



Final session

Very beginning of Neural-Network

R-CCS International HPC Summer School 2022
Lecture: Dr. Toshiyuki Imamura, RIKEN R-CCS
Assistant: Dr. Takeshi Terao, RIKEN R-CCS



What do you image AI or deep-machine-learning?

Honestly, ...


I am not an AI expert.

It is a good experience to learn it from non-AI persons.

Textbook:

N. Buduma, Fundamentals of Deep Learning, designing next-generation machine intelligence algorithms, O'Reiley

<https://www.oreilly.com/library/view/fundamentals-of-deep/9781491925607/>



- **AI (Artificial Intelligence) = Intelligent Machine**
- **Q. Can we emulate human brain on our computer system?**



Brain is, inherently, what makes us intelligent.

- **Dream of building intelligent machine with brains like ours is ...**
- **We have to develop a radically difference programming a computer using techniques largely developed over the past decade.**

The limits of Traditional ways

- **Traditional Computer program**

- can do the trick for determined or non-vague issues.
- How to recognize a messy 0 from a 6 ?
- How do we write a program for that ?



← 6 or 0 ?

You can document the rules one after another. For example, points, angles, and rounds (radii) of the writing/drawing character.

Very famous MNIST handwritten letter test

Different approach

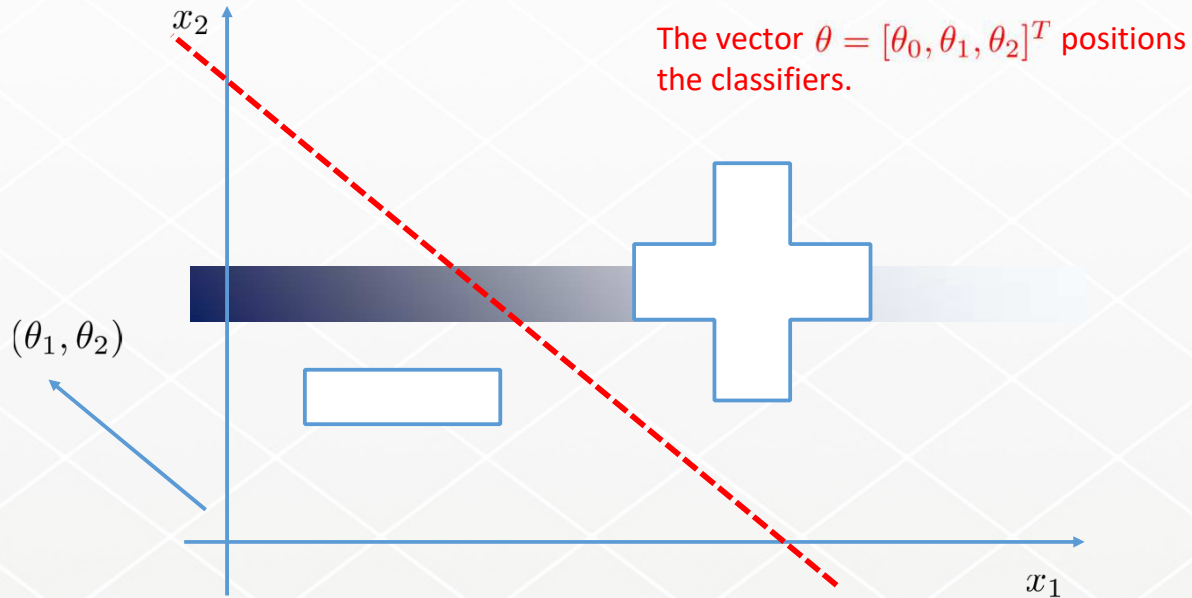
- **As human learn a lot of things at school,**
 - how to multiply numbers, solve equations, derivatives, and further
 - The things we find most natural, are learned by **EXAMPLES**, not by formula.
 - Deep Learning = subset of a more general field of AI, machine learning.

Linear Perceptron

$$h(\mathbf{x}, \theta) = \begin{cases} -1 & \text{if } \mathbf{x}^T \begin{bmatrix} \theta_1 \\ \theta_2 \end{bmatrix} + \theta_0 < 0 \\ +1 & \text{if } \mathbf{x}^T \begin{bmatrix} \theta_1 \\ \theta_2 \end{bmatrix} + \theta_0 \geq 0 \end{cases}$$

Linear Perceptron


$$h(\mathbf{x}, \theta) = \begin{cases} -1 & \text{if } \mathbf{x}^T \begin{bmatrix} \theta_1 \\ \theta_2 \end{bmatrix} + \theta_0 < 0 \\ +1 & \text{if } \mathbf{x}^T \begin{bmatrix} \theta_1 \\ \theta_2 \end{bmatrix} + \theta_0 \geq 0 \end{cases}$$



The Neuron

- ***The foundational unit of the human brain is the neuron. A tiny piece of the brain, about the size of grain of rice, contains over 10,000 neurons, each of which forms an average of 6,000 connections with other neurons. It's this massive biological network that enables us to experience the world around us.***


N. Buduma, Fundamentals of Deep Learning, designing next-generation machine intelligence algorithms



https://en.wikipedia.org/wiki/File:Blausen_0657_MultipolarNeuron.png

Neurons ← Linear Perceptrons

- **Modeling a neuron as a network [1943, Warren, McCulloch, Pitts], $y=f(z)$, $z=(w,x)+b$.**

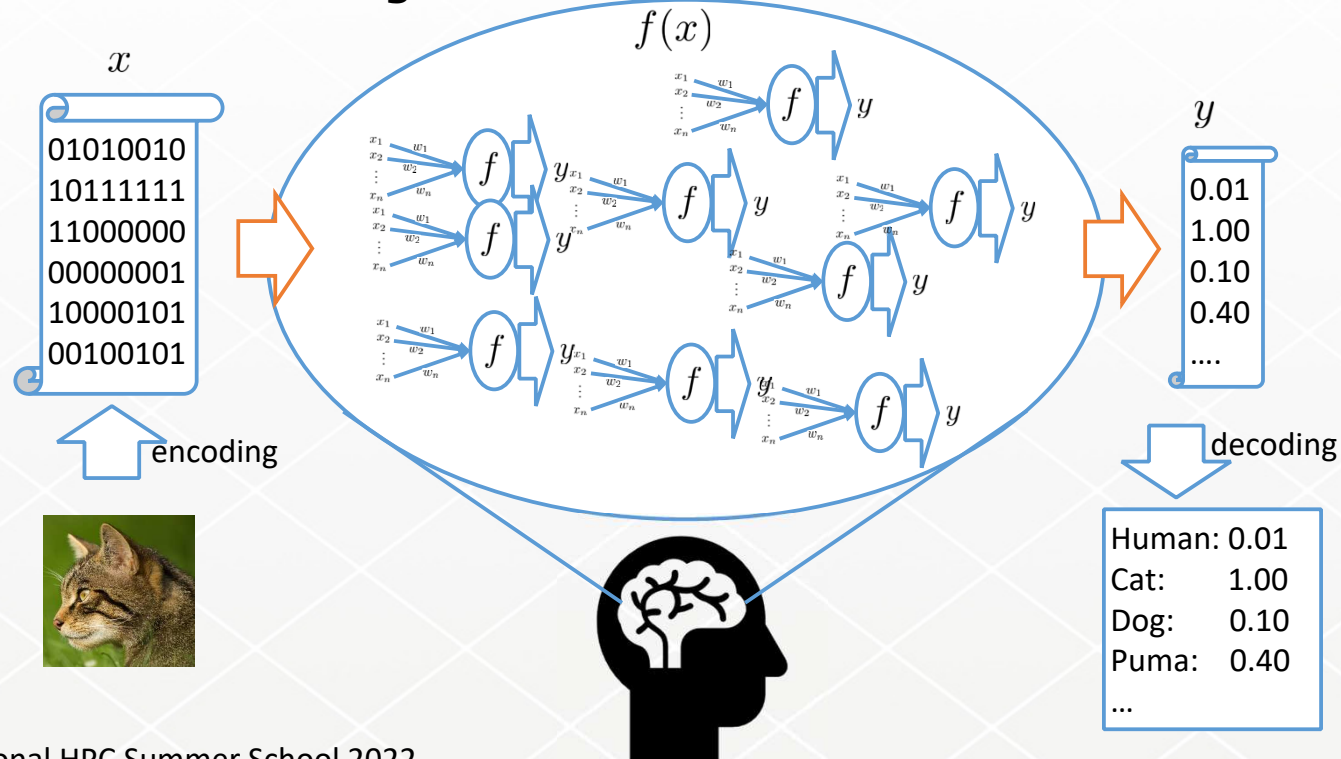


$$f(z) = \begin{cases} -1 & \text{if } z < 0 \\ +1 & \text{if } z \geq 0 \end{cases}$$

- **Neurons in the human brain are layered.**
 - The human cerebral cortex → six layers.
 - The simplest network is called feed-forward network
 - Linear Perceptron has easiness but limitation in the capability to express the hidden layers, which are sandwiched between the first input and the last output layers.
 - We need to employ some sort of non-linearity in order to learn the complex relationships.

Short Summary of Neural Network (NN)

- We need to **find a very complex function, or its parameter set** that approximates the target model.




Short break

- Which function should we use or are available for neuron $f(z)$ (namely, activation function) ?

Sigmoid	Tanh
$\frac{1}{1 + e^{-z}}$	$\tanh z$

ReLU	softmax
$\max(0, z)$	$\frac{e^{z_j}}{\sum e^{z_j}}$



Training Feed-Forward network

- **Supposed:**

- We have a large number set of training examples.
- Also, we can calculate the return value/vector of the neural-network when we input a training set.

- **Definition of Error:**

- By SSE (or MSE) function as follows

$$E = \frac{1}{2} \sum_i \left(t^{(i)} - y^{(i)} \right)^2$$

$t^{(i)}$: i-th training dataset

$y^{(i)}$: corresponding output value

- **Gradient descent direction:**

- If we plot a contour map of the Error function, and walk down



Step down to the descent direction

$$-\nabla E$$

We can minimize \square step by step

$$\square \leftarrow \square - \eta \frac{\partial E}{\partial \square} = \square - \Delta \square$$

Training Feed-Forward network

- **How to Change weight parameters:**

- Calculate partial derivative of the error functions with respect to each of the weights.

$$\Delta w_k = -\epsilon \frac{\partial E}{\partial w_k} = \sum_i \epsilon \left(t^{(i)} - y^{(i)} \right) \frac{\partial y^{(i)}}{\partial w_k}$$

$$E = \frac{1}{2} \sum_i \left(t^{(i)} - y^{(i)} \right)^2$$

$$\square \leftarrow \square - \eta \frac{\partial E}{\partial \square} = \square - \Delta \square$$

$$z = \sum_k w_k x_k, y = f(z) \Rightarrow \frac{\partial y}{\partial w_k} = \frac{\partial y}{\partial z} \frac{\partial z}{\partial w_k} = f'(z) x_k$$

- If we use sigmoidal neurons; $y=f(z)=1/(1+\exp(-z))$

$$\Delta w_k = -\epsilon \frac{\partial E}{\partial w_k} = \sum_i \epsilon \left(t^{(i)} - y^{(i)} \right) \left(y^{(i)} (1 - y^{(i)}) \right) x_k^{(i)}$$

- By using the obtained rule for the weights, we modify all the weights step by step.

Framework of the backward calculation

- If only output layer is considered for the Error:

$$\frac{\partial E}{\partial y_j} = -(t_j - y_j)$$

- Next, at j-th layer. We calculate how the outgoing j-th layer y_j affects incoming i-th layer (i→j)

- Similar to the output layer, we have following relation

$$\frac{\partial E}{\partial y_i} = \sum_j \frac{dz_j}{dy_i} \frac{\partial E}{\partial z_j} = \sum_j w_{ij} \frac{\partial E}{\partial z_j} = \sum_j w_{ij} f'_j(z_j) \frac{\partial E}{\partial y_j}$$

- Next, $\partial E/\partial w_{ij}$ is obtained by compound expressions with chain rule

$$\frac{\partial E}{\partial w_{ij}} = \frac{\partial z_j}{\partial w_{ij}} \frac{\partial E}{\partial z_j} = \frac{\partial z_j}{\partial w_{ij}} \frac{\partial y_j}{\partial z_j} \frac{\partial E}{\partial y_j} = y_{ij} f'_j(z_j) \frac{\partial E}{\partial y_j}$$

BP (Back-Propagation)


- As last slide, We know only the relation of the output data and test data.

$$\frac{\partial E}{\partial y_j} = -(t_j - y_j)$$

- Therefore, the method of calculating the error in the middle layer from the error in the final output layer (backward) to initial is significant, called “backpropagation”

Chain rule:
$$\frac{\partial f}{\partial u} = \frac{\partial x}{\partial u} \frac{\partial f}{\partial x} + \frac{\partial y}{\partial u} \frac{\partial f}{\partial y} + \dots$$

$$f = f(x(u, v, \dots), y(u, v, \dots), \dots)$$



$$\frac{\partial E}{\partial z_i} = \sum_j \frac{da_{1j}}{dz_i} \frac{\partial E}{\partial a_{1j}} = \sum_j \frac{da_{1j}}{dz_i} \frac{\partial y_j}{\partial a_{1j}} \frac{\partial E}{\partial y_j} = \sum_j W_{1ji} \nabla f_{1j}(a_1) (y_j - t_j)$$

$$\nabla_z E = W_1 \Delta y, \quad \Delta y = \nabla f_1(a_1) \odot (y - t)$$

$$\frac{\partial E}{\partial x_i} = \sum_j \frac{da_{0j}}{dx_i} \frac{\partial E}{\partial a_{0j}} = \sum_j \frac{da_{0j}}{dx_i} \frac{\partial z_j}{\partial a_{0j}} \frac{\partial E}{\partial z_j} = \sum_j W_{0ji} \nabla f_{0j}(a_0) \frac{\partial E}{\partial z_j}$$

$$\nabla_x E = W_0 \Delta z, \quad \Delta z = \nabla f_0(a_0) \odot \nabla_z E$$

- We obtained the gradient of E in terms of z and x.

$$\nabla_z E = W_1 \Delta y, \quad \Delta y = \nabla f_1(a_1) \odot (y - t)$$

$$\nabla_x E = W_0 \Delta z, \quad \Delta z = \nabla f_0(a_0) \odot \nabla_z E$$

- From them, we update weighted factors {W, b}.

$$\Delta W_{1ij} = -\epsilon \frac{\partial E}{\partial W_{1ij}} = -\epsilon z_i \Delta y_j \quad \Rightarrow \quad \Delta W_1 = -\epsilon z \Delta y^\top$$

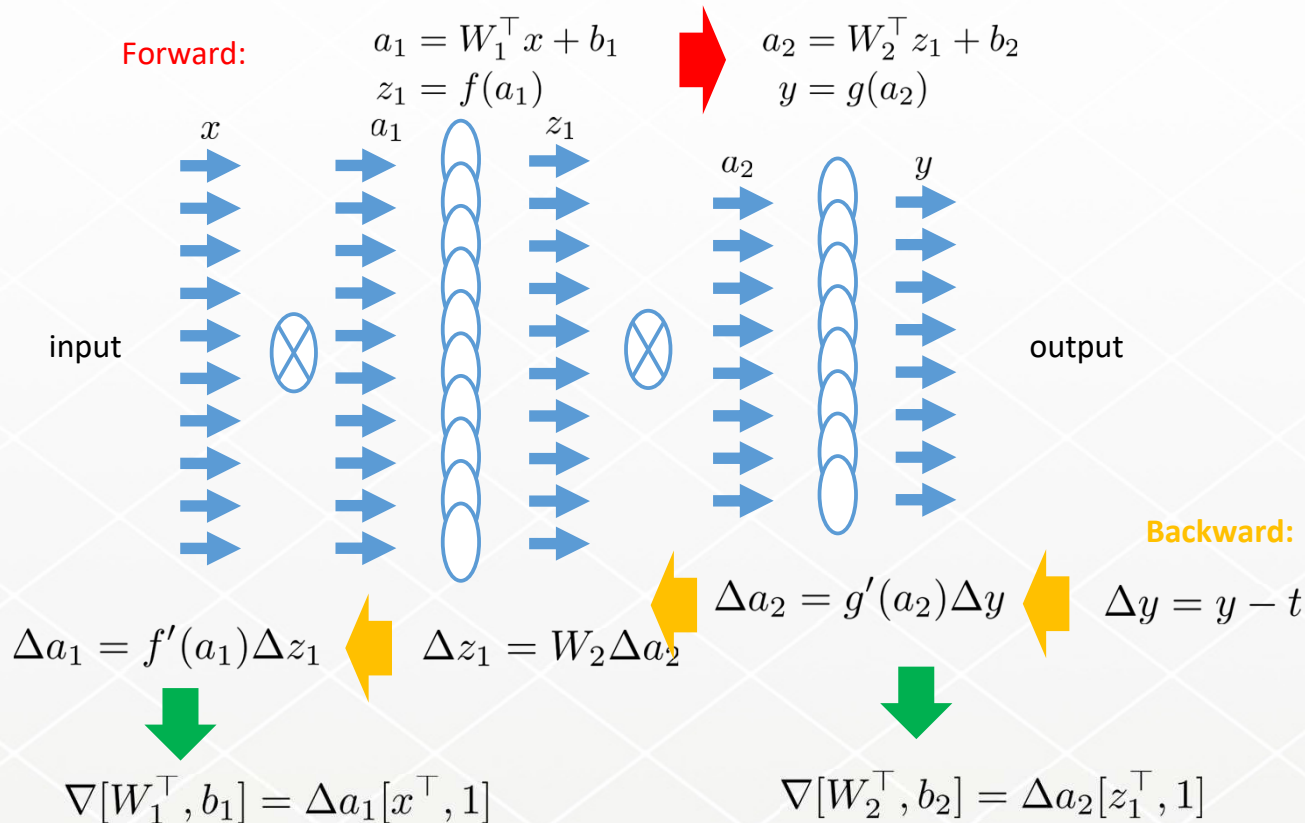
$$\Delta W_{0ij} = -\epsilon \frac{\partial E}{\partial W_{0ij}} = -\epsilon x_i \Delta z_j \quad \Rightarrow \quad \Delta W_0 = -\epsilon x \Delta z^\top$$

$$\Delta b_{1i} = -\epsilon \frac{\partial E}{\partial b_{1i}} = -\epsilon \Delta y_i \quad \Rightarrow \quad \Delta b_1 = -\epsilon \Delta y$$

$$\Delta b_{0i} = -\epsilon \frac{\partial E}{\partial b_{0i}} = -\epsilon \Delta z_i \quad \Rightarrow \quad \Delta b_0 = -\epsilon \Delta z$$

← Please confirm them by yourselves

Summary [Schematics on a simple I-O layer case]



Short break and demonstration

--MNIST handwritten letters and fashion--

The MNIST database

<http://yann.lecun.com/exdb/mnist/>

THE MNIST DATABASE
of handwritten digits

Yann LeCun, Courant Institute, NYU
Corinna Cortes, Google Labs, New York
Christopher J.C. Burges, Microsoft Research, Redmond

The MNIST database of handwritten digits, available from this page, has a training set of 60,000 examples, and a test set of 10,000 examples. It is a subset of a larger set available from NIST. The digits have been size-normalized and centered in a fixed-size image.

It is a good database for people who want to try learning techniques and pattern recognition methods on real-world data while spending minimal efforts on preprocessing and formatting.

Four files are available on this site:

```
train-images-idx3-ubyte.gz: training set images (9812422 bytes)
train-labels-idx1-ubyte.gz: training set labels (28881 bytes)
t10k-images-idx3-ubyte.gz: test set images (1648877 bytes)
t10k-labels-idx1-ubyte.gz: test set labels (16488 bytes)
```

please note that your browser may uncompress these files without telling you. If the files you downloaded have a larger size than the above, they have been uncompressed by your browser. Simply rename them to remove the `gz` extension. Some people have asked me "my application can't open your image files". These files are not in any standard image format. You have to write your own (very simple) program to read them. The file format is described at the bottom of this page.


The original black and white (bilevel) images from NIST were size normalized to fit in a 28x28 pixel box while preserving their aspect ratio. The resulting images contain grey levels as a result of the anti-aliasing technique used by the normalization algorithm. the images were centered in a 28x28 image by computing the center of mass of the pixels, and translating the image so as to position this point at the center of the 28x28 field.

With some classification methods (particularly template-based methods, such as SVM and K-nearest neighbors), the error rate improves when the digits are centered by bounding box rather than center of mass. If you do this kind of pre-processing, you should report it in your publications.

The MNIST database was constructed from NIST's Special Database 3 and Special Database 1 which contain binary images of handwritten digits. NIST originally designated SD-3 as their training set and SD-1 as their test set. However, SD-3 is much cleaner and easier to recognize than SD-1. The reason for this can be found on the fact that SD-3 was collected among Census Bureau employees, while SD-1 was collected among high-school students. Drawing sensible conclusions from learning experiments requires that the result be independent of the choice of training set and test among the complete set of samples. Therefore it was necessary to build a new database by mixing NIST's datasets.

The MNIST training set is composed of 30,000 patterns from SD-3 and 30,000 patterns from SD-1. Our test set was composed of 5,000 training set contained examples from approximately 250 writers. We made sure that the sets of writers of the training set and test set were

https://en.wikipedia.org/wiki/MNIST_database




Please access Google Colab.

Python program




Input figure



Input label

Neural Network



Output data

```
# forward  
a1 = np.dot(x, W1) + b1  
z1 = sigmoid(a1)  
a2 = np.dot(z1, W2) + b2  
y = softmax(a2)
```

Python code of
forward and back propagation
(Look at **two_layer_net.py**)

```
# back  
dy = (y - t) / batch_num  
grads['W2'] = np.dot(z1.T, dy)  
grads['b2'] = np.sum(dy, axis=0)  
dz1 = np.dot(dy, W2.T)  
da1 = sigmoid_grad(a1) * dz1  
grads['W1'] = np.dot(x.T, da1)  
grads['b1'] = np.sum(da1, axis=0)
```

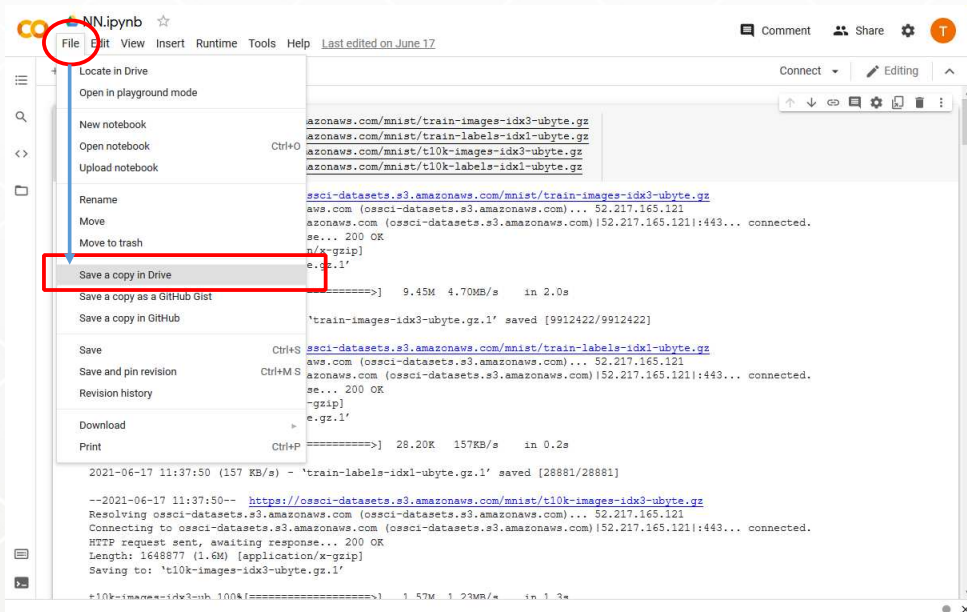
```
NUM_DATAS = 60000 # The number of input data  
NUM_TESTS = 10000 # The number of test data  
max_iter = 10000 # The number of iteration  
batch_size = 100 # The batch size mini batchlearning  
lr=0.1 # learning ratio
```

Look at **NN.py**

Link and copy

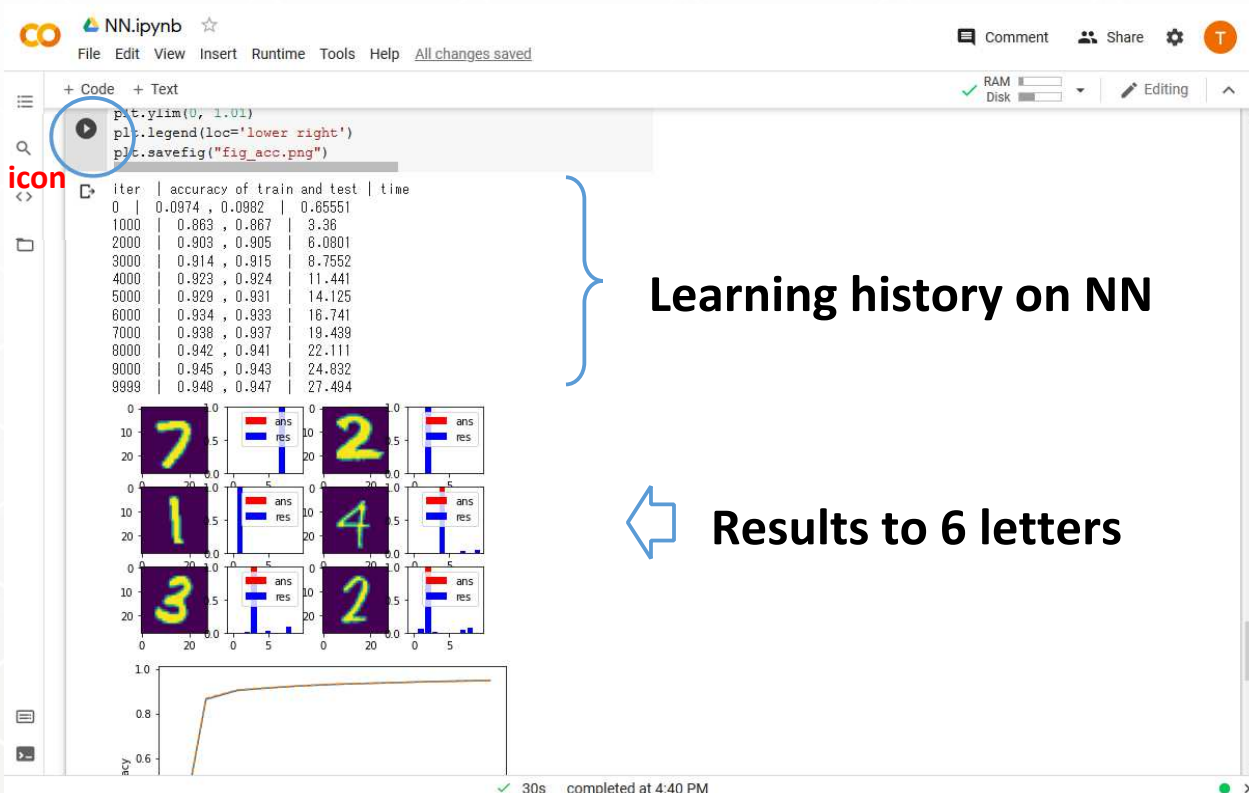
- **Click the link on Google Drive**

- <https://colab.research.google.com/drive/1sE42v9k0oJjweiowAJtRPb-h2R7h06K4?usp=sharing> , then copy it on your drive, in order to make the code editable by yourself.



Run = push the "play" icon

Play icon



The screenshot shows a Jupyter Notebook with the following code in the first cell:

```
plt.ylim(0, 1.01)
plt.legend(loc='lower right')
plt.savefig("fig_acc.png")
```

The second cell contains a table of training and testing accuracy over 10,000 iterations:

iter	accuracy of train	accuracy of test	time
0	0.0974	0.0982	0.65551
1000	0.863	0.867	3.36
2000	0.903	0.905	6.0801
3000	0.914	0.915	8.7552
4000	0.923	0.924	11.441
5000	0.929	0.931	14.125
6000	0.934	0.933	16.741
7000	0.938	0.937	19.439
8000	0.942	0.941	22.111
9000	0.945	0.943	24.832
9999	0.948	0.947	27.494

Below the table, there are six subplots arranged in a 3x2 grid. Each subplot shows a handwritten digit on the left and a bar chart on the right. The bar chart compares the predicted answer ('ans') and the actual result ('res'). The digits shown are 7, 1, 3, 2, 4, and 2. A blue arrow points from the text 'Results to 6 letters' to these subplots.

Learning history on NN

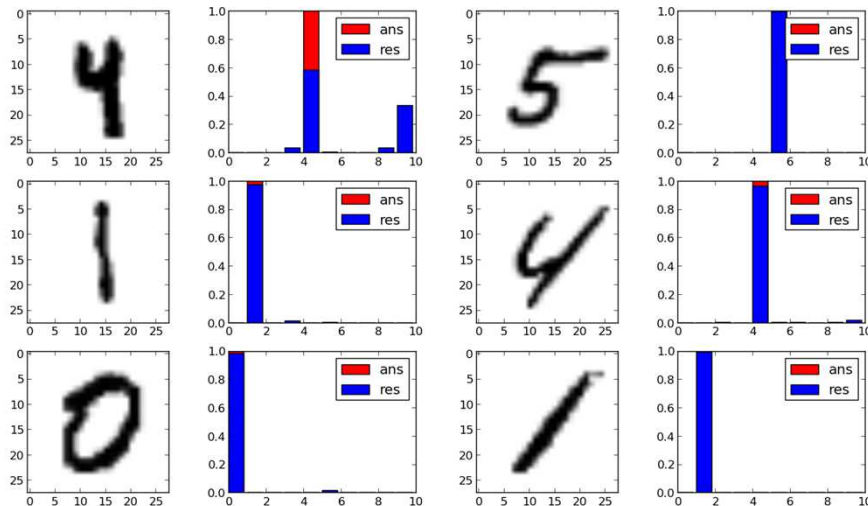
Results to 6 letters

30s completed at 4:40 PM

Example of result

Recognition fails

ans -> 0
max(res) -> 7



Success

ans = max(res)



Let run the Python program as follows:
\$ python NN.py

<https://www.nist.gov/itl/products-and-services/ernist-dataset>

<https://www.tensorflow.org/datasets/catalog/ernist>


Exercise 1:

- Please modify `load_mnist.py` as make available next lines and commented out other lines: `key_file` ...

```
"""  
key_file = {  
    'x_train': 'Fashion/train-images-idx3-ubyte.gz',  
    't_train': 'Fashion/train-labels-idx1-ubyte.gz',  
    'x_test': 'Fashion/t10k-images-idx3-ubyte.gz',  
    't_test': 'Fashion/t10k-labels-idx1-ubyte.gz'  
}  
"""
```

Another case (fashion MNIST)

- 0: T-shirt
- 1: Trouser
- 2: Pullover
- 3: Dress
- 4: Coat
- 5: Sandal
- 6: Shirt
- 7: Sneaker
- 8: Bag
- 9: Ankle boot



<https://github.com/zalando-research/fashion-mnist>

<https://www.tensorflow.org/tutorials/keras/classification>

Exercise 2 (Kuzushiji):

- Please modify `load_mnist.py` as make it available on the next lines and commented out other lines: `key_file`

```
# read at Kuzushiji data
key_file = {
    'x_train': 'Kuzushiji/train-images-idx3-ubyte.gz',
    't_train': 'Kuzushiji/train-labels-idx1-ubyte.gz',
    'x_test': 'Kuzushiji/t10k-images-idx3-ubyte.gz',
    't_test': 'Kuzushiji/t10k-labels-idx1-ubyte.gz'
}
```

Please see more information at <http://codh.rois.ac.jp/kmnist/>

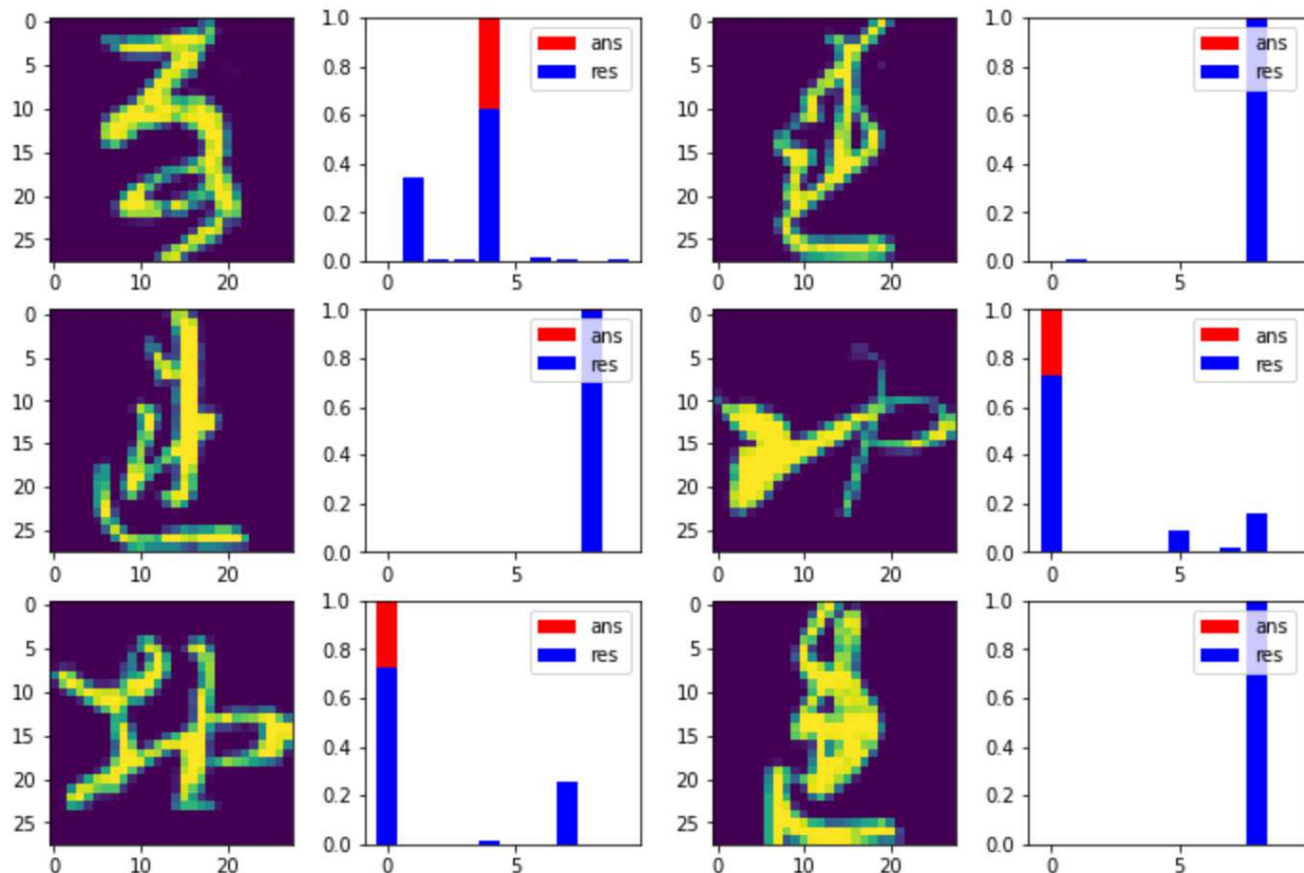
Advanced Exercise 1:

Please modify `TwoLayerdNetwork` as Three-or-more-layer or further complicated one, then confirm the accuracy.

Advanced Exercise 2:

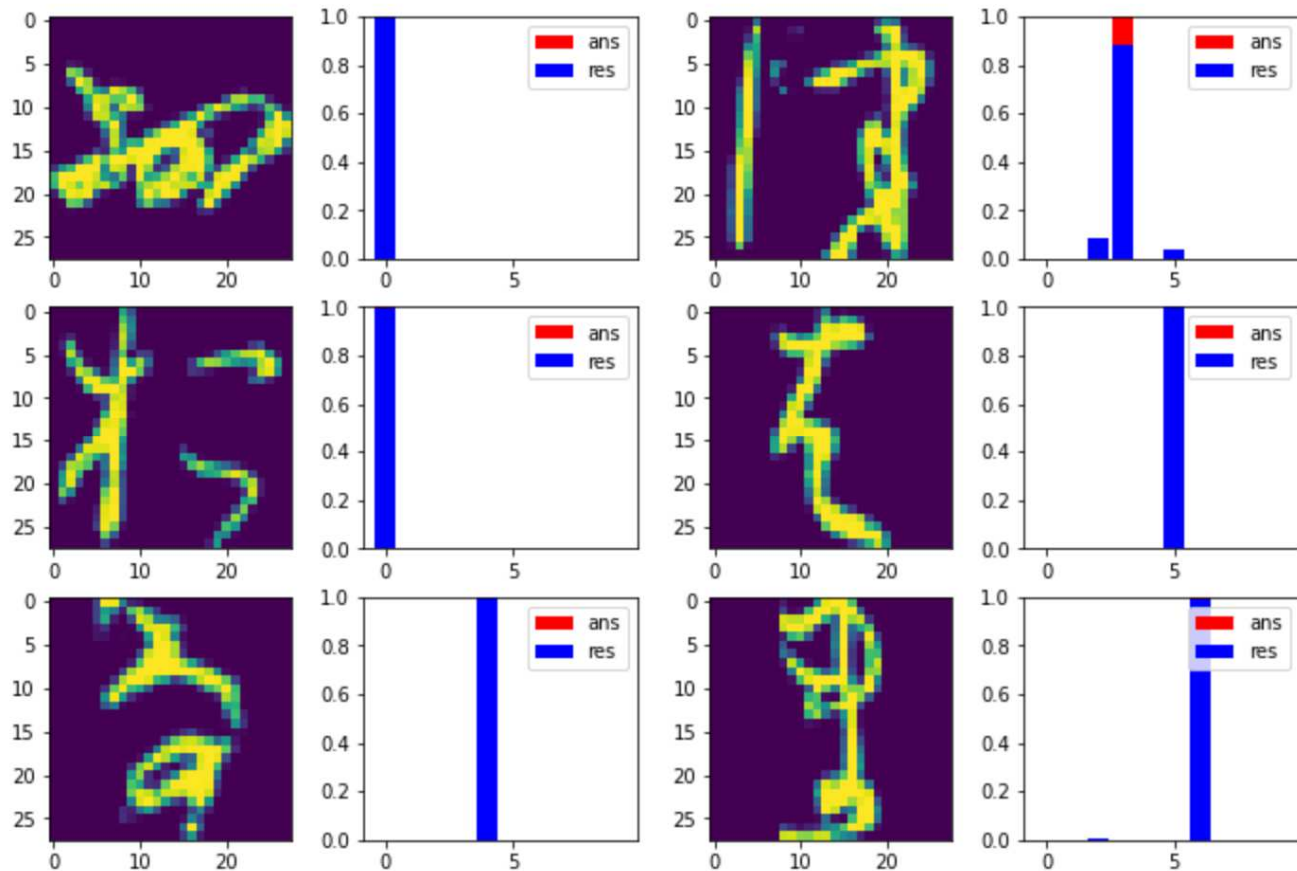
Please scan your handwritten digit characters, and put them into the test letters.

Kuzushi-ji recognition (sigmoid)




0: お
 1: き
 2: す
 3: つ
 4: な
 5: は
 6: ま
 7: や
 8: れ
 9: を

Kuzushi-ji recognition (relu-relu)



0: お
1: き
2: す
3: つ
4: な
5: は
6: ま
7: や
8: れ
9: を

Kuzushi-ji recognition (relu-relu)



- 0: お
- 1: き
- 2: す
- 3: つ
- 4: な
- 5: は
- 6: ま
- 7: や
- 8: ら
- 9: を

Exercise 2 (Kuzushiji):

- Please modify `load_mnist.py` as make available next lines and commented out other lines: `key_file ...`.

```
# read at Kuzushiji data
key_file = {
    'x_train': 'Kuzushiji/train-images-idx3-ubyte.gz',
    't_train': 'Kuzushiji/train-labels-idx1-ubyte.gz',
    'x_test': 'Kuzushiji/t10k-images-idx3-ubyte.gz',
    't_test': 'Kuzushiji/t10k-labels-idx1-ubyte.gz'
}
```

Please see more information at <http://codh.rois.ac.jp/kmnist/>

Advanced Exercise 1 (MIDEUM):

Please modify `TwoLayerdNetwork` as Three-or-more-layer or further complicated one, then confirm the accuracy.

Advanced Exercise 2 (HARD):

Please scan your handwritten digit characters, and use them into the test letters.

Further Optimization of $\{w\}$

- **Stochastic / batched approach**

- For every iteration, we update W by a full batch unit,

$$\Delta w_{ij} = -\epsilon \sum_{k \in \text{dataset}} y_i^{(k)} y_j^{(k)} (1 - y_j^{(k)}) \frac{\partial E^{(k)}}{\partial y_j^{(k)}}$$

- However, it might fail in local minimum or stagnation while GDM performs.
- On stochastic SDM, at each iteration, error surface is estimated only with respect to a single example.
- Major pitfall is a significant amount of iteration = time!
- Another Mini-batch SD, which prevents such pros, divides data set as a bunch of batch sets.

$$\Delta w_{ij} = -\epsilon \sum_{k \in \text{minibatch}} y_i^{(k)} y_j^{(k)} (1 - y_j^{(k)}) \frac{\partial E^{(k)}}{\partial y_j^{(k)}}$$


Training, Test, Validation, etc.

- **One problem:**

- The model becomes very complicated.
- What if we have a very complex model and a lot of training data sets, it is quite easy to fit the model.
- However, new complex model does not generalize well.


- **Over-fitting**

- Biggest challenge for the machine learning
- *On large network, overfitting is commonplace.*



MNIT-fashion

- Result of stagnated-learning history due to Overtraining.



Another Error function in NN.py

- **SSE: sum of squared error: explained**
- **Cross-Entropy**

$$E = - \sum_k (t_k \log y_k + (1 - t_k) \log(1 - y_k))$$

- **In the case of classification problem, E should be simplified, t is one-hot vector and normalized $|\mathbf{y}|=1$**

$$E = - \sum_k t_k \log y_k \quad \Rightarrow \quad \frac{\partial E}{\partial y_k} = -\frac{1}{y_k}$$

- **In NN.py, taking $\mathbf{y}=\mathbf{g}(\mathbf{z})$ as softmax()**

$$\frac{\partial y_j}{\partial a_{1j}} \frac{\partial E}{\partial y_j} = (y_j(t_j - y_j))(-1/y_j) = y_j - t_j = (\Delta y)_j$$

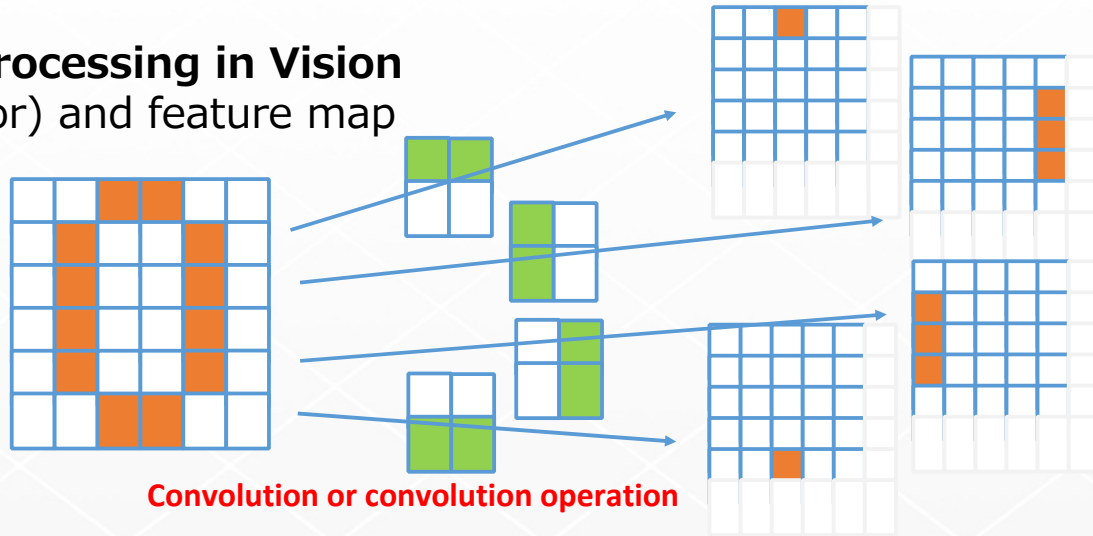
- **Cross-entropy error shows fast learning speed and better accuracy than SSE when classification.**

**For more advanced,
image recognition case for
practical uses**


-- a Bridge and introduction

Convolution Neural Networks

- **Information processing in Vision**
 - Filter (detector) and feature map

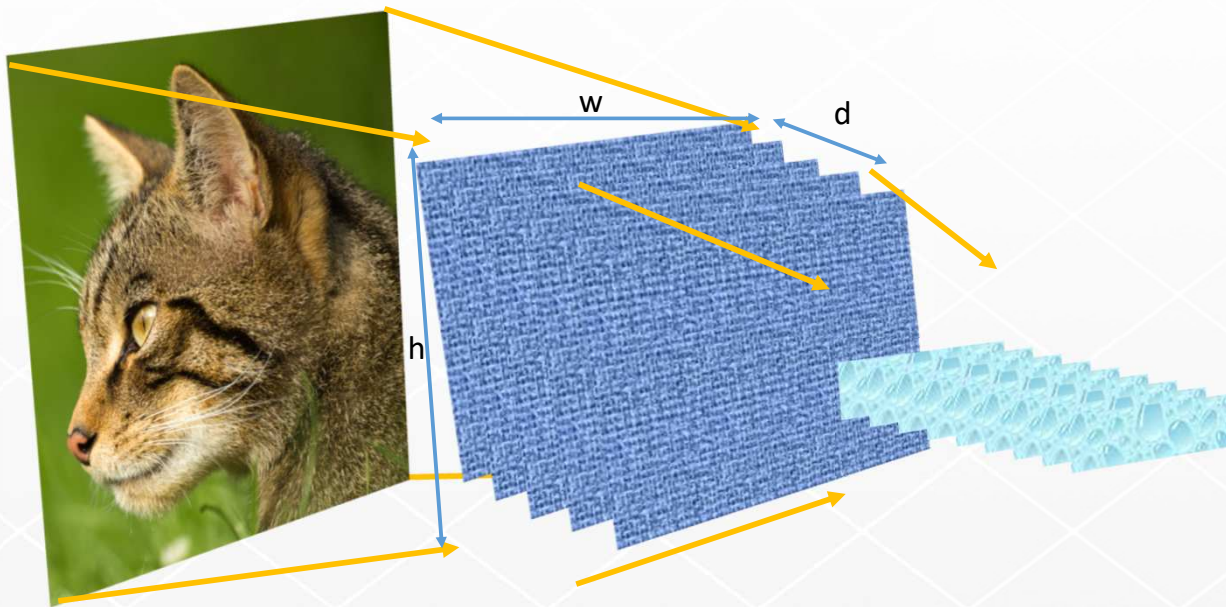


$$m_{ij}^k = f \left(\overbrace{(W * x)_{ij}}^{\text{bitop \& sum}} + b^k \right)$$



Schematics

- (w,h,d,p) : [width, height, depth, zero padding]



This network consists of a lot of parameters, part of them are called hyper-parameters. These are also updated by similar techniques as back-propagation and descend sweep. More flexible and user-friendly packages should be recommended for AI users.

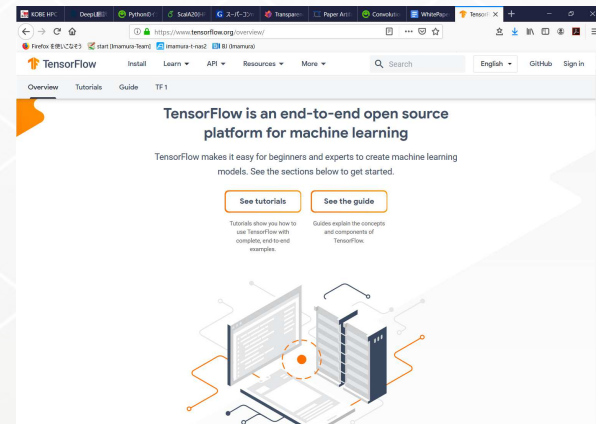
Appendix, More about One of the DL software

- **TensorFlow (AI-engine) and Keras (high-level APIs)**
 - <http://www.tensorflow.org/>

Building a deep and large deep learning model from scratch ..., one of the best or primary tool sets is TensorFlow by Google, 2015.

- Open Source Software
- Rich in tutorials
- Many examples
- A lot of available platforms

- Accelerated by specific HWs




PyTorch

- **What is PyTorch?**

- PyTorch provides a machine learning framework for several key target problems. It can provide a path to the cutting edge, starting from a prototype of AI research, especially with the concerted efforts of major AI vendors to optimize it.

- On Fugaku, Fujitsu and RIKEN have been collaborating to enhance the performance to take advantage of the brand-new processor A64FX.

<https://pytorch.org/>



Note: from this page to the end of the ppt, please do not distribute it for other than this school participants due to copyright restrictions and so on.

Prerequisite for PyTorch environment on Fugaku

- **PyTorch ver.1.7.0**

- Horovod ver.0.20.3
- oneDNN ver.2.1.0L01_aarch64

- **Python ver.3.8.2**

- mpi4py ver.3.0.3
- pandas ver.1.2.2
- numpy ver.1.19.0
- scipy ver.1.5.2
- h5py ver.2.8.0
- fapp ver.1.0.0

- If you want to build the PyTorch environment on your PC, read the Fugaku AI user guide!

- https://www.fugaku.r-ccs.riken.jp/doc_root/en/user_guides/FugakuAIGuide/
- https://www.fugaku.r-ccs.riken.jp/doc_root/en/user_guides/FugakuAIGuide/pytorch.html

Tutorial

- **This is a translation from the Japanese document into English since only the Japanese one is provided on the Fugaku portal site (please tell me if you find bugs and wrong interpretations).**
- **Below are the instructions for running the Resnet image recognition sample using PyTorch on Fugaku. To create an account and log in to Fugaku, please refer to the Startup Guide. We assume you are logged in and working in ~/tutorial.**
- **PyTorch is pre-installed below. Please follow the tutorial step by step!**
 - /vol0004/apps/oss/PyTorch-1.7.0
- **Next page is to copy an example to your workspace. The running example, 95_output, is large in size, so it is excluded from the copying. Please refer to 95_output if you want to see the preliminary result.**

Copy the examples to your workspace!

```
$ cd # move to $HOME
$ mkdir -p tutorial/PyTorch # make a working directory
$ cd /vol0004/apps/oss/PyTorch-1.7.0/example # Change the workspace into the PyTorch
$ tar -cf - --exclude=95_output --exclude=95_log . | (cd ~/tutorial/PyTorch ; tar -xvf -) # Copy examples excluding run results
$ cd # Back to $HOME
$ cd tutorial/PyTorch # Change directory to the working space
$ ls # Look at the working space
01_resnet 03_Bert env.src README.md
02_OpenNMT 04_Mask-R-CNN env.src.spack
```

- **Image recognition (Resnet), natural language processing (OpenNMT, Bert), and object detection (Mask-R-CNN) are included. The env.src file provides environment variables and other settings necessary to run PyTorch. If you run PyTorch on your own, you must run before running PyTorch as follows. (Do not forget type '.'(dot), it is a significant shortcut command on a linux-shell)**

\$. env.src

Run a Resnet example

- **Sample codes are for image recognition using the ResNet-50 network. Dummy data is used, and a lot of examples are provided such as for inference, training, distributed data parallelism, performance acquisition, use with Spack, and use of LLIO.**
- **The sample model setup is not necessarily the fastest. You can compare the results of the preliminary runs in Fugaku in `/vol0004/apps/oss/PyTorch-1.7.0/example/0*/95_output`.**

```
$ pjsub submit_train.sh
$ less submit_train.sh.(job_id).out
>> script option: Namespace(batch=256, itr=20, lr=0.001, momentum=0.9, trace=False, type='cpu_mklntensor', weight_decay=0.0)
## Start Training
[ 1] loss: 7.269 time: 2.555 s
[ 2] loss: 5.051 time: 2.320 s
[ 3] loss: 1.229 time: 2.303 s
[ 4] loss: 0.012 time: 2.303 s
[ 5] loss: 0.000 time: 2.304 s
[ 6] loss: 0.000 time: 2.304 s
[ 7] loss: 0.000 time: 2.303 s
[ 8] loss: 0.000 time: 2.303 s
[ 9] loss: 0.000 time: 2.303 s
(omitted)
```

- The larger the batch size, the more efficient the performance will be as available memory is used. However, since the peak CPU performance is not high (about 6.6TF for single precision), multi-process execution on many nodes is often required. `submit_train_multi.sh`, and `submit_val_multi.sh` are examples of data parallel learning and inference with 4 processes.
- They run a different model than the one used in the one-process test, but with a different parameter. `submit_val_multi.sh` performs inference in four parallel processes, but in the case of inference, each process runs asynchronously. The following is an example of the output of `submit_train_multi.sh` on Fugaku (2.2GHz). The four processes in total are running at more than 100img/sec.

(omitted)

Running benchmark...

Iter #0: 28.5 img/sec per CPU # Rank0 performance

Iter #1: 28.5 img/sec per CPU

Iter #2: 28.5 img/sec per CPU

Iter #3: 28.5 img/sec per CPU

Iter #4: 28.5 img/sec per CPU

Img/sec per CPU: 28.5 +-0.0

Total img/sec on 4 CPU(s): 114.0 +-0.1 # Total of four processes

LLIO_transfer for Python modules, etc.

- In the example, we only use a single node. However, we advise you to move the data to the first-level storage beforehand because the image data will increase. The Python modules take longer to load once the number of nodes exceeds 1,000 parallelisms in 3 racks.
- To reduce access to the second-level FEFS, run_tar.sh offers a sample script that tarballs a complete Python environment and extracts it on a compute node. However, since the file I/O associated with the tar expansion process is slow on the compute node, we recommend using llio_transfer to transfer only the necessary files.

```
$ rm -Rf strace_log          # 不要なファイルを消去します
$ pjsub submit_llio_pre.sh   # 転送すべきファイルをstraceを用いて検索します
$ egrep -v '=\$-1 ENOENT|O_DIRECTORY' ./strace_log/strace.0.* | egrep O_RDONLY | \
cut -d\$" -f 2 | egrep ^/vol.... >> llio_transfer.list # 出力されたstraceの結果から、転送するファイルリストを作成します
$ pjsub submit_llio_main.sh   # llio_transferを用いた本計算です
```

- **In `submit_llip_pre.sh`, `strace` is run only for rank 0 to retrieve the necessary files and output the accessed files. `strace` processing takes a bit of time, so the batch size and number of iterations are set low. The `strace` log creates a list of files to be transferred using `grep`. Here is an example of the created file list.**

```
$ cat llio_transfer.list
/vol0004/apps/oss/PyTorch-1.7.0/bin/python3
/vol0004/apps/oss/PyTorch-1.7.0/lib/libtcmalloc.so
/vol0004/apps/oss/PyTorch-1.7.0/lib/libpython3.8.so.1.0
/vol0004/apps/oss/PyTorch-1.7.0/lib/python3.8/encodings/__pycache__/__init__.cpython-38.pyc
(omitted)
```

- **In `submit_llio_main.sh`, the transfers the 1,000 or so files in the list to the first level before the calculation.**

```
$ cat ./llio_transfer.list | xargs -L 100 llio_transfer
```

Note on Fugaku AI platform

- **For acceleration, `dnnl_aarch64` (oneDNN-2.1.0L1), under Fujitsu Laboratories' development, is used. It is designed to speed up image recognition, natural language processing analysis, object detection, and other processes. Depending on the network used, there is a possibility that some functions may be faulty or may not work fast enough.**
- **Any requests for speed-up, supported versions, modules, etc., should be directed to the support desk of your user account with the network scripts you are using. Please note that it may take several months to respond to your request.**
- **These are commented on by the support Desk and Fujitsu Laboratories.**

Thank you for all.

**Please enjoy the rest of the time
for your review exercises in
MNIST and Resnet examples!**