Contribution

Yanwen Sun and Peihao Sun contributed equally to this work.

1. Introduction

The successful operations of free-electron laser (FEL) facilities across the globe [1, 2] have made it possible to produce extremely short X-ray pulses with unprecedented brightness: a typical pulse delivers more than 10^{11} coherent photons within 10^{-13} seconds. This has enabled many new X-ray tools to be used at the forefront of research in physics, chemistry, life science, material and energy sciences [3–10]. For many of these applications, it is vitally important to understand the X-ray beam profile and have the ability to control it. In particular, higher-order modes appear in some shots but not others [11] (see Fig. 1), and we would like to suppress (or enhance) these higher-order modes in order to address experimental needs.

Because FEL X-rays are generated by accelerated electron bunches from an undulator, we should in principle be able to control the X-ray profile by changing the electron bunch parameters. In practice, however, because the FEL generation process is inherently stochastic [12], the X-ray beam profile varies from pulse to pulse and is difficult to predict from electron bunch parameters using physical principles. Thus, an experiment was carried out where

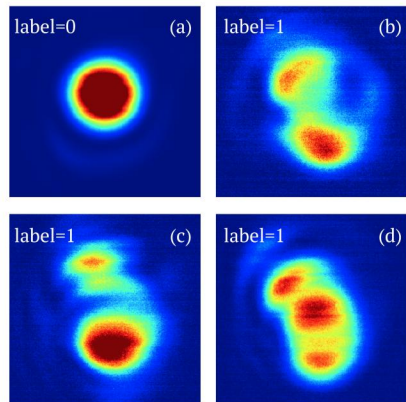


Fig. 1: Examples of beam profiles. (a) 0th-order mode, labeled “0”. (b)-(d) higher-order modes, labeled “1”.

a large amount of X-ray pulses were recorded along with various electron bunch parameters, with the aim to understand the correlations between these parameters and the X-ray beam profiles. To achieve this goal, one major task is to classify the X-ray profile images. However, this classification is hard to define mathematically, especially because FEL pulses very often deviate from ideal Gaussian modes. Therefore, our project aims to use deep neural networks to classify the X-ray images. The project is divided into two parts: (1) a supervised learning part, where the inputs are $100 \times 100 \times 1$ X-ray beam profiles, and a convolutional neural network (CNN) is trained to categorize them as “0” or “1” (see Fig. 1); (2) an unsupervised learning part, where we use an autoencoder neural network trained on the same input to learn a 32-dimensional representation of the images. We compare results from the two ap-

* phsun@stanford.edu

† yws@stanford.edu

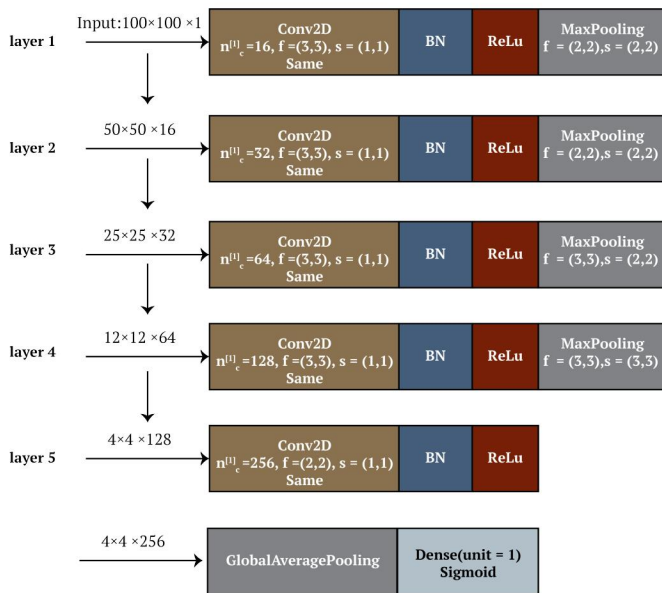


Fig. 2: Supervised learning CNN architecture.




Fig. 3: Loss history for the training and dev sets in supervised CNN training. The brown dashed line indicates the position of early stopping.

proaches, and use them to understand the influence of various external parameters.

2. Dataset

Our dataset comes from Experiment X227 at the Linac Coherent Light Source (LCLS), an FEL facility at SLAC National Accelerator Laboratory. During this experiment, more than one million X-ray profiles were recorded along with various external parameters, such as electron bunch charge, electron beam energy, undulator length, etc. The X-ray profiles are originally $1080 \times 1280 \times 1$ grayscale images denoting the X-ray intensity at each pixel, but are cropped and downsized to $100 \times 100 \times 1$ in order to expedite learning. The downsized images are used as the input X .




Fig. 4: Examples of ambiguous images which are hard to label as “0” or “1”.

3. Supervised learning

3.A. CNN architecture and training

Our goal for supervised learning is to train a neural network to distinguish 0th-order beam profiles, labeled as “0”, from higher-order ones, labeled as “1” (see Fig. 1). We manually classify more than 4,000 images into these two categories, and choose 2,000 0th-order and 2,000 higher-order mode profiles as our training data. Of these 4,000 images, 400 are selected as the dev set. The test set consists of 200 images, with 100 images in each class.

Our neural network architecture, presented in Fig. 2, is inspired by the VGG neural network [13]. However, we have simplified the original VGG network because our image size is smaller and more importantly, because our images contain simpler features. This facilitates learning and allows us to train the model within a few minutes. The modified CNN contains six weighted layers, including five Conv2D layers and one final fully-connected layer before applying a sigmoid activation. Adam gradient descent and cross entropy loss are used for the optimization process with a mini batch size of 32. The training is done using the Keras API.

The training history is shown in Fig. 3. Based on the dev set loss, we apply early stopping at epoch 27 (indicated by the brown dashed line). The model trained at the end of this epoch has a train/dev/test set accuracy of 92%. This is close to the human labeling accuracy estimated in the following way: we both labeled the same $N = 857$ images, and found that $n_d = 74$ of them were labeled differently. Thus, we estimate our human error to be around $n_d/N \approx 8\%$. Careful examination of these 74 images leads to the conclusion that around 60 of them are difficult to label as 0 or 1. Examples of such images are shown in Fig. 4, where they either contain a barely visible side lobe or have a central peak which deviates from ideal Gaussian shapes. We thus estimate that around 7% of the images are ambiguous when labeling with human vision.




Fig. 5: Percentage of the higher-order mode under different beam parameters.

3.B. Results and discussion

With the trained model, we are able to calculate the percentage of higher-order modes (label “1”) under different electron/X-ray beam conditions. The results are plotted in Fig. 5, where the vertical axis shows higher-order mode percentage and the horizontal axis denotes electron bunch charges. The four different colors represent four silicon crystal reflections represented by the parenthesized numbers Si(hkl) (also known as “Miller indices”). The order of crystal reflection can be calculated as $\sqrt{h^2 + k^2 + l^2}$, and the larger this number is, the smaller the bandwidth of the X-ray photon energy.

From Fig. 5 we can see that, in general, higher-order X-ray modes appear more frequently in higher-order crystal reflections. Since increasing hkl means decreasing window size on the X-ray bandwidth, our results suggest that the 0th-order beam profiles from lower-order crystal reflections may actually be the sum of higher-order modes lasing at different X-ray energies within the bandwidth. This is also confirmed by the decrease of the percentage of higher-order modes with the increase of the bunch charge using Si(800) reflection. However, the opposite is observed when we select a broader bandwidth with Si(220), which suggests that some round-shaped beams, especially those with low bunch charge (80 pC) and high bandwidth (Si(220)), can be the result of real 0th-order Gaussian modes dominating at lasing saturation.

Although we can already draw several conclusions from the results above, there exist a few drawbacks to this supervised training method. As mentioned before, our human error in labeling these images is

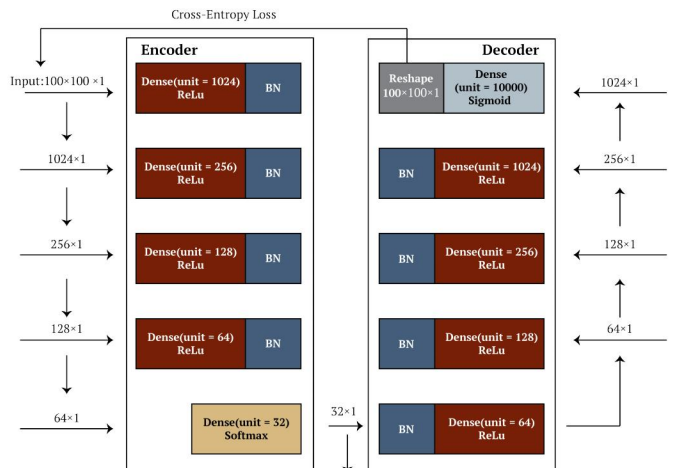


Fig. 6: Autoencoder neural network architecture.





Fig. 7: Examples of the input and output images of the autoencoder.

relatively large—around 7% to 8%. Moreover, since we have a large dataset containing tens of thousands of images, manual labeling becomes rather tedious for the whole dataset. In addition, although humans may be good at distinguishing the 0th-order mode from higher-order modes, we are not so good at classifying among the latter. In particular, because multiple peaks appear in these higher-order modes, the relative position and intensity of the peaks may contain significant information but are difficult to determine and classify using human vision. All of the reasons above point to the use of unsupervised learning, which constitutes the second part of our project.


4. Unsupervised learning

4.A. Autoencoder architecture and training


As discussed above, unsupervised learning is a great tool to address the ambiguity in labeling and utilize the large available (unlabeled) dataset. It allows for automatic learning of image feature extraction and significantly reduces the data dimension, thus facilitating further analysis. Inspired by [14], we use an autoencoder network for unsupervised learning. The architecture is shown in Fig. 6. The neural network can be separated into two parts: (1) the



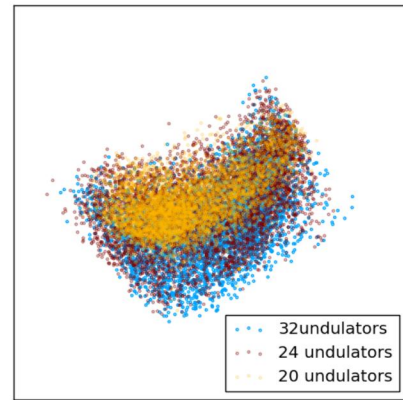
(a) PCA results showing 2D embedding of the 32-dimensional vector “codes”.



(a) Density map of beam profiles at different bunch charges.



(b) Beam profiles picked randomly from each bin on the 2D map.



(b) Density map of beam profiles at different number of undulators.

Fig. 8: Visualization of beam profiles on the 2D map.

Fig. 9: Density maps of beam profiles for various electron bunch parameters. Each dot represents one beam profile, and 4800 beam profiles are used for each condition.

“encoder” part, which uses several fully-connected layers to gradually reduce the ($100 \times 100 \times 1 =$) 10,000-dimensional input to a 32-dimensional vector “code”, and (2) the “decoder” part, which mirrors the “encoder” and gradually increases the dimension to 10,000. When training this model, we force the output to be the same as the input. After training, we use the trained weights to obtain the values in the “code” layer for each input image. The extracted features of each of the 32 codes are shown in A.3.

The autoencoder network is trained on all 76,800 images, with a minibatch size of 64. We find that training one epoch is sufficient to obtain desired results.

4.B. Results

To visualize the 32-dimensional “code” for each beam profile, we use principal component analysis (PCA) to further reduce its dimension to two. In order to understand this 2D map, we show in Fig. 8a the distribution of our pre-labelled data used in supervised training. Because we are plotting equal amounts of data (2000 each) with label “0” and “1”, the density distribution reflects that higher-order-mode images cluster in the upper left corner of the 2D space. This is confirmed by dividing the 2D map into smaller bins and observing that images

which fall into the same bin share similarities, while images from different bins contain different features. One example is shown in Fig. 8b (see Fig. A.1 for 20 more examples), where we plot a randomly picked image from each bin. This demonstrates that our autoencoder network has learned to recognize the important beam profile features, such as beam size, the number of lobes, intensity level, beam center location, orientations, etc.

Now that we know how these features are encoded in the 2D mapping, we can obtain more information on the X-ray beam profile than a mere percentage in Fig. 5. We illustrate in Fig. 9 two such examples, where Fig. 9a uses the same profile data as the cardinal red bar in Fig 5. While we have already obtained from supervised learning that the percentage of higher-order modes drops with increasing number of undulators (33%, 28%, 15% for 20, 24, 32 undulators, respectively), the density plot in Fig 9b provides further information: the dots are inclined to scatter downwards with the increase of undulators, suggesting that the y axis of this 2D space (with positive direction pointing down) corresponds approximately to the time axis of the lasing process.

5. Conclusion and future work

Both supervised and unsupervised learning are used to categorize X-ray beam profiles. For supervised learning, a 6-weight-layer CNN is trained on 4,000 human-labeled images to separate 0th-order beam profiles from higher-order ones. The trained model is then used to predict the percentage of higher-order modes when different electron bunch parameters and crystal reflection orders (X-ray bandwidth) are selected. We find that higher-order beam profiles appear more frequently when lower bandwidth is selected. However, opposite trends are observed with fixed bandwidth but varying electron bunch charges.

For unsupervised learning, an autoencoder network is built, where the hidden “bottleneck” (or “code”) layer outputs a 32-dimensional code representing each input image. PCA is then used to visualize these codes on a 2D map. We find that positions on this map indeed encode features of the beam profiles. This allows us to extract more information from the X-ray beam images, which can be used for more sophisticated FEL analysis. Compared with supervised learning, our unsupervised learning model overcomes the difficulties in human labeling, and allows for easy transfer of the same architecture to similar analysis tasks.

However, we also note that this project is not yet complete. Although our autoencoder network is

able to find 32-dimensional representations of the images, it may be possible to further reduce the dimension with careful optimization of the architecture and evaluation of the interpretability of the “codes”. In addition, although we are able to draw qualitative conclusions from the 2D mapping of images, more work needs to be done to analyze the codes quantitatively.

References

- [1] P. Emma *et al.*, *Nature* **4**, 641 (2010).
- [2] T. Ishikawa *et al.*, *Nature Photonics* **6**, 540544 (2012).
- [3] L. Young, E. Kanter, B. Krässig, Y. Li, A. March, S. Pratt, R. Santra, S. Southworth, N. Rohringer, L. DiMauro, *et al.*, *Nature* **466**, 56 (2010).
- [4] T. Glover, D. Fritz, M. Cammarata, T. Allison, S. Coh, J. Feldkamp, H. Lemke, D. Zhu, Y. Feng, R. Coffee, *et al.*, *Nature* **488**, 603 (2012).
- [5] M. Fuchs, M. Trigo, J. Chen, S. Ghimire, S. Shwartz, M. Kozina, M. Jiang, T. Henighan, C. Bray, G. Ndabashimiye, *et al.*, *Nature Physics* **11**, 964 (2015).
- [6] M. Minitti, J. Budarz, A. Kirrander, J. Robinson, D. Ratner, T. Lane, D. Zhu, J. Glowia, M. Kozina, H. Lemke, *et al.*, *Physical review letters* **114**, 255501 (2015).
- [7] H. N. Chapman, P. Fromme, A. Barty, T. A. White, R. A. Kirian, A. Aquila, M. S. Hunter, J. Schulz, D. P. DePonte, U. Weierstall, *et al.*, *Nature* **470**, 73 (2011).
- [8] M. M. Seibert, T. Ekeberg, F. R. Maia, M. Svenda, J. Andreasson, O. Jönsson, D. Odić, B. Iwan, A. Rocker, D. Westphal, *et al.*, *Nature* **470**, 78 (2011).
- [9] S. Gerber, H. Jang, H. Nojiri, S. Matsuzawa, H. Yasumura, D. Bonn, R. Liang, W. Hardy, Z. Islam, A. Mehta, *et al.*, *Science* **350**, 949 (2015).
- [10] M. Harmand, A. Ravasio, S. Mazevet, J. Bouchet, A. Denoëud, F. Dorchies, Y. Feng, C. Fourment, E. Galtier, J. Gaudin, *et al.*, *Physical Review B* **92**, 024108 (2015).
- [11] M. Xie, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **445**, 59 (2000).
- [12] A. Kondratenko and E. Saldin, *Part. Accel.* **10**, 207 (1980).
- [13] K. Simonyan and A. Zisserman, *ArXiv e-prints* (2014), arXiv:1409.1556 [cs.CV].
- [14] G. E. Hinton and R. R. Salakhutdinov, **313**, 504 (2006).

Appendix A




Fig. A.1: 20 examples of beam profiles picked randomly from each bin on the 2D map.




Fig. A.2: Loss history in autoencoder training.

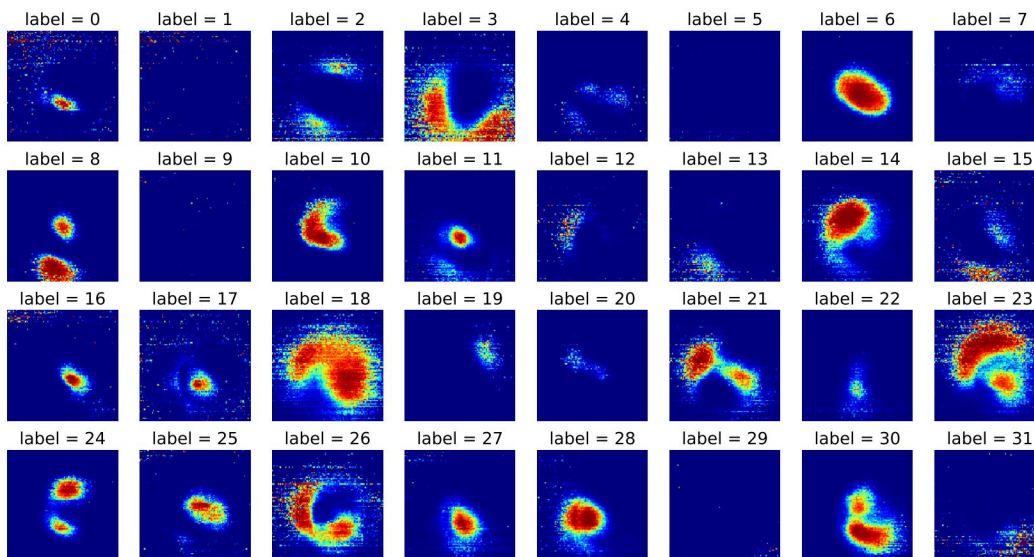


Fig. A.3: Visualization of each component of the 32-dimensional code.