# Unsupervised Deep Learning with Stacked Autoencoders on Chameleon

**Dr. Pablo Rivas*** **Ezequiel Rivas^** **Deep Dand***

* *Marist College,* School of Computer Science and Mathematics

^ *Nogales Institute of Technology*, Research and Grad. Studies

MARIST

# The Problem

- American Sign Language (ASL)

# Existing Solutions

- ***Learning*** the American Sign Language (ASL)

# Existing Solutions

- ***Learning*** the American Sign Language (ASL)

MARIST

# Existing Solutions

- ***Learning*** the American Sign Language (ASL)

MARIST

# Existing Solutions

| Method | Class type | # of class | # of subj. | Test w/ diff. | Input | Accur.(%) |
|---|---|---|---|---|---|---|
| Nagi *et al.* [8] | Gesture | 6 | - | No | Color | 96 |
| Van den Bergh *et al.* [14] | Gesture | 6 | - | No | Color & Depth | 99.54 |
| Isaacs *et al.* [3] | Alphabets | 24 | - | - | Color | 99.9 |
| Pugeault *et al.* [10] | Alphabets | 24 | 5 | - | Color | 73 |
| Pugeault *et al.* [10] | Alphabets | 24 | 5 | - | Depth | 69 |
| Pugeault *et al.* [10] | Alphabets | 24 | 5 | - | Color & Depth | 75 |
| Kuznetsova *et al.* [6] (50/50)% | Alphabets | 24 | 5 | No | Depth | 87 |
| Kuznetsova *et al.* [6] (4/1) | Alphabets | 24 | 5 | Yes | Depth | 57 |
| Dong *et al.* [2] (50/50)% | Alphabets | 24 | 5 | No | Depth | 90 |
| Dong *et al.* [2] (4/1) | Alphabets | 24 | 5 | Yes | Depth | 70 |
| Ours (re-training) (50/25/25)% | Alph. & Digit | 31 | 5 | No | Depth | 99.99 |
| Ours (re-training) (3/1/1) | Alph. & Digit | 31 | 5 | Yes | Depth | 75.18 |
| Ours (re-training) (4/1) | Alph. & Digit | 31 | 5 | Yes | Depth | 78.39 |
| Ours (fine-tuning) (3/1/1) | Alph. & Digit | 31 | 5 | Yes | Depth | 83.58 |
| Ours (fine-tuning) (4/1) | Alph. & Digit | 31 | 5 | Yes | Depth | 85.49 |

- Kang, B., Tripathi, S., Nguyen, T.Q.: Real-time sign language fingerspelling recognition using convolutional neural networks from depth map. In: Pattern Recognition (ACPR), 2015 3rd IAPR Asian Conference on, pp. 136–140. IEEE (2015)

# Existing Solutions

- Learning the American Sign Language (ASL) ***with CNNs***

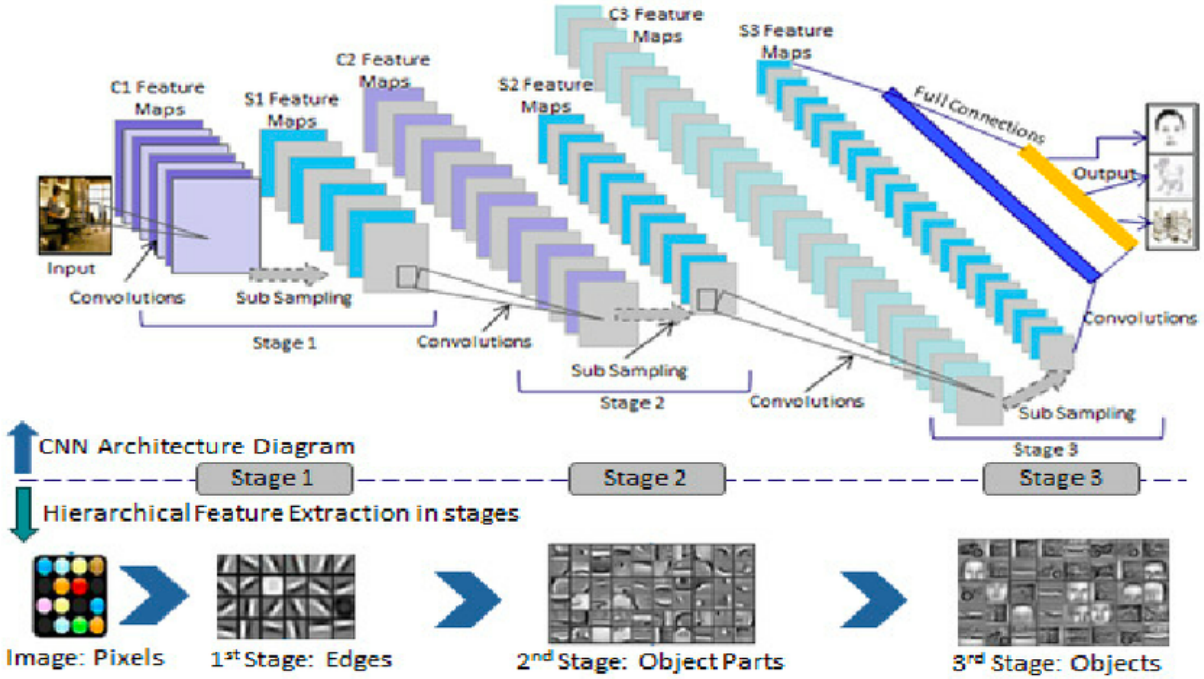

# The Proposal

- Learning the American Sign Language (ASL) ***with Auto-encoders***
  - ***Simpler than CNN***
  - ***More efficient than CNN***
  - ***Faster to train than CNN***
  - ***Similar performance to a CNN***
  - ***CNNs are not the panacea in pattern recognition on images or computer vision***

# The Proposal

- ***Lets talk about Auto-encoders***

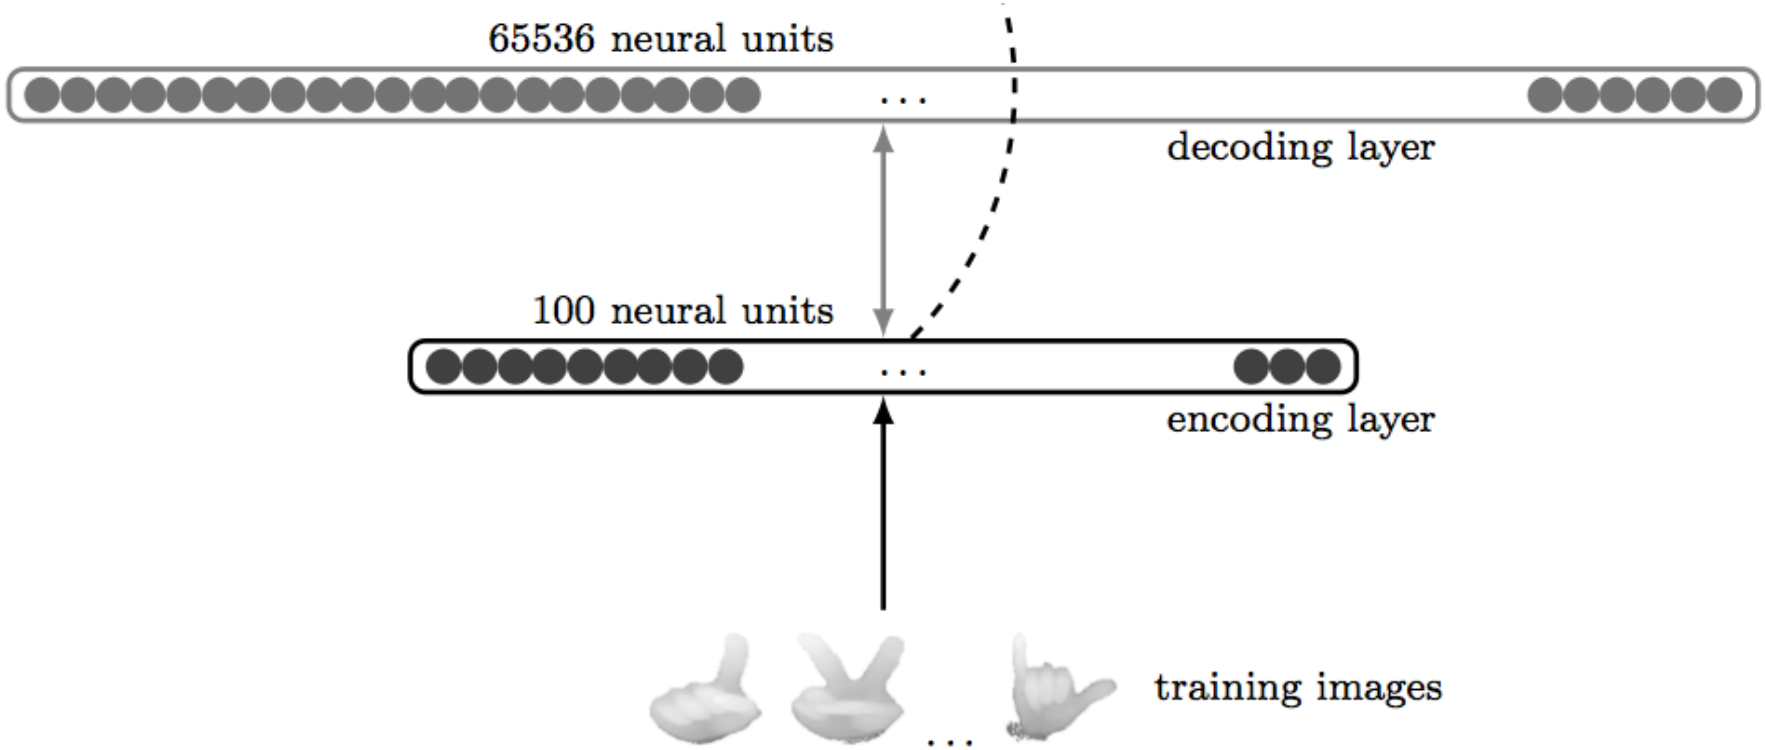


MARIST

# The Proposal

- ***Lets talk about Auto-encoders***

$$L = \frac{1}{N} \left\| \mathbf{x}_n - \hat{\mathbf{x}}_n \right\|_2^2 + \theta_w \frac{1}{2} \sum_{l=1}^{L} \left\| \mathbf{w}^l \right\|_2^2 + \theta_s \sum_{m=1}^{M} KL\left(\theta_\alpha \left\| \bar{\alpha}_m \right.\right)$$

MARIST

# The Proposal

- ***Lets talk about Auto-encoders***

$$L = \frac{1}{N} \left\| \mathbf{x}_n - \hat{\mathbf{x}}_n \right\|_2^2 + \theta_w \frac{1}{2} \sum_{l=1}^{L} \left\| \mathbf{w}^l \right\|_2^2 + \theta_s \sum_{m=1}^{M} KL\left(\theta_\alpha \left\| \bar{\alpha}_m \right.\right)$$

$$\sum_{m=1}^{M} KL\left(\theta_\alpha \left\| \bar{\alpha}_m \right.\right) = \sum_{m=1}^{M} \theta_\alpha \log\left(\frac{\theta_\alpha}{\bar{\alpha}_m}\right) + (1-\theta_\alpha) \log\left(\frac{1-\theta_\alpha}{1-\bar{\alpha}_m}\right)$$

$$\bar{\alpha}_m = \frac{1}{N} \sum_{n=1}^{N} \psi\left(\mathbf{w}_m^{(l)T} \mathbf{x}_n + b_m^{(l)}\right)$$

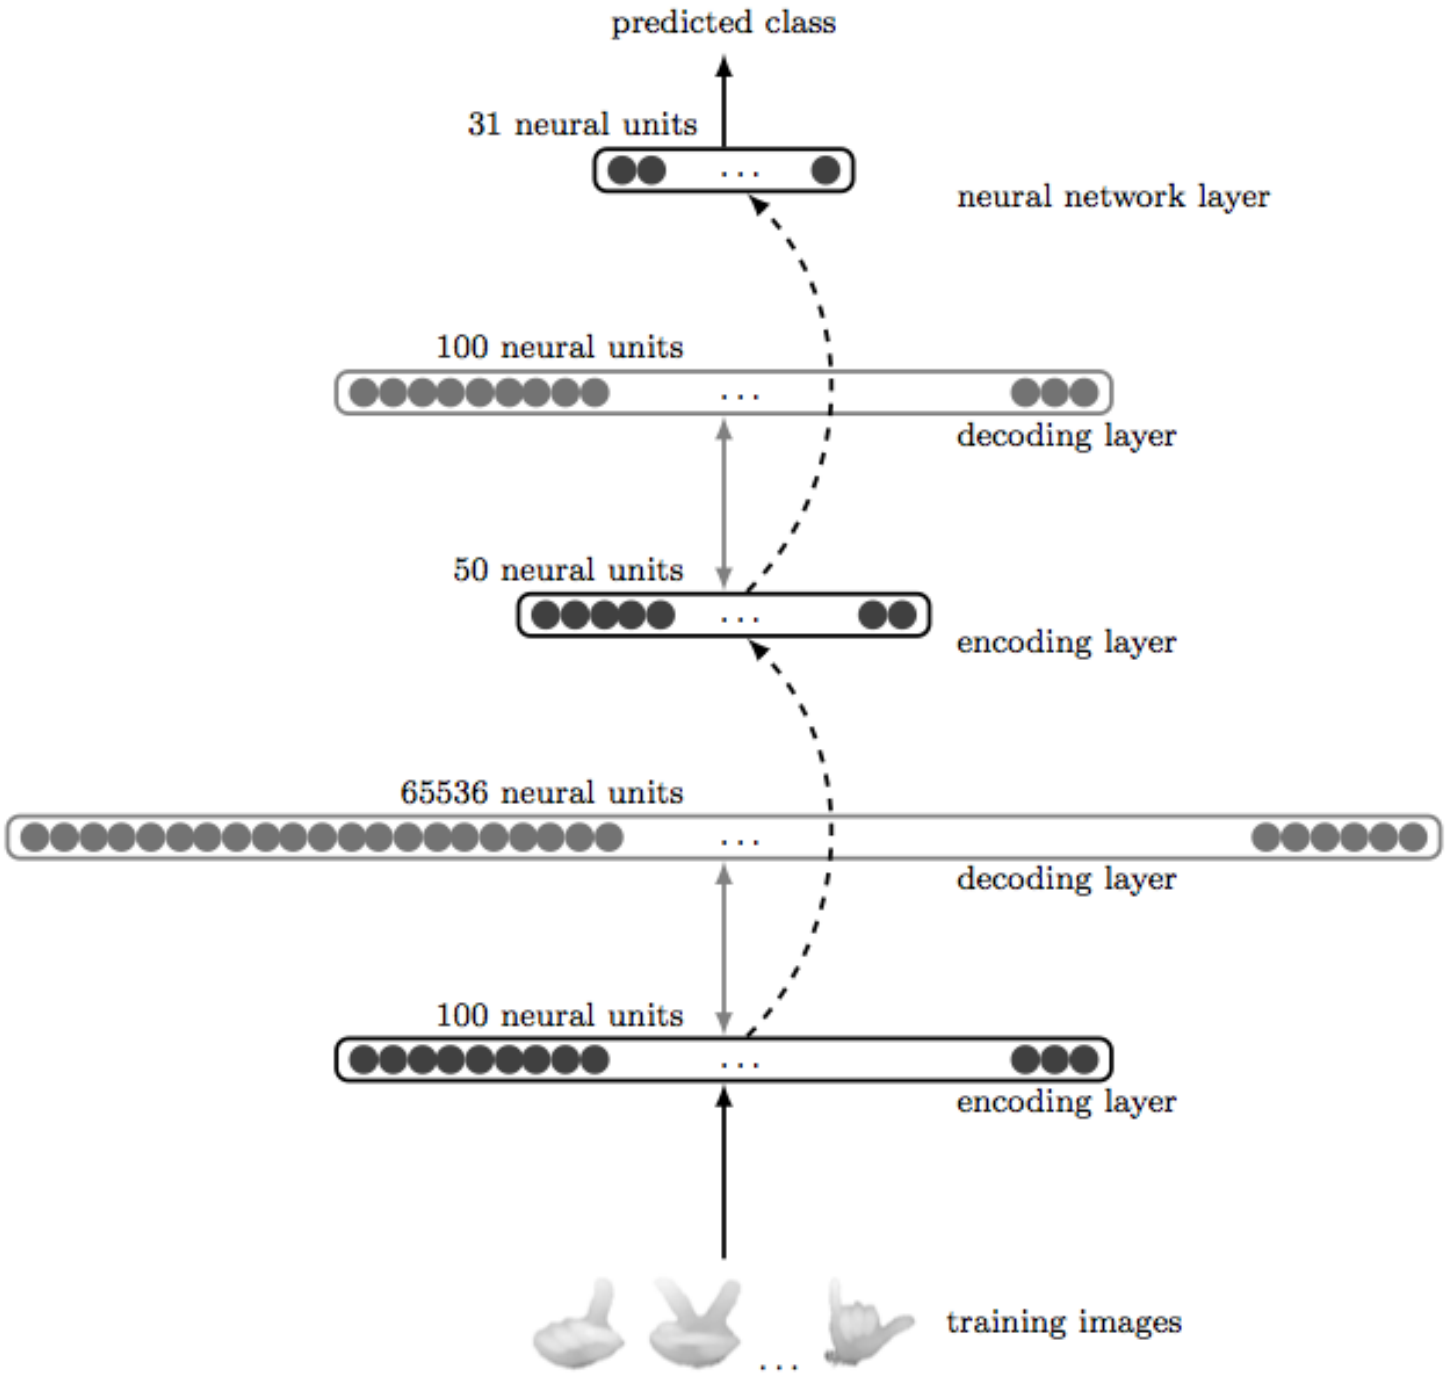
predicted class
31 neural units
neural network layer
100 neural units
decoding layer
50 neural units
encoding layer
65536 neural units
decoding layer
100 neural units
encoding layer
training images

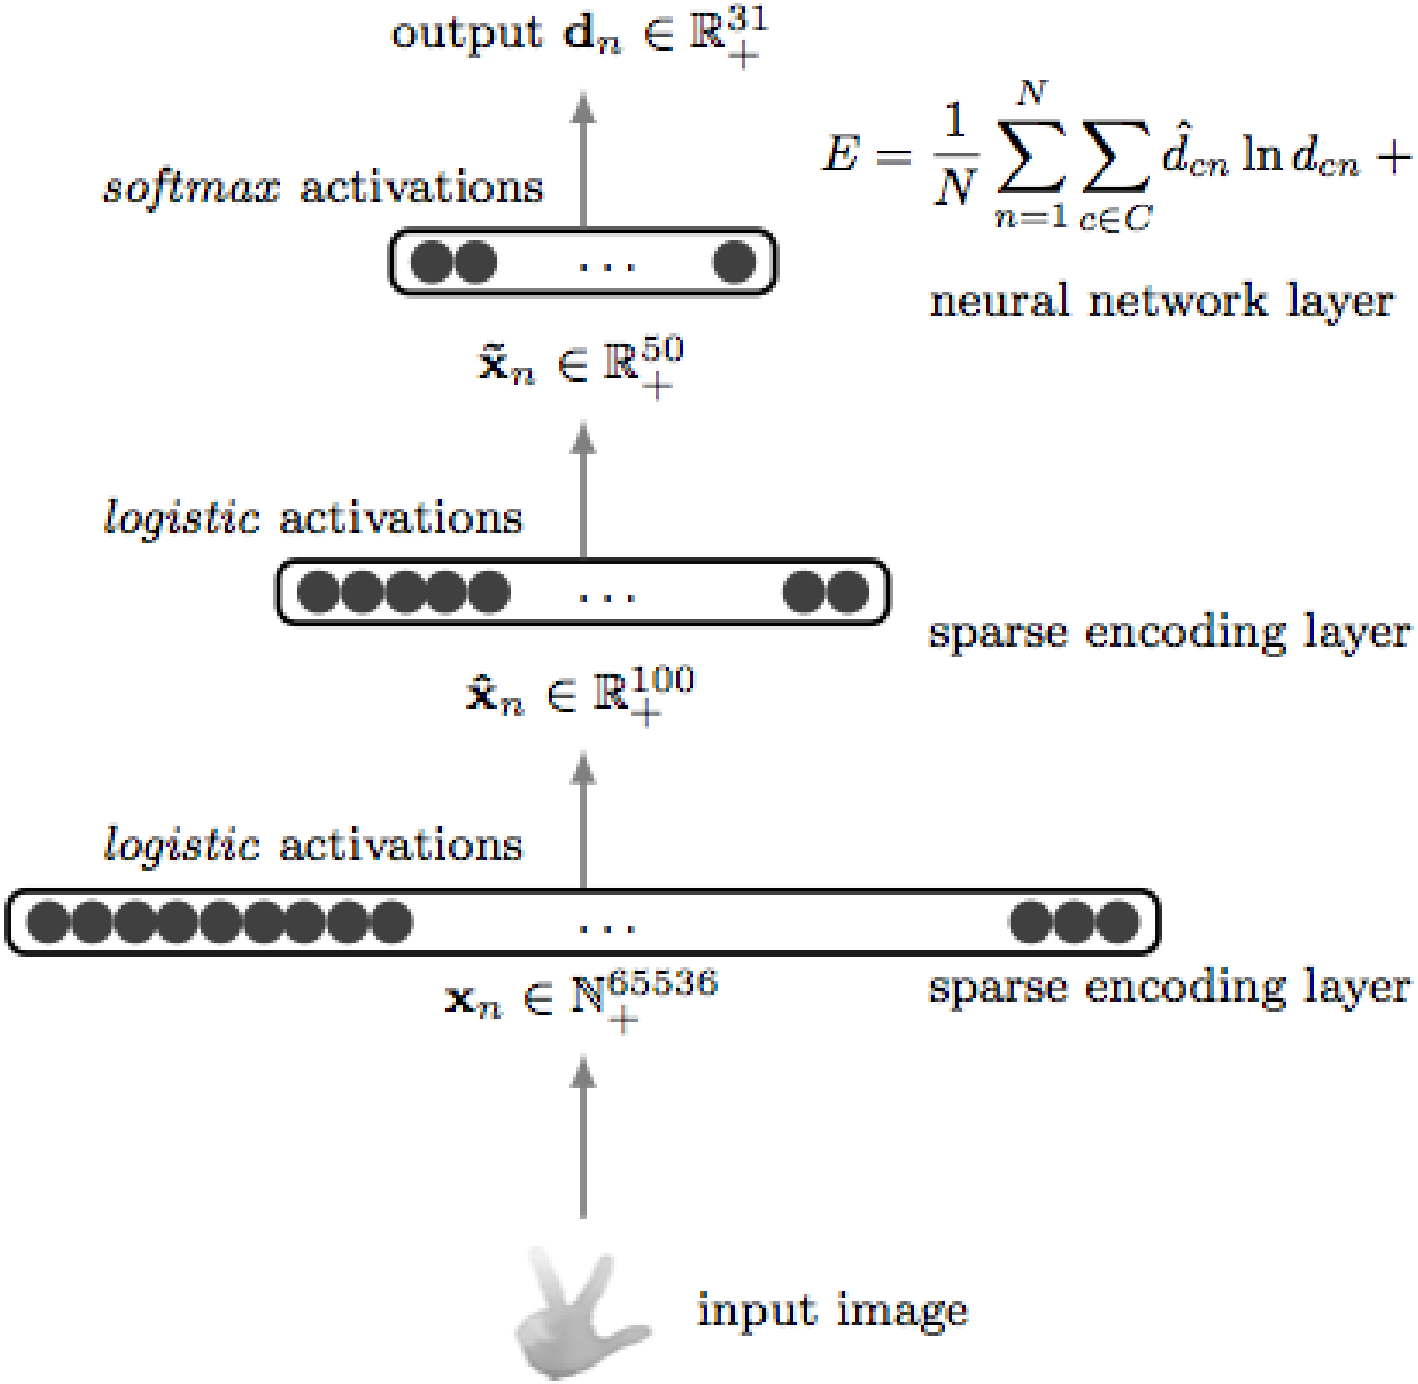


$$E = \frac{1}{N} \sum_{n=1}^{N} \sum_{c \in C} \hat{d}_{cn} \ln d_{cn} + (1 - \hat{d}_{cn}) \ln(1 - d_{cn})$$

neural network layer

# Results

- ***of the Auto-encoders***

| | S1 | S2 | S3 | S4 | S5 | Avg. |
|---|---|---|---|---|---|---|
| ACC | 0.9748 | 0.9923 | 0.9935 | 0.9929 | 0.9910 | 0.9889 |
| SPC | 0.9991 | 0.9997 | 0.9998 | 0.9998 | 0.9997 | 0.9996 |
| MAE | 0.1483 | 0.0640 | 0.0373 | 0.0347 | 0.0494 | 0.0667 |

# Results

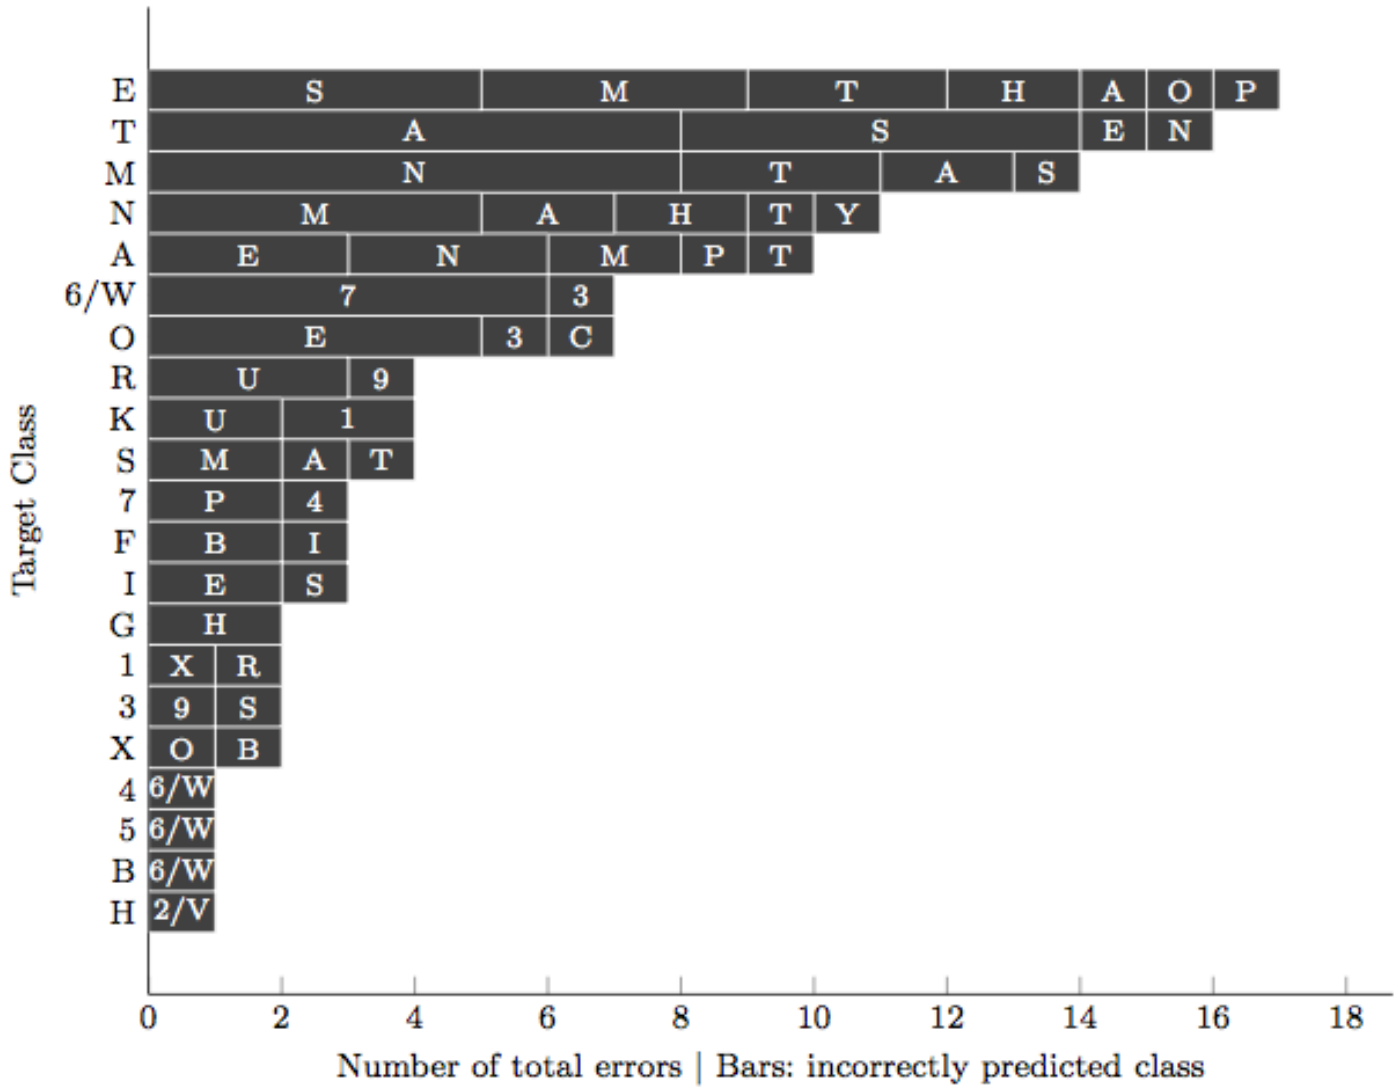
Target Class
E S M T H A O P
T A S E N
M N T A S
N M A H T Y
A E N M P T
6/W 7 3
O E 3 C
R U 9
K U 1
S M A T
7 P 4
F B I
I E S
G H
1 X R
3 9 S
X O B
4 6/W
5 6/W
B 6/W
H 2/V
0 2 4 6 8 10 12 14 16 18
Number of total errors | Bars: incorrectly predicted class

MARIST

# Results

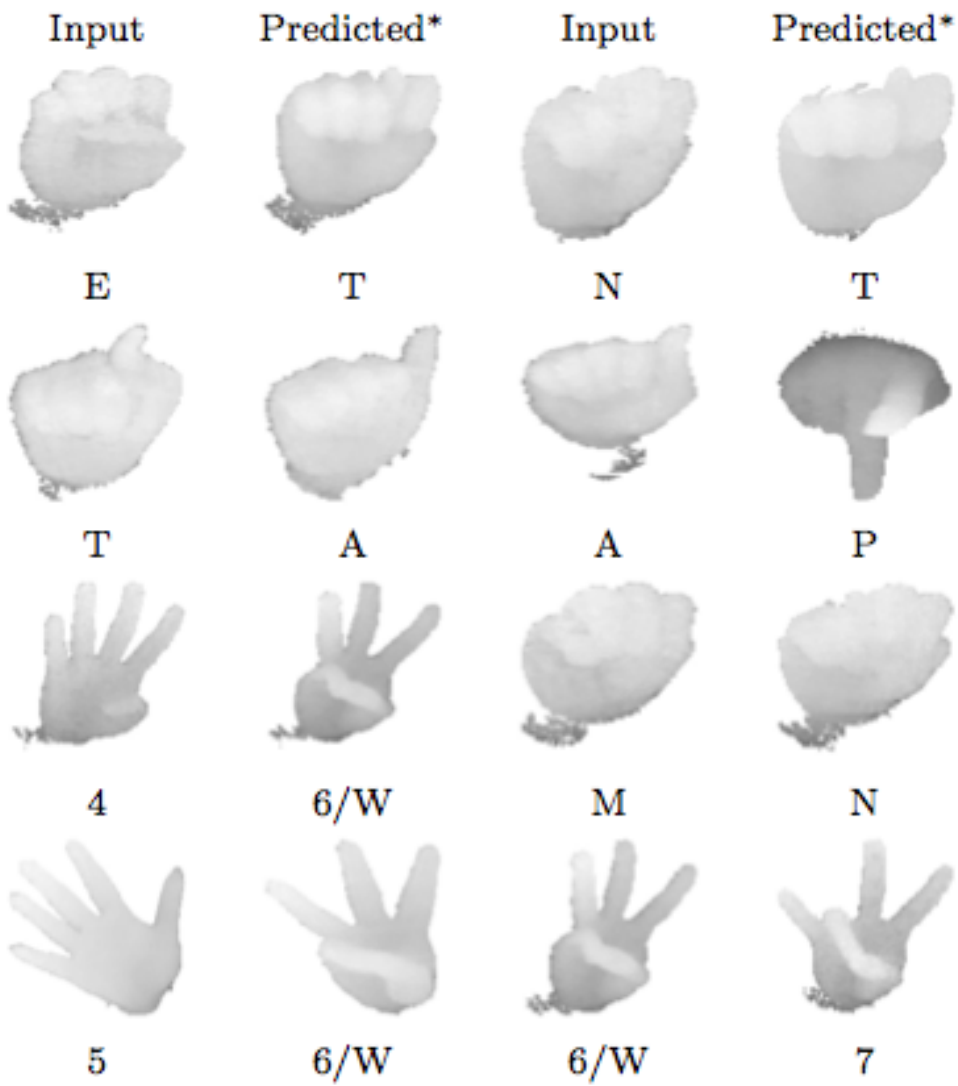

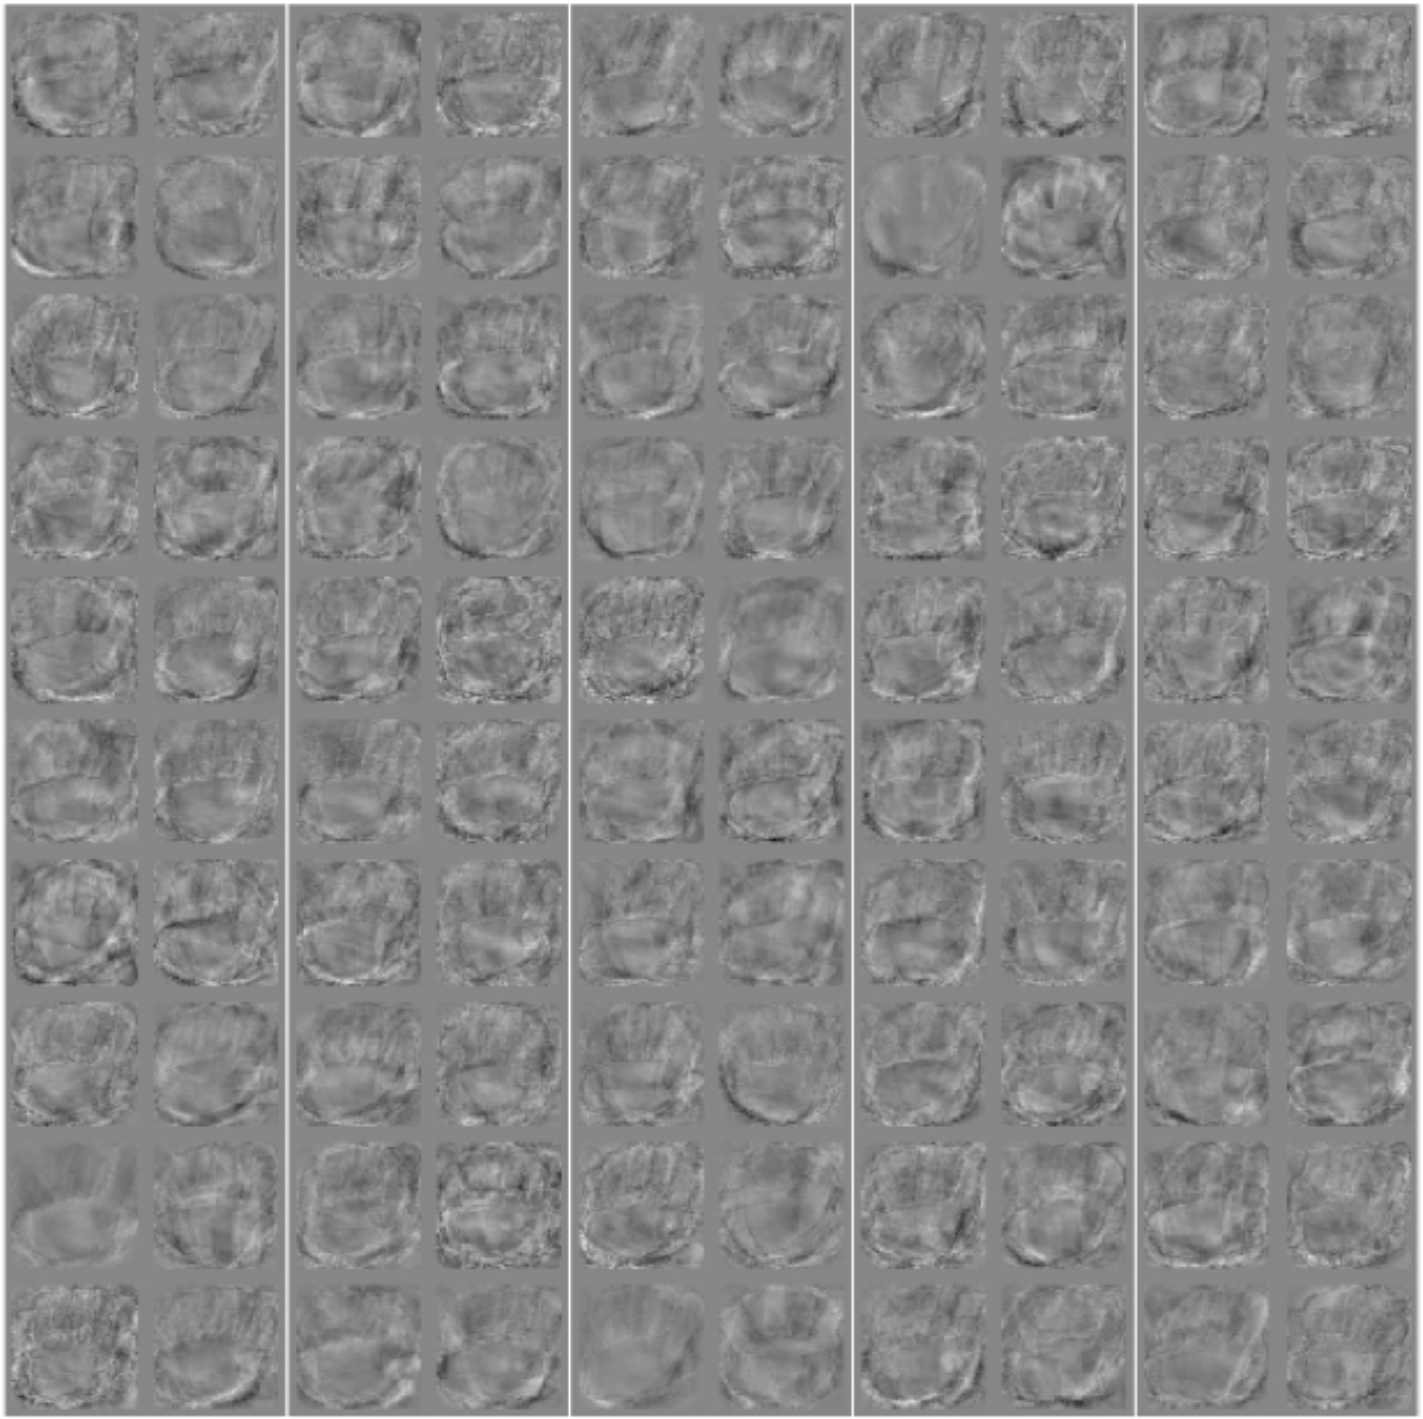

# Results

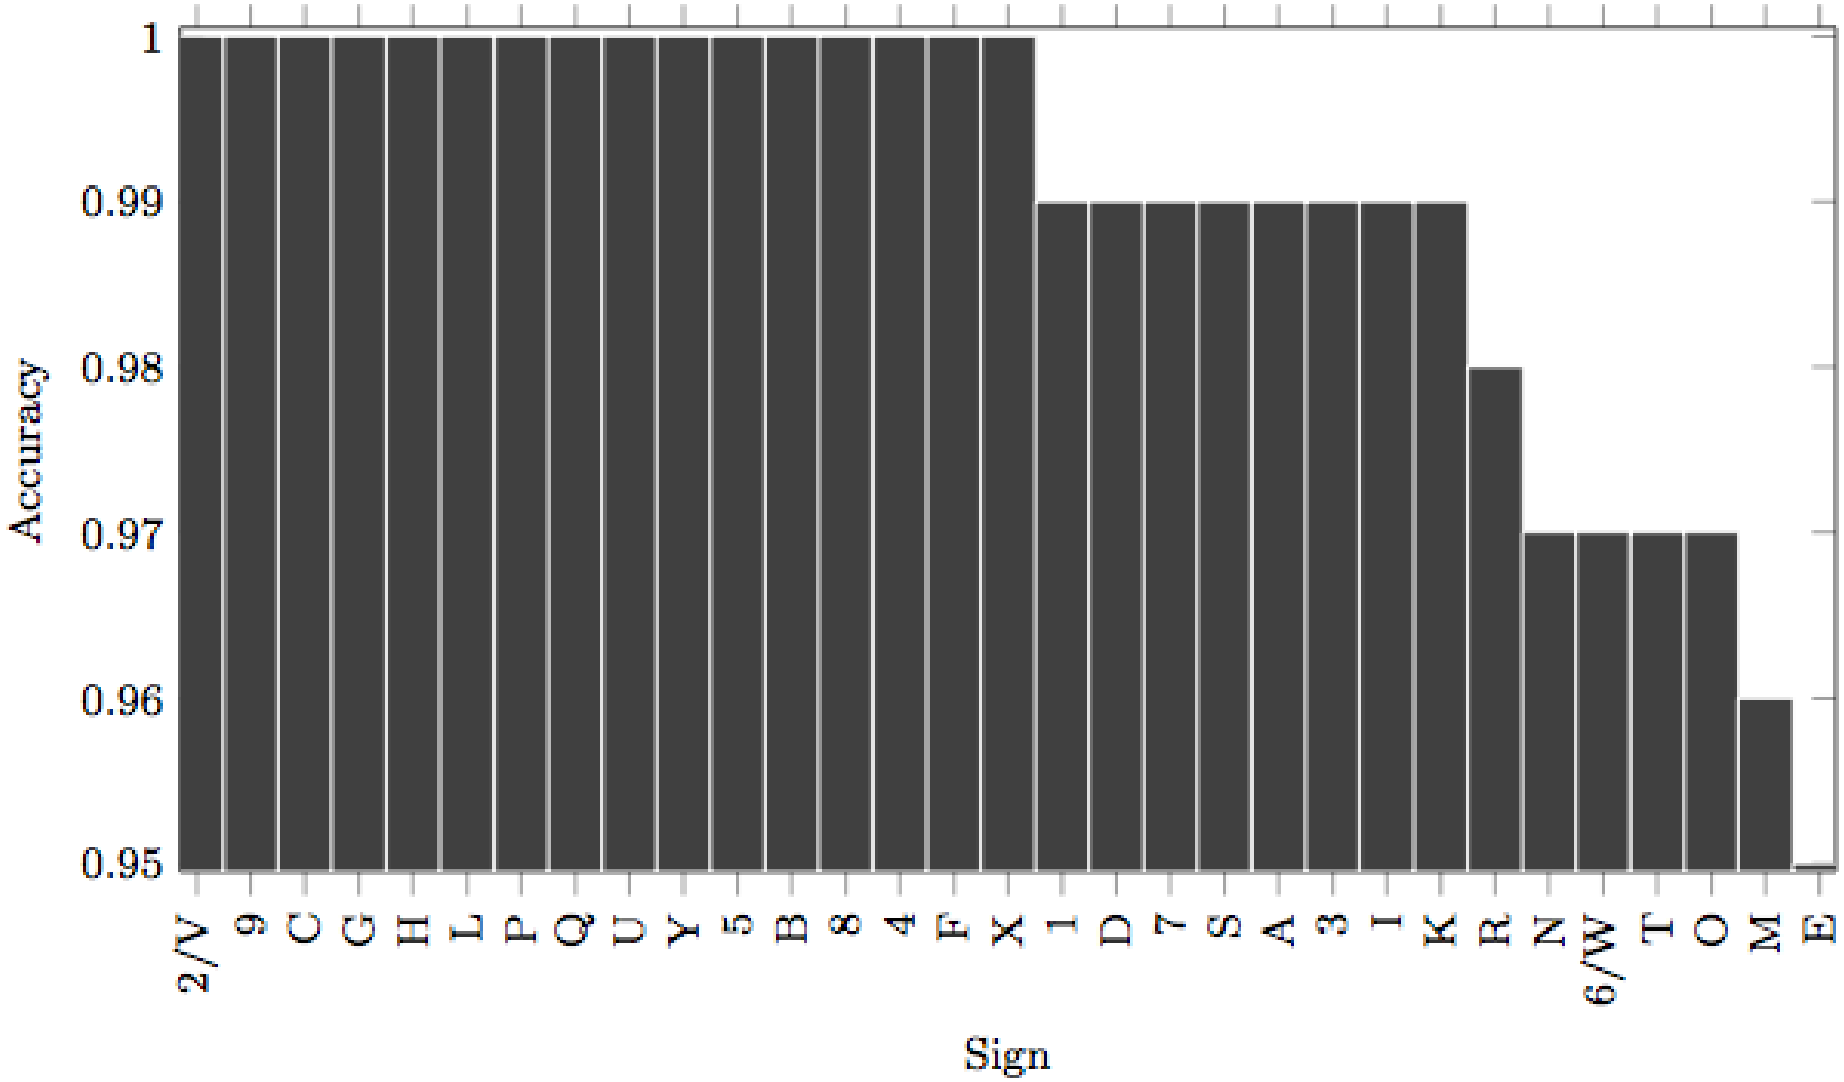
Accuracy
1
0.99
0.98
0.97
0.96
0.95
2/V
9
C
G
H
L
P
Q
U
Y
5
B
8
4
F
X
1
D
7
S
A
3
I
K
R
N
6/W
T
O
M
E
Sign

MARIST

# Chameleon Setup

- ***Resource type:*** bare metal/CHI@TACC
- ***Lease:*** GPU P100
- ***Image:*** CC-Ubuntu16.04-CUDA8
- ***Libraries:***
  - `cuDNN`
  - `libatlas-dev`

Chameleon

MARIST

# Libraries: **cuDNN**

cc@fox: ~ — ssh — 90×20 — ⌘1

```
Connection closed by 129.114.109.140
Pablos-MacBook-Air:.ssh rivas_perea$ ssh cc@129.114.109.140
Welcome to Ubuntu 16.04.2 LTS (GNU/Linux 4.4.0-72-generic x86_64)

 * Documentation:  https://help.ubuntu.com
 * Management:     https://landscape.canonical.com
 * Support:        https://ubuntu.com/advantage

  Get cloud support with Ubuntu Advantage Cloud Guest:
    http://www.ubuntu.com/business/services/cloud

0 packages can be updated.
0 updates are security updates.


Last login: Tue Sep 12 16:35:00 2017 from 204.210.149.122
cc@fox:~$ sudo cp cuda/include/cudnn.h /usr/local/cuda/include
cc@fox:~$ sudo cp cuda/lib64/libcudnn* /usr/local/cuda/lib64
cc@fox:~$ sudo chmod a+r /usr/local/cuda/include/cudnn.h /usr/local/cuda/lib64/libcudnn*
cc@fox:~$
```

MARIST

# Chameleon Setup

- ***Resource type:*** bare metal/CHI@TACC
- ***Lease:*** GPU P100
- ***Image:*** CC-Ubuntu16.04-CUDA8
- ***Packages:***
  - `gcc, gfortran`
  - `python-{numpy scipy matplotlib}`
  - `tensorflow`
  - `glances, nvidia-ml-py, screen`

# Packages: **tensorflow**

cc@fox: ~ — ssh — 80×20 — ⌘1

```
11.0)
  Downloading funcsigs-1.0.2-py2.py3-none-any.whl
Collecting pbr>=0.11 (from mock>=2.0.0->tensorflow==0.11.0)
  Downloading pbr-3.1.1-py2.py3-none-any.whl (99kB)
    100% |████████████████████████████████| 102kB 11.7MB/s
Requirement already up-to-date: setuptools in /usr/local/lib/python2.7/dist-pack
ages (from protobuf==3.0.0->tensorflow==0.11.0)
Installing collected packages: funcsigs, pbr, mock, protobuf, tensorflow
  Found existing installation: funcsigs 0.4
    Uninstalling funcsigs-0.4:
      Successfully uninstalled funcsigs-0.4
  Found existing installation: pbr 1.8.0
    Uninstalling pbr-1.8.0:
      Successfully uninstalled pbr-1.8.0
  Found existing installation: mock 1.3.0
    Uninstalling mock-1.3.0:
      Successfully uninstalled mock-1.3.0
Successfully installed funcsigs-1.0.2 mock-2.0.0 pbr-3.1.1 protobuf-3.0.0 tensor
flow-0.11.0
cc@fox:~$
```

MARIST

cc@fox: ~ — ssh — 107×20 — ⌘1

```
cc@fox:~/research/code-grid-search/summaries$ tensorboard --logdir=./
I tensorflow/stream_executor/dso_loader.cc:111] successfully opened CUDA library libcublas.so locally
I tensorflow/stream_executor/dso_loader.cc:111] successfully opened CUDA library libcudnn.so locally
I tensorflow/stream_executor/dso_loader.cc:111] successfully opened CUDA library libcufft.so locally
I tensorflow/stream_executor/dso_loader.cc:111] successfully opened CUDA library libcuda.so.1 locally
I tensorflow/stream_executor/dso_loader.cc:111] successfully opened CUDA library libcurand.so locally
Starting TensorBoard 29 on port 6006
(You can navigate to http://127.0.0.1:6006)
```

MARIST

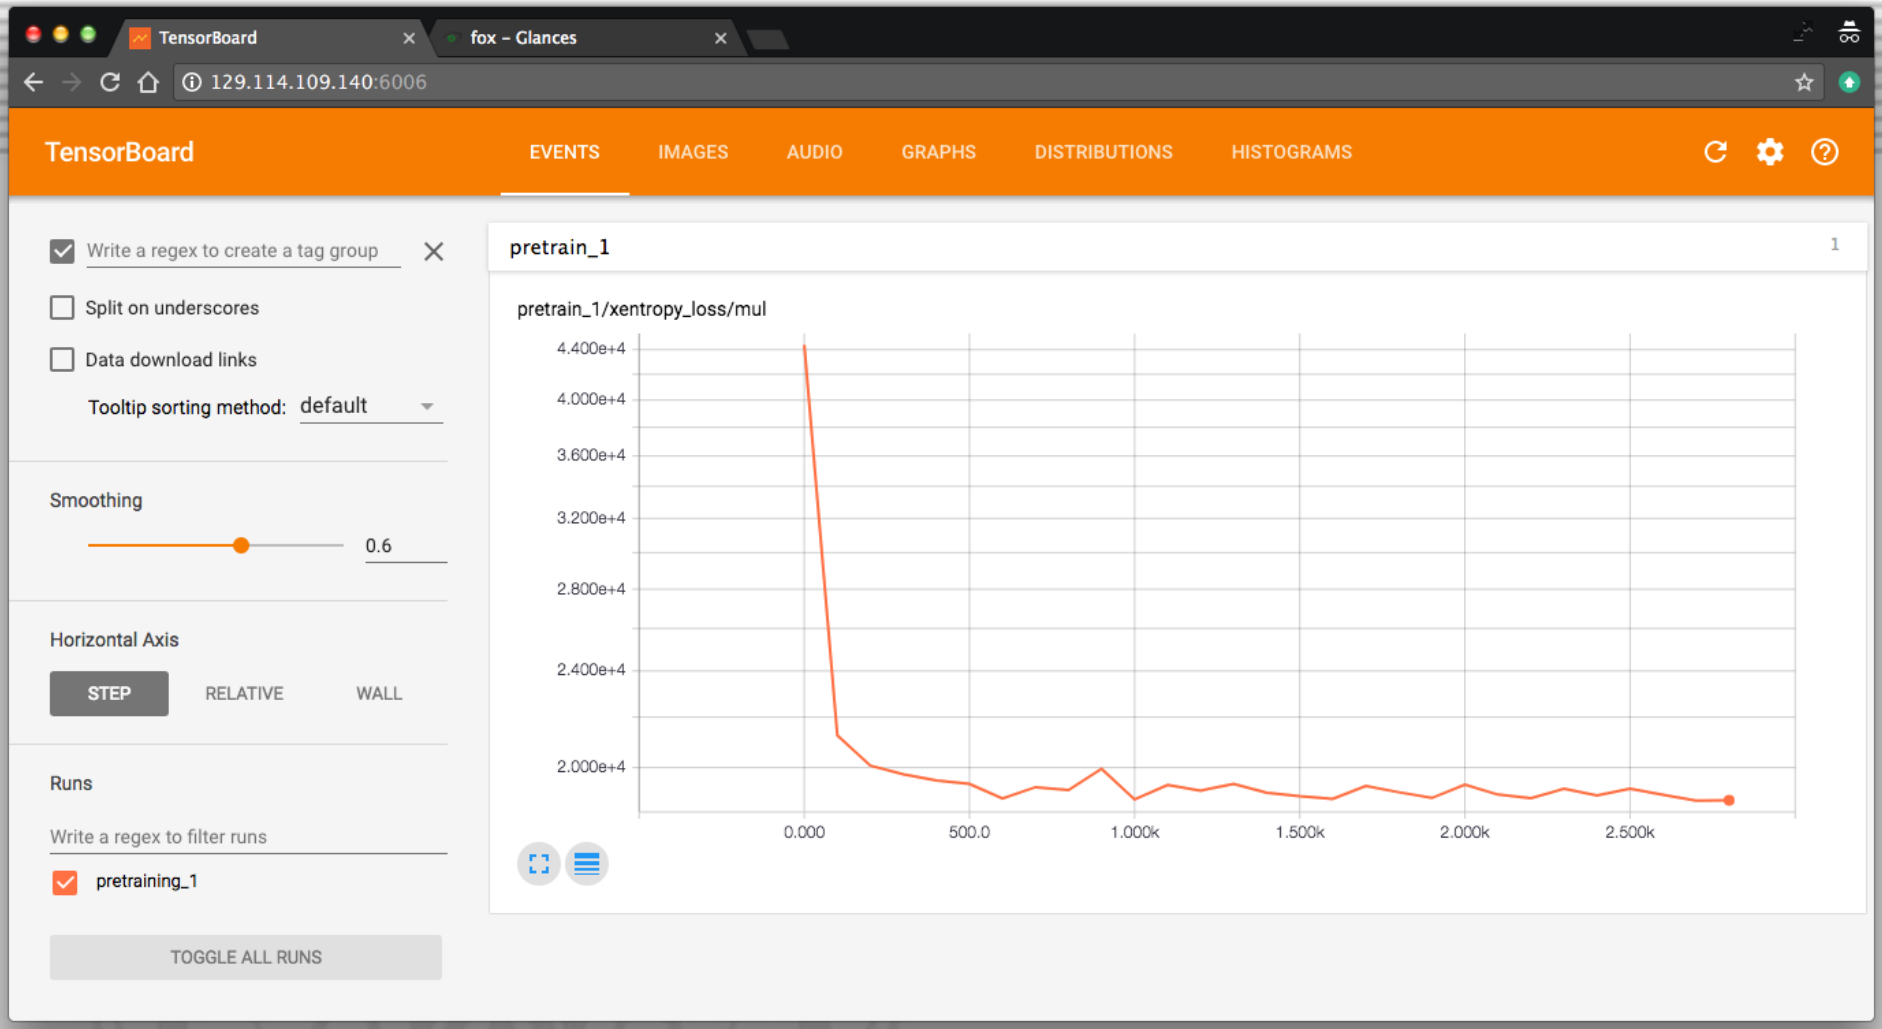
TensorBoard
fox – Glances
129.114.109.140:6006
TensorBoard
EVENTS
IMAGES
AUDIO
GRAPHS
DISTRIBUTIONS
HISTOGRAMS
Write a regex to create a tag group
Split on underscores
Data download links
Tooltip sorting method: default
Smoothing
0.6
Horizontal Axis
STEP
RELATIVE
WALL
Runs
Write a regex to filter runs
pretraining_1
TOGGLE ALL RUNS
pretrain_1
1
pretrain_1/xentropy_loss/mul
4.400e+4
4.000e+4
3.600e+4
3.200e+4
2.800e+4
2.400e+4
2.000e+4
0.000
500.0
1.000k
1.500k
2.000k
2.500k

MARIST

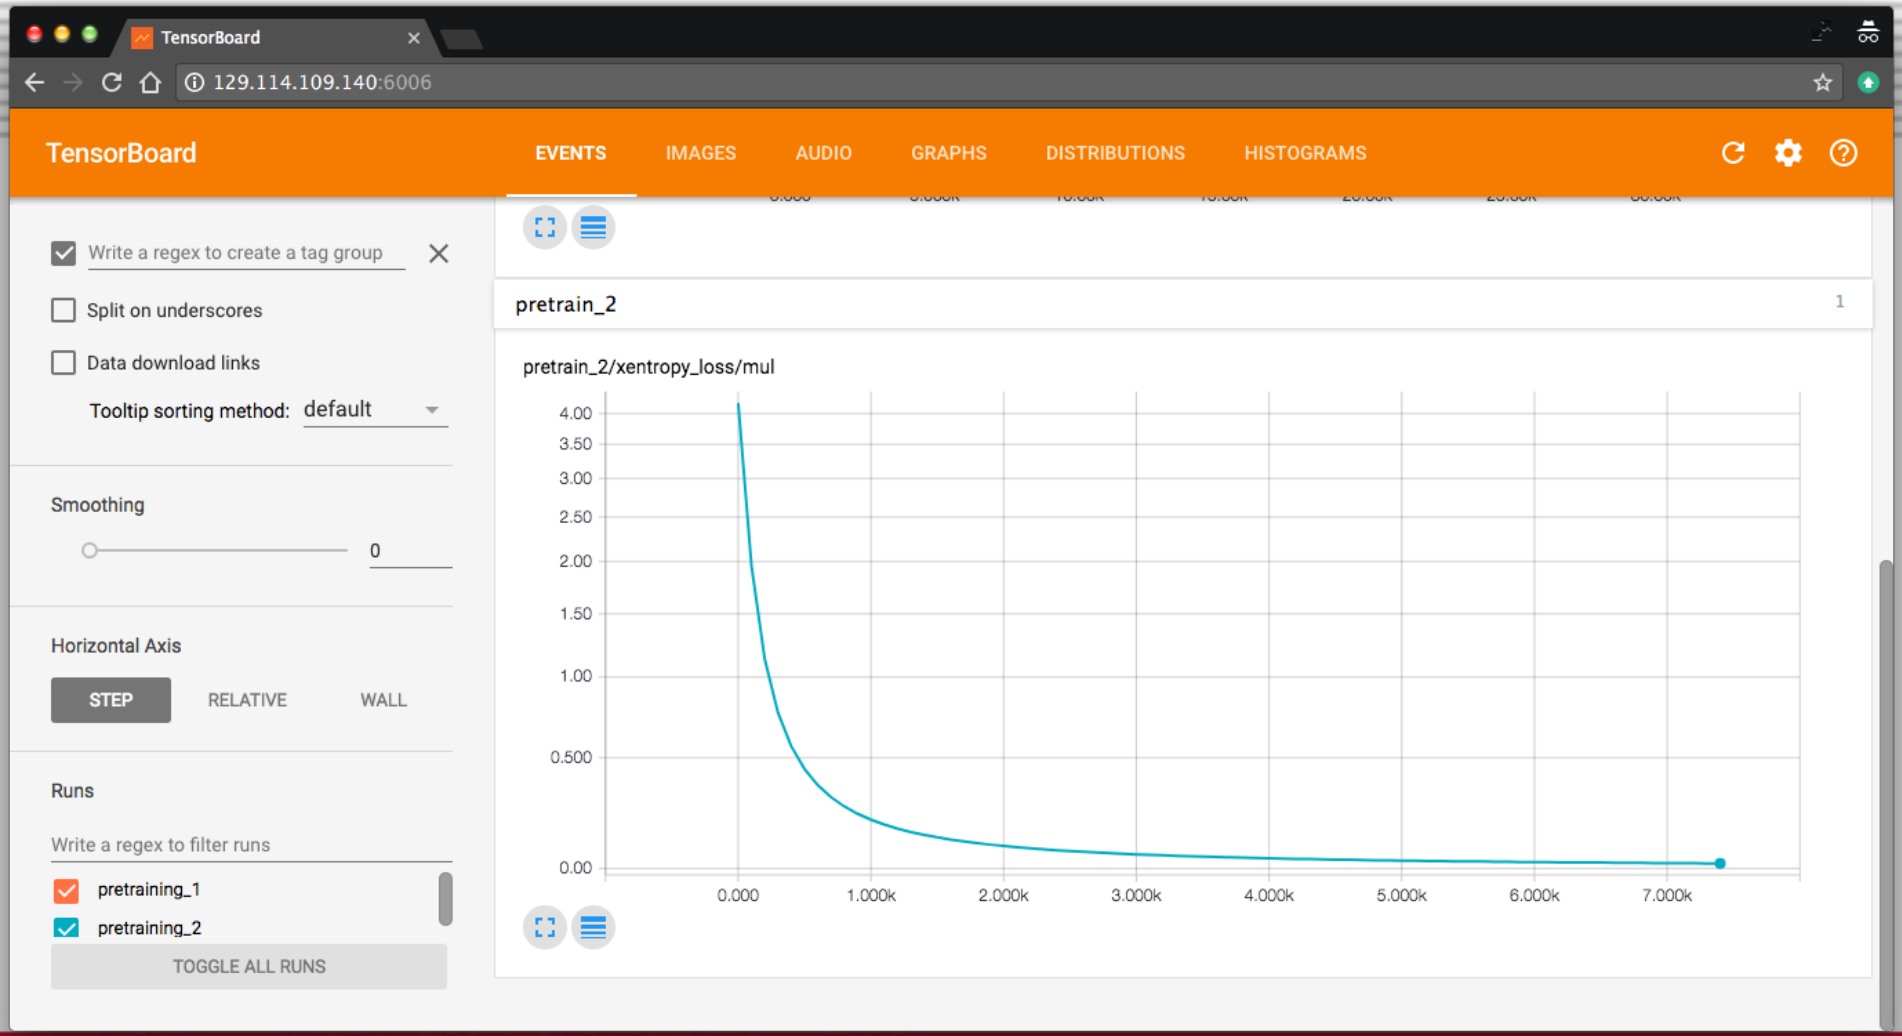
TensorBoard
129.114.109.140:6006
TensorBoard
EVENTS
IMAGES
AUDIO
GRAPHS
DISTRIBUTIONS
HISTOGRAMS
Write a regex to create a tag group
Split on underscores
Data download links
Tooltip sorting method: default
Smoothing
0
Horizontal Axis
STEP
RELATIVE
WALL
Runs
Write a regex to filter runs
pretraining_1
pretraining_2
TOGGLE ALL RUNS
pretrain_2
1
pretrain_2/xentropy_loss/mul
4.00
3.50
3.00
2.50
2.00
1.50
1.00
0.500
0.00
0.000
1.000k
2.000k
3.000k
4.000k
5.000k
6.000k
7.000k

MARIST

# Packages: `glances`

```
cc@fox: ~ — ssh — 80×20 — ⌘1

fox - IP 10.40.0.33/24 Pub 129.114.109.140                    Uptime: 2:38:42
AWS EC2

1.40/2.30GHz    CPU -     2.3% GPU Tesla M40     MEM -     3.5% SWAP -    0.0%
CPU  [  1.9%]   user:     1.4% proc:        36% total:    126G total:       0
MEM  [  3.6%]   system:   0.7% mem:         96% used:    4.45G used:        0
SWAP [  0.0%]   idle:    97.9%                   free:    121G free:        0

NETWORK      Rx/s   Tx/s   TASKS 458 (581 thr), 2 run, 456 slp, 0 oth
eno1          4Kb   784b
lo           184b   184b     CPU%  MEM%    PID USER        NI S Command
                             90.9   2.2  27647 cc           0 R python -u autoenc
DISK I/O      R/s    W/s      5.8   0.0  28738 cc           0 R /usr/bin/python /
sda1            0   912K      0.2   0.0      8 root         0 S rcu_sched
sr0             0      0      0.0   0.0   2982 root       -20 S kworker/20:1H
                              0.0   0.0    226 root         0 S kworker/43:0
FILE SYS     Used  Total      0.0   0.0    132 root       -20 S kworker/24:0H
/ (sda1)   7.92G   218G      0.0   0.0   1171 root       -20 S bioset

2017-09-12 19:19:26        No warning or critical alert detected
```

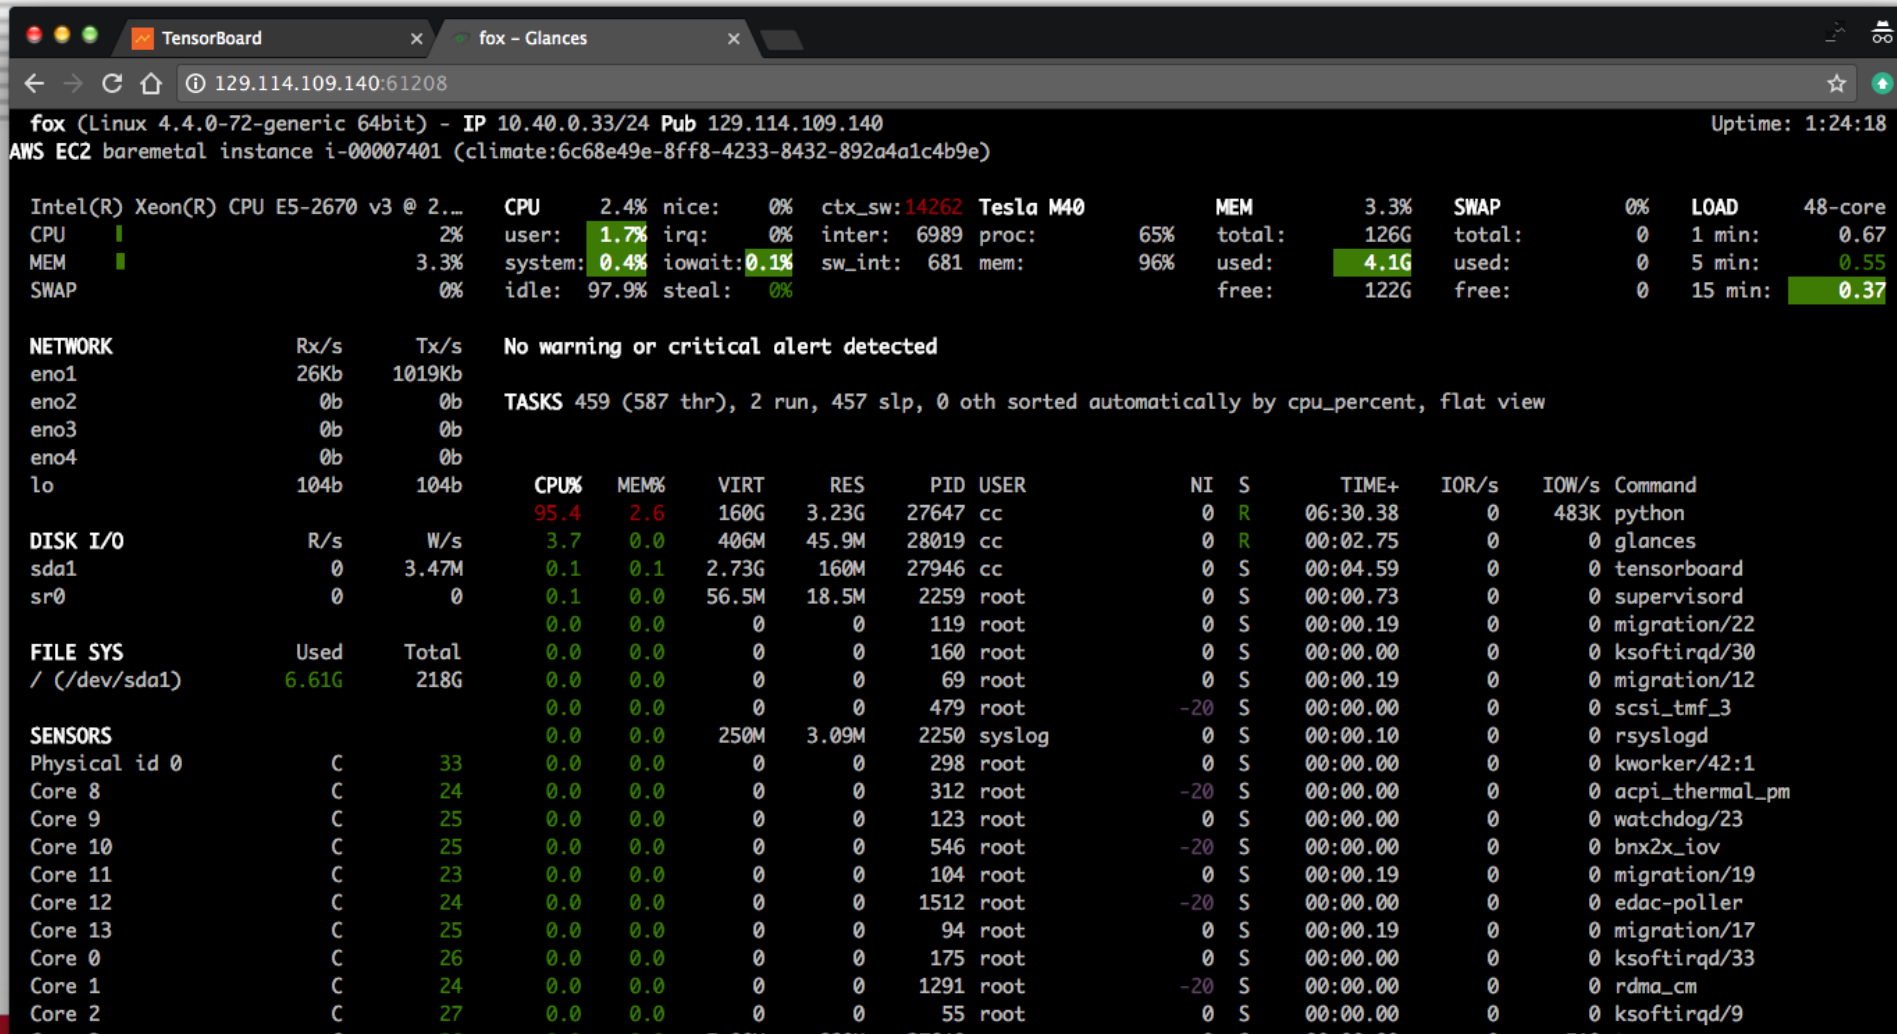
TensorBoard
fox – Glances
129.114.109.140:61208
fox (Linux 4.4.0-72-generic 64bit) - IP 10.40.0.33/24 Pub 129.114.109.140
Uptime: 1:24:18
AWS EC2 baremetal instance i-00007401 (climate:6c68e49e-8ff8-4233-8432-892a4a1c4b9e)
Intel(R) Xeon(R) CPU E5-2670 v3 @ 2.…
CPU 2%
MEM 3.3%
SWAP 0%
CPU 2.4% nice: 0% ctx_sw:14262
user: 1.7% irq: 0% inter: 6989
system: 0.4% iowait:0.1% sw_int: 681
idle: 97.9% steal: 0%
Tesla M40
proc: 65%
mem: 96%
MEM 3.3%
total: 126G
used: 4.1G
free: 122G
SWAP 0%
total: 0
used: 0
free: 0
LOAD 48-core
1 min: 0.67
5 min: 0.55
15 min: 0.37
NETWORK Rx/s Tx/s
eno1 26Kb 1019Kb
eno2 0b 0b
eno3 0b 0b
eno4 0b 0b
lo 104b 104b
DISK I/O R/s W/s
sda1 0 3.47M
sr0 0 0
FILE SYS Used Total
/ (/dev/sda1) 6.61G 218G
SENSORS
Physical id 0 C 33
Core 8 C 24
Core 9 C 25
Core 10 C 25
Core 11 C 23
Core 12 C 24
Core 13 C 25
Core 0 C 26
Core 1 C 24
Core 2 C 27
No warning or critical alert detected
TASKS 459 (587 thr), 2 run, 457 slp, 0 oth sorted automatically by cpu_percent, flat view
CPU% MEM% VIRT RES PID USER NI S TIME+ IOR/s IOW/s Command
95.4 2.6 160G 3.23G 27647 cc 0 R 06:30.38 0 483K python
3.7 0.0 406M 45.9M 28019 cc 0 R 00:02.75 0 0 glances
0.1 0.1 2.73G 160M 27946 cc 0 S 00:04.59 0 0 tensorboard
0.1 0.0 56.5M 18.5M 2259 root 0 S 00:00.73 0 0 supervisord
0.0 0.0 0 0 119 root 0 S 00:00.19 0 0 migration/22
0.0 0.0 0 0 160 root 0 S 00:00.00 0 0 ksoftirqd/30
0.0 0.0 0 0 69 root 0 S 00:00.19 0 0 migration/12
0.0 0.0 0 0 479 root -20 S 00:00.00 0 0 scsi_tmf_3
0.0 0.0 250M 3.09M 2250 syslog 0 S 00:00.10 0 0 rsyslogd
0.0 0.0 0 0 298 root 0 S 00:00.00 0 0 kworker/42:1
0.0 0.0 0 0 312 root -20 S 00:00.00 0 0 acpi_thermal_pm
0.0 0.0 0 0 123 root 0 S 00:00.00 0 0 watchdog/23
0.0 0.0 0 0 546 root -20 S 00:00.00 0 0 bnx2x_iov
0.0 0.0 0 0 104 root 0 S 00:00.19 0 0 migration/19
0.0 0.0 0 0 1512 root -20 S 00:00.00 0 0 edac-poller
0.0 0.0 0 0 94 root 0 S 00:00.19 0 0 migration/17
0.0 0.0 0 0 175 root 0 S 00:00.00 0 0 ksoftirqd/33
0.0 0.0 0 0 1291 root -20 S 00:00.00 0 0 rdma_cm
0.0 0.0 0 0 55 root 0 S 00:00.00 0 0 ksoftirqd/9

MARIST

cc@fox: ~ — ssh — 80×20 — ⌘1

```
fox - IP 10.40.0.33/24 Pub 129.114.109.140                    Uptime: 3:21:41
AWS EC2

2.60/2.30GHz    CPU /    61.2% GPU Tesla M40     MEM -    5.0% SWAP -    0.0%
CPU  [ 74.9%]   user:    54.3% proc:        39% total:   126G total:       0
MEM  [  4.9%]   system:  16.3% mem:         96% used:   6.24G used:        0
SWAP [  0.0%]   idle:    29.4%                   free:    119G free:        0

NETWORK     Rx/s   Tx/s   TASKS 462 (680 thr), 2 run, 460 slp, 0 oth
eno1        240b    1Kb
lo          176b   176b     CPU%  MEM%    PID USER        NI S Command
                           3069.0   1.7  31154 cc           0 S python -u autoenc
DISK I/O     R/s    W/s      98.5   2.1  30914 cc           0 S python -u autoenc
sda1           0     7K       8.1   0.0  31306 cc           0 R /usr/bin/python /
sr0            0      0       0.2   0.0   3071 root         0 S kworker/u288:0
                              0.2   0.0    119 root         0 S migration/22
FILE SYS    Used  Total       0.0   0.0   2982 root       -20 S kworker/20:1H
/ (sda1)   6.58G   218G       0.0   0.0    226 root         0 S kworker/43:0

2017-09-12 20:02:25        No warning or critical alert detected
```

MARIST

cc@fox: ~ — ssh — 150×35 — ⌘1

```
fox (Ubuntu 16.04 64bit / Linux 4.4.0-72-generic) - IP 10.40.0.33/24 Pub 129.114.109.140                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                
2.61/2.30GHz  CPU \  30.1%  nice:  0.0%  ctx_sw:  51K  GPU Tesla M40  MEM -  5.1%  active:  5.56G  SWAP -  0.0%  LOAD  48-core
CPU [ 68.1%]  user:  43.9%  irq:  0.0%  inter:  18608  proc:  62%  total:  126G  inactive:  1.90G  total:  0  1 min:  31.31
MEM [ 5.0%]  system:  12.6%  iowait:  0.0%  sw_int:  11551  mem:  96%  used:  6.40G  buffers:  97.5M  used:  0  5 min:  12.45
SWAP [ 0.0%]  idle:  43.5%  steal:  0.0%  free:  119G  cached:  2.74G  free:  0  15 min:  4.69

NETWORK  Rx/s  Tx/s  TASKS 463 (682 thr), 1 run, 462 slp, 0 oth sorted automatically by cpu_percent, flat view
eno1  6Kb  6Kb
lo  176b  176b  CPU%  MEM%  VIRT  RES  PID  USER  NI  S  TIME+  R/s  W/s  Command
3246.9  1.7  10.8G  2.10G  31154  cc  0  S  1h04:08  0  0  python -u autoencoder.py
DISK I/O  R/s  W/s  98.5  2.1  160G  2.67G  30914  cc  0  R  5:41.78  0  869K  python -u autoencoder.py
sda1  0  0  8.0  0.0  411M  42.0M  31306  cc  0  R  0:11.20  0  0  /usr/bin/python /usr/local/bin/glances
sr0  0  0  0.2  0.0  19.3M  2.47M  2453  root  0  S  0:04.26  0  0  /usr/sbin/irqbalance --pid=/var/run/irqbalance.pid
0.2  0.0  56.5M  18.5M  2259  root  0  S  0:01.63  0  0  /usr/bin/python /usr/bin/supervisord -n -c /etc/super
FILE SYS  Used  Total  0.2  0.0  0  0  3071  root  0  S  0:01.88  0  0  kworker/u288:0
/ (sda1)  6.59G  218G  0.0  0.0  0  0  2982  root  -20  S  0:00.00  0  0  kworker/20:1H
0.0  0.0  0  0  226  root  0  S  0:00.00  0  0  kworker/43:0
SENSORS  0.0  0.0  0  0  132  root  -20  S  0:00.00  0  0  kworker/24:0H
Physical id 0  45C  0.0  0.0  0  0  1171  root  -20  S  0:00.00  0  0  bioset
Core 8  39C  0.0  0.0  0  0  130  root  0  S  0:00.40  0  0  ksoftirqd/24
Core 9  40C  0.0  0.0  0  0  29409  root  0  S  0:00.16  0  0  kworker/u289:1
Core 10  40C  0.0  0.0  0  0  151  root  0  S  0:00.00  0  0  kworker/28:0
Core 11  39C  0.0  0.0  0  0  2499  root  -20  S  0:00.00  0  0  kworker/14:1H
Core 12  37C  0.0  0.0  0  0  2966  root  -20  S  0:00.00  0  0  kworker/11:1H
Core 13  38C  0.0  0.0  0  0  208  root  0  S  0:00.20  0  0  watchdog/40
Core 0  40C  0.0  0.0  0  0  485  root  0  S  0:00.00  0  0  scsi_eh_5
Core 1  38C  0.0  0.0  0  0  52  root  -20  S  0:00.00  0  0  kworker/8:0H
Core 2  41C  0.0  0.0  0  0  50  root  0  S  0:00.00  0  0  ksoftirqd/8
Core 3  39C  0.0  0.0  0  0  496  root  -20  S  0:00.00  0  0  scsi_tmf_10
Core 4  42C  0.0  0.0  0  0  493  root  0  S  0:00.00  0  0  scsi_eh_9
Core 5  40C  0.0  0.0  0  0  89  root  0  S  0:00.20  0  0  migration/16
Physical id 1  36C
2017-09-12 20:03:05  No warning or critical alert detected
```

MARIST

# Conclusions

- Learning the American Sign Language (ASL) ***with Auto-encoders***
  - ***Simpler than CNN***
  - ***More efficient than CNN***
  - ***Faster to train than CNN***
  - ***Similar performance to a CNN***
  - ***CNNs are not the panacea in pattern recognition on images or computer vision*** *(no free lunch theorem)*

Chameleon

MARIST

# Unsupervised Deep Learning with Stacked Autoencoders on Chameleon

**Interested in code? Check Deep's repo:**

https://github.com/DeepDand/research

MARIST