BIG DATA ANALYSIS 2019

PREDICTING BONE AGE FROM X-RAY IMAGES

ALISA AGAFONOVA, JON RAUNKJÆR SØNDERGAARD,
YIFAN LIU VESTERGAARD & MADS EHRHORN KJÆR

https://imgur.com/gallery/c1ZUTUx

PROBLEM

- Medical images low hanging fruit for ML
- Machines might be able to better classify images

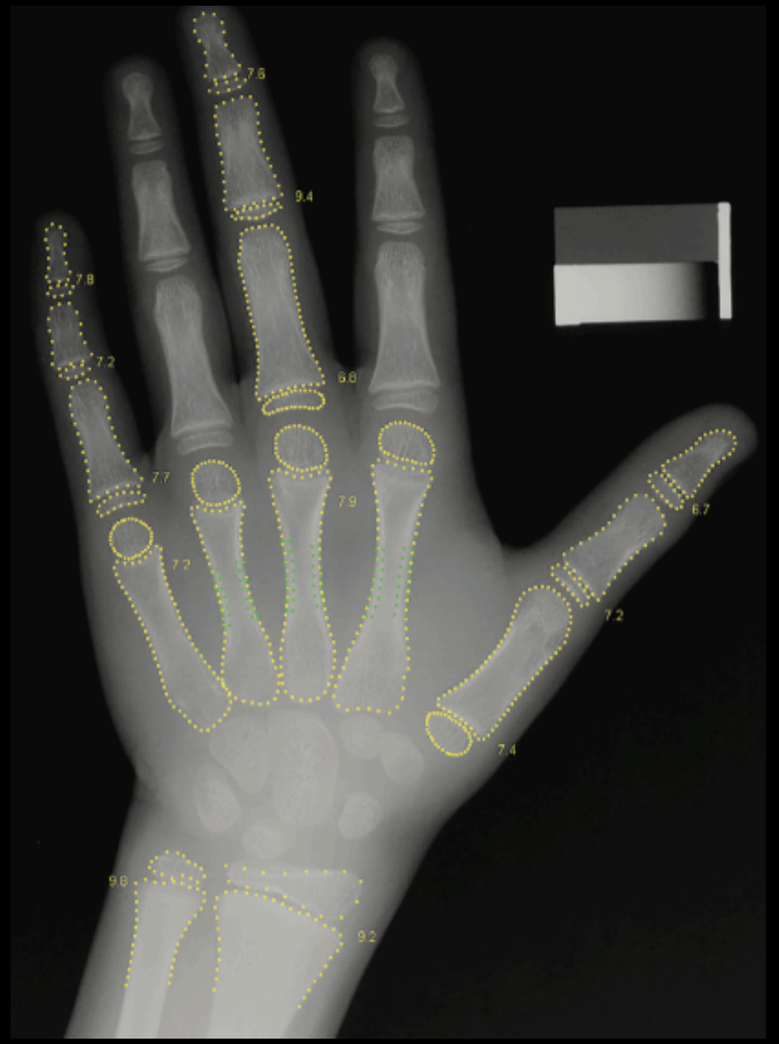

https://doi.org/10.1148/radiol.2018180736

“If you work as a radiologist you're like the coyote that's already over the edge of the cliff but hasn't yet looked down so it doesn't realize there's no ground underneath him.”

–PROF. GEOFFREY HINTON

https://www.vox.com/2015/3/5/8157519/chuck-jones-rules-for-roadrunner-coyote

BIG F*#K'IN DATA

- Images in full resolution ~10 GB
- Difficult to fit in RAM
- Need GPU for training
- Ie. you're not gonna use your laptop
- Solution:
  - Kaggle Kernels; GPU, commit history
  - Google Colab; GPU, 8-core TPU

https://www.amazon.com/Sarcasm-Engineer-Profession-Typography-Coffee/dp/B077L52266

RSNA BONE AGE CHALLENGE

- ~12600 X-ray images
- Data given as (plus files) —>
- Radiologists MAE ~7 months
- Winners MAE ~4 months

id	boneage	male
1377	180	False
1378	12	False
1379	94	False
1380	120	True
1381	82	False
1382	138	True
1383	150	True
1384	156	True

DISTRIBUTION

- 18.2% more data for male than female
- Age ranges randomly underrepresented (data is not beautiful)
- Splitting strategy:
  - Stratify (Validation/Training)
  - Multiple sampling (to achieve uniform distribution also in gender)

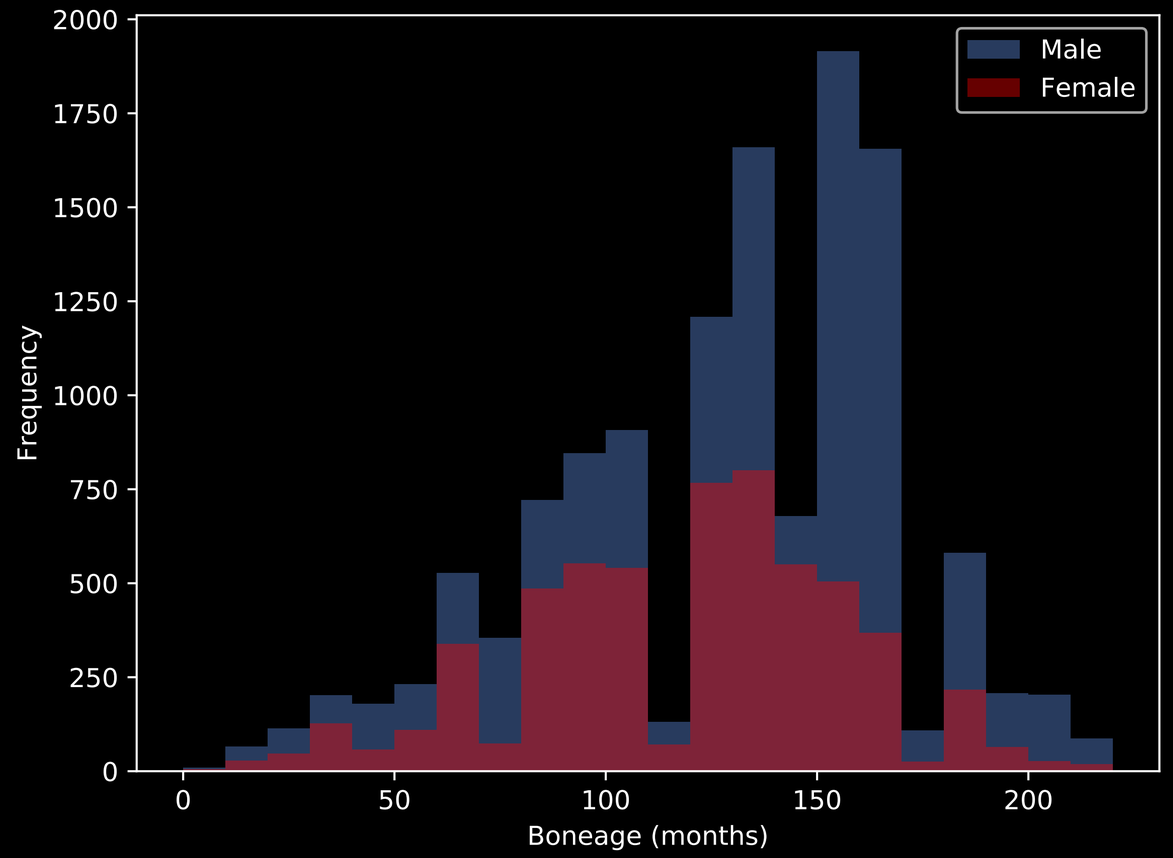

PROCESSING AND FLOW

- No Pre-processing (data is ugly)
- Transform
  - Rescale (299 by 299)
  - Rotate (20 degrees)
  - Shift (20% of total width and height)
  - Zoom (0.8-1.2x)
  - Flip (horizontal)
  - Shear (0.2 degrees)
- Hard to fit all images in RAM
  - Use flow; NN is fed images a batch at a time
  - Current method takes 10+GB of RAM out of 13GB

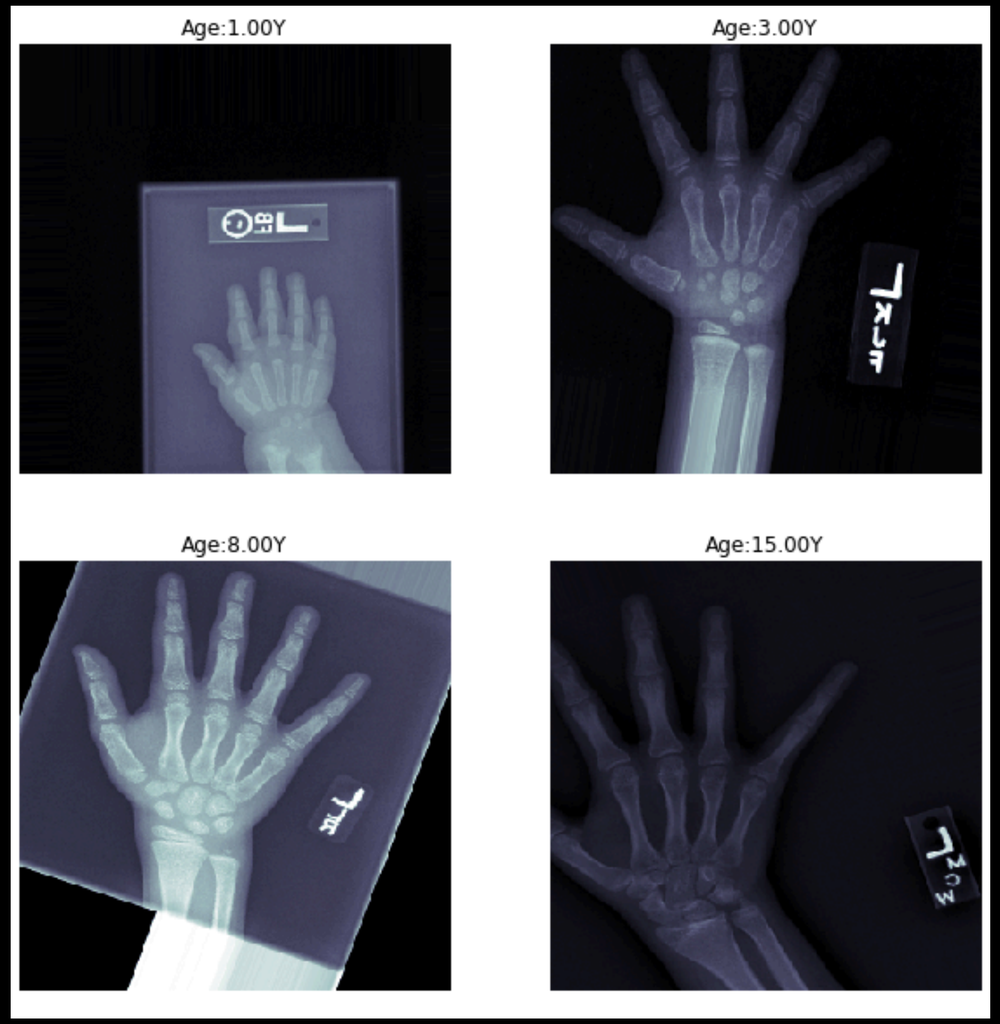

OUR MODEL

- Convolutional Neural Network
  - Multiple inputs
  - Using pretrained network
  - Attention

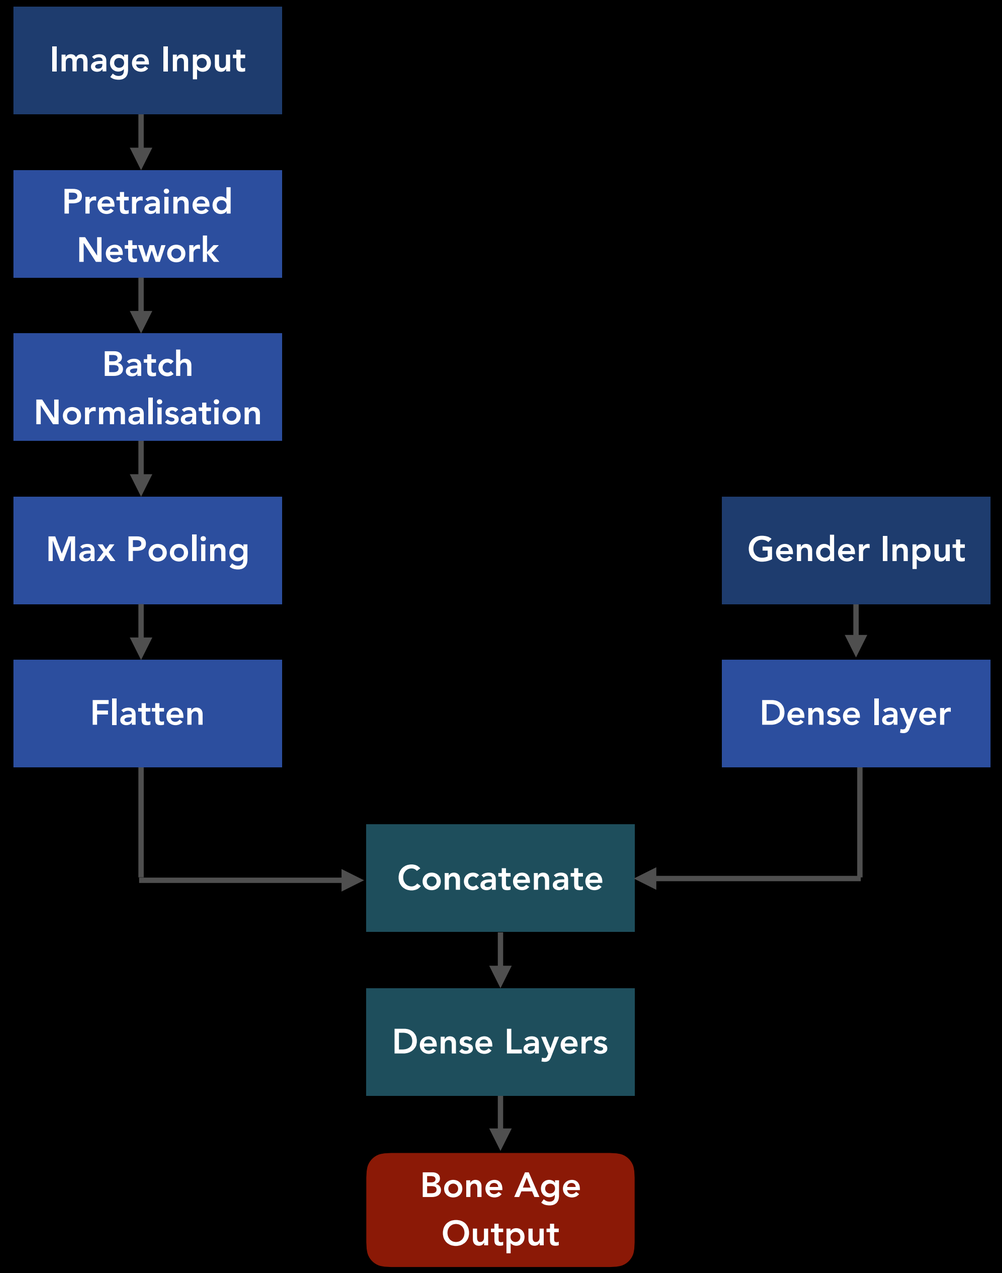

ATTENTION

- Focus on most important parts
- Inspiration from the human vision
- Light attention implementation
- Not faster, not performing better

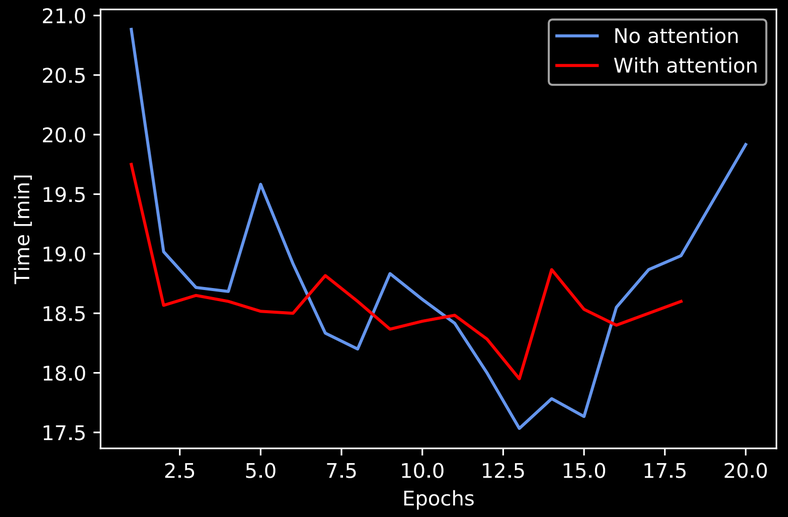

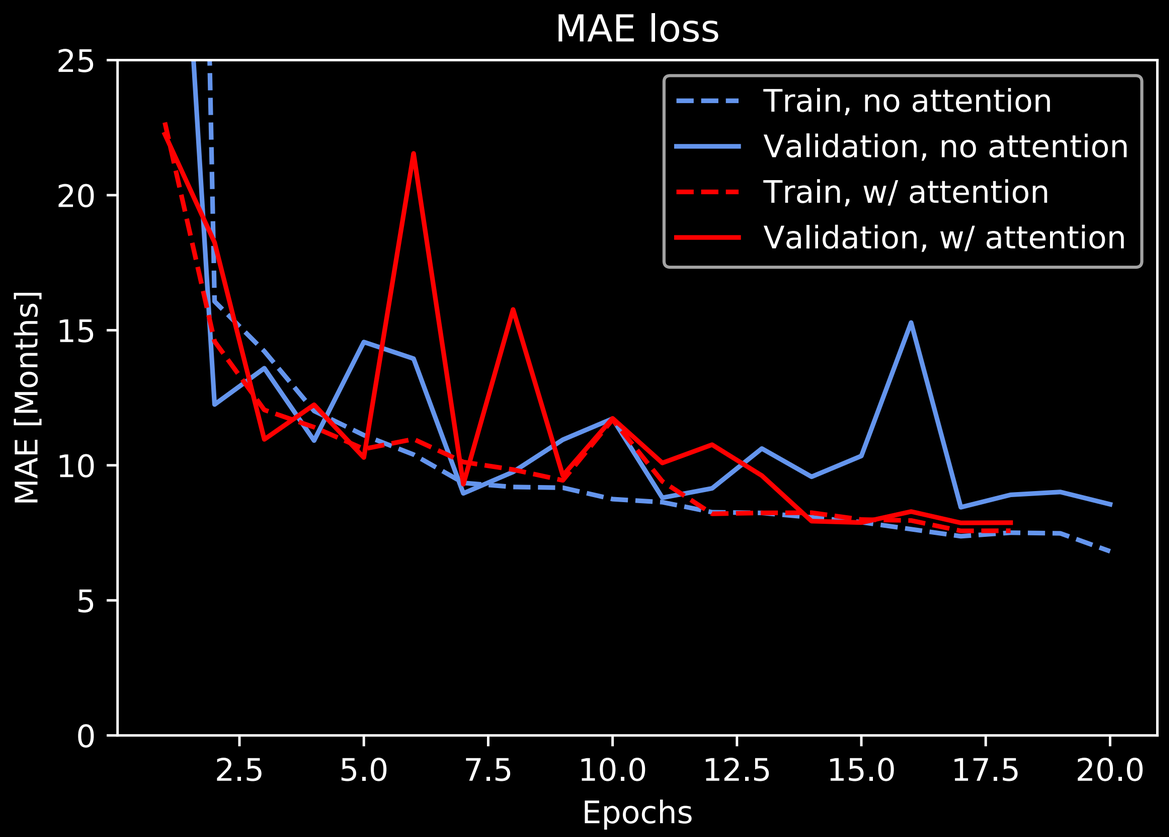

PRE-TRAINED MODELS

- Pre-trained for classification between 1000 objects in images
- Also offers preset filters without using trained weights
- Inception V3 (23.9 million parameters): Winner's pick
- VGG19 (144 million parameters): Largest is not always the best
- Xception (22.9 million parameters): our best result

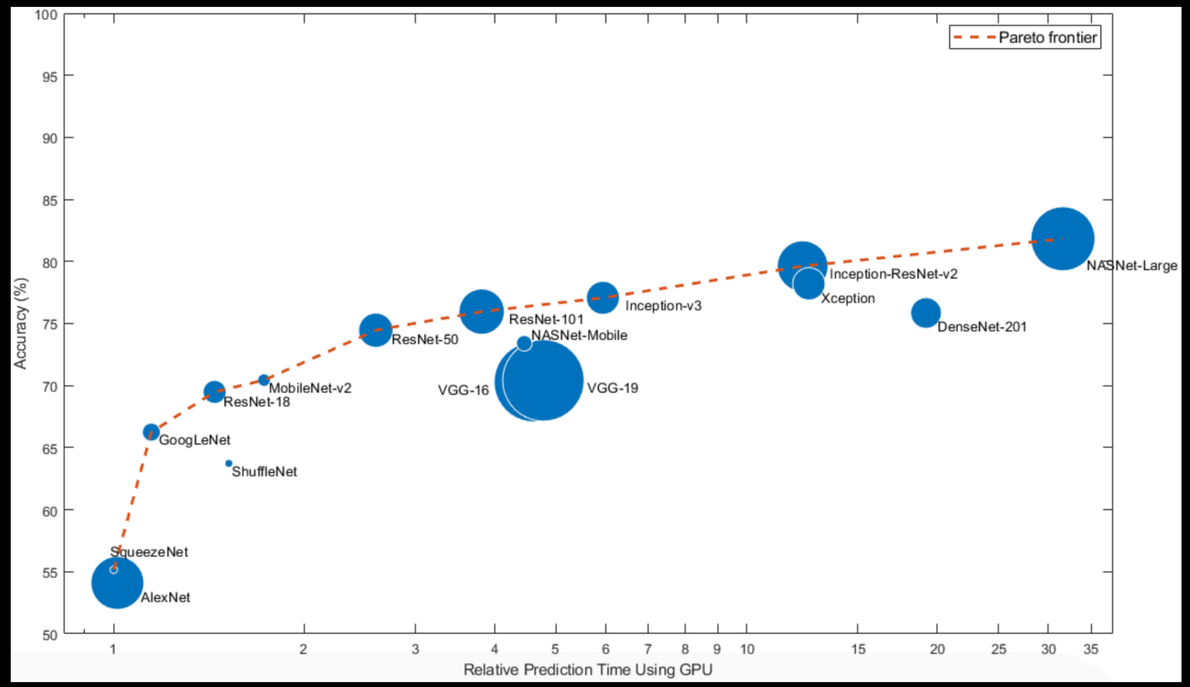

https://www.mathworks.com/help/deeplearning/ug/pretrained-convolutional-neural-networks.html

RESULTS

- Best result: MAE ~7 months
  - No attention
  - Gender as secondary input
  - Xception
  - Two dense layers

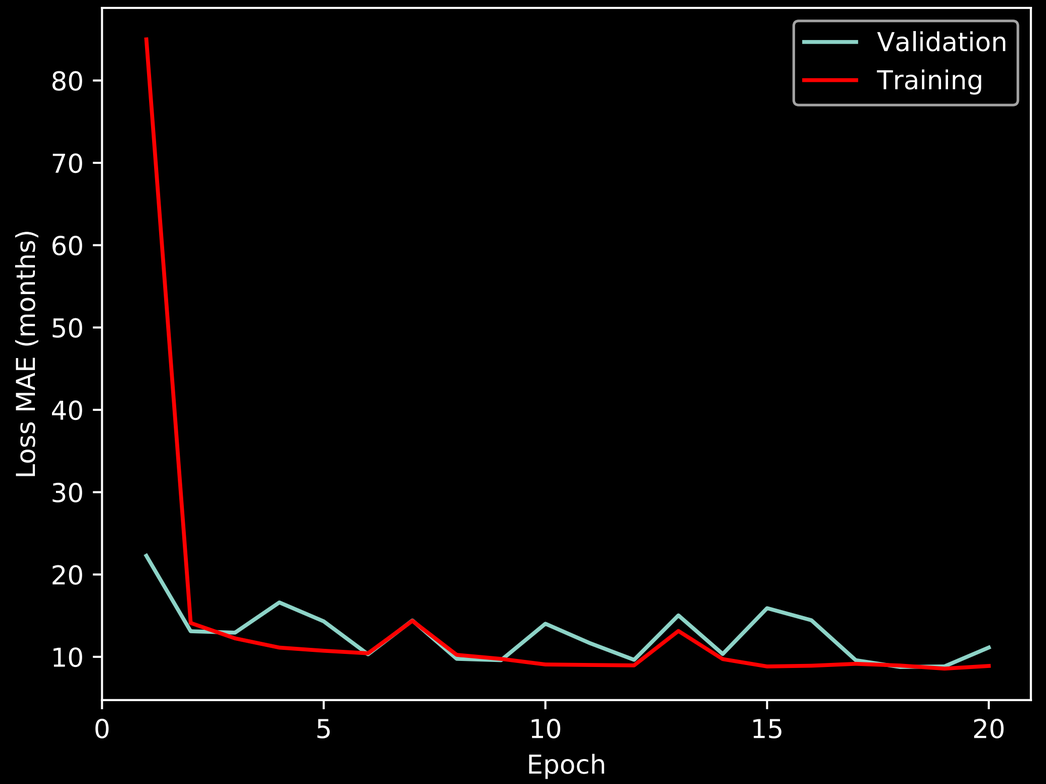

REFERENCE

1. Halabi SS, Prevedello LM, Kalpathy-Cramer J et al. The RSNA pediatric bone age machine learning challenge. **Radiology** 2019;290(2):498–503. Link, Google Scholar

2. Our Kaggle Kernel: https://www.kaggle.com/ehrhorn2019/ku-bda-2019-boneage-project

3. Code from https://www.kaggle.com/kmader/attention-on-pretrained-vgg16-for-bone-age

4. Code from https://www.kaggle.com/sinkie/keras-data-augmentation-with-multiple-inputs