

An introduction to Tensorflow


Jinming Hu & Simei He

2021.4.24

An intro to DL


What is machine learning?

- Roughly speaking, learn a mapping function from data
- inputs x to outputs y
 - e.g image classification -> map images to labels




What is deep learning/neural network?

- Consist of many **layers** of **neurons**
- For neurons, $f(\mathbf{x}) = \mathbf{g}(\mathbf{w}^T \mathbf{x} + b)$



What is a layer of neural network?

- Multiple neurons together, with the same input:




$$f(\mathbf{x}) = \mathbf{g}(\mathbf{w}^T \mathbf{x} + b)$$

$$\mathbf{y} = F(\mathbf{x}) = \mathbf{G}(\mathbf{W}^T \mathbf{x} + \mathbf{b})$$

What is neural network?

- Multiple layers stacked together



$$\mathbf{y} = F(\mathbf{x}) = \mathbf{G}(\mathbf{W}^T \mathbf{x} + \mathbf{b})$$

$$\mathbf{z} = F(\mathbf{y}) = \mathbf{G}(\mathbf{W}^T \mathbf{y} + \mathbf{b})$$


How to get the output with an existing NN?

- This is called predicting/testing.
- Just do some matrix multiplication:

$$\mathbf{y} = F(\mathbf{x}) = \mathbf{G}(\mathbf{W}^T \mathbf{x} + \mathbf{b})$$


$$\mathbf{z} = F(\mathbf{y}) = \mathbf{G}(\mathbf{W}^T \mathbf{y} + \mathbf{b})$$

This is called feedforward



How are the parameters learned?

- This is called training
- Get some “ground truth” labeled data, a set of (x, y) i.e. training data
- Feedforward: $y' = f(x)$, calculate loss: $L(y', y)$
- Gradient Descent $W = W - \lambda \frac{\partial L}{\partial W}$




How are the parameters learned?

- **Backward: SGD** $W = W - \lambda \frac{\partial L}{\partial W}$
- How to get gradient?
- Backpropagation.

$$\frac{\partial L}{\partial w_{kj}} = \frac{\partial L}{\partial z_k} \cdot \frac{\partial z_k}{\partial net_k} \cdot \frac{\partial net_k}{\partial w_{kj}}$$

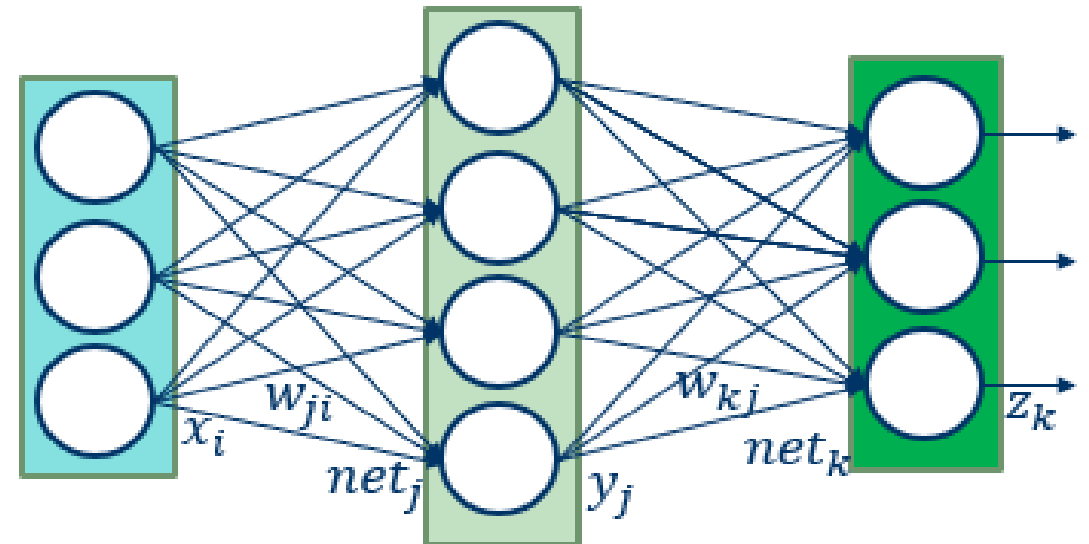
$$= (z_k - t_k) \cdot f'(net_k) \cdot y_j$$

$$\frac{\partial L}{\partial W} = \frac{\partial L}{\partial z} \cdot \frac{\partial z}{\partial net} \cdot \frac{\partial net}{\partial W}$$



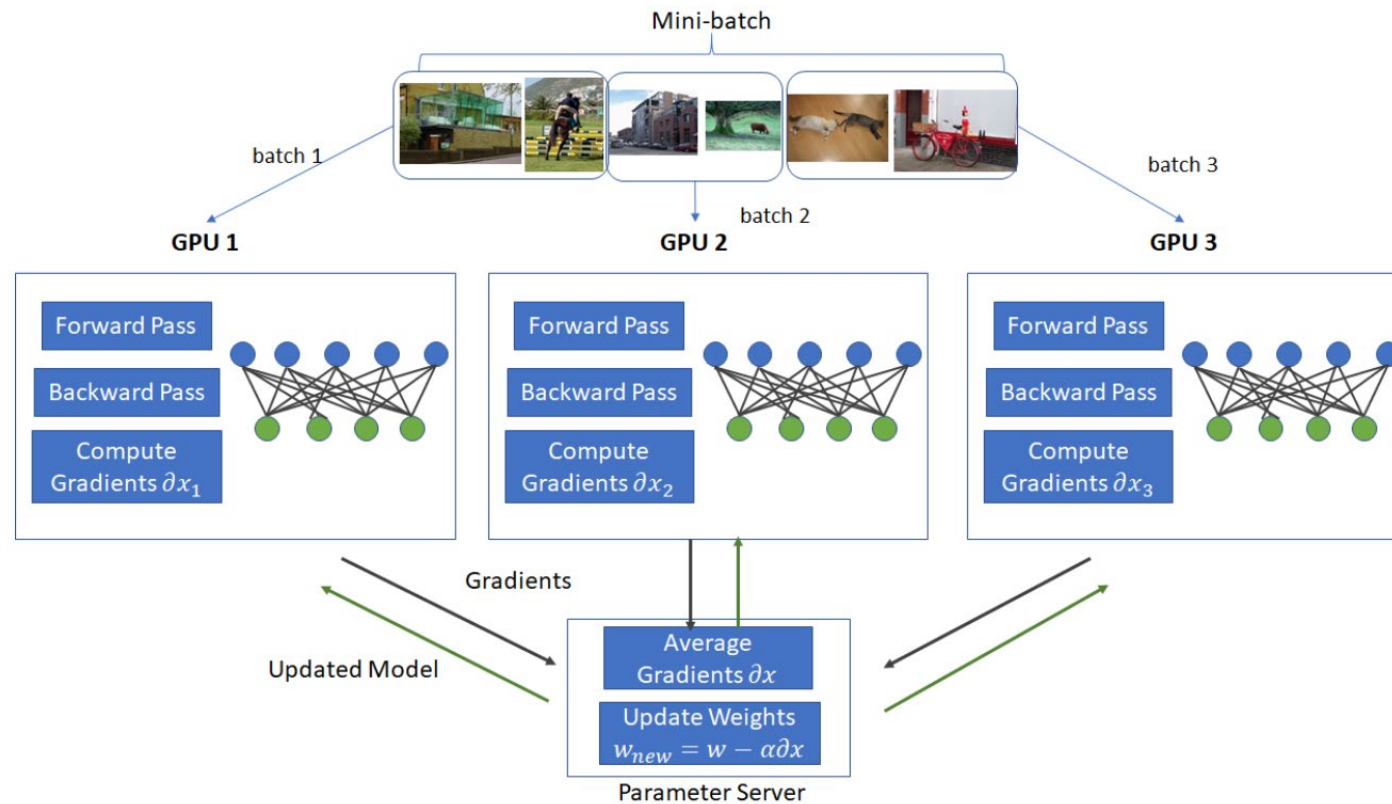
How are the parameters learned?

- For each **step/iteration**, we sample a **batch** of data from training data(usually sth like 64, 128, 256...)
- Feedforward + backward for this batch, get gradient for each neuron
- Apply some kinds of gradient descent for each neurons:
 - SGD: $W = W - \lambda \frac{\partial L}{\partial W}$
- Also some other methods



How to parallelize the computation?


- Data Parallel: computation for **every sample** in **one batch** can be parallelized




How to parallelize the computation?

- Model parallel: computation for different parts of one neural network can be parallelized

Data Parallel



Model Parallel



So what do we want for a DL system?

- A lot of matrix computation, so we want for efficient matrix computation.
 - Thus run it efficiently on GPU
- Efficient to represent the computation for NN.
 - Use computational graph!
 - Automatically compute gradients.
- Parallelize computation easily.
 - Enable two types of parallelism efficiently
- Flexible to add new types of layers/optimization methods.
 - Can develop new layers/optimizations with basic operations
 - Use a handy wrapper language: usually python
- Once trained, deploy everywhere.
-


System Part

Design Goals


- To unify large-scale and small-scale machine learning abstraction
- To allow expressive paralleling progressing

Programming Models

- Graph: stateful dataflow computation
- Node: operation
- Tensor: values that flow through normal edges



Automatic Gradient Computation

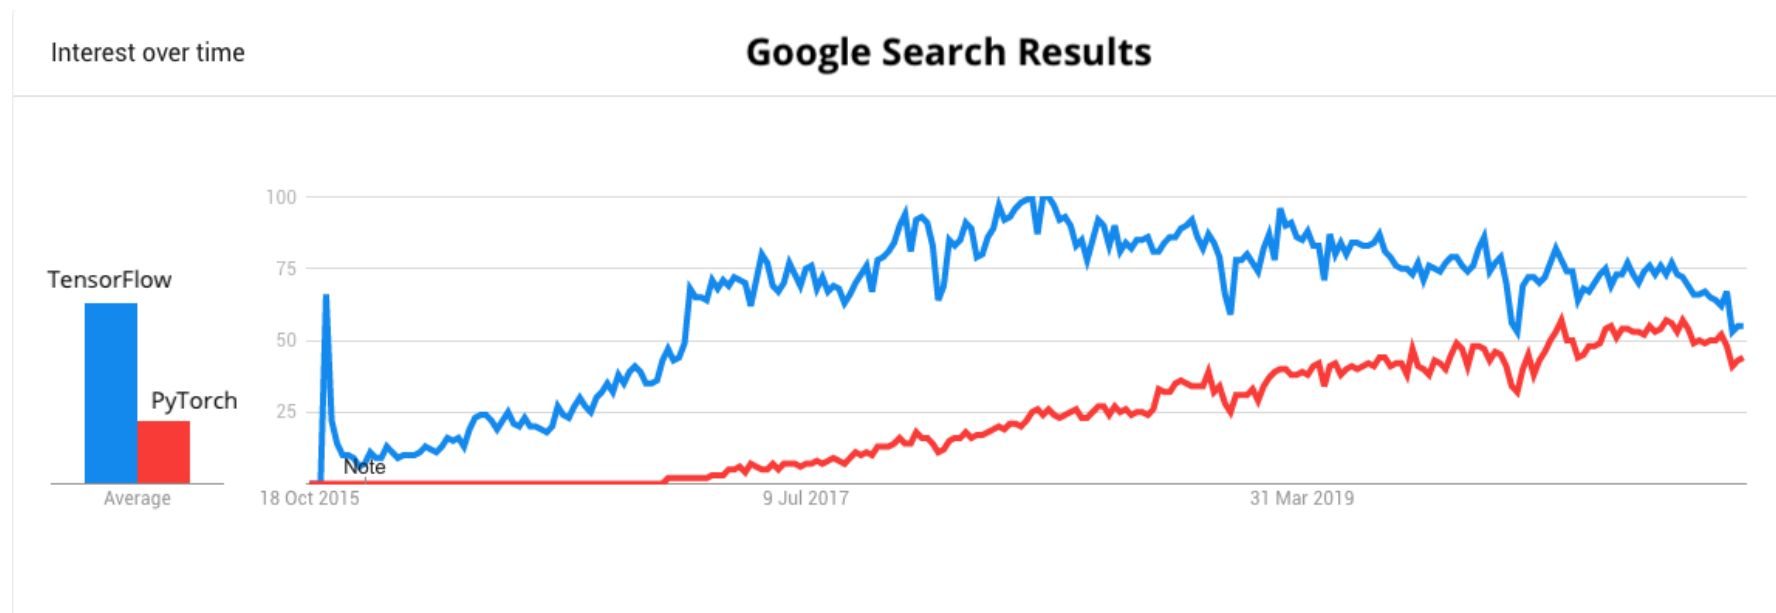


Graph Execution (TF1)


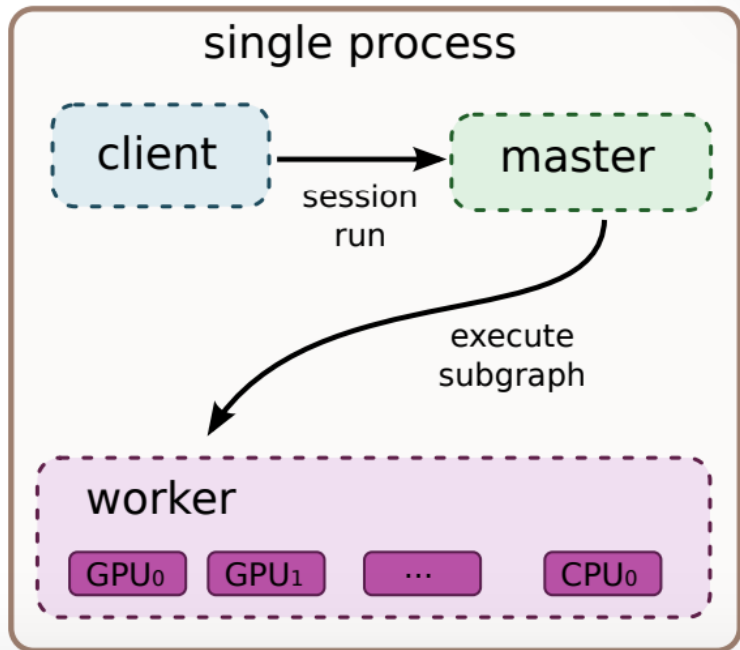
- Fast and efficient;
- Runs in parallel
- Easy to optimize
- With GPU & TPU acceleration capability

Eager Execution (TF2)


- PyTorch - dynamic computation graphs



Client, Master, Worker Layered Cake

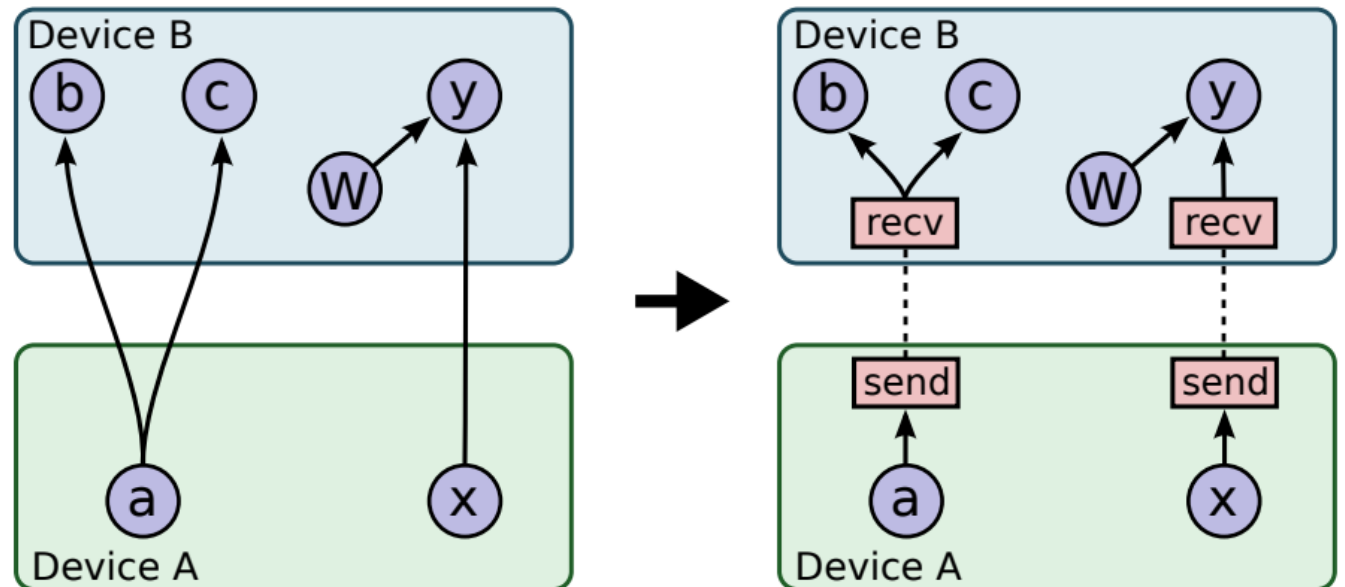


Partial Execution



Multi-Device Executions

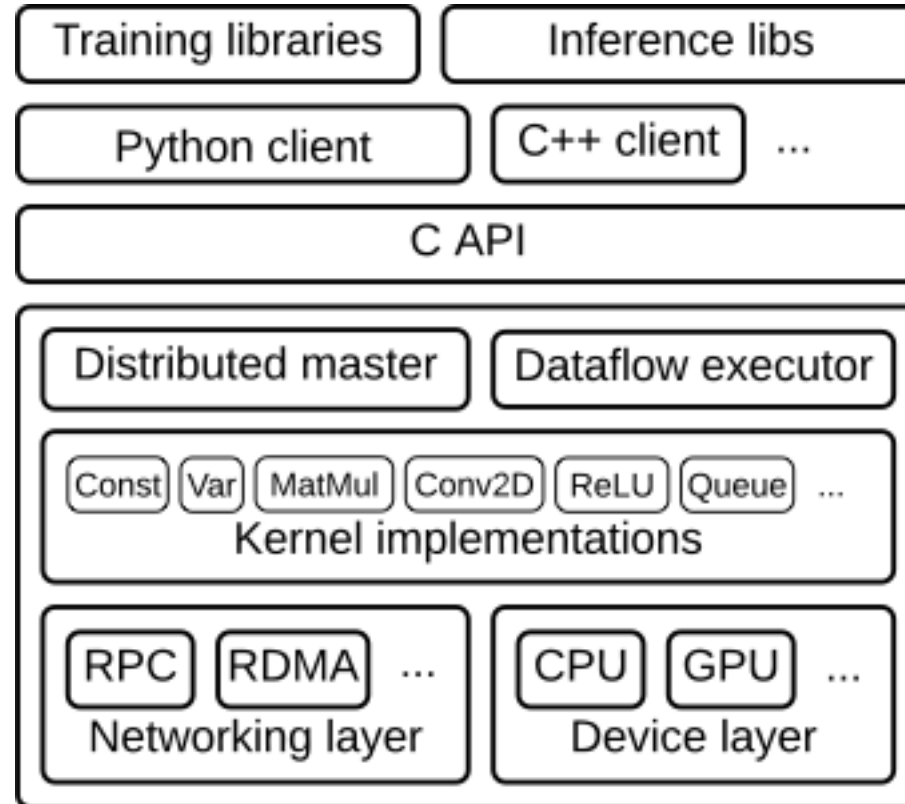
- Node placement: heuristics placement algo; device constraints
- Cross-device communication



Distributed Executions

- Checkpoint support: to connect Variable node to Save node/Restore node

General Architecture



Eh?

badabummbadabing 1 month ago · edited 1 month ago 🤝 😏 🐻

I used Tensorflow 1.x for years and felt like I was quite the expert at it. Then, due to a technical limitation, I had to implement a project in Pytorch (which I had never used before), which I had tried and failed to implement in Tensorflow for a long time (some graph manipulation). Not only did the implementation only take a few days, it didn't even take a month before I felt that I was as good in Pytorch as I was in Tensorflow. Turns out, most of my supposed expert knowledge in Tensorflow revolved around dealing with the many quirks and weird behaviours of Tensorflow, while things just work without arcane knowledge in Pytorch.

And Tensorflow has no design philosophy.