CPEN 455: Deep Learning

Tutorial 6: Introduction to Project (Conditional PixelCNN++)

Qihang Zhang & Sadegh Mahdavi

University of British Columbia

Winter, Term 2, 2024

Project Description

- The project is available at https://github.com/DSL-Lab/CPEN455HW-2024W2
- You will modify the original code of PixelCNN++ to create a conditional PixelCNN++ for Image Classification.
- The details of how PixelCNN++ works will be covered in subsequent tutorials, but you can read the README, and the PixelCNN++ paper to get yourself familiar with the code template.

Overview of the Dataset

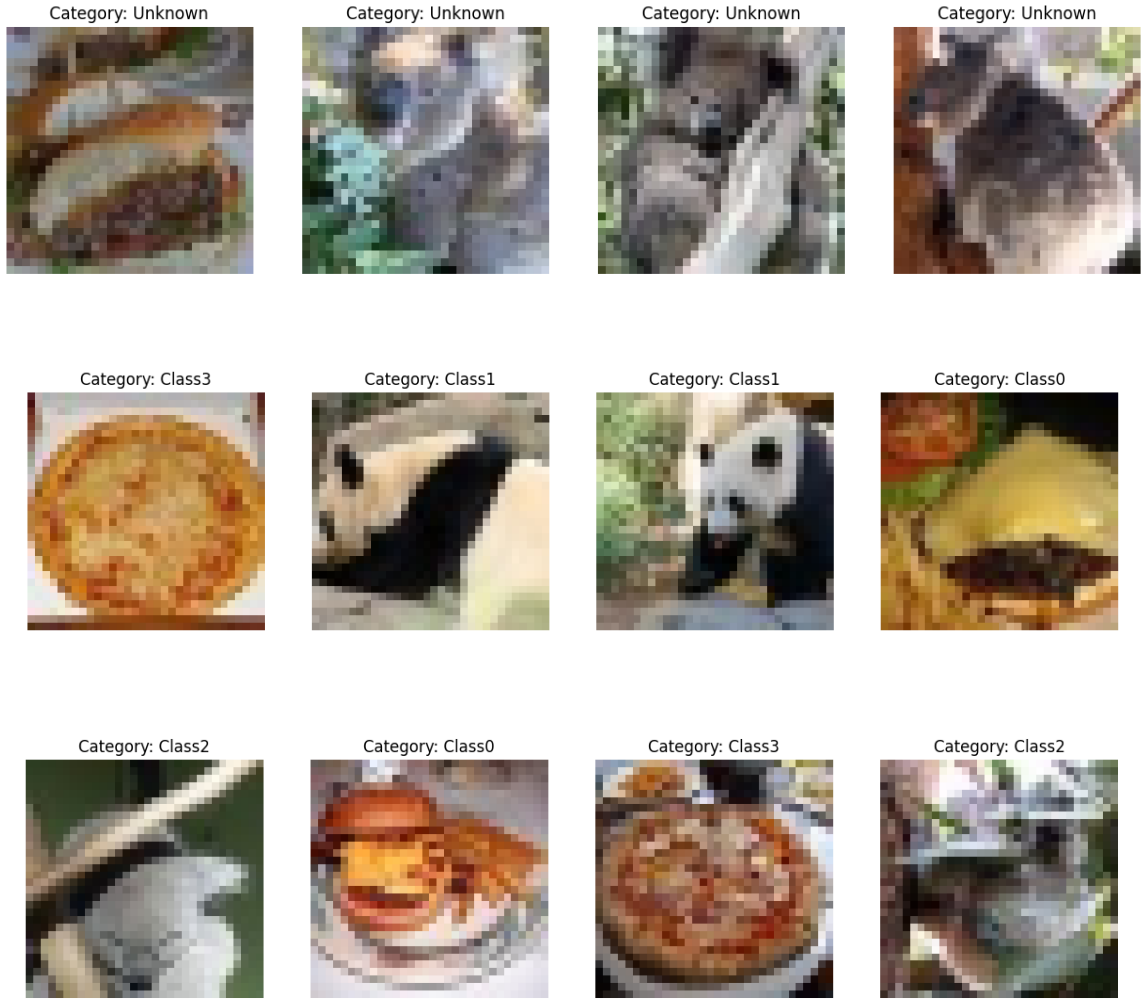

Project Description

- TL;DR: Teach classification to a generative model!
- You will implement:
  - A conditional PixelCNN++
  - Use it to both **generate** and **classify** images.

$$p(c = i|x) = \frac{p_\theta(x|c = i) \times p(c = i)}{\sum_{k=1}^{4} p_\theta(x|c = k)p(c = k)}$$

Running On Colab

- Put the project files into your google drive
- Mount your Google Drive on Colab
- Edit the files
- Run the notebook!

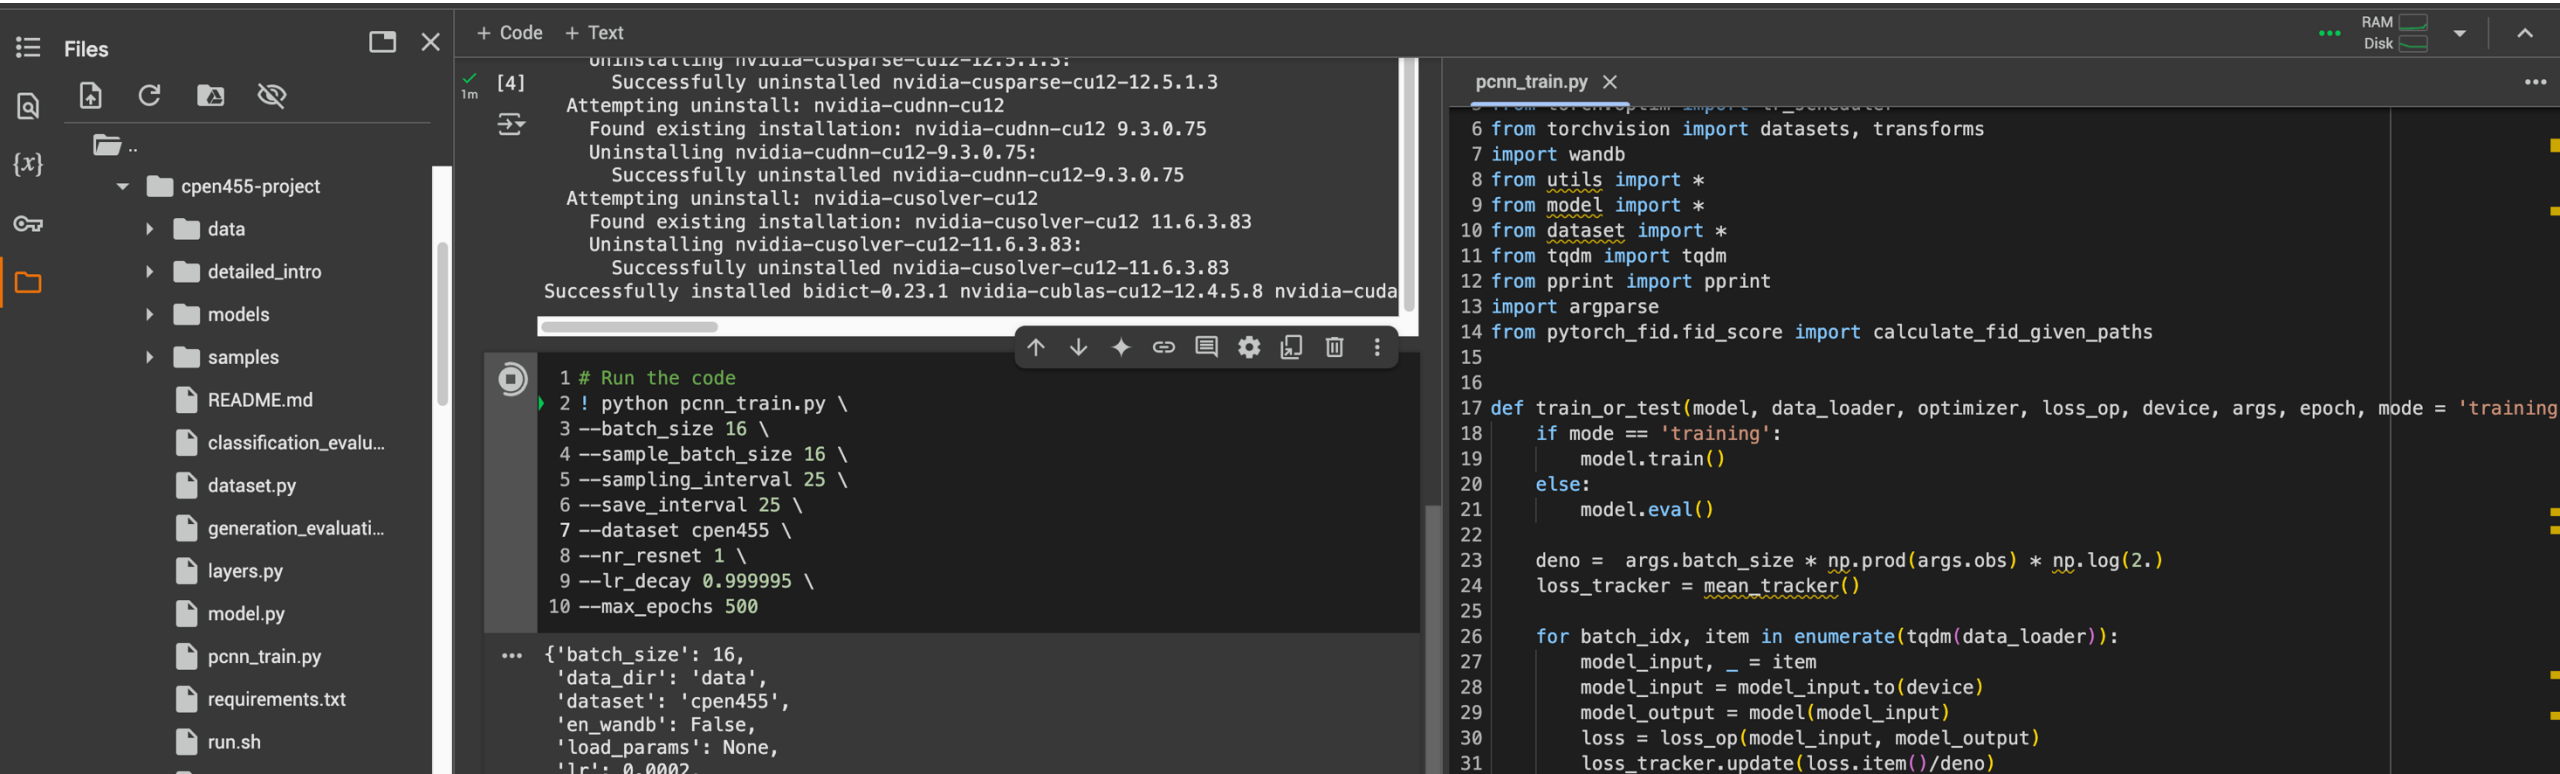

Running On Colab

- If you work with Colab:
  - Make sure you always use GPU on colab (otherwise code is too slow)
  - Make sure you periodically take a snapshot of your code (no easy Git support on colab)

Conda Env (Running Locally)

- Install pytorch
- pip install –r requirements.txt

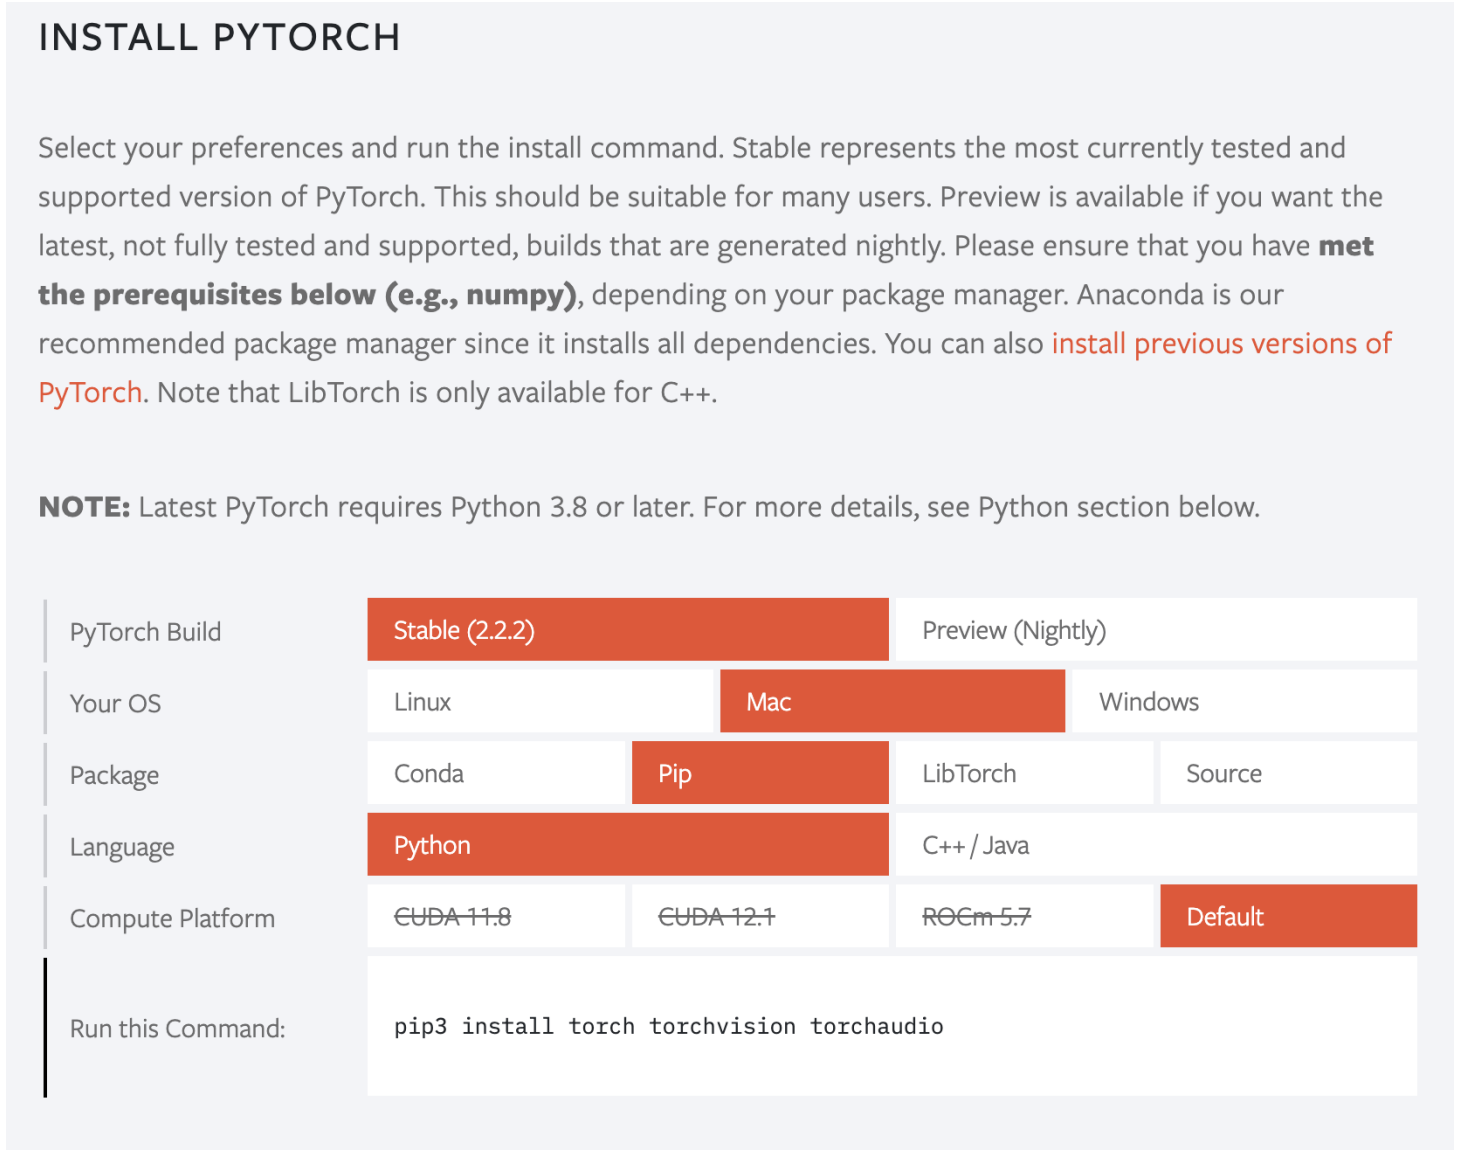

Conda Env

- Install pytorch
- pip install –r requirements.txt

nvidia-smi

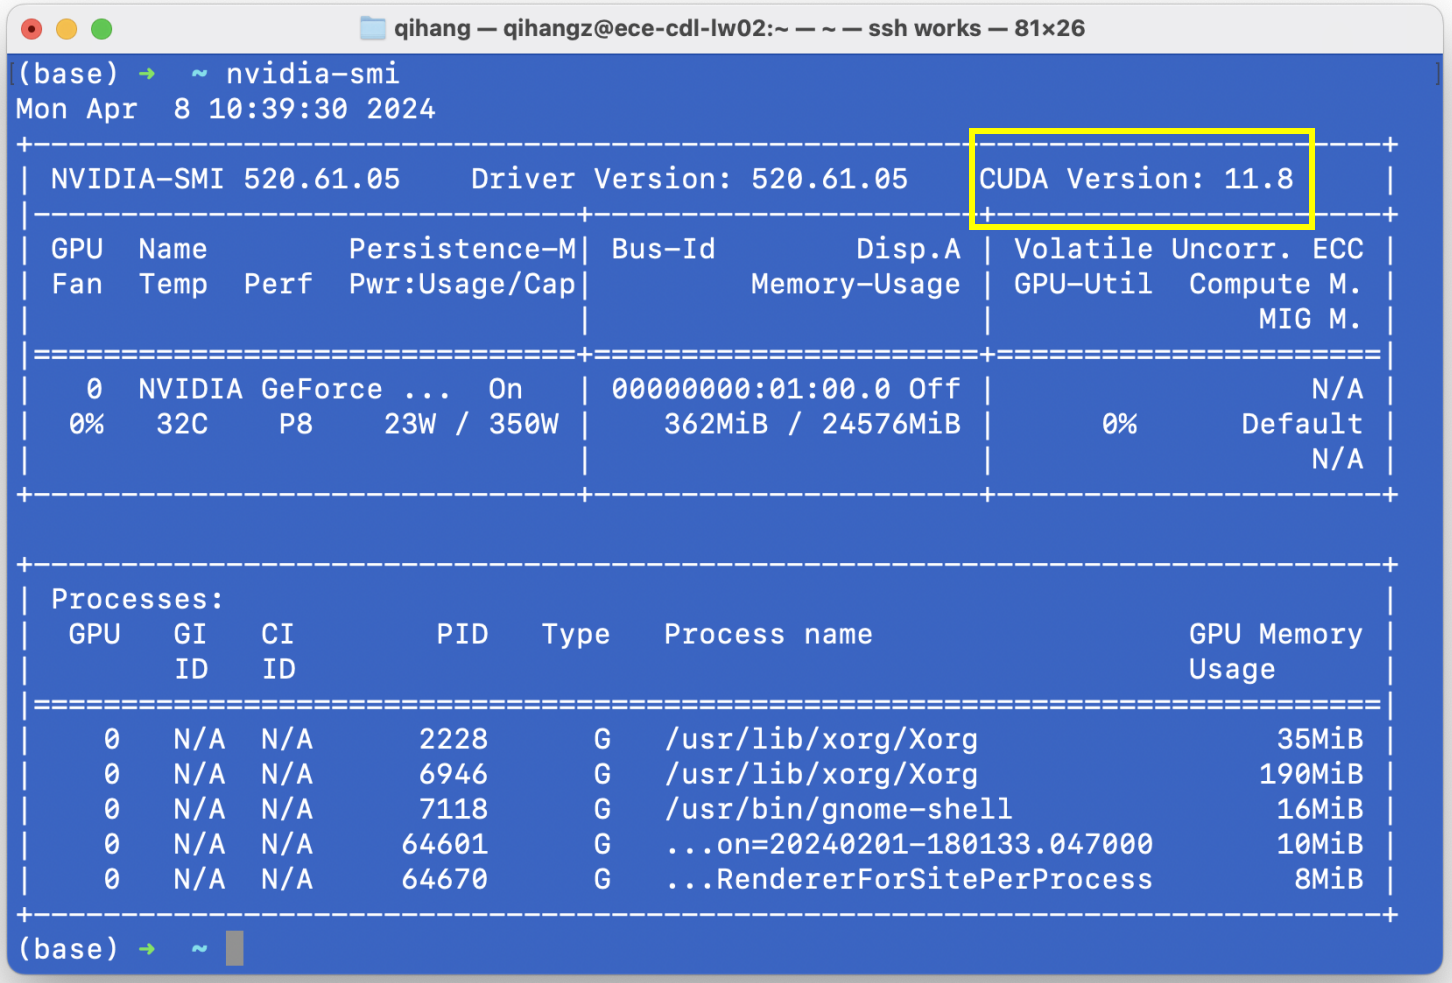

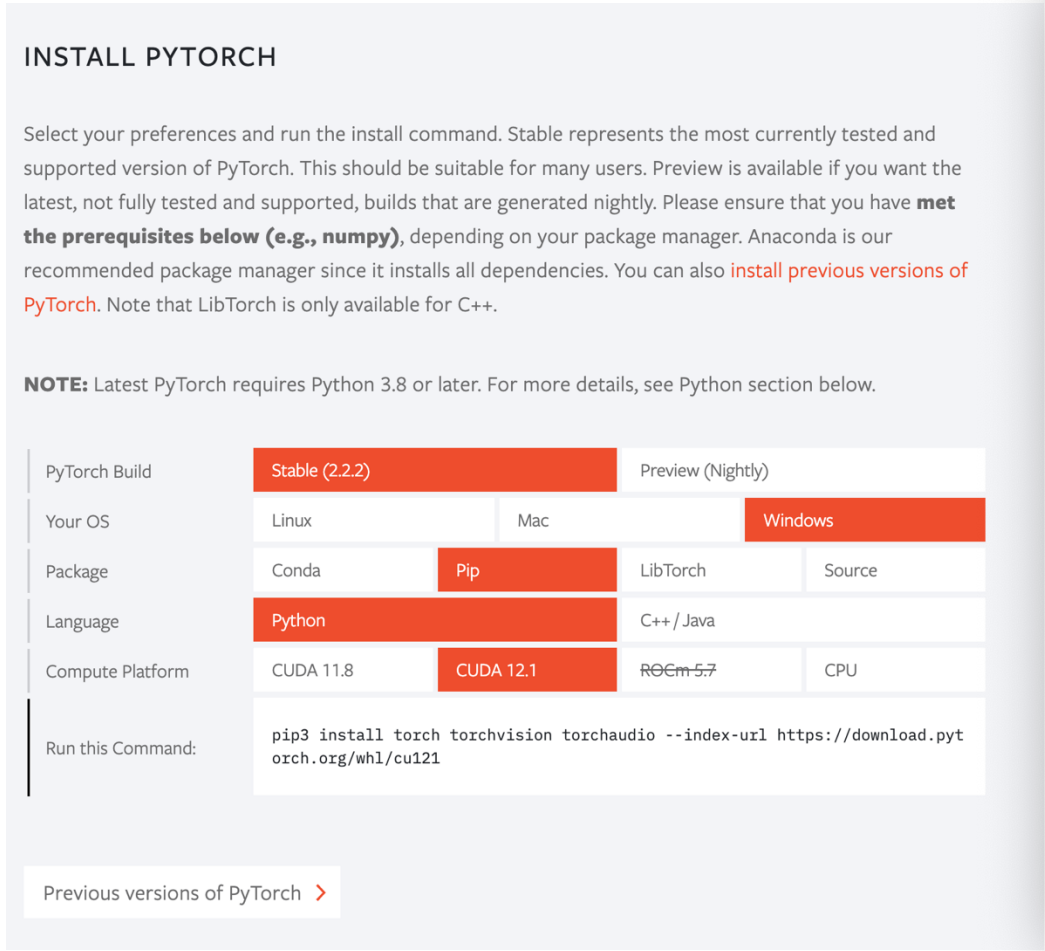

Pro Tip: Potential numerical issue you may encounter

- **Never use "log" and "softmax" functions sequentially** (like log(softmax()) function).
- Every deep learning scientist has learned this the hard way (~days of debugging!)
- This may lead to severe numerical issues.
- Instead, the following torch functions are numerically stable (depending on your need):

torch.logsumexp

torch.logsumexp(*input*, *dim*, *keepdim=False*, ***, *out=None*)

Returns the log of summed exponentials of each row of the `input` tensor in the given dimension `dim`. The computation is numerically stabilized.

For summation index j given by *dim* and other indices i, the result is

$$\text{logsumexp}(x)_i = \log \sum_j \exp(x_{ij})$$

If `keepdim` is `True`, the output tensor is of the same size as `input` except in the dimension(s) `dim` where it is of size 1. Otherwise, `dim` is squeezed (see `torch.squeeze()`), resulting in the output tensor having 1 (or `len(dim)`) fewer dimension(s).

CrossEntropyLoss

CLASS torch.nn.CrossEntropyLoss(*weight=None*, *size_average=None*, *ignore_index=-100*, *reduce=None*, *reduction='mean'*, *label_smoothing=0.0*) [SOURCE]

This criterion computes the cross entropy loss between input logits and target.

It is useful when training a classification problem with *C* classes. If provided, the optional argument `weight` should be a 1D *Tensor* assigning weight to each of the classes. This is particularly useful when you have an unbalanced training set.

The *input* is expected to contain the unnormalized logits for each class (which do *not* need to be positive or sum to 1, in general). *input* has to be a Tensor of size (C) for unbatched input, $(minibatch, C)$ or $(minibatch, C, d_1, d_2, ..., d_K)$ with $K \geq 1$ for the *K*-dimensional case. The last being useful for higher dimension inputs, such as computing cross entropy loss per-pixel for 2D images.

LogSoftmax

CLASS torch.nn.LogSoftmax(*dim=None*) [SOURCE]

Applies the $\log(\text{Softmax}(x))$ function to an n-dimensional input Tensor.

The LogSoftmax formulation can be simplified as:

$$\text{LogSoftmax}(x_i) = \log\left(\frac{\exp(x_i)}{\sum_j \exp(x_j)}\right)$$

Shape:

- Input: $(*)$ where * means, any number of additional dimensions
- Output: $(*)$, same shape as the input

Leaderboard and Bonus Points [Subject to change]

- Bonus points for:
  - If your model outperforms our baseline (on a Huggingface Leaderboard).
  - Detailed analysis of your model or generated results.

Hugging Face

Questions?