

DL/DLOps (2023)
Lab Assignment 9: Parallel Training [100 Marks]
Deadline: 23/04/2023, 23:59:59

There will be a 25% penalty for each day of late submission.

Programming instructions:

1. Programming language: Python
2. Use of PyTorch is compulsory. Marks shall not be given for TensorFlow implementation.

Reporting instructions:

1. Please submit all your working codes as .py or .ipynb files, which should be named appropriately with your complete roll number(ex:"XYZ_Lab_Assignment_9.py").
2. A single report (PDF) file should be submitted containing all relevant information including data pre-processing, observations, results, and analysis across the problem. Do not put snapshots of code in the report.
3. The report should be detailed and clearly explain every step you have followed. All the intermediate outputs, their inferences should be present in the report. The PDF file should be properly named with your complete roll number XYZ ("XYZ_Lab_Assignment_9.pdf"), with your name and roll number mentioned inside the report as well.
4. Record a video of your terminal/ipynb file which shows that your written codes are working. (max. video size - 3 minutes)
5. Mention any resources/articles/GitHub links that you may have used as a reference to solve any question of the assignment in the references section of the report.
6. Make sure all the submission files, along with the working codes, are included in a single zip file.

General instructions:

1. DO NOT plagiarize from the internet or your peers. The institute's plagiarism policy will be strictly enforced.
 2. The assignment will be evaluated out of 50% of the total marks in case a report is not submitted.
 3. We highly suggest using the DGX server with GPU runtimes for this assignment.
-

Question 1

[35 marks]

Train a ResNet18 model on SVHN dataset using DataParallelism and DistributedDataParallelism for 30 epochs. You can use any library of your choice for parallel training of your model in PyTorch.

[35 marks]

- A. Use a batch size of 32 and train on a single GPU node. **[5 Marks]**
- B. Use a batch size of 64 and train and two GPU nodes first using DataParallelism and then using DistributedDataParallelism. **[15 * 2 = 30 Marks]**

1. Report all the hyper-parameters used for Parallel Training.
2. Compare the time (in seconds) and report the speed up. Show this speed up using a graphical representation.
3. Describe your observations in terms of memory usage for multi-node training.

Instructions to use multiple GPUs on server 172.25.0.15

Please follow the following instructions to enable multiple GPUs in the Slurm environment on the server:

1. Navigate to the target directory and edit the batch script file.
2. Using vim or nano, edit line from "`#SBATCH --gres=gpu`" to "`#SBATCH --gres=gpu:2`" (line 7 in `test-gpu.sh`).
3. Save this file and submit your job using `sbatch` command. This will allocate two GPUs for your job.
4. You can monitor the memory usage of each GPU using `nvidia-smi` command.