

Weakly Supervised Learning for Findings Detection in Medical Images

HaoCheng Kao


Outline

- The Task
- A Naïve method
- The Dataset
- Grading
- Submission
- Rules
- Contacts

Task



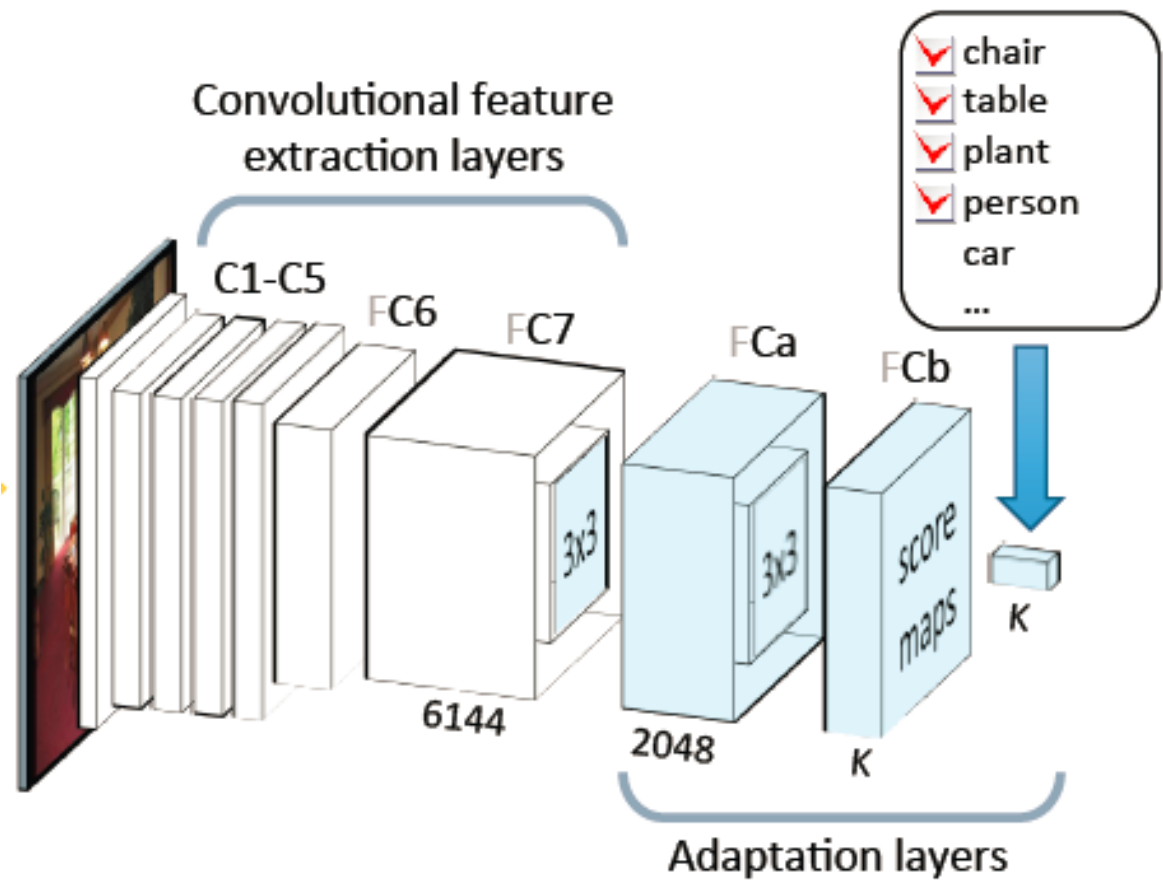
Task



Task

- 14 Findings (classes)
- Predict bounding box + class
- Training data has only class labels

Naïve method



Dataset

- We use NIH ChestX-ray 14 dataset
<https://nihcc.app.box.com/v/ChestXray-NIHCC>
- Training set (111240)
- Validation set (440)
- Testing set (440)


Download

- `Data_Entry_2017_v2.csv`
Image-level annotation for all images.
- `Bbox_List_2017.csv`
Bounding box annotation for validation / testing images.
- `train.txt / valid.txt / test.txt`
Lists of images in each set.

Grading

- 60% accuracy
- 40% novelty & report or poster

Intersection over Union



Accuracy

- For each ground truth box in a test image, we'll check if there is an output box correctly locate this box.
 - Class label is correctly predicted.
 - IoU of ground truth box and predicted box $\geq T(\text{IoU})$.
- The ratio of successful ground truth boxes in an image is the score of the image.
- The average score of all images is the score of a specific $T(\text{IoU})$.
- The final score is the average of $T(\text{IoU})=0.25$ and $T(\text{IoU})=0.5$.
- You must output up to 10 bounding boxes per image.

Typical workflow of a submission

1. Call `get_file_names()` to get the list of files.
2. Load your model.
3. For each image, inference the answer.
4. Call `get_output_file_object()` to get the output file object.
5. Write the output.
6. Call `judge()` to get the result.

Note that the timestamp created when `judge()` is called is used for judging if the submission is on time.

Submission

- Submission for evaluating on testing data
 - During this competition, you can submit your model to be evaluated on testing data.
- Final submission
 - After you complete your project, you should submit your whole project so that we can verify your result.

Final Submission

- Your final submission should contain the following:
 - Trained model & Whole Code
 - Document & report

Task Restrictions

- Keras / PyTorch / TensorFlow / MXNet / CNTK
- You can only submit a single archive containing all the codes and model. This archive must not exceed 1GB in size.

Restrictions on Dataset

- Stick to the dataset provided by link.
- The bounding boxes of validation / testing set are publicly disclosed, but you may not use them to train your model.
- You are allowed to use only the following extra context:
Patient Age / Patient Gender / View Position

Pre-trained Model

- The pre-trained model should only be trained by the ImageNet dataset for classification and localization.
- You can download any available pre-trained model from the internet
 - Make sure that it is only trained by the ImageNet dataset for classification and localization.
 - Make sure you have the rights to use the model.
- You can create your own pre-trained model, but it should be trained from randomly-initialized parameters.

Rules

- The final model must only rely on the training / validation set. The bounding boxes of the validation set can only be used to validate your model.
- For grading, you must submit your model / algorithm to the submission site before the deadline following the submission rules.
- All submitted materials must be created by the team members or the team members must have proper rights to use them.
- All novelty must be created solely by the team members. No assistance from outside of the team is allowed.
- Course regulations, the contract, and school regulations must be followed.

Contacts

- HaoCheng_Kao@htc.com
- In general, we will announce your questions and our answers to all students instead of replying to each email.