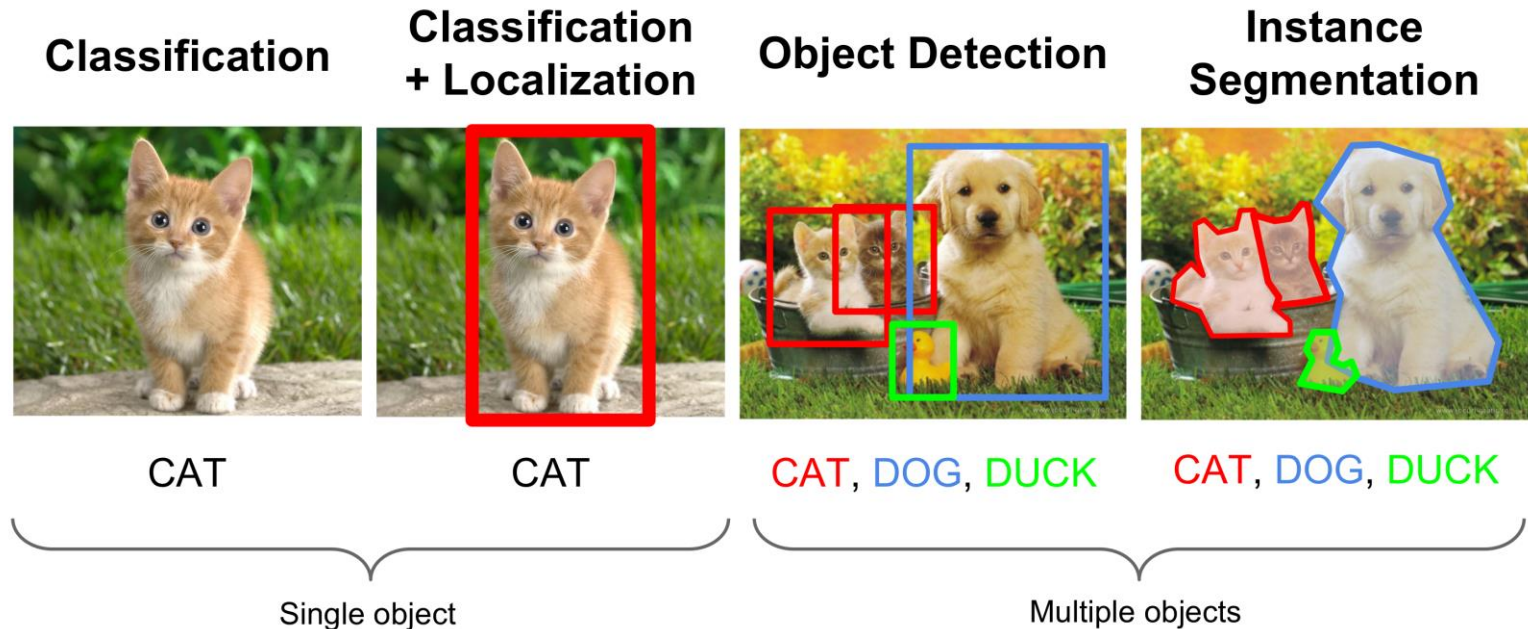


UT Southwestern
Medical Center
Lyda Hill Department of Bioinformatics

BioHPC


Distributed Deep Learning
BioHPC – 6/22/2022

Given examples, can we train a computer to do:



We want the DL model to LEARN features! We want to TRAIN the DL model!

Machine Learning on Images




Source - <http://cs231n.github.io/>


Artificial Neural Networks

- What are **neurons**?
- How are they **connected**?

} **Neural Networks**




Artificial Neural Networks



- Signal goes in, via input layer
- Weighted links transfer input values to neurons in hidden layers
- Signals are summed at hidden neurons and passed through transfer/activation function
- Processed signal arrives at output layer
- Decisions made using output signal(s)



What's in an (Artificial) Neuron?

Closer look at a neuron and a single layer




Sources - <http://cs231n.github.io/>
<https://www.v7labs.com/blog/neural-networks-activation-functions>

What do we expect from the hidden layers?



<https://www.3blue1brown.com/lessons/neural-networks>


What we expect from Deep Learning?



More layers can encapsulate more knowledge.

More weights to train – need more data, need more computation


Where's the Knowledge?



How the Neural Network learns?


- Weights encapsulate the knowledge of a network
- Network learns using an algorithm that optimize weights given **training data**.
- Minimize cost function
 - Cost function: summation over the differences between the NN prediction and the actual value of a class

Minimizing the cost function by using gradient descent algorithm



<https://www.3blue1brown.com/lessons/neural-networks>

Backpropagation




Q: X and Y are fed to the network and Z is the model predication - How good is this prediction or Z? How much error is associated with X and Y?

A: Calculate the derivative of the loss function with respect to X and Y (dL/dx and dL/dy) - Then subtract X and Y by dL/dx and dL/dy

<https://becominghuman.ai/back-propagation-in-convolutional-neural-networks-intuition-and-code-714ef1c38199>

Overview of the training process



Training process review:

- 1. Forward pass input data through the network*
- 2. Get the output from the network*
- 3. Measure the error/loss between the network output and the true labels of data*
- 4. Perform the backpropagation*
- 5. Update network weights with the gradients*

GPUs to the Rescue!



GPU cards are exceptionally well suited to Neural Network Mathematics

Orders of magnitude faster than CPU-based training



Keras

- High-level, open-source Python API
- *“Being able to go from idea to result with the least possible delay is key to doing good research”*
- Interface for TensorFlow, Microsoft Cognitive Toolkit, and Theano

Installing a Conda Environment for Keras and TensorFlow with Jupyter Support

```
$ module load python/3.6.4-anaconda
```

```
$ conda create --name py3.6-keras python=3.6 ipykernel keras  
tensorflow-gpu pillow matplotlib
```

```
$ ipython kernel install --user --name py3.6-tfgpu --display-  
name="Keras (GPU)"
```

BioHPC

[Comment on this page](#)
Logged in as: s178337

Status Training Guides FAQs Cloud Services BioHPC OnDemand Links Software Careers

Demand - Jupyter Notebook

The Jupyter project provides web-based notebooks for including Python, R, and Julia. Create beautiful notebooks

OnDemand DIGITS

OnDemand Jupyter

OnDemand BisQue

OnDemand RStudio

languages,
ts, tables, and


Basic concepts in distributed and parallel training Deep Learning: Performance metrics

- Speedup
- Throughput
- Scalability


How to speedup the training process?

Basic concepts in distributed and parallel training Deep Learning: Types of parallelism

Model Parallelism

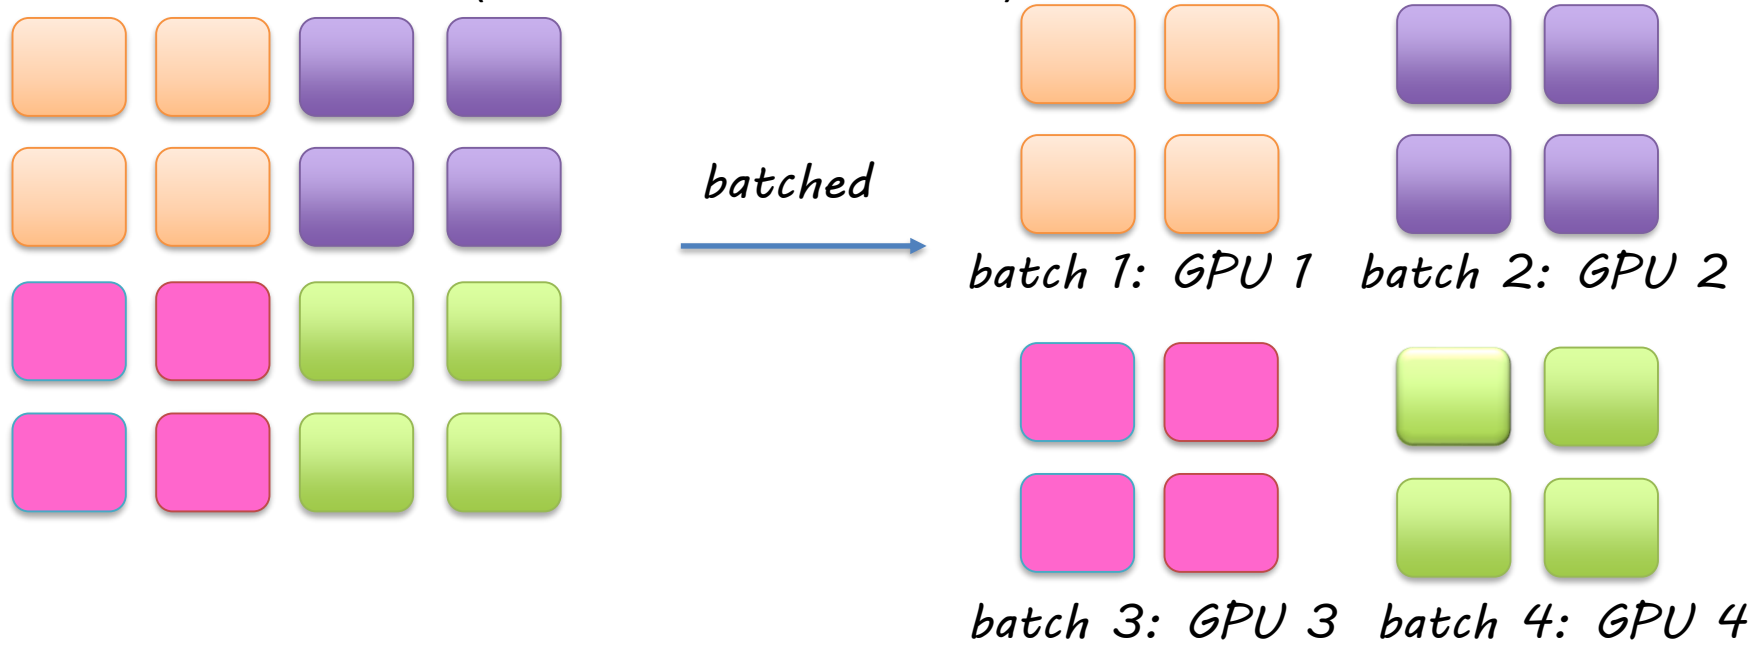


Data Parallelism



Basic concepts in distributed and parallel training Deep Learning: Batch size

The scale of data parallelism corresponds to the *batch size*.
(for the most common NNs)



Divide dataset into batches -> stochastic gradient descent -> takes longer to converge

Example: `python my_dl_model.py --model resnet50 --batch-size 128`

But what are the limits of the data parallelism approach, and when should we expect to see large speedups?

“while simple data parallelism can provide large speedups for some workloads at the limits of today's hardware (e.g. Cloud TPU Pods), and perhaps beyond, some workloads require moving beyond simple data parallelism in order to benefit from the largest scale hardware that exists today...”

Maximum useful batch size depends on:

- *Model architecture*
- *The dataset*
- *The optimizer*

Read more: <https://ai.googleblog.com/2019/03/measuring-limits-of-data-parallel.html>

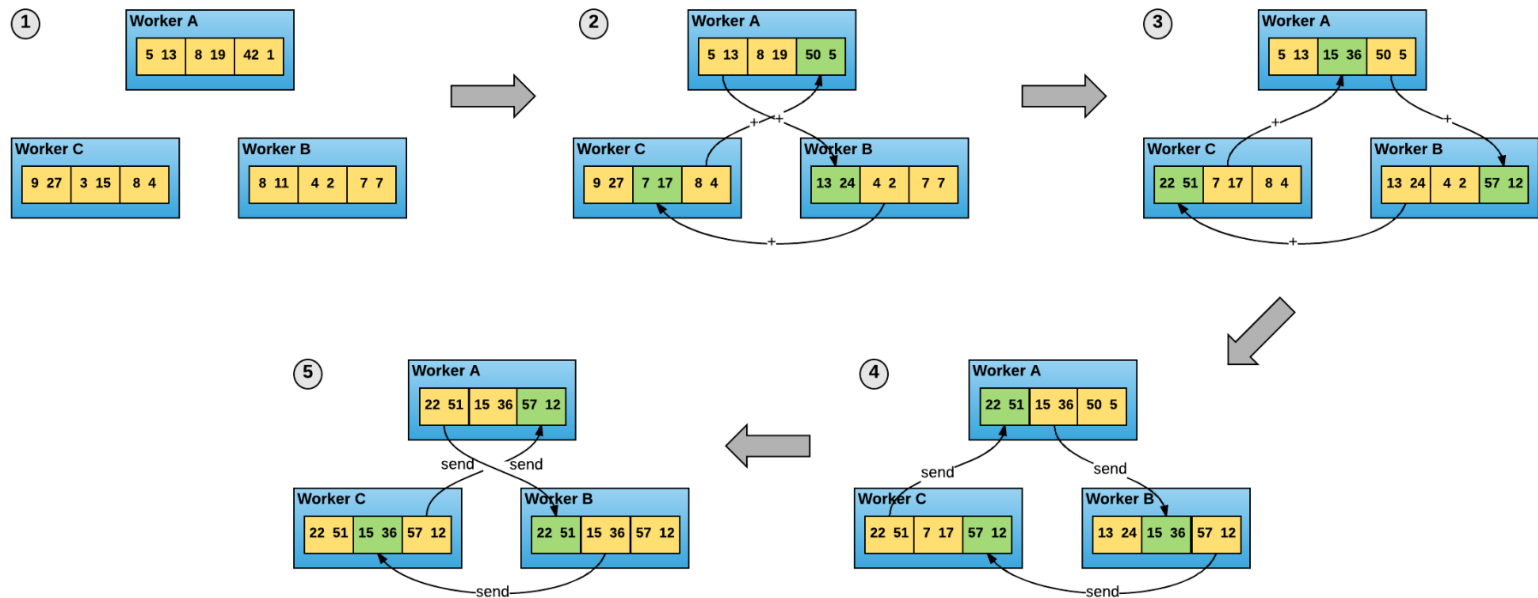
Distributed Deep Learning through Horovod

- Distributed DL training framework for:
 - Tensorflow/Keras
 - Pytorch
 - MXNet
- Open-source
- Separates infrastructure from DL/ML
- Installation: `pip install horovod` (not recommend)
 - Use existing Docker image and convert to Singularity
- Uses bandwidth-optimal communication protocols (e.g. Infiniband) if available
- Some terminology: Horovod master and worker nodes
 - Master sends the variables to the worker during the initialization
 - Worker nodes do the training job!



Example: `python my_dl_model_horovod.py --model resnet50 --batch-size 128`

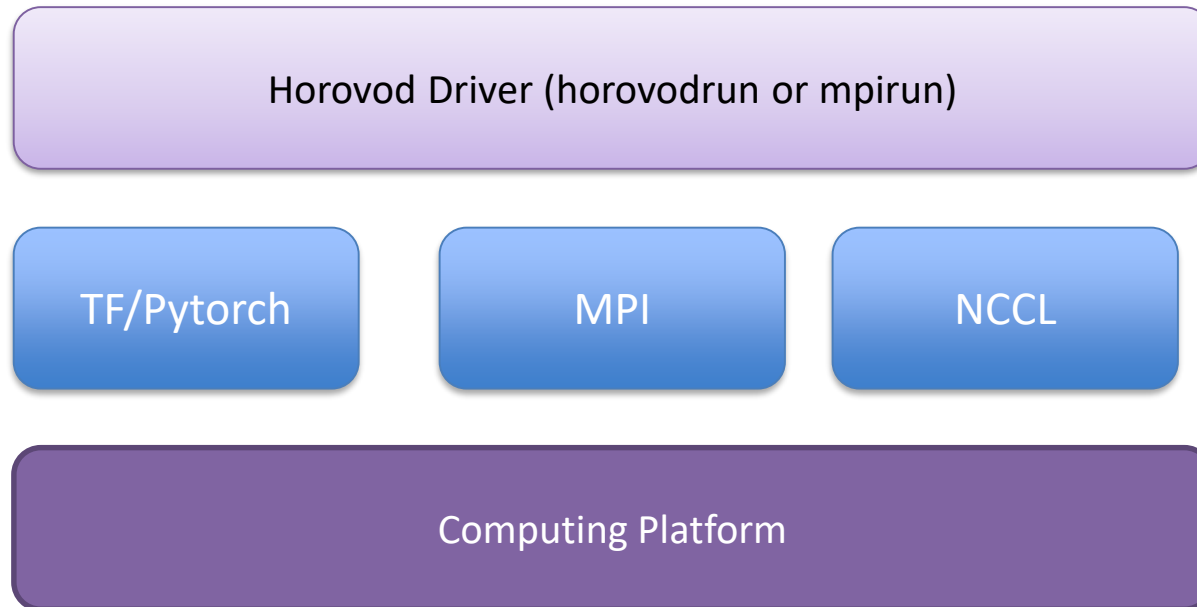
Distributed Deep Learning through Horovod (cont.)



Horovod Algorithm: Ring All-reduce

<https://arxiv.org/pdf/1802.05799.pdf>

Distributed Deep Learning through Horovod (contd.)



Horovod Stack

NCCL: Used for GPU-2-GPU communications

MPI: Used for CPU-2CPU communications

Distributed Deep Learning through Horovod (cont.)

Simplified example of showing how to use Horovod with Tensorflow:

1. Initialize Horovod:

```
import horovod.tensorflow as hvd  
  
hvd.init()
```

2. Pin GPU to be used to process local rank (one GPU per process):

```
config = tf.ConfigProto()  
  
config.gpu_options.visible_device_list = str(hvd.local_rank())
```

3. Add Horovod Distributed Optimizer and scale the learning rate:

```
opt = tf.train.AdagradOptimizer(0.01 * hvd.size())  
  
opt = hvd.DistributedOptimizer(opt)
```

<https://horovod.ai/getting-started/>

Distributed Deep Learning through Horovod (cont.)

4. *Broadcast variables from rank 0 to all other processes during initialization.*

```
hooks = [hvd.BroadcastGlobalVariablesHook(0)]
```

5. *Save checkpoints only on worker 0 to prevent other workers from corrupting them.*

```
checkpoint_dir = '/tmp/train_logs' if hvd.rank() == 0 else None
```

Read more: <https://towardsdatascience.com/distributed-deep-learning-with-horovod-2d1eea004cb2>

1. Either install it via pip OR pull the Docker image of Horovod and convert to Singularity as following:

```
module load singularity/3.0.2
singularity pull docker://horovod/horovod:latest horovod_latest.sif
```

2. Get inside the Singularity container and check the mpi version:

```
singularity run horovod_latest.sif
mpirun --version
```

3. Based on the MPI version in the container, install the same version of openmpi in your account (instruction uploaded BioHPC Portal -> Training -> Training Slides & Hands out)

4. Submit the job to the cluster – Demo

(code is uploaded BioHPC Portal -> Training -> Training Slides & Hands out)

Try on GPU A100 or GPU V100s