

MACHINE LEARNING

WHAT IS ML?

- A subfield of Artificial Intelligence
- On developing algorithms or systems for automated prediction and classification
- Examples:
 - Robotics, visual perception
 - Recommendation engine
 - Stock trading
 - Medical image analysis
 - Fraud: credit card, mortgage
 - and many others ...

WHAT IS ML?

Image



[201, 185, 125, ..., 233]

Algorithms

Everyday operations

‘+’, ‘-’, ‘*’, ‘/’

Classification Task

Classes:
82% cat
15% dog
2% hat
1% mug

WHAT IS ML?



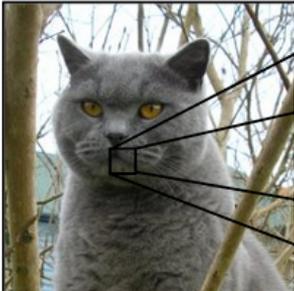
08	02	22	97	38	15	00	40	00	75	04	05	07	78	52	12	50	77	81	80
49	49	99	40	17	81	18	57	60	87	17	40	98	43	69	48	04	56	62	00
81	49	31	73	55	79	14	29	93	71	40	67	58	88	30	03	49	13	36	65
52	70	95	23	04	60	11	42	69	21	68	56	01	32	56	71	37	02	36	91
22	31	16	71	51	67	83	89	41	92	36	54	22	40	40	28	66	33	13	80
24	47	38	60	99	03	45	02	44	75	33	53	78	36	84	20	35	17	12	50
32	98	81	28	64	23	67	10	26	38	40	67	59	54	70	66	18	38	64	70
67	26	20	68	02	62	12	20	95	63	94	39	63	08	40	91	66	49	94	21
24	55	58	05	66	73	99	26	97	17	78	78	96	83	14	88	34	89	63	72
21	36	23	09	75	00	76	44	20	45	35	14	00	61	33	97	34	31	33	95
78	17	53	28	22	75	31	67	15	94	03	80	04	62	16	14	09	53	56	92
16	39	05	42	96	35	31	47	55	58	88	24	00	17	54	24	36	29	85	57
86	56	00	48	35	71	89	07	05	44	44	37	44	60	21	58	51	54	17	58
19	80	81	68	05	94	47	69	28	73	92	13	86	52	17	77	04	89	55	40
04	52	08	83	97	35	99	16	07	97	57	32	16	26	26	79	33	27	98	66
55	86	68	87	57	62	20	72	03	46	33	67	46	55	12	32	63	93	53	69
04	42	16	73	38	85	39	11	24	94	72	18	08	46	29	32	40	62	76	36
20	69	36	41	72	30	23	88	34	82	89	69	82	67	59	85	74	04	36	16
20	73	35	29	78	31	90	01	74	31	49	71	48	88	81	16	23	57	05	54
01	70	54	71	83	51	54	69	16	92	33	48	61	43	52	01	89	29	67	48

What the computer sees

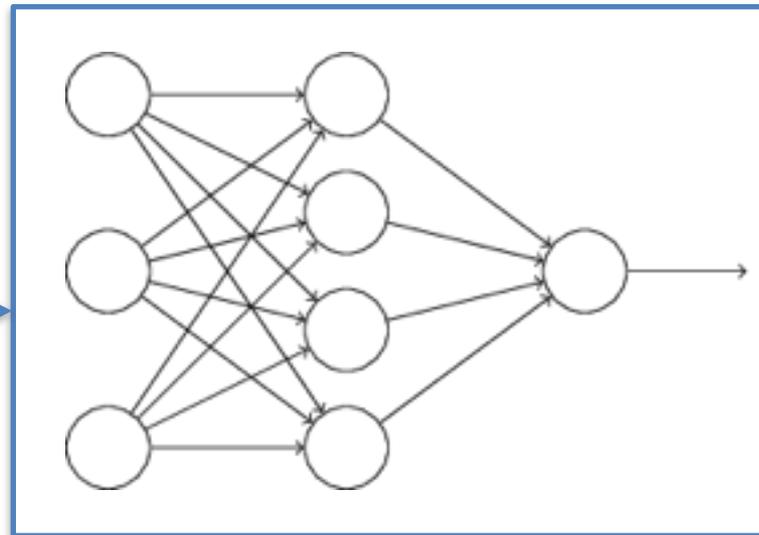
image classification → 82% cat
15% dog
2% hat
1% mug

NEURAL NETWORKS

Image



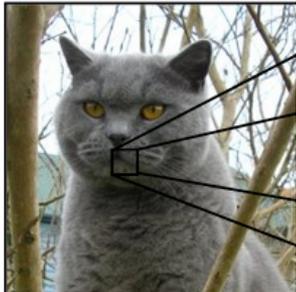
[201, 185, 125, ..., 233]



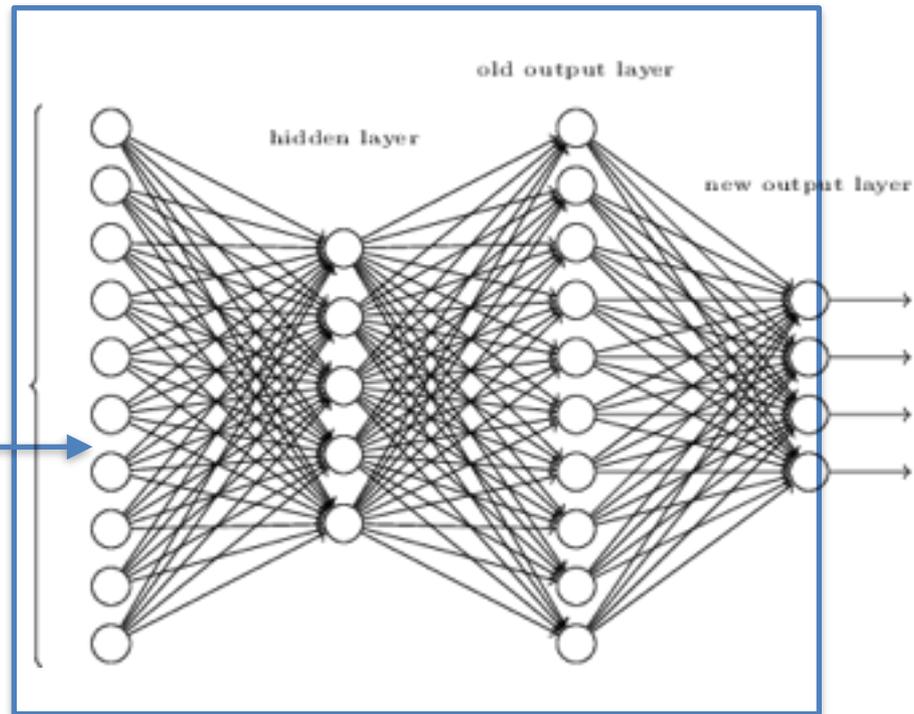
Classes:
82% cat
15% dog
2% hat
1% mug

DEEP NEURAL NETWORKS

Image



[201, 185, 125, ..., 233]



Classes:
82% cat
15% dog
2% hat
1% mug

NN STUDY LINKS

- <https://www.coursera.org/course/neuralnets>
- <http://cs231n.stanford.edu/>
- <http://karpathy.github.io/neuralnets/>
- <http://neuralnetworksanddeeplearning.com/chap1.html>
- <http://scikit-learn.org/stable/>

DEEP LEARNING

MIT
Technology
Review

10 BREAKTHROUGH TECHNOLOGIES 2013

Deep Learning

With massive amounts of computational power, machines can now recognize objects and translate speech in real time. Artificial intelligence is finally getting smart. →

Temporary Social Media

Messages that quickly self-destruct could enhance the privacy of online communications and make people freer to be spontaneous. →

Prenatal DNA Sequencing

Reading the DNA of fetuses will be the next frontier of the genomic revolution. But do you really want to know about the genetic problems or musical aptitude of your unborn child? →

Memory Implants

A maverick neuroscientist believes he has deciphered the code by which the brain

Smart Watches

Ultra-Efficient Solar Power

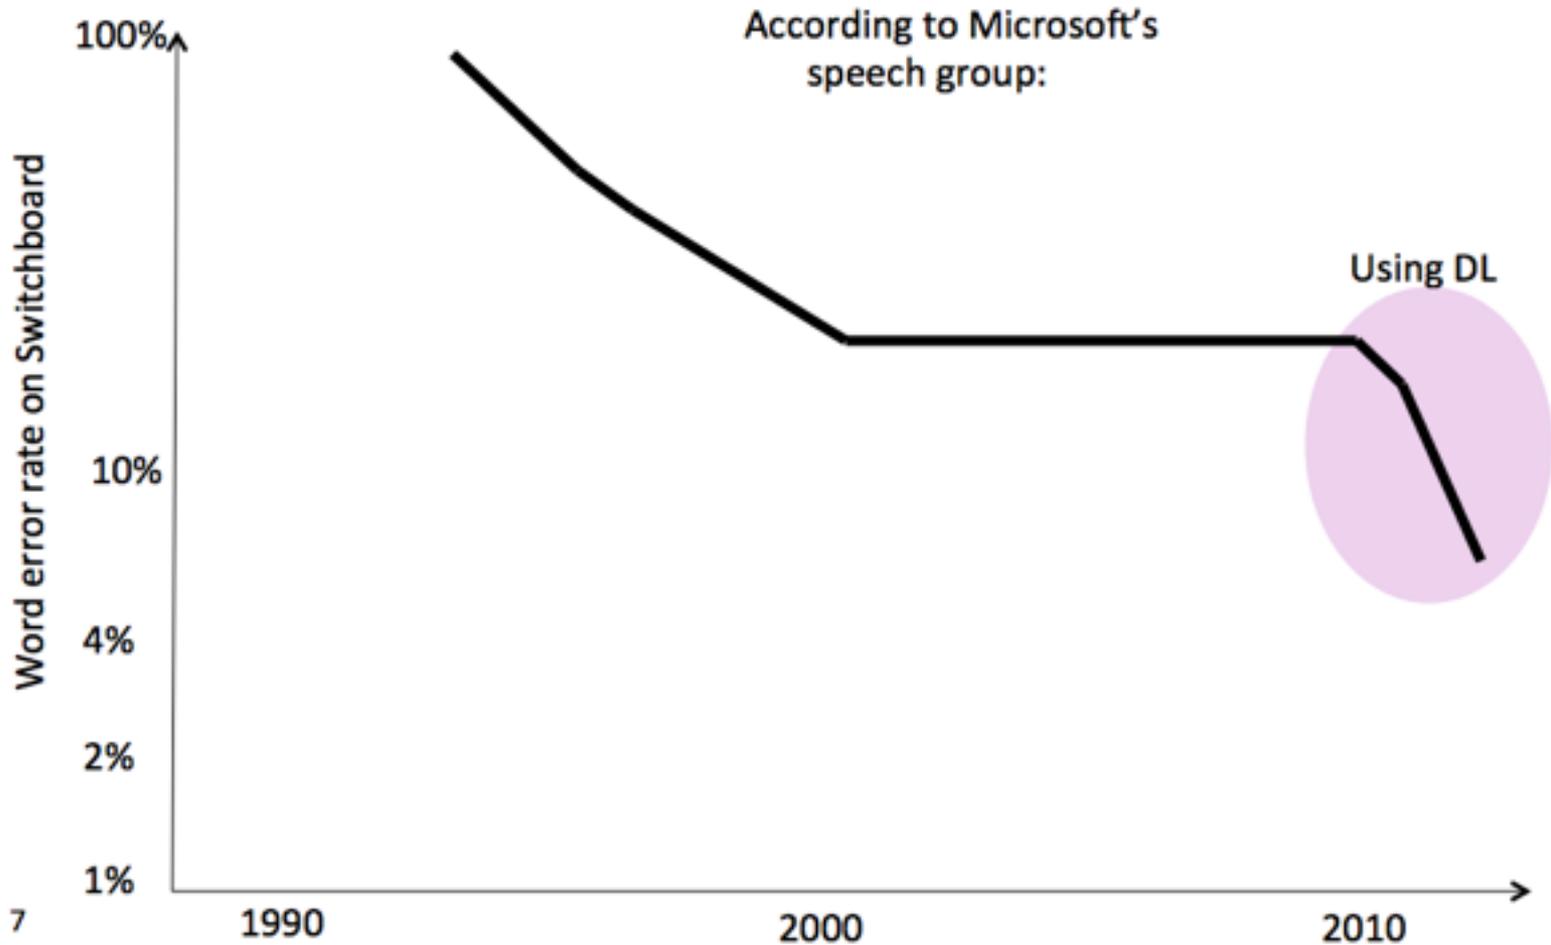
Doubling the efficiency of a solar cell would completely

Discovering Structure

Dramatic increase in both computational power and the amount of data available from web, video cameras, laboratory measurements.

- Develop statistical models that can discover underlying structure, semantic relations, constraints, or invariances from data.
- Robust, adaptive models that can deal with missing measurements, nonstationary distributions, multimodal data.
- Deep Learning can also discover structure at multiple features levels: from low level pixels to high level semantic features

Speech Recognition is Improved



NETFLIX

The Netflix Tech Blog



Netflix uses:

- Restricted Boltzmann machines
- Probabilistic Matrix Factorization

- From their blog:

To put these algorithms to use, we had to work to overcome some limitations, for instance that they were built to handle 100 million ratings, instead of the more than 5 billion that we have, and that they were not built to adapt as members added more ratings. But once we overcame those challenges, we put the two algorithms into production, where they are still used as part of our recommendation engine.

Key Computational Challenges

Scaling up our deep learning algorithms:

- Learning from billions of (unlabeled) data points
- Developing new parallel algorithms
- Scaling up Computation using clusters of GPUs and FPGAs

Building bigger models using more data improves performance of deep learning algorithms!

Deep Learning in Action

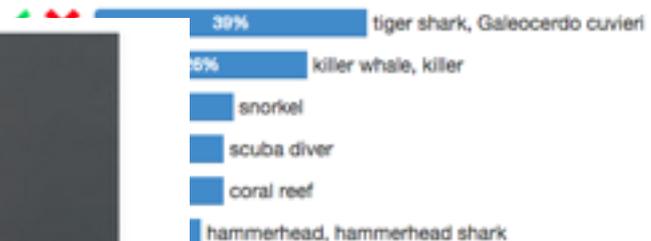
- Achieves state-of-the-art on many object recognition tasks!
Try it at deeplearning.cs.toronto.edu!



Possible tags:



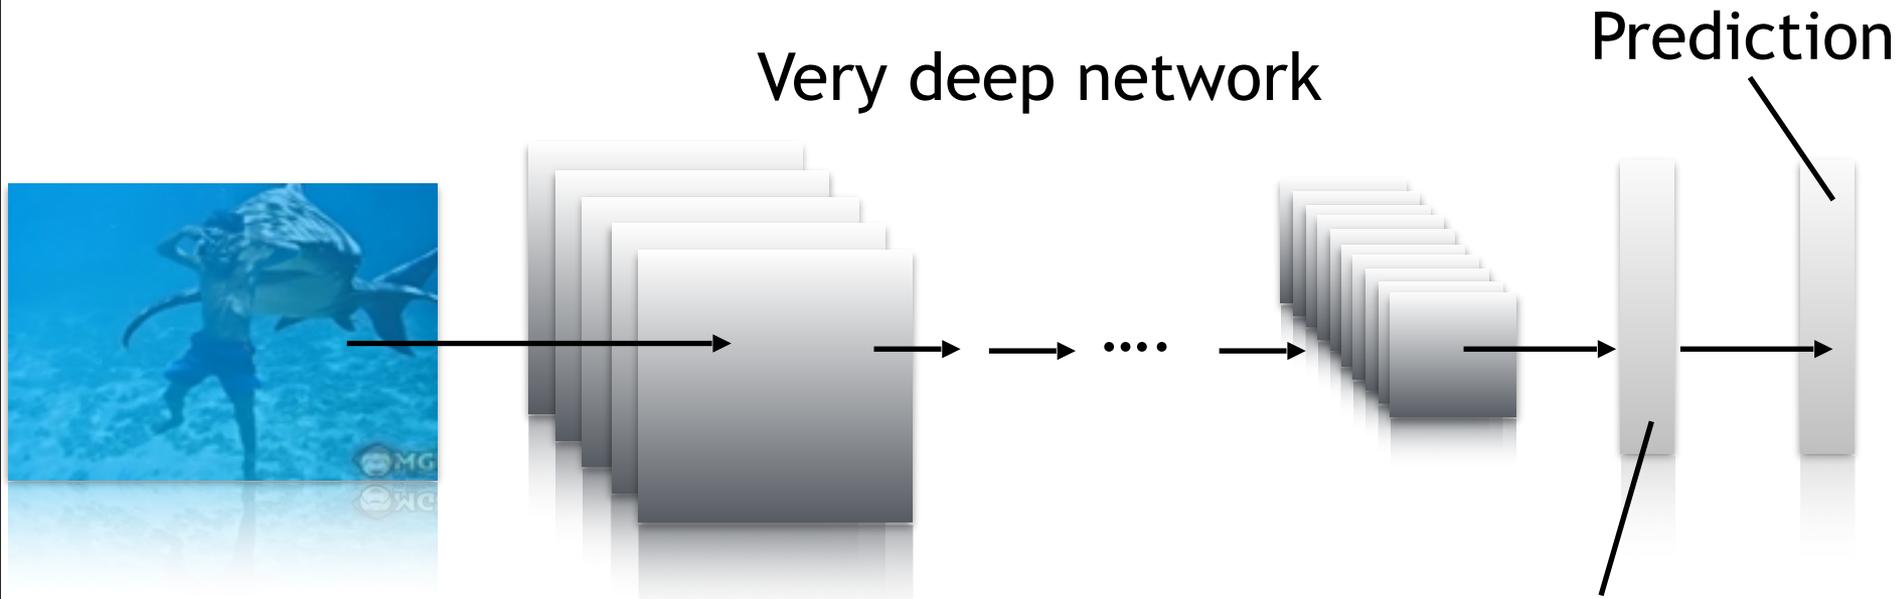
Possible tags:



Possible tags:



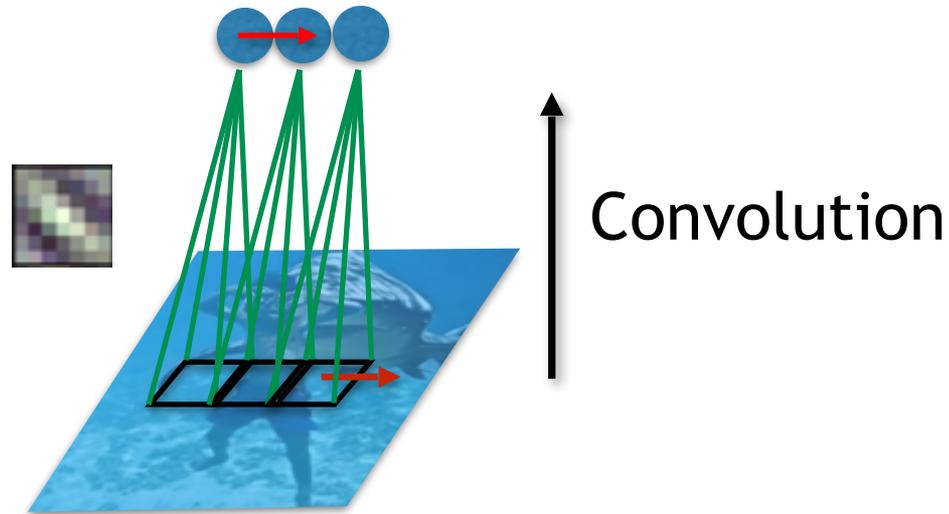
Deep Convolutional Nets



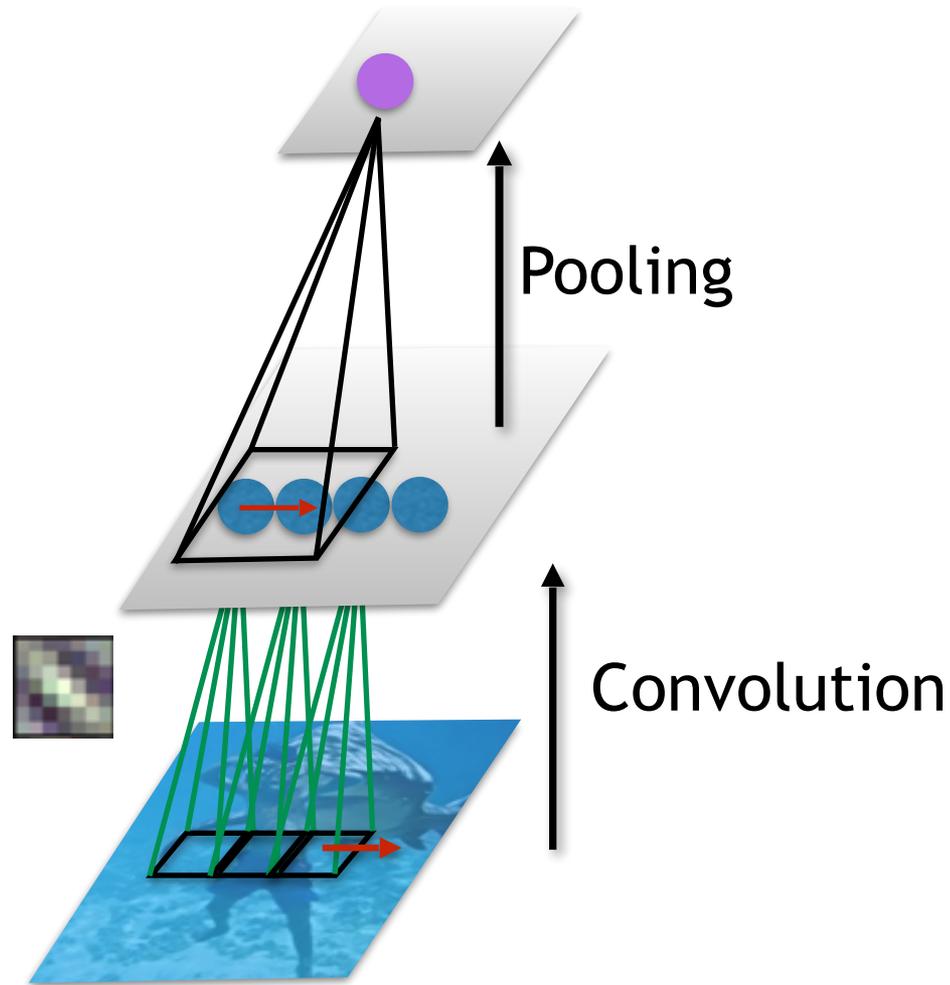
- Convolution
- Pooling
- Normalization
- Densely connected

High-level feature space

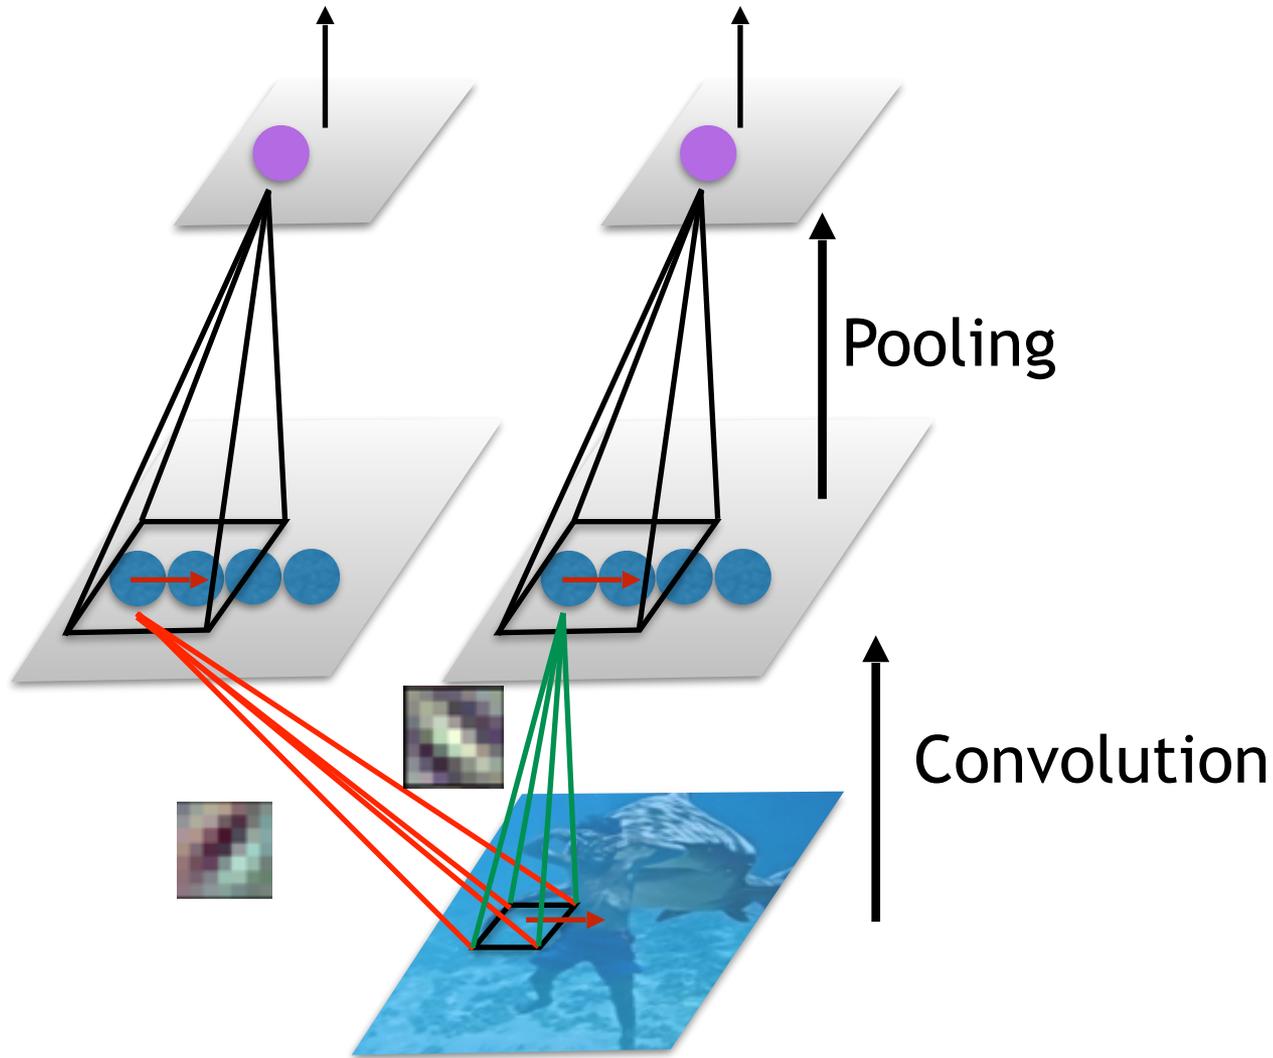
Deep Convolutional Nets



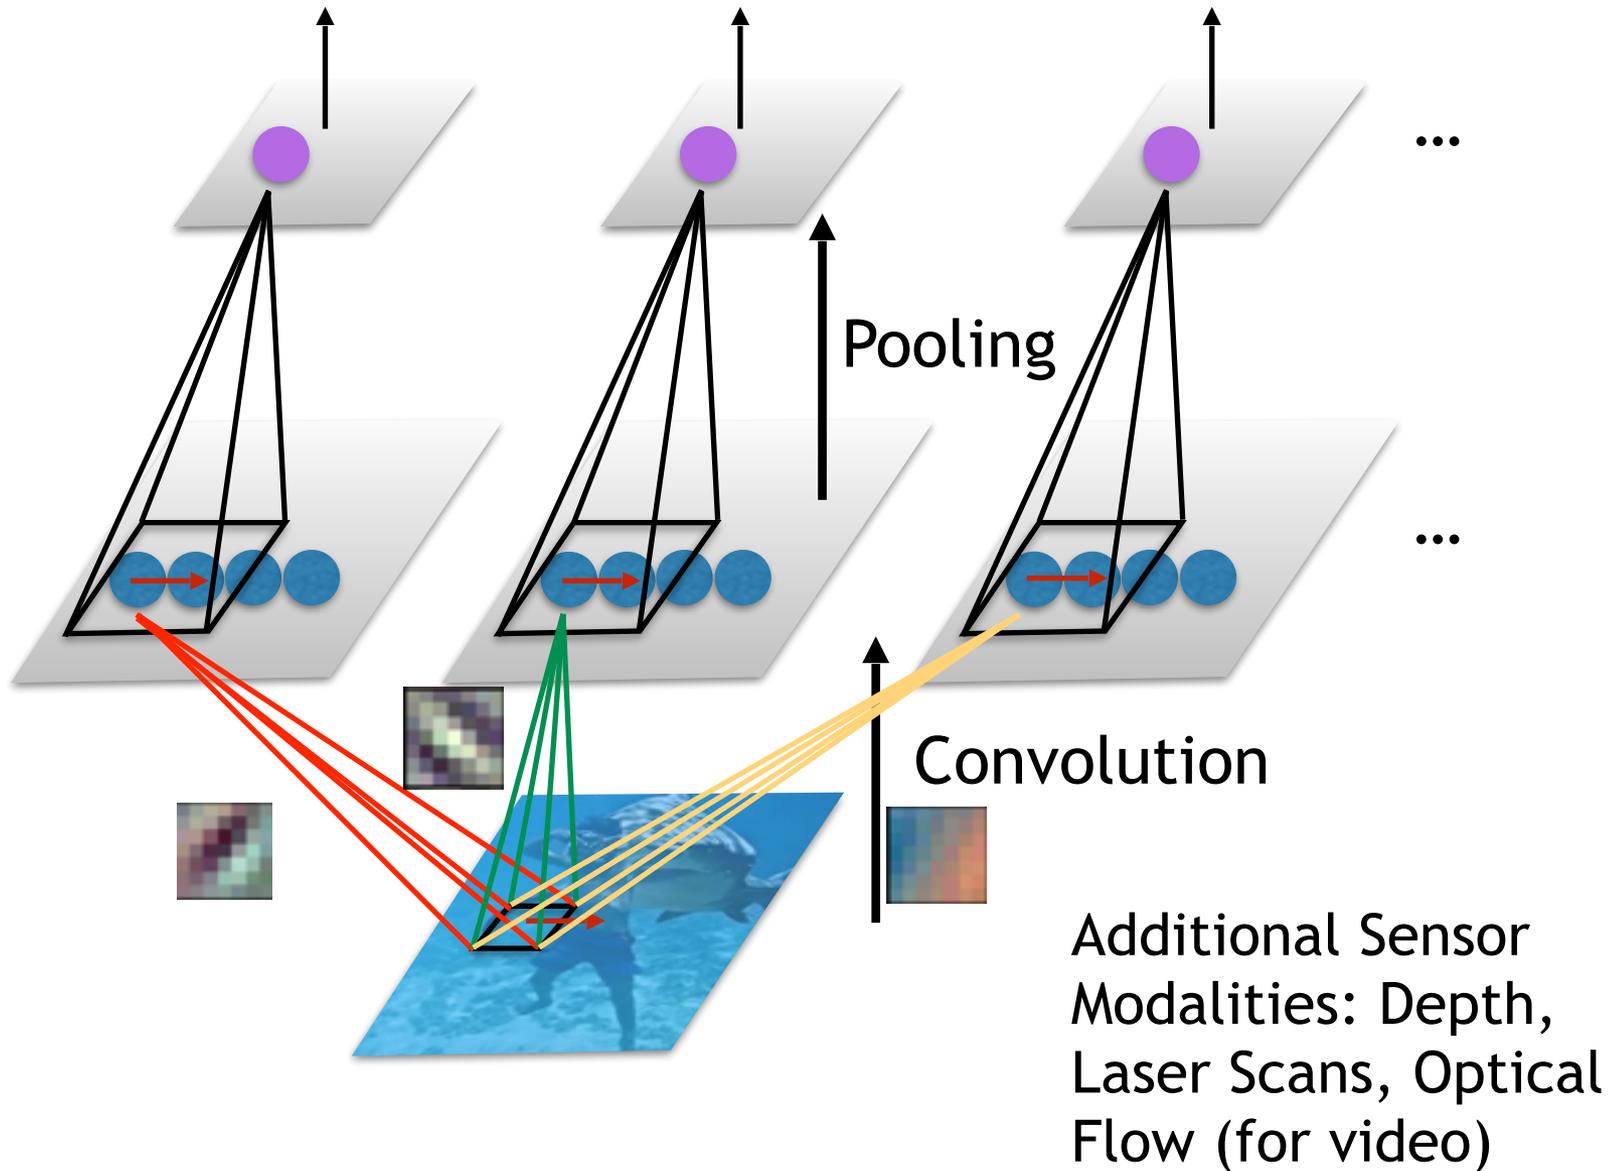
Deep Convolutional Nets



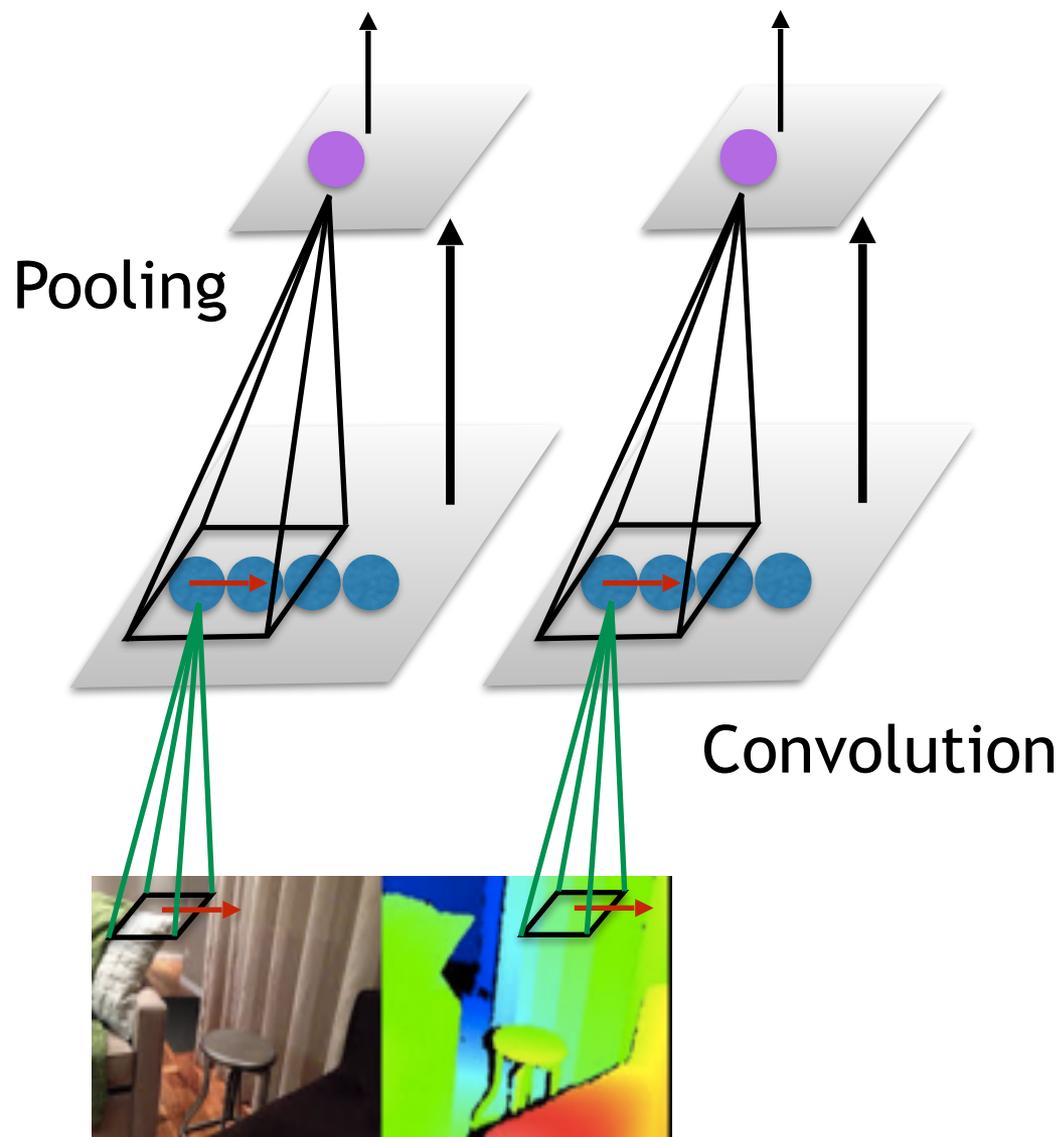
Deep Convolutional Nets



Deep Convolutional Nets



Multimodal Deep Convolutional Nets



Visual Search

Query Image



- 0.39 tiger shark
- 0.26 killer whale
- 0.10 snorkel
- 0.08 scuba diver
- 0.08 coral reef



Visual Search

Query Image



0.93 agaric

0.07 mushroom



Learning Structured Outputs

- More challenging problem.
- How can we generate complete descriptions of images?

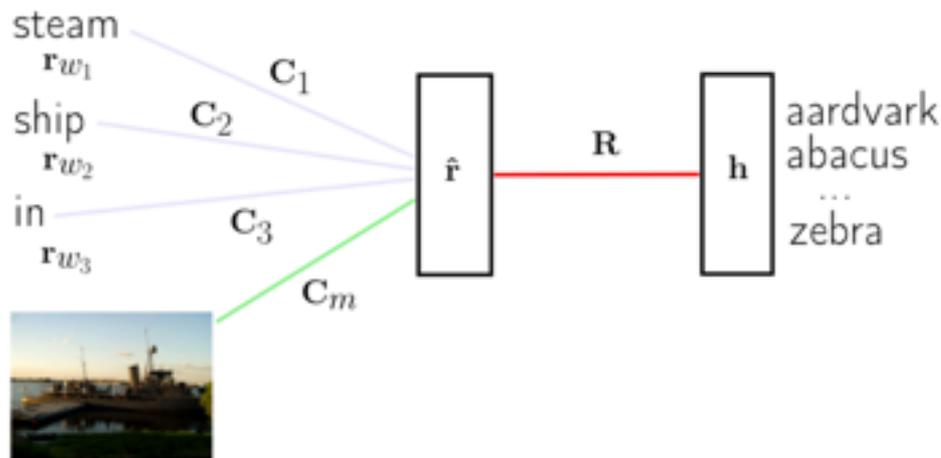
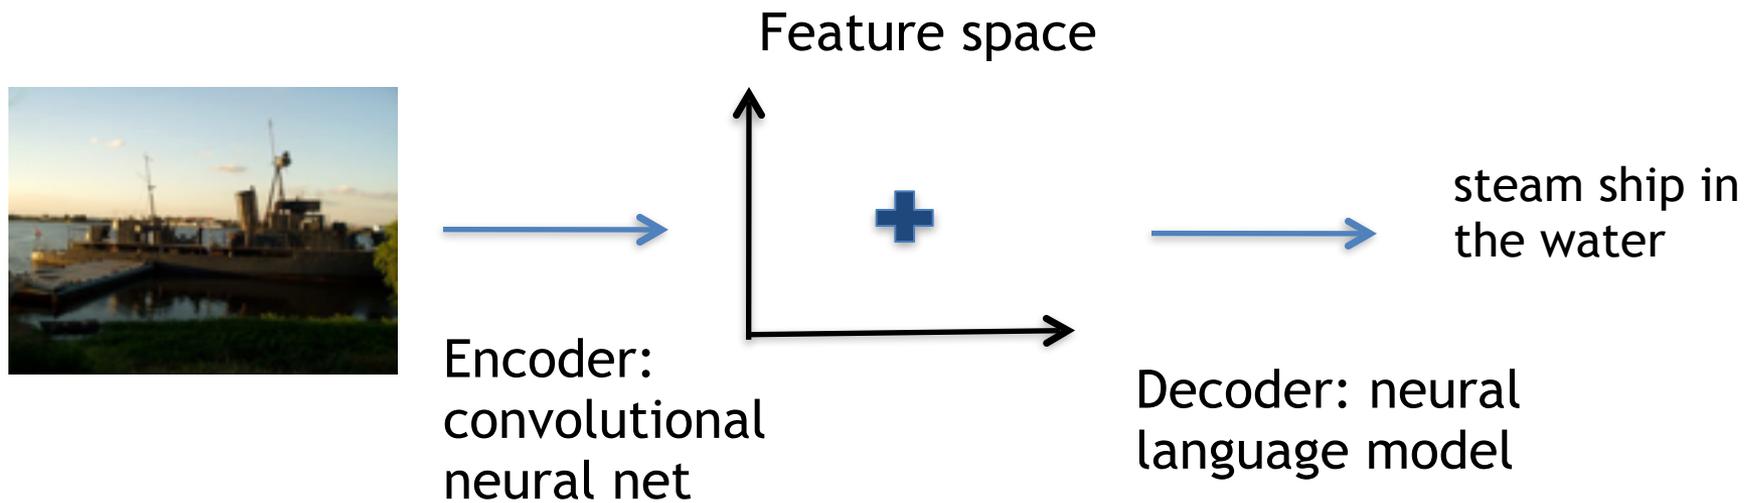
Input



Output

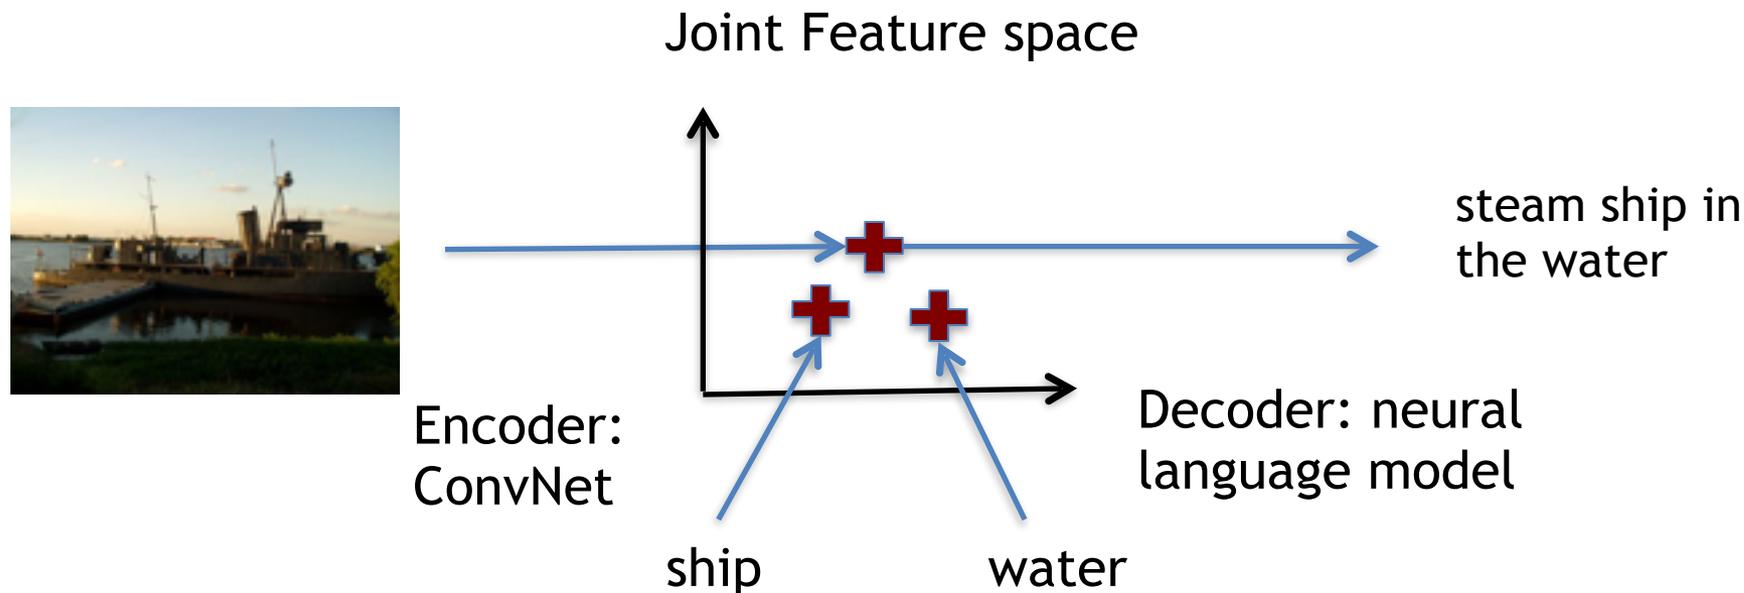
A man skiing down the snow covered mountain with a dark sky in the background.

An Image-Text Encoder-Decoder



Use the image features to
additively
bias the prediction of the next
word
representation

An Image-Text Encoder-Decoder



- Learn a joint embedding space of images and text:
 - Can condition on anything (images, words, phrases, etc)
 - Natural definition of a scoring function (inner products in the joint space)
 - Use a new language model that incorporates additional structure

Tagging and Retrieval



mosque, tower,
building,
cathedral,
dome, castle



ski, skiing,
skiers, skiers,
snowmobile

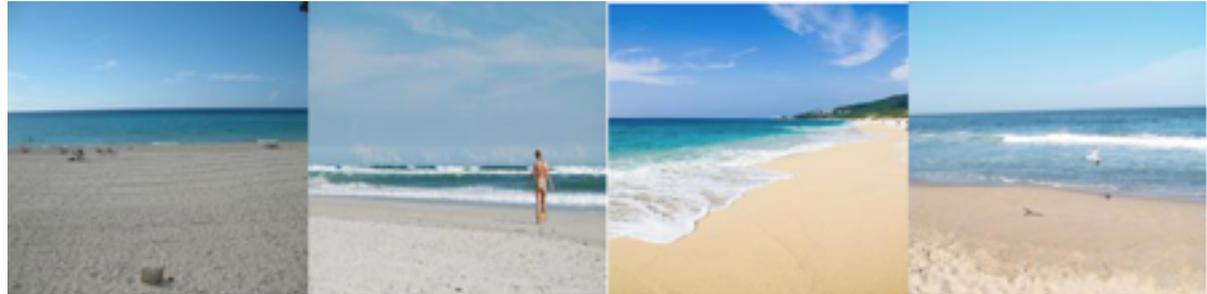


kitchen, stove,
oven,
refrigerator,
microwave

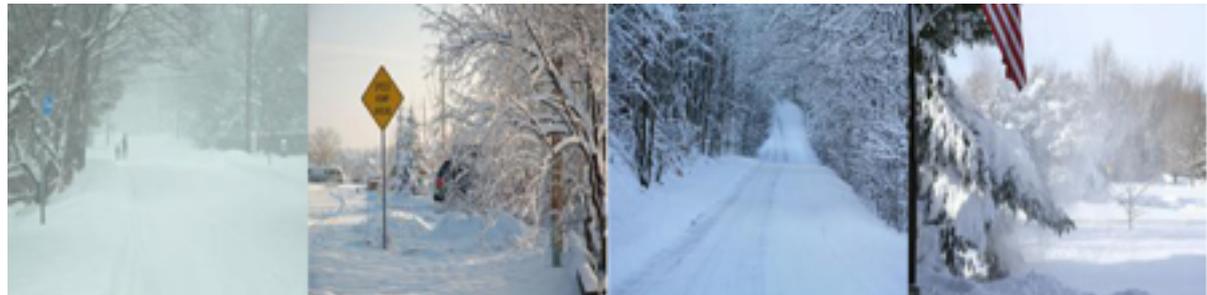


bowl, cup,
soup, cups,
coffee

beach



snow



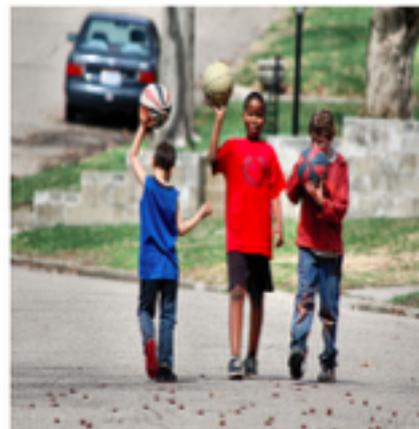
Caption Generation



a car is parked
in the middle
of nowhere .



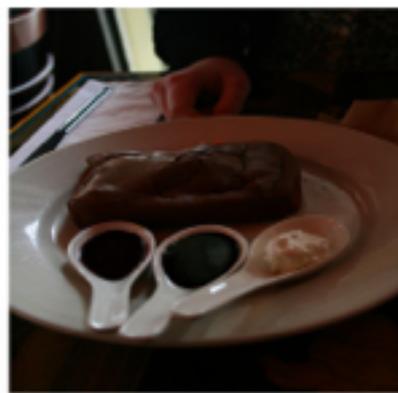
a ferry boat on
a marina with a
group of people .



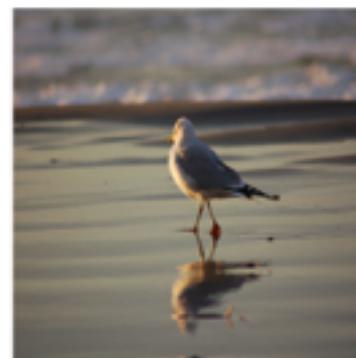
a little boy with
a bunch of friends
on the street .



there is a cat
sitting on a shelf .



a plate with a fork
and a piece of cake .



the two birds are
trying to be seen
in the water .

Understanding Images



cat, ferret, hamster,
weasel, puppy

cute, furry, cuddly, adorable,
naughty



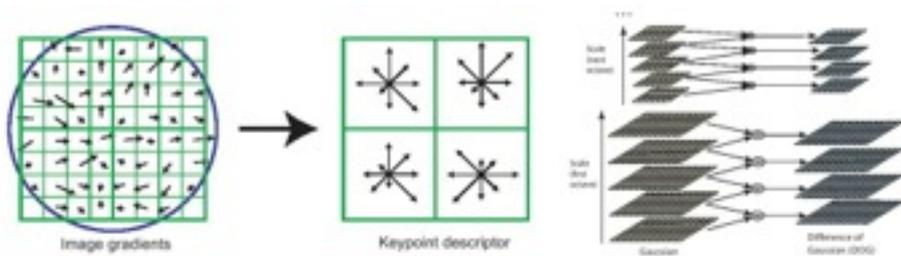
cup, bowl, coffee, soup,
cups

yummy, delicious, plastic,
foamy, savory

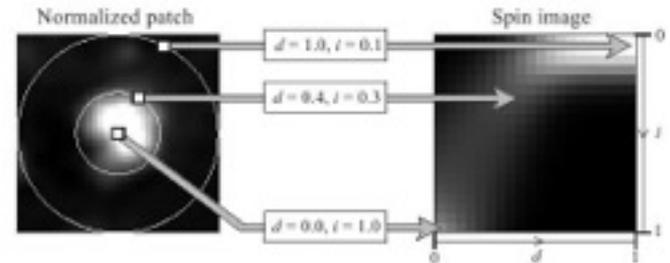
Checkout demo at

deeplearning.cs.toronto.edu

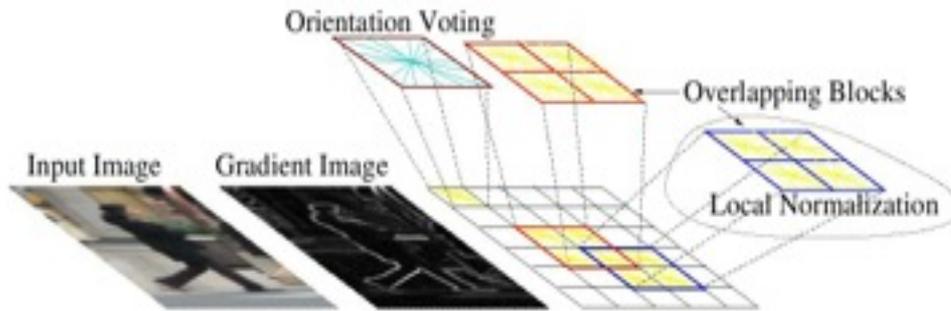
Computer Vision Features



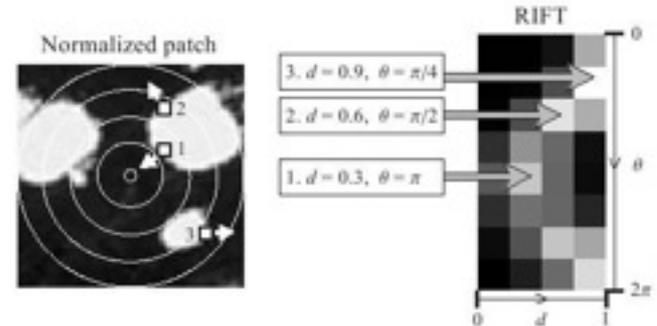
SIFT



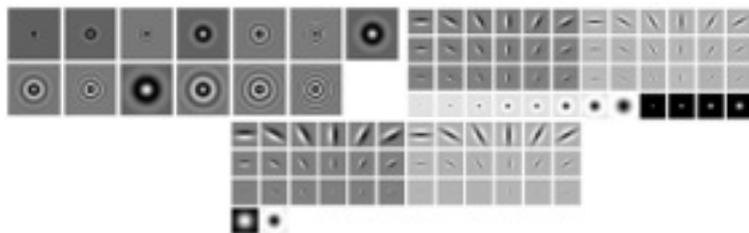
Spin image



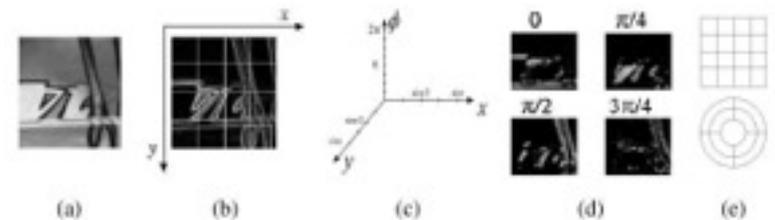
HoG



RIFT



Textons



GLOH

Computer Vision Features

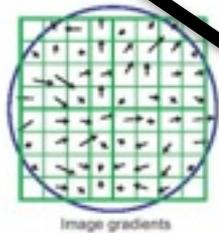
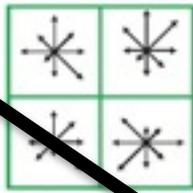
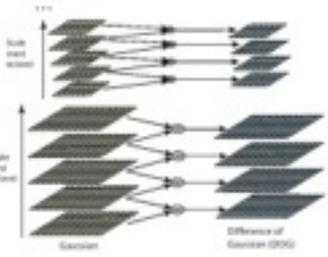


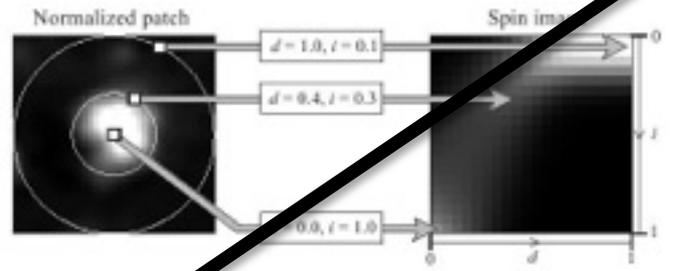
Image gradients



Keypoint detection



Gaussian Difference of Gaussian (DOG)



Normalized patch

Spin image

SIFT

Spin image

Deep Learning

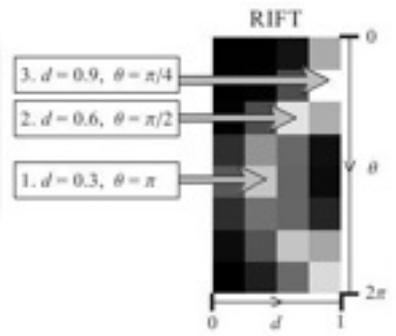


Input Image

Gradient Image

Oriented

HoG



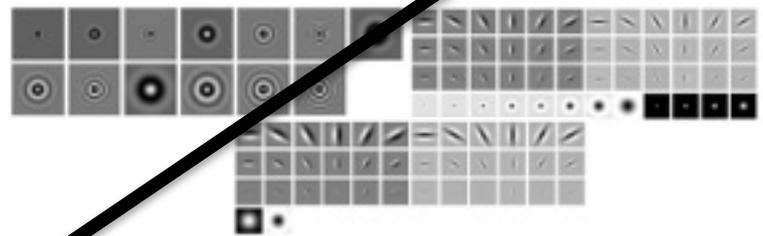
RIFT

3. $d = 0.9, \theta = \pi/4$

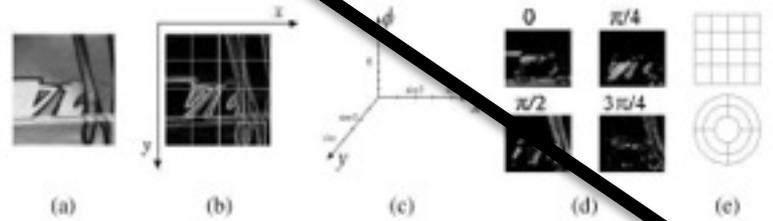
2. $d = 0.6, \theta = \pi/2$

1. $d = 0.3, \theta = \pi$

RIFT



Textons



(a)

(b)

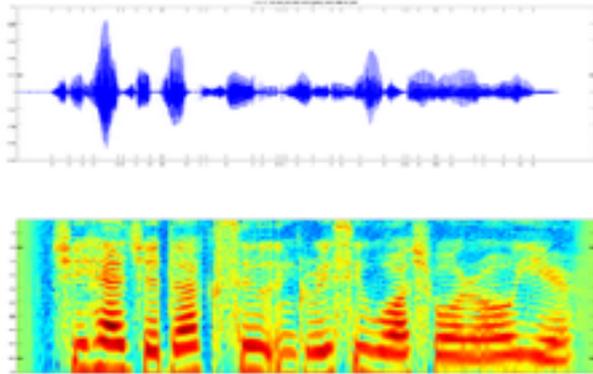
(c)

(d)

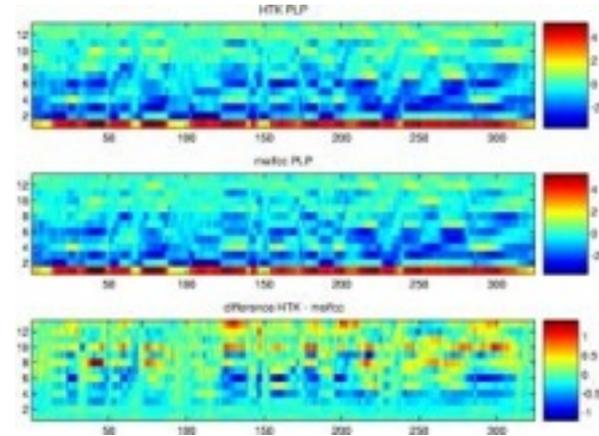
(e)

GLOH

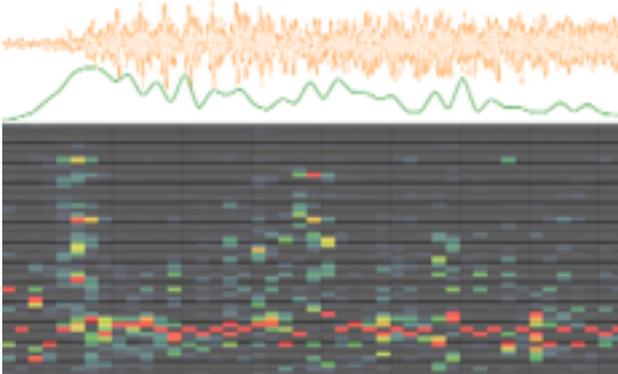
Audio Features



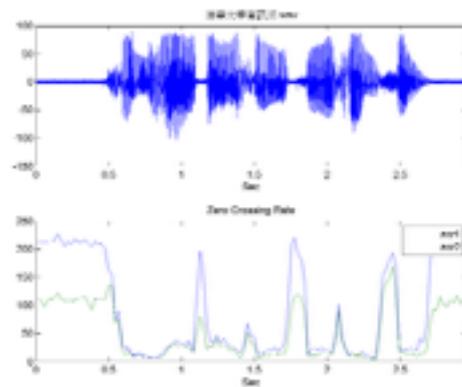
Spectrogram



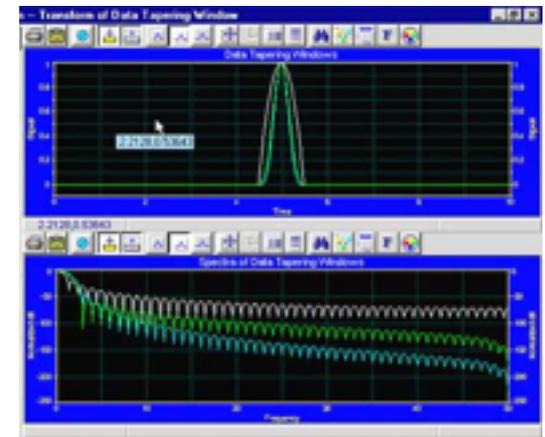
MFCC



Flux

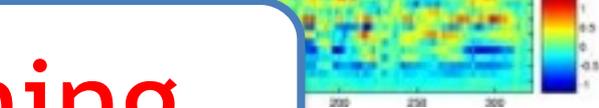
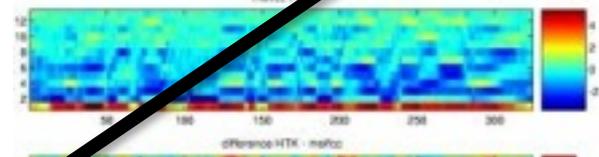
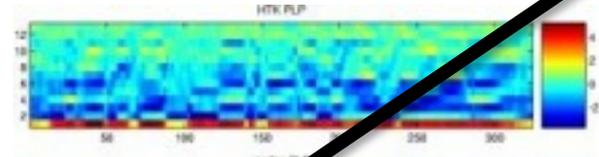
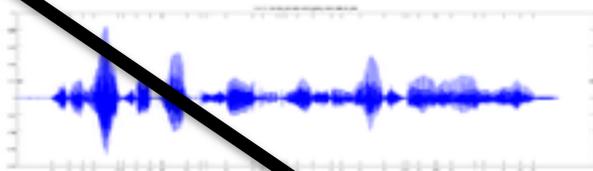


ZCR



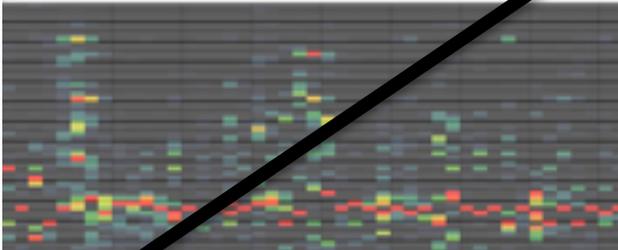
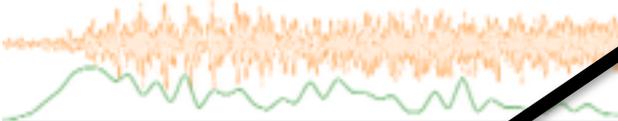
Rolloff

Audio Features

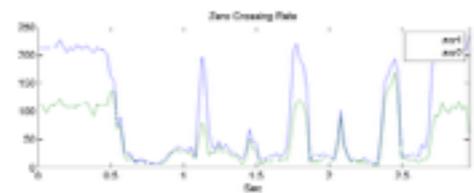
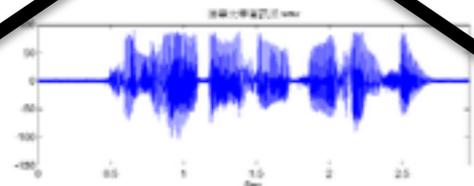


Deep Learning

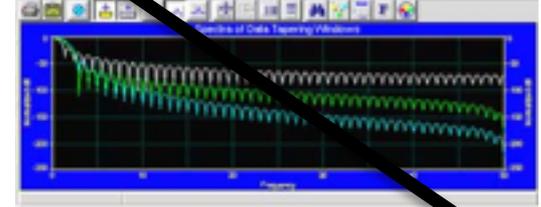
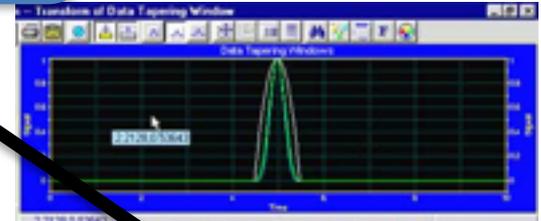
CC



Flux



ZCR



Rolloff

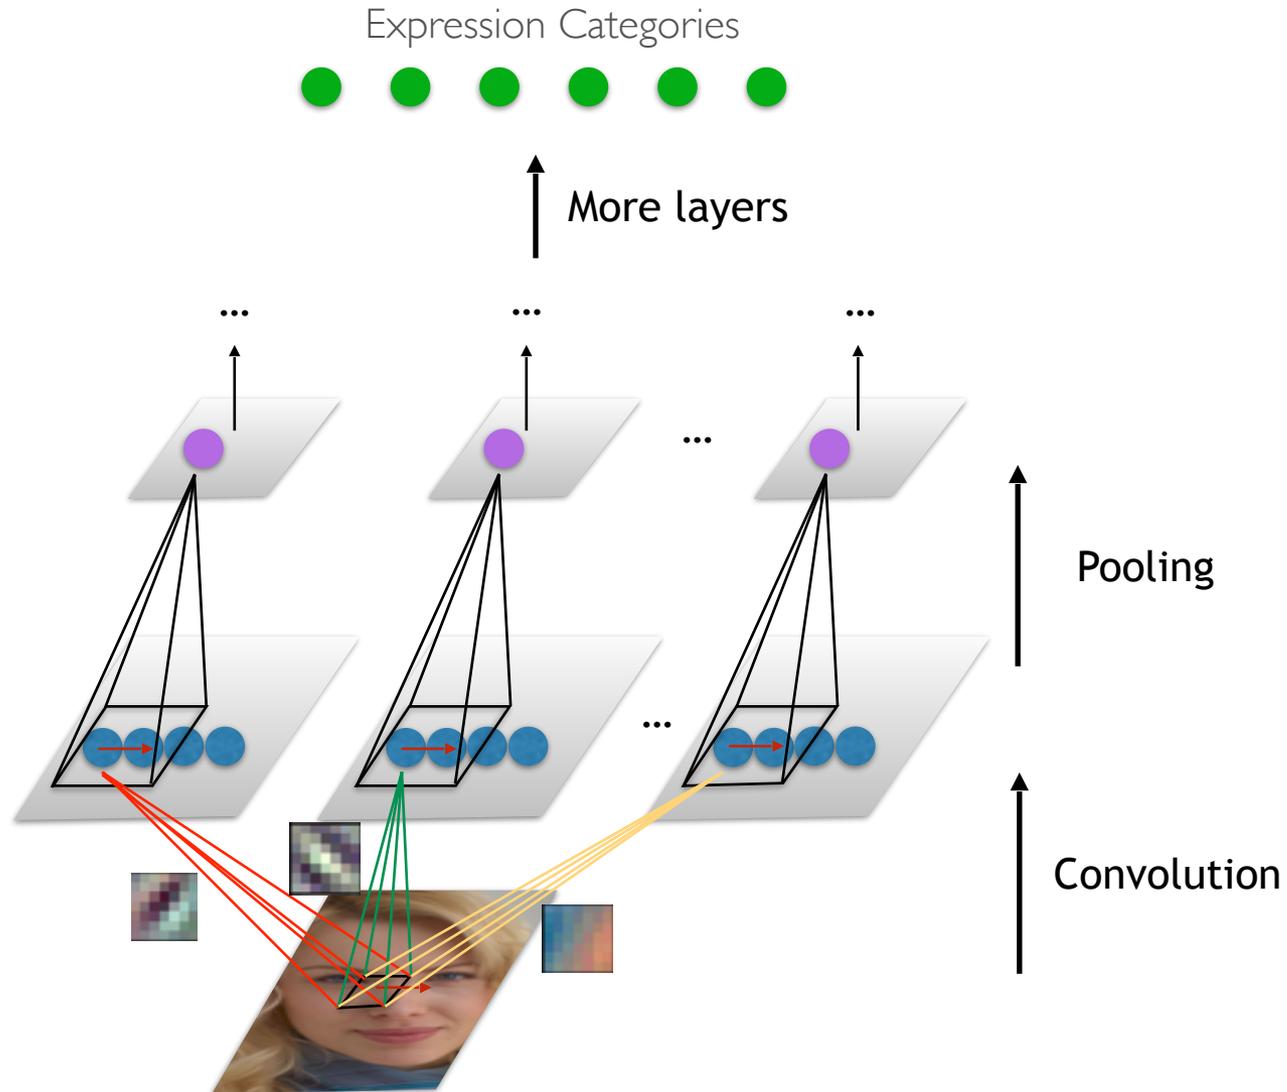
Expression Recognition



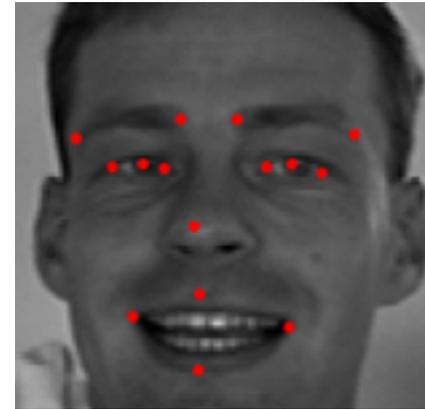
7 classes: neutral, happy, anger, sad surprised, disgust, and fear

- **Surpassed human judge performances by 2%**

Expression Recognition



Facial Keypoints



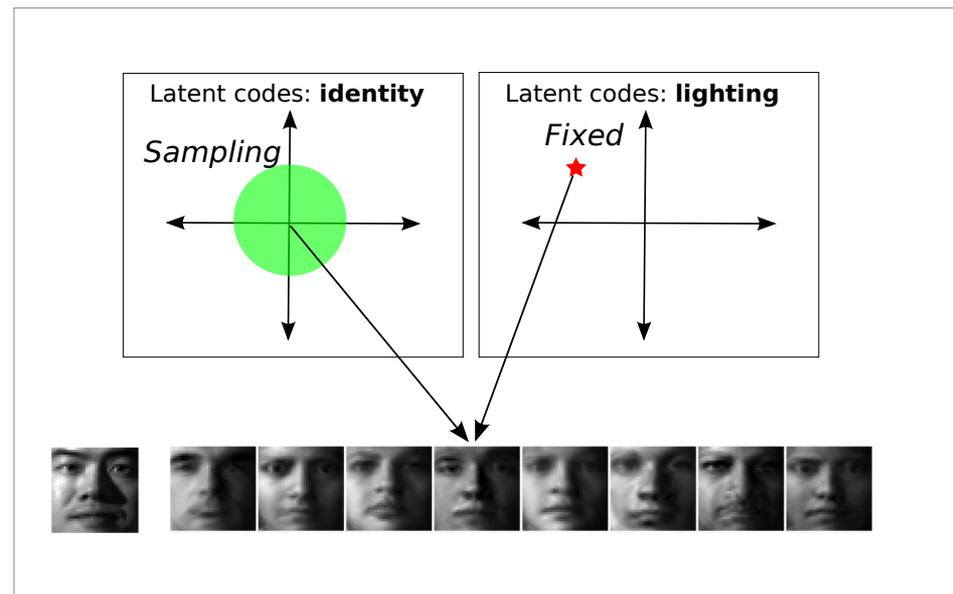
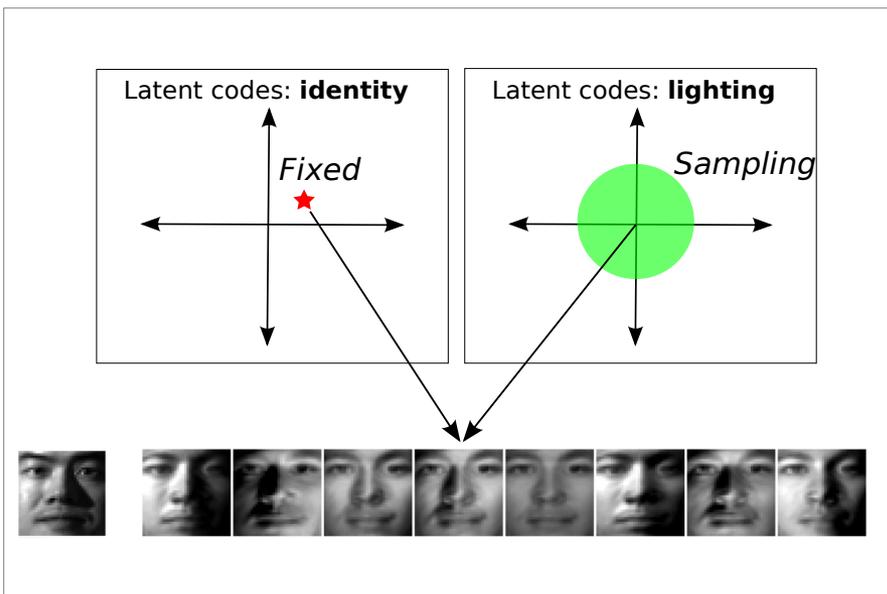
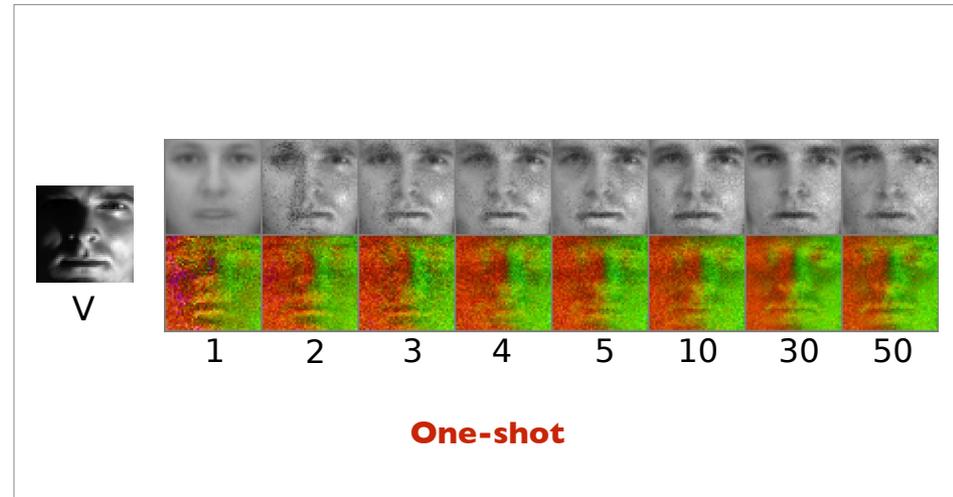
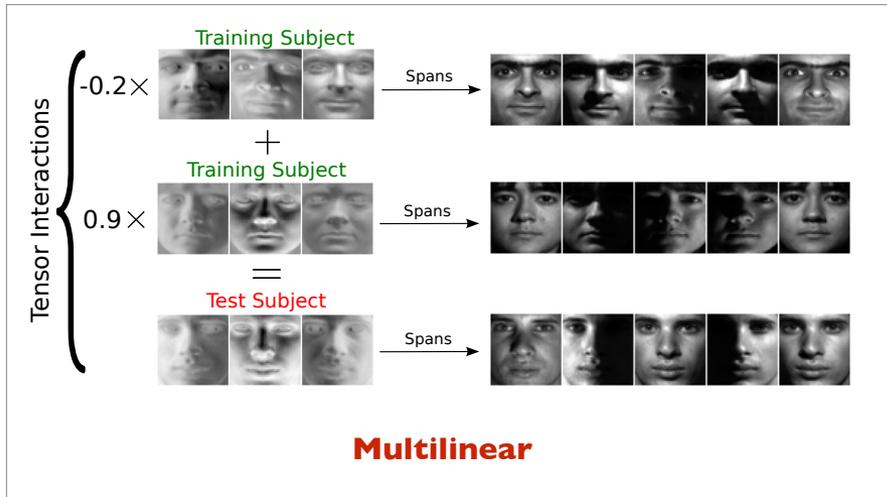
Facial Keypoints

Tuesday, May 7, 2013

**Kaggle Competition:
predict points locations**

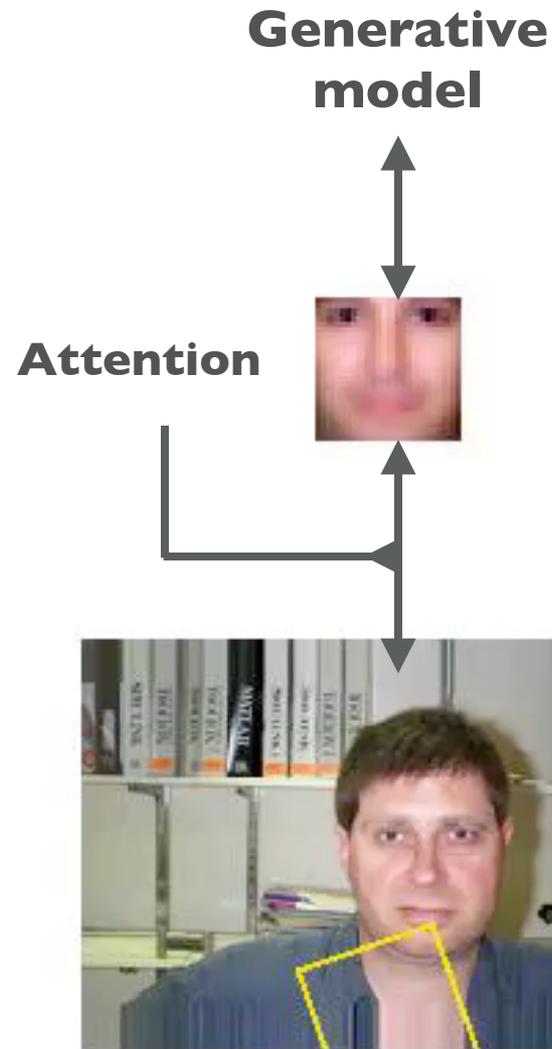
<https://www.kaggle.com/c/facial-keypoints-detection/>

Illumination Variations



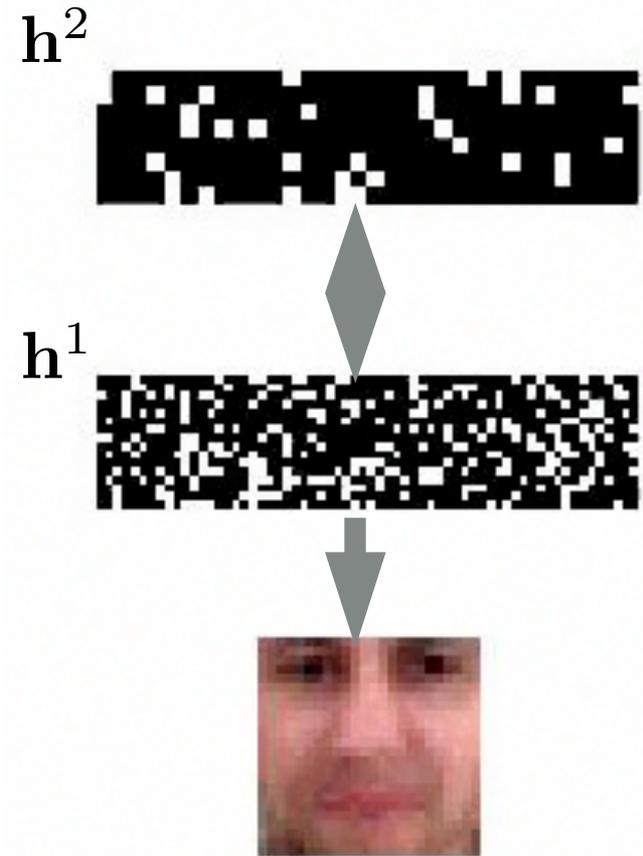
Localization using Attention

Introduce a probabilistic framework that uses visual attention to learn generative models of objects of interest

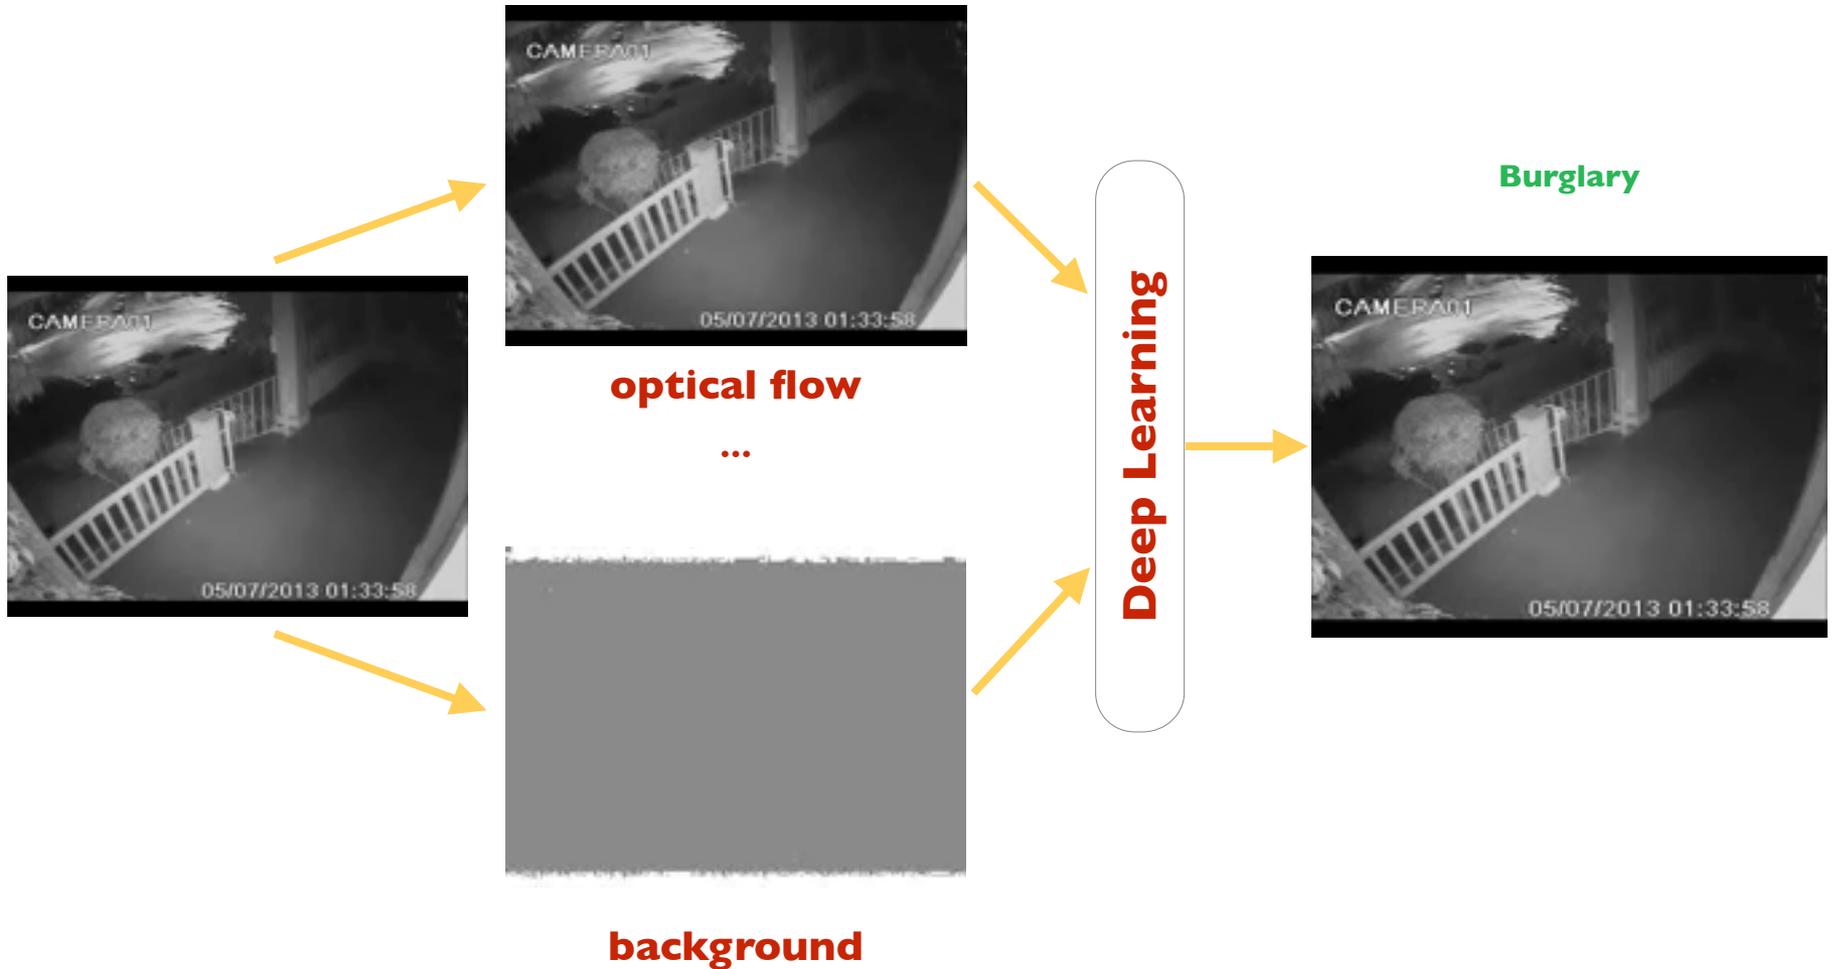


Deep Generative Face Model

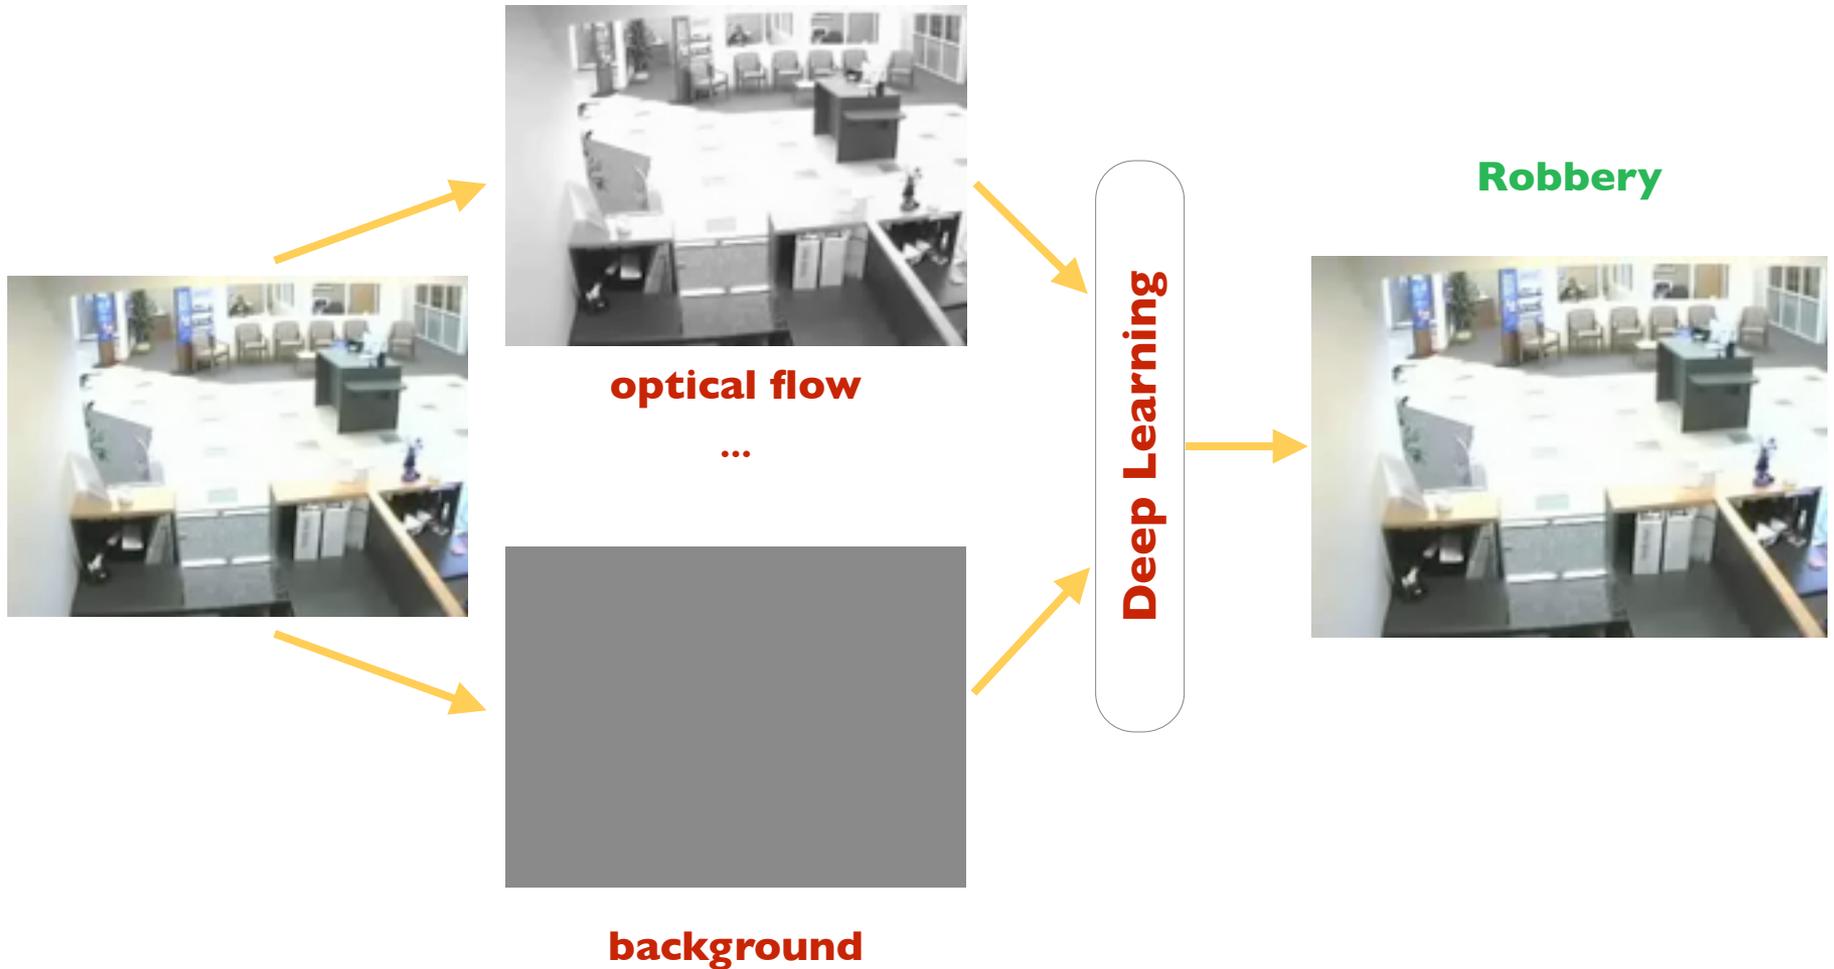
- Gaussian Deep Belief Net is a type of deep generative model
- Hybrid undirected-directed graphical model
- Good generative model of objects and faces



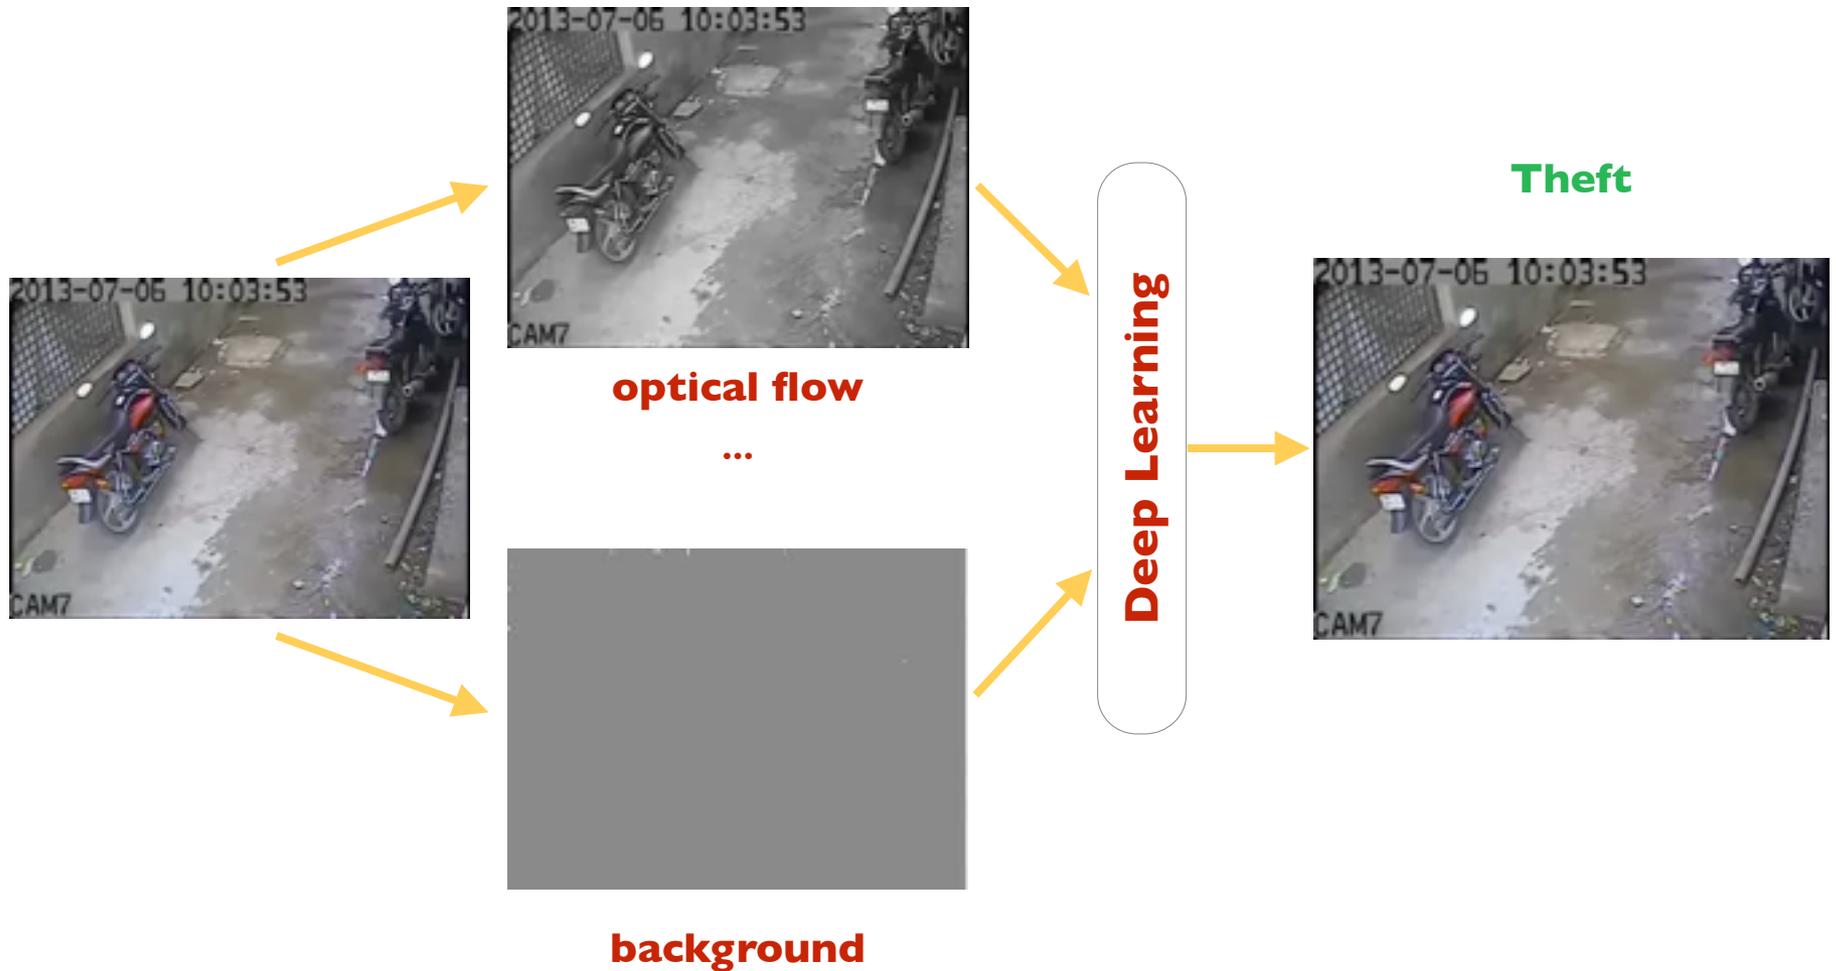
Human Activity Analytics



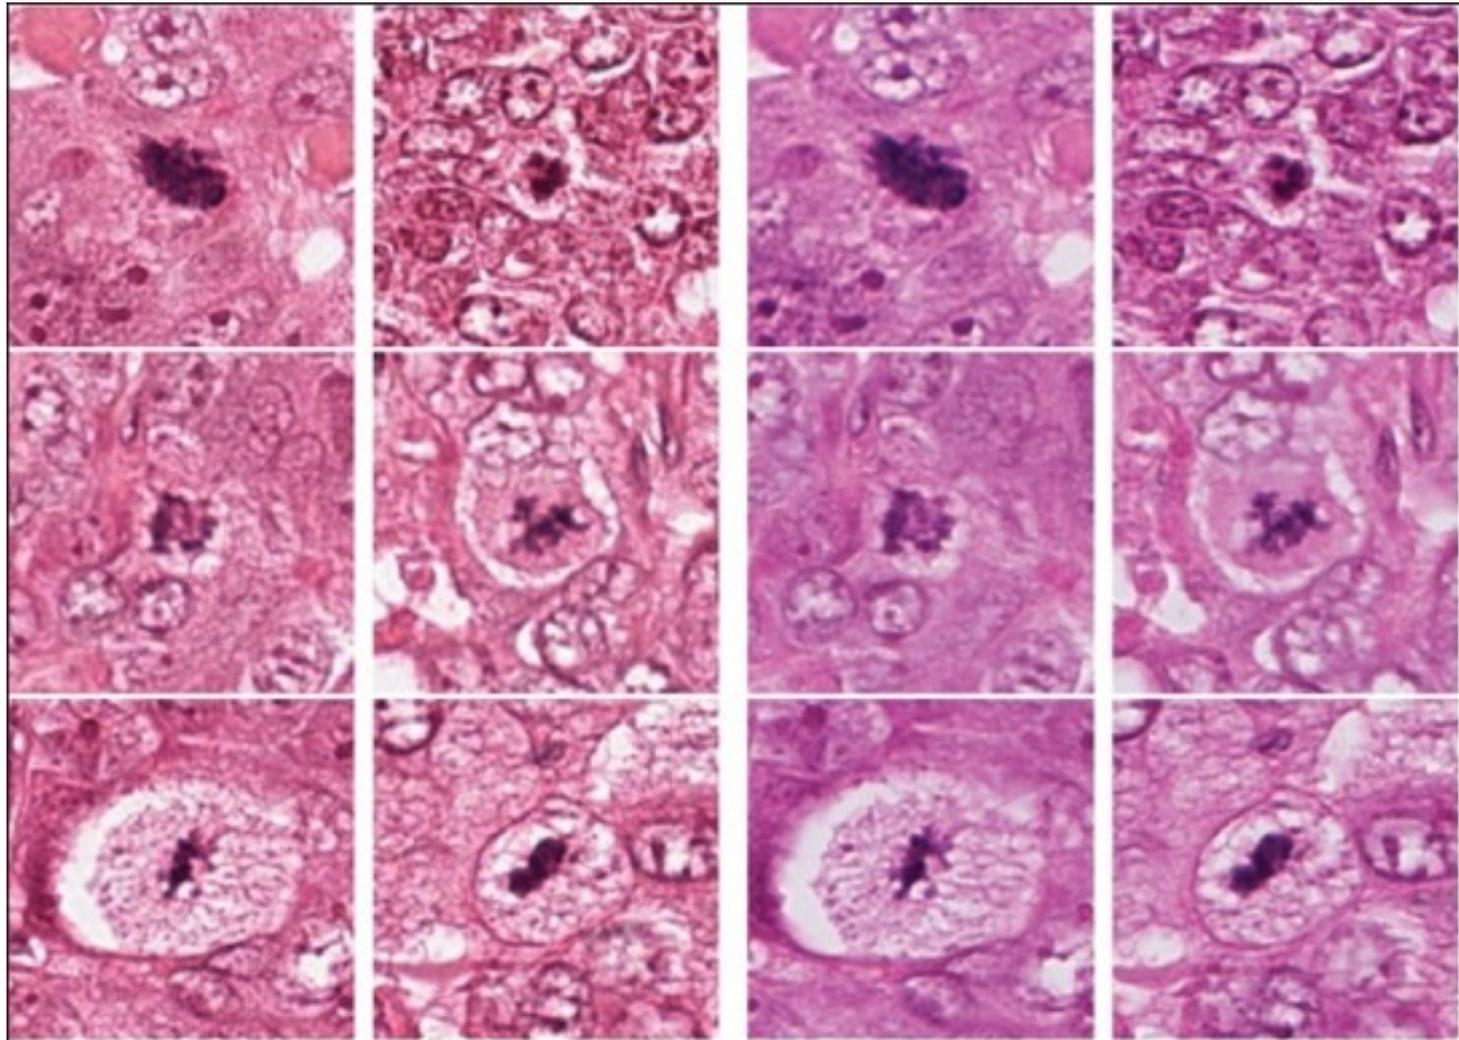
Human Activity Analytics



Human Activity Analytics



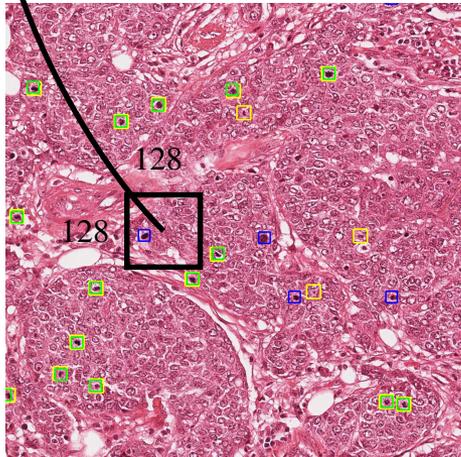
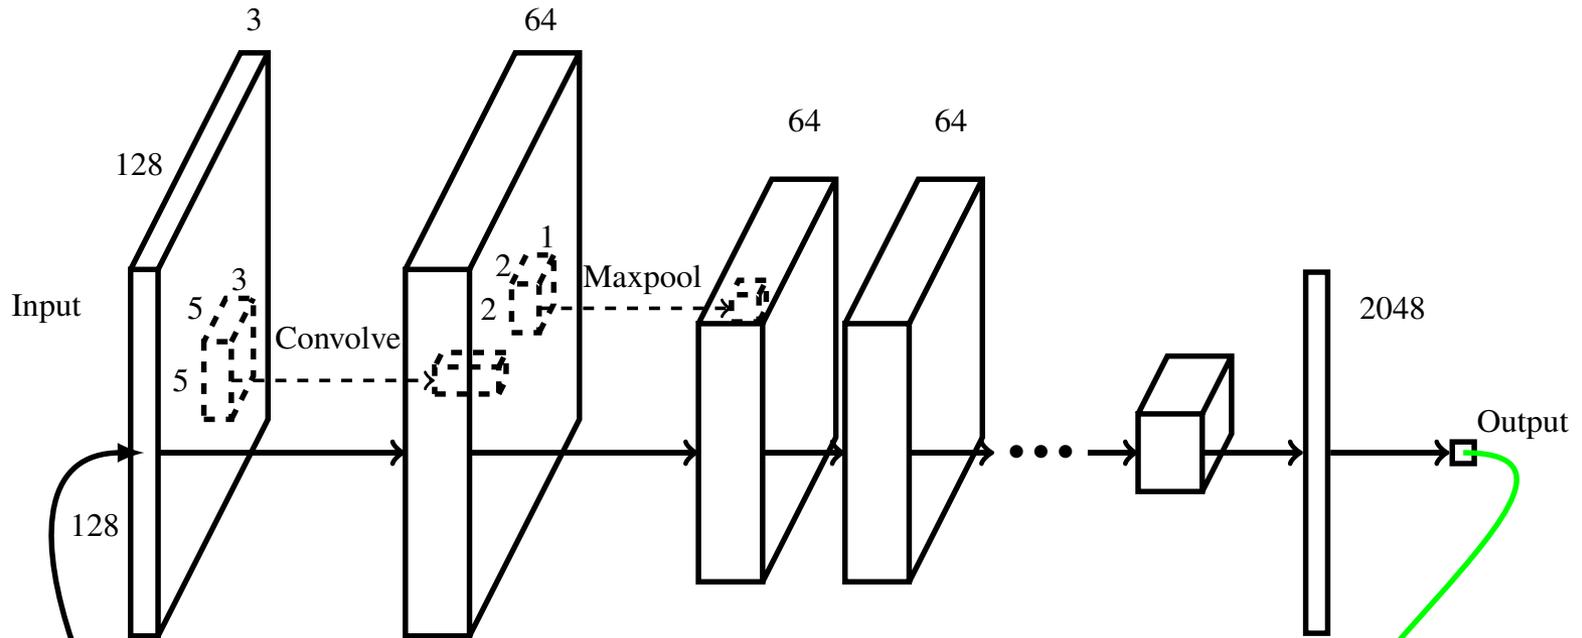
Brest Cancer - mitotic cells



(a) Scanner A

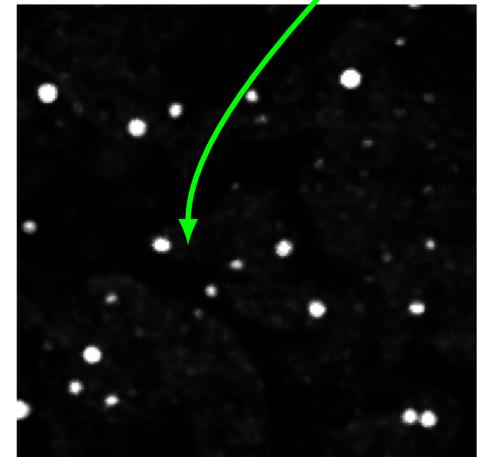
(b) Scanner H

Detection Algorithm



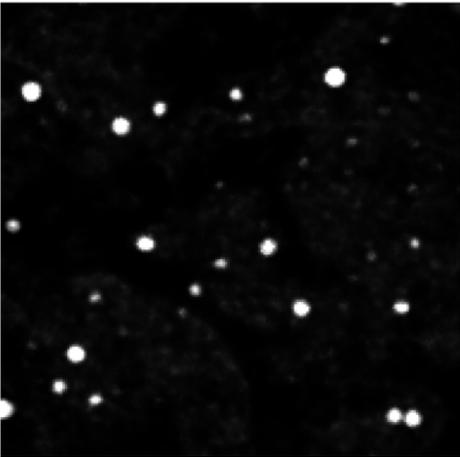
Input Image

Output Probability Map

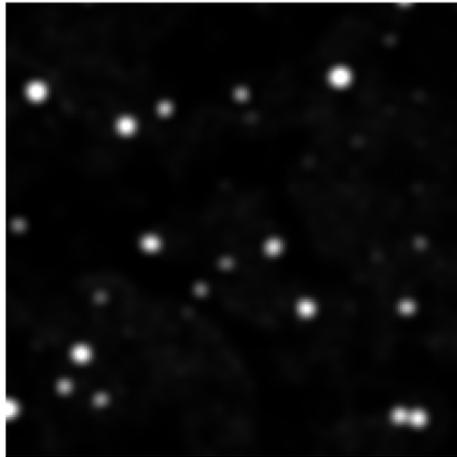


Algorithm

Step 1: Raw probability map



Step 2: After smoothing



Step 3: Non-maximal suppression

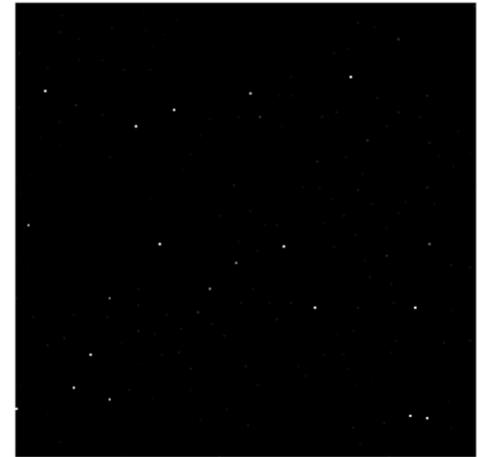


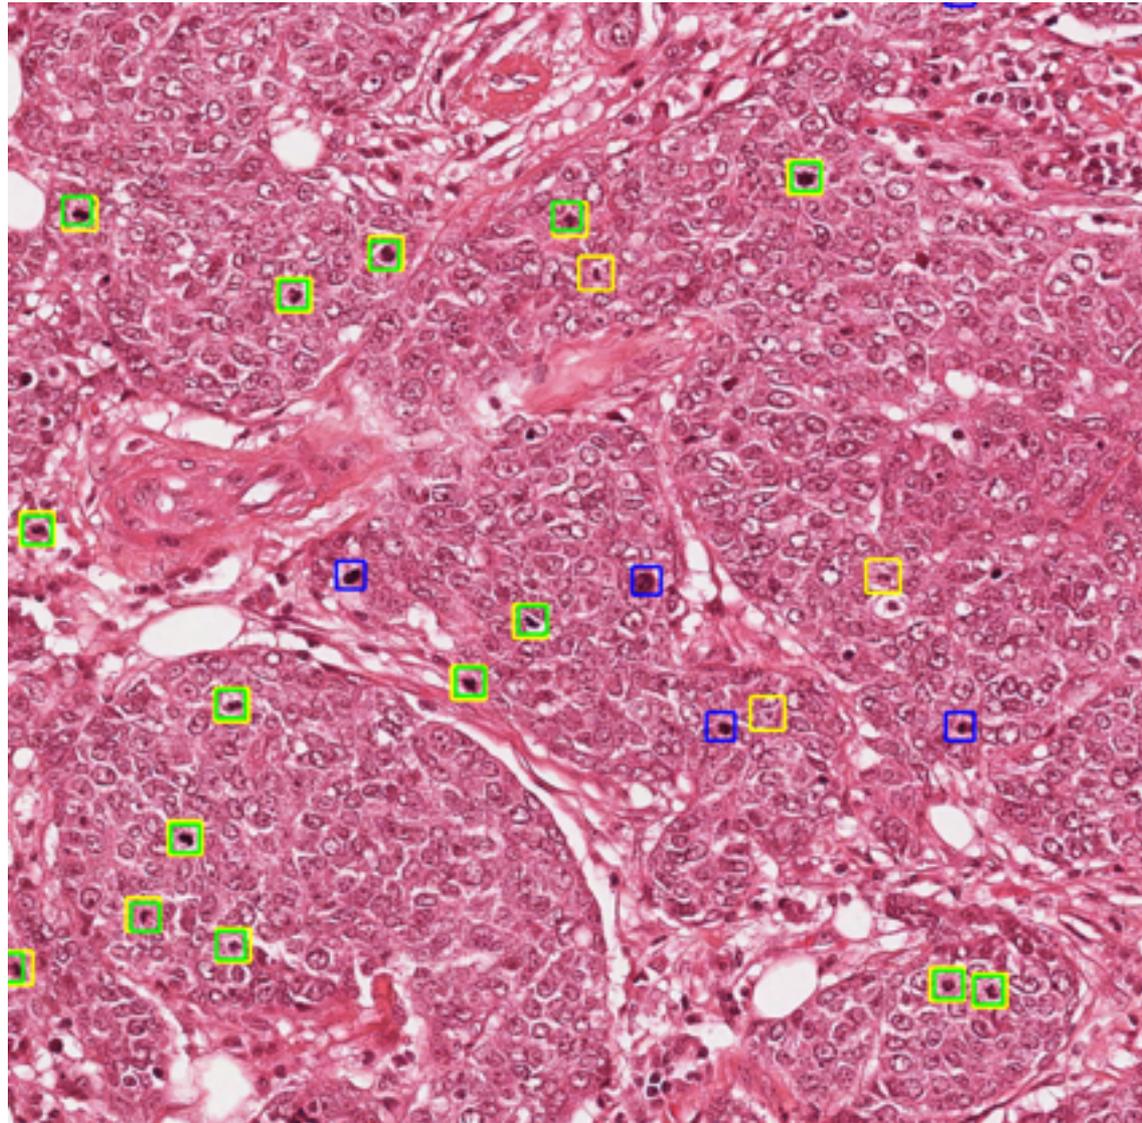
Fig. 3. *The detection process. Left: Convolutional Net outputs a raw probability map. Middle: Gaussian low-pass smoothing prepares the map for non-maximal suppression. Right: Non-maximal suppression suppresses the pixels of the map which are not local maximum. Zoom in to see the "white dots".*

Results

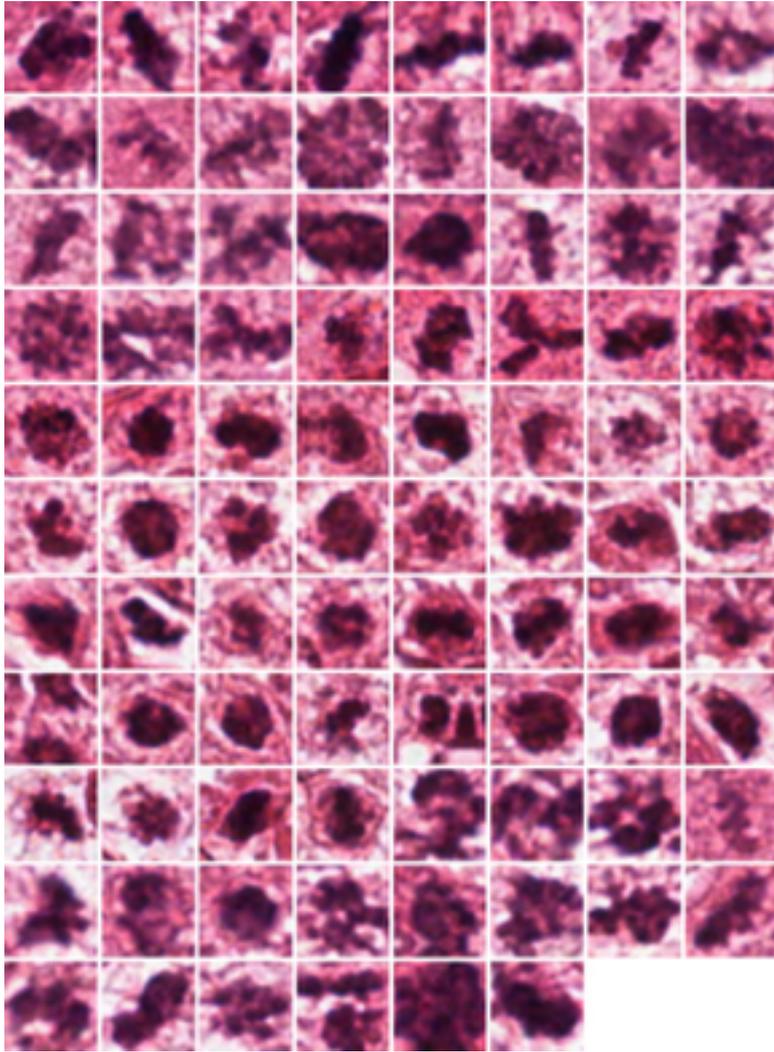
- Learning performs well even with a small dataset
- 50 images with 300 mitosis cells
- Better than untrained human
- State-of-the-art across all literature

Method	Precision	Recall	F-measure
This paper	0.75	0.85	0.800
DNN [3]	0.88	0.70	0.782
IPAL [11]	0.69	0.74	0.718
SUTECH	0.70	0.72	0.709
NEC [15]	0.74	0.59	0.659
UTRECHT [22]	0.51	0.68	0.583
WARWICK [13]	0.46	0.57	0.513
NUS	0.63	0.40	0.490
ISIK [21]	0.28	0.68	0.397
ETH-HEILDERBERG [20]	0.14	0.80	0.374
OKAN-IRISA-LIAMA	0.78	0.22	0.343
IITG	0.17	0.46	0.255
DREXEL	0.14	0.21	0.172

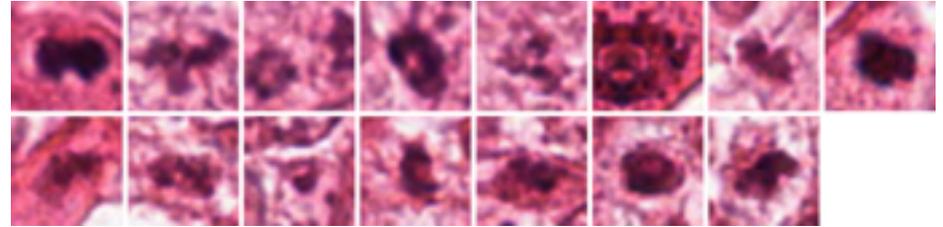
Table 2. Detection performance on the test set images. Comparison with previously published results on the exact same test set. Our result is state-of-the-art.



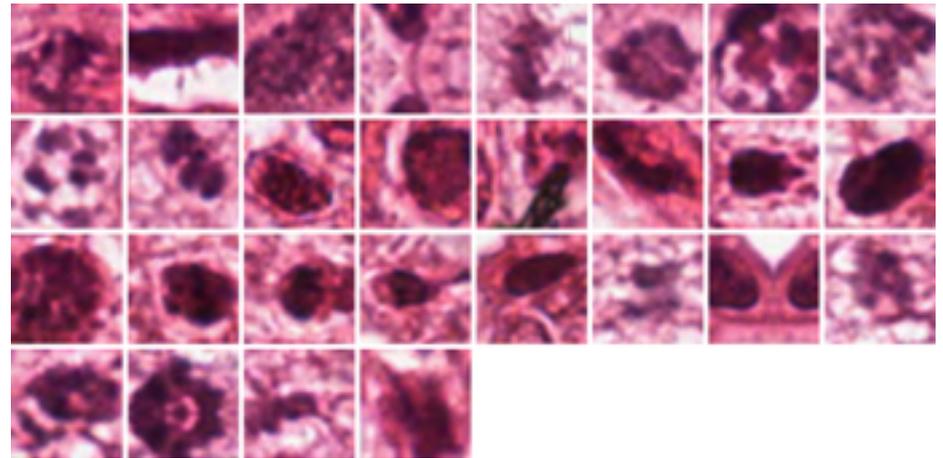
Results



True Positives



False Negatives



False Positives