RSNA 2017 Pediatric Bone Age Challenge

Alexandr Kalinin, University of Michigan
December 27, 2017

About myself and our ODS team

@alxndrkalinin, Alexandr Kalinin, University of Michigan

- BSc, MSc in Applied Math and Informatics from NSTU, 2004-2010
- Fulbright Scholar in Statistics at University of California, Los Angeles, 2013
- 5th year PhD Candidate in Bioinformatics at University of Michigan

- @ternaus, Vladimir Iglovikov, Lyft Inc.
- @rakhlin, Alexander Rakhlin, National Research University of Electronic Technology
- @shvetsiya, Alexey Shvets, MIT

Goals of the project

Pediatric bone age assessment – is a common clinical practice to diagnose endocrine and metabolic disorders in child development.

Goal: Develop an algorithm which can most accurately determine skeletal age on a validation set of pediatric hand radiographs.

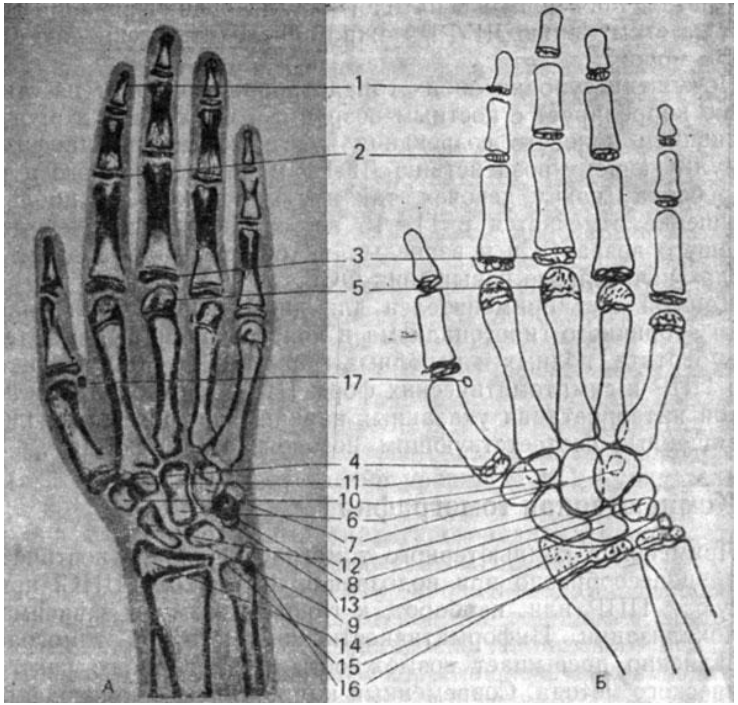

Challenge organization and logistics

Radiological Society of North America (RSNA) – international society of radiologists, medical physicists and other medical professionals with more than 54,000 members from 136 countries. Custom competition platform.

- Training phase start: Aug. 5, 2017, midnight
- Leaderboard phase start: Sept. 1, 2017, midnight
- Test phase start: Oct. 7, 2017, midnight
- Competition end: ~~Oct. 15~~ Oct. 17, 2018, 3:55 a.m.

The Challenge winners invited to join the Challenge panel at the RSNA Annual Meeting in Chicago, Nov 26. Data made publicly available.

Data description

Data structure:

- X-ray images of a left hand
- patient sex
- [outcome] bone age in month

Number of patients:

- Training: 12.6k
- Validation set: 1425
- Test: 200

EDA: train set patient sex distribution

- Approx 50/50 sex ratio
- up to 20 years old
- most ages divisible by 6

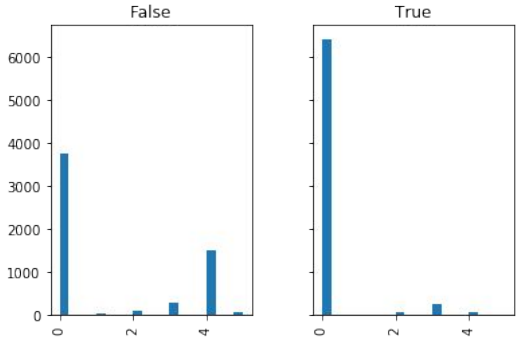

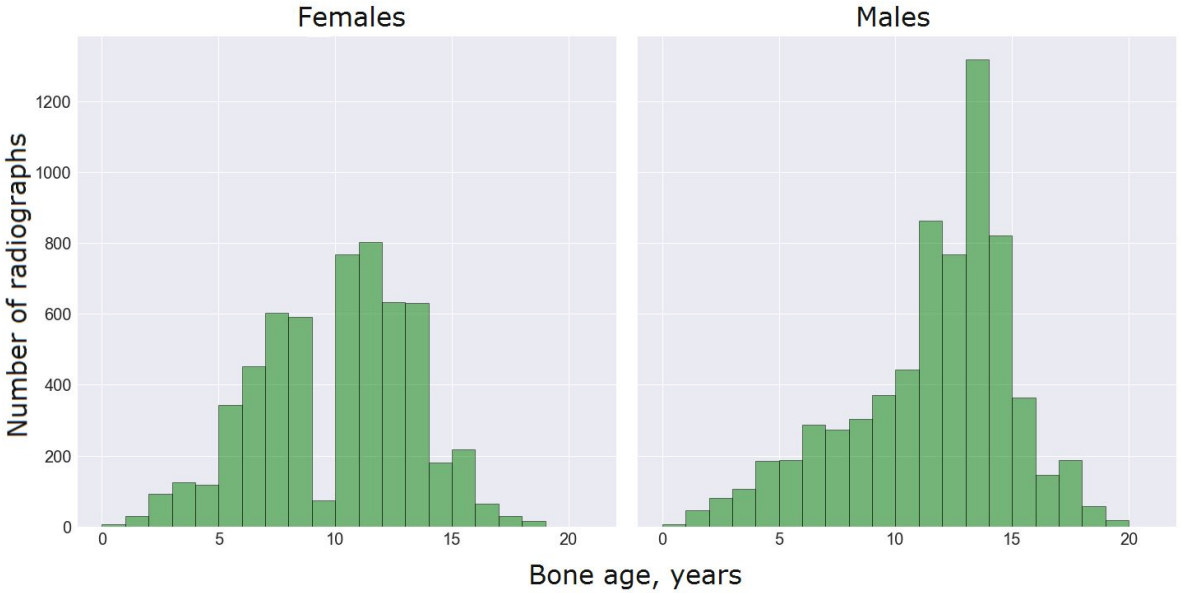

EDA: varying quality of radiographs

- annotations
- very light/dark areas
- rotated images
- contrast issues
- etc

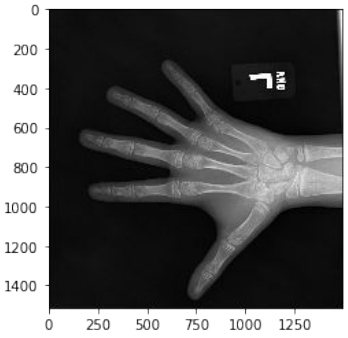

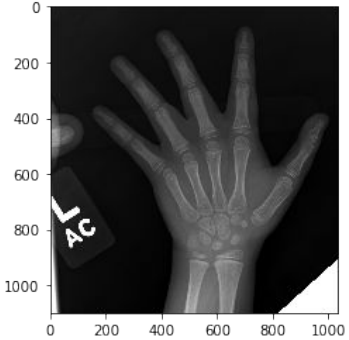

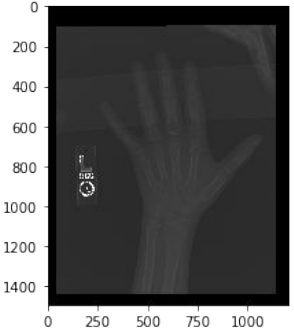

average:

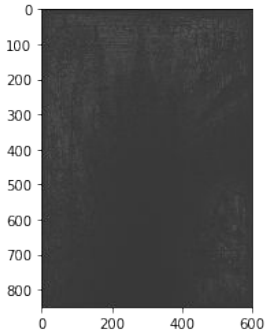

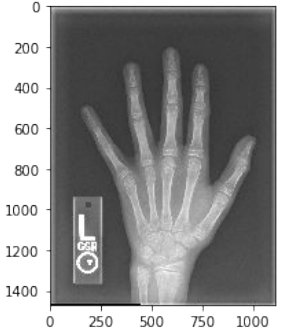

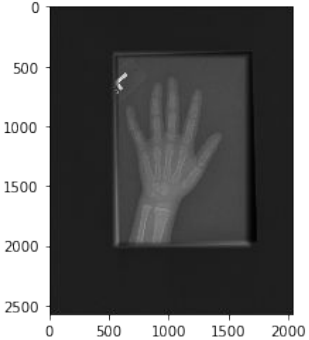

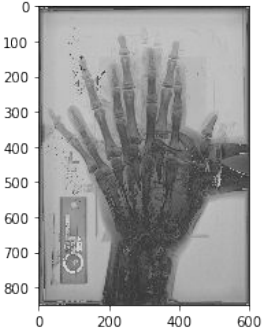

Preprocessing: segmentation by @ternaus

U-Net positive mining:

1. manually label 100 masks with Supervisely
2. train U-Net
3. predict the rest of training set
4. in every image leave 1 biggest CC
5. visually inspect all predicted masks to keep those of good quality
6. repeat steps 2-5 x5 more times
7. manually label 100 worst masks

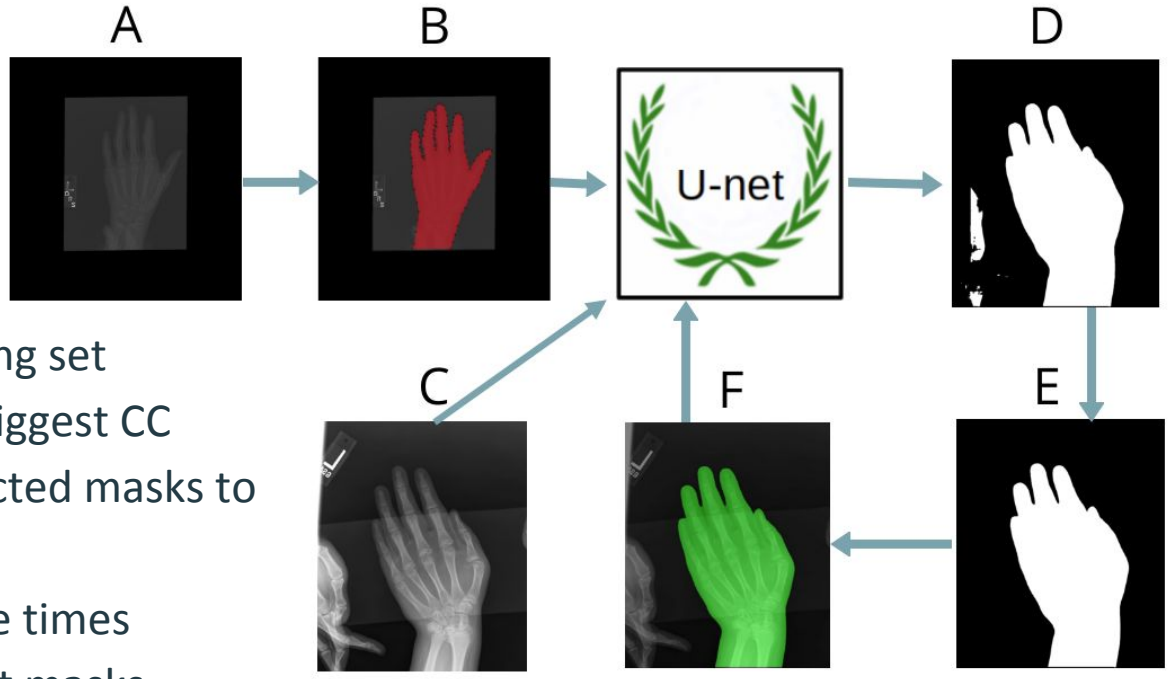

Preprocessing: registration by @rakhlin

Needed to evaluate different areas of the hand (A,B,C)

1. manually label 800 radiographs with keypoints
2. train VGG-like net on downsides images
3. predict coordinates on the rest
4. register with affine transform

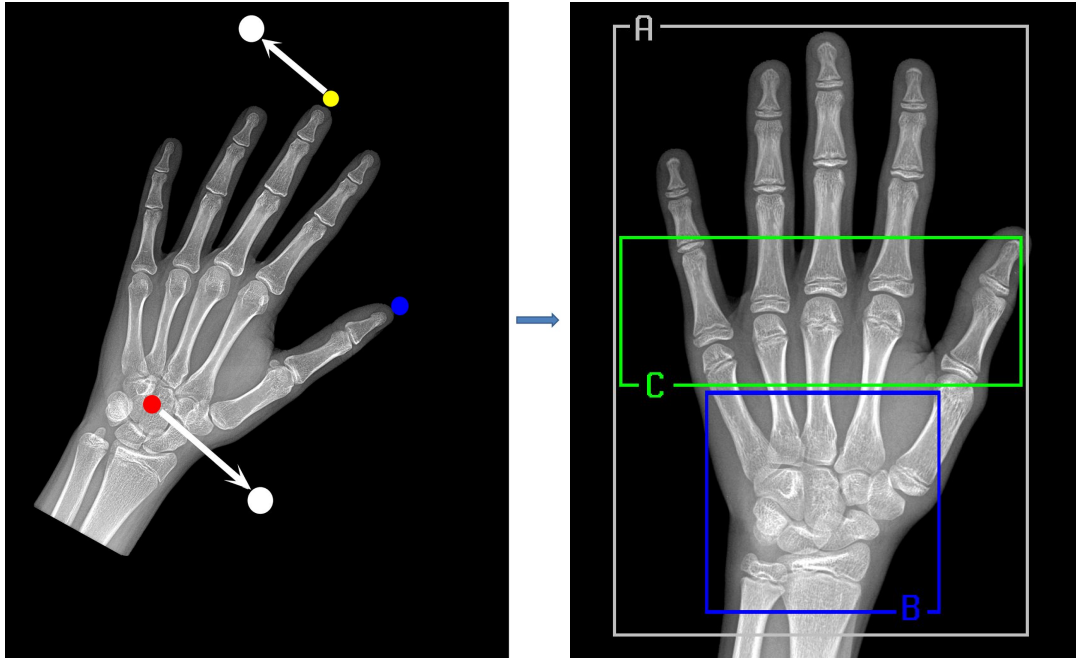

Preprocessing: registration by @rakhlin

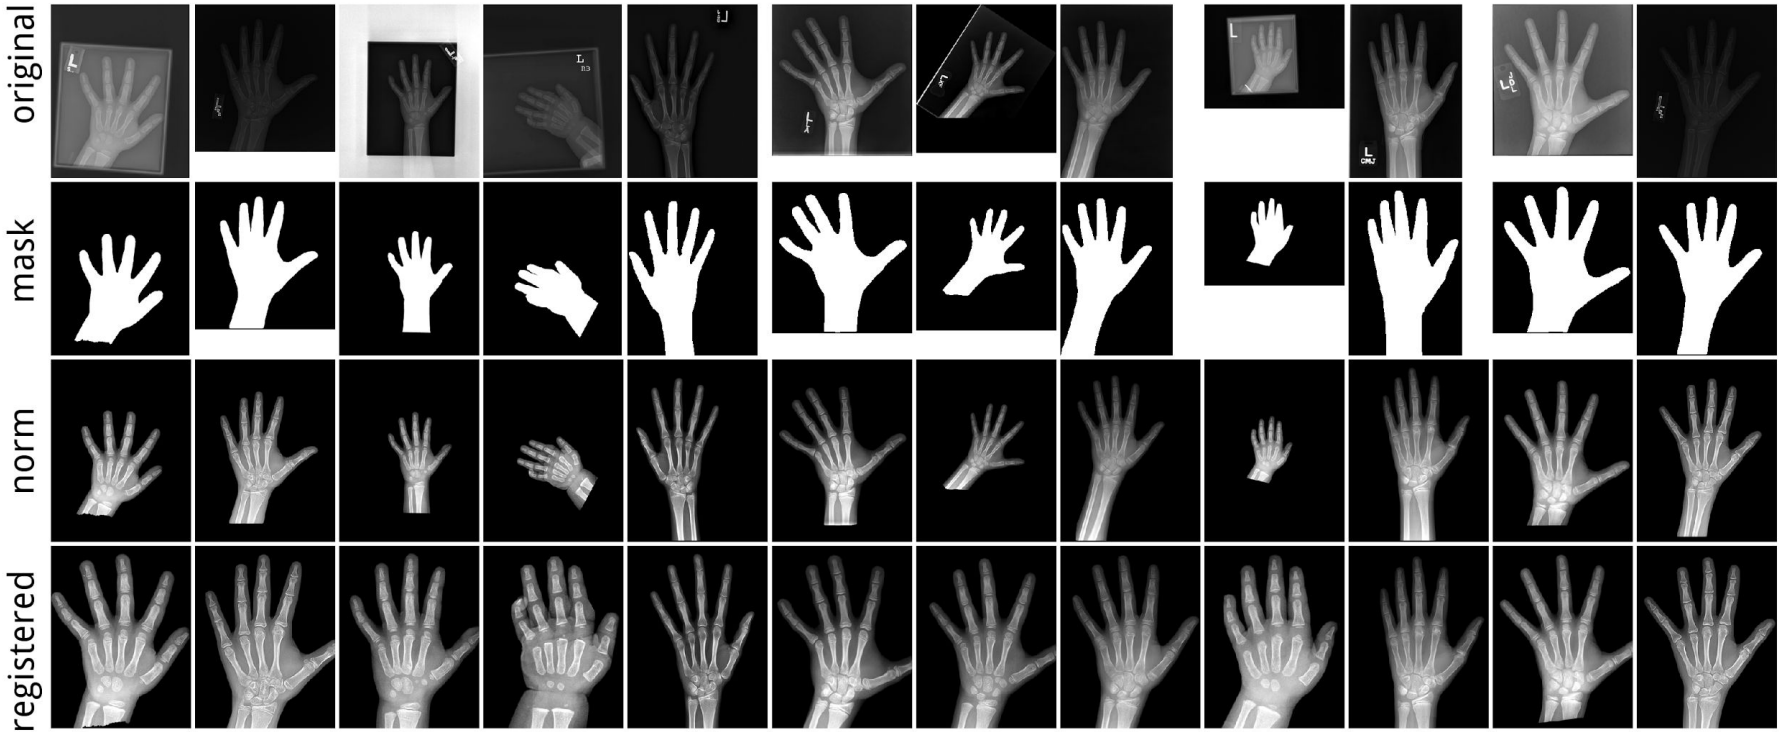

VGG-like architectures

- ELU
- Adam

Regression:

- MAE loss
- age in [−1, 1]

Classification:

- BCE loss
- 240 age bins

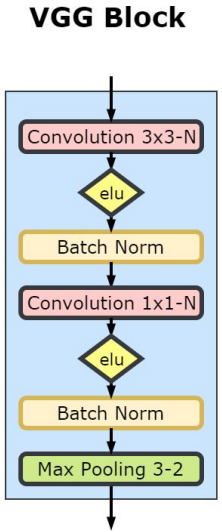

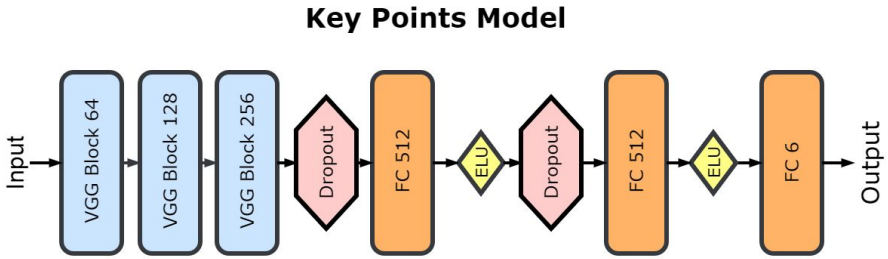

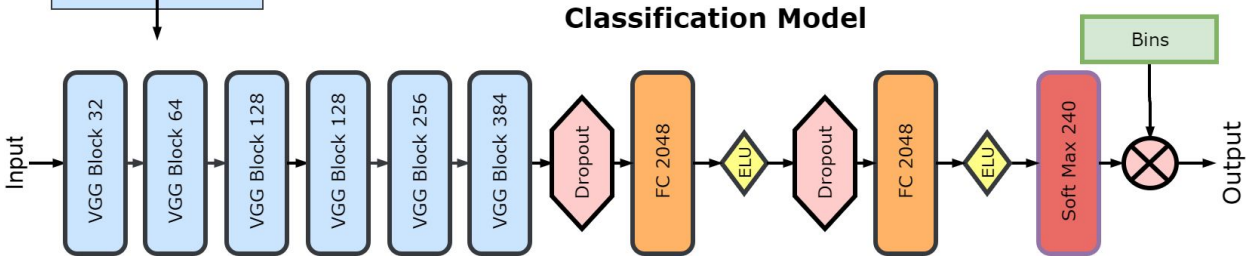

Results: different models, bones, and sex

Sex is important

MAE(B) > MAE(C) > MAE(A) > MAE (ensemble)

Classification perform slightly better than regression

Ensembling improves results

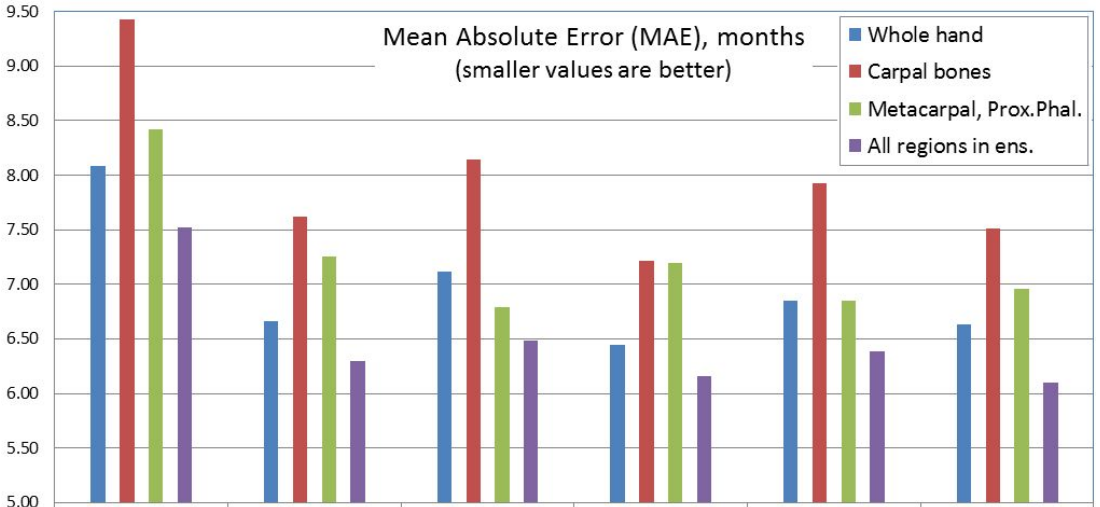

	Regression mixed	Regression males	Regression females	Classific. males	Classific. females	Regression Classific. ensemble
Whole hand	8.08	6.67	7.12	6.44	6.85	6.63
Carpal bones	9.43	7.62	8.15	7.22	7.92	7.52
Metacarpal, Prox.Phal.	8.42	7.25	6.79	7.19	6.85	6.96
All regions in ens.	7.52	6.30	6.49	6.16	6.39	6.10

Results: skeletal development stages

Unlike previous studies, we found carpal bones not the most important compared to others

More data per age – higher accuracy

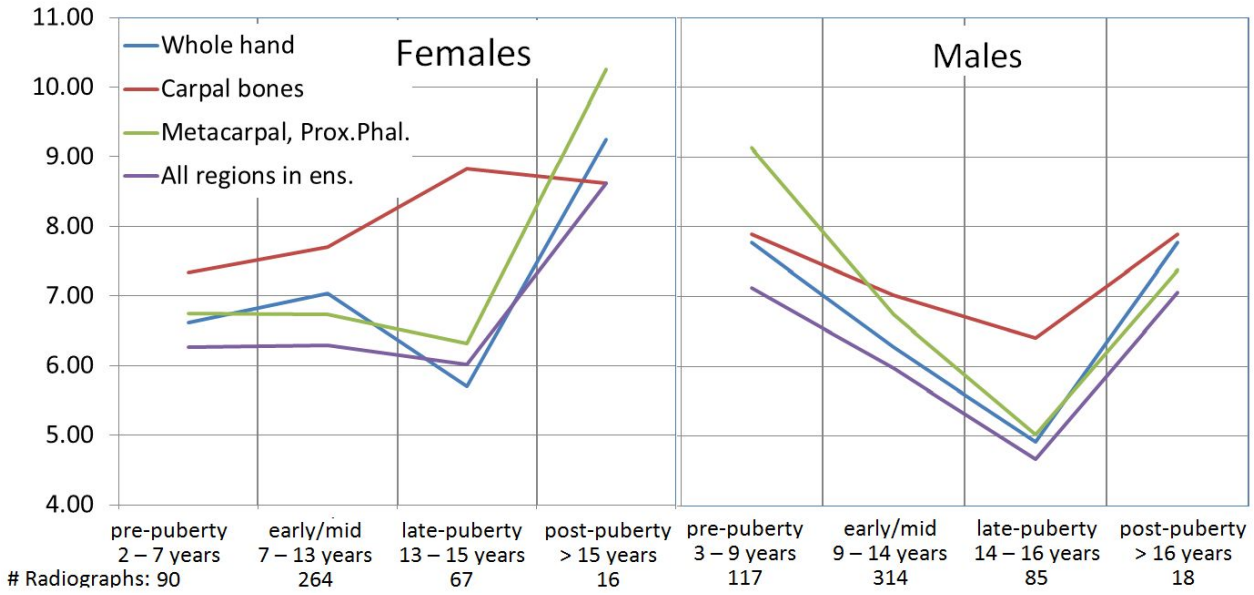

Results: winning solution

- based on Inception v3
- images resized to 500x500
- no preprocessing
- 5 networks in final ensemble
- TTA-10

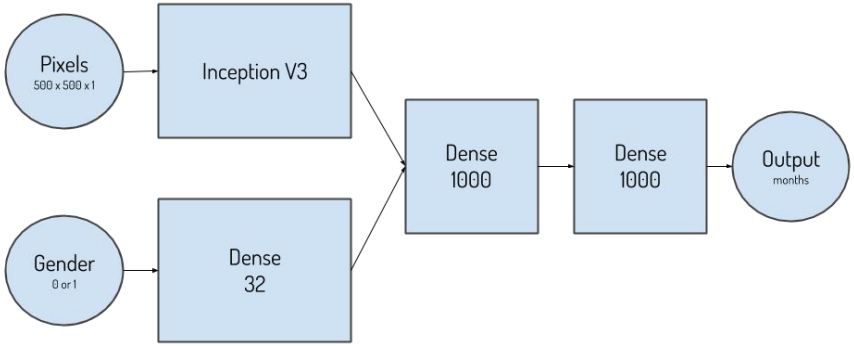

Conclusions

- we joined late & didn't win the competition

BUT

- we learned a lot
- we produced potentially useful solution
- we published a preprint that was downloaded >300 times in 12 days

Pediatric Bone Age Assessment Using Deep Convolutional Neural Networks

Vladimir Iglovikov, Alexander Rakhlin, Alexandr Kalinin, Alexey Shvets

doi: https://doi.org/10.1101/234120

This article is a preprint and has not been peer-reviewed [what does this mean?].

Abstract | Info/History | **Metrics** | Preview PDF

ARTICLE USAGE

Article lifetime | Last 6 months | This month

Article usage: December 2017 to December 2017

Show by month	**Abstract**	**Pdf**
Total	1,173	315

Framework

- dynamic, imperative, fast
- aims to be a drop-in replacement for NumPy with GPU support and autodiff
- provides convenient DataLoader with out-of-the-box multiprocessing
- provides many DL abstractions
- supports multi-GPU out-of-the-box
- has decent tutorials, forum, and not-too-steep learning curve