

*Conceitos Básicos de Deep
Learning e o seu uso através do
NVIDIA® DIGITS™ (parte 2)*

Cristina Nader Vasconcelos

Cronograma

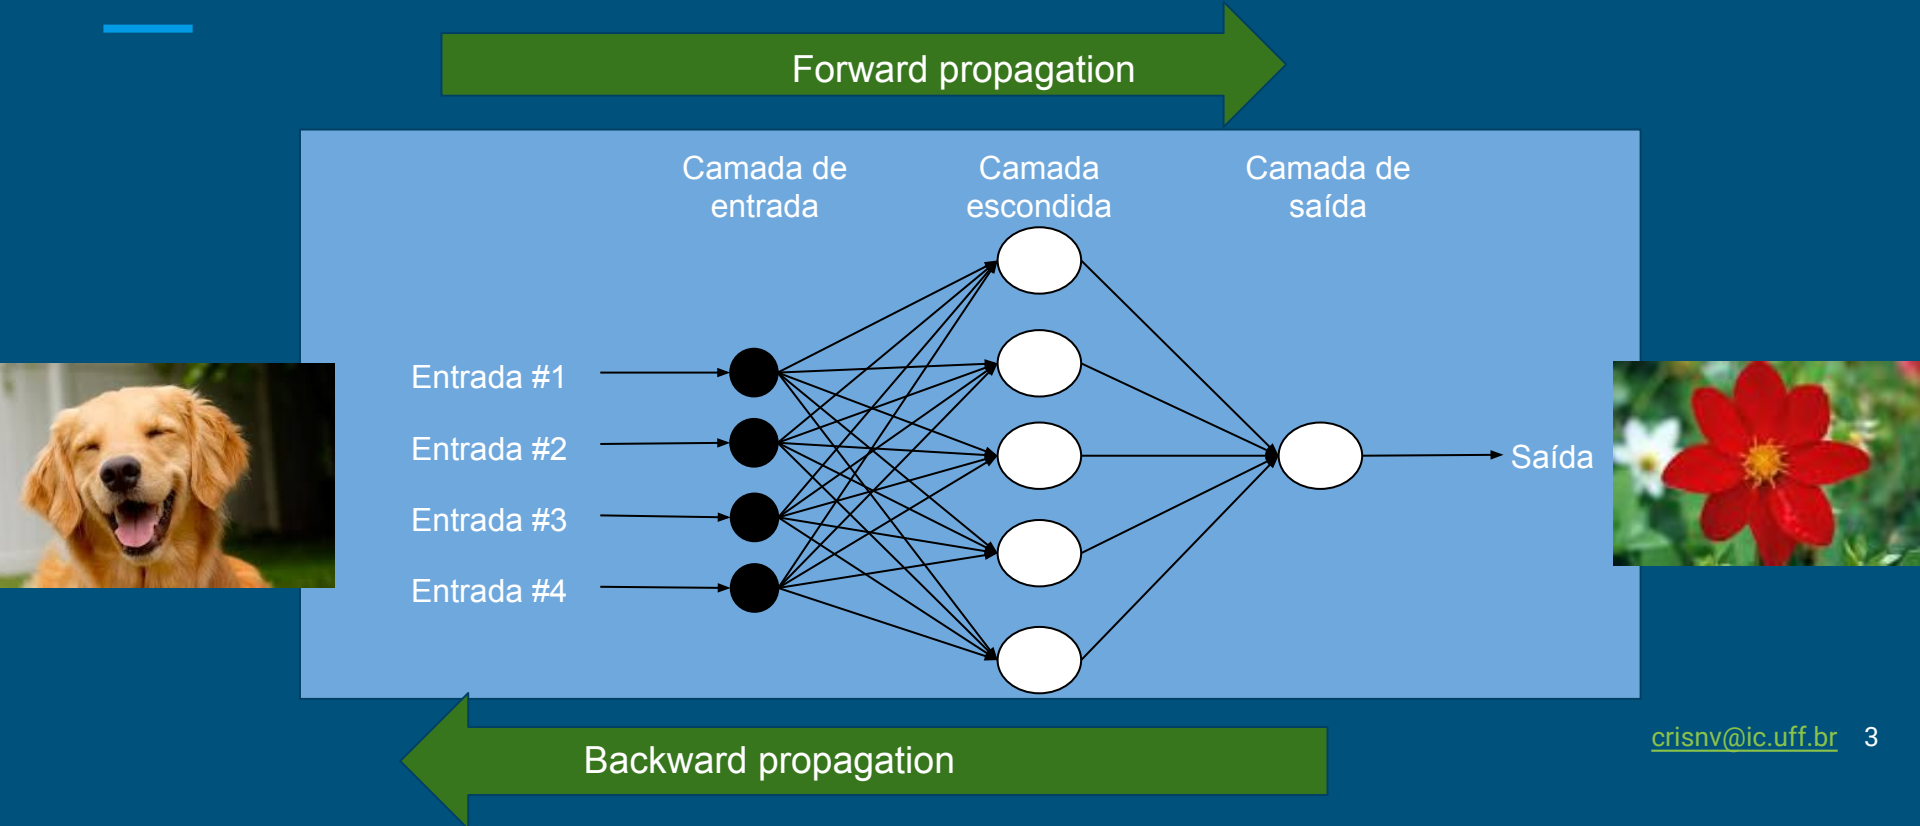
Webinar I

- Motivação
- Abordagem Clássica para reconhecimento de padrões em sinais
- Exemplos: features, descriptors, classifiers
- Desafios existentes
- ML: modelos de aprendizado
- Redes neurais: formulação clássica
- CNNs
- Aplicações

Webinar II

- Treinamento
- GPUs e DNN
- Hierarquia de programação Deep Learning
- CuDNN, Caffe, Digits
- Digits
- Exemplo: Alexnet
- Perguntas

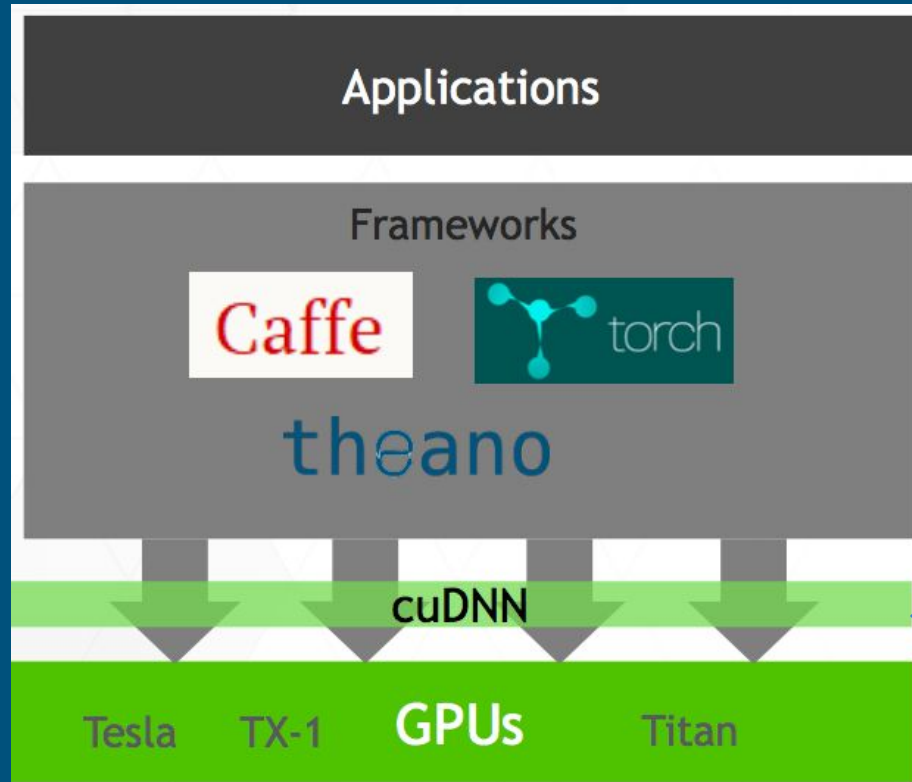
Treinamento



Aprendizado profundo e GPUs

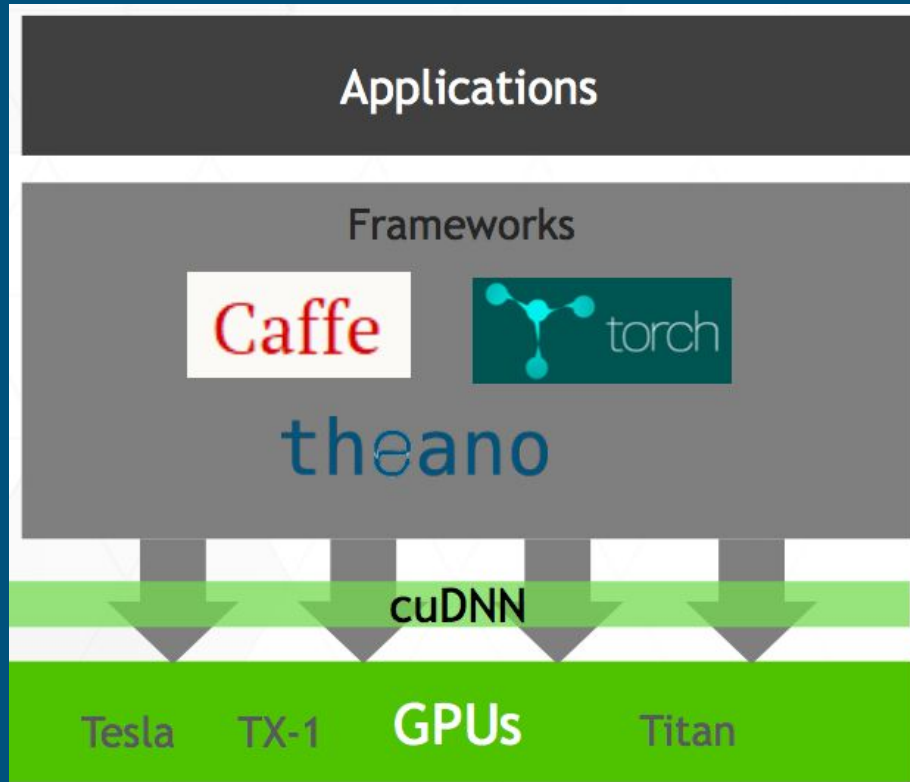
	redes neurais	GPUs
inerentemente paralela	sim	sim
operações de matriz	sim	sim
FLOPS	sim	sim

Hierarquia de programação



cuDNN:
biblioteca de
primitivas de
deep learning

Hierarquia de programação

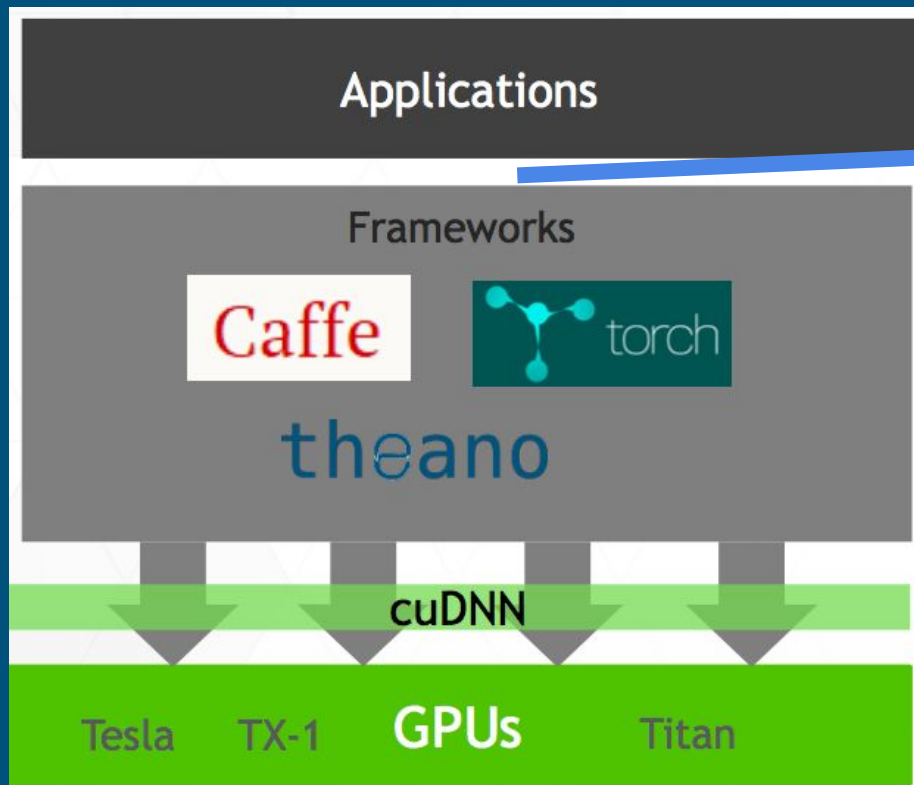


Frameworks:

- Caffe



Hierarquia de programação



- Digits

NVIDIA DIGITS:



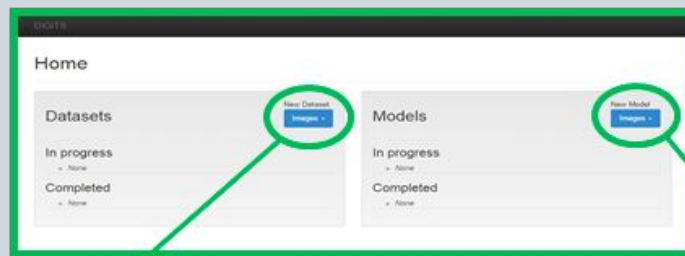
Interactive Deep Learning GPU Training System

Permite:

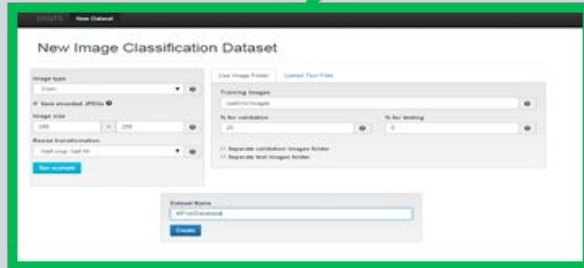
- visualizar a topologia de DNNs e a ativação por dados de treinamento;
- gerenciar o treinamento de conjuntos de DNNs em paralelo em sistemas multi-GPU
- importar diferentes formatos e fontes de imagens
- monitorar o treinamento da rede em tempo real
- open source (pode ser personalizado e estendido)

Interface Digits

Main Console

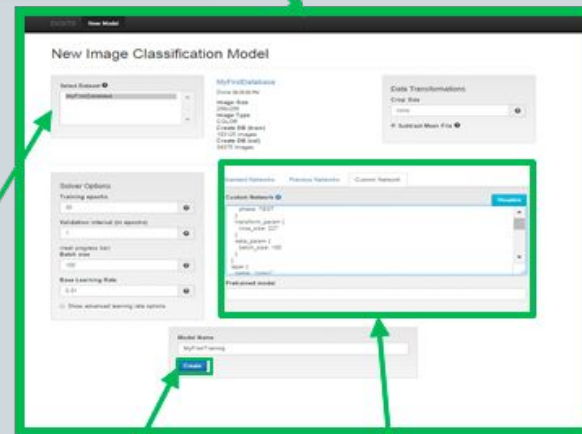


Create your Dataset



Choose your dataset

Configure your Network



Start Training

Choose a default network, modify one, or create your own

Interface Digits: página principal

DIGITS

Home

Datasets

New Dataset
Images ▾

In progress
◦ None

Completed

- Database3** Delete
Submitted: Thu Mar 12, 10:52:38 PM
Status: Done after 5 minutes, 7 seconds
- Database2** Delete
Submitted: Tue Mar 10, 01:42:18 PM
Status: Done after 8 minutes, 17 seconds
- Database1** Delete
Submitted: Tue Mar 10, 12:46:39 PM
Status: Done after 13 minutes, 50 seconds

Models

New Model
Images ▾

In progress

- TrainingRun6** Delete
Submitted: 01:26:02 PM (8 seconds ago)
Status: Running
- TrainingRun5** Delete
Submitted: 12:45:38 PM (40 minutes, 32 seconds ago)
Status: Running

Completed

- TrainingRun4** Delete
Submitted: Tue Mar 10, 09:30:50 PM
Status: Aborted after 7 hours, 54 minutes
- TrainingRun3** Delete
Submitted: Tue Mar 10, 09:10:26 PM
Status: Aborted after 8 hours, 14 minutes
- TrainingRun2** Delete
Submitted: Tue Mar 10, 05:04:29 PM
Status: Done after 9 hours, 48 minutes

crisnv@icuff.br 10

Interface Digits: página principal

DIGITS

Home

Datasets

New Dataset
Images -

In progress
◊ None

Completed

Database3 Submitted: Thu Mar 12, 10:52:38 PM Status: Done after 5 minutes, 7 seconds	Delete
Database2 Submitted: Tue Mar 10, 01:42:18 PM Status: Done after 8 minutes, 17 seconds	Delete
Database1 Submitted: Tue Mar 10, 12:46:39 PM Status: Done after 13 minutes, 50 seconds	Delete

Models

New Model
Images -

In progress

TrainingRun6 Submitted: 01:26:02 PM (8 seconds ago) Status: Running	Delete
TrainingRun5 Submitted: 12:45:38 PM (40 minutes, 32 seconds ago) Status: Running	Delete

Completed

TrainingRun4 Submitted: Tue Mar 10, 09:30:50 PM Status: Aborted after 7 hours, 54 minutes	Delete
TrainingRun3 Submitted: Tue Mar 10, 09:10:26 PM Status: Aborted after 8 hours, 14 minutes	Delete
TrainingRun2 Submitted: Tue Mar 10, 05:04:29 PM Status: Done after 9 hours, 48 minutes	Delete

crisnv@icuff.br 11

Criando uma base

DiGiTS New Dataset

New Image Classification Dataset

DiGiTS can automatically create your training and validation set

OR

Insert the path to your train and validation set

Use Image Folder **Upload Text Files**

Training Images: /path/to/images

% for validation: 25 % for testing: 0

Separate validation images folder
 Separate test images folder


Image type: Color

Save encoded JPEGs

Image size: 256 x 256

Resize transformation: Squash

See example



Input image parameter options

Set	Text file	Image folder (optional)
Training	<input type="button" value="Choose File"/> train.txt	/path/to/train/
<input checked="" type="checkbox"/> Validation	<input type="button" value="Choose File"/> val.txt	/path/to/val/
<input type="checkbox"/> Test	<input type="button" value="Choose File"/> No file chosen	

Labels: synsets.txt

Dataset Name: Database3

Create your dataset

crisnv@ic.uff.br 12

Interface Digits: página principal

The screenshot displays the DIGITS interface home page. The top navigation bar is dark grey with the text 'DIGITS' in white. Below the navigation bar, the page is titled 'Home'. There are two main panels: 'Datasets' on the left and 'Models' on the right. The 'Models' panel is highlighted with a red border. Each panel has a 'New Dataset' or 'New Model' button with a dropdown arrow. The 'In progress' section in both panels is currently empty. The 'Completed' section in the 'Datasets' panel lists three datasets: Database3, Database2, and Database1, each with submission and completion details and a 'Delete' button. The 'Completed' section in the 'Models' panel lists three training runs: TrainingRun4, TrainingRun3, and TrainingRun2, each with submission and completion details and a 'Delete' button. The 'In progress' section in the 'Models' panel lists two training runs: TrainingRun6 and TrainingRun5, each with submission and completion details and a 'Delete' button.

DIGITS

Home

Datasets

New Dataset
Images ▾

In progress
◦ None

Completed

Database3 Submitted: Thu Mar 12, 10:52:38 PM Status: Done after 5 minutes, 7 seconds	Delete
Database2 Submitted: Tue Mar 10, 01:42:18 PM Status: Done after 8 minutes, 17 seconds	Delete
Database1 Submitted: Tue Mar 10, 12:46:39 PM Status: Done after 13 minutes, 50 seconds	Delete

Models

New Model
Images ▾

In progress

TrainingRun6 Submitted: 01:26:02 PM (8 seconds ago) Status: Running	Delete
TrainingRun5 Submitted: 12:45:38 PM (40 minutes, 32 seconds ago) Status: Running	Delete

Completed

TrainingRun4 Submitted: Tue Mar 10, 09:30:50 PM Status: Aborted after 7 hours, 54 minutes	Delete
TrainingRun3 Submitted: Tue Mar 10, 09:10:26 PM Status: Aborted after 8 hours, 14 minutes	Delete
TrainingRun2 Submitted: Tue Mar 10, 05:04:29 PM Status: Done after 9 hours, 48 minutes	Delete

crisov@icuff.br 13

New Image Classification Model

Select Dataset

- Database2
- Database1

Database1

Done Tue Mar 10, 01:00:29 PM

Image Size
256x256

Image Type
COLOR

Create DB (train)
137500 images

Create DB (val)
6066 images

Data Transformations

Crop Size
none

Subtract Mean File

Standard Networks Previous Networks Custom Network

Network	Pretrained Model	
<input type="radio"/> TrainingRun5 View	None	Customize
<input type="radio"/> TrainingRun4 View	None	Customize
<input type="radio"/> TrainingRun3 View	None	Customize

OR choose a previously used one

Solver Options

Training epochs
30

Validation interval (in epochs)
1

(neat progress bar)

Batch size
100

Base Learning Rate
0.01

Show advanced learning rate options

Standard Networks Previous Networks Custom Network

Network	Description	Intended image size	
<input type="radio"/> LeNet	Yan LeCunn's network. Home page	28x28	Customize
<input checked="" type="radio"/> AlexNet	Alex Krizhevsky's network. Original paper	256x256	Customize

Choose a preconfigured network

OR add it here

Standard Networks Previous Networks Custom Network

Custom Network

Visualize

Insert your network here

Pretrained model

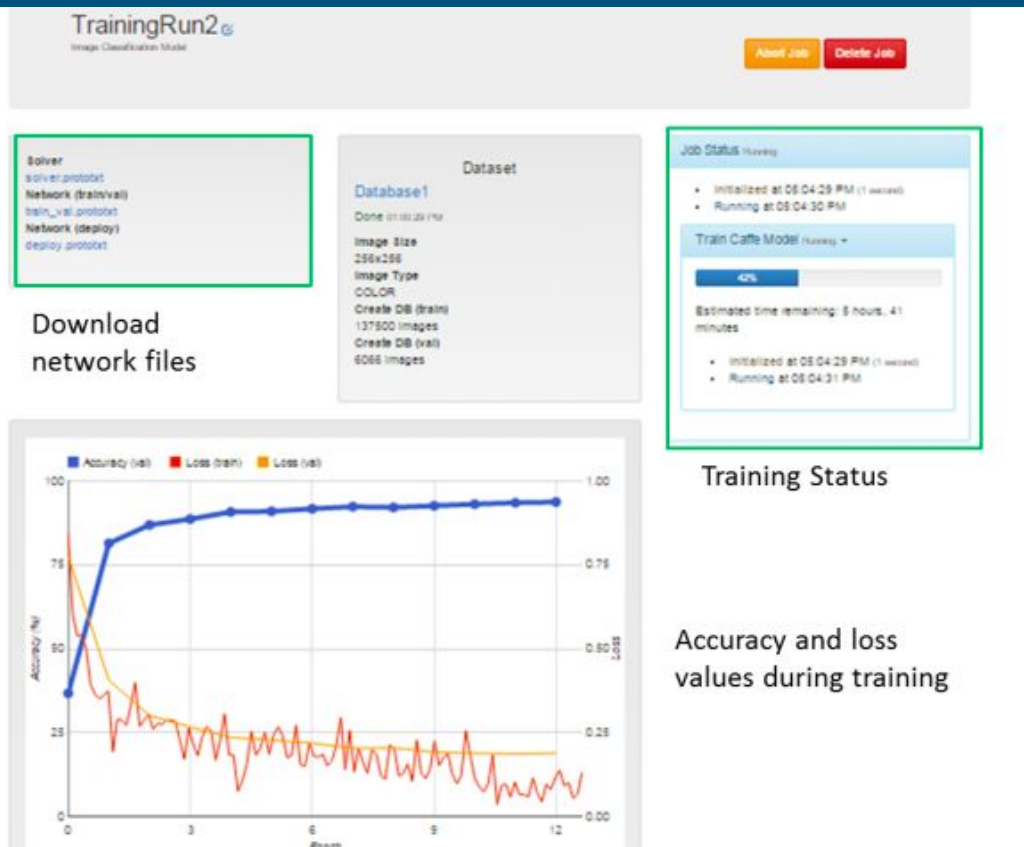
Model Name

TrainingRun7

Create

Start training

Interface Digits: treinamento



Interface Digits: treinamento

Select Model

Epoch#12

Download

Image URL

Upload Image

Choose File No file chosen

Test one image

Upload Image List

Choose File No file chosen

Accept a list of filenames or urls (you can use your own list file)

Number of images use from the file

100

Leave blank to use all

Number of images to show per category

9

Test several images This takes a while, be patient.

Classification on the fly with most recent or previous network snapshots



Predictions

ship

100.0%

no ship



0.0%

Layer

Activations

Weights

conv1





conv1

Activation


Data shape: (96, 55, 55)

Mean: 6.01675


Std deviation: 11.7991



AlexNet



Alex Krizhevsky Ilya Sutskever Geoffrey E. Hinton. ImageNet Classification with Deep Convolutional Neural Networks. NIPS 2012



```


layer {
  name: "conv1"
  type: "Convolution"
  bottom: "data"
  top: "conv1"
  param {
    lr_mult: 1
    decay_mult: 1
  }
  param {
    lr_mult: 2
    decay_mult: 0
  }
}

```

```

convolution_param {
  num_output: 96
  kernel_size: 11
  stride: 4
  weight_filler {
    type: "gaussian"
    std: 0.01
  }
  bias_filler {
    type: "constant"
    value: 0
  }
}

```



```

layer {
  name: "relu1"
  type: "ReLU"
  bottom: "conv1"
  top: "conv1"
}

```

```


layer {
  name: "norm1"
  type: "LRN"
  bottom: "conv1"
  top: "norm1"
  lrn_param {
    local_size: 5
    alpha: 0.0001
    beta: 0.75
  }
}

```

```

layer {
  name: "pool1"
  type: "Pooling"
  bottom: "norm1"
  top: "pool1"
  pooling_param {
    pool: MAX
    kernel_size: 3
    stride: 2
  }
}

```



```


layer {
  name: "conv2"
  type: "Convolution"
  bottom: "pool1"
  top: "conv2"
  param {
    lr_mult: 1
  }
  decay_mult: 1
}
param {
  lr_mult: 2
  decay_mult: 0 }

```

```

convolution_param {
  num_output: 256
  pad: 2  kernel_size: 5  group: 2
  weight_filler {
    type: "gaussian"
    std: 0.01
  }
  bias_filler {
    type: "constant"
    value: 0.1
  }
}

```



```

layer {
  name: "relu2"
  type: "ReLU"
  bottom: "conv2"
  top: "conv2"
}

```

```


layer {
  name: "norm2"
  type: "LRN"
  bottom: "conv2"
  top: "norm2"
  lrn_param {
    local_size: 5
    alpha: 0.0001
    beta: 0.75
  }
}

```

```

layer {
  name: "pool2"
  type: "Pooling"
  bottom: "norm2"
  top: "pool2"
  pooling_param {
    pool: MAX
    kernel_size: 3
    stride: 2
  }
}

```

```

layer {
  name: "conv3" ...
  convolution_param {
    num_output: 384
    pad: 1
    kernel_size: 3 ...}
  layer { name: "relu3" ...}

```

```


layer {
  name: "conv4" ...
  convolution_param {
    num_output: 384
    pad: 1
    kernel_size: 3
    group: 2 ...
  }
  layer { name: "relu4" ...}

```

```

layer {
  name: "conv5"
  convolution_param {
    num_output: 256
    pad: 1
    kernel_size: 3
    group: 2... }
  layer { name: "relu5" ...}
  layer { name: "pool5" ...
    kernel_size: 3
    stride: 2}

```

```

layer { name: "fc6"
  type: "InnerProduct"
  bottom: "pool5"
  top: "fc6"
  ...
  inner_product_param {
    num_output: 4096
  }
  ...
}

```

```


layer {
  name: "relu6"
  type: "ReLU" ..}
layer {
  name: "drop6"
  type: "Dropout"...
  dropout_param {
    dropout_ratio: 0.5
  }
}

```

```

layer { name: "fc7"
  type: "InnerProduct"
  ... num_output: 4096}
layer {
  name: "relu7" ...
}
layer {
  name: "drop7"...
}
layer { name: "fc8"
  type: "InnerProduct"
  ...}

```



```

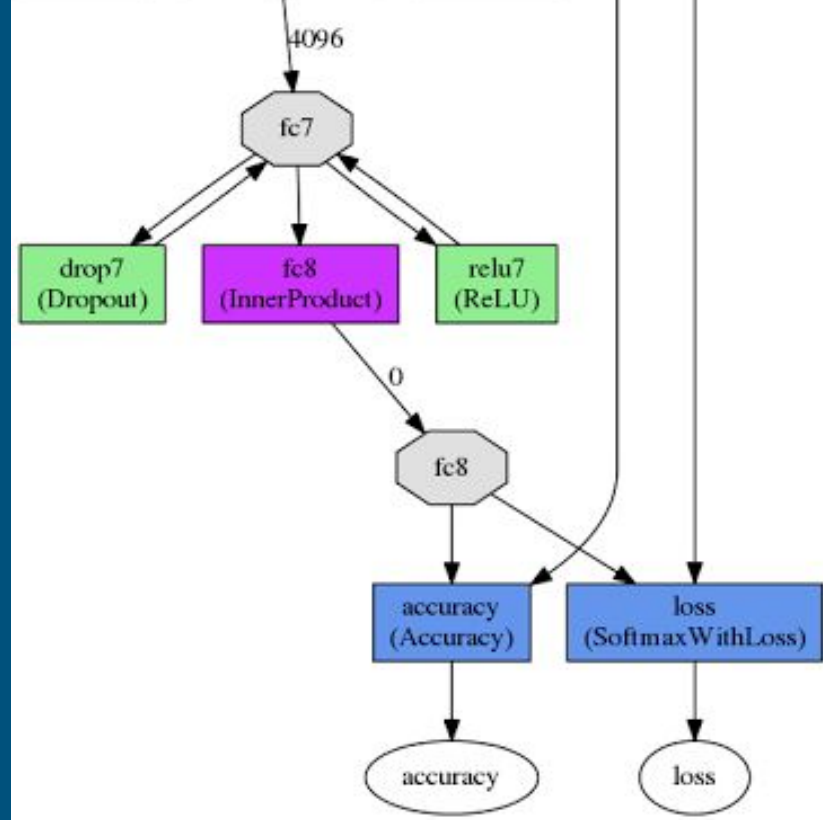
layer {
  name: "accuracy"
  type: "Accuracy"
  bottom: "fc8"
  bottom: "label"
  top: "accuracy"
  include {
    phase: TEST
  }
}

```

```

layer {
  name: "loss"
  type: "SoftmaxWithLoss"
  bottom: "fc8"
  bottom: "label"
  top: "loss"}

```





slides em:

www.ic.uff.br/~crisnv

obrigada!

crisnv@ic.uff.br