Course Logistics

CS 698O: Visual Recognition

Vinay P. Namboodiri
vinaypn@iitk.ac.in

Instructor Details

- Vinay P. Namboodiri
- # RM 406, RM building, CSE
- Office Hours: Tuesday and Thursday 12noon - 1 pm
- Preferably email: vinaypn@iitk, Subject: CS6980….

About the Course

- Visual Recognition
- A new graduate elective
- Objectives: Obtain different perspectives on understanding visual recognition
  - Problems
  - Approaches
  - Advances

Lectures

- Monday, Wednesday and Friday
- M 10-10:50, W 10-10:50, F 12 -12:50
- Venue: KD 101

Grading

- Weightage (Tentative):
- Quizzes 10%
- Mid-Sem - 20%
- End-Sem 20%
- Assignments 25% (Programming, Paper Review, Paper Presentation)
- Project 25% (atleast 2 stages)

Course Outline

- Introduction
- Exact instance retrieval
- Classification
- Detection
- Segmentation
- Weak Supervision
- Active Learning
- Domain Adaptation
- Unsupervised Representation learning
- Dynamic Temporal Aspects

Course Outline

- Introduction
- Exact instance retrieval
- Classification
- Detection
- Segmentation
- Weak Supervision
- Active Learning
- Domain Adaptation
- Unsupervised Representation learning
- Dynamic Temporal Aspects

Course Outline

- Introduction
- Exact instance retrieval
- Classification
- Detection
- Segmentation

Traditional Feature Based

- Weak Supervision
- Active Learning
- Domain Adaptation
- Unsupervised Representation learning
- Dynamic Temporal Aspects

Course Outline

- Introduction
- Exact instance retrieval
- Classification
- Detection
- Segmentation
- Weak Supervision
- Active Learning
- Domain Adaptation
- Unsupervised Representation learning
- Dynamic Temporal Aspects

Deep learning based

Course Outline

- Introduction
- Exact instance retrieval
- Classification
- Detection
- Segmentation

Deep learning based

- Weak Supervision
- Active Learning
- Domain Adaptation
- Unsupervised Representation learning
- Dynamic Temporal Aspects

Tentative set of advanced topics

Course Material

- Lecture slides that will be posted online
- Course will be based mainly on research papers

- Reference books:
- Computer Vision: Algorithms and Applications by Richard Szeliski Available online
- Computer Vision: Models, Learning, and Inference by Simon J.D. Prince Available online
- Deep Learning by Ian Goodfellow, Yoshua Bengio and Aaron Courville Available online
- Computer Vision: A Modern Approach by Forsyth and Ponce Indian edition available

Introduction

CS 698O: Visual Recognition

Vinay P. Namboodiri
vinaypn@iitk.ac.in

What is Visual Recognition?

Semantic Representation

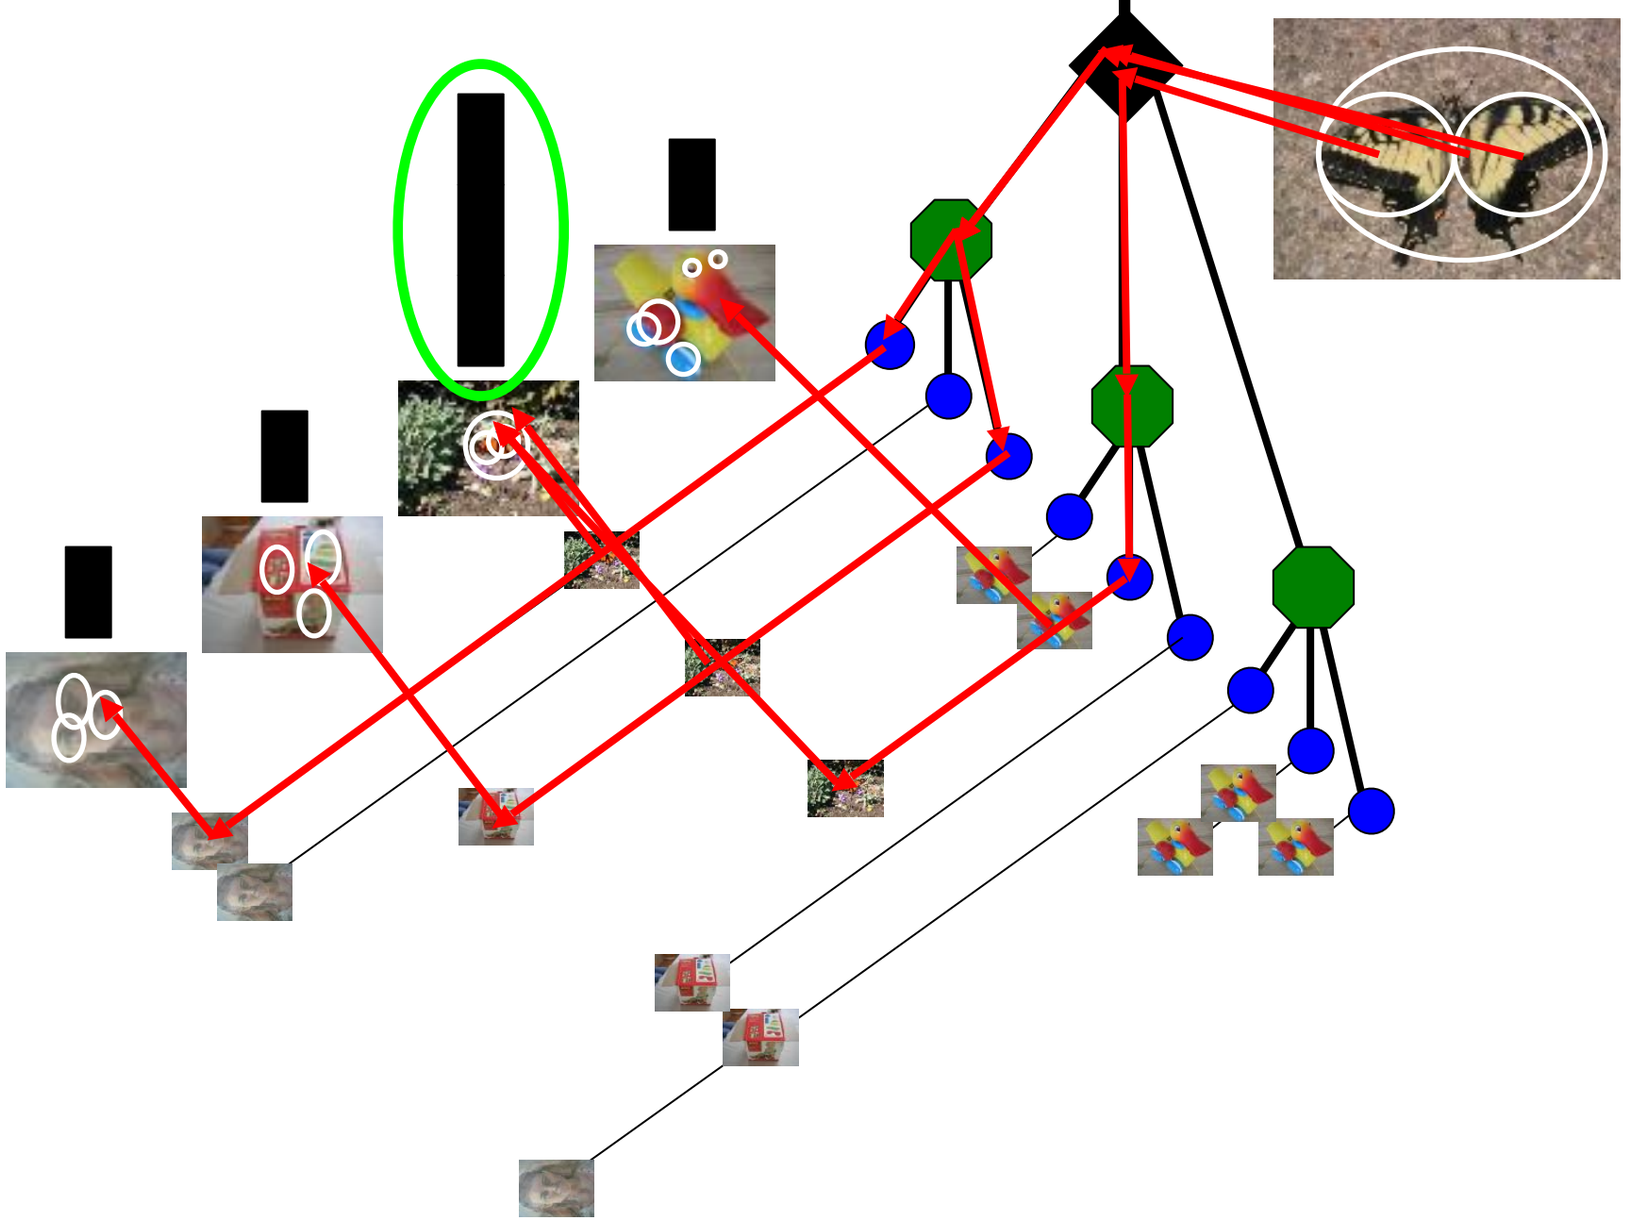

Instance Recognition

Scalable recognition with a Vocabulary tree
Nister, Strewenius, CVPR 2006

Semantic Representation

Object Classification
ImageNet

Image credit: Karpathy

Semantic Representation

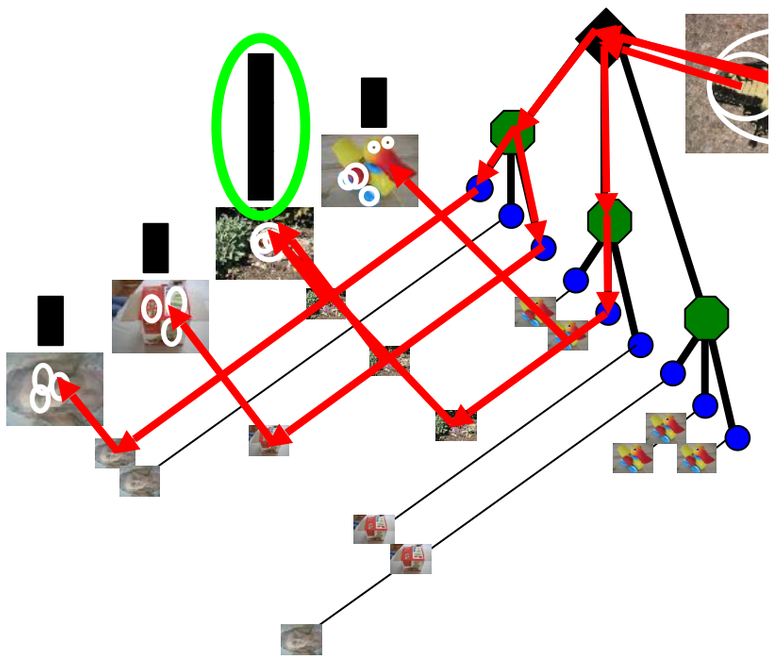

Instance

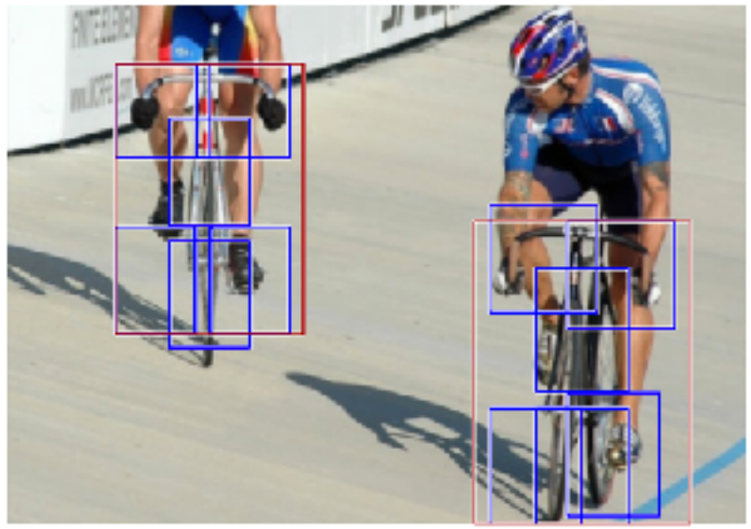

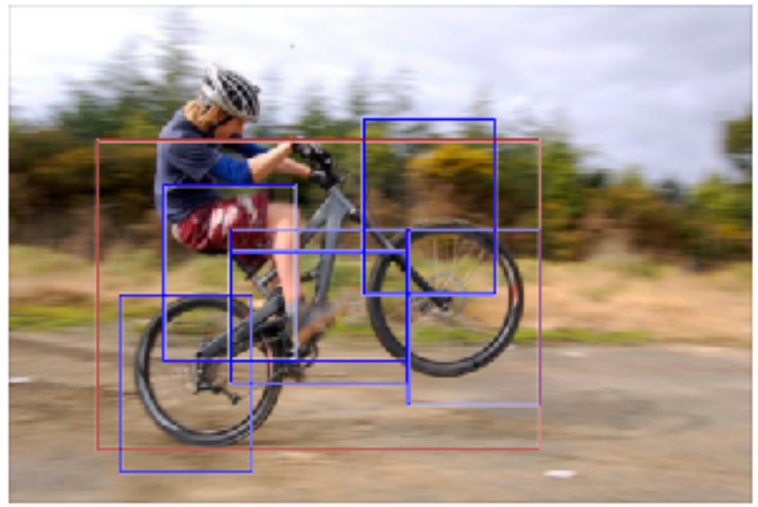

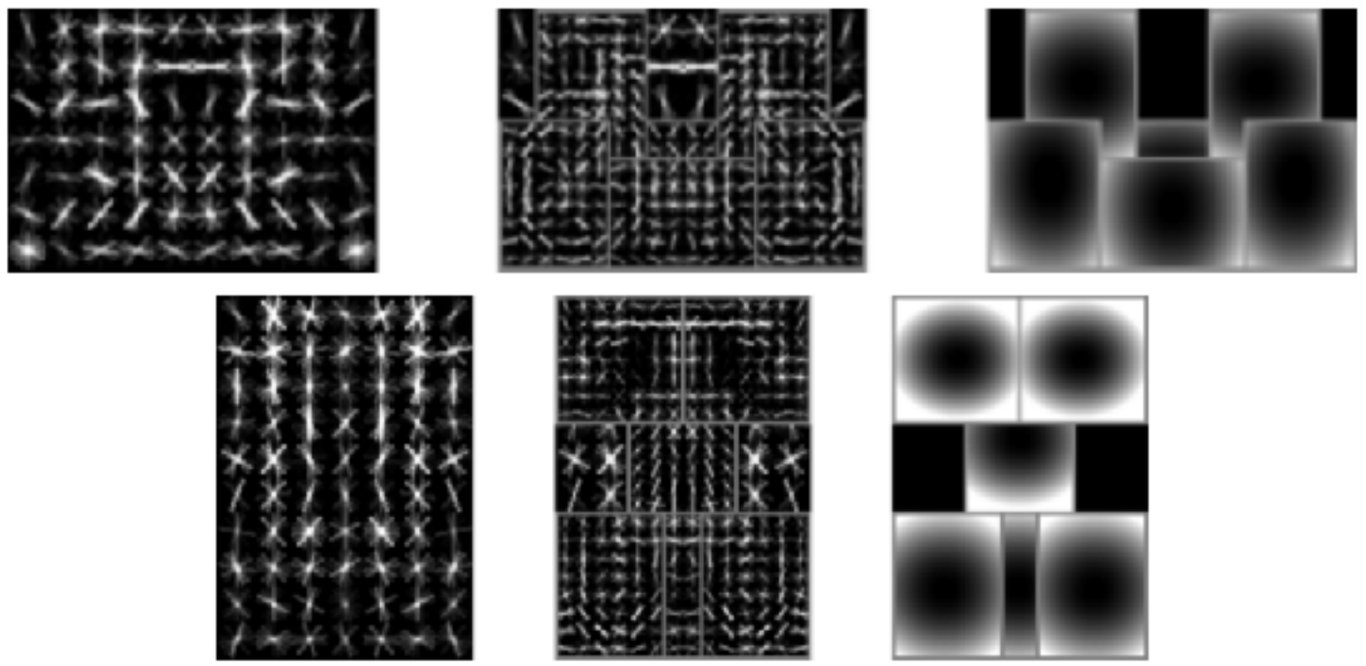

Object Detection

Object Detection with Discriminatively Trained Part Based Models
P. Felzenszwalb, R. Girshick, D. McAllester, D. Ramanan
PAMI 2010

Semantic Representation

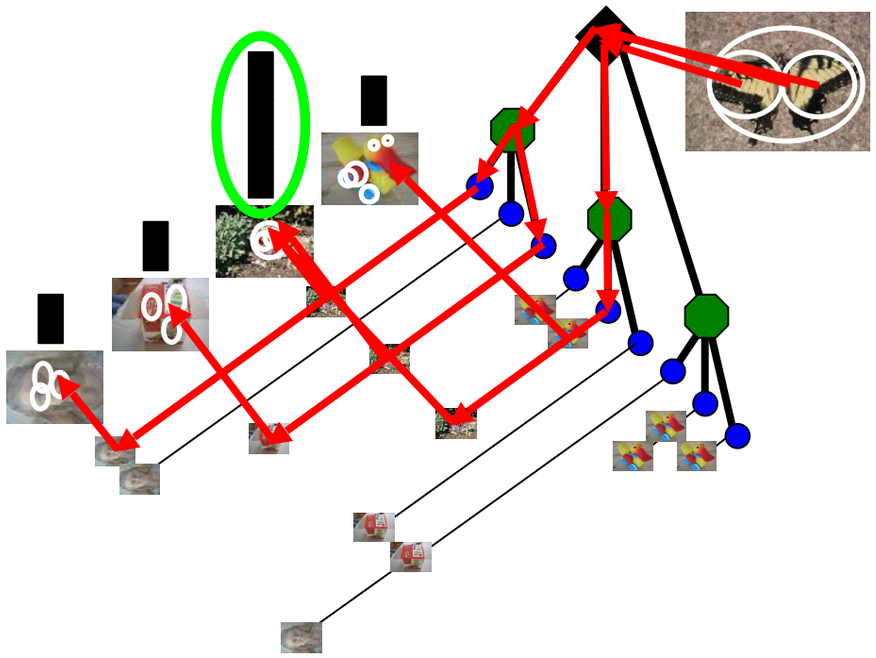

Instance

Object

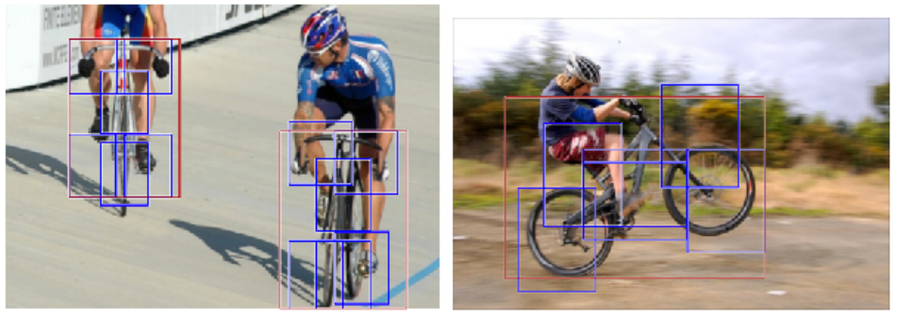

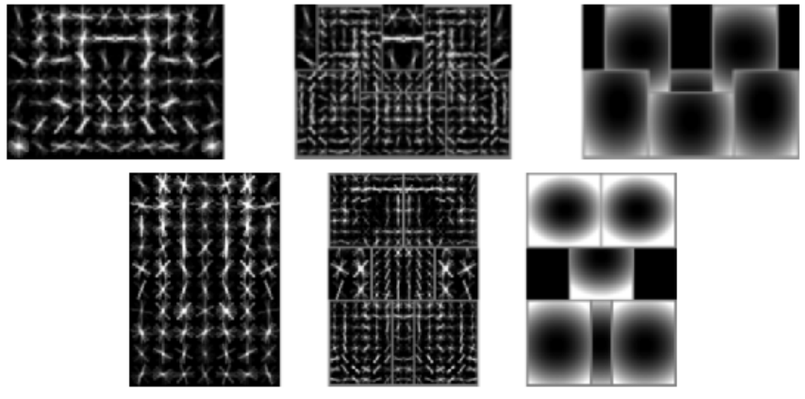

Object Detection

Visual Recognition

Why understand Visual Recognition?

Motivation

- Intellectual curiosity
- Algorithms for general visual perception (also enable general machine learning methods)
- Applications

Intellectual Challenge

Intellectual Challenge

Figure credit: Hakan Bilen

Intellectual Challenge

- Making machines *see*

- Extracting semantic information from signals

Table 1

3	120	23	33
6	34	45	56
1	59	67	90
90	99	23	84
200	121	89	55

Figure credit: Hakan Bilen

Algorithms

- Segmentation (Graph partitioning, Non-parametric density estimation)
- Denoising (L1 norm based denoising)
- Template Matching
- Deep Neural networks

Applications

Self driving cars

Applications

Surveillance

Applications

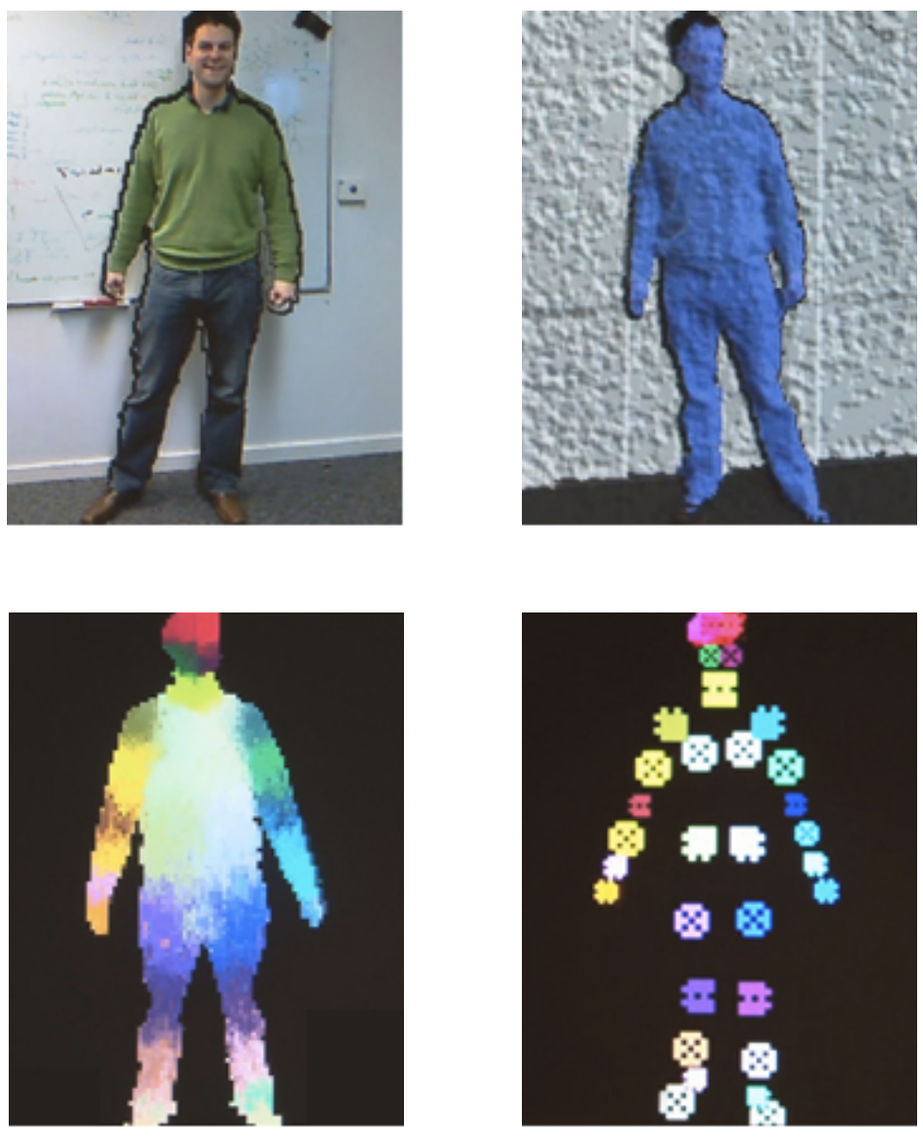

Human-Computer Interfaces

Challenges

Challenges

Challenges

Challenges

Recent Successes

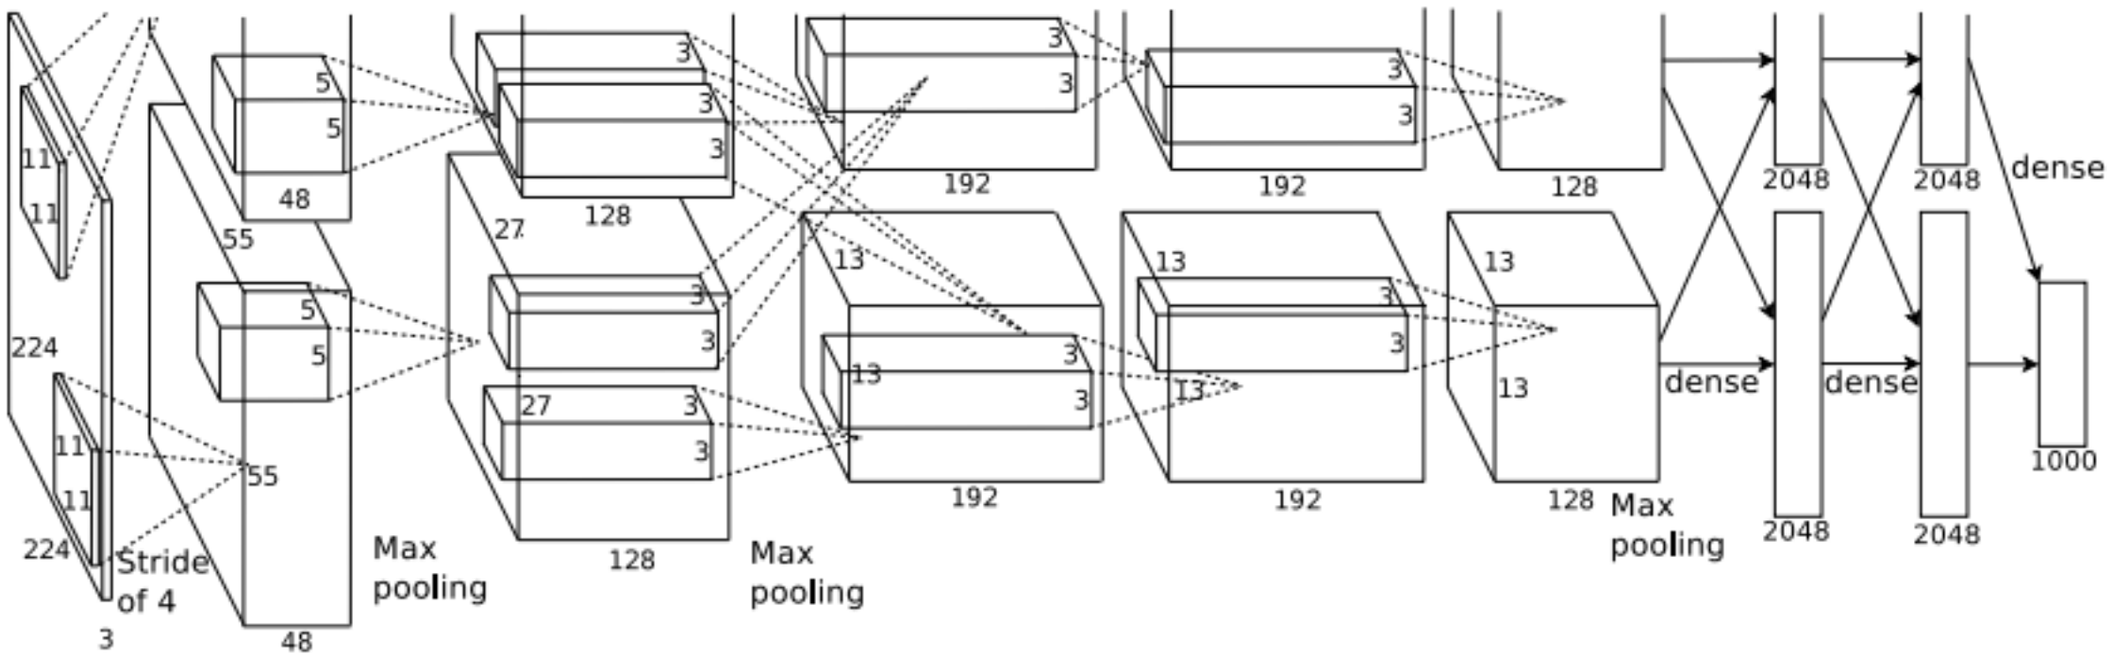

ImageNet Classification with Deep Convolutional Neural Networks
Alex Krizhevsky, Ilya Sutskever, Geoffrey E. Hinton
NIPS 2012

Recent Successes

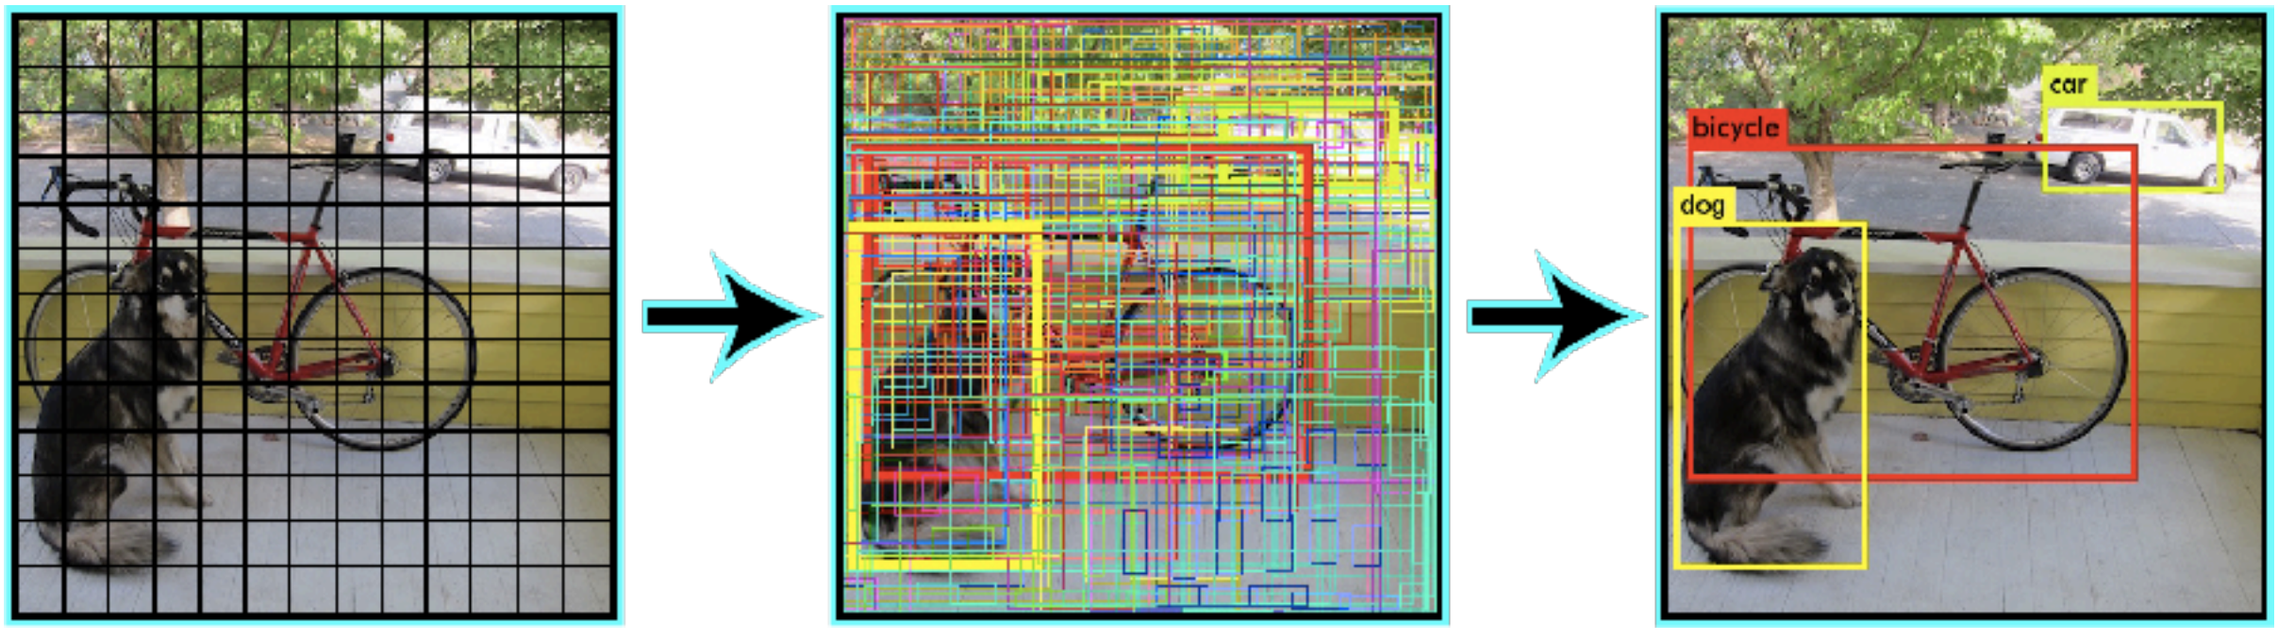

You Only Look Once: Unified, Real-Time Object Detection
Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi
CVPR 2016

Recent Successes

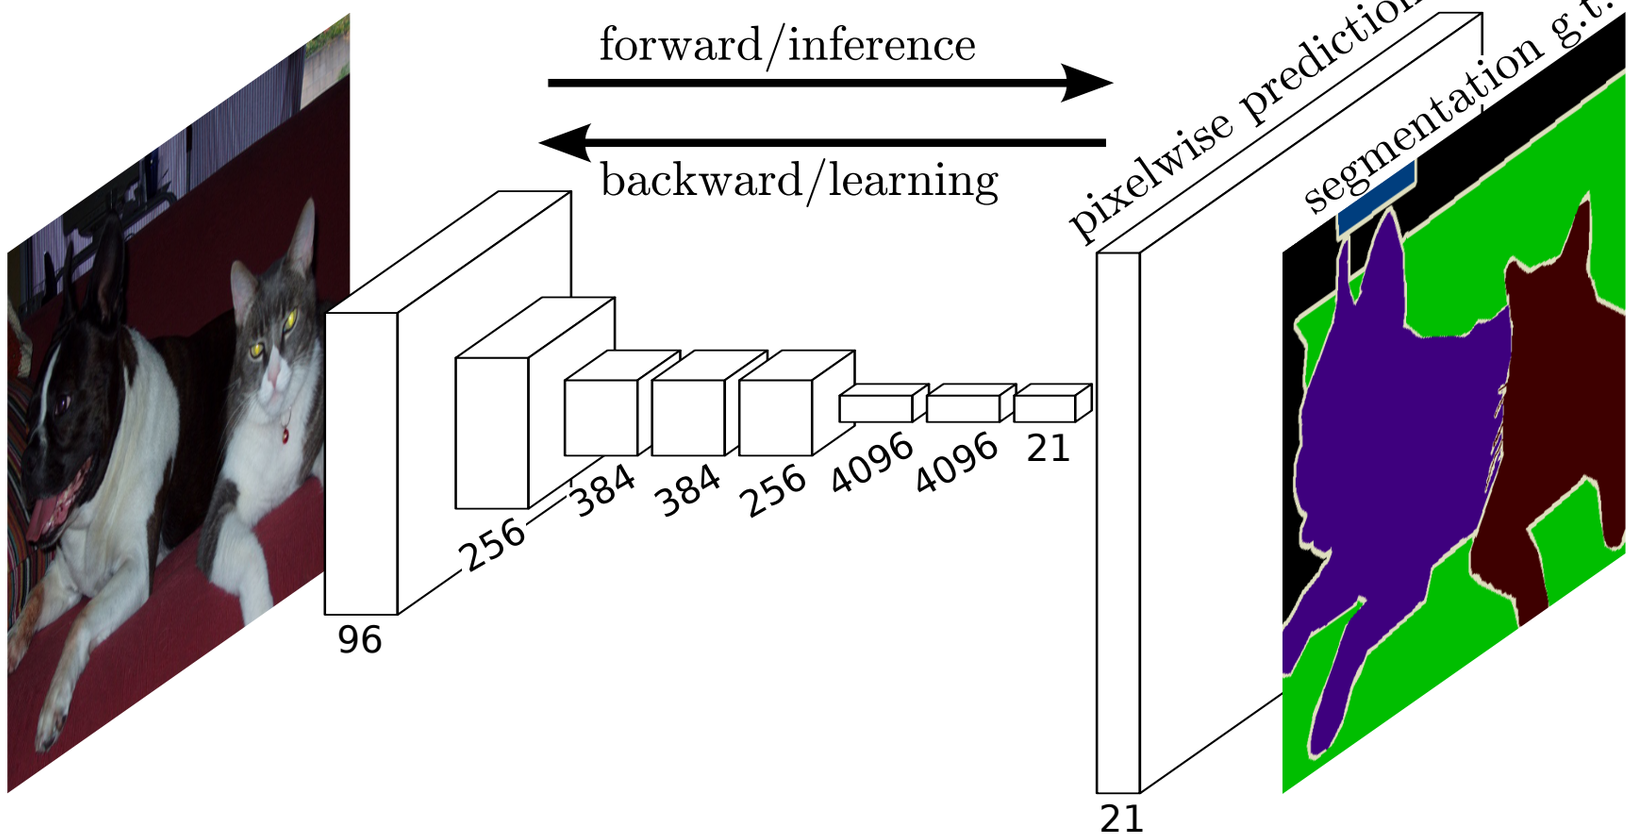

Fully Convolutional Networks for Semantic Segmentation

Jon Long*, Evan Shelhamer*, Trevor Darrell

CVPR 2015

Conclusion

- Study of visual recognition is one of the classical and interesting problems that is fascinating
- Solving this enables many applications
- This could enable us to move towards real developments in *AI*