



Introduction to Artificial Intelligence

COSC 4550 / COSC 5550

Professor Cheney
11/6/17

sample exam posted

2 weeks until project half-way check-in

paper readings

withdraw deadline this week!

many slides today borrowed
from Xavier Giro-i-Nieto

(sorry for lack of pixels...)

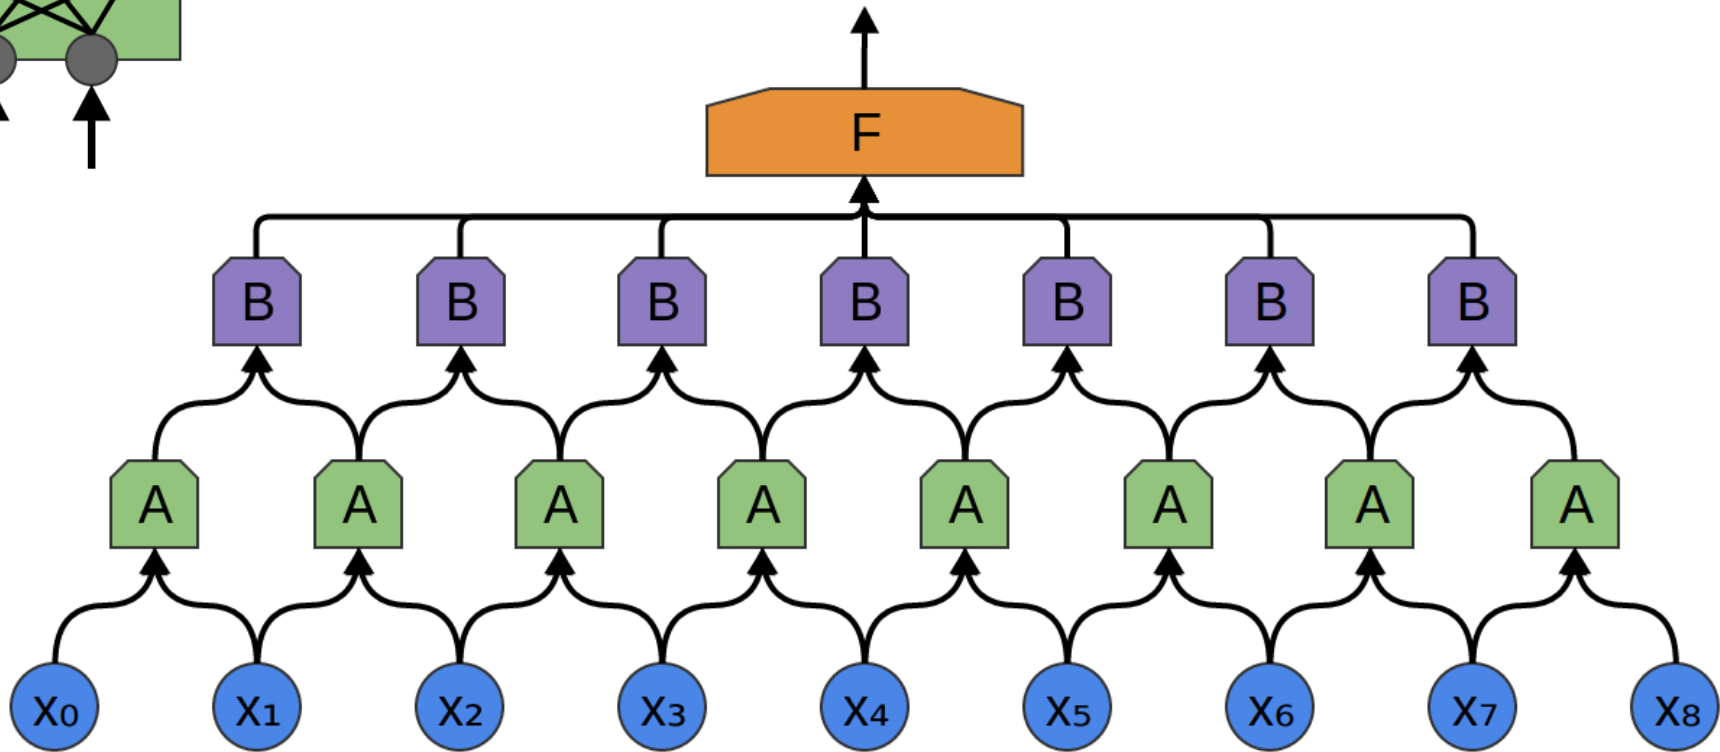



UNIVERSITAT POLITÈCNICA DE CATALUNYA
BARCELONATECH

Department of Signal Theory
and Communications

Image Processing Group


end-to-end differentiable!



AlexNet



UNIVERSITY OF
TORONTO



A Krizhevsky, I Sutskever, GE Hinton “[Imagenet classification with deep convolutional neural networks](#)” Part of: [Advances in Neural Information Processing Systems 25 \(NIPS 2012\)](#)

neural networks have been around for decades – why now?

recent advances in deep learning are due to:

huge labeled datasets

deeper and deeper architectures

some simple, but clever, tricks to better train these networks

demo! ...?

Deep Visualization Toolbox

yosinski.com/deepvis

#deepvis



Jason Yosinski



Jeff Clune



Anh Nguyen



Thomas Fuchs



Hod Lipson



yosinski.com/deepvis

huge labeled datasets

ImageNet Challenge

IMAGENET


- 1,000 object classes (categories).
- Images:
 - 1.2 M train
 - 100k test.






☐ I'm not a robot

reCAPTCHA

Type the text







Verify




☐ I'm not a robot

reCAPTCHA

Select all images below that match this one:







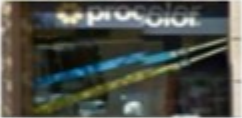





Verify



Lo




Select all images with a store front













VERIFY




amazon mechanical turk



MNIST

(Modified National Institute of Standards and Technology)

60,000 training, 10,000 testing



(256 object categories, 30607 images)

dog



A horizontal collage of nine small images related to basketball. From left to right: 1. A basketball player in a white jersey with the number 32 is shown from behind, jumping towards the hoop. 2. A basketball hoop with a black net against a dramatic, purple and orange sunset sky. 3. A basketball hoop with a white net, with a basketball and a pink egg on the floor in front of it. 4. A basketball hoop with a white net, with a basketball on the floor in front of it. 5. A basketball hoop with a white net, with a basketball on the floor in front of it. 6. A basketball hoop with a white net, with a basketball on the floor in front of it. 7. A basketball hoop with a white net, with a basketball on the floor in front of it. 8. A basketball hoop with a white net, with a basketball on the floor in front of it. 9. A basketball player in a black jersey is shown from behind, jumping towards the hoop.

and many more...

including industrial datasets

(e.g. Netflix Prize – \$1,000,000 prize to create rating prediction algorithm better than Netflix's current best)


Kaggle.com

Data.gov

...


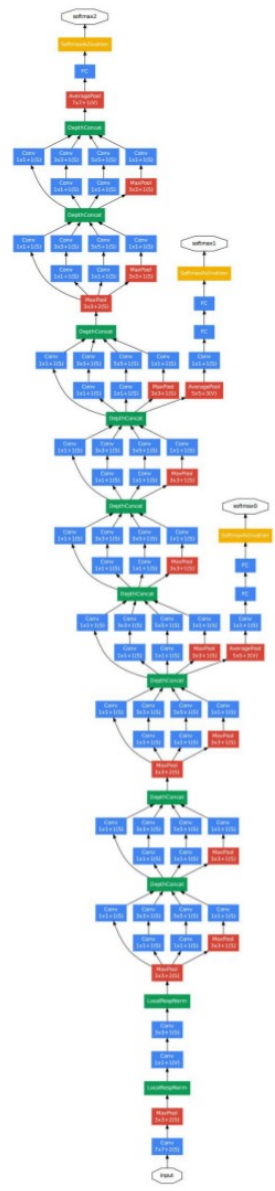


deeper and deeper architectures

ILSVRC top-5 error on ImageNet



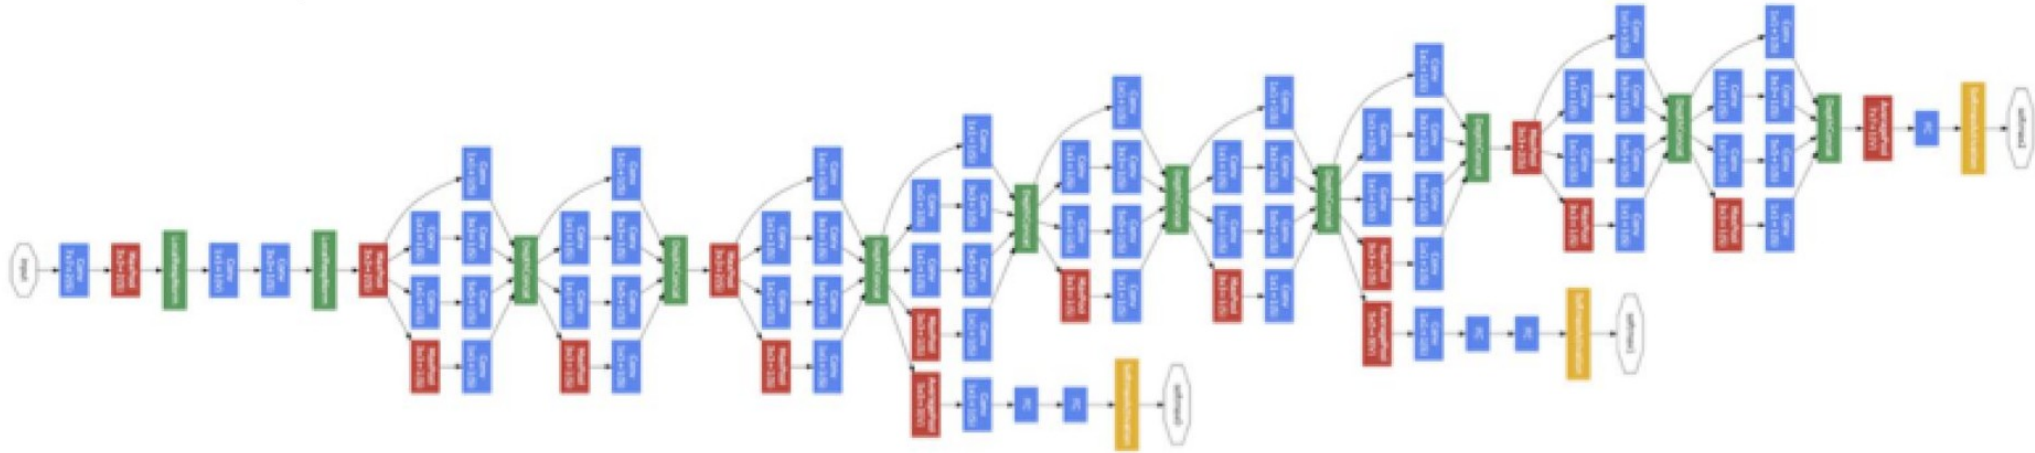
ImageNet Challenge: 2014

AlexNet



GoogLeNet (Inception)


- 22 layers, but 12 times fewer parameters than AlexNet.



Convolution
Pooling
Softmax
Other

Szegedy, Christian, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. ["Going deeper with convolutions."](#)

GoogLeNet (Inception)

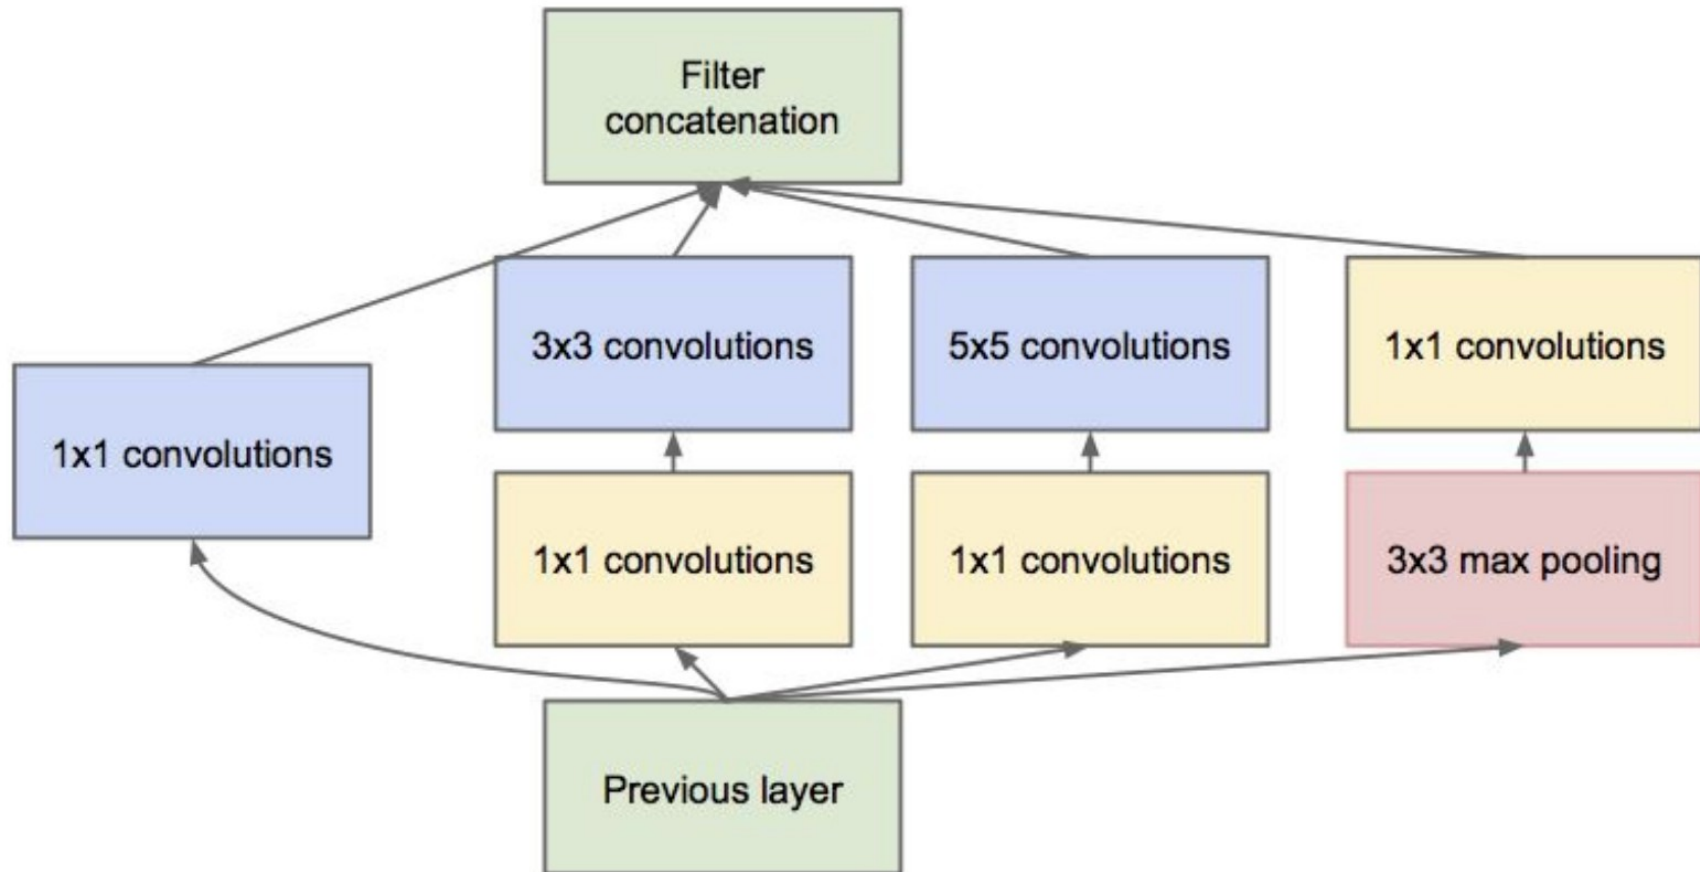


9 **Inception** modules

Network in a network in a network...

Convolution
Pooling
Softmax
Other

GoogLeNet (Inception)




Lin, Min, Qiang Chen, and Shuicheng Yan. ["Network in network."](#) *ICLR 2014*.

GoogLeNet (NiN)

3x3 convolutions

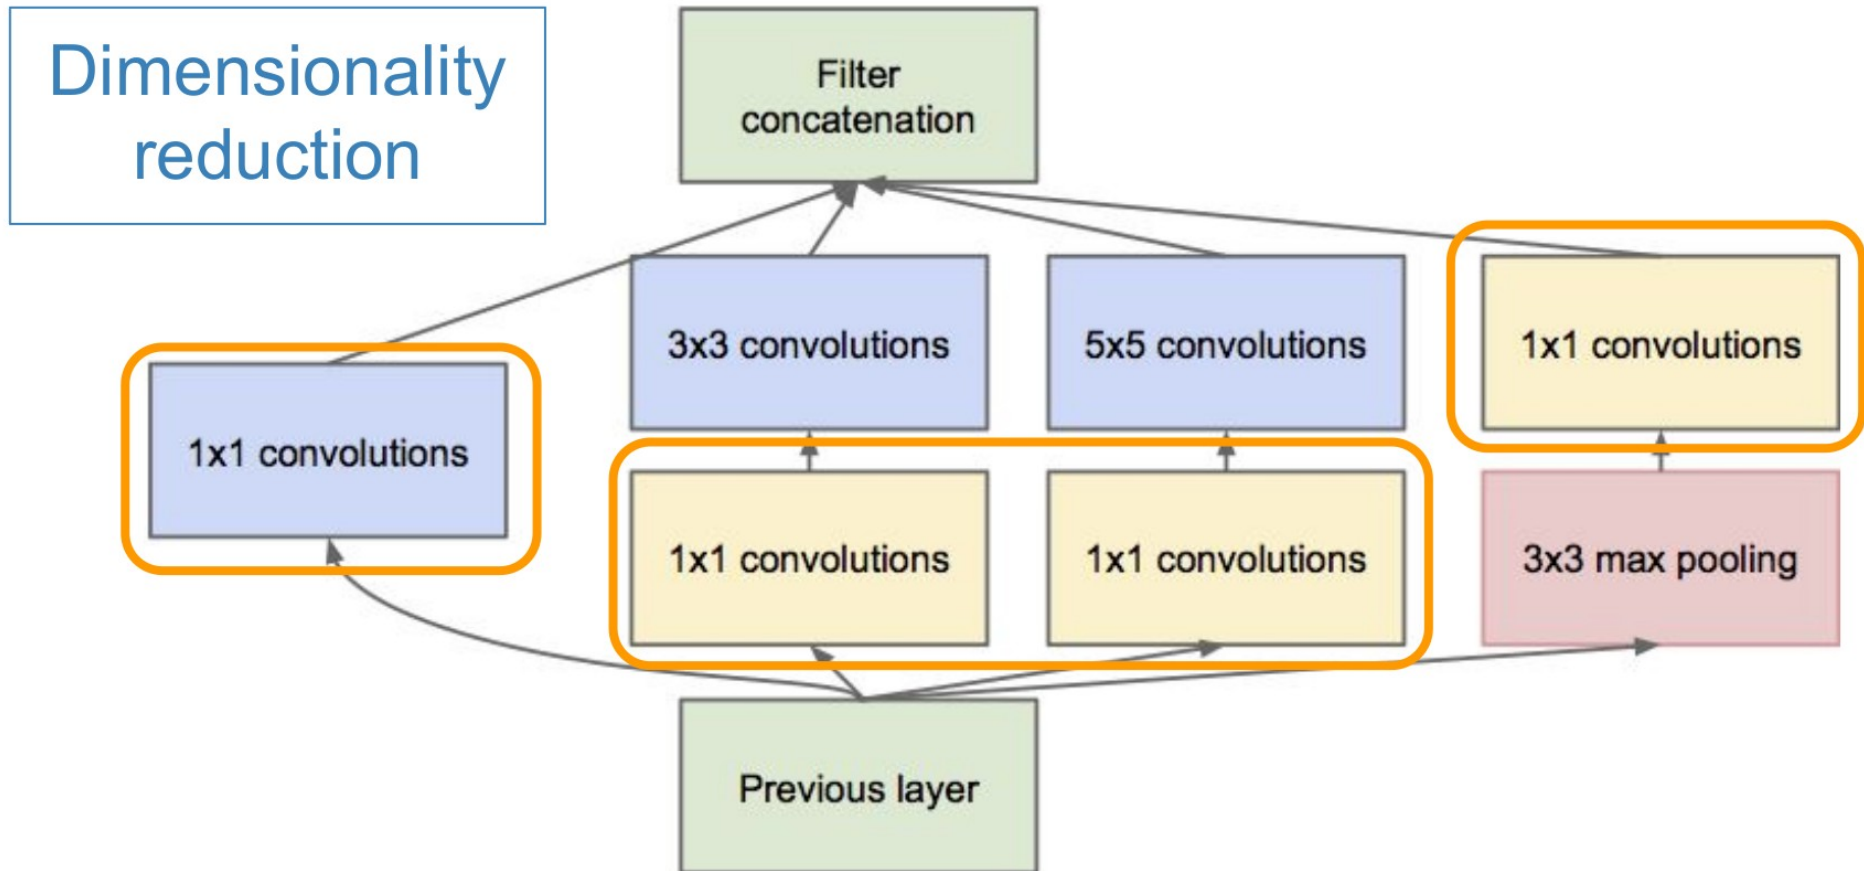
5x5 convolutions

3x3 and 5x5 convolutions deal with different scales.



Lin, Min, Qiang Chen, and Shuicheng Yan. ["Network in network."](#) ICLR 2014. [\[Slides\]](#)

GoogLeNet (Inception)

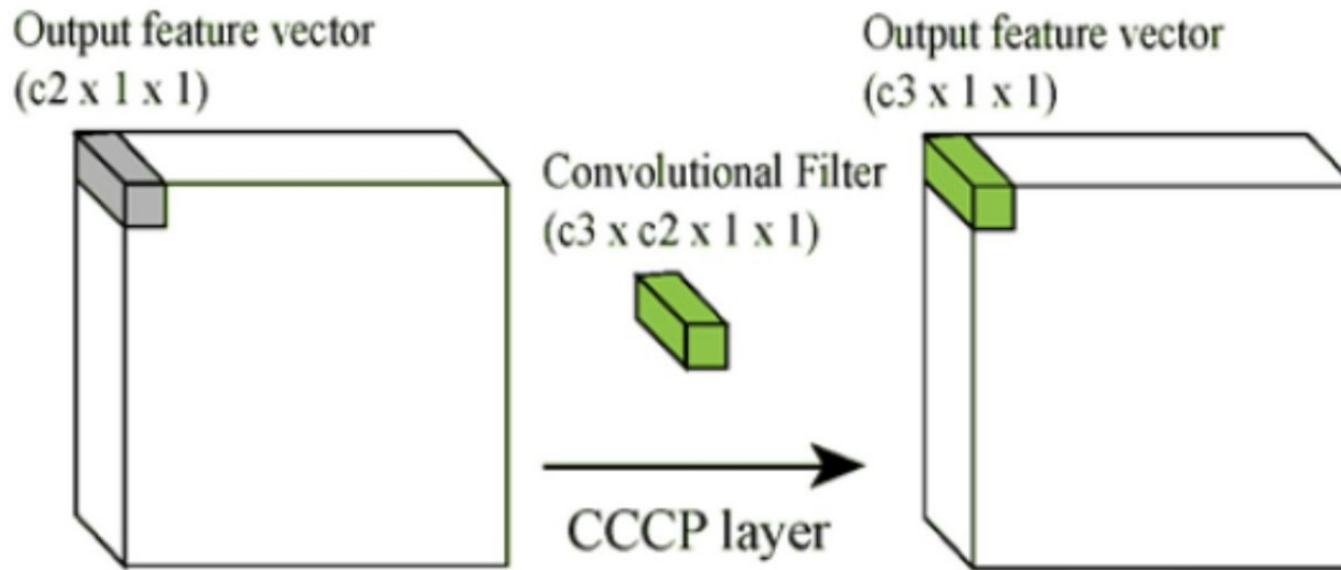


Lin, Min, Qiang Chen, and Shuicheng Yan. ["Network in network."](#) ICLR 2014.

GoogLeNet (Inception)

1x1 convolutions


1x1 convolutions does dimensionality reduction ($c_3 < c_2$) and accounts for rectified linear units (ReLU).



Lin, Min, Qiang Chen, and Shuicheng Yan. ["Network in network."](#) ICLR 2014. [\[Slides\]](#)

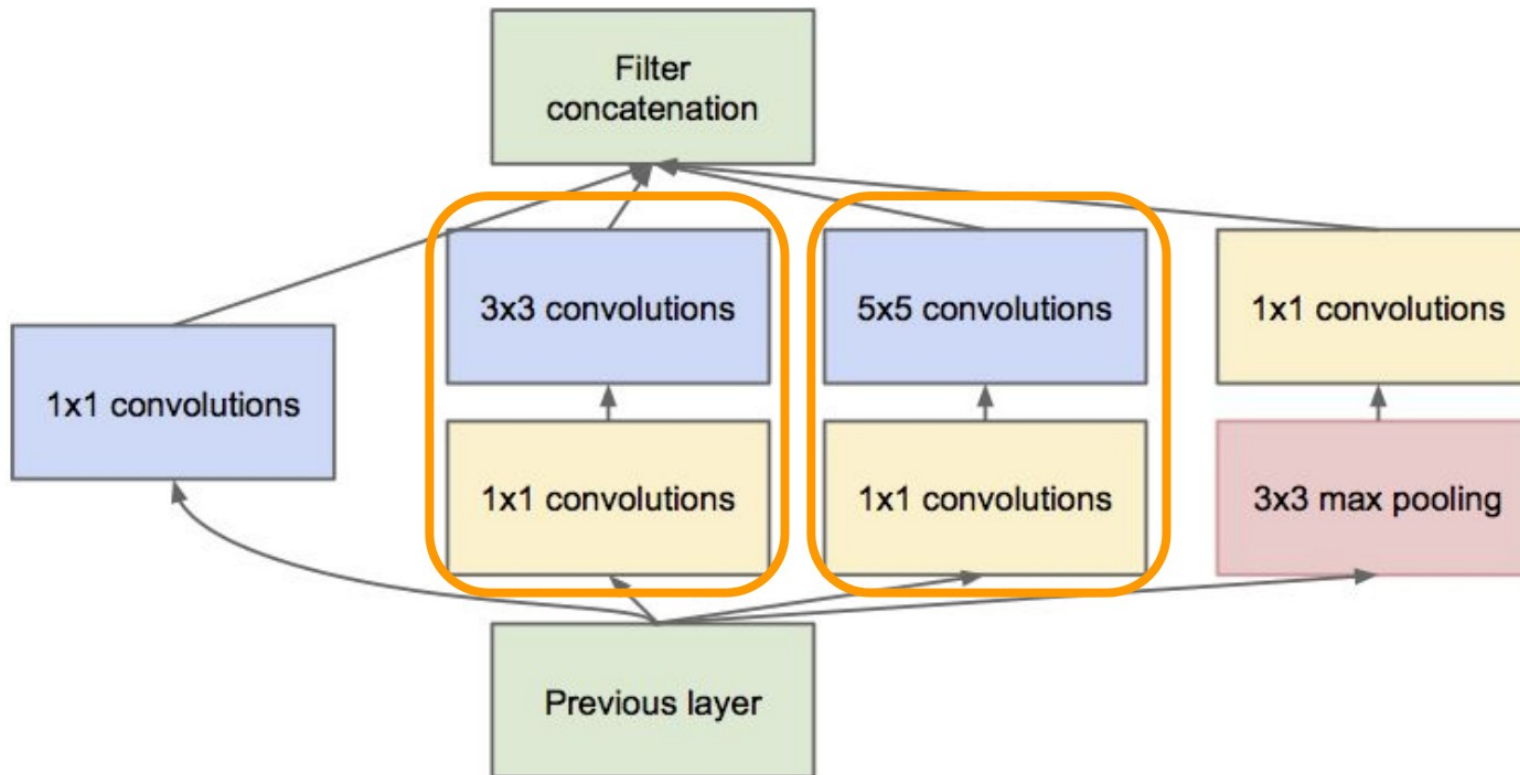
GoogLeNet (Inception)

In GoogLeNet, the **Cascaded 1x1 Convolutions** compute reductions before the expensive 3x3 and 5x5 convolutions.

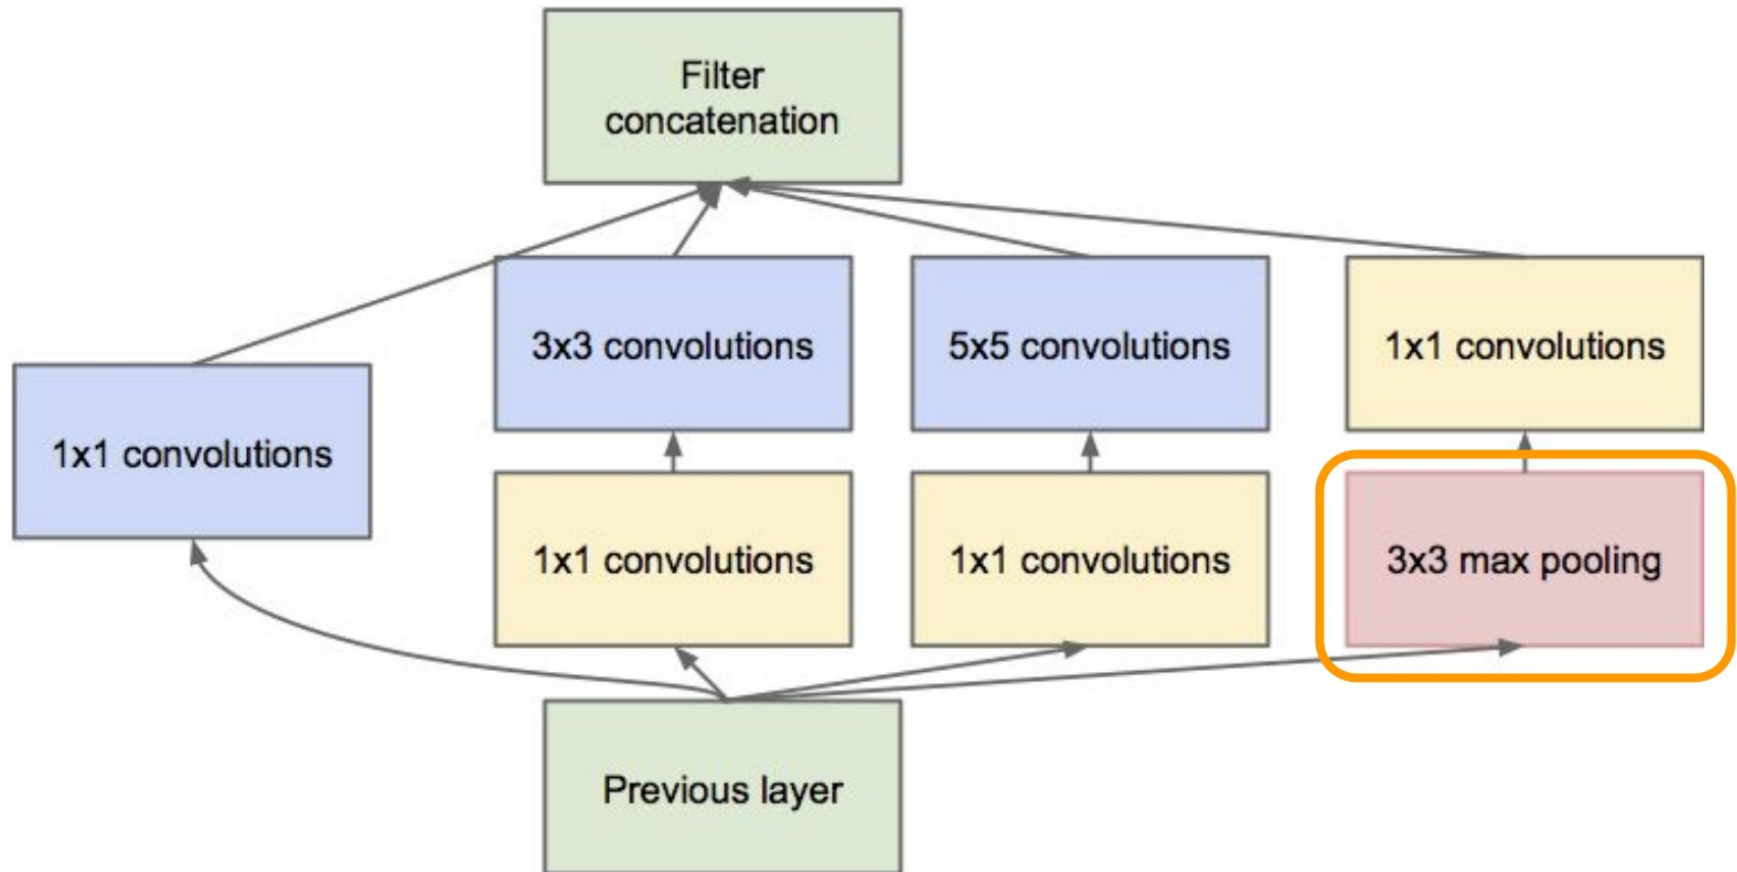


GoogLeNet (Inception)

In GoogLeNet, the **Cascaded 1x1 Convolutions** compute reductions before the expensive 3x3 and 5x5 convolutions.




GoogLeNet (Inception)




Lin, Min, Qiang Chen, and Shuicheng Yan. ["Network in network."](#) *ICLR 2014*.


GoogLeNet (Inception)



They somewhat spatial invariance, and has proven a beneficial effect by adding an alternative parallel path.




Convolved
feature



Pooled
feature

Conceiving the Inception Module



GoogLeNet (Inception)

Two Softmax Classifiers at intermediate layers combat the vanishing gradient while providing regularization at training time.

