2018학년도 2학기 [딥러닝(영강)] 강의계획서

◆ 수업정보 ◆

[수업정보]

시간/강의실	월(7-8) 정보통신관 205호 수(7) 정보통신관 205호		
학점	3학점	학수번호(분반)	COSE474(00)
이수구분	전공선택		

[강의담당자]

성명	주재걸	소속	컴퓨터학과
E-mail			
Homepage			
연구실호실		연락처	
면담시간			

[조교정보]

성명		소속	
E-mail			
연구실		연락처	

◆ 수업운영 ◆

[수업방법]

활동유형	강의, 발표, 토론, 실습, 협동학습, 집단지도, 퀴즈, QnA

[평가방법]

항목	점수	항목	점수
수시과제	20 점	텀프로젝트	25 점
퀴즈	5 점	참여도	5 점
중간시험	20 점	기말시험	25 점
총점	100 점		
평가점수공개여부	비공개		

COSE474(00): Special Topics: Deep Learning

Lecture 0. Course Overview

Instructor: Jaegul Choo (주재걸)

주재걸 교수 소개

- 2015~: 고려대학교 컴퓨터학과 조교수
- 서울대학교 전기공학부 학사 (2001)
- Georgia Tech ECE 석사 (2005), CSE 박사 (2013)
- Main Research

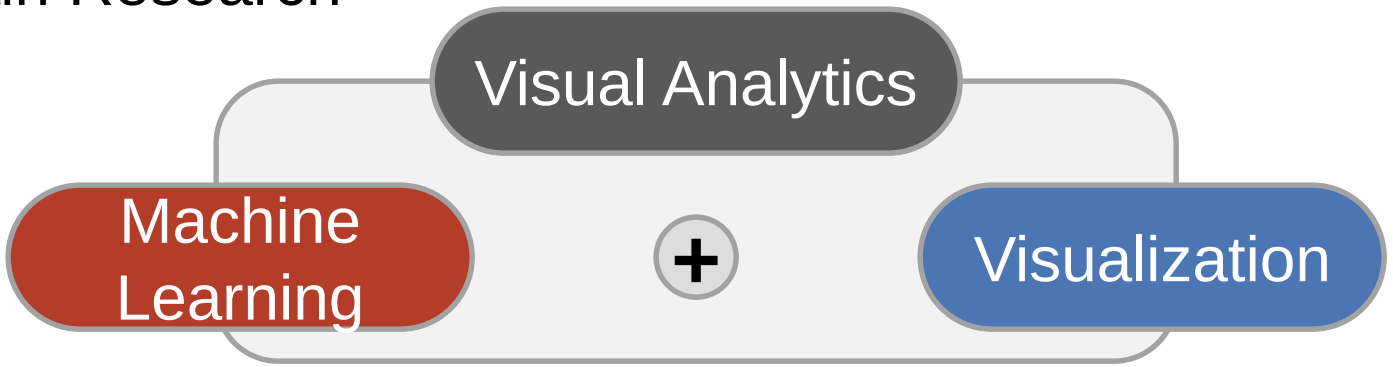

- Published >80 research articles (~900 citations)
  - **ML/DM**: KDD, WWW, WSDM, CVPR, AAAI, IJCAI, ICDM, TKDD, DMKD, ICWSM, SDM
  - **Vis/HCI**: IEEE VIS, CHI, TVCG, EuroVIS, CGF, CG&A
  - Best Student Paper Award at ICDM'16
  - 네이버 신진교수상, 2015
  - Best Poster Award at IEEE Vis'14

Today's Lecture

- Basic Course Information
- Course Schedule
- Grading Policy

Basic Course Information

- Time and Location
  - Monday 5:00-6:15pm
  - Woojeong Information Building 205 (우정정보통신관 205)
- Instructor: Jaegul Choo (주재걸), Ph.D
  - Office: Room# 502, Woojeong Information Building (우정정보통신관 502)
- Course Website:
  - We will use blackboard.

  (All the class materials including the slides and homework assignments will be available here.)

Goal of This Course

- This course will study the theory and the application of deep neural networks.
- In detail, we will cover deep neural networks techniques including fully connected networks, convolutional neural networks, recurrent neural networks and long short-term memory, attention models, memory networks, generative adversarial networks.
- In addition, we will discuss various techniques necessary to effectively train these models, such as initialization, normalization, gradient descent, pre-training, regularization.

Flipped Class

Our lecture will be flipped classes. You have to watch one or two of the following lectures per week before the class:

- Stanford cs231n (computer vision)
  - http://cs231n.stanford.edu/ (slides + videos)

During our classes, I'll take questions and give in-depth discussions on what you learned from the above lectures.

This is a demanding course. You are advised to take the course only if you can devote a substantial amount of your time.

FYI, my advice on how to study machine learning and deep learning is shown in the following:

- http://blog.naver.com/joyfull1/221004891456

Textbook (Optional)

- Deep Learning, MIT Press, Ian Goodfellow, Yoshua Bengio, Aaron Courville
- Freely available:
  - http://www.deeplearningbook.org/

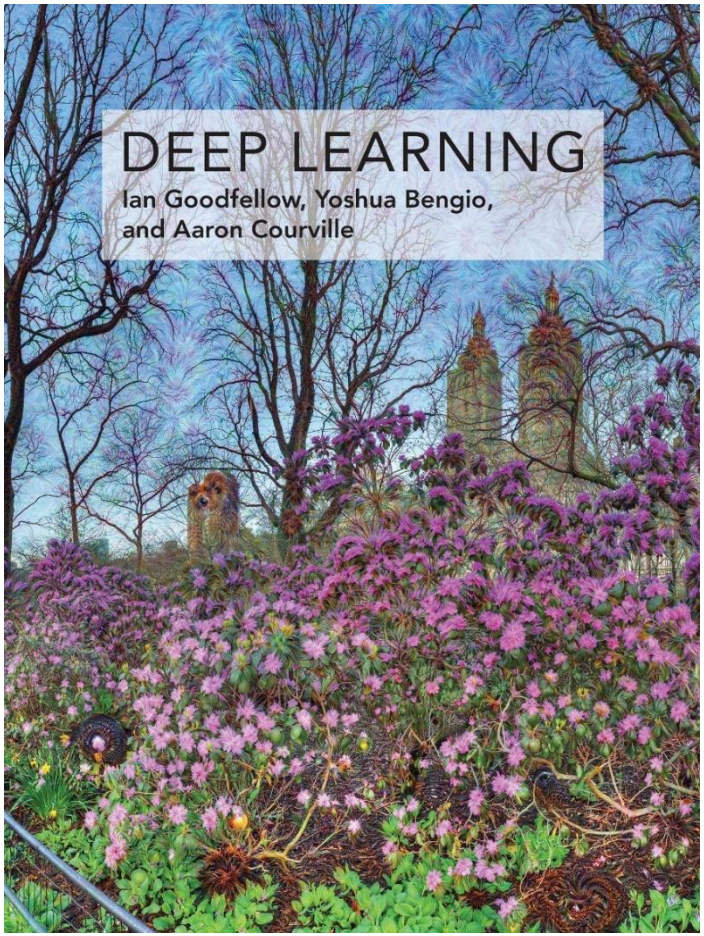

Pre-requisite

- Proficiency in Python
  - All the assignments will be in Python (and numpy)
  - I encourage you to use Python for your term projects.
  - We will mainly use PyTorch as a deep learning library.
- Calculus
- Linear algebra
- Probability and statistics
- Machine Learning
  - I assume that students are familiar with the coursera ML course at https://www.coursera.org/learn/machine-learning
  - I also encourage you to do its homework with Python.

Course Schedule (Tentative)

- Week 01 – Basics of machine learning, linear and logistic regression
- Week 02 – Feed-forward neural networks and computatio nal graph
- Week 03 – Convolutional neural networks
- Week 04 – Training neural networks: initialization, batch n ormalization, dropout, data augmentation, and transfer lea rning and fine-tuning
- Week 05 – Optimization algorithms: stochastic gradient d escent, momentum, AdaGrad, RMSProp, Adam
- Week 06 – Deep learning software, Caffe, Torch, Theano, TensorFlow, Keras, and PyTorch

Course Schedule (Tentative) – cont'd

- Week 07 – CNN architectures: AlexNet, VGG, GoogLeNet , and ResNet
- Week 08 – Midterm
- Week 09 – Word embedding: word2vec and GloVe
- Week 10 – Recurrent neural networks, long short-term m emory, gated recurrent units
- Week 11 – Attention model, image captioning, and seque nce-to-sequence model
- Week 12 – Detection and segmentation
- Week 13 – Visualizing and understanding
- Week 14 – Generative models: generative adversarial net works and variational autoencoders

Course Schedule (Tentative) – cont'd

- Week 15 – Project presentation
- Week 16 – Final Exam

Grading Policy

- 20%: Homework assignments
  - Problem sets
  - Programming tasks
- 20%: Midterm
- 25%: Final
- 25%: Term project
- 5%: In-class quiz
- 5%: Class participation
- Attendance: For each missed class, 0.5% point will be deducted.

Background Knowledge Check

Poll

- Who took artificial intelligence?
- Who took machine learning?
- Anyone who took data science?
- Who has some background knowledge in deep learning?
- Who currently conducts research related to deep learning?

Things to Know in Machine Learning

- Linear regression and logistic regression
- Ensemble learning
- Bias/variance tradeoff
- Overfitting vs. underfitting
- Regularization
- Gradient descent optimization
- Training/validation/test sets
- Cross-validation
- Model selection and hyper-parameter tuning
- Loss function
- Evaluation metric
  - accuracy, precision, recall, F-1 measure, precision-recall curve, area under the ROC curve, …

Deep Learning

Deep learning refers to artificial neural networks that are composed of many layers.

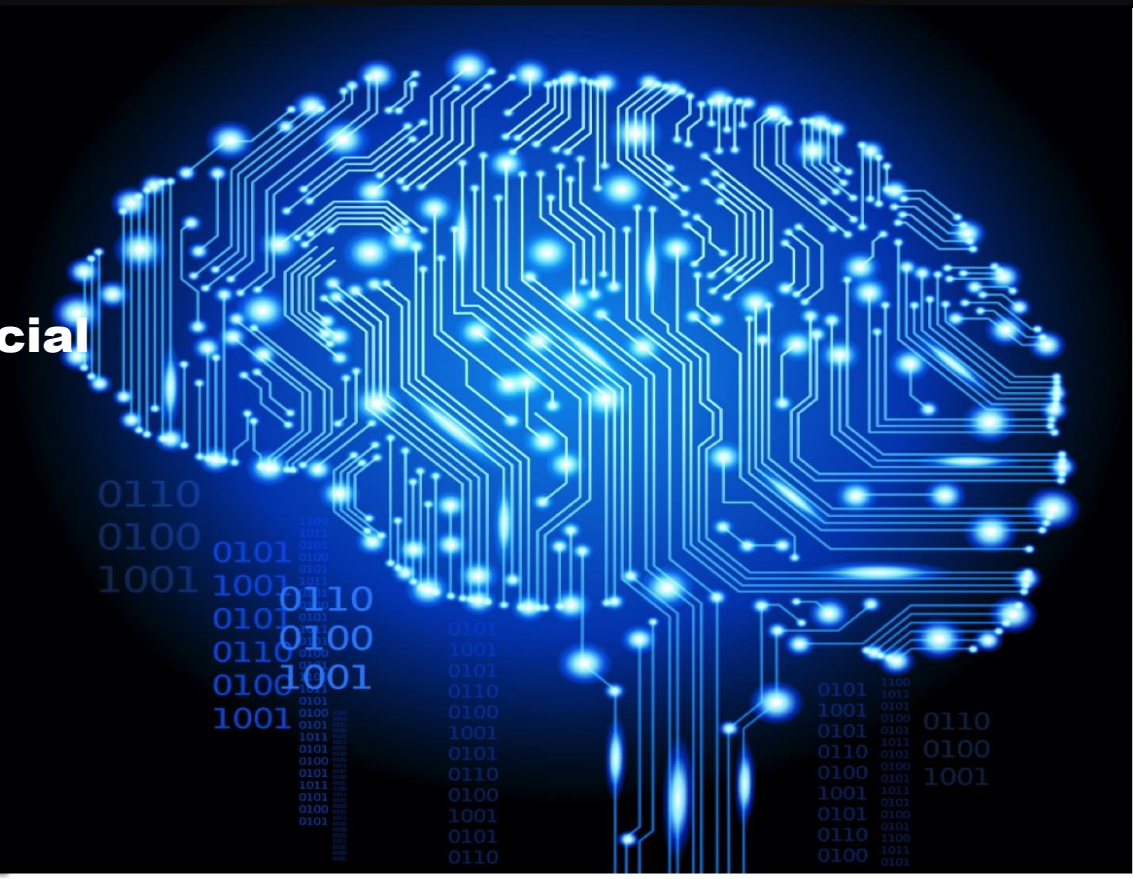

Deep Learning

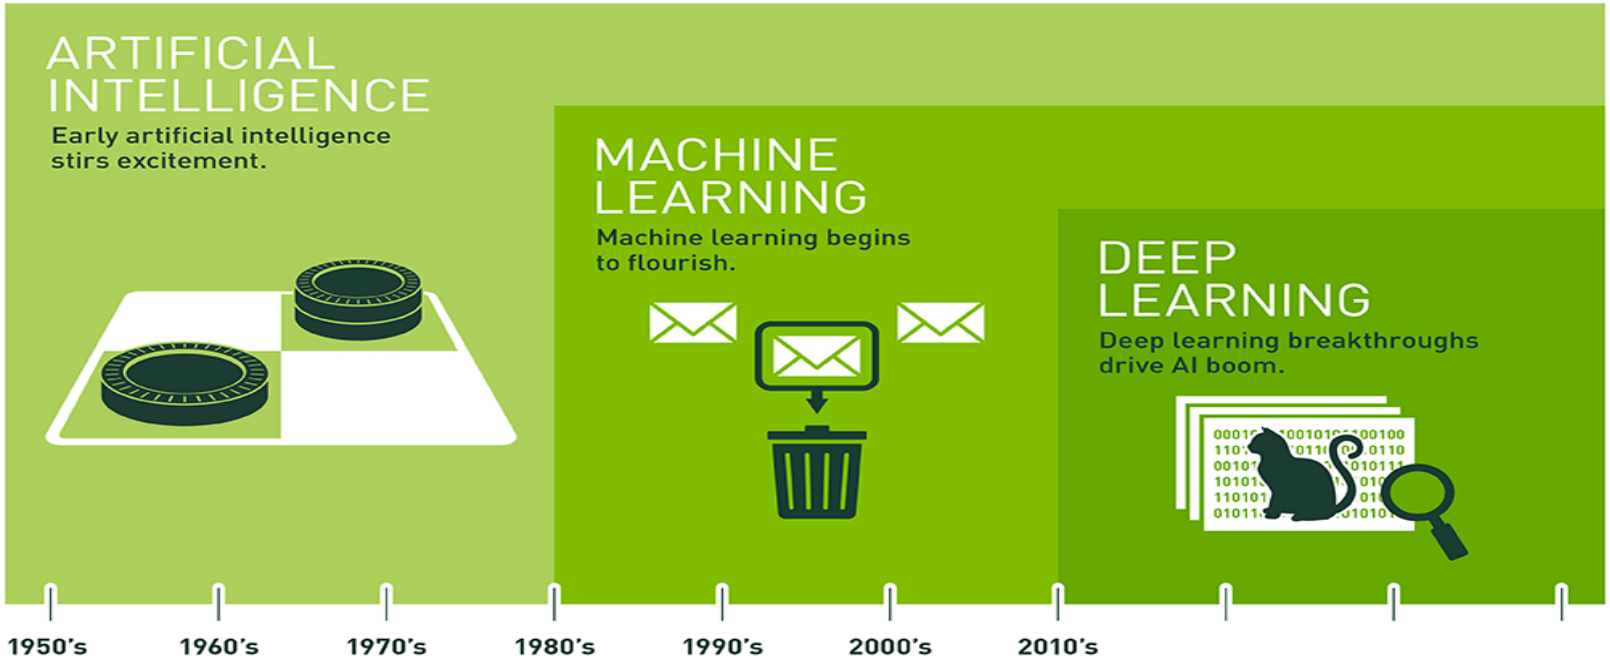

Since an early flush of optimism in the 1950s, smaller subsets of artificial intelligence – first machine learning, then deep learning, a subset of machine learning – have created ever larger disruptions.

Factors for Success of Deep Learning

- Big Data
  - Accumulated data and their labels
- Improved training algorithm
  - Various techniques to properly train deep neural networks
    - ReLU, batch normalization, dropout, …
- Fast hardware
  - NVIDIA’s GPU

Fast Hardware

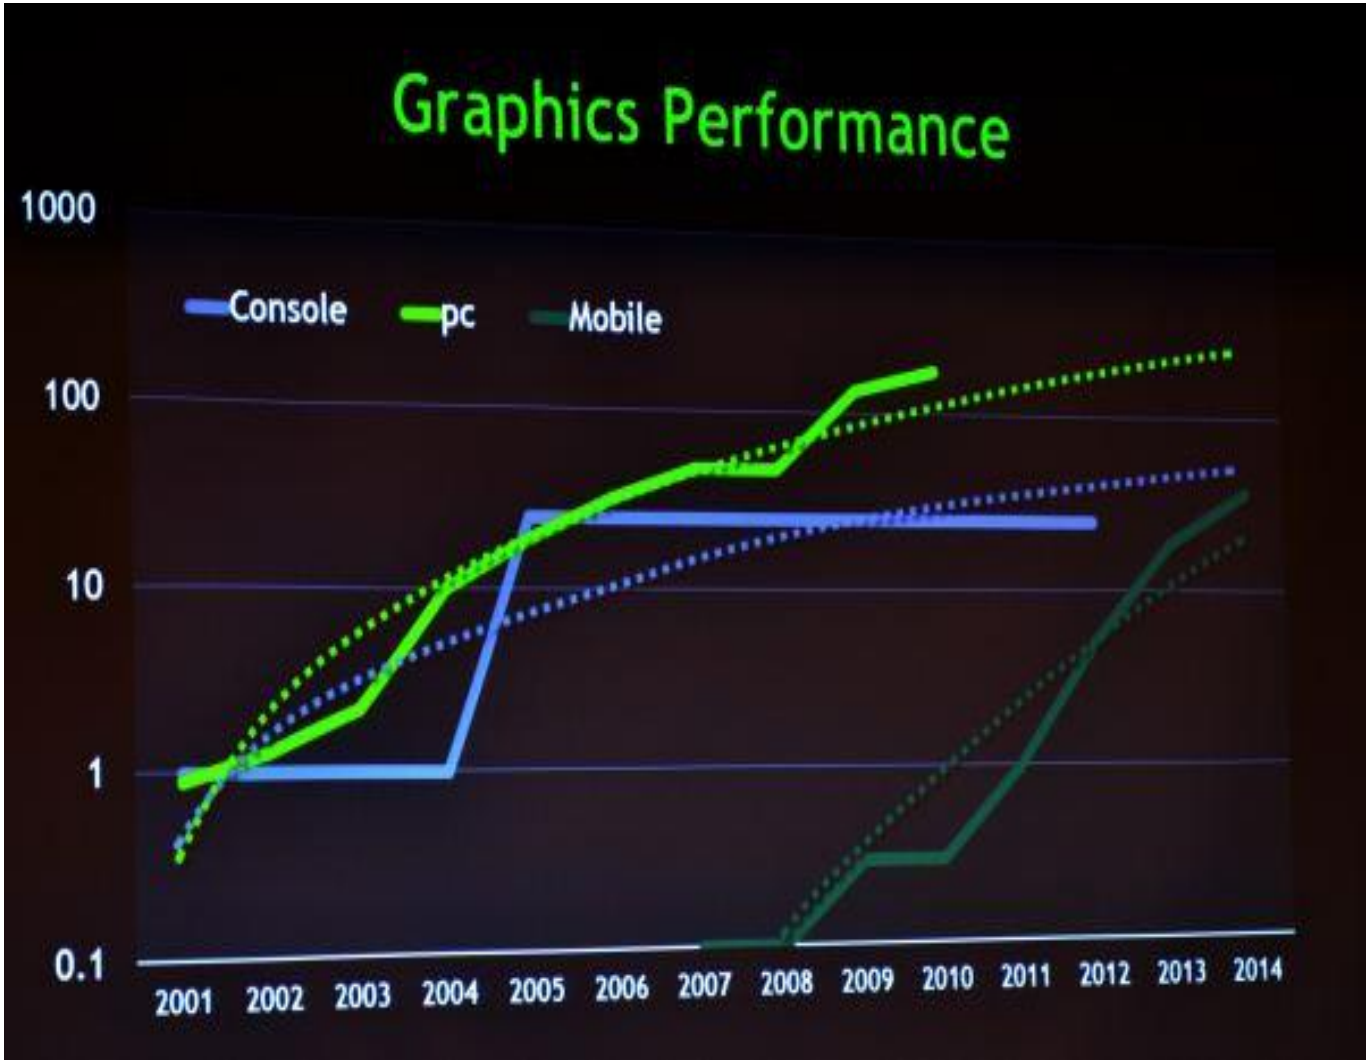

Big Data

ImageNet Challenge

IMAGENET

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

Deep Learning Heroes

Geoffrey Hinton

Yann LeCun

Andrew Ng

Yoshua Bengio

Game AI

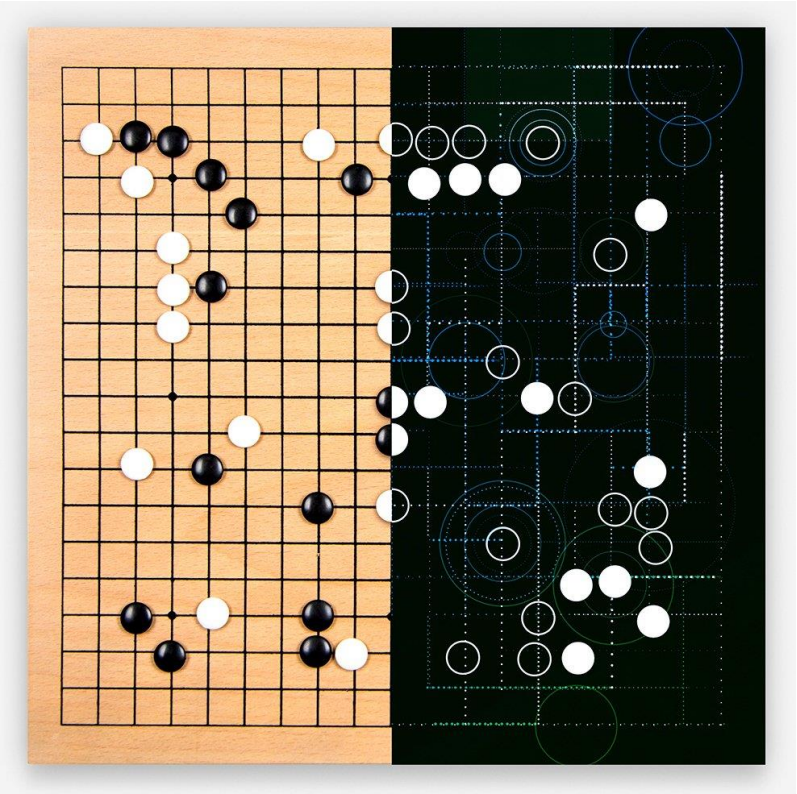

Face Detection & Recognition

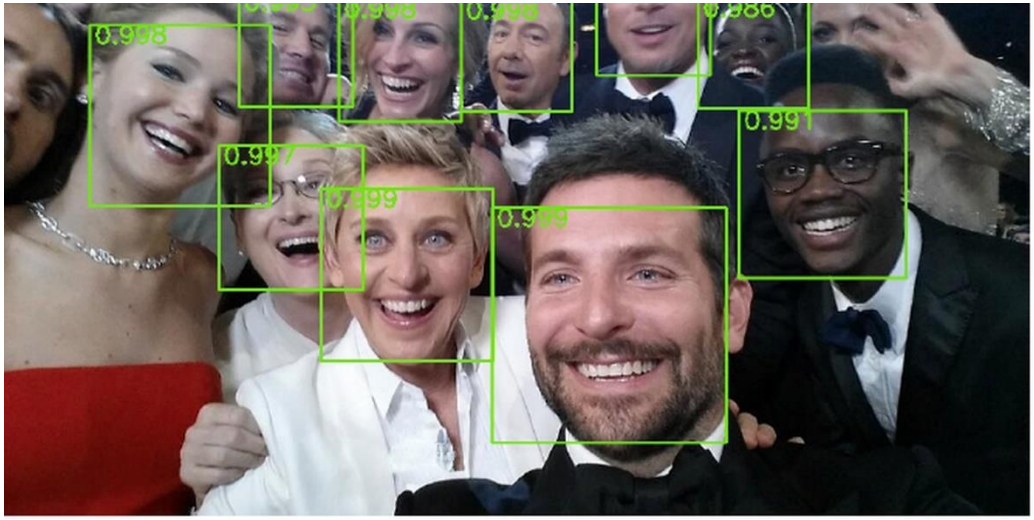

Object Detection & Recognition

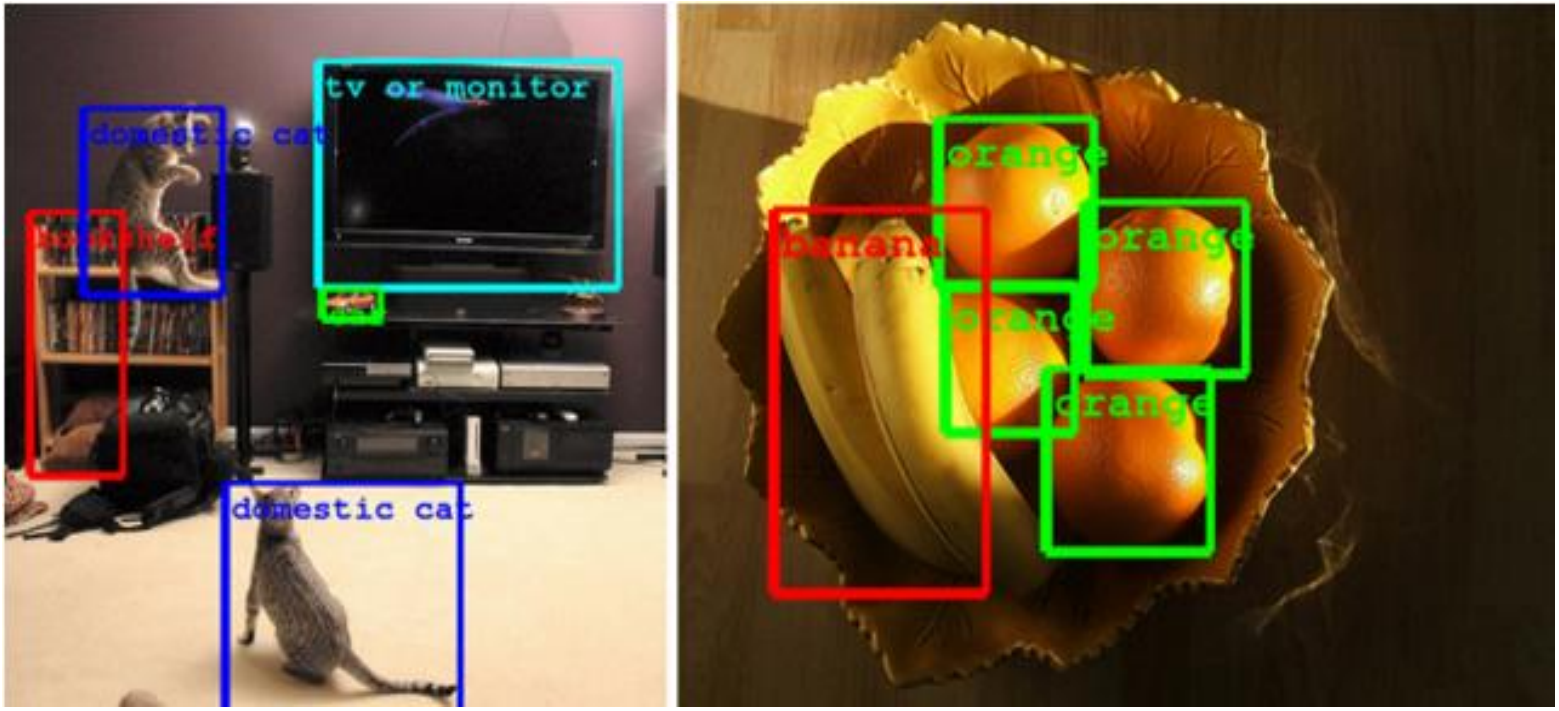

Image Captioning

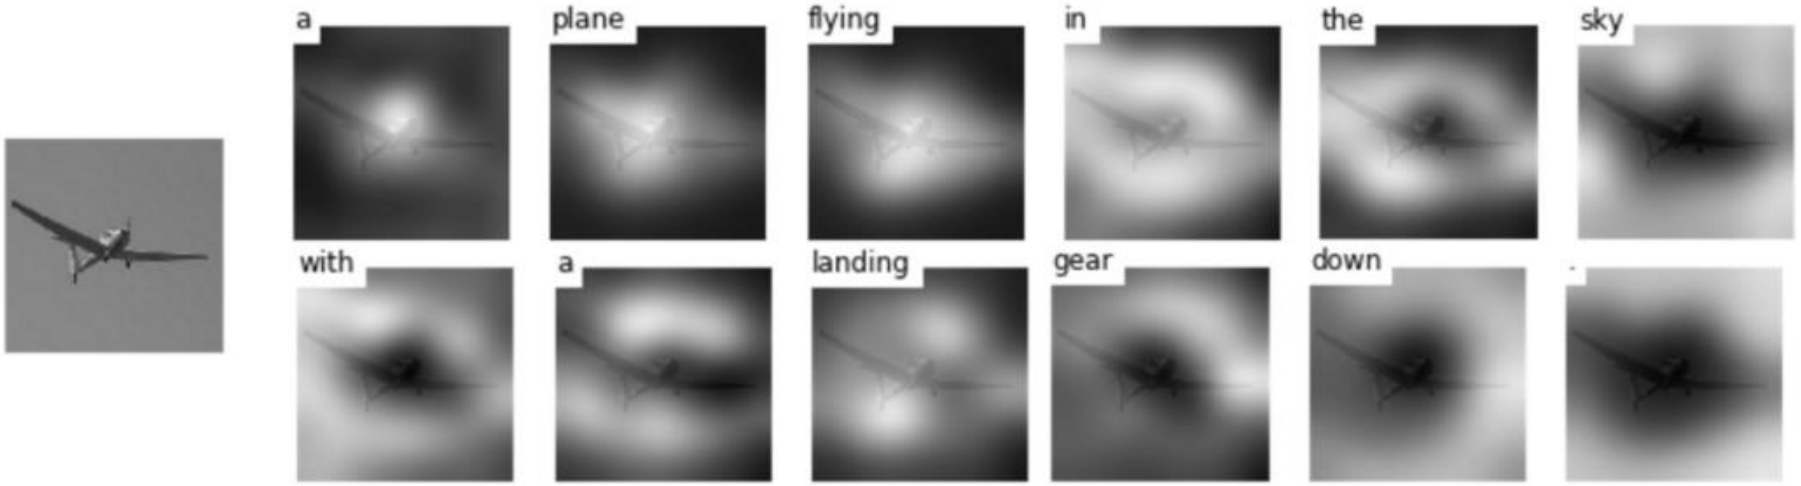

https://github.com/yunjey/show-attend-and-tell

Image Captioning

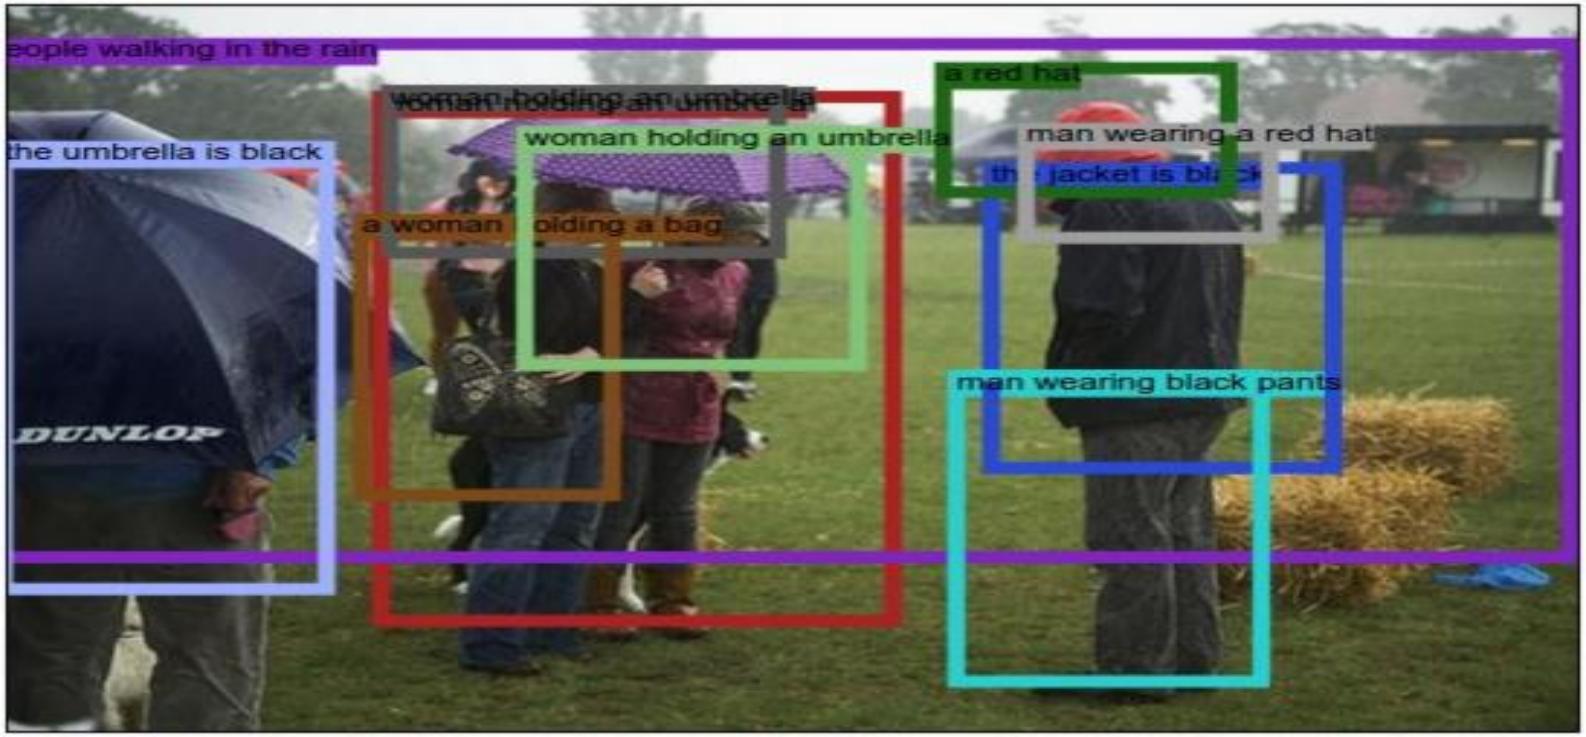

[DenseCap: Fully Convolutional Localization Networks for Dense Captioning]

Machine Translation

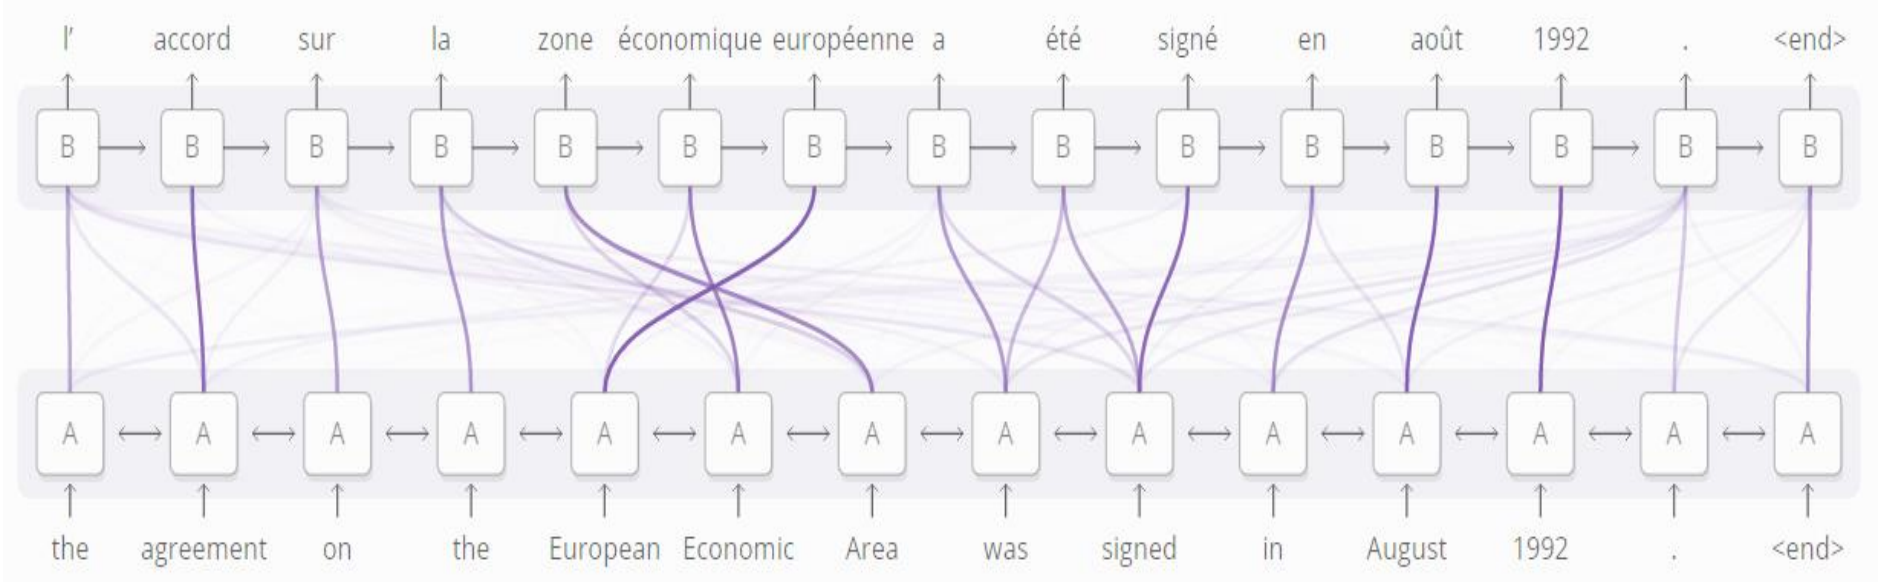

Speech Synthesis

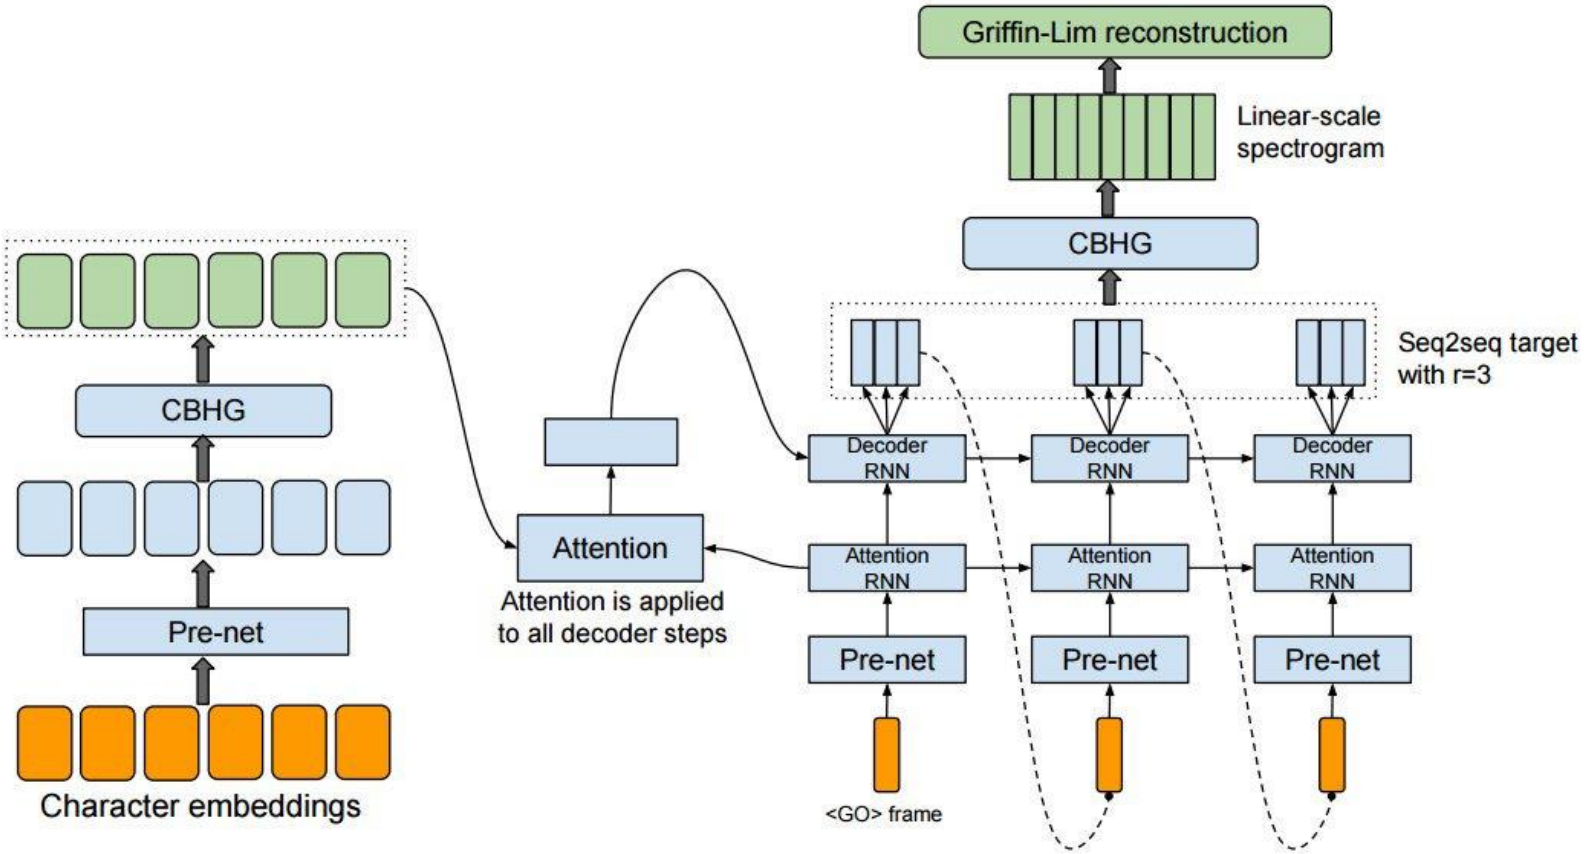

Figure 1: *Model architecture. The model takes characters as input and outputs the corresponding raw spectrogram, which is then fed to the Griffin-Lim reconstruction algorithm to synthesize speech.*

Question Answering

Task 1: Single Supporting Fact

Mary went to the bathroom.
John moved to the hallway.
Mary travelled to the office.
Where is Mary? A:office

Task 2: Two Supporting Facts

John is in the playground.
John picked up the football.
Bob went to the kitchen.
Where is the football? A:playground

Task 3: Three Supporting Facts

John picked up the apple.
John went to the office.
John went to the kitchen.
John dropped the apple.
Where was the apple before the kitchen? A:office

Task 4: Two Argument Relations

The office is north of the bedroom.
The bedroom is north of the bathroom.
The kitchen is west of the garden.
What is north of the bedroom? A: office
What is the bedroom north of? A: bathroom

Task 5: Three Argument Relations

Mary gave the cake to Fred.
Fred gave the cake to Bill.
Jeff was given the milk by Bill.
Who gave the cake to Fred? A: Mary
Who did Fred give the cake to? A: Bill

Task 6: Yes/No Questions

John moved to the playground.
Daniel went to the bathroom.
John went back to the hallway.
Is John in the playground? A:no
Is Daniel in the bathroom? A:yes

Task 7: Counting

Daniel picked up the football.
Daniel dropped the football.
Daniel got the milk.
Daniel took the apple.
How many objects is Daniel holding? A: two

Task 8: Lists/Sets

Daniel picks up the football.
Daniel drops the newspaper.
Daniel picks up the milk.
John took the apple.
What is Daniel holding? milk, football

Task 9: Simple Negation

Sandra travelled to the office.
Fred is no longer in the office.
Is Fred in the office? A:no
Is Sandra in the office? A:yes

Task 10: Indefinite Knowledge

John is either in the classroom or the playground.
Sandra is in the garden.
Is John in the classroom? A:maybe
Is John in the office? A:no

Visual Question Answering

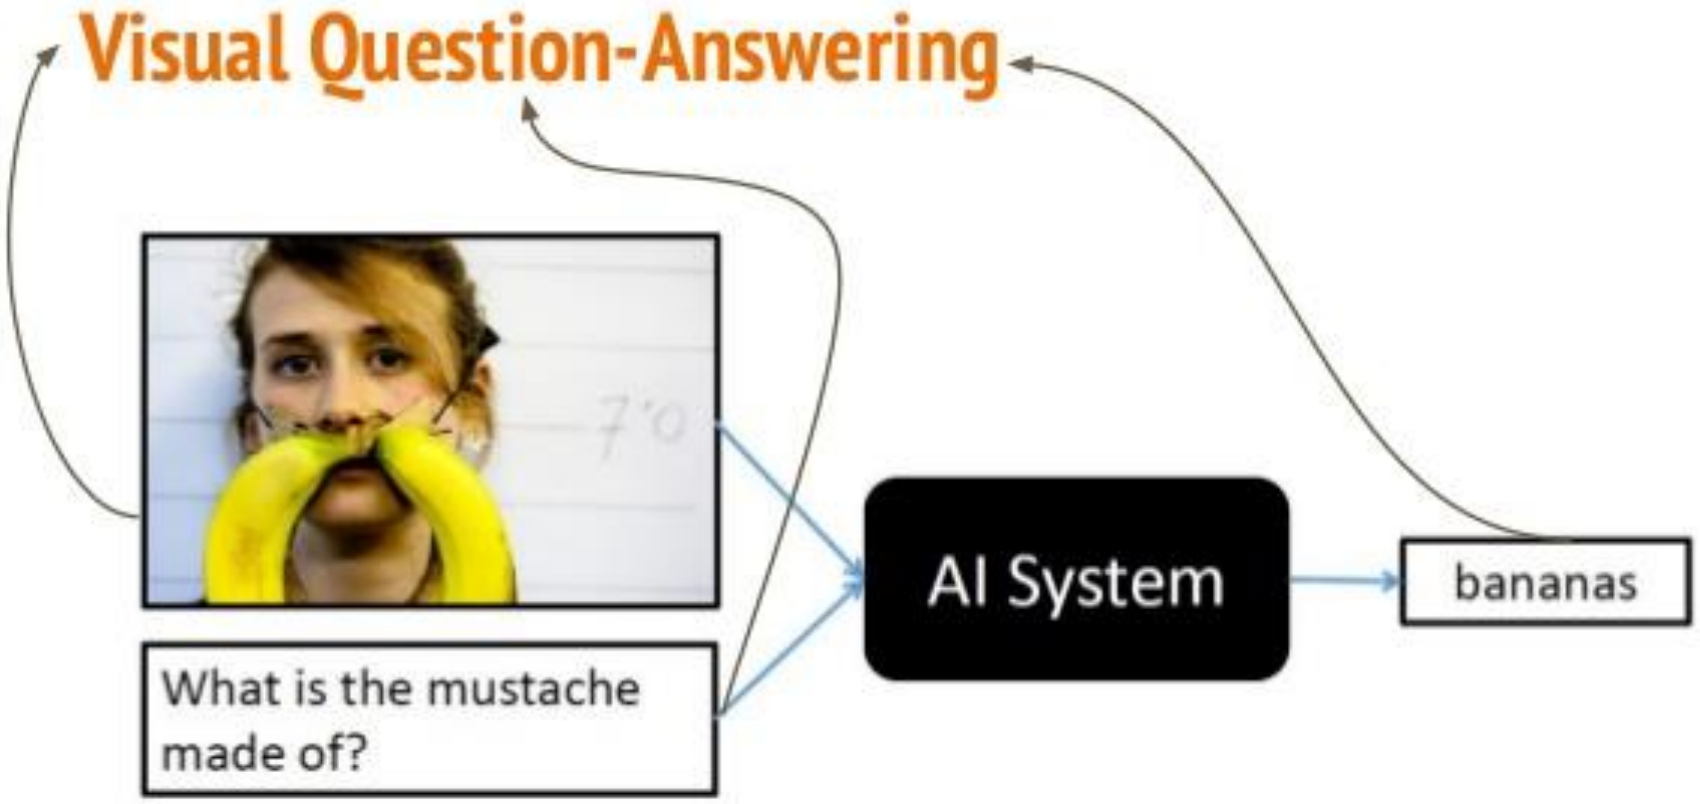

Antol, S., Agrawal, A., Lu, J., Mitchell, M., Batra, D., Lawrence Zitnick, C., & Parikh, D. (2015). Vqa: Visual question answering. In Proceedings of the IEEE International Conference on Computer Vision (pp. 2425-2433).

Self-Driving Cars

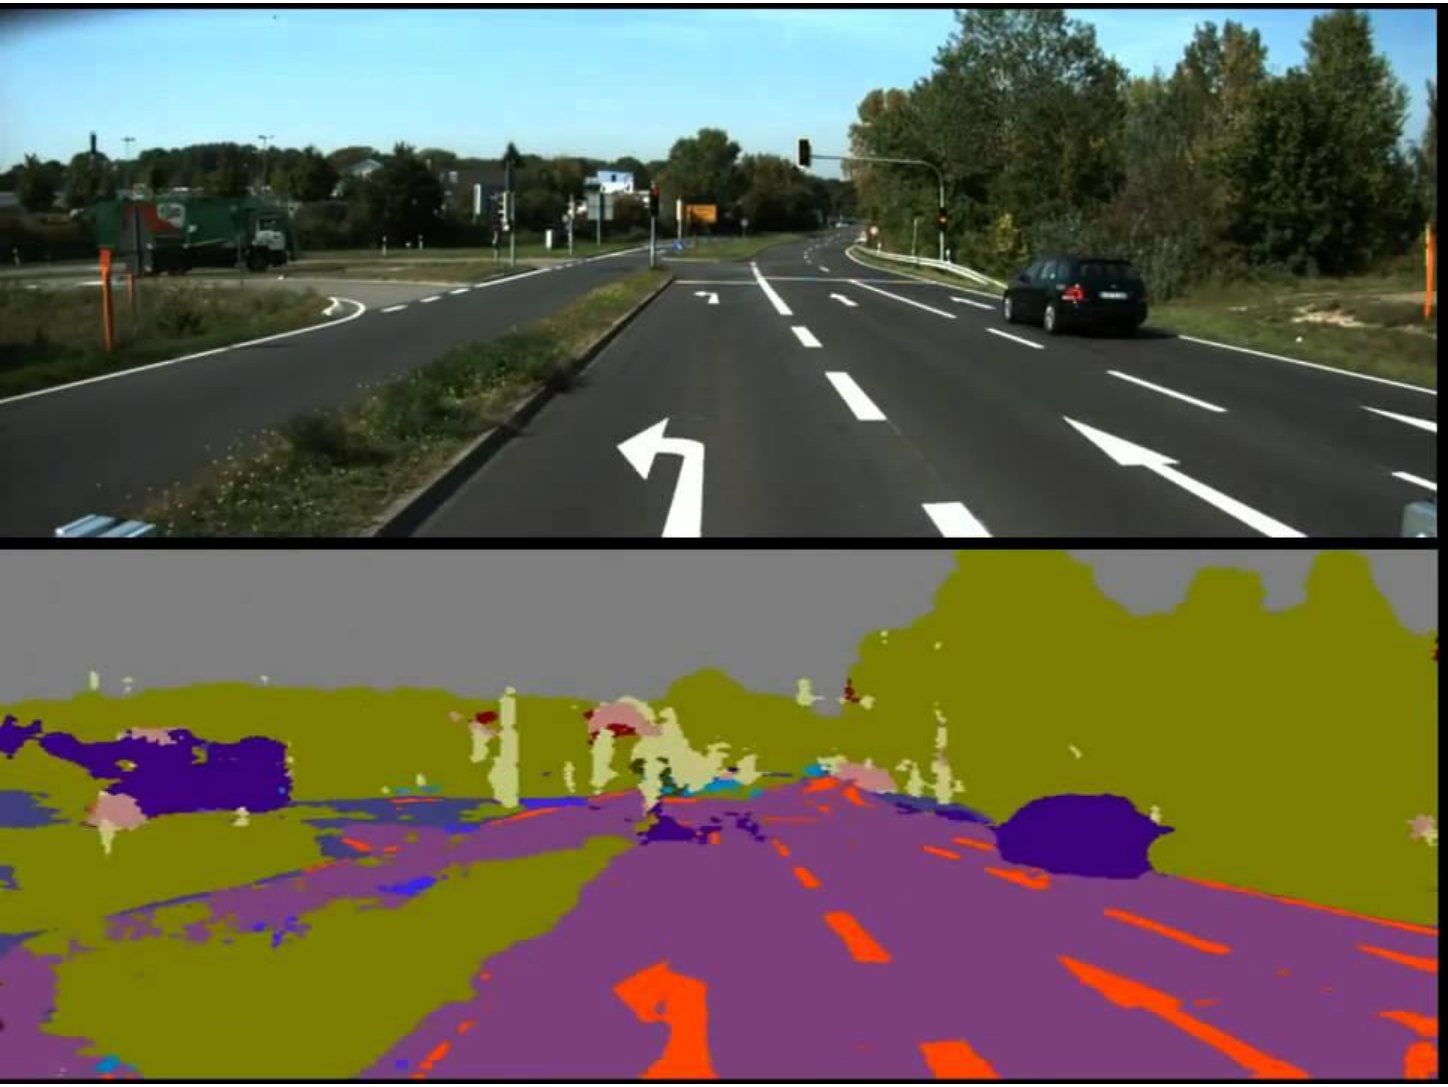

Style Transfer

https://medium.com/element-ai-research-lab/stabilizing-neural-style-transfer-for-video-62675e203e42