Rapid peak detection for diffraction images

Po-Nan Li
Dept. of Electrical Engineering, Stanford & Biosciences Div., SLAC
liponan@stanford.edu

Motivation

- Protein structure determination relies on accurate diffraction peak positions
- The “peak finding” algorithm currently being used (Droplet) takes ~1.4 seconds to process a single image; while the data are generated at a rate of 65Hz.

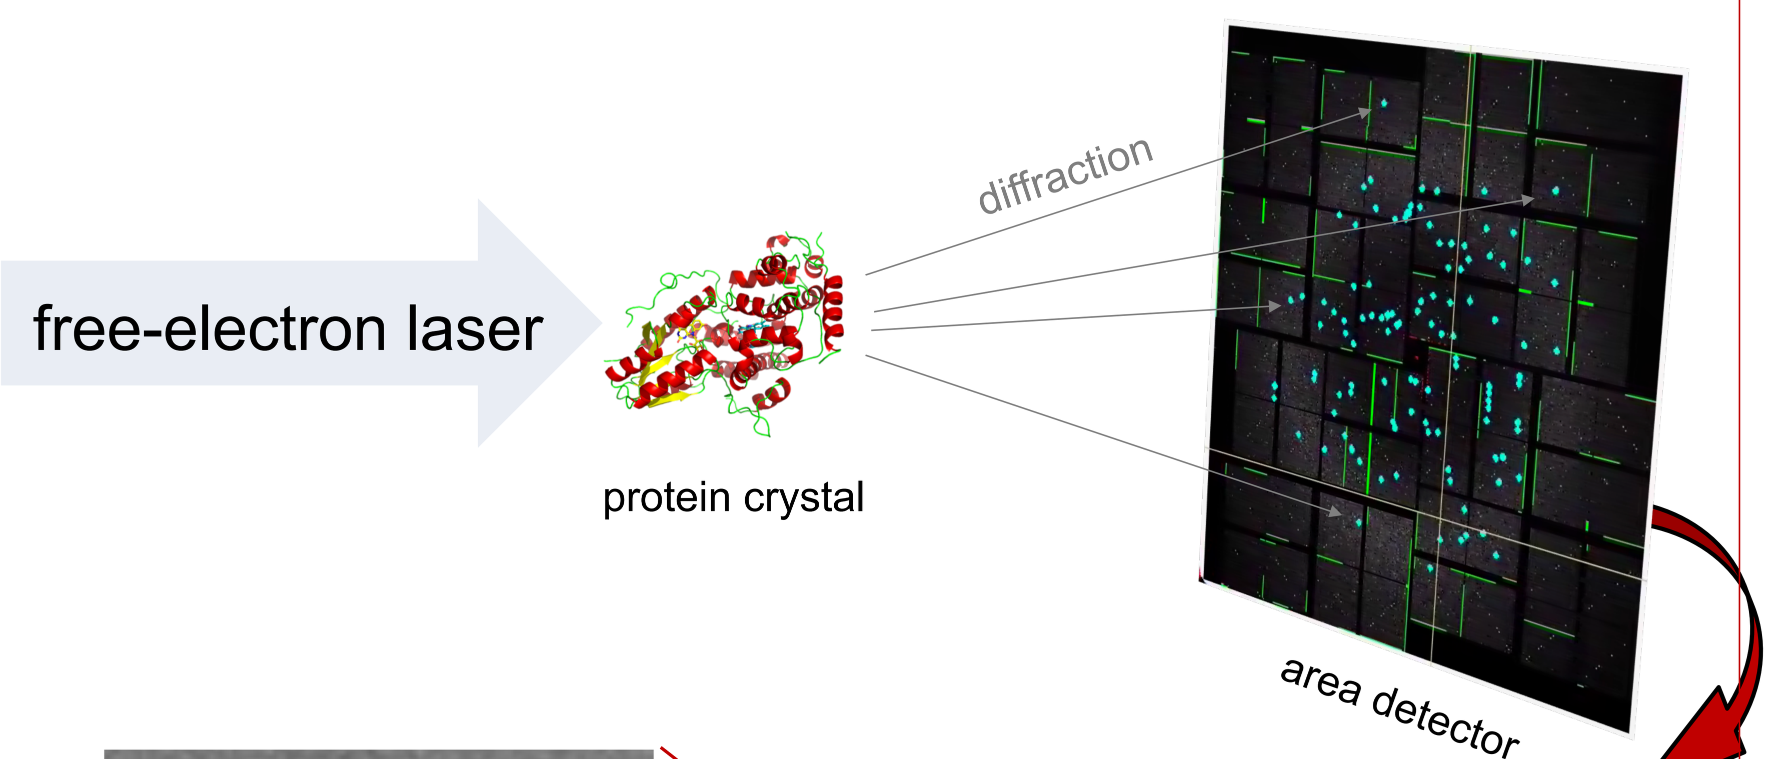

A “peak” is made of only a few bright pixels.

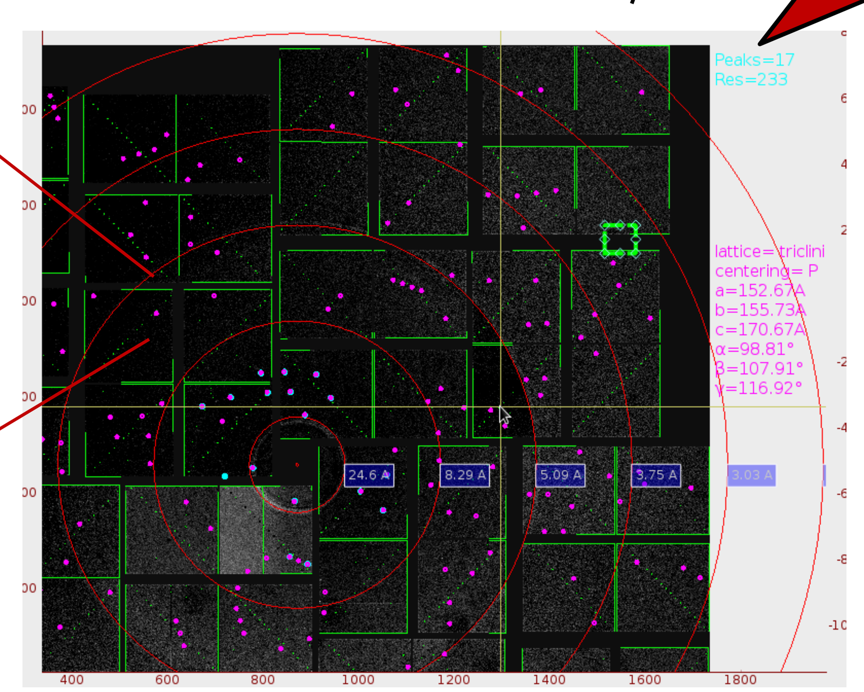

Crystal parameters can be retrieved from precise peak coordinates

Data

Diffraction images from CXILP9515, a serial femtosecond crystallography (SFX) [1] experiment conducted at LCLS’s CXI beamline, are used, including

- 137,076 images for training, and
- 5,188 images for validation

Each sample is a 1552x1480 gray-scale image, which is binned to 776x740 to fit the memory of GPU.
Labels are peaks that were previously found by Droplet.

Method

A peak detection pipeline, namely PSNET, is built, which is based off *darknet*, an open-source convolutional neural network (CNN) framework [2] and *YOLO*, a fast object detection system [3,4]. The net is trained with 137K diffraction images with labels of peak coordinates and a variable box size.

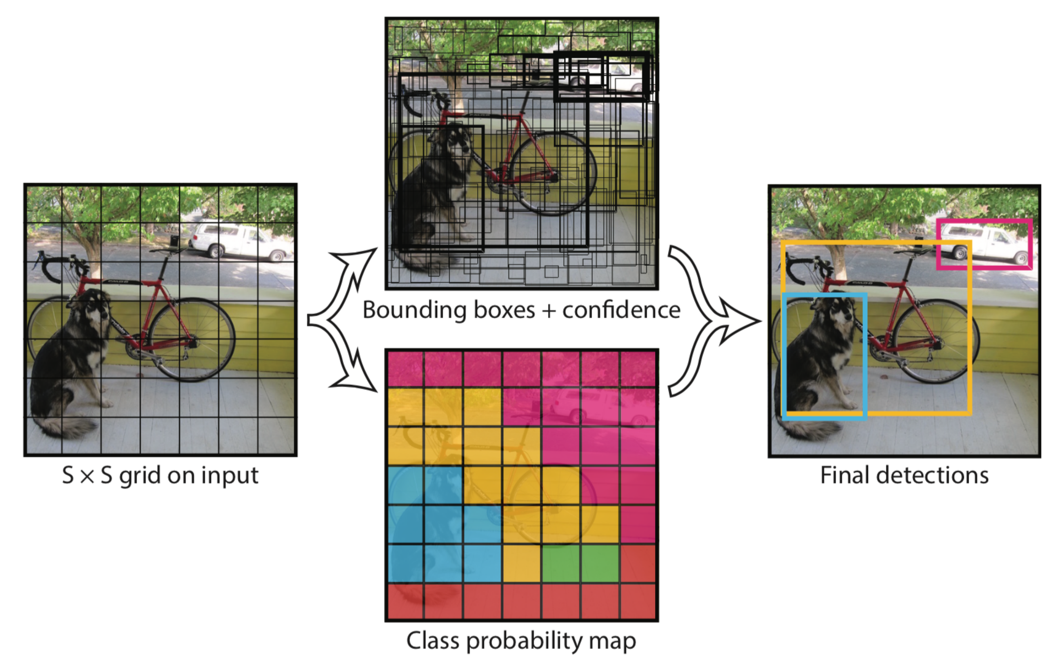

YOLO divides an image into grid cells, where the probability of existence of an object is evaluated. The position is treated as a regression problem [3].

Type	Filters	Size/Stride
Convolutional	32	3 × 3
Maxpool		2 × 2/2
Convolutional	64	3 × 3
Maxpool		2 × 2/2
Convolutional	128	3 × 3
Convolutional	64	1 × 1
Convolutional	128	3 × 3
Maxpool		2 × 2/2
Convolutional	256	3 × 3
Convolutional	128	1 × 1
Convolutional	256	3 × 3
Maxpool		2 × 2/2
Convolutional	512	3 × 3
Convolutional	256	1 × 1
Convolutional	512	3 × 3
Convolutional	256	1 × 1
Convolutional	512	3 × 3
Maxpool		2 × 2/2
Convolutional	1024	3 × 3
Convolutional	512	1 × 1
Convolutional	1024	3 × 3
Convolutional	512	1 × 1
Convolutional	1024	3 × 3
Convolutional	1000	1 × 1
Avgpool		Global
Softmax		

Architecture of CNN used in this project.

Results

PSNET, a pipeline dedicated for diffraction peak detection for crystallography, returns locations of peaks found in an image in less than 200ms.

```
peaksLocations = PSNET( expID, runID, imgID, thresh )
```

The PSNET was trained with two kinds of box size, 51x51 and 11x11. It’s astounding that PSNET is capable of detecting even very small/faint peaks. Location accuracy needs to be improved.

Box size	Dataset	# Peaks	Recall	*Avg. IOU.	*Avg. Dist.
51x51	Training	1,472,839	0.81	0.32	24.44 px
	Validation	186,438	0.94	0.43	18.06 px
11x11	Training	1,472,839	0.10	0.18	7.44 px
	Validation	186,438	0.30	0.19	7.70 px

IOU: intersection over union. Dist.: center-to-center distance. *Conditioned on positives.

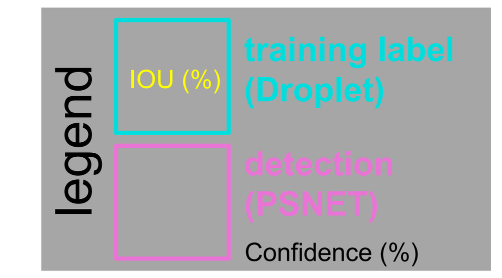

Trained with 51x51 label boxes

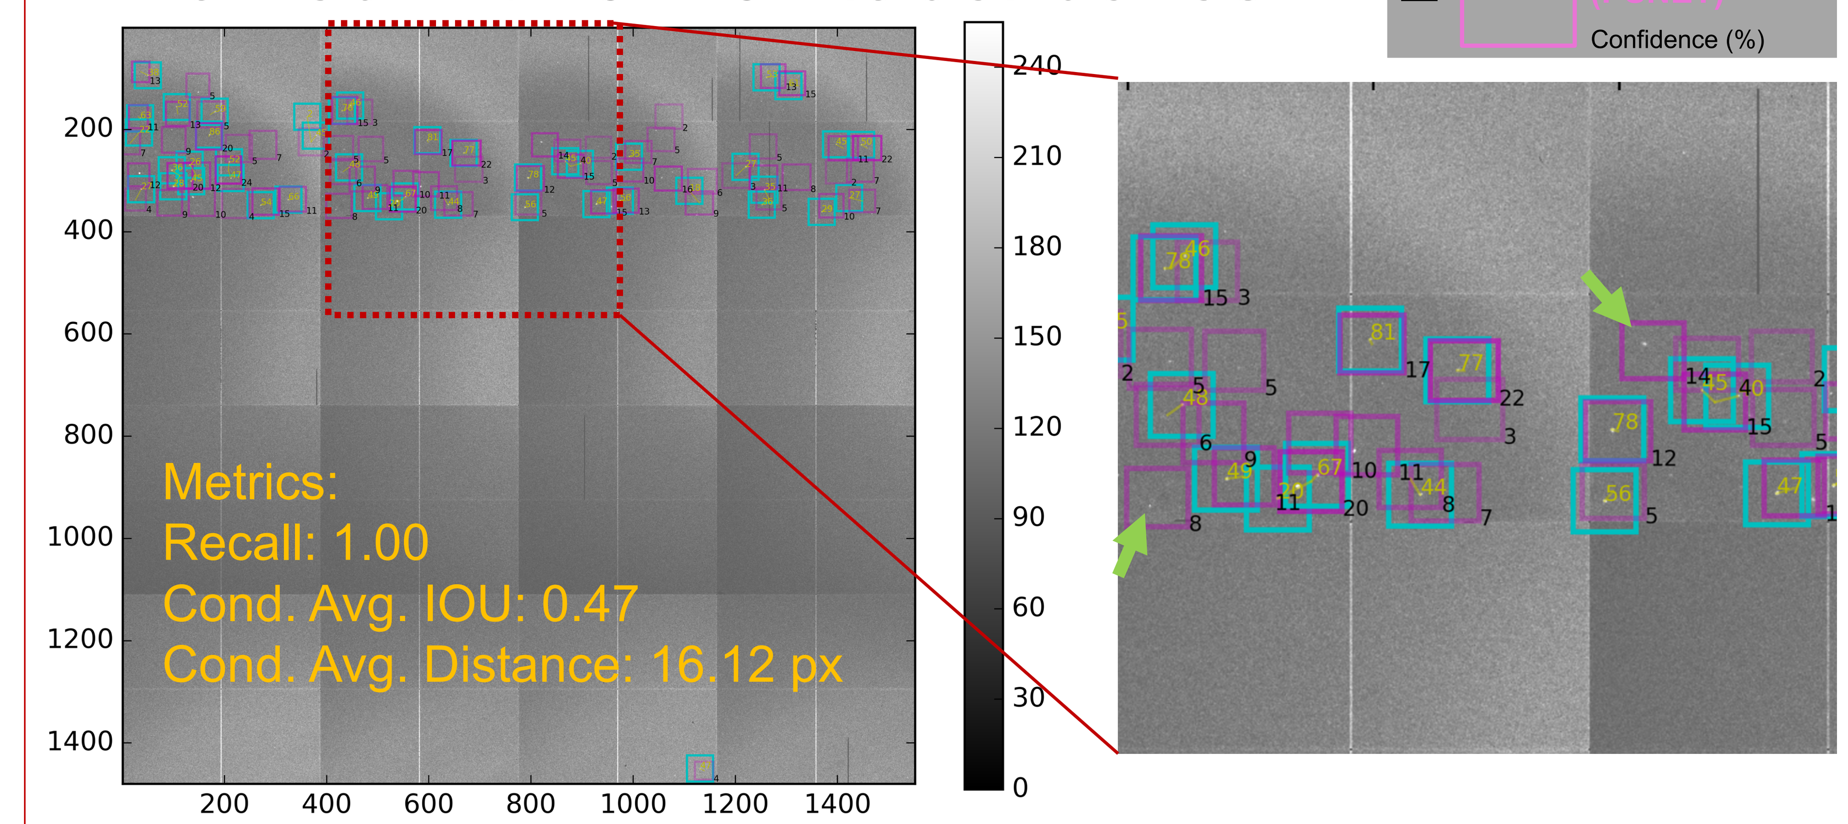

A “peak” is usually only a few pixels big, which makes it very difficult to be trained and detected. We started by intentionally using relatively large label boxes. Most of the labelled peaks were detected; the PSNET even found quite a few true peaks that were not found by Droplet previously (indicated by green arrows).

Trained with 11x11 label boxes

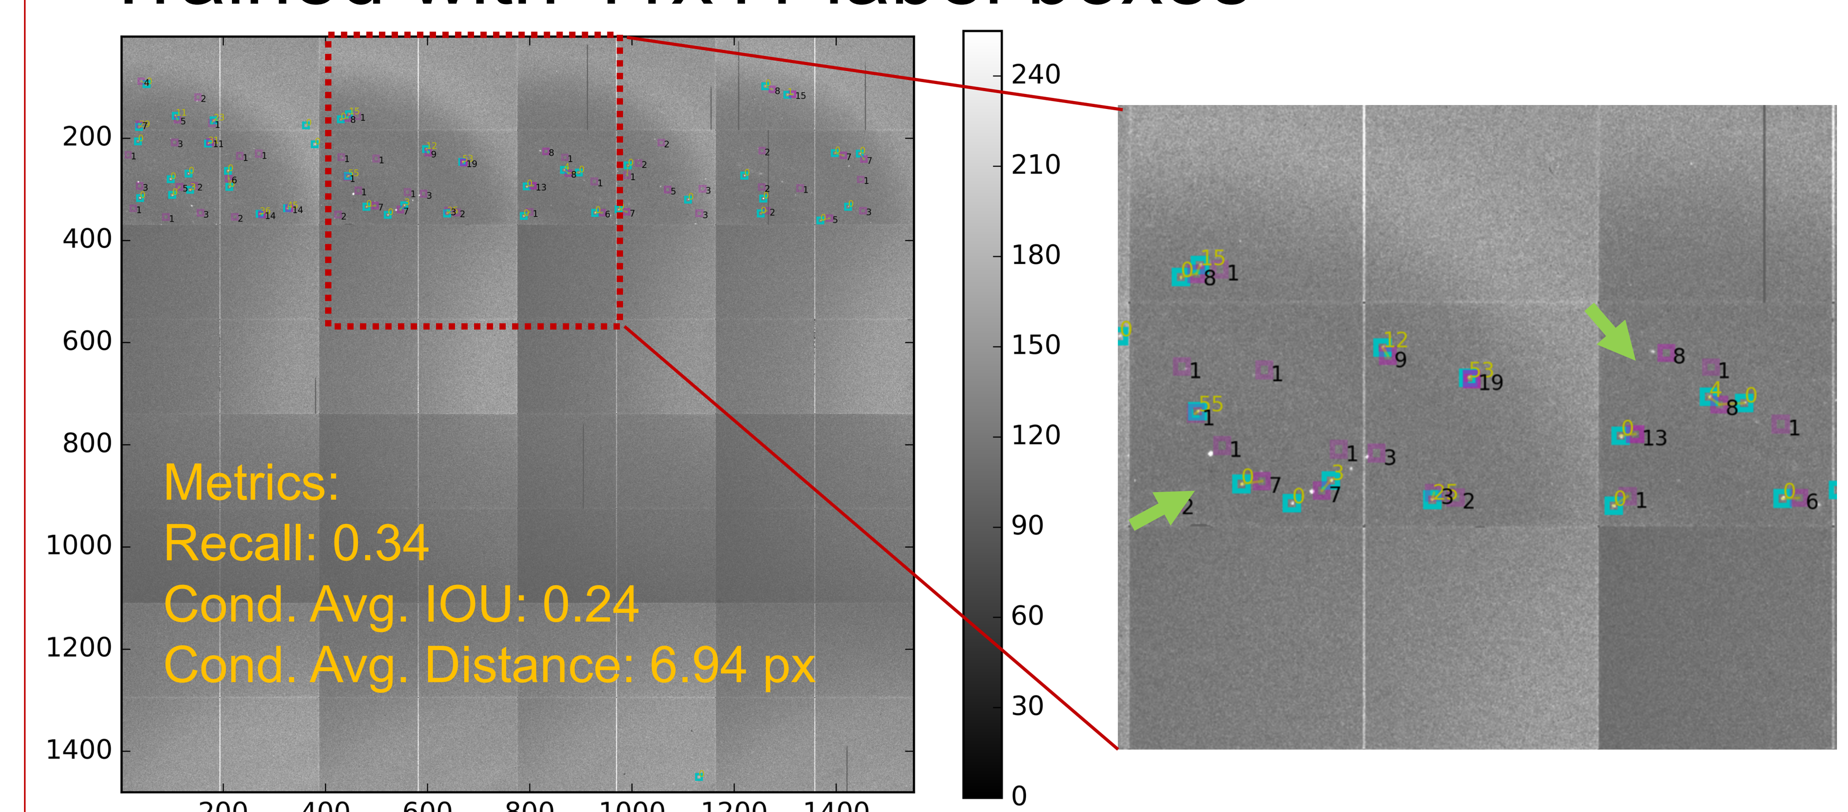

We then re-trained the PSNET with smaller label boxes, which reduces the error of prediction positions. However, this is at the cost of lower accuracy (TP/(TP+FP)) and lower confidence levels.

[1] J. Tenboer *et al.* Science 324, p. 1246, 2014.
[2] J. Redmon. Darknet: Open source neural networks in c. http://pjreddie.com/darknet/, 2013—2016.
[3] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi. You only look once: Unified, real-time object detection. *arXiv preprint arXiv:1506.02640*, 2015.
[4] J. Redmon and A. Farhadi. YOLO9000: Better, faster, stronger. *arXiv preprint arXiv:1612.08242*, 2016.