CSEP 573: Artificial Intelligence

Applications

Hanna Hajishirzi

Many slides over the course adapted from Pieter Abbeel, Luke Zettlemoyer and Dan Klein.

Applications of AI

- Web
- NLP
- Vision
- Robotics
- Games
- Predictions
- Diagnosis
- …

Web Search

- Information retrieval:
  - Given information needs, produce information
  - Includes, e.g. web search, question answering, and classic IR

- Web search: not exactly classification, but rather ranking

x = “Apple Computers”

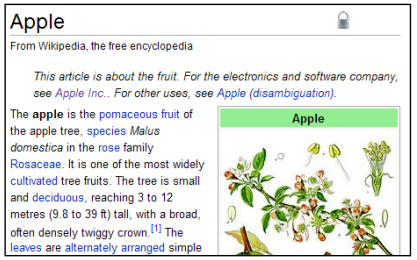

Feature-Based Ranking

x = “Apple Computer”

$f(x,$ 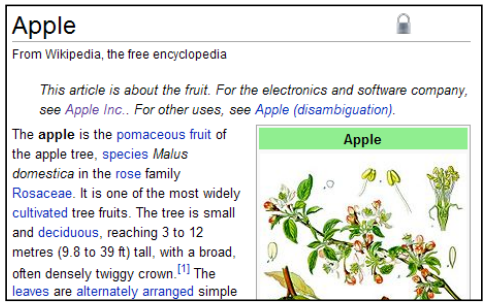

$) = [0.3\ 5\ 0\ 0\ \ldots]$

$f(x,$ 

$) = [0.8\ 4\ 2\ 1\ \ldots]$

Perceptron For Ranking

- Inputs x
- Candidates y
- Many feature vectors: $f(x, y)$
- One weight vector: w
  - Prediction:

$$y = \arg\max_y \; w \cdot f(x, y)$$

  - Update (if wrong):

$$w = w + f(x, y^*) - f(x, y)$$

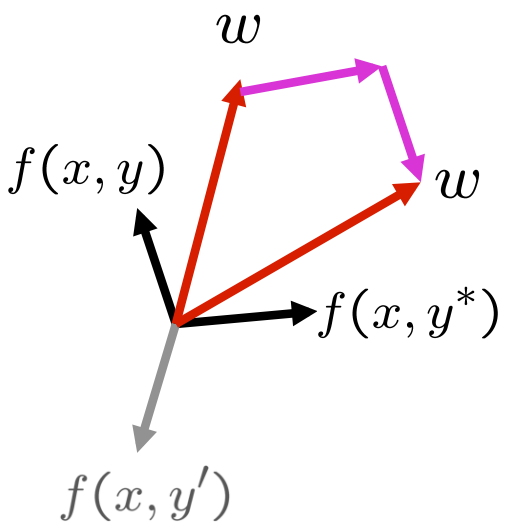

NLP

- Headlines:
  - Enraged Cow Injures Farmer With Ax
  - Hospitals Are Sued by 7 Foot Doctors
  - Ban on Nude Dancing on Governor's Desk
  - Iraqi Head Seeks Arms
  - Local HS Dropouts Cut in Half
  - Juvenile Court to Try Shooting Defendant
  - Stolen Painting Found by Tree
  - Kids Make Nutritious Snacks

- Why are these funny?

- Fundamental goal: analyze and process human language, broadly, robustly, accurately…
- End systems that we want to build:
  - Ambitious: speech recognition, machine translation, information extraction, dialog interfaces, question answering…
  - Modest: spelling correction, text categorization…

Parsing

Hershey bars protest

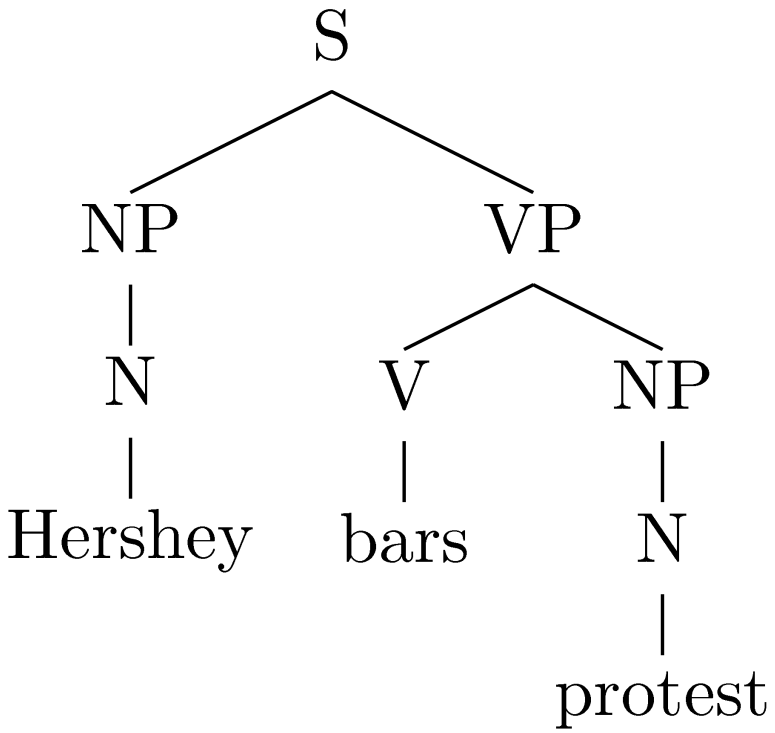

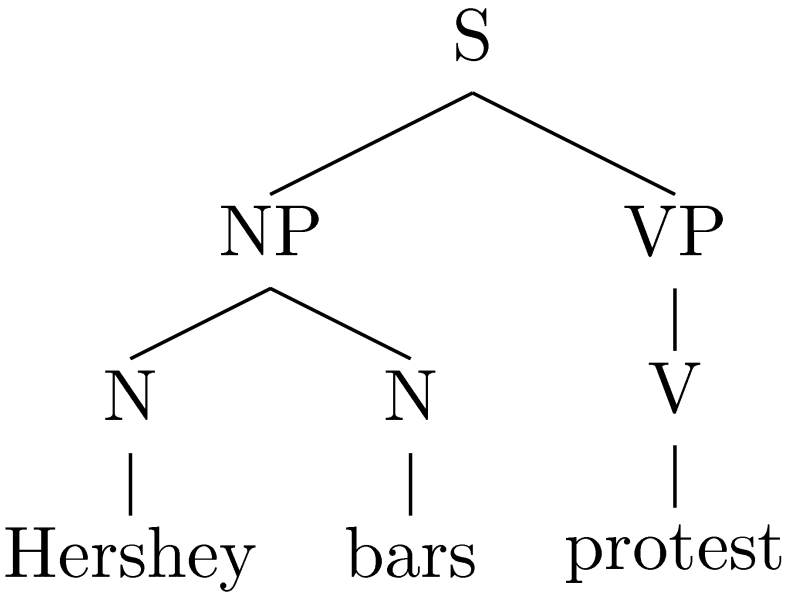

Grammar

- Natural language grammars are very ambiguous!
- PCFGs are a formal probabilistic model of trees
  - Each "rule" has a conditional probability (like an HMM)
  - Tree's probability is the product of all rules used
- Parsing: Given a sentence, find the best tree – search!

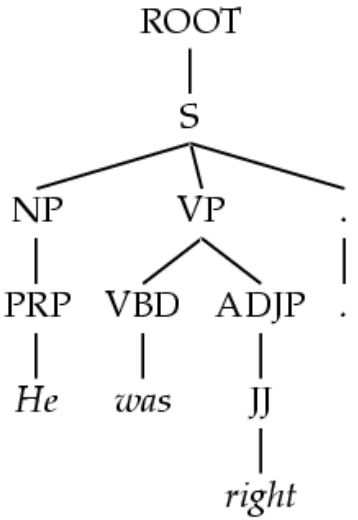

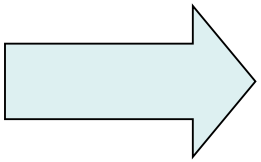

ROOT → S	375/420
S → NP VP .	320/392
NP → PRP	127/539
VP → VBD ADJP	32/401
…..	

Dialogue Systems

- Watson
  - A question-answering system (IBM, 2011)
  - Designed for the game of Jeopardy
  - How does it work:
    - Sophisticated NLP: deep analysis of questions, noisy matching of questions to potential answers
    - Lots of data: onboard storage contains a huge collection of documents (e.g. Wikipedia, etc.), exploits redundancy
    - Lots of computation: 90+ servers
  - Can beat all of the people all of the time?

Machine Translation

"Il est impossible aux journalistes de rentrer dans les régions tibétaines"

Bruno Philip, correspondant du "Monde" en Chine, estime que les journalistes de l'AFP qui ont été expulsés de la province tibétaine du Qinghai "n'étaient pas dans l'illégalité".

Les faits Le dalaï-lama dénonce l'"enfer" imposé au Tibet depuis sa fuite, en 1959

Vidéo Anniversaire de la rébellion tibétaine : la Chine sur ses gardes

"It is impossible for journalists to enter Tibetan areas"

Philip Bruno, correspondent for "World" in China, said that journalists of the AFP who have been deported from the Tibetan province of Qinghai "were not illegal."

Facts The Dalai Lama denounces the "hell" imposed since he fled Tibet in 1959

Video Anniversary of the Tibetan rebellion: China on guard

- Translate text from one language to another
- Recombines fragments of example translations
- Challenges:
  - What fragments? [learning to translate]
  - How to make efficient? [fast translation search]

Problem with Dictionary Lookups

顶部	/**top**/roof/
顶端	/summit/peak/**top**/apex/
顶头	/coming directly towards one/**top**/end/
盖	/lid/**top**/cover/canopy/build/Gai/
盖帽	/surpass/**top**/
极	/extremely/pole/utmost/**top**/collect/receive/
尖峰	/peak/**top**/
面	/fade/side/surface/aspect/**top**/face/flour/
摘心	/**top**/topping/

Example from Douglas Hofstadter

Data-Driven Approach

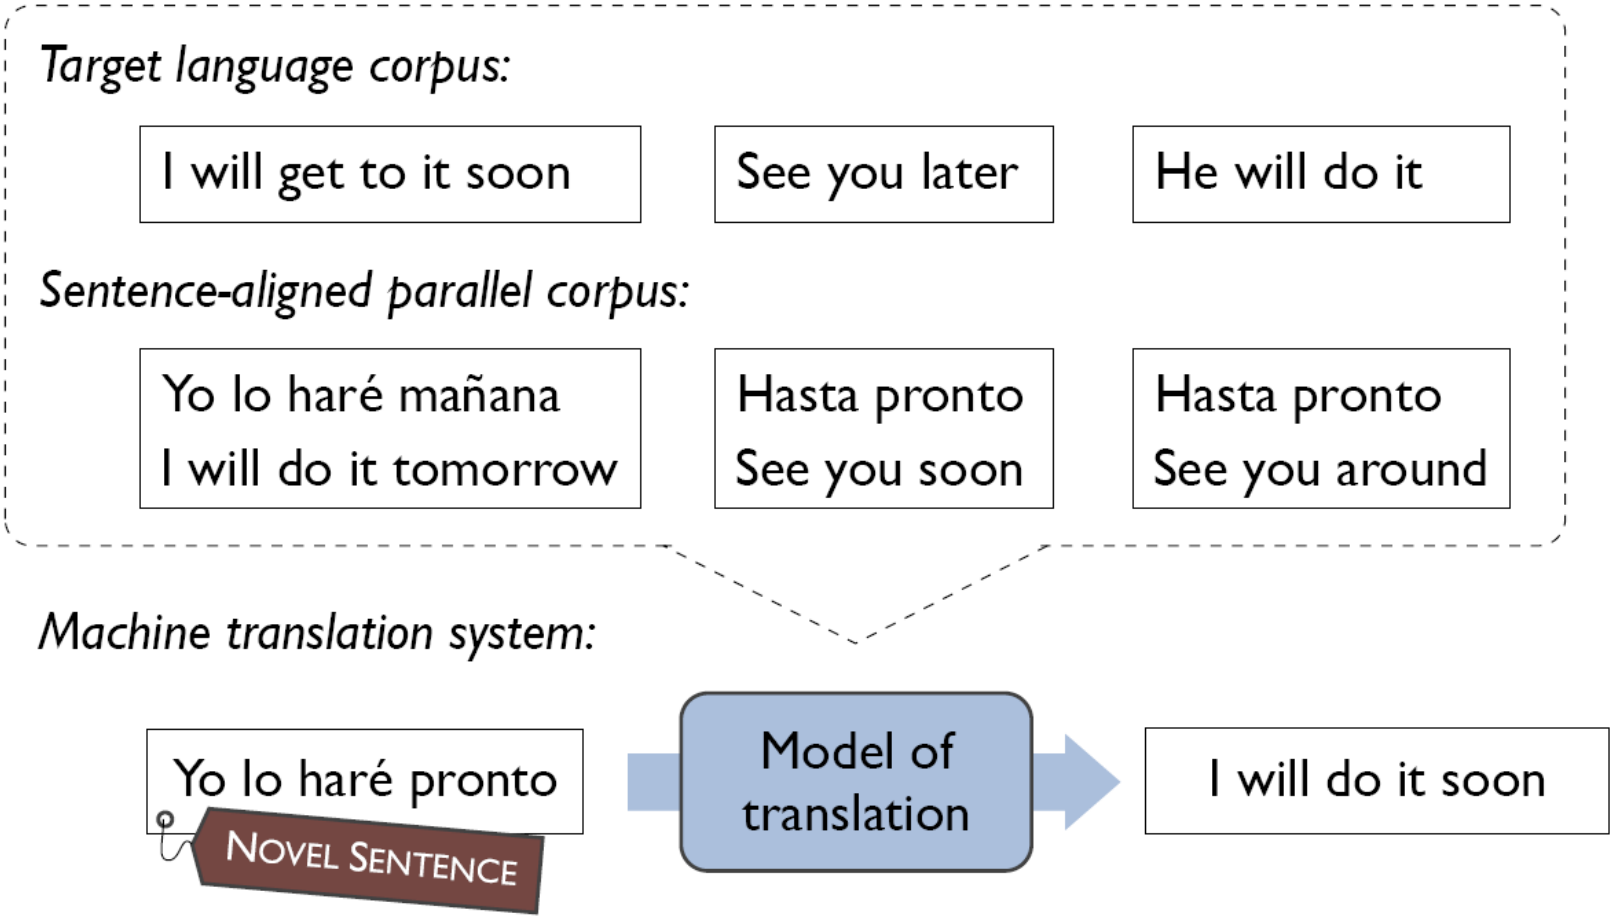

Learning to Translate

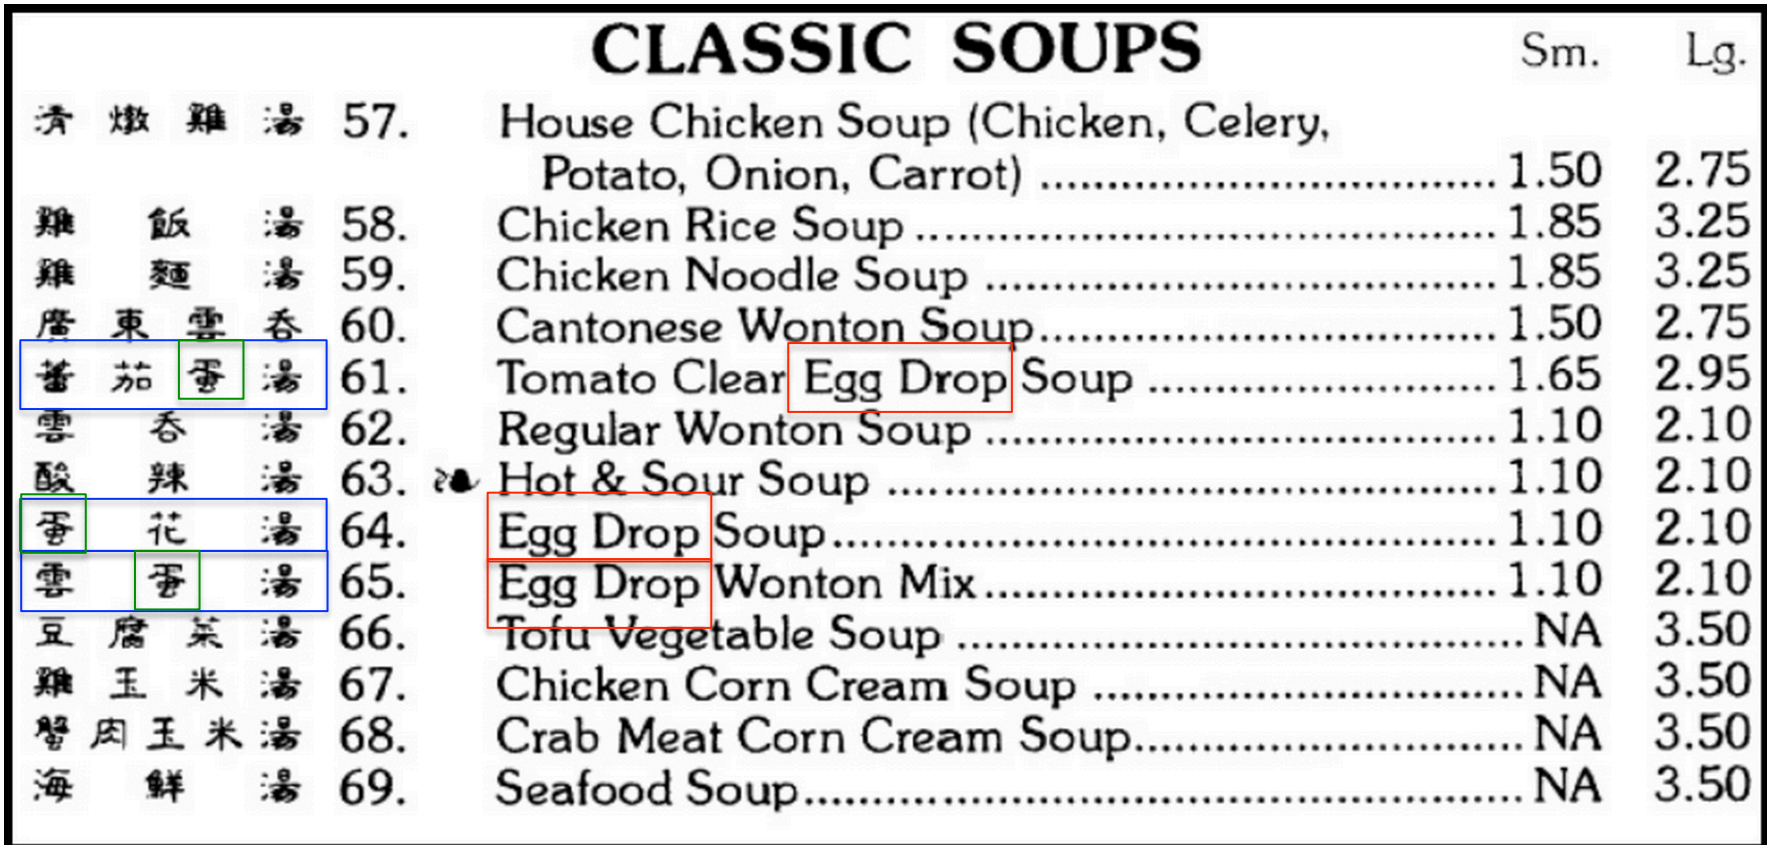

CLASSIC SOUPS

			Sm.	Lg.
清燉雞湯	57.	House Chicken Soup (Chicken, Celery, Potato, Onion, Carrot)	1.50	2.75
雞飯湯	58.	Chicken Rice Soup	1.85	3.25
雞麵湯	59.	Chicken Noodle Soup	1.85	3.25
廣東雲吞	60.	Cantonese Wonton Soup	1.50	2.75
番茄蛋湯	61.	Tomato Clear Egg Drop Soup	1.65	2.95
雲吞湯	62.	Regular Wonton Soup	1.10	2.10
酸辣湯	63.	Hot & Sour Soup	1.10	2.10
蛋花湯	64.	Egg Drop Soup	1.10	2.10
雲蛋湯	65.	Egg Drop Wonton Mix	1.10	2.10
豆腐菜湯	66.	Tofu Vegetable Soup	NA	3.50
雞玉米湯	67.	Chicken Corn Cream Soup	NA	3.50
蟹肉玉米湯	68.	Crab Meat Corn Cream Soup	NA	3.50
海鮮湯	69.	Seafood Soup	NA	3.50

An HMM Translation Model

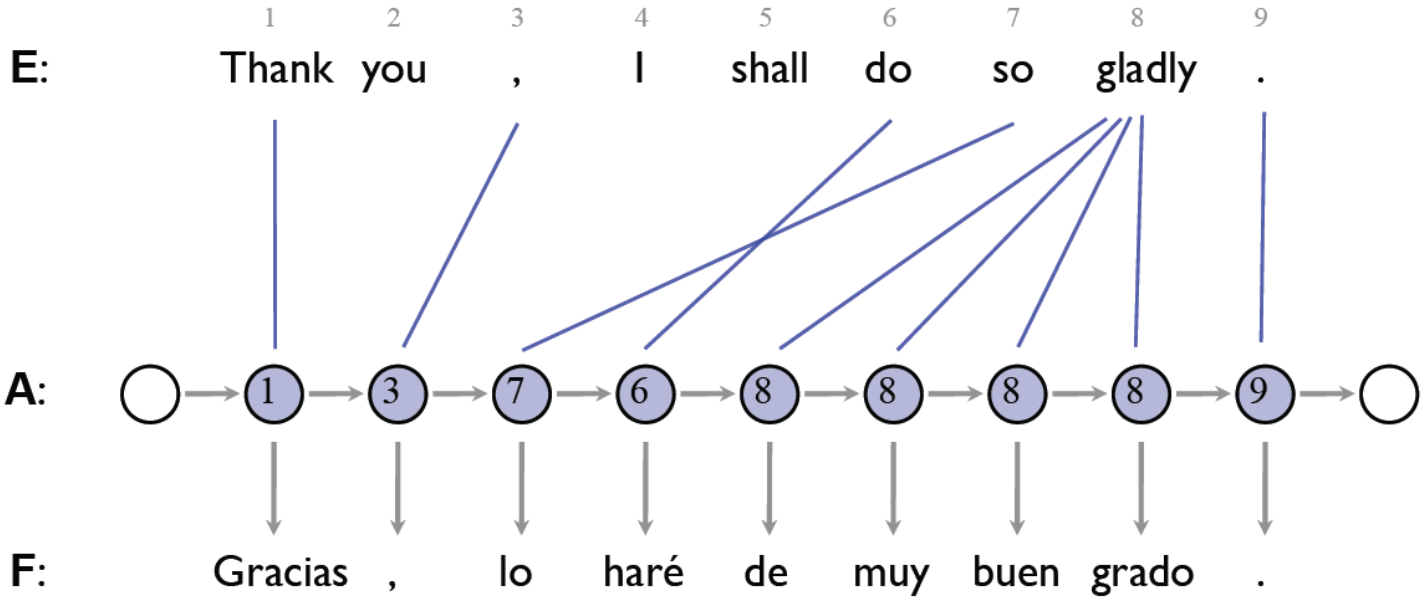

Model Parameters

Emissions: $P(F_1 = \text{Gracias} \mid E_{A1} = \text{Thank})$ *Transitions:* $P(A_2 = 3 \mid A_1 = 1)$

Levels of Transfer

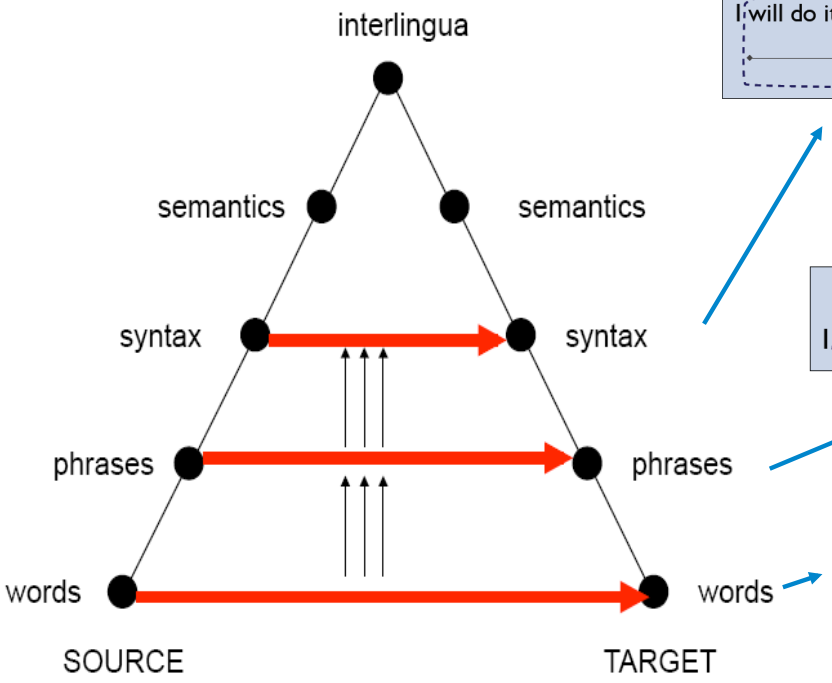

Yo lo haré mañana
I will do it tomorrow
NP
VP

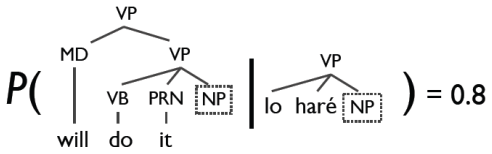

Yo lo haré mañana
I will do it tomorrow

English (E)	*P*(E | lo haré)
will do it	0.8
will do so	0.2

Yo lo haré mañana
I will do it tomorrow

English (E)	*P*(E | mañana)
tomorrow	0.7
morning	0.3

Coreference Resolution

[Michael Eisner] and [Donald Tsang] announced the grand opening of [Hong Kong Disneyland] yesterday. [Eisner] thanked [Mr. Tsang] and welcomed [fans] to [the park].

- Coreference resolution:
  - Determine when two mentions refer to same individual

- Require semantic knowledge to better coreference

Named Entity Linking

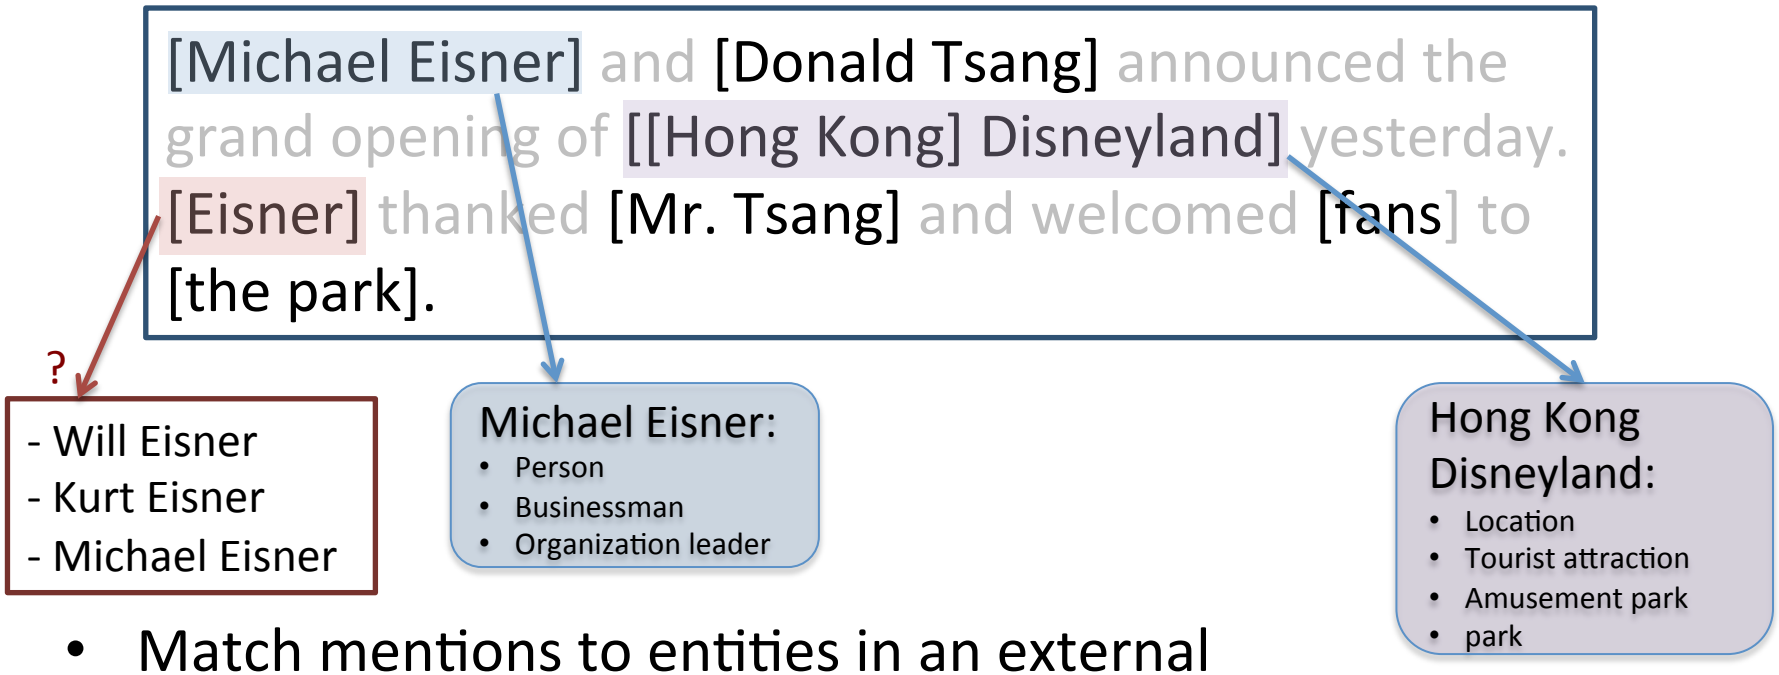

- Match mentions to entities in an external knowledge base (Freebase, Wikipedia)
  - Use entity attributes as semantic knowledge
- NEL is challenging

Solving Arithmetic Word Problems

Liz had 9 black kittens. She gave some of her kittens to Joan. Joan has now 11 kittens. Liz has 5 kitten left and 3 has spots. How many kittens did Joan get?

Equation: $9 - x = 5$

Solution: $x = 5$ kittens

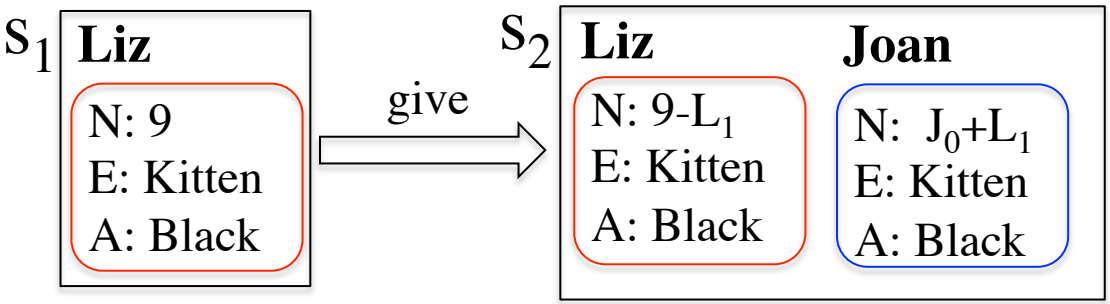

Liz gave some of her kittens to **Joan**.

Vision

- Search
- Detection
- Surveillance
- Recognition

Mobile visual search: Google Goggles

Google Goggles in Action

Click the icons below to see the different ways Google Goggles can be used.

Google goggles labs

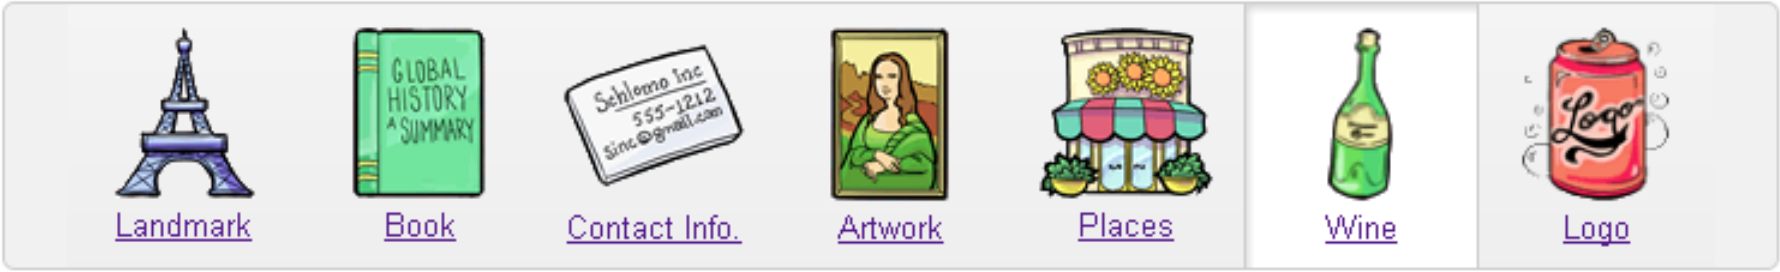

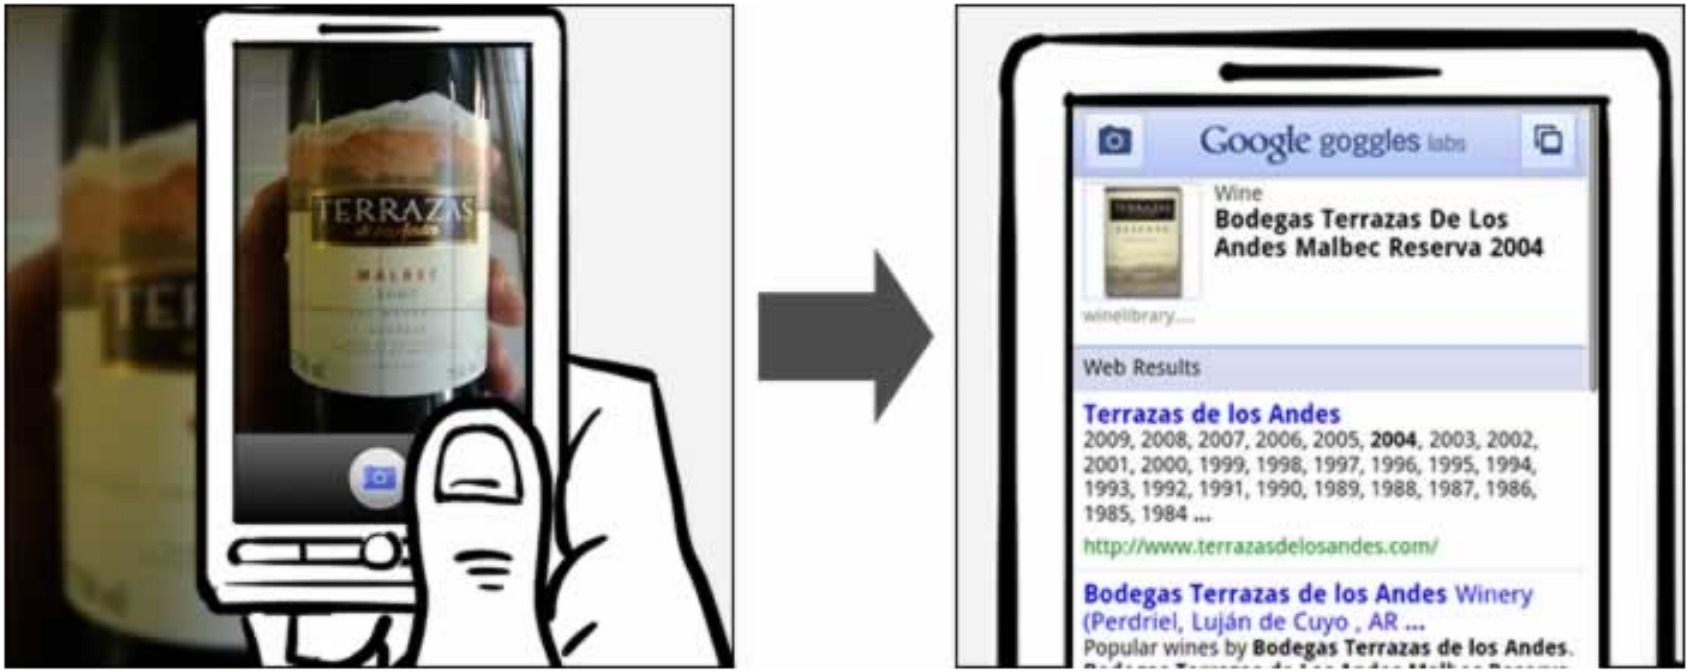

Face detection

- Many new digital cameras now detect faces
  - Canon, Sony, Fuji, …

Source: S. Seitz

Smile detection

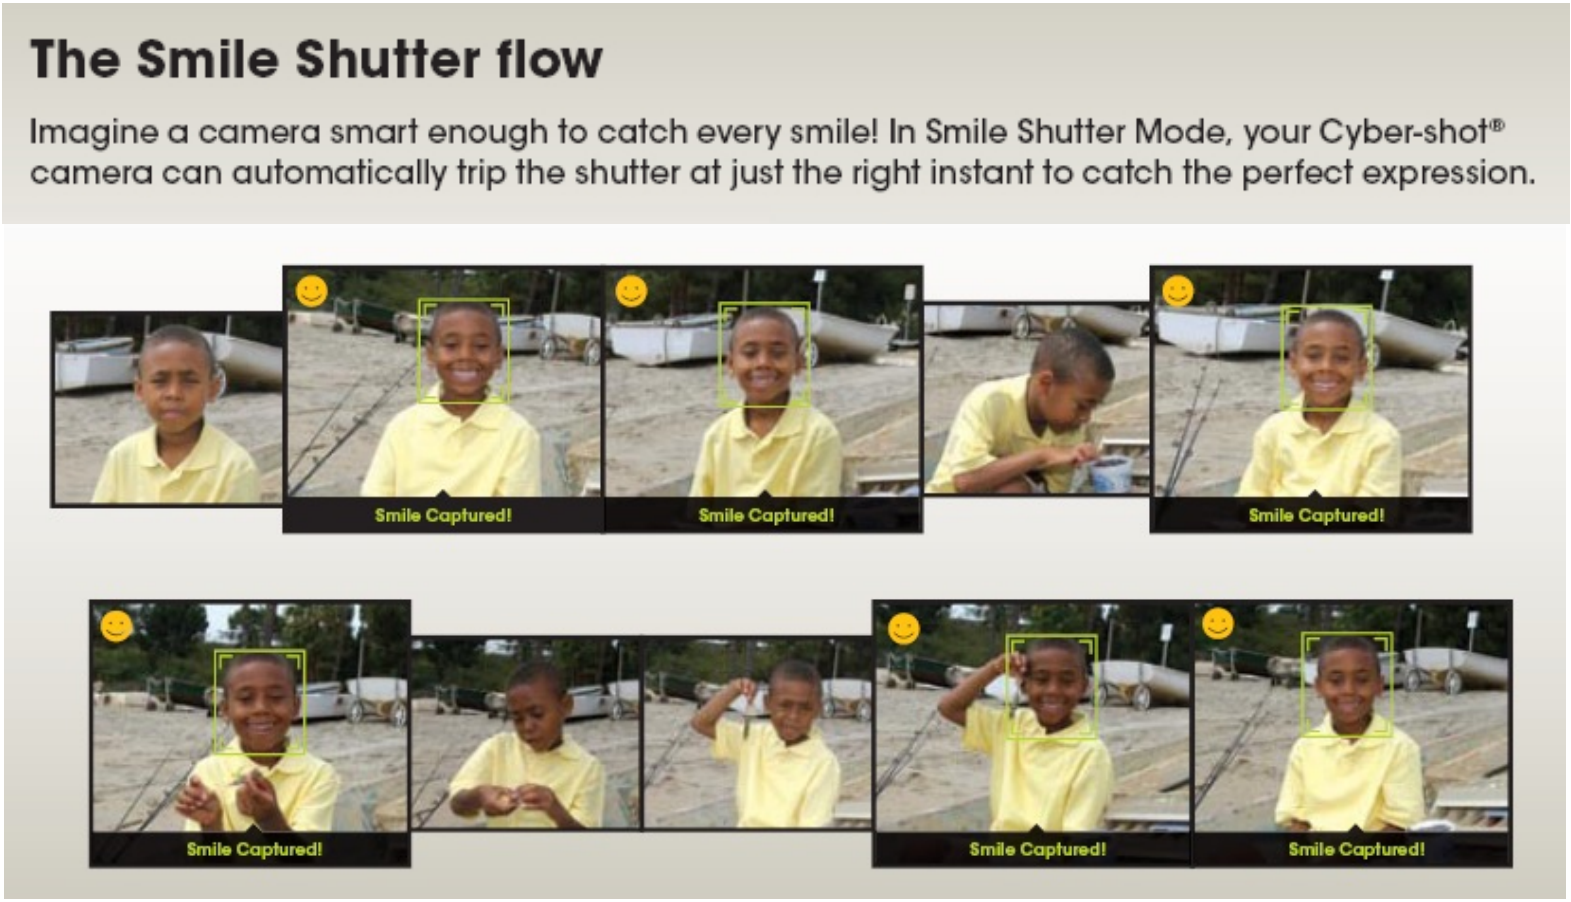

Sony Cyber-shot® T70 Digital Still Camera

Source: S. Seitz

Face recognition: Apple iPhoto, Facebook, Google, etc

Object recognition (in supermarkets)

LaneHawk by EvolutionRobotics

“A smart camera is flush-mounted in the checkout lane, continuously watching for items. When an item is detected and recognized, the cashier verifies the quantity of items that were found under the basket, and continues to close the transaction. The item can remain under the basket, and with LaneHawk,you are assured to get paid for it… “

News Front Page

Africa
Americas
Asia-Pacific
Europe
Middle East
South Asia
UK
England
Northern Ireland
Scotland
Wales
UK Politics
Education
Magazine
Business
Health
Science &

Last Updated: Wednesday, 31 August 2005, 05:44 GMT 06:44 UK

E-mail this to a friend

Printable version

Computer alert for drowning girl

A 10-year-old girl has been saved from drowning by a computer system designed to raise the alarm when swimmers get into difficulties.

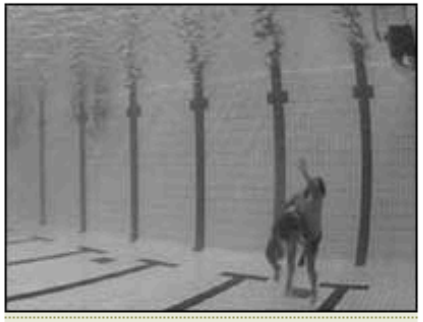

VIDEO Watch the rescue

The girl, from Rochdale, was at the deep end of the pool in Bangor, north Wales, when she sank to the bottom.

The £65,000 system, called Poseidon, detected her on the pool floor and sounded the alarm. A lifeguard pulled her out and she recovered in hospital.

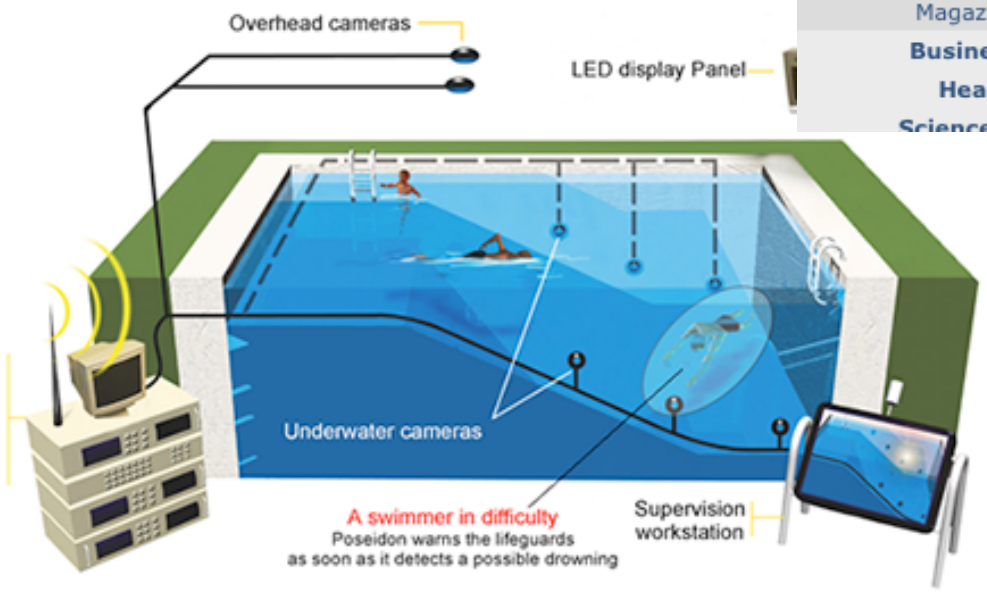

Security

Local

Cameras help confirm Scott suicide ruling

Friday, December 04, 2009

TAGS: local, paul meincke

Comment Now Email Print Report a typo

Paul Meincke
More: Bio, News Team

December 4, 2009 (CHICAGO) (WLS) -- Chicago police have closed the case in the death of Chicago School Board President Michael Scott.

Police Supt. Jody Weis says investigators used police cameras in the city to trace Scott's last steps in the hours before his body was found in November.

Scott's death has been ruled a suicide. The medical examiner's office concluded --not long after Scott's body was found -- that he had committed suicide. Police did not dispute the finding but wanted to pursue all the investigative leads they could. They say they have done that and have now reached the same conclusion.

Share this Story

Recommend Be the first of your friends to recommend this.

Tweet 0 +1 Recommend this on Google

News Headlines Video abcNEWS

- 2 suspects arrested in volleyball star's murder 47 min ago
- BP Gas Recall: BP finds, fixes source of bad gas
- Teachers union, board resume negotiating
- Back to School
- 5 injured in South Side shooting 49 min ago
- Pastor: Stacy Peterson said she lied for Drew

Automotive safety

- Mobileye: Vision systems in high-end BMW, GM, Volvo models
  - Pedestrian collision warning
  - Forward collision warning
  - Lane departure warning
  - Headway monitoring and warning

Source: A. Shashua, S. Seitz

Intelligent Suspension system

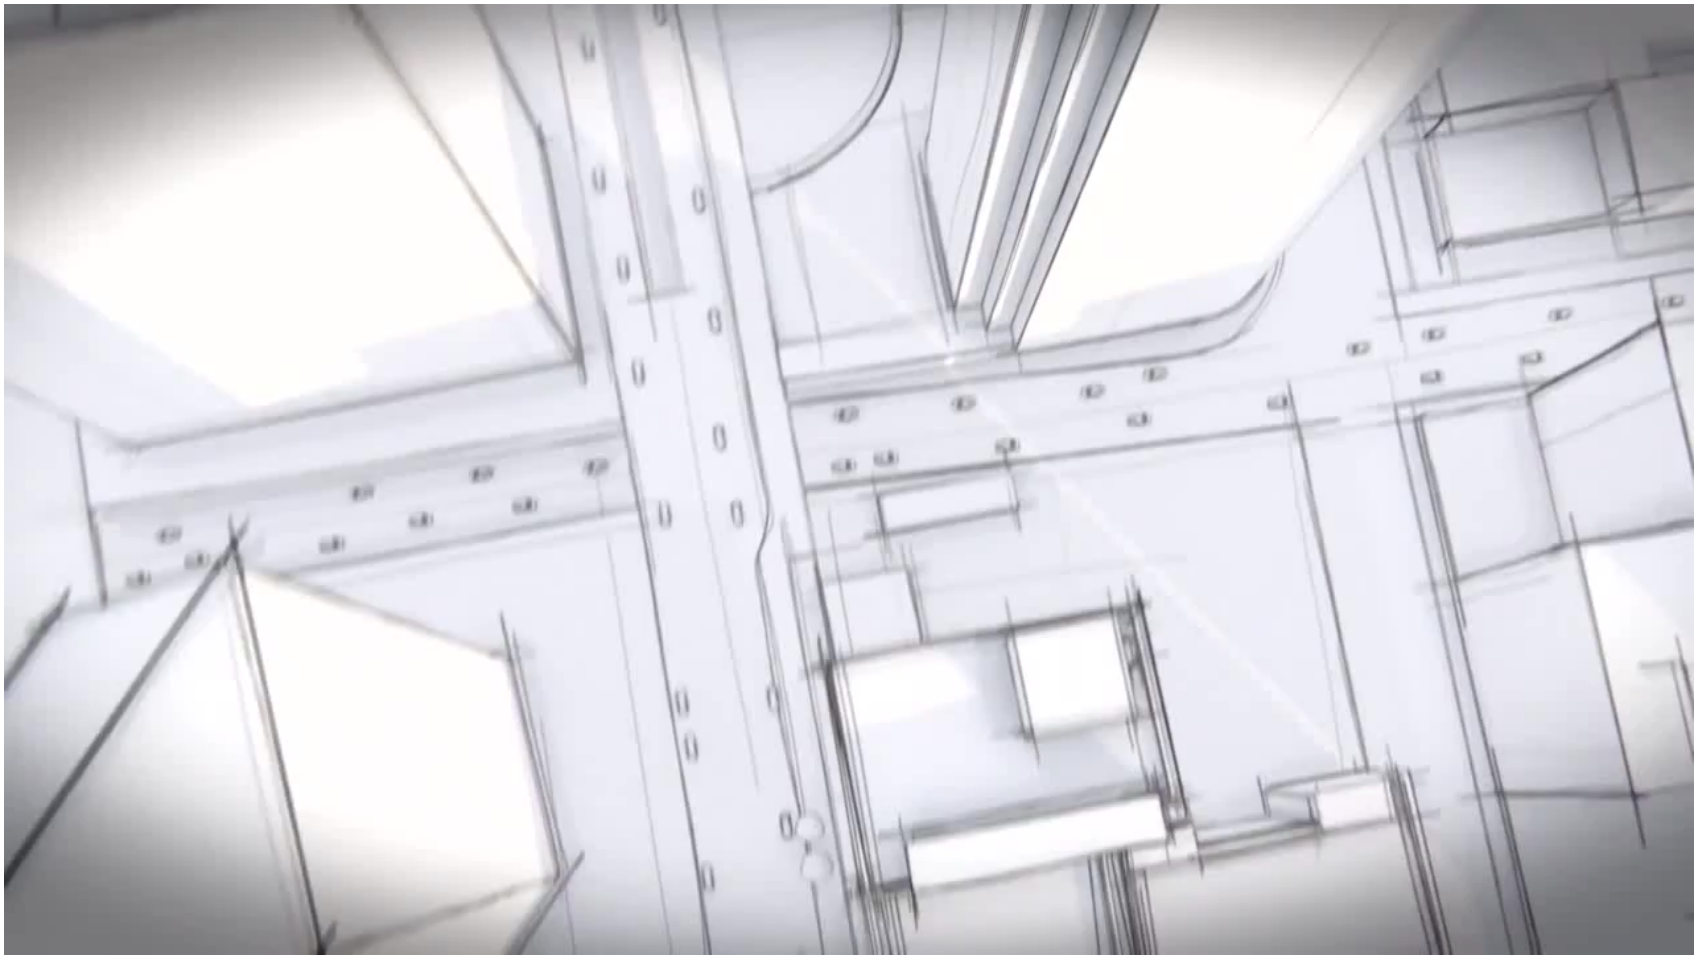

Kinect Fusion

Object Detection

Person model

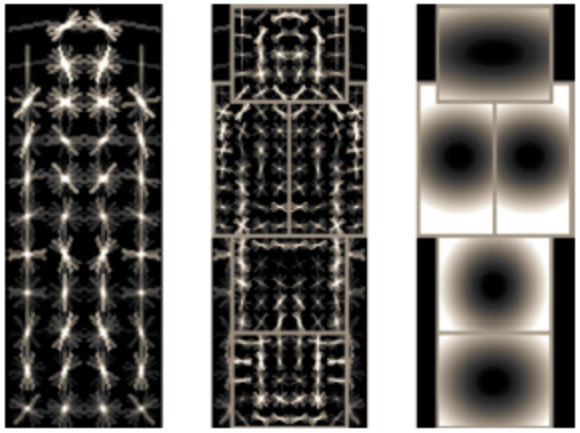

root filters
coarse resolution

part filters
finer resolution

deformation
models

Person detections

high scoring true positives

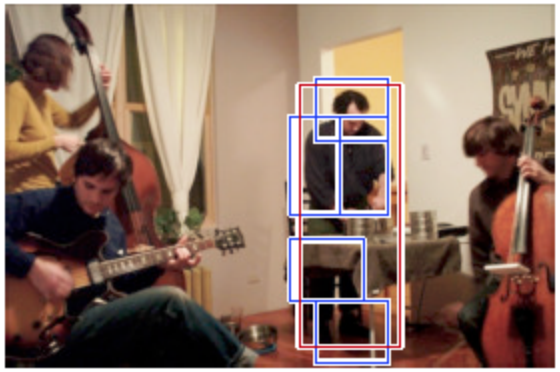

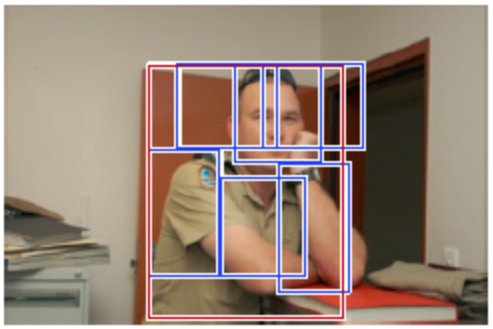
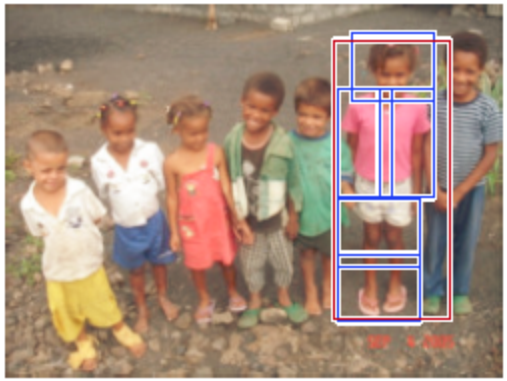

high scoring false positives
(not enough overlap)

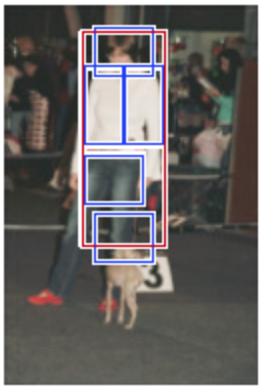


Car

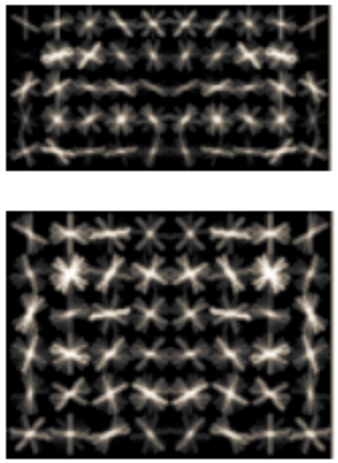

root filters
coarse resolution

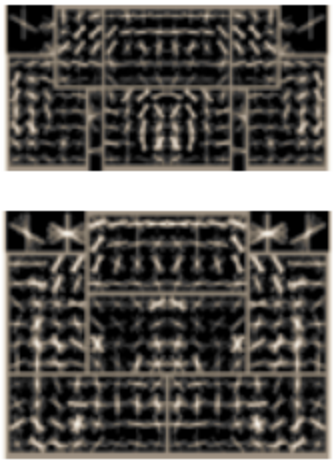

part filters
finer resolution

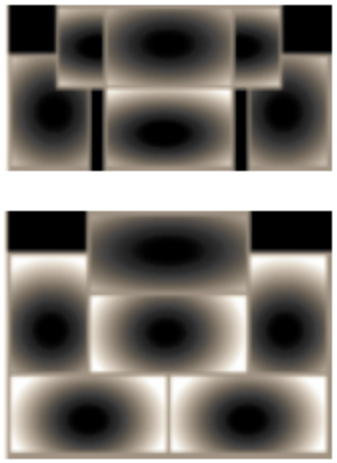

deformation
models

Car detections

high scoring true positives

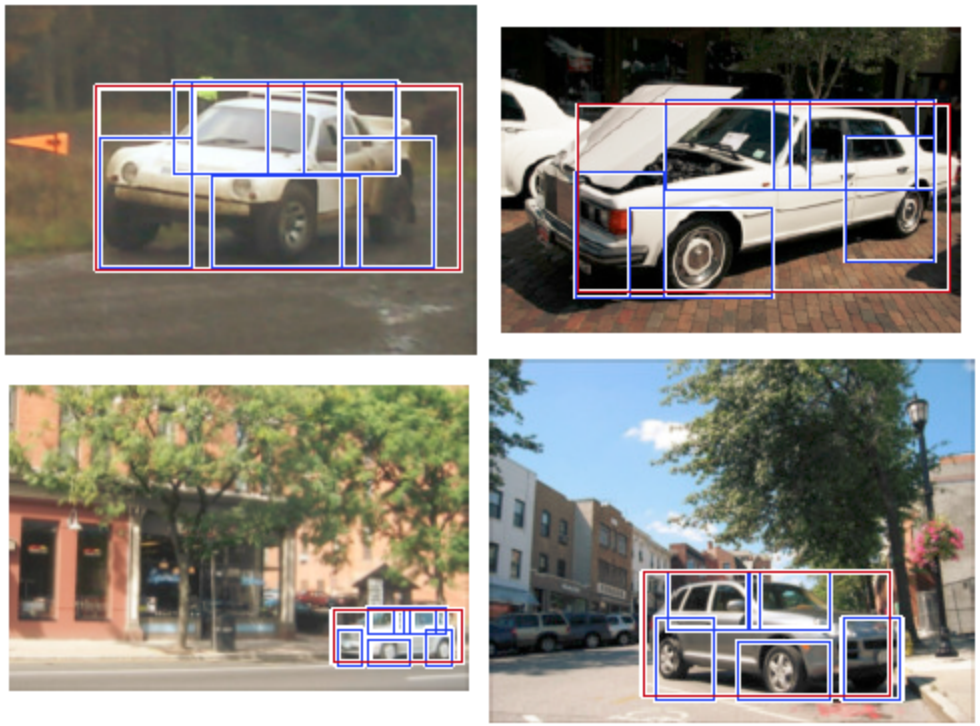

high scoring false positives

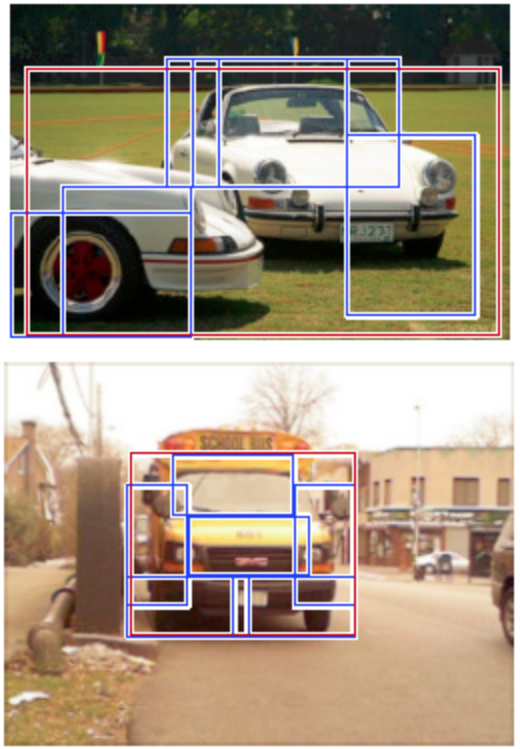

Cat

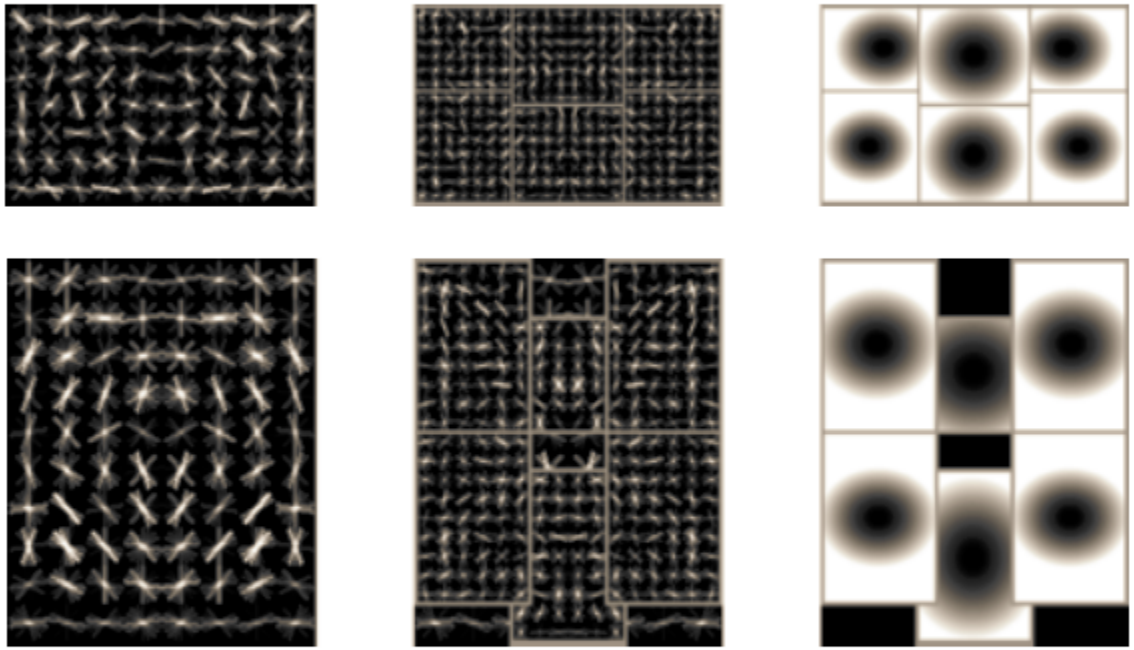

root filters
coarse resolution

part filters
finer resolution

deformation
models

Cat detections

high scoring true positives

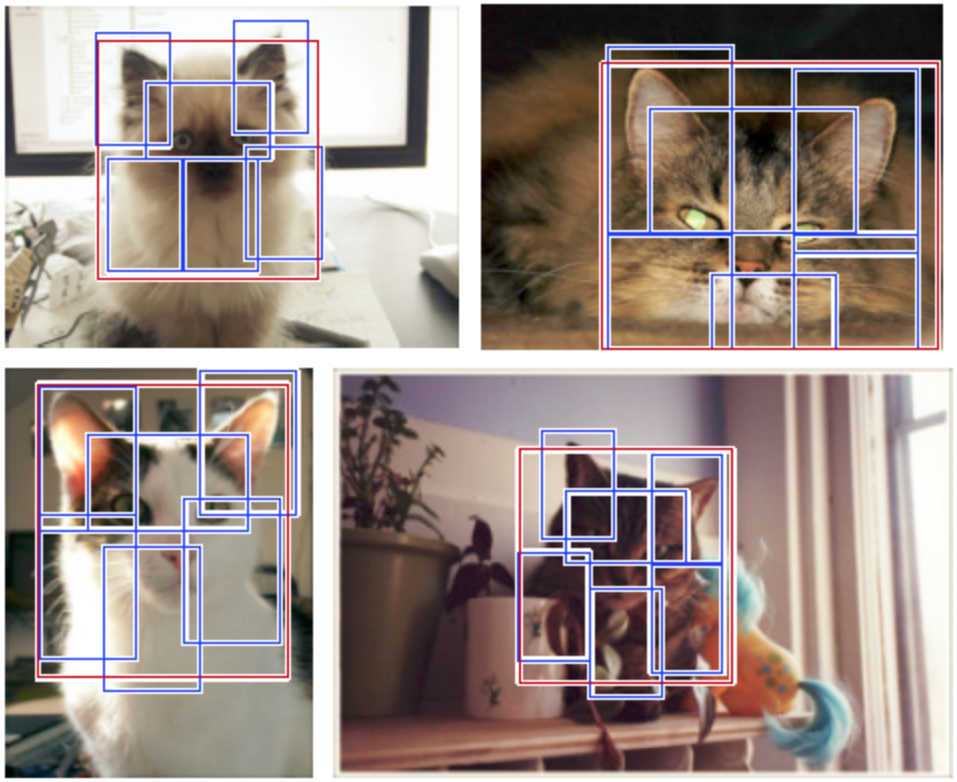

high scoring false positives
(not enough overlap)

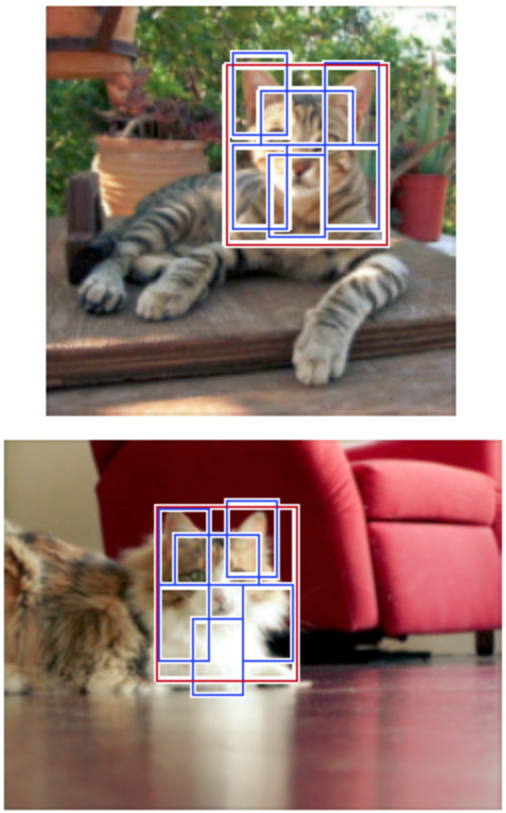

Autonomous Driving

Failures

An Autonomous Car

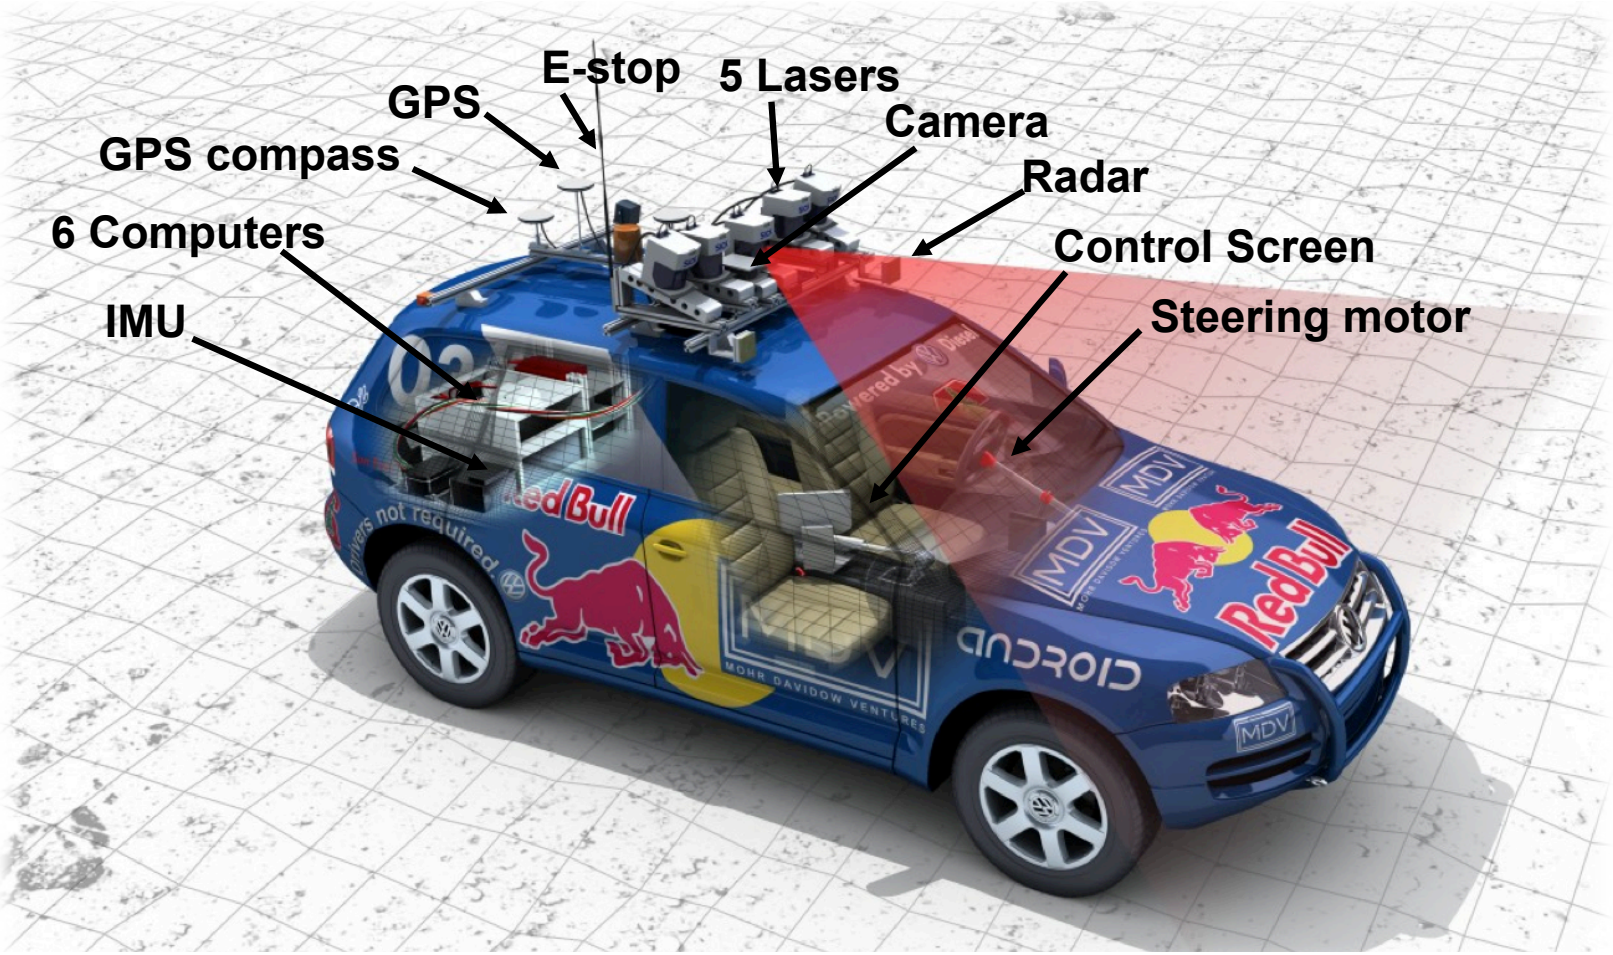

Actions: Steering Control

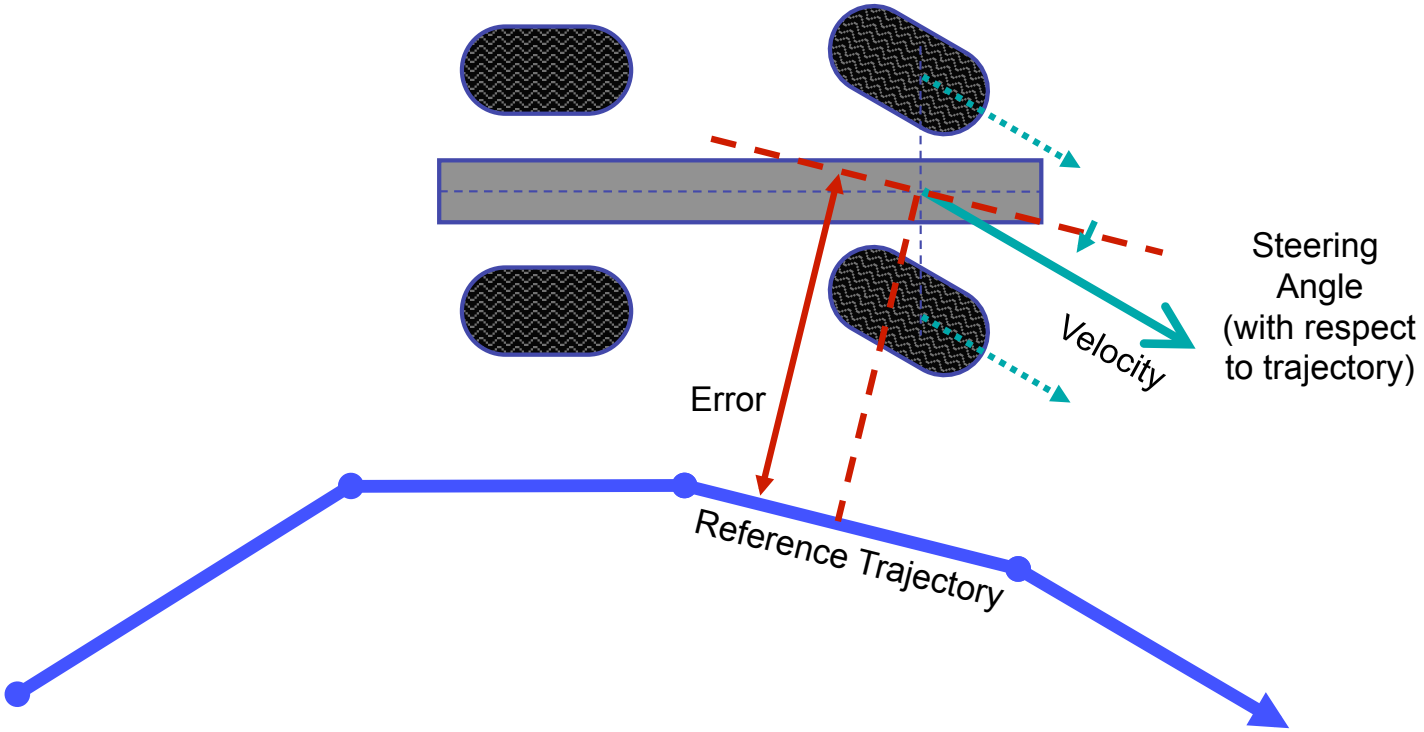

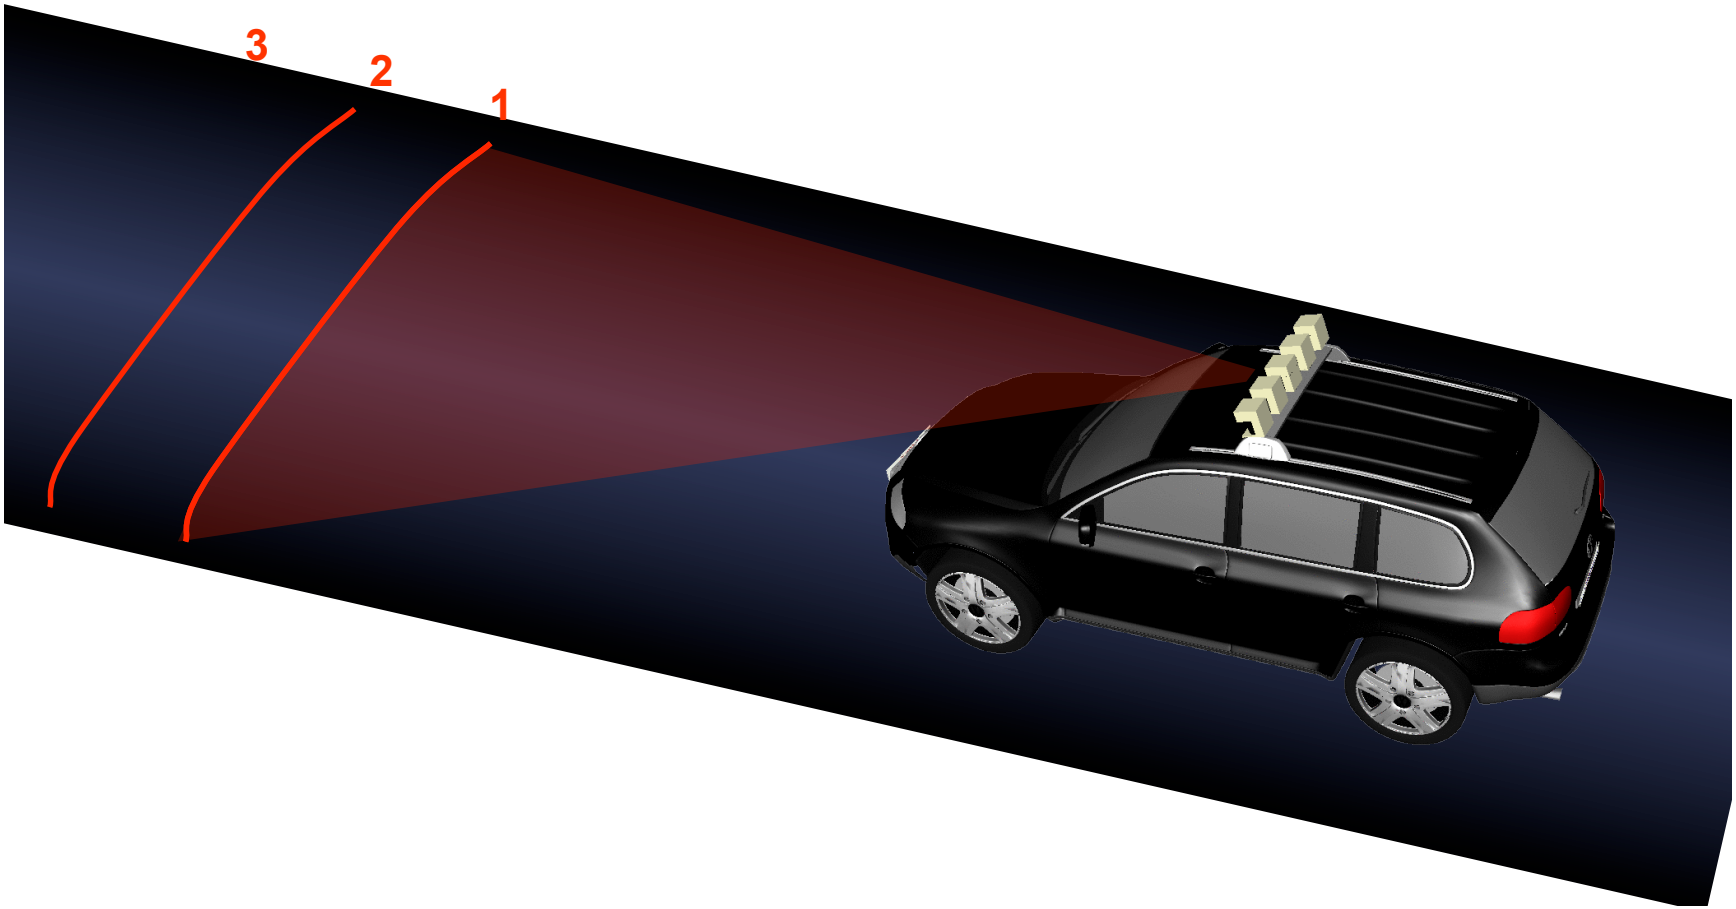
3
2
1

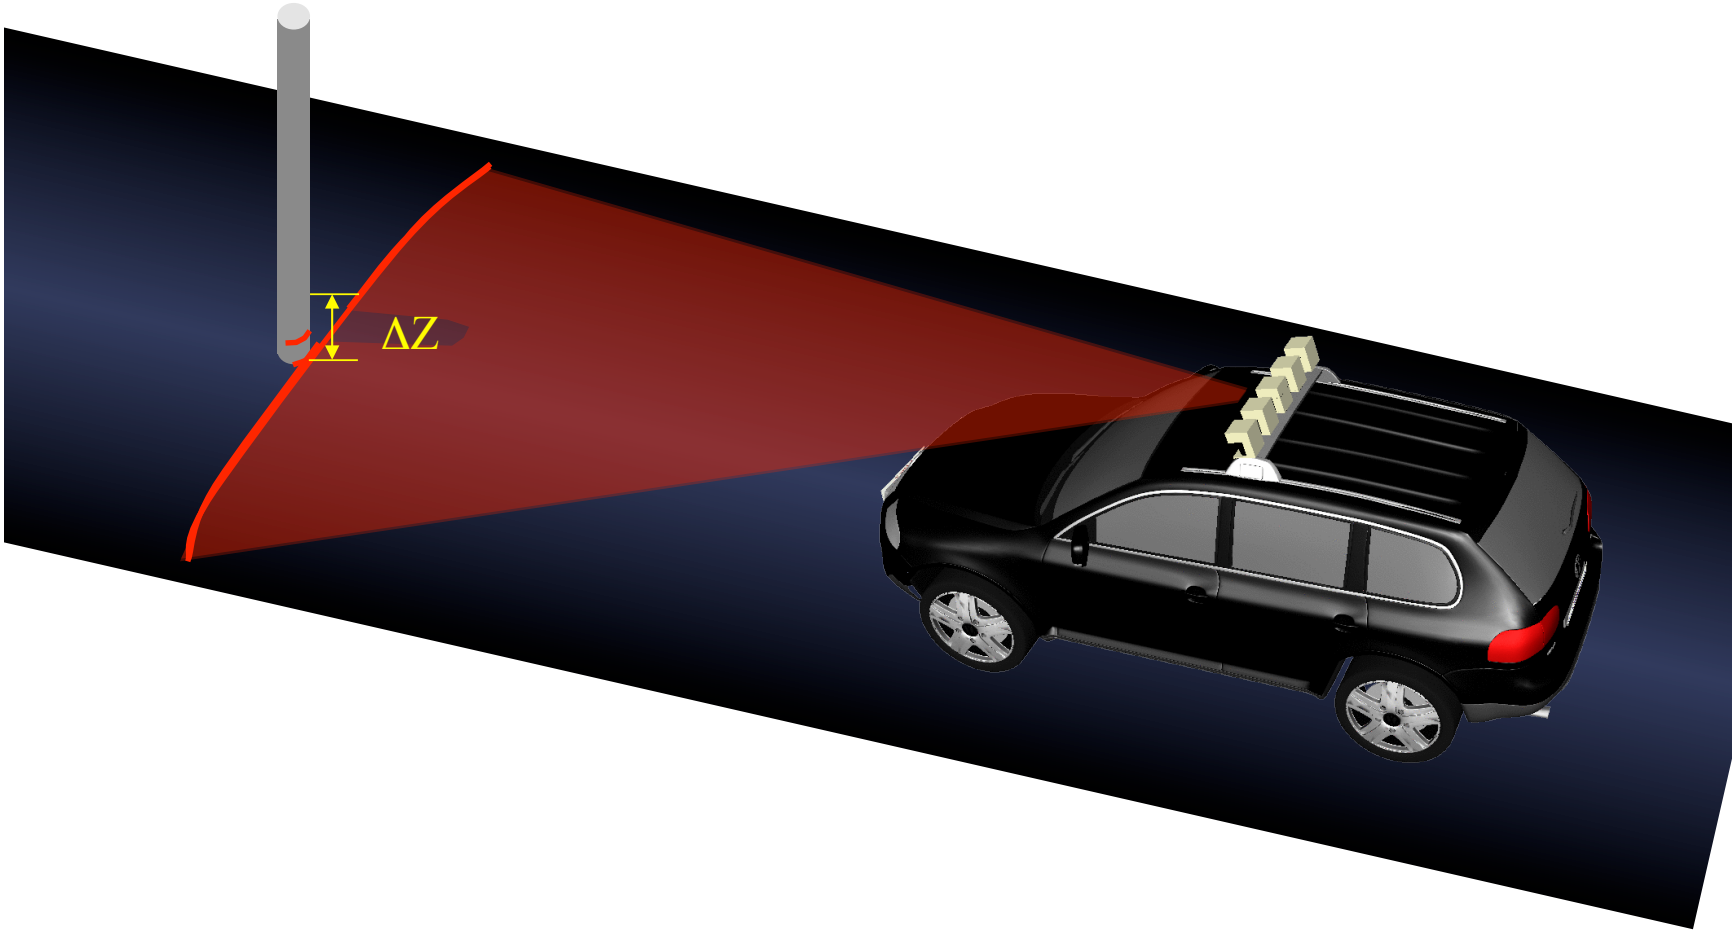
ΔZ

Obstacle Detection

Trigger if $|Z^i - Z^j| > 15\text{cm}$ for nearby z^i, z^j

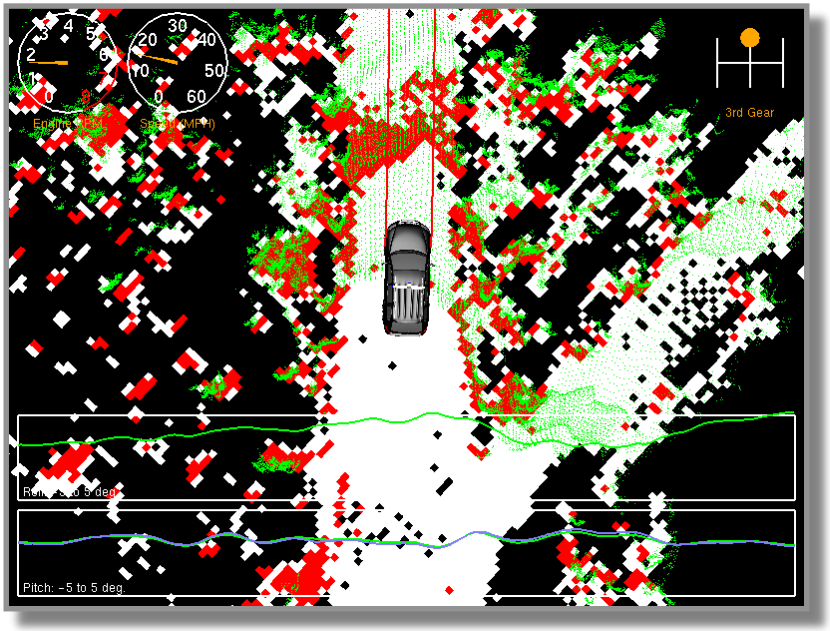

Raw Measurements: 12.6% false positives

Probabilistic Error Model

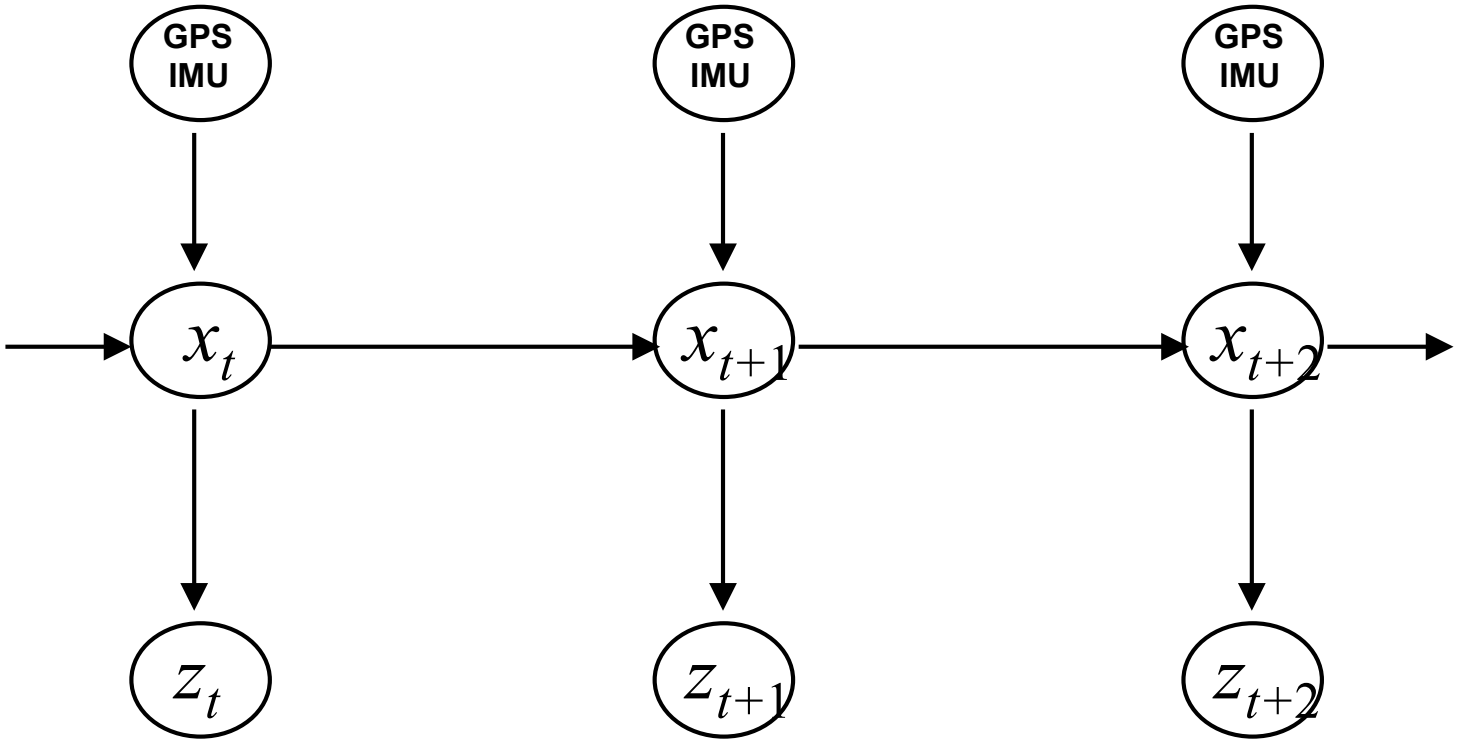

HMMs for Obstacle Detection

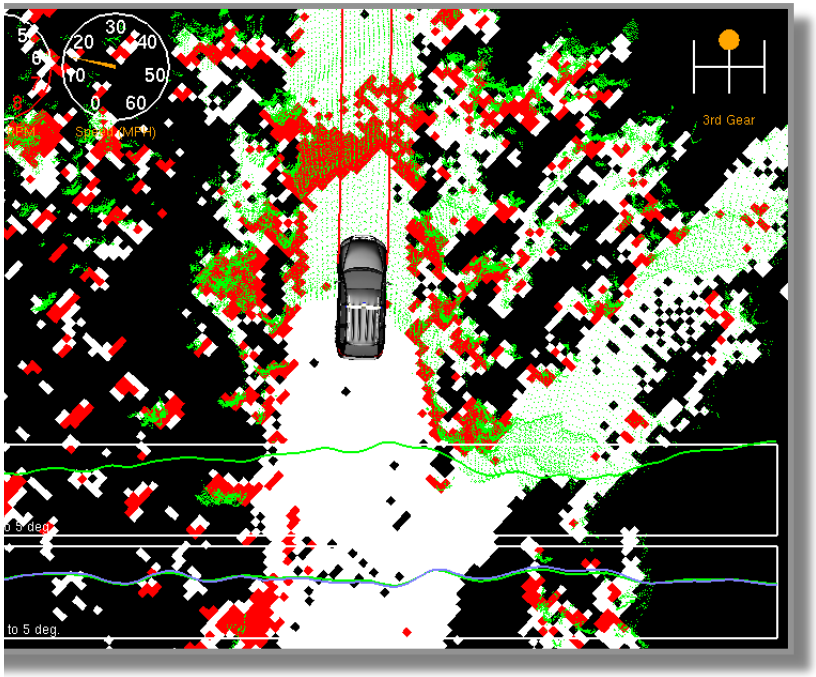

Measurements: 12.6% false positives

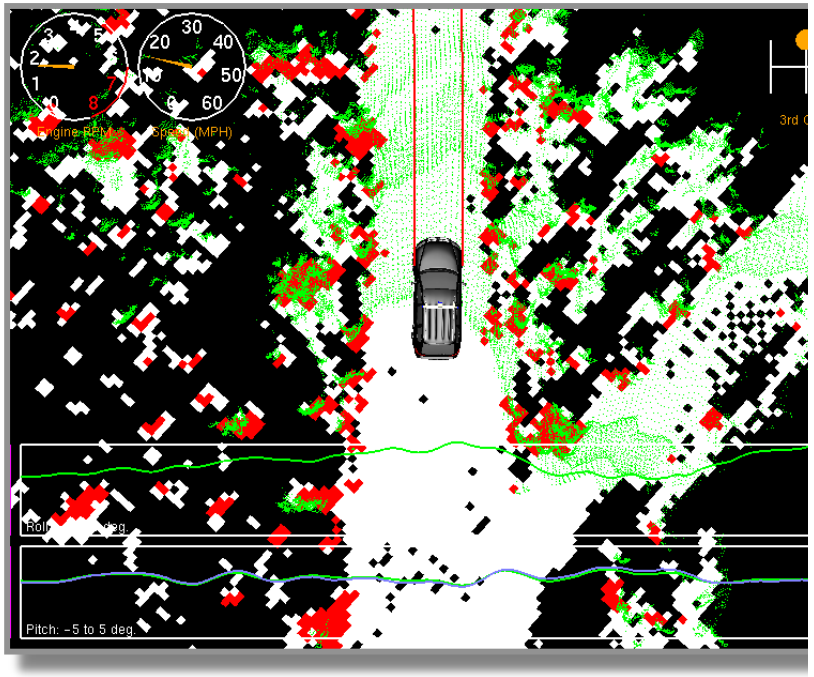

HMM Inference: 0.02% false positive

Road Detection

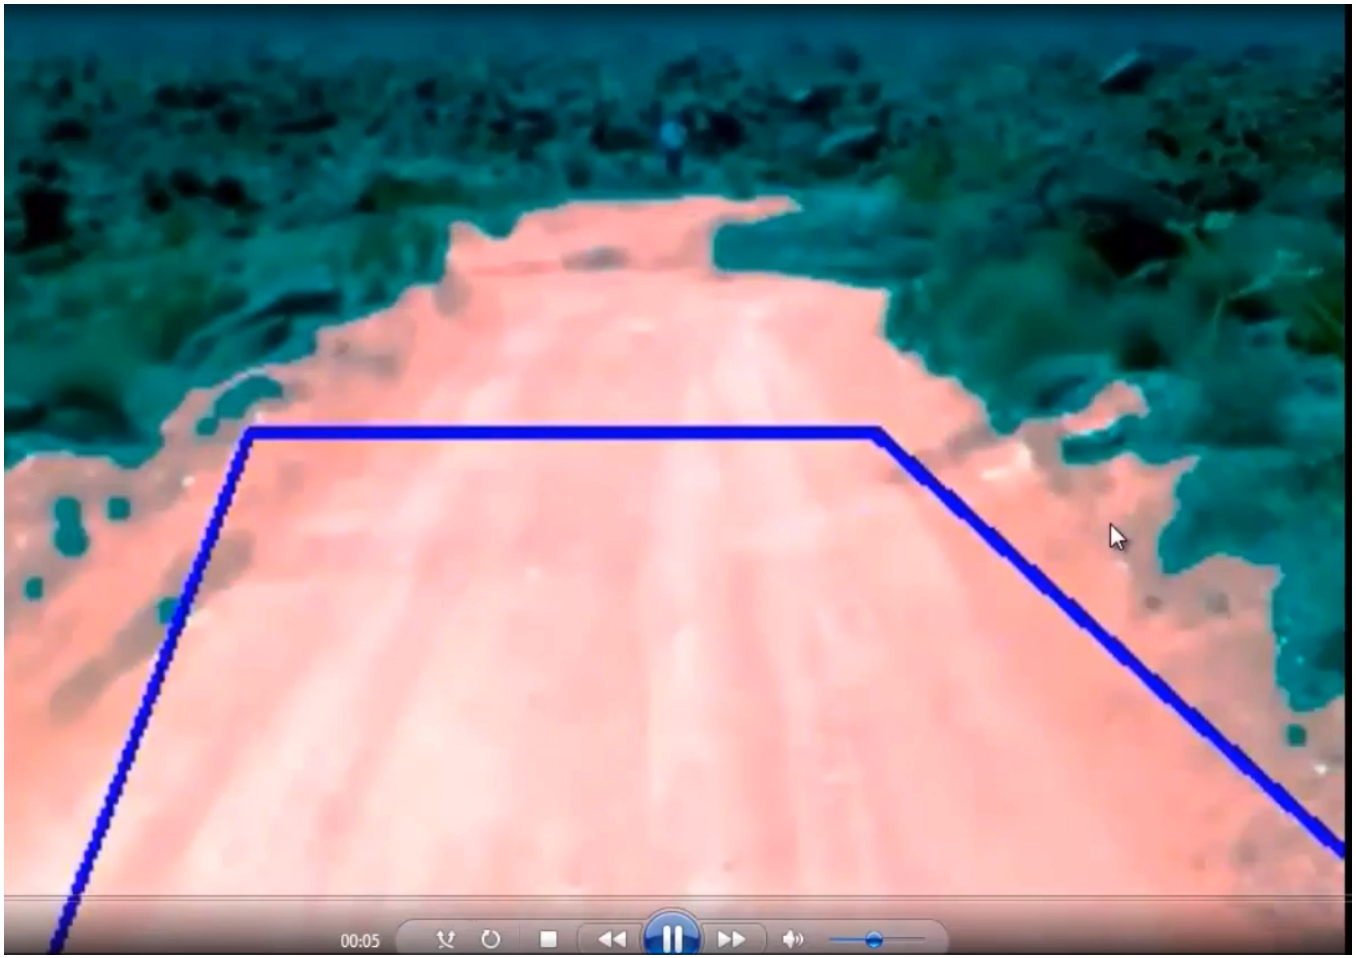

Now on the Streets