中國醫藥大學附設醫院
China Medical University Hospital

Bone Age Estimation from Hand X-Rays Images Using NVIDIA DGX-1

Ken Ying-Kai Liao

Big Data Engineer
Artificial Intelligence Center for Medical Diagnosis
China Medical University Hospital
(T32209@mail.cmuh.org.tw)

Outline

1. Introduction: Medical Imaging and Artificial Intelligence
2. Sharing a DGX-1 using JupyterHub
3. NVIDIA GPU Cloud Docker Images to Jupyter Notebook Images
4. Case Study: Bone Age Estimation from X-Ray Images
    a. Clinical Problem Definition
    b. Data Cleaning and Preprocessing
    c. Hyperparameter Tuning
    d. DGX-1 Speedup
    e. Current Results
5. Conclusion: Future Work

Medical Imaging and Artificial Intelligence

Advantages:

- Accurate?
- Fast
- Objective

Disadvantages:

- Accurate?
- What is the AI actually learning? (Black Box?)

Medical Imaging and Artificial Intelligence

中國醫藥大學附設醫院
China Medical University Hospital

Sharing a DGX-1 using JupyterHub

- To speed up AI research at CMUH, our director purchased a DGX-1

- Problem:
  - DGX-1 best practices state to use Docker Images for running deep learning applications
  - Our research team is used to directly coding on Jupyter Notebooks
- Solution:
  - JupyterHub + Jupyter Notebook Docker Images with GPU Selection

Sharing a DGX-1 using JupyterHub

- Set up JupyterHub to allow selection of GPUs
- This allows GPU sharing between engineers

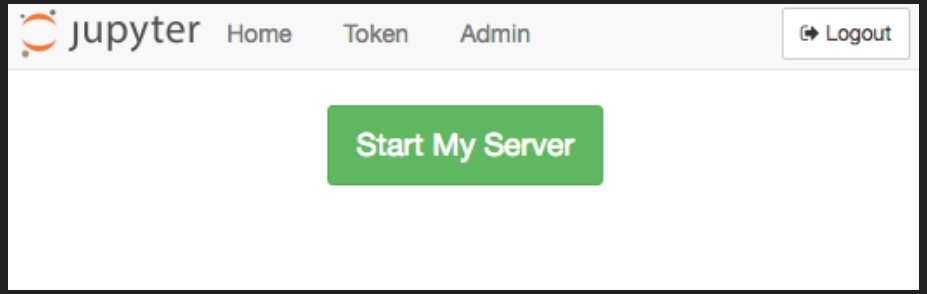

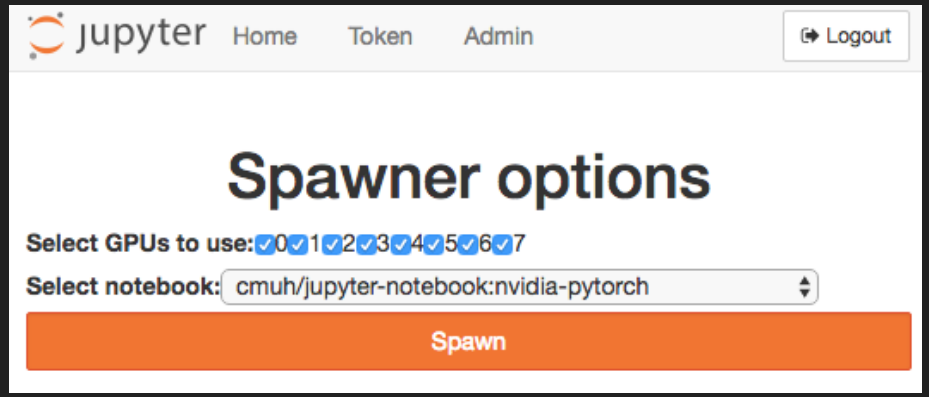

NVIDIA GPU Cloud Docker Images to Jupyter Notebook Images

- Use NGC Containers as Base
- Add Jupyter
- NGC + Jupyter = Fast Data Science Prototyping!

```
FROM nvcr.io/nvidia/tensorflow:18.04-py3

RUN python3 -m pip install --upgrade pip
RUN python3 -m pip install --upgrade keras opencv-python pyt
    numpy scikit-learn scikit-image jupyter jupyterlab jupyt
RUN python3 -m pip install --upgrade h5py pycuda

# Configure container startup
ENTRYPOINT ["tini", "--"]
CMD ["start-notebook.sh"]
EXPOSE 8888
```

Case Study: Bone Age Estimation from X-Ray Images

Clinical Problem Definition

- VP Dr. Tsai wants us to use AI to look at bone age
  - Actual Age?
  - Bone Age?
  - Comparison can show us whether there is a growth problem!
  - Estimate Ultimate Height

Clinical Problem Definition

- Doctors compare X-Ray images to an Atlas
  - Slow
  - Time Consuming
  - Results are Subjective

Clinical Problem Definition

- With AI Bone Age Reading
  - Fast (Especially with NVIDIA GPU!)
  - Objective rather than subjective

Clinical Problem Definition

- We have our problem, now what?
  - Hospital IT Department:
    - Here's some excel sheets and images!

Good Images

OK Image...

??????????????????????


RJP


L
pt


15

Our First Image Cleaning Task...

- Go through and remove all the images without hands
- This means looking at 19829 images manually…
- We ended up with 7107 that has bone age and "hand" images

Data Cleaning and Preprocessing

Image Cleaning

- Black on White? White on Black?
  - Normalize to White on Black!
- Different Images Sizes
  - Resize and pad with black to keep aspect ratio

Data Cleaning and Preprocessing

Label Cleaning

Boneage: 16
Boneage: 12-13
Boneage: 5
Boneage: :
Boneage: 8
Boneage: 13 +
Boneage: 16
Boneage: 4
Boneage: 8
Boneage: 13.5
Boneage: 11-12
Boneage: 10
Boneage: 13.5-14
Boneage: 2
Boneage: 10-11
Boneage: 12
Boneage: 15.5
Boneage: 13.5
Boneage: 16
Boneage: 3-4
Boneage: 15+
Boneage: -- 8-9
Boneage: 16
Boneage: 2
Boneage: 12
Boneage: 7-8
Boneage: 11
Boneage: N/A
Boneage: 13
Boneage: 6-7
Boneage: 13.5
Boneage: 10
Boneage: --15.5-16
Boneage: 14
Boneage: 11
Boneage: 9
Boneage: > 16
Boneage: 16
Boneage: 13.5
Boneage: (2009/09/24)
Boneage: 5-6
Boneage: 7-8
Boneage: 9
Boneage: 9-10

Task 2, Label Cleaning

- Ask Doctors if they can help us provide clean labels
  - They said OK, but only if they have a nice interface to do it on

Data Cleaning and Preprocessing

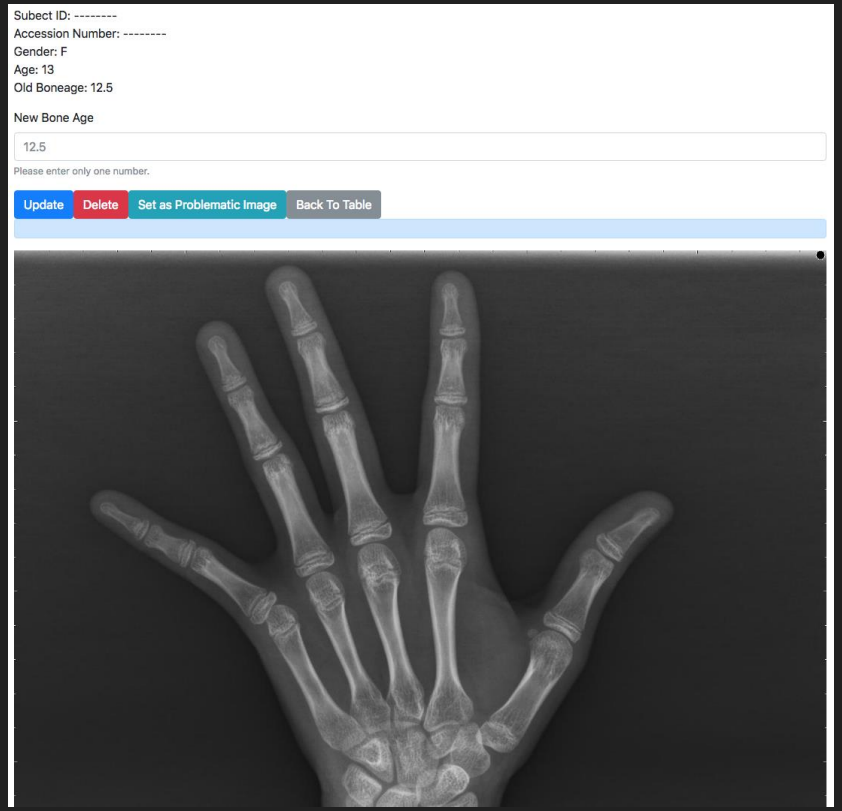

A few days later ...

Label Cleaning

- Dr. Tsai has gone through more than 5000 images, and added clean labels!
- Now we can start hyperparameter tuning!

Hyperparameter Tuning

Base Network?

- ResNet
- Inception
- Inception-ResNet
- DenseNet
- NASNet

Layers?

- #

Optimizer?

- SGD
- RMSProp
- Adam
- Adadelta

Learning Rate?

- Constant
- Decay
- Schedule

Image Size?

- 128x128
- 256x256
- 512x512

Loss Function?

- Log Loss (Categorical)
- MAE
- MSE

...More?

DGX-1 Speedup

Using InceptionResNetV2 as an example

CPU	8 x GPU (DGX-1)
Per Epoch: 980s	Per Epoch: 20s
Epochs to Convergence: 128	Epochs to Convergence: 128
Time to Convergence: 35 Hours	Time to Convergence: 42 Minutes

DGX-1 Speedup

- DGX-1 Speedup = 50x FASTER!
- We can try 50 different hyperparameters in the same amount of time!

Current Results

Models: Separate models for Male and Female subjects

Cleaned Data: 2739 Male, 4368 Female

- Using InceptionResNet-V2 Base Network
  - MSE as Loss Function
  - Adadelta as Optimizer

Current Results

Female

- Label Cleaning 76% complete by CMUH doctors
- Train / Validation: 80% / 20%
- Validation Score:
  - Prediction Accuracy within 36 Months: 99%
  - Prediction Accuracy within 24 Months: 97%
  - Prediction Accuracy within 12 Months: 75%

Current Results

Male

- Label Cleaning 100% complete by CMUH doctors
- Train / Validation: 80% / 20%
- Validation Score:
  - Prediction Accuracy within 36 Months: 100%
  - Prediction Accuracy within 24 Months: 99%
  - Prediction Accuracy within 18 Months: 97%
  - Prediction Accuracy within 12 Months: 90%
  - Prediction Accuracy within 06 Months: 65%

Current Results

- Cleaner Data = Higher Accuracy
- More Data = Higher Accuracy
- Training can be done REALLY FAST through NVIDIA DGX-1
- We can also visualize the results

Current Results (Saliency Map)

Bone Age: 5.0

Predicted 4.967061996459961

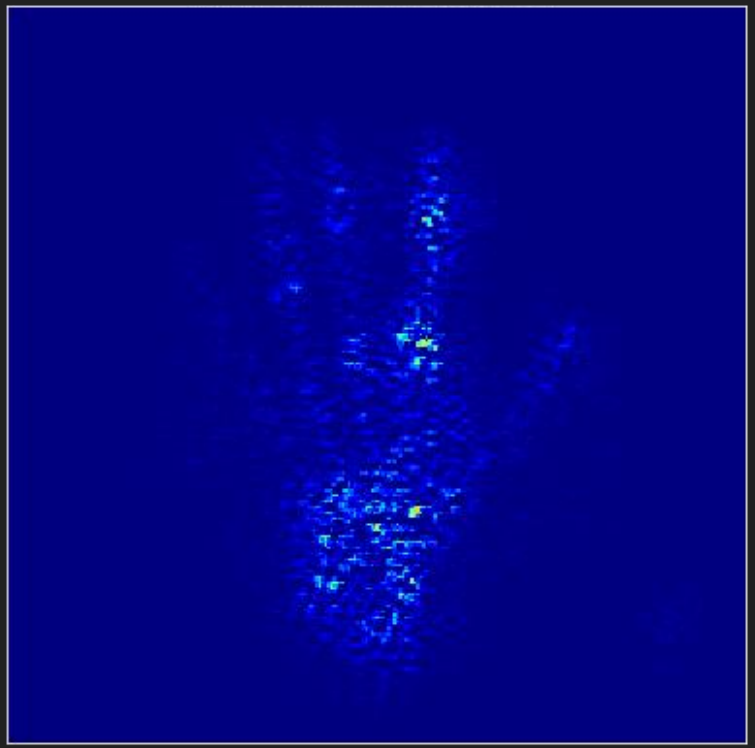

Current Results (Saliency Map)

Bone Age: 8.0

Predicted 7.9445905685424805

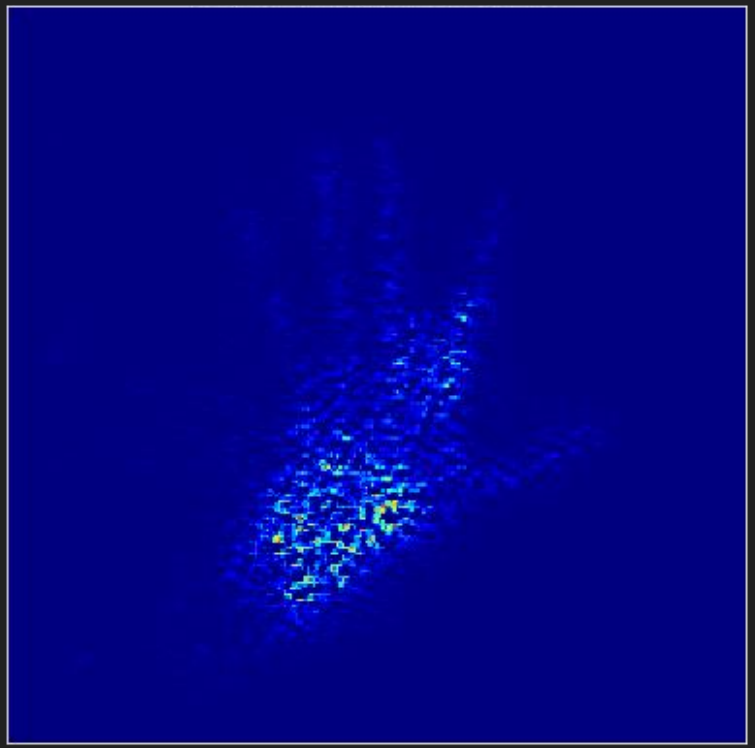

Current Results (Saliency Map)

Bone Age: 14.5

Predicted 14.389472007751465

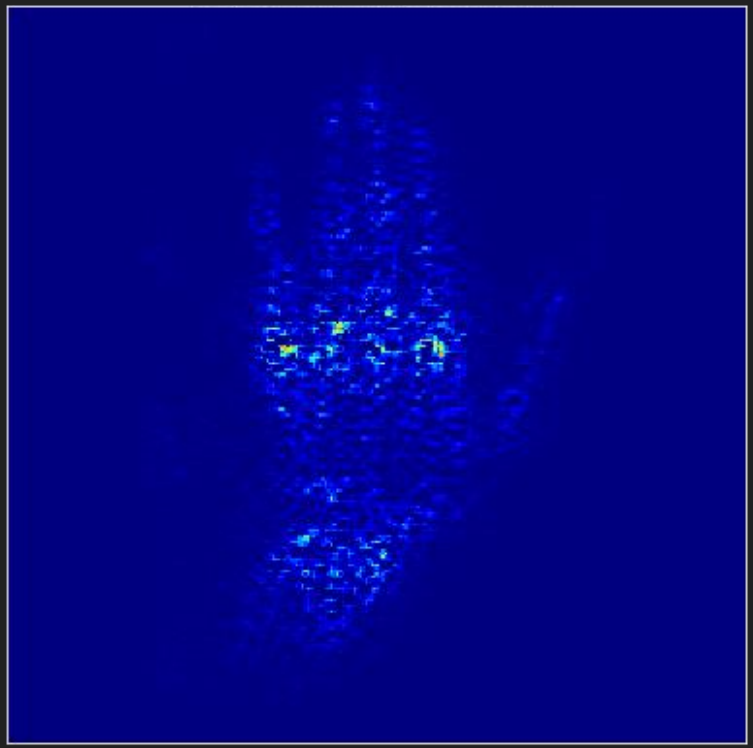

Conclusion: Future Work

For GPU sharing:

- More DGX-1 → JupyterHub Docker → Kubernetes

To increase accuracy:

- Clean Data!
- More Data!
- More Hyperparameters!

中國醫藥大學附設醫院
China Medical University Hospital

Thank You

Ken Ying-Kai Liao

Big Data Engineer
Artificial Intelligence Center for Medical Diagnosis
China Medical University Hospital
(T32209@mail.cmuh.org.tw)