Detection on Kaggle

@tiulpin

2019

Detection?

No,

- RSNA Intracranial Hemorrhage Detection — multi-label classification task
- Severstal: Steel Defect Detection — instance segmentation task

RSNA Intracranial Hemorrhage Detection

Competition rules

- In this competition, the challenge is to build an algorithm to detect acute intracranial hemorrhage and its subtypes.
- The training data is provided as a set of image Ids and multiple labels (epidural, intraparenchymal, intraventricular, subarachnoid, subdural), one for each of five sub-types of hemorrhage + for any).
- This is a two-stage competition.

Evaluation metric

Submissions are evaluated using **a weighted multi-label logarithmic loss.**

Each hemorrhage sub-type is its own row for every image, and you are expected to predict a probability for that sub-type of hemorrhage.

There is also an any label, which indicates that a hemorrhage of ANY kind exists in the image. The any label is weighted **more highly** than specific hemorrhage sub-types.

Data: images, DICOM

- With hemorrhage: 674258
    1. Any: 97103
    2. Epidural: 2761
    3. Intraparenchymal: 32564
    4. Intraventricular: 23766
    5. Subarachnoid: 32122
    6. Subdural: 42496

Before / after windowing

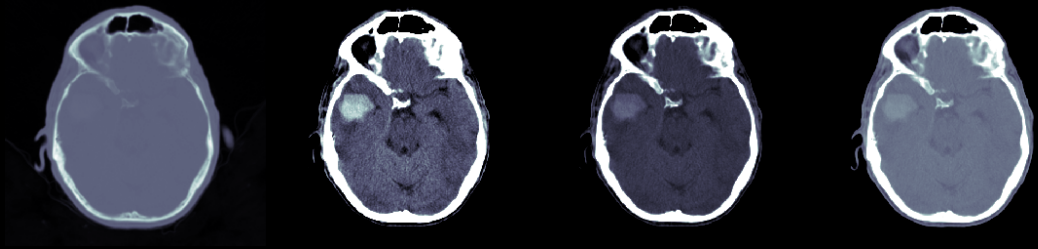

Figure 1: before window | brain window | blood/subdural window | soft tissues window

This is a two-stage competition

In order to be eligible for Stage 2, each team's Stage 1 submission must include the model uploaded.

This model should match that which will be used to generate the 2 final submissions you will select for scoring in Stage 2.

The deadline for model upload is firmly the end of Stage 1.

Our Team

1push [ods.ai] — 48/1345, Top 4%

1. Viktor Tiulpin, @tiulpin
2. Aindriú, @andtem2000
3. Aleksei Tiulpin, @alekseit

Private LB score: 0.05461

48th place solution

Thanks to @Appian.

- Data: 5 folds, 512x512 images, 3 types of windows: brain, blood and soft tissues.
- Model: ImageNet pretrained se-resnext50-32x4d, 4 epochs
- Optimizer: Adam, loss: BCEWithLogitsLoss
- Augmentations: HorizontalFlip, RandomResizedCrop, RandomBrightnessContrast, Rotate, Cutout
- TTA (5 times)
- 5-folds-averaging for submission

Competition pain

- @andtem2000 magic stacking didn't work: validation score was increasing, public score was decreasing. Why? Folding was not based on time series!
- The deadline was extended (didn't help at all)

2nd place solution — TL;DR

Apply windowing + scipy.ndimage.minimum_filter

Classify images — resnext101 — extract pre logit layer (GAP layer) at inference time (3 folds, 5 epochs)

Create Sequences — extract metadata from DCM: Sequence images on Patient, Study and Series

LSTM — feed in the embeddings in sequence on above key - Patient, Study and Series (3 * 5 LSTMs), then average predictions

Q&A

Our workers are spending a little extra time on your submission. Grab a coffee and refresh later to see your score...

Sevestal: Steel Defect Detection

Competition rules

- In this competition we were predicting the location and type of defects found in steel manufacturing.
- Each image may have no defects, a defect of a single class, or defects of multiple classes. For each image you must segment defects of each class (ClassId = [1, 2, 3, 4]).
- This is a Kernels-only competition.

Evaluation metric

This competition is evaluated on the mean Dice coefficient.

$$2 * \frac{|X \cap Y|}{|X| + |Y|}$$

where X is the predicted set of pixels and Y is the ground truth.

The Dice coefficient **is defined to be 1** when both X and Y are empty.

The leaderboard score is the mean of the Dice coefficients for each image-class pair in the test set.

Data: images, JPEG

- **No defects**: 5902 images
- **With defects**: 6666 images, by classes:
  1. 897 images,
  2. 247 images,
  3. 5150 images,
  4. 801 images.


THE
SHAKEUP
IS
NEAR

Training example

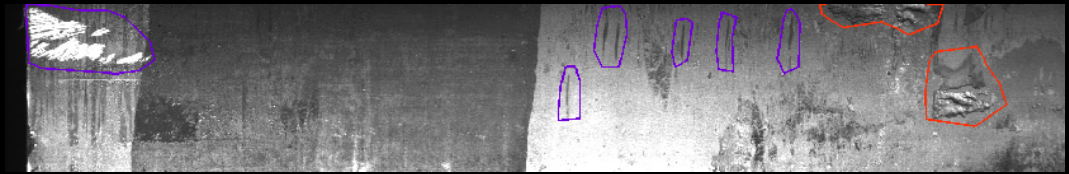

Figure 2: Image with multiple defects (3, 4)

This is a Kernels-only competition

Kernel requirements:

- CPU or GPU Kernels $\leq$ 1 hour run-time
- Internet must be turned off

A submission that results in an error will count against your daily submission limit.

Our Team

1push [ods.ai] — 30/2433, Top 2%

1. Aindriú, @andtem2000
2. ZFTurbo, @zfturbo
3. Andrew L, @jiy3ahko
4. Mikhail Liz, @lizmisha
5. Viktor Tiulpin, @tiulpin

Private LB score: 0.90431, BUT we had 0.90611 (gold)...


ONE DOESNT TOP THE LEADERBOARD
WITHOUT ENSEMBLING

Gold-we-didn't-get solution (< 1 hour inference time) TL;DR

Classification:

- by @ZFTurbo: DenseNet121 — pseudo-labelling
- public: ResNet34 by @hengck23

Segmentation (U-Net):

- by @andtem2000: with ResNet34
- public: with MobileNet2, ResNet34, SEResNeXt50 by @lightforever

Both:

- by @lizmisha: 2x U-Net (EfficientNetB5 (additional head for classification) for class 1, 2
- by @aluzan: FPN ResNet34 ensemble (additional head for classification) — pseudo-labelling, heavy augmentations

Things that will follow you in your nightmares...

```
1  from numba import cuda
2
3  cuda.select_device(0)
4  cuda.close()
```

That’s how you clean your GPU memory after working with Keras to use PyTorch.

Q&A

Viktor Tiulpin

mail: viktor@tiulp.in

socials: **@tiulpin**