Introduction to deep learning with Pytorch

Tutors: Roman Vaxenburg (vaxenburgr@janelia.hhmi.org)
Habib Bukhari (bukharih@janelia.hhmi.org)

Session 1: 2020-11-30 13:00 UTC – 2020-12-01 17:00 UTC
Session 2: 2020-12-01 21:00 UTC – 2020-12-02 01:00 UTC

I2K 2020 tutorial

Title: Introduction to deep learning with Pytorch

Abstract:

This tutorial will introduce you to Pytorch, a popular open-source Python package for deep learning. The course will start with basic introduction to tensors, tensor operations, main components for building neural networks, and complete models. We will demonstrate the entire machine learning project pipeline by solving a classification task with several simple models: linear classifier, multilayer perceptron, convolutional neural network. Then we will work on a real-world classification and segmentation project. We will utilize pre-trained deep learning models, such as VGG, ResNet18, U-Net, and tailor them to our specific problem using transfer learning. Attendees are encouraged to bring their own data to apply the neural networks to their own classification/segmentation problems. This is a great preparation for the Decode tutorial.

How many people: 15

Requirements:

- Attendees:
    - Machine learning experience is not required
    - Comfortable with script programming
    - Optional: python and image analysis experience
- Technical requirements:
    - Access to a machine with Nvidia GPU, either locally, or via Google Colab (Google account is needed for Google Colab)
    - Additional directions for setting up the working environment and data format will be provided

About the tutors:

Habib Bukhari & Roman Vaxenburg: We are members of the Scientific Computing team at Janelia. We work with Turaga Lab on deep learning algorithms for neuroscience image analysis, neuron activity inference, and modeling the drosophila visual system.

Srini Turaga: A group leader of the deep learning research lab at Janelia. I develop machine learning algorithms to map neural connectivity, characterize neural activity and models to relate activity to connectivity.

Preferred time of the day: Any

Tutorial outline:

1. Main Pytorch components and concepts:

- Main datatype: **torch.Tensor** and its main properties
- Quick overview of model building blocks in **torch.nn** and **torch.nn.functional**
- The base class for all neural network modules: **torch.nn.Module**
- Loss functions, forward and backward pass
- Optimizers and gradient update with **torch.optim**
- CUDA semantics and **torch.cuda**

2. Basic machine learning pipeline with Pytorch:

- Building datasets and dataloaders
- Working with CPU and GPU
- Building simple models
- Training loop
- Model validation and testing

3. From basics to project:

- How to set up a project (download data, explore data, create labels, data augmentation)
- How to choose architecture
- Solving classification and segmentation tasks with pre-trained convolutional neural networks using transfer learning
- Attendees are welcome to bring their own data compatible with these two tasks