



# Neural Networks and Deep Learning: Residual Architectures

**Nicolas Thome**

Conservatoire National des Arts et Métiers (Cnam)  
Département Informatique


# Deep Learning since 2012

ILSVRC since 2012: larger & larger networks



- ImageNet'14: VGG & GoogLeNet ~ 20 layers

# Issue: Training Deeper Networks




► **BUT**: deeper nets have worse performances

- Ex: plain nets: stacking  $3 \times 3$  conv
- Ex: deeper VGG  $\Rightarrow$  VGG56  $<$  VGG20


# Issue: Training Deeper Networks

- ⚠️ ⚠️ Not a generalization issue, training error also higher!
- General phenomenon, observed in many datasets


ImageNet-1000



CIFAR-10




# Issue: Training Deeper Networks



- ▶ 18 layers (left) vs 34 layers (right)
- ▶ **Deeper counterpart: richer solution space**  
⇒ should not have higher training error
- ▶ Construction: copy from a shallower model
  - ▶ Extra layers: set as identity

# Identity Mapping with Plain Blocks


- ▶ **BUT:** optimization challenge  $\Rightarrow$  solvers cannot find this solution when going deeper...
- ▶  $H(x)$  desired mapping to fit with a 2 weight layers
- ▶ **Multiple non-linear layers  $\Rightarrow$  identity mapping difficult**




Plain Block

# Identity Mapping with Residual Blocks

- ▶  $H(x)$  desired mapping, 2 layers fit  $F(x) \Rightarrow H(x) = F(x) + x$
- ▶ If identity optimal, easy to set weights as 0
- ▶ If optimal mapping closer to identity, easier to find small fluctuations



ImageNet'15: Residual Networks (ResNet) [He et al., 2016]




- ▶ Ex for VGG-style: 3x3 conv
- ▶ **Simple design, just deeper**
  - ▶ Trained from scratch
  - ▶ Using Batch Normalization




# ResNet: Results

CIFAR-10 plain nets



CIFAR-10 ResNets




## CIFAR-10 results


- ▶ Deep ResNets trained without difficulty:  
Deeper Models  $\Rightarrow$  training & testing error lower

# ResNet: Results

ImageNet plain nets



ImageNet ResNets




## ImageNet

- ▶ Deep ResNets trained without difficulty:  
Deeper Models  $\Rightarrow$  training & testing error lower

# ResNet: Conclusion

- ▶ ResNet: training deeper models
- ▶ Performance below 4% error at ImageNet Large Scale Visual Recognition Competition'15!
- ▶ Other modern deep learning modules?  $\Rightarrow$  following!



# References I



He, K., Zhang, X., Ren, S., and Sun, J. (2016).  
Deep residual learning for image recognition.  
In *CVPR*.