

Fast R-CNN

Object detection with Caffe

Ross Girshick

Microsoft Research

[arXiv](#) [code](#)

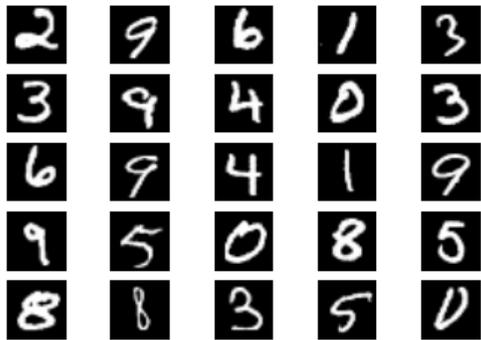
Latest roasts

Goals for this section

- Super quick **intro to object detection**
- Show **one way to tackle obj. det. with ConvNets**
- Highlight some **more sophisticated uses of Caffe**
 - Python layers
 - Multi-task training with multiple losses
 - Batch sizes that change dynamically during `Net::Forward()`
- Pointers to open source code so you can **explore, try, and understand!**

Image classification (mostly what you've seen)

- K classes
- Task: Assign the correct class label to the whole image



Digit classification (MNIST)

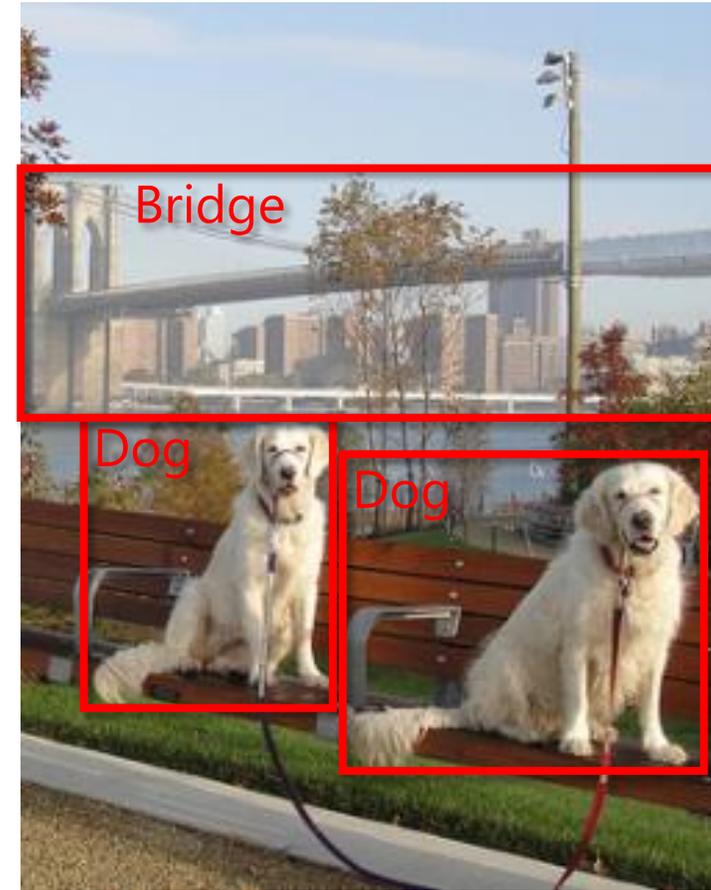


Object recognition (Caltech-101, ImageNet, etc.)

Classification vs. Detection



Easyish, these days



Still quite a lot harder

Problem formulation

The Visual World $\approx K$ object classes

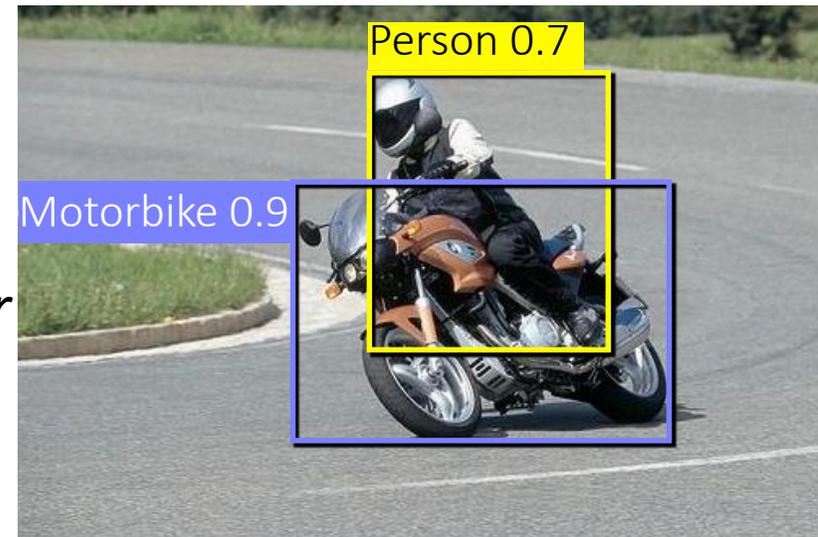
{airplane, bird, motorbike, person, sofa, bg}



Input



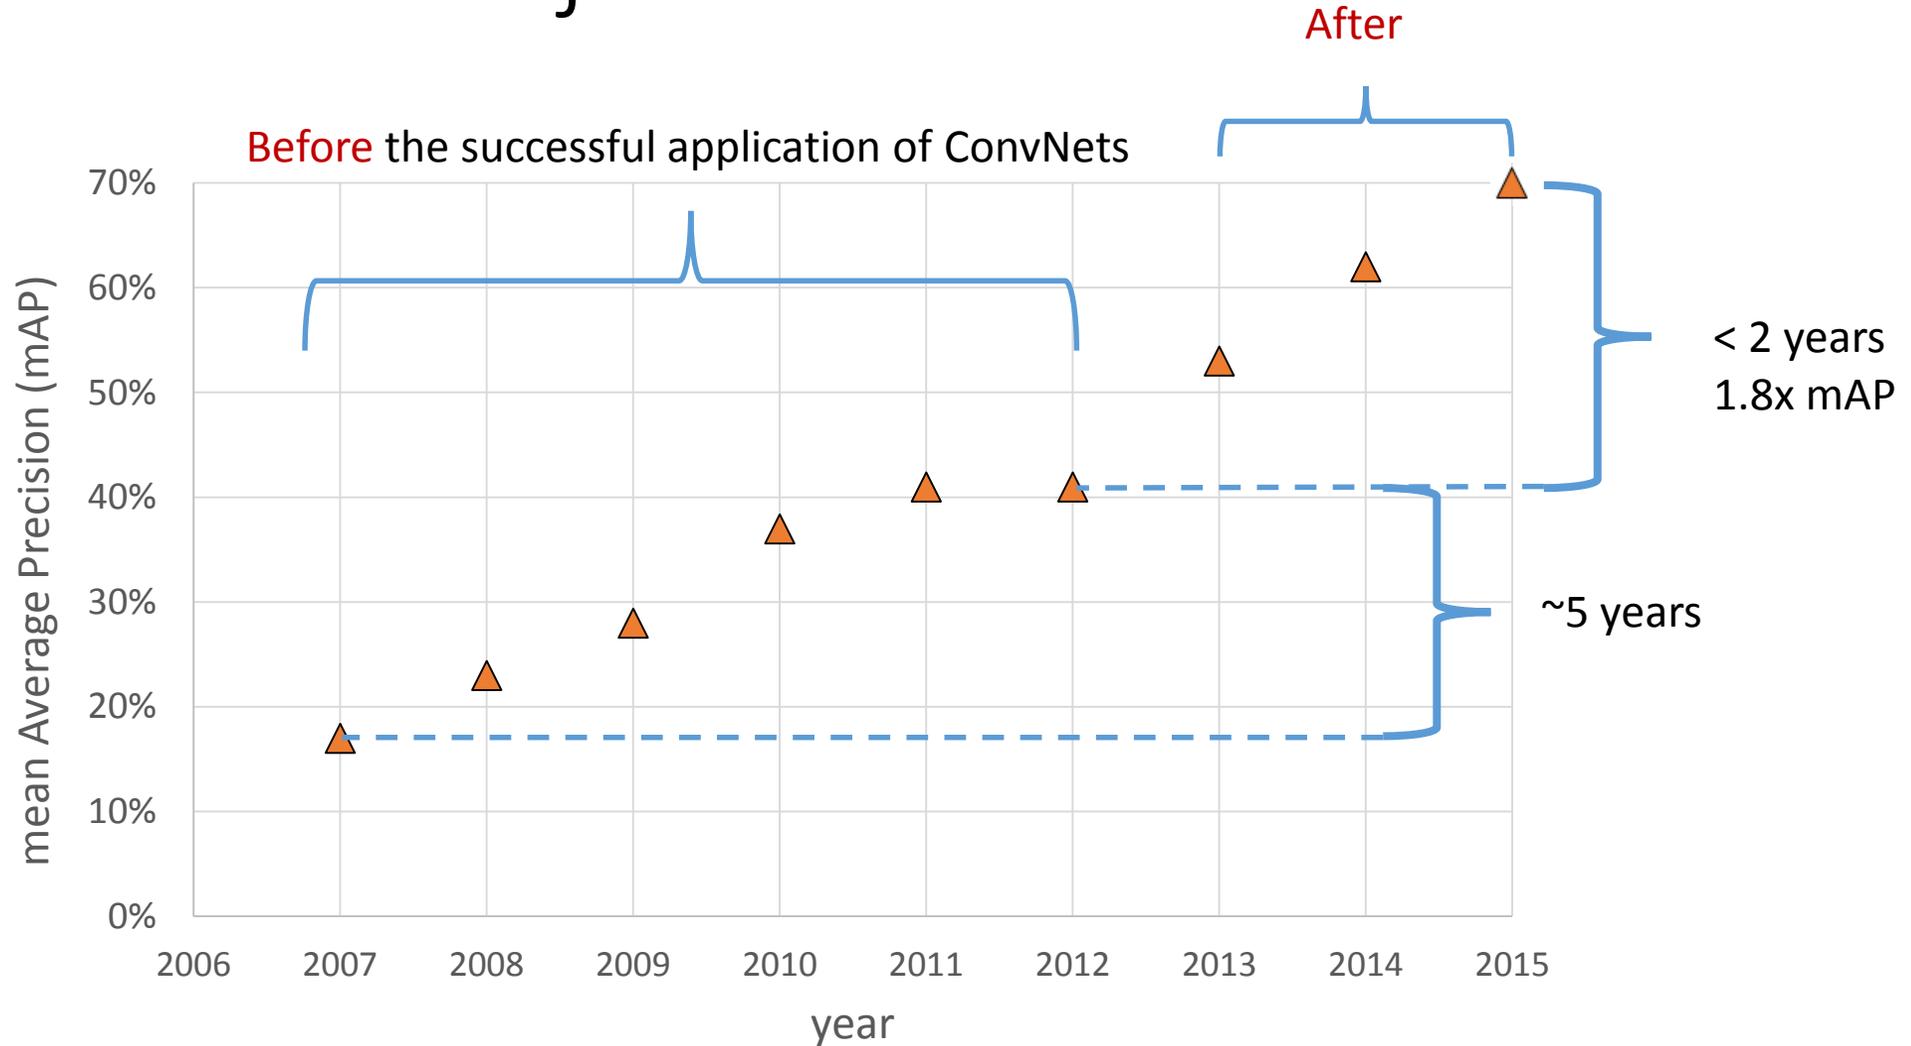
YODA:
→
*Yet another
Object
Detection
Algorithm*



Desired output

*Actual results may vary

PASCAL VOC object detection



Precision: higher is better

Fast R-CNN (Region-based Convolutional Networks)

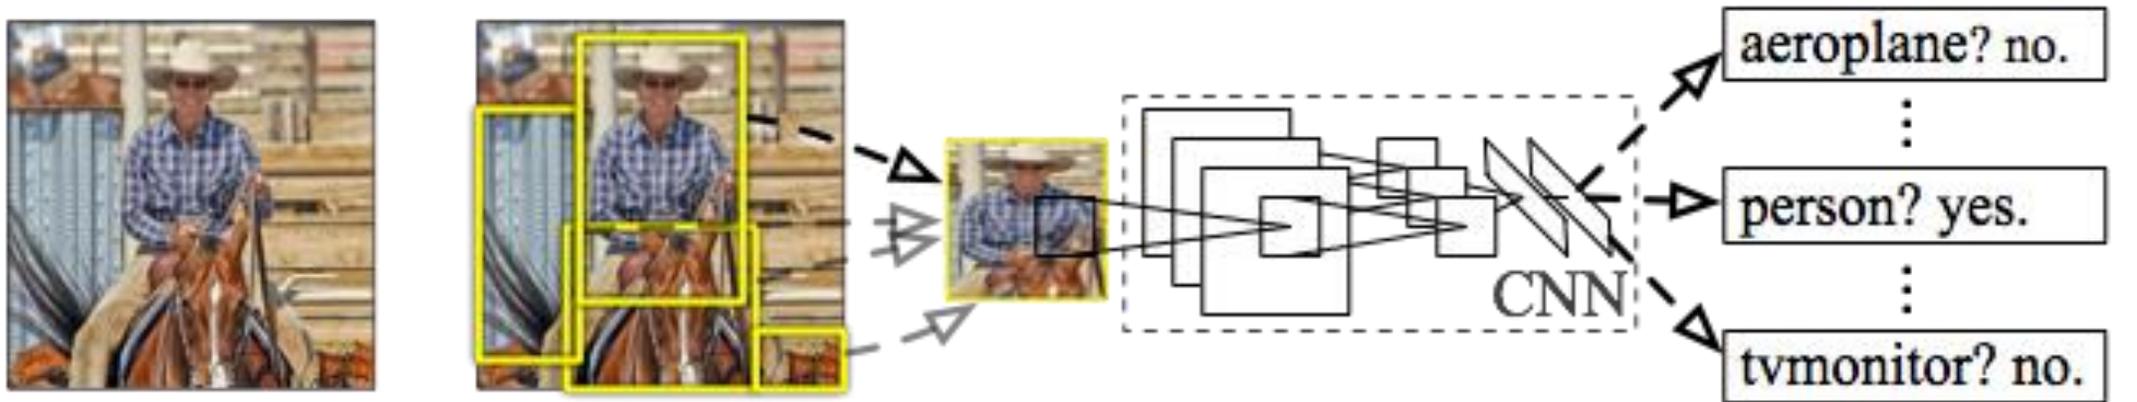
A fast object detector implemented with Caffe

- Caffe fork on GitHub that adds two new layers (ROIPoolingLayer and SmoothL1LossLayer)
- Python (using pycaffe) / more advanced Caffe usage
- A type of Region-based Convolutional Network (R-CNN)

Let's see how it works!

Quick background

Region-based Convolution Networks (R-CNNs)



Input
image

Extract region
proposals (~2k / image)
e.g., selective search
[van de Sande, Uijlings et al.]

Compute CNN
features on
regions

Classify and refine
regions

Fast R-CNN (test-time detection)



Object box proposals (N)
e.g., selective search

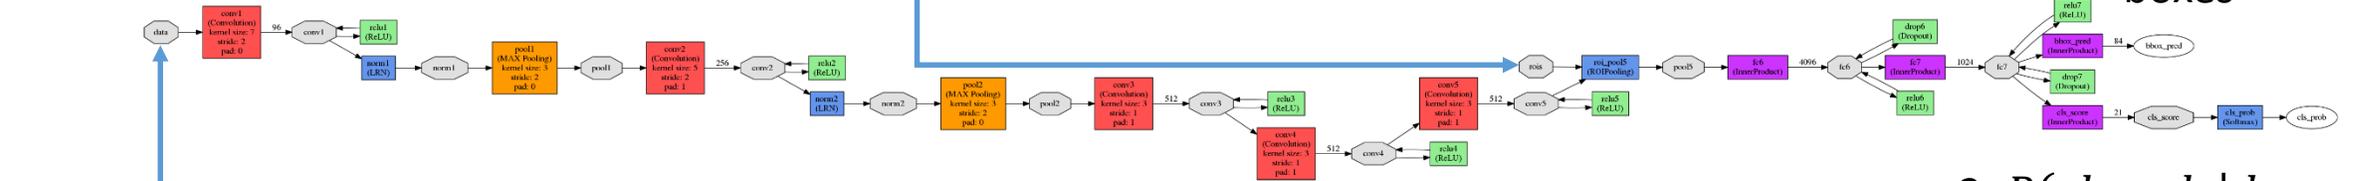
A Fast R-CNN network
(VGG_CNN_M_1024)

Two output types:

1. NK regressed object boxes

2. $P(\text{cls} = k \mid \text{box} = n, \text{image})$

for each NK boxes



Net::Forward() takes 60 to 330ms

Given an image and object proposals,
detection happens with a single call to the Net::Forward()



Image

Fast R-CNN (test-time detection)



Object box proposals (N)
e.g., selective search

A Fast R-CNN network
(VGG_CNN_M_1024)

Two output types:

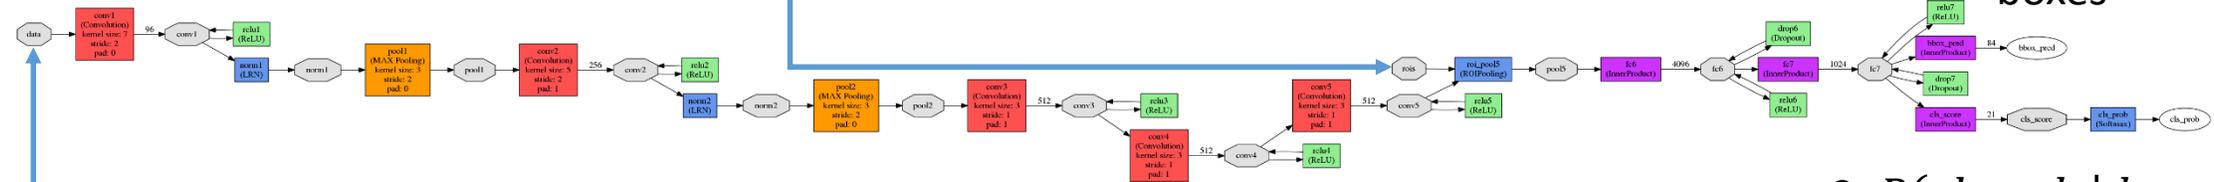
1. NK regressed object boxes

2. $P(\text{cls} = k \mid \text{box} = n, \text{image})$

for each NK boxes



Image



Minimal post-processing:

- Non-maximum suppression (NMS)

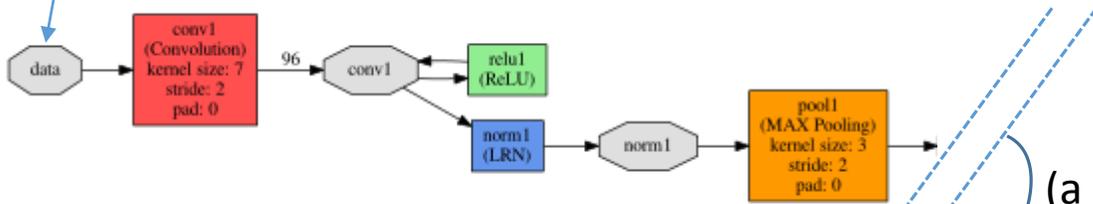
Object proposals comes from:

- Selective Search (2s / image) [van de Sande/Uijlings et al.]
- EdgeBoxes (0.2s / image) [Zitnick & Dollar]
- MCG (30s / image) [Arbelaez et al.]
- Etc.

Zooming into the net



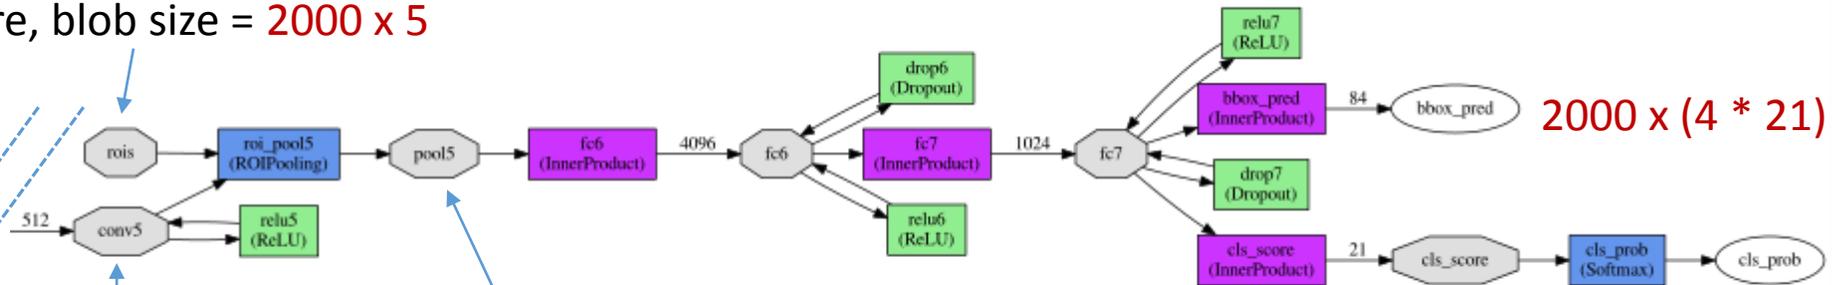
image comes in here, blob size = $S \times 3 \times H \times W$ (e.g., $S = 1$ or 5 , $H = 600$, $W = 1000$)



(a bunch of conv layers and whatnot)



2000 image regions come in here, blob size = 2000×5



$2000 \times (4 * 21)$

Pool5 blob size = $2000 \times 512 \times 6 \times 6$

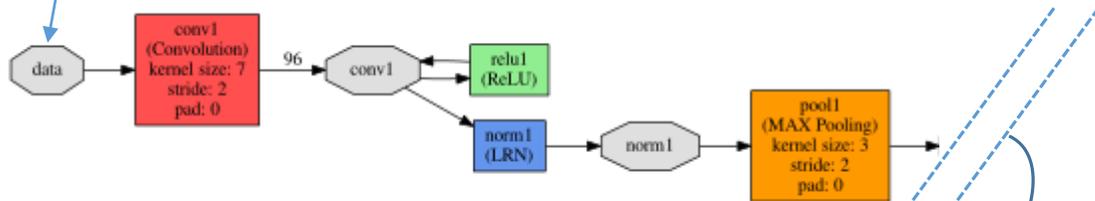
2000×21

conv5 feature map blob size = $S \times 512 \times H/16 \times W/16$

Zooming into the net



image comes in here, blob size = $S \times 3 \times H \times W$ (e.g., $S = 1$ or 5 , $H = 600$, $W = 1000$)

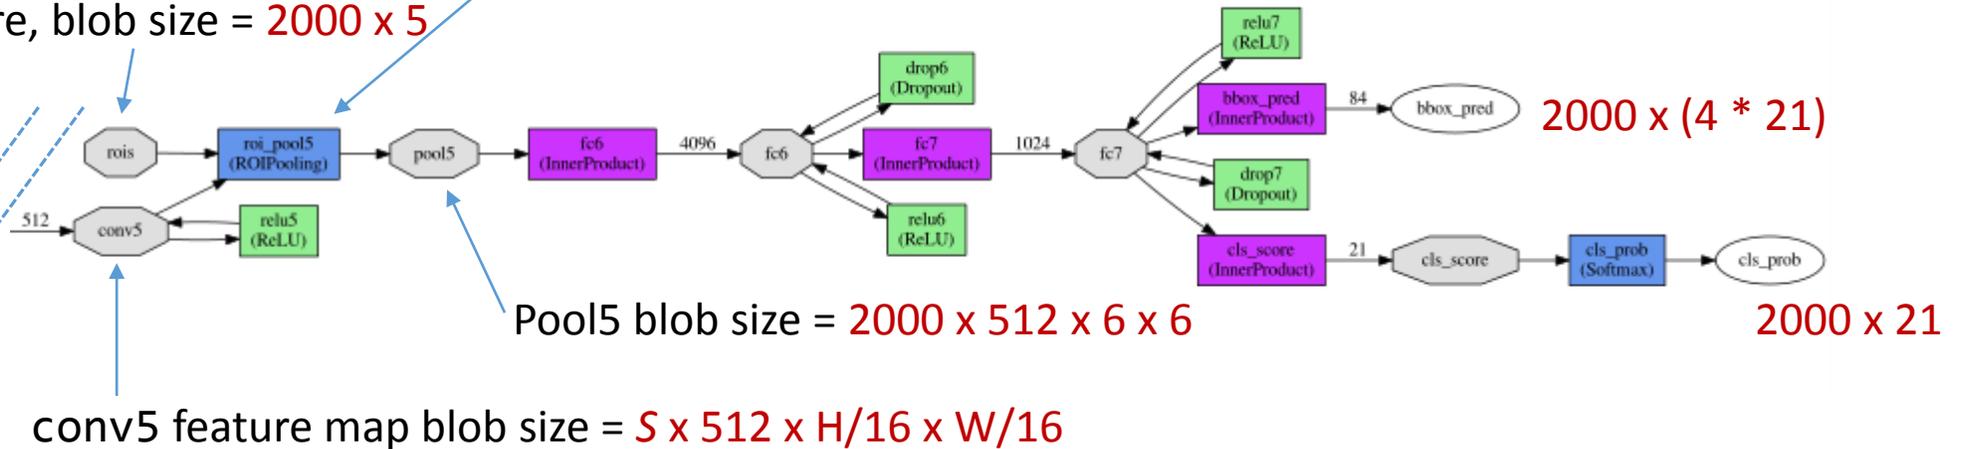


RoI Pooling Layer:

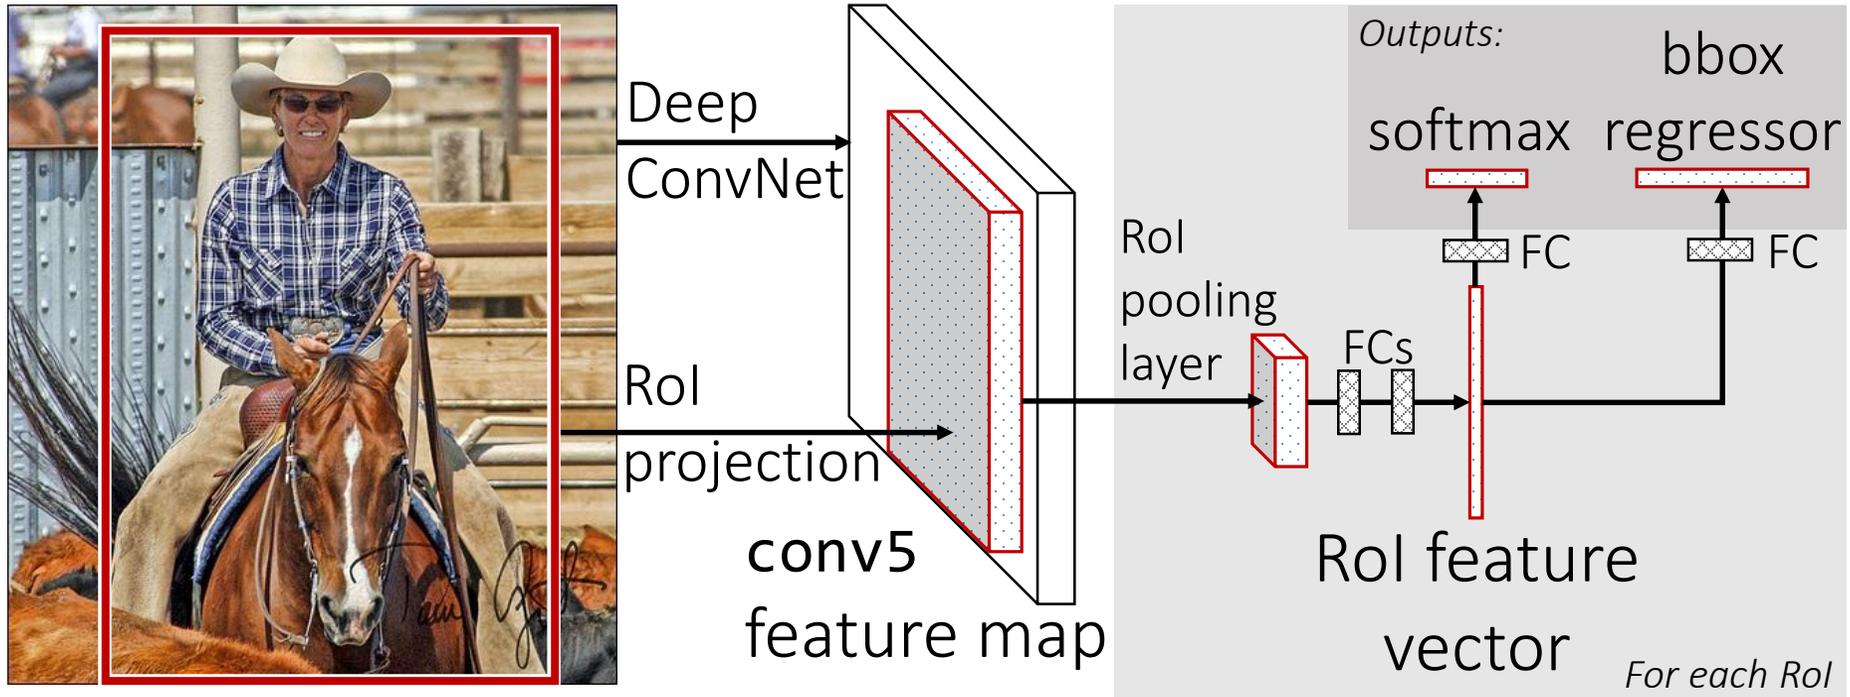
- *adaptive max pooling layer*
- *dynamically expands batch from S to R (e.g., 2000)*



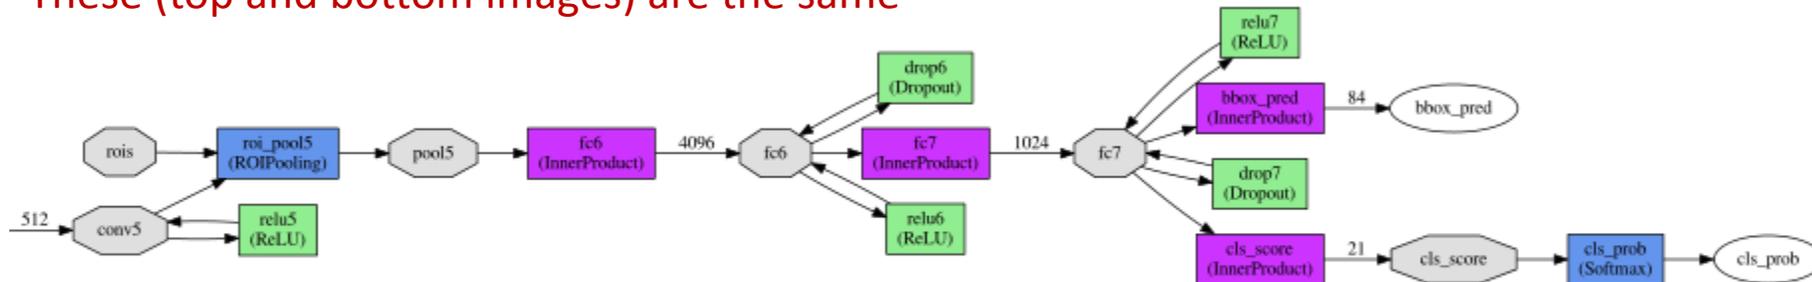
2000 image regions come in here, blob size = 2000×5



Another view of the same thing



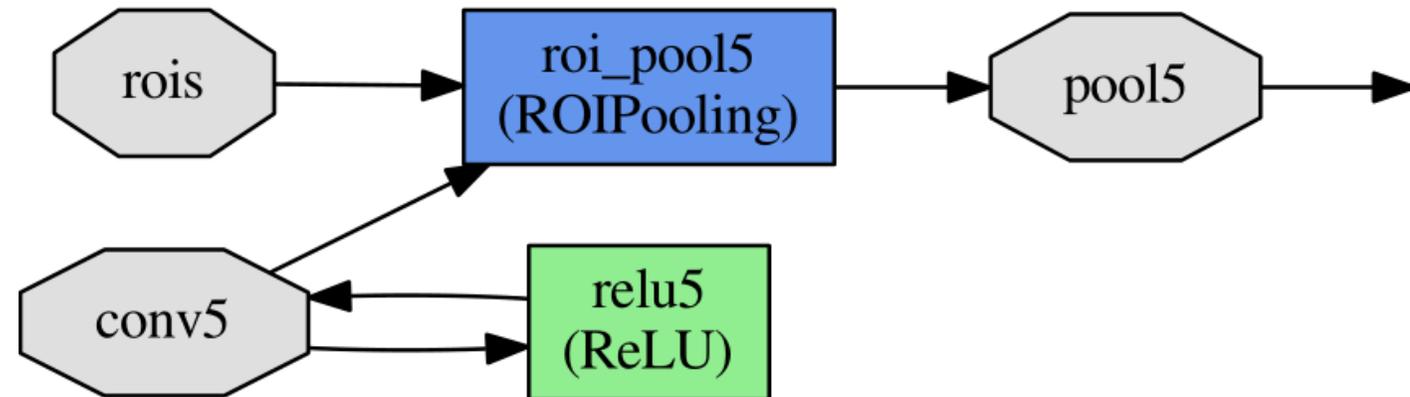
These (top and bottom images) are the same



RoI Pooling Layer

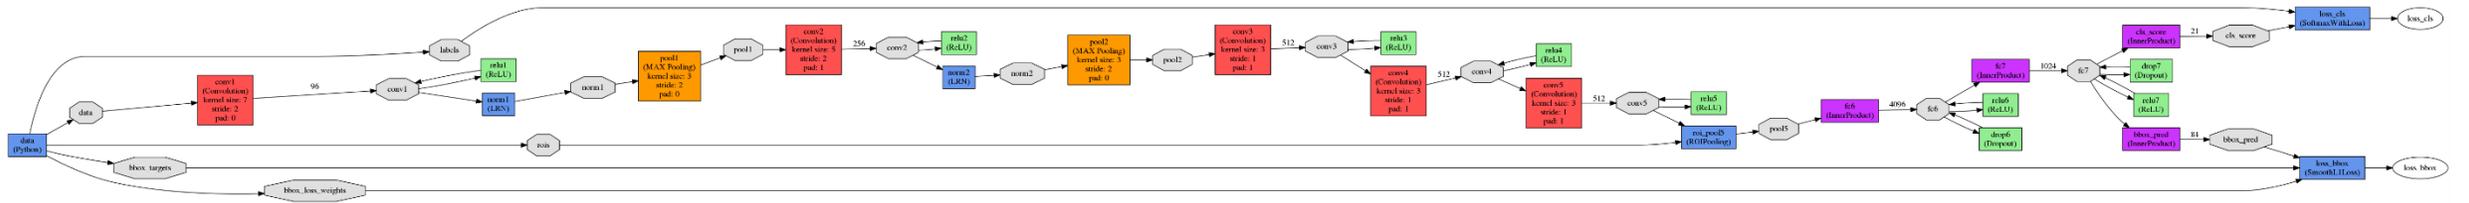
- Special case of SPPnet's SPP layer [He et al. ECCV'14]
- Two inputs ("bottoms")
 - Conv feature map: $S \times 512 \times H \times W$
 - Regions of Interest: $R \times 5$
 - 5 comes from $[r, x1, y1, x2, y2]$, where r in $[0, R - 1]$ specifies an image batch index

```
188 layer {
189   name: "roi_pool5"
190   type: "ROIPooling"
191   bottom: "conv5"
192   bottom: "rois"
193   top: "pool5"
194   roi_pooling_param {
195     pooled_w: 6
196     pooled_h: 6
197     spatial_scale: 0.0625 # 1/16
198   }
199 }
```



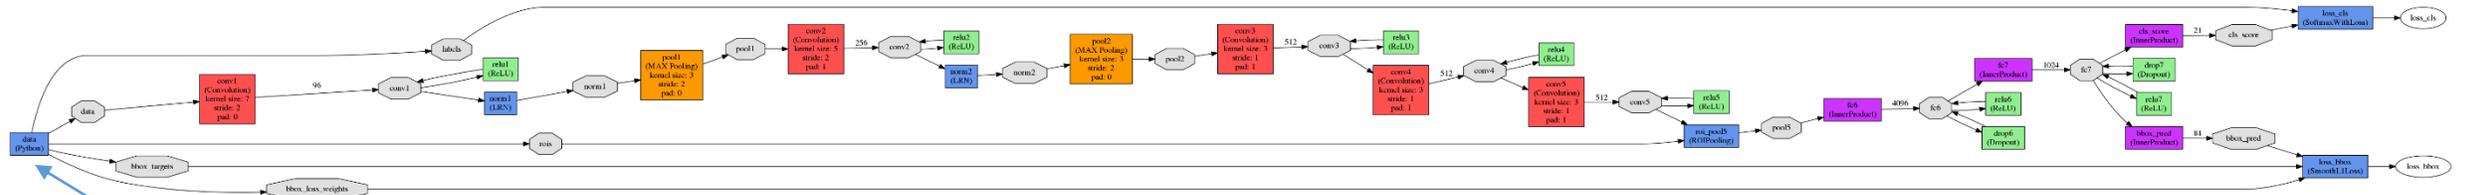
The train-time net

Single fine-tuning operation all in Caffe



Even more boxes and arrows
Let's look at them

The train-time net (exotic data layers)



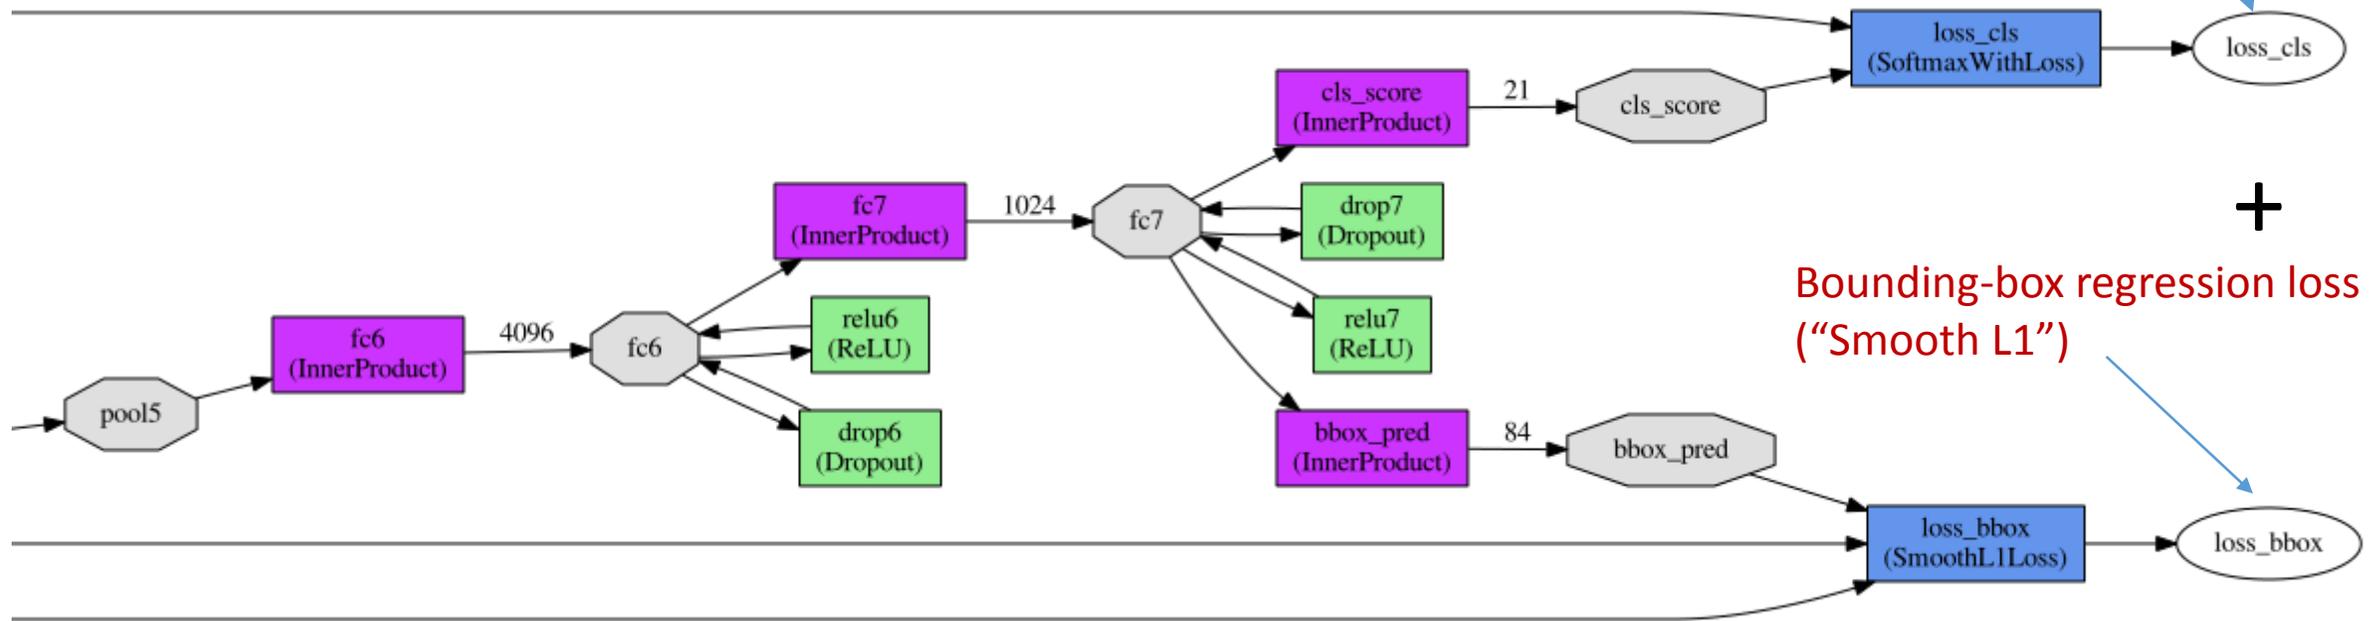
Custom Python data layer

- Samples 2 images
- From each sampled image, takes 64 RoIs
- Input batch is initially 2 elements
- Gets expanded by the RoI Pooling Layer to 128 elements
- Outputs 5 “tops”
 - data [images]
 - rois [regions of interest]
 - labels [class labels for the rois]
 - bbox_targets [box regression targets]
 - bbox_loss_weights [...details...]

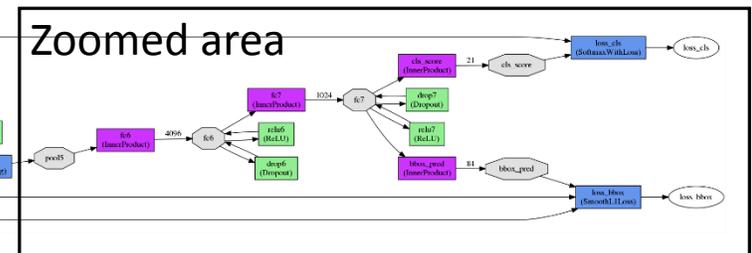
```
2 layer {
3   name: 'data'
4   type: 'Python'
5   top: 'data'
6   top: 'rois'
7   top: 'labels'
8   top: 'bbox_targets'
9   top: 'bbox_loss_weights'
10  python_param {
11    module: 'roi_data_layer.layer'
12    layer: 'RoIDataLayer'
13    param_str: "'num_classes': 21"
14  }
15 }
```

The train-time net (multi-task losses)

Classification loss
(Cross-entropy)



Bounding-box regression loss
("Smooth L1")



Fast R-CNN — Edit

172 commits 1 branch 0 releases 2 contributors

branch: master fast-rcnn / +

Merge pull request #8 from drozdvdym/fix_cpu_mode

rbgirshick authored on May 6 latest commit b0758d0a67

caffe-fast-rcnn @ bcd9b4e	Update caffe submodule (no roi pooling CPU test)	a month ago
data	add mat files for boxes, remove pickled version (and update demo.py)	a month ago
experiments	Make with python layer	a month ago
lib	add args to tools/reval.py; fix comp_mode abstraction	a month ago
matlab	rm accidentally added file	a month ago
models	Improve readmes	a month ago
output	Improve readmes	a month ago
tools	fixed cpu mode	a month ago
.gitignore	ignore ipython notebook checkpoints	3 months ago
.gitmodules	change to https urls for easier installation	a month ago
LICENSE	update license	2 months ago
README.md	Add pointers for computing object proposals	a month ago
todo.txt	update todos	a month ago

- Code
- Issues 11
- Pull requests 2
- Wiki
- Pulse
- Graphs
- Settings

HTTPS clone URL https://github.com/

You can clone with HTTPS, SSH, or Subversion.

Clone in Desktop

Download ZIP

Fast R-CNN: Fast Region-based Convolutional Networks for object detection

Created by Ross Girshick at Microsoft Research, Redmond.

Code is on
GitHub
(MIT License,
Runs on Linux)

A brief tour of some of the code

Caffe fork

→	📁 caffe-fast-rcnn @ bcd9b4e	Update caffe submodule (no roi pooling CPU test)	a month ago
	📁 data	add mat files for boxes, remove pickled version (and update demo.py)	a month ago
	📁 experiments	Make with python layer	a month ago
Python modules	→ 📁 lib	add args to tools/reval.py; fix comp_mode abstraction	a month ago
	📁 matlab	rm accidentally added file	a month ago
	📁 models	Improve readmes	a month ago
	📁 output	Improve readmes	a month ago
Train, test	→ 📁 tools	fixed cpu mode	a month ago
	📄 .gitignore	ignore ipython notebook checkpoints	3 months ago
	📄 .gitmodules	change to https urls for easier installation	a month ago
	📄 LICENSE	update license	2 months ago
	📄 README.md	Add pointers for computing object proposals	a month ago
	📄 todo.txt	update todos	a month ago

A brief tour of some of the code (Caffe bits)

Caffe fork

	 caffe-fast-rcnn @ bcd9b4e	Update caffe submodule (no roi pooling CPU test)	a month ago
	 data	add mat files for boxes, remove pickled version (and update demo.py)	a month ago
	 experiments	Make with python layer	a month ago
	 lib	add args to tools/reval.py; fix comp_mode abstraction	a month ago
	 matlab	rm accidentally added file	a month ago
	 models	Improve readmes	a month ago
	 output	Improve readmes	a month ago
	 tools	fixed cpu mode	a month ago
	 .gitignore	ignore ipython notebook checkpoints	3 months ago
	 .gitmodules	change to https urls for easier installation	a month ago
	 LICENSE	update license	2 months ago
	 README.md	Add pointers for computing object proposals	a month ago
	 todo.txt	update todos	a month ago

Python modules

Train, test

branch: fast-rcnn | caffe-fast-rcnn / src / caffe / layers / roi_pooling_layer.cu

rbgirshick on Apr 1 add spatial scale message field; fix some variable naming

1 contributor

189 lines (163 sloc) | 7.633 kB

Raw Blame History

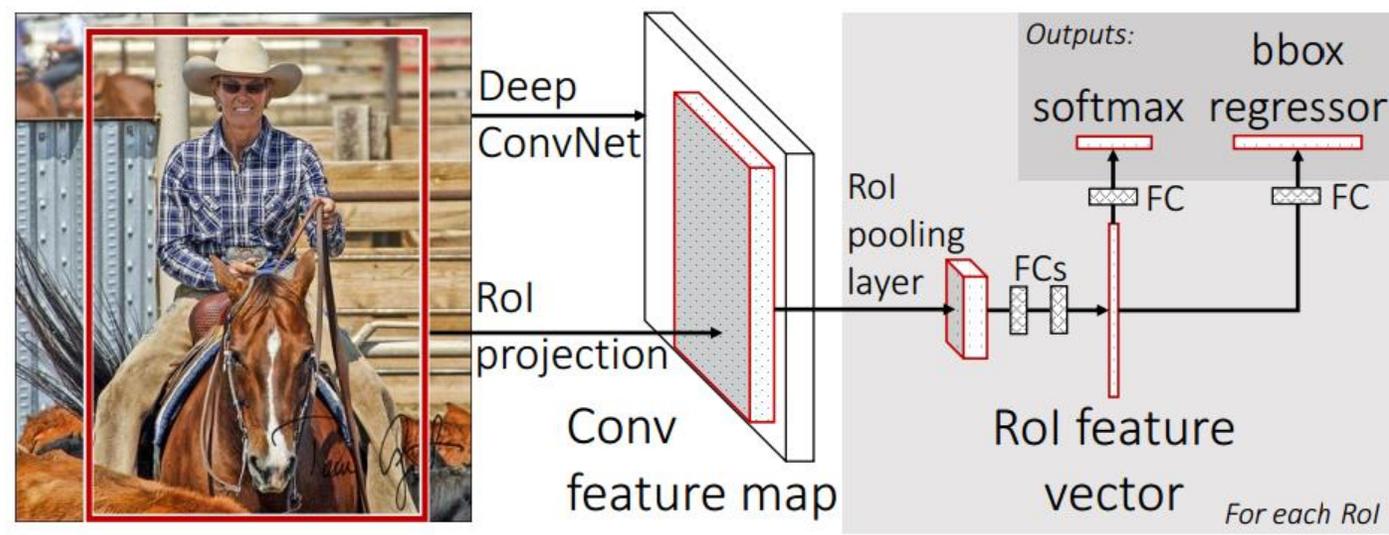
```

1 // -----
2 // Fast R-CNN
3 // Copyright (c) 2015 Microsoft
4 // Licensed under The MIT License [see fast-rcnn/LICENSE for details]
5 // Written by Ross Girshick
6 // -----
7
8 #include <float>
9
10 #include "caffe/fast_rcnn_layers.hpp"
11
12 using std::max;
13 using std::min;
14
15 namespace caffe {
16
17 template <typename Dtype>
18 __global__ void ROIPoolForward(const int nthreads, const Dtype* bottom_data,
19     const Dtype spatial_scale, const int channels, const int height,
20     const int width, const int pooled_height, const int pooled_width,
21     const Dtype* bottom_rois, Dtype* top_data, int* argmax_data) {
22     CUDA_KERNEL_LOOP(index, nthreads) {
23         // (n, c, ph, pw) is an element in the pooled output
24         int pw = index % pooled_width;
25         int ph = (index / pooled_width) % pooled_height;
26         int c = (index / pooled_width / pooled_height) % channels;
27         int n = index / pooled_width / pooled_height / channels;
28     }
29 }

```

Region of Interest (RoI) Pooling Layer

Expands a small batch into a big batch



rbgirshick on Mar 30 add license blurb; split out fast r-cnn layers into separate header file
1 contributor

91 lines (81 sloc) | 2.809 kB
Raw Blame History

```

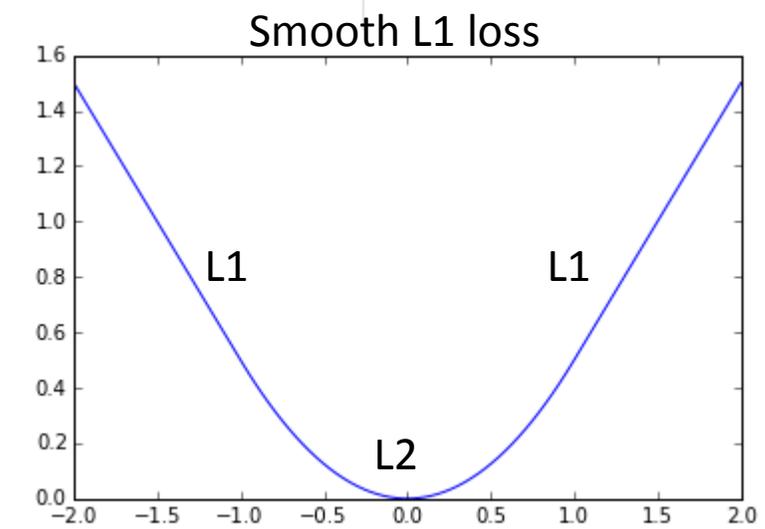
1 // -----
2 // Fast R-CNN
3 // Copyright (c) 2015 Microsoft
4 // Licensed under The MIT License [see fast-rcnn/LICENSE for details]
5 // Written by Ross Girshick
6 // -----
7
8 #include "caffe/fast_rcnn_layers.hpp"
9
10 namespace caffe {
11
12 template <typename Dtype>
13 __global__ void SmoothL1Forward(const int n, const Dtype* in, Dtype* out) {
14 // f(x) = 0.5 * x^2    if |x| < 1
15 //         |x| - 0.5   otherwise
16 CUDA_KERNEL_LOOP(index, n) {
17     Dtype val = in[index];
18     Dtype abs_val = abs(val);
19     if (abs_val < 1) {
20         out[index] = 0.5 * val * val;
21     } else {
22         out[index] = abs_val - 0.5;
23     }
24 }
25 }
26

```

Smooth L1 Loss Layer

Robust to outliers
Optimizer friendly

Per-dimension loss weights



A brief tour of some of the code (Python bits)

Caffe fork

→	caffe-fast-rcnn @ bcd9b4e	Update caffe submodule (no roi pooling CPU test)	a month ago
	data	add mat files for boxes, remove pickled version (and update demo.py)	a month ago
	experiments	Make with python layer	a month ago
Python modules	lib	add args to tools/reval.py; fix comp_mode abstraction	a month ago
	matlab	rm accidentally added file	a month ago
	models	Improve readmes	a month ago
	output	Improve readmes	a month ago
Train, test	tools	fixed cpu mode	a month ago
	.gitignore	ignore ipython notebook checkpoints	3 months ago
	.gitmodules	change to https urls for easier installation	a month ago
	LICENSE	update license	2 months ago
	README.md	Add pointers for computing object proposals	a month ago
	todo.txt	update todos	a month ago

rbgirshick on Apr 29 cleanup script files

1 contributor

161 lines (131 sloc) | 5.93 kB

Raw Blame History

Python data layer for Fast R-CNN

Reshapes blobs on-the-fly

```

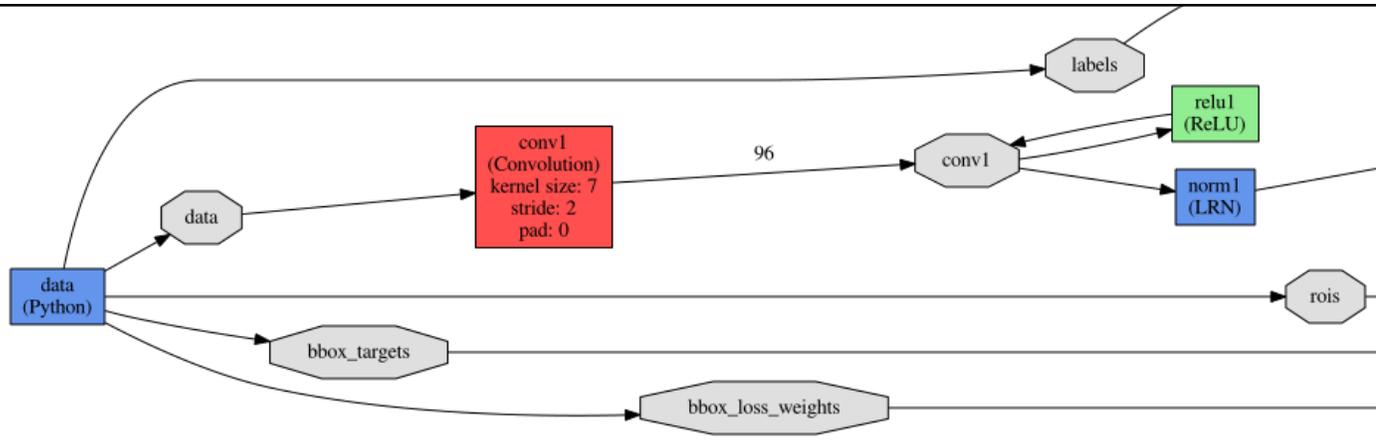
1 # -----
2 # Fast R-CNN
3 # Copyright (c) 2015 Microsoft
4 # Licensed under The MIT License [see LICENSE for details]
5 # Written by Ross Girshick
6 # -----
7
8 """The data layer used during training to train a Fast R-CNN
9
10 RoIDataLayer implements a Caffe Python layer.
11 """
12
13 import caffe
14 from fast_rcnn.config import cfg
15 from roi_data_layer.minibatch import get_minibatch
16 import numpy as np
17 import yaml
18 from multiprocessing import Process, Queue
19
20 class RoIDataLayer(caffe.Layer):
21     """Fast R-CNN data layer used for training."""
22
23     def _shuffle_roidb_inds(self):
24         """Randomly permute the training roidb."""
25         self._perm = np.random.permutation(np.arange(len(self._roidb)))
26         self._cur = 0
27
28     def _get_next_minibatch_inds(self):
29         """Return the roidb indices for the next minibatch."""
30         if self._cur + cfg.TRAIN.IMS_PER_BATCH >= len(self._roidb):

```

```

106     def forward(self, bottom, top):
107         """Get blobs and copy them into this layer's top blob vector."""
108         blobs = self._get_next_minibatch()
109
110         for blob_name, blob in blobs.iteritems():
111             top_ind = self._name_to_top_map[blob_name]
112             # Reshape net's input blobs
113             top[top_ind].reshape(*(blob.shape))
114             # Copy data into net's input blobs
115             top[top_ind].data[...] = blob.astype(np.float32, copy=False)

```



```
1 # -----
2 # Fast R-CNN
3 # Copyright (c) 2015 Microsoft
4 # Licensed under The MIT License [see LICENSE for details]
5 # Written by Ross Girshick
6 # -----
7
8 """Train a Fast R-CNN network."""
9
10 import caffe
11 from fast_rcnn.config import cfg
12 import roi_data_layer.roidb as rdl_roidb
13 from utils.timer import Timer
14 import numpy as np
15 import os
16
17 from caffe.proto import caffe_pb2
18 import google.protobuf as pb2
19
20 class SolverWrapper(object):
21     """A simple wrapper around Caffe's solver.
22     This wrapper gives us control over the snapshotting process, which
23     we use to unnormalize the learned bounding-box regression weights.
24     """
25
26     def __init__(self, solver_prototxt, roidb, output_dir,
27                 pretrained_model=None):
28         """Initialize the SolverWrapper."""
29         self.output_dir = output_dir
```

```
84     def train_model(self, max_iters):
85         """Network training loop."""
86         last_snapshot_iter = -1
87         timer = Timer()
88         while self.solver.iter < max_iters:
89             # Make one SGD update
90             timer.tic()
91             self.solver.step(1)
92             timer.toc()
93             if self.solver.iter % (10 * self.solver_param.display) == 0:
94                 print 'speed: {:.3f}s / iter'.format(timer.average_time)
95
96             if self.solver.iter % cfg.TRAIN.SNAPSHOT_ITERS == 0:
97                 last_snapshot_iter = self.solver.iter
98                 self.snapshot()
99
100         if last_snapshot_iter != self.solver.iter:
101             self.snapshot()
```

Python training code

Custom solver loop
with custom snapshot
method

A brief tour of some of the code (CLI tools)

Caffe fork

→	📁 caffe-fast-rcnn @ bcd9b4e	Update caffe submodule (no roi pooling CPU test)	a month ago
	📁 data	add mat files for boxes, remove pickled version (and update demo.py)	a month ago
	📁 experiments	Make with python layer	a month ago
Python modules	→ 📁 lib	add args to tools/reval.py; fix comp_mode abstraction	a month ago
	📁 matlab	rm accidentally added file	a month ago
	📁 models	Improve readmes	a month ago
	📁 output	Improve readmes	a month ago
Train, test	→ 📁 tools	fixed cpu mode	a month ago
	📄 .gitignore	ignore ipython notebook checkpoints	3 months ago
	📄 .gitmodules	change to https urls for easier installation	a month ago
	📄 LICENSE	update license	2 months ago
	📄 README.md	Add pointers for computing object proposals	a month ago
	📄 todo.txt	update todos	a month ago

fixed cpu mode

drozdvadym authored on May 6

latest commit 962672b41e

..

README.md	Improve readmes	a month ago
_init_paths.py	minor refactoring	a month ago
compress_net.py	tool script docstrings	a month ago
demo.py	fixed cpu mode	a month ago
reval.py	add args	
test_net.py	add args	
train_net.py	fix the rng	
train_svms.py	tool script	

README.md

Tools for training, testing, and

```
rbg@rbgk40: ~/working/fast-rcnn/fast-rcnn
rbg@rbgk40:~/working/fast-rcnn/fast-rcnn$ ./tools/demo.py -h
usage: demo.py [-h] [--gpu GPU_ID] [--cpu]
              [--net {vgg16,caffenet,vgg_cnn_m_1024}]

Train a Fast R-CNN network

optional arguments:
  -h, --help            show this help message and exit
  --gpu GPU_ID          GPU device id to use [0]
  --cpu                 Use CPU mode (overrides --gpu)
  --net {vgg16,caffenet,vgg_cnn_m_1024}
                        Network to use [vgg16]
rbg@rbgk40:~/working/fast-rcnn/fast-rcnn$ ./tools/demo.py --gpu 7
```

Figure 1@rbgk40

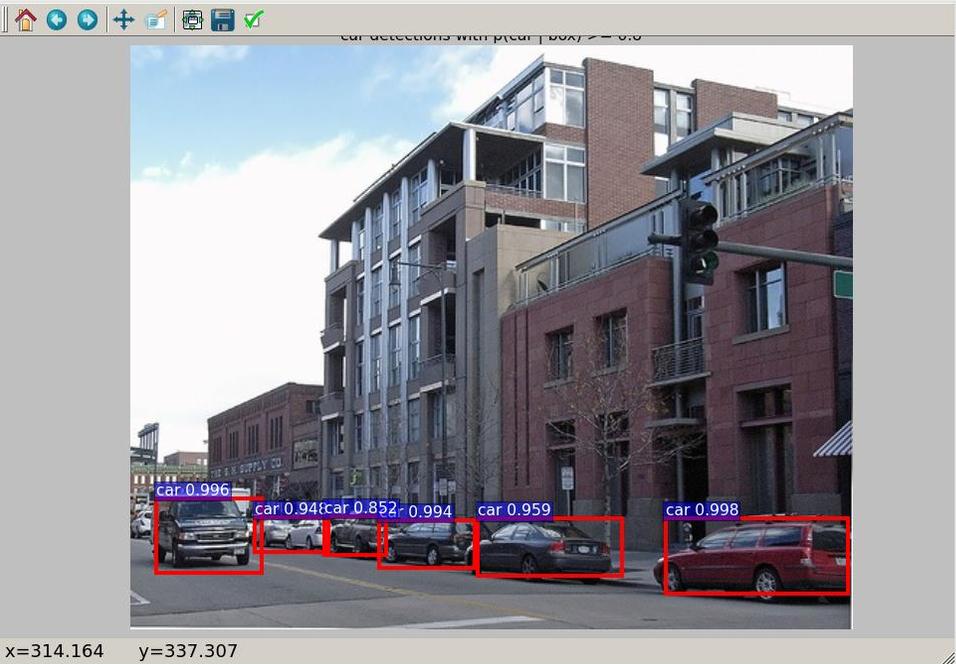


Figure 2@rbgk40

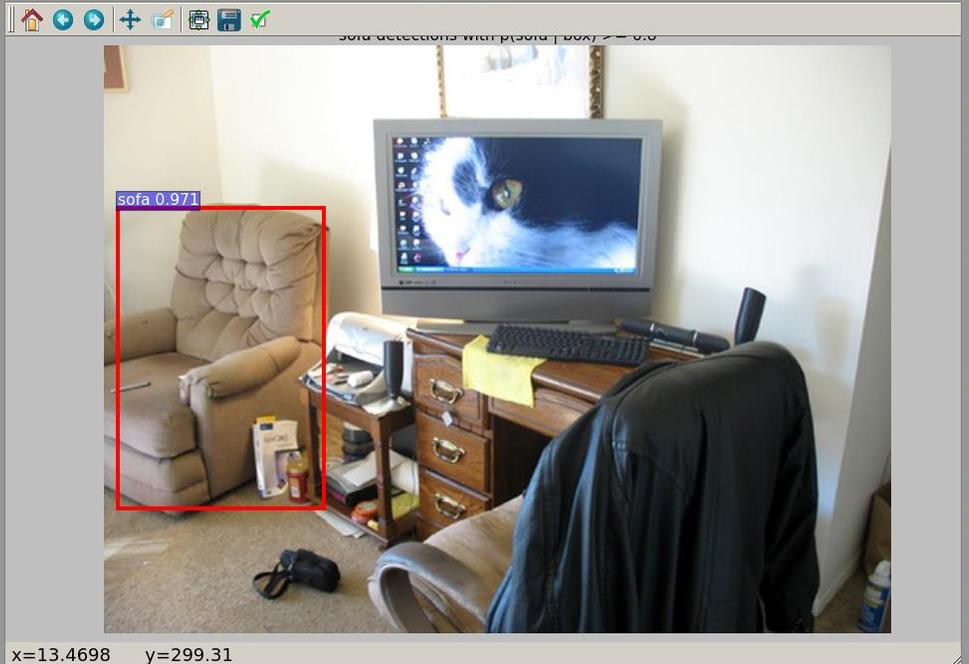
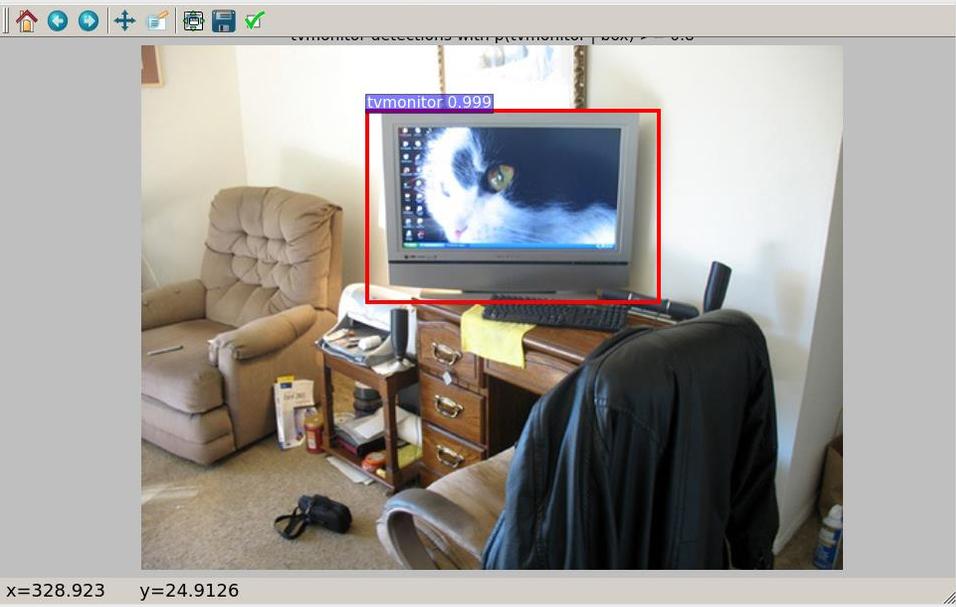


Figure 3@rbgk40



```

fast-rcnn/fast-rcnn
2_1 does not need backward computatio
/2_1 does not need backward computatio
1 does not need backward computation.
1_2 does not need backward computatio
/1_2 does not need backward computatio
1_1 does not need backward computatio
/1_1 does not need backward computatio
s network produces output bbox_pred
s network produces output cls_prob
lecting Learning Rate and Weight Decay
work initialization done.
ory required for data: 114633208
d_stream.cc:505] Reading dangerously l
large protocol message. If the message turns out to be larger than 2147483647 byte
s, parsing will be halted for security reasons. To increase the limit (or to disa
ble these warnings), see CodedInputStream::SetTotalBytesLimit() in google/protobuf
/io/coded_stream.h.
[libprotobuf WARNING google/protobuf/io/coded_stream.cc:78] The total number of by
tes read was 538766130

Loaded network /mnt/data/rbg/fast-rcnn/fast-rcnn/data/fast_rcnn_models/vgg16_fast_
rcnn_iter_40000.caffemodel
~~~~~
Demo for data/demo/000004.jpg
Detection took 0.578s for 2888 object proposals
All car detections with p(car | box) >= 0.8
~~~~~
Demo for data/demo/001551.jpg
Detection took 0.364s for 2057 object proposals
All sofa detections with p(sofa | box) >= 0.8
All tvmonitor detections with p(tvmonitor | box) >= 0.8

```

Teaser: *Faster R-CNN*

Shaoqing Ren, Kaiming He, Ross Girshick, Jian Sun. Microsoft Research

- The detection network *also proposes objects*
- Marginal cost of proposals: 10ms
- VGG16 runtime $\sim 200\text{ms}$ *including all steps*
- Higher mAP, faster
- Open-source Caffe code coming later this summer

Region Proposal Network shares conv layers with Fast R-CNN object detection network

