
Introduction to Recognition

Saurabh Gupta

Computer Vision

To extract “meaning” from pixels






person, motorcycle, car, chair

Meaning can take different forms:

- Geometric Inferences
- Semantic Inferences
- Inferences about actions
- ...

Computer vision is easy for humans

- Effortlessly analyze images for a variety of tasks
- Infer semantics even from severely ablated
- Can also make precise inference about certain geometric properties



Yet has proven very hard for computers

- Computer vision research easily goes back 60 years ...

MASSACHUSETTS INSTITUTE OF TECHNOLOGY
PROJECT MAC

Artificial Intelligence Group July 7, 1966
Vision Memo. No. 100.

THE SUMMER VISION PROJECT
Seymour Papert

The summer vision project is an attempt to use our summer workers effectively in the construction of a significant part of a visual system. The particular task was chosen partly because it can be segmented into sub-problems which will allow individuals to work independently and yet participate in the construction of a system complex enough to be a real landmark in the development of "pattern recognition".




- Entirely true as of 2014 (or so) when this [xkcd](#) was published

Why is computer vision hard?


- Images are a lossy projection of the world

3D world



Point of observation

2D image



Geometry information is lost



Why is computer vision hard?

- Images are a lossy projection of the world

What color is the dress?

- A) Black and blue
- B) White and gold?

Appearance
information is
also lost



Why is computer vision hard?

- Images are a lossy projection of the world



Might cause
objects to blend

Why is computer vision hard?

- Images are a lossy projection of the world (geometry, appearance, ... are lost)
- Visual world is diverse



Viewpoint variation



Shape variation

Why is computer vision hard?

- Images are a lossy projection of the world (geometry, appearance, ... are lost)
- Visual world is diverse



Background clutter



Occlusion

Why is computer vision hard?

- Images are a lossy projection of the world (geometry, appearance, ... are lost)
 - need some priors to interpret what you are seeing
- Visual world is diverse
 - can't write down these priors by hand



John's Diner with John's Chevelle, 2007

Enter machine learning

Why machine learning?

- Good old-fashioned AI (GOF AI) answer:
Program expertise into the agent




Figure 3. Structure of description built by the system.

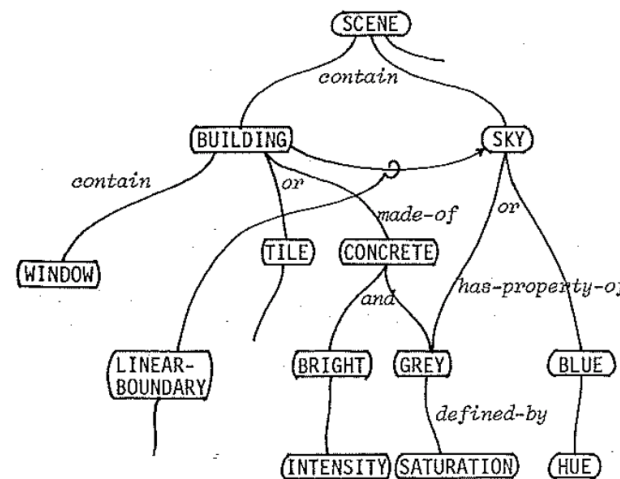


Figure 2. Semantic network for knowledge representation.

lock as a set of
; it must satisfy
The rules have




Figure 5-a. Digitized color scene.

5-b. Result of preliminary segmentation.

5-c. Plan image.

5-d. Result of semantic segmentation.

Why machine learning?

- Good old-fashioned AI (GOF AI) answer:
Program expertise into the agent

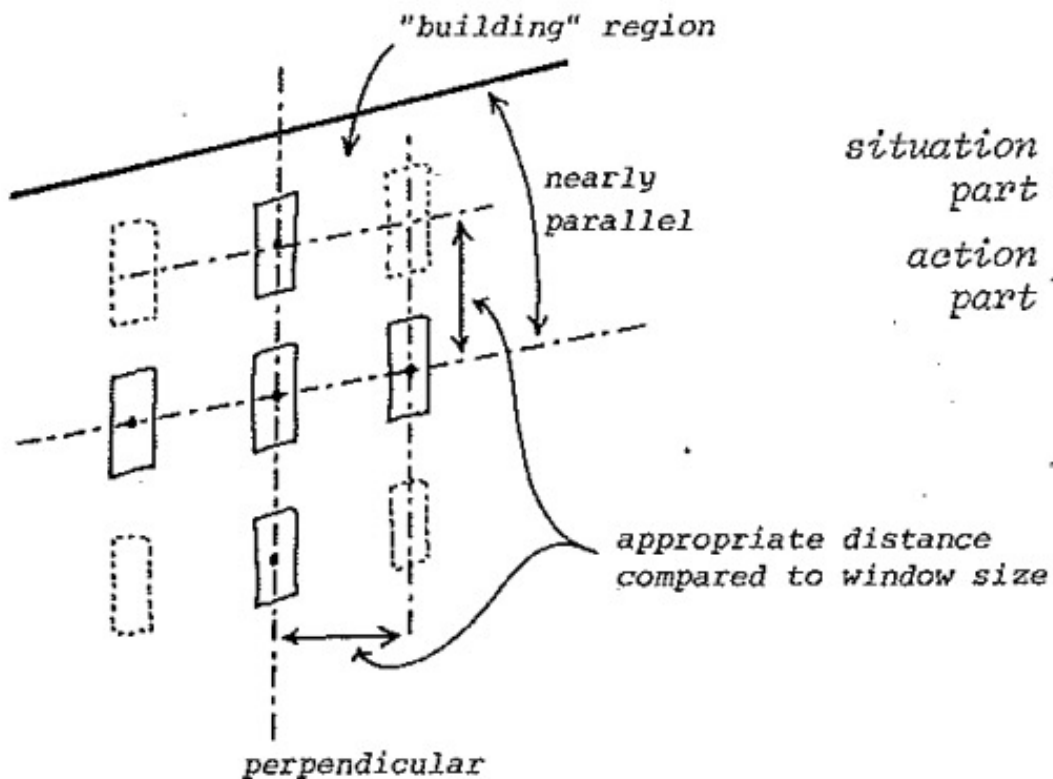


Figure 4-a. "Building" region and "windows".

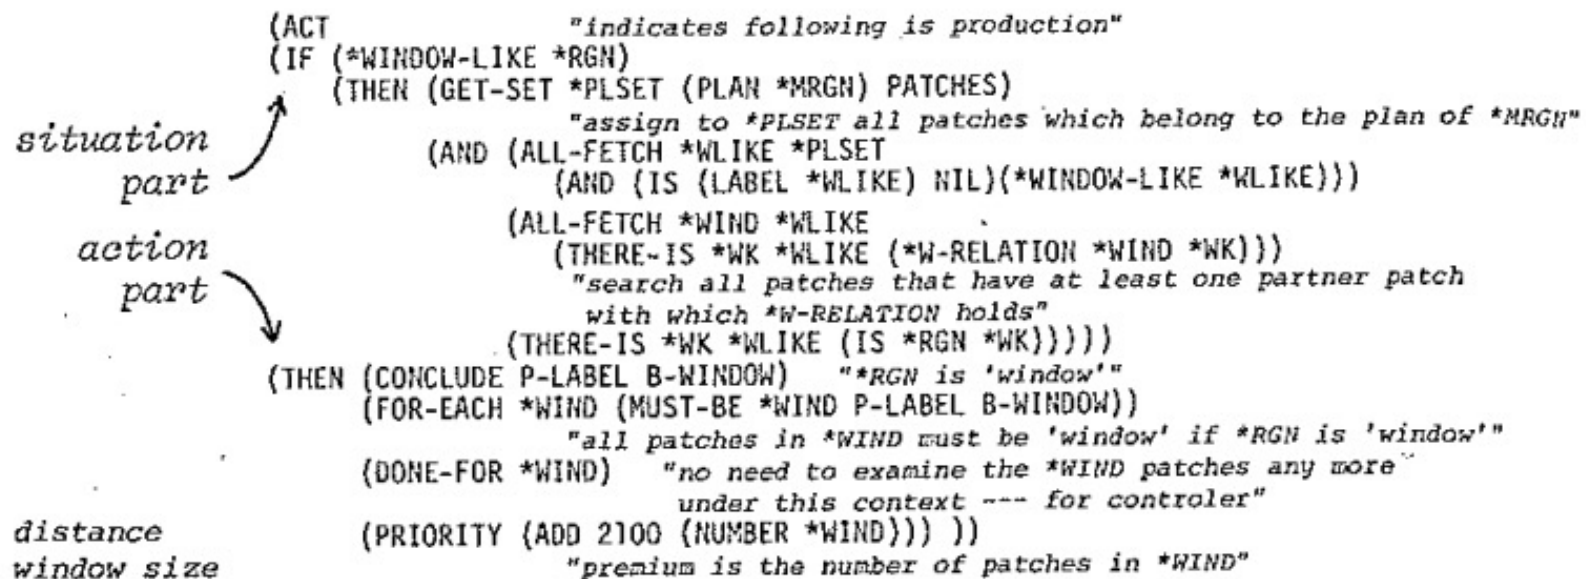


Figure 4-b. The production for analyzing "windows".

Why machine learning?

- Good old-fashioned AI (GOF AI) answer: Program expertise into the agent

Appendix-B Complete Listing of the Model

```
!SCENE knowledge-block-of-scene
(OBJECTS (ASKY #TREE #BUILDING #ROAD #WINDOW)
 SUB-OBJECTS (#B-WINDOW #CAR #C-SHADOW)
 KEY-PATCH-IS ((GREATERP (AREA #POCH 300) #POCH))

PLAN-IMAGE-GENERATION (GVY (BOUNDARY-LENGTH #POCH #POCH)
 (MULT (R-LS-DIFFERENCE #POCH #POCH)
 (BOUNDARY-CONTRACT #POCH #POCH))
 (#POCH #POCH))

IF-PLAN-IS-MODIFIED (IF-DONE (
 rule-for-horizon-detection
 (FACT (IF (IS (OF HORIZON (SCENE)) NIL)
 (ALL-FETCH #RSON #PLAN-REGIONS
 (IF (AND (NOT (PROBABLY ROAD #RSON)
 (NOT (TOUCHING #RSON (LS-SIDE))
 (ALL-FETCH #RSON #PLAN-REGIONS
 (IF (AND (MAY-BE ROAD #RSON)
 #ASCIC #RSON #RSON)
 (NOT (#SAME-COLOR #RSON #RSON))
 (#FACING HORIZONTALLY #RSON #RSON))
 (MULT (SUB (FACING HORIZONTALLY #RSON #RSON) 0.5)
 (SUB (RTR (ASK-VALUE ROAD #RSON) 0.5)
 (ASK-VALUE ROAD #RSON) ))))
 (VALUE #RSON #RSON) ))
 (THEN (MEMO (SCENE) ROAD-ZONE
 (WITH (MFR-LOA-SIDE #RSON 256 1 256))
 (MEMO (SCENE) HORIZON (MFR-LOA-SIDE #RSON))
 (EXECUTE PLAN-EVALUATION) )) ))

P-SELECT (TO-DO (
 rule-for-initial-start
 (FACT (AND (PROBABLY BUILDING #POCH (NOTFOUND BUILDING))

 rule-for-tree-occlusion
 (FACT (AND (#DARK #POCH (#UPPER #POCH)
 (OR (TOUCHING #POCH (UP-SIDE) (TOUCHING #POCH (SIDE))
 (THERE-IS #TR #REGIONS
 (AND (IS (LABEL #TR) TREE)
 (#ADVE #POCH #TR)
 (TOUCHING #TR (SIDE)
 (#WITH-THZ #POCH (V-ZONE #TR))))))
 (THEN (CONCLUDE P-LABEL TREE)
 (CONCLUDE D-HERSE (WITH (OCCLUDE #TR FRAME))
 (SCORE-IS 1.0))) (#POCH) ))

 rule-for-tree-garbage
 (FACT (PROBABLY TREE #POCH)
 (THEN (CONCLUDE P-LABEL TREE)
 (SCORE-IS (ASK-VALUE TREE #POCH))) (#POCH) ))

P-LABEL (IF-DONE (
 if-done-rule-to-be-activated-when-keypatch-is-labeled
 (FACT (NOT (IS (OF PLAN #POCH) NIL))
 (THEN (EXECUTE PLAN-EVALUATION))) (#POCH) ))

!SKY knowledge-block-of-sky
(PROPERTY-RULES (
 (GEN (NOT (#LARGER #RSON (1.0 . 0.5)) (#RSON)
 (GEN (#SHINING #RSON (1.0 . 0.2)) (#RSON))
 (GEN (OR (#BLUE #RSON (#GREY #RSON)) (1.0 . 0.2)) (#RSON)
 (GEN (NOT (#TEXTURAL #RSON) (1.0 . 0.7)) (#RSON)
 (STR (TOUCHING #RSON (UP-SIDE) (0.7 . 0.2)) (#RSON) ))

 RELATION-RULES (
 (STR (AND (#LINEAR-BOUNDARY #RSON #RSON)
 (IF (#LINEAR-BOUNDARY (POSITION (OAN #RSON #RSON))
 (0.0 . 0.5) FOR SKY) (#RSON #RSON))
 (STR (IF (NOT (IS (OF BUILDING-ZONE (SCENE)) NIL))
 (#FUZZY (O-RATIO #RSON (OF BUILDING-ZONE (SCENE))) 0.5 0.3)
 (0.0 . 0.5) FOR SCENE) (#RSON))

!BUILDING knowledge-block-of-building
(MADE-OF (OR (#CONCRETE #TITLE #BRICK)
 SUB-OBJECTS (##AR #WINDOW))

 PROPERTY-RULES (
 (GEN (#MIDDLE #RSON (0.5 . 0.3)) (#RSON)
 (STR (#SMOOTH #RSON (0.5 . 0.2)) (#RSON)
 (STR (#HEAVY-TEXTURE #RSON (0.5 . 0.5)) (#RSON) ))

 RELATION-RULES (
 (GEN (AND (#LINEAR-BOUNDARY #RSON #RSON)
 (IF (#LINEAR-BOUNDARY (NOT (POSITION (UP #RSON #RSON))
 (0.0 . 0.4) FOR SKY) (#RSON #RSON))
 (STR (IF (NOT (IS (OF BUILDING-ZONE (SCENE)) NIL))
 (AND (O-RATIO #RSON (OF BUILDING-ZONE (SCENE)))
 (#FUZZY (R-0.3) FOR SCENE) (#RSON))
 (0.5 . 0.3) FOR SCENE) (#RSON) ))

P-SELECT (
 TO-DO (
 (FACT (AND (MAY-BE BUILDING #POCH (SAME-ZONE #POCH #RSON)
 (THEN (CONCLUDE P-LABEL BUILDING)
 (CONCLUDE R-HERSE #RSON)
 (SCORE-IS (ADD 2.0 (ASK-VALUE BUILDING #POCH))) (#POCH #RSON)
 (FACT (AND (NOT (IS-PLAN #POCH #RSON) (SAME-ZONE #POCH #RSON)
 (MAY-BE BUILDING (PLAN #POCH))
 (THEN (CONCLUDE P-LABEL BUILDING)
 (CONCLUDE R-HERSE #RSON)
 (SCORE-IS (ADD 1.50 (ASK-VALUE BUILDING (PLAN #POCH))))
 (#POCH #RSON))

 rule-for-window-extraction
 (FACT (IF (AND (IS-PLAN #POCH #RSON) (SAME-ZONE #POCH #RSON)
 (VERTICALLY-LONG #POCH (#CONTACT #POCH (PLAN #RSON))
 (THEN (GET-SET #PSET (PLAN #RSON) #PATCHES)
 (AND (#ALL-FETCH #VALUE #PSET
 (AND (IS (LABEL #VALUE) NIL)
 (SAME-ZONE #VALUE #RSON)
 (VERTICALLY-LONG #VALUE)
 (#CONTACT #VALUE (PLAN #RSON))))

 THERE-IS #SK #VALUE (##-RELATION #POCH #RSON)
 (ALL-FETCH #RSON #VALUE
 (##-RELATION #VALUE #RSON))
 (THEN (CONCLUDE P-LABEL B-WINDOW)
 (FOR-EACH #RSON (AND (MULT-SE #AND P-LABEL B-WINDOW)
 (#SAME-FOR #RSON))
 (SCORE-IS (ADD 2.1 (GVY (NUMBER-OF #RSON) (#RSON))))
 (#POCH #RSON))

 (FACT (AND (IS-PLAN #POCH #RSON) (SAME-ZONE #POCH #RSON)
 (THEN (CONCLUDE P-LABEL BUILDING)
 (CONCLUDE R-HERSE #RSON)
 (SCORE-IS 2.0))) (#POCH #RSON))

 D-HERSE (IF-DONE (
 (FACT #R (#DESCRIBE-BUILDING (REGION #POCH))) (#POCH) ))

 O-CREATE (IF-DONE (
 (FACT #R (#EXTRACT-BUILDING-SHAPE (REGION #POCH)
 (#DESCRIBE-BUILDING (REGION #POCH))
 (EXECUTE PLAN-EVALUATION))) (#POCH) ))

 #PRIORI-VALUE-IS 0.2

!ROAD knowledge-block-of-road
(MADE-OF (OR (#GRAVEL #CONCRETE)
 SUB-OBJECTS (##AR #SHADOW)

 PROPERTY-RULES (
 (GEN (#LOSER #RSON (0.0 . 0.4)) (#RSON)
 (GEN (#HORIZONTALLY-LONG #RSON (0.7 . 0.2)) (#RSON)
 (STR (TOUCHING #RSON (USER-SIDE) (0.0 . 0.2)) (#RSON) ))

 RELATION-RULES (
 (STR (AND (#SAME-COLOR #RSON #RSON) (TOUCHING #RSON #RSON)
 (0.0 . 0.2) FOR ROAD) (#RSON #RSON))
 (STR (IF (NOT (IS (OF HORIZON (SCENE)) NIL))
 (O-RATIO #RSON (OF ROAD-ZONE (SCENE)))
 (1.0 . 0.3) FOR SCENE) (#RSON))
```

```
(THEN (CONCLUDE P-LABEL BUILDING)
 (SCORE-IS (ADD 4.0 (CONFIDENCE-VALUE #POCH))) (#POCH)
 (FACT (AND (PROBABLY ROAD #POCH (NOTFOUND ROAD))
 (THEN (CONCLUDE P-LABEL ROAD)
 (SCORE-IS (ADD 4.0 (CONFIDENCE-VALUE #POCH))) (#POCH)
 (FACT (AND (PROBABLY SKY #POCH (NOTFOUND SKY))
 (THEN (CONCLUDE P-LABEL SKY)
 (SCORE-IS (ADD 4.0 (CONFIDENCE-VALUE #POCH))) (#POCH)
 (FACT (AND (PROBABLY TREE #POCH)
 (NOT (THERE-IS #TR #REGIONS
 (AND (IS (LABEL #TR) TREE)
 (OR (TOUCHING (PLAN #POCH) (PLAN #TR)
 (#WITH-THZ (PLAN #POCH)
 (V-ZONE 00 (PLAN #TR)))))))))
 (THEN (CONCLUDE P-LABEL TREE)
 (SCORE-IS (ADD 4.0 (CONFIDENCE-VALUE #POCH))) (#POCH)

 rule-for-adjacent-wall-of-building
 (FACT (AND (MAY-BE BUILDING #POCH)
 (THERE-IS #BL #REGIONS
 (AND (IS (LABEL #BL) BUILDING)
 (NOT (IS (OF SHAPE VIEW (OBJECT #BL) ))
 (IS (OF ADJACENT (OBJECT #BL) NIL)
 (DIFFERENT-ZONE #POCH #BL)))
 (THEN (CONCLUDE P-LABEL BUILDING)
 (CONCLUDE D-HERSE (WITH (ADJACENT #BL)
 (SCORE-IS (ADD 5.0 (ASK-VALUE BUILDING #POCH))) (#POCH)

 rule-for-building-occlusion
 (FACT (AND (MAY-BE BUILDING #POCH)
 (THERE-IS #BL #REGIONS
 (AND (IS (LABEL #BL) BUILDING)
 (SAME-ZONE #POCH #BL)
 (#SAME-COLOR #POCH #BL)
 (THERE-IS #TR #KEYPATCHES
 (AND (BETWEEN #TR #POCH #BL)
 (OR (IS (LABEL #TR) TREE)
 (AND (IS (LABEL #TR) BUILDING)
 (NOT (IS (OBJECT #BL)
 (OBJECT #TR)))))))))
 (THEN (CONCLUDE P-LABEL BUILDING)
 (CONCLUDE D-HERSE (WITH (OCCLUDE #BL (REGION #TR))
 (SCORE-IS (ADD 1.0 (ASK-VALUE BUILDING #POCH))) (#POCH)

P-SELECT (
 TO-DO (
 (FACT (MAY-BE SKY #POCH)
 (THEN (SCORE-IS (ADD 2.0 (ASK-VALUE SKY #POCH))) (#POCH)
 (FACT (AND (IS-PLAN #POCH #RSON) (#BRIGHT #POCH)
 (THEN (SCORE-IS 2.0)) (#POCH #RSON))
 (FACT (#BRIGHT #POCH) (THEN (SCORE-IS 0.05))) (#POCH) ))

 IF-DONE (
 (FACT #R (#THEN (CONCLUDE P-LABEL SKY)
 (CONCLUDE R-HERSE (#MASTER #POCH))) (#POCH) ))

 #PRIORI-VALUE-IS 0.1

!TREE knowledge-block-of-tree
(MADE-OF (#LEAVES)

 PROPERTY-RULES (
 (GEN (#MIDDLE #RSON (0.5 . 0.3)) (#RSON)
 (STR (#HEAVY-TEXTURE #RSON (0.8 . 0.2)) (#RSON) ))

P-SELECT (
 TO-DO (
 (FACT (MAY-BE TREE #POCH)
 (THEN (SCORE-IS (ADD 2.0 (ASK-VALUE TREE #POCH))) (#POCH)
 (FACT (AND (IS-PLAN #POCH #RSON) (NOT (#SHINING #POCH))
 (THEN (SCORE-IS 3.0))) (#POCH #RSON))

 IF-DONE (
 (FACT #R (#THEN (CONCLUDE P-LABEL TREE)
 (CONCLUDE R-HERSE (#MASTER #POCH))) (#POCH) ))

 #PRIORI-VALUE-IS 0.2
```

```
!BUILDING knowledge-block-of-building
(MADE-OF (OR (#CONCRETE #TITLE #BRICK)
 SUB-OBJECTS (##AR #WINDOW))

 PROPERTY-RULES (
 (GEN (#MIDDLE #RSON (0.5 . 0.3)) (#RSON)
 (STR (#SMOOTH #RSON (0.5 . 0.2)) (#RSON)
 (STR (#HEAVY-TEXTURE #RSON (0.5 . 0.5)) (#RSON) ))

 RELATION-RULES (
 (GEN (AND (#LINEAR-BOUNDARY #RSON #RSON)
 (IF (#LINEAR-BOUNDARY (NOT (POSITION (UP #RSON #RSON))
 (0.0 . 0.4) FOR SKY) (#RSON #RSON))
 (STR (IF (NOT (IS (OF BUILDING-ZONE (SCENE)) NIL))
 (AND (O-RATIO #RSON (OF BUILDING-ZONE (SCENE)))
 (#FUZZY (R-0.3) FOR SCENE) (#RSON))
 (0.5 . 0.3) FOR SCENE) (#RSON) ))

P-SELECT (
 TO-DO (
 (FACT (AND (MAY-BE BUILDING #POCH (SAME-ZONE #POCH #RSON)
 (THEN (CONCLUDE P-LABEL BUILDING)
 (CONCLUDE R-HERSE #RSON)
 (SCORE-IS (ADD 2.0 (ASK-VALUE BUILDING #POCH))) (#POCH #RSON)
 (FACT (AND (NOT (IS-PLAN #POCH #RSON) (SAME-ZONE #POCH #RSON)
 (MAY-BE BUILDING (PLAN #POCH))
 (THEN (CONCLUDE P-LABEL BUILDING)
 (CONCLUDE R-HERSE #RSON)
 (SCORE-IS (ADD 1.50 (ASK-VALUE BUILDING (PLAN #POCH))))
 (#POCH #RSON))

 rule-for-window-extraction
 (FACT (IF (AND (IS-PLAN #POCH #RSON) (SAME-ZONE #POCH #RSON)
 (VERTICALLY-LONG #POCH (#CONTACT #POCH (PLAN #RSON))
 (THEN (GET-SET #PSET (PLAN #RSON) #PATCHES)
 (AND (#ALL-FETCH #VALUE #PSET
 (AND (IS (LABEL #VALUE) NIL)
 (SAME-ZONE #VALUE #RSON)
 (VERTICALLY-LONG #VALUE)
 (#CONTACT #VALUE (PLAN #RSON))))

 THERE-IS #SK #VALUE (##-RELATION #POCH #RSON)
 (ALL-FETCH #RSON #VALUE
 (##-RELATION #VALUE #RSON))
 (THEN (CONCLUDE P-LABEL B-WINDOW)
 (FOR-EACH #RSON (AND (MULT-SE #AND P-LABEL B-WINDOW)
 (#SAME-FOR #RSON))
 (SCORE-IS (ADD 2.1 (GVY (NUMBER-OF #RSON) (#RSON))))
 (#POCH #RSON))

 (FACT (AND (IS-PLAN #POCH #RSON) (SAME-ZONE #POCH #RSON)
 (THEN (CONCLUDE P-LABEL BUILDING)
 (CONCLUDE R-HERSE #RSON)
 (SCORE-IS 2.0))) (#POCH #RSON))

 D-HERSE (IF-DONE (
 (FACT #R (#DESCRIBE-BUILDING (REGION #POCH))) (#POCH) ))

 O-CREATE (IF-DONE (
 (FACT #R (#EXTRACT-BUILDING-SHAPE (REGION #POCH)
 (#DESCRIBE-BUILDING (REGION #POCH))
 (EXECUTE PLAN-EVALUATION))) (#POCH) ))

 #PRIORI-VALUE-IS 0.2

!ROAD knowledge-block-of-road
(MADE-OF (OR (#GRAVEL #CONCRETE)
 SUB-OBJECTS (##AR #SHADOW)

 PROPERTY-RULES (
 (GEN (#LOSER #RSON (0.0 . 0.4)) (#RSON)
 (GEN (#HORIZONTALLY-LONG #RSON (0.7 . 0.2)) (#RSON)
 (STR (TOUCHING #RSON (USER-SIDE) (0.0 . 0.2)) (#RSON) ))

 RELATION-RULES (
 (STR (AND (#SAME-COLOR #RSON #RSON) (TOUCHING #RSON #RSON)
 (0.0 . 0.2) FOR ROAD) (#RSON #RSON))
 (STR (IF (NOT (IS (OF HORIZON (SCENE)) NIL))
 (O-RATIO #RSON (OF ROAD-ZONE (SCENE)))
 (1.0 . 0.3) FOR SCENE) (#RSON))
```


Why machine learning?

- Good old-fashioned AI (GOF AI) answer:
Program expertise into the agent
 - Never worked (in general)

Why machine learning?

- Good old-fashioned AI (GOF AI) answer:
Program expertise into the agent
- Modern answer: Program into the agent the *ability to improve performance based on experience*
 - Experience should come from *training data* or *demonstrations*
 - We want to optimize the performance of the agent on the training data, with the hope that it will *generalize* to unseen inputs
 - This is the *statistical learning* viewpoint

The basic ML framework (for supervised learning)



The basic ML framework (for supervised learning)

$$y = f(x)$$

output prediction function input



- **Training (or learning):** given a *training set* of labeled examples $\{(x_1, y_1), \dots, (x_N, y_N)\}$, instantiate a predictor f
- **Testing (or inference):** apply f to a new *test example* x and output the predicted value $y = f(x)$
- Rather than hand-defining how 2D projections of apples are different from pears, f will learn this from the data.

Deep Learning

- A general way to model function f as composition (layers) of simple functions, very loosely inspired by the brain.

Lecture overview


- Different recognition problems in computer vision
- Supervised classification
- Taxonomy of learning problems

Different Recognition Problems




This image by Nikita is licensed under [CC-BY 2.0](https://creativecommons.org/licenses/by/2.0/)


Classification: Assign image to one of a fixed set of categories



Object Detection: Put a bounding box around each instance of a class



Instance Segmentation: Mark pixels for each instance of a class




Semantic Segmentation: Label each pixels with its category

Different Recognition Problems




Image Captioning: Man riding a horse on a beach



Depth Prediction: how far is each pixel in the image

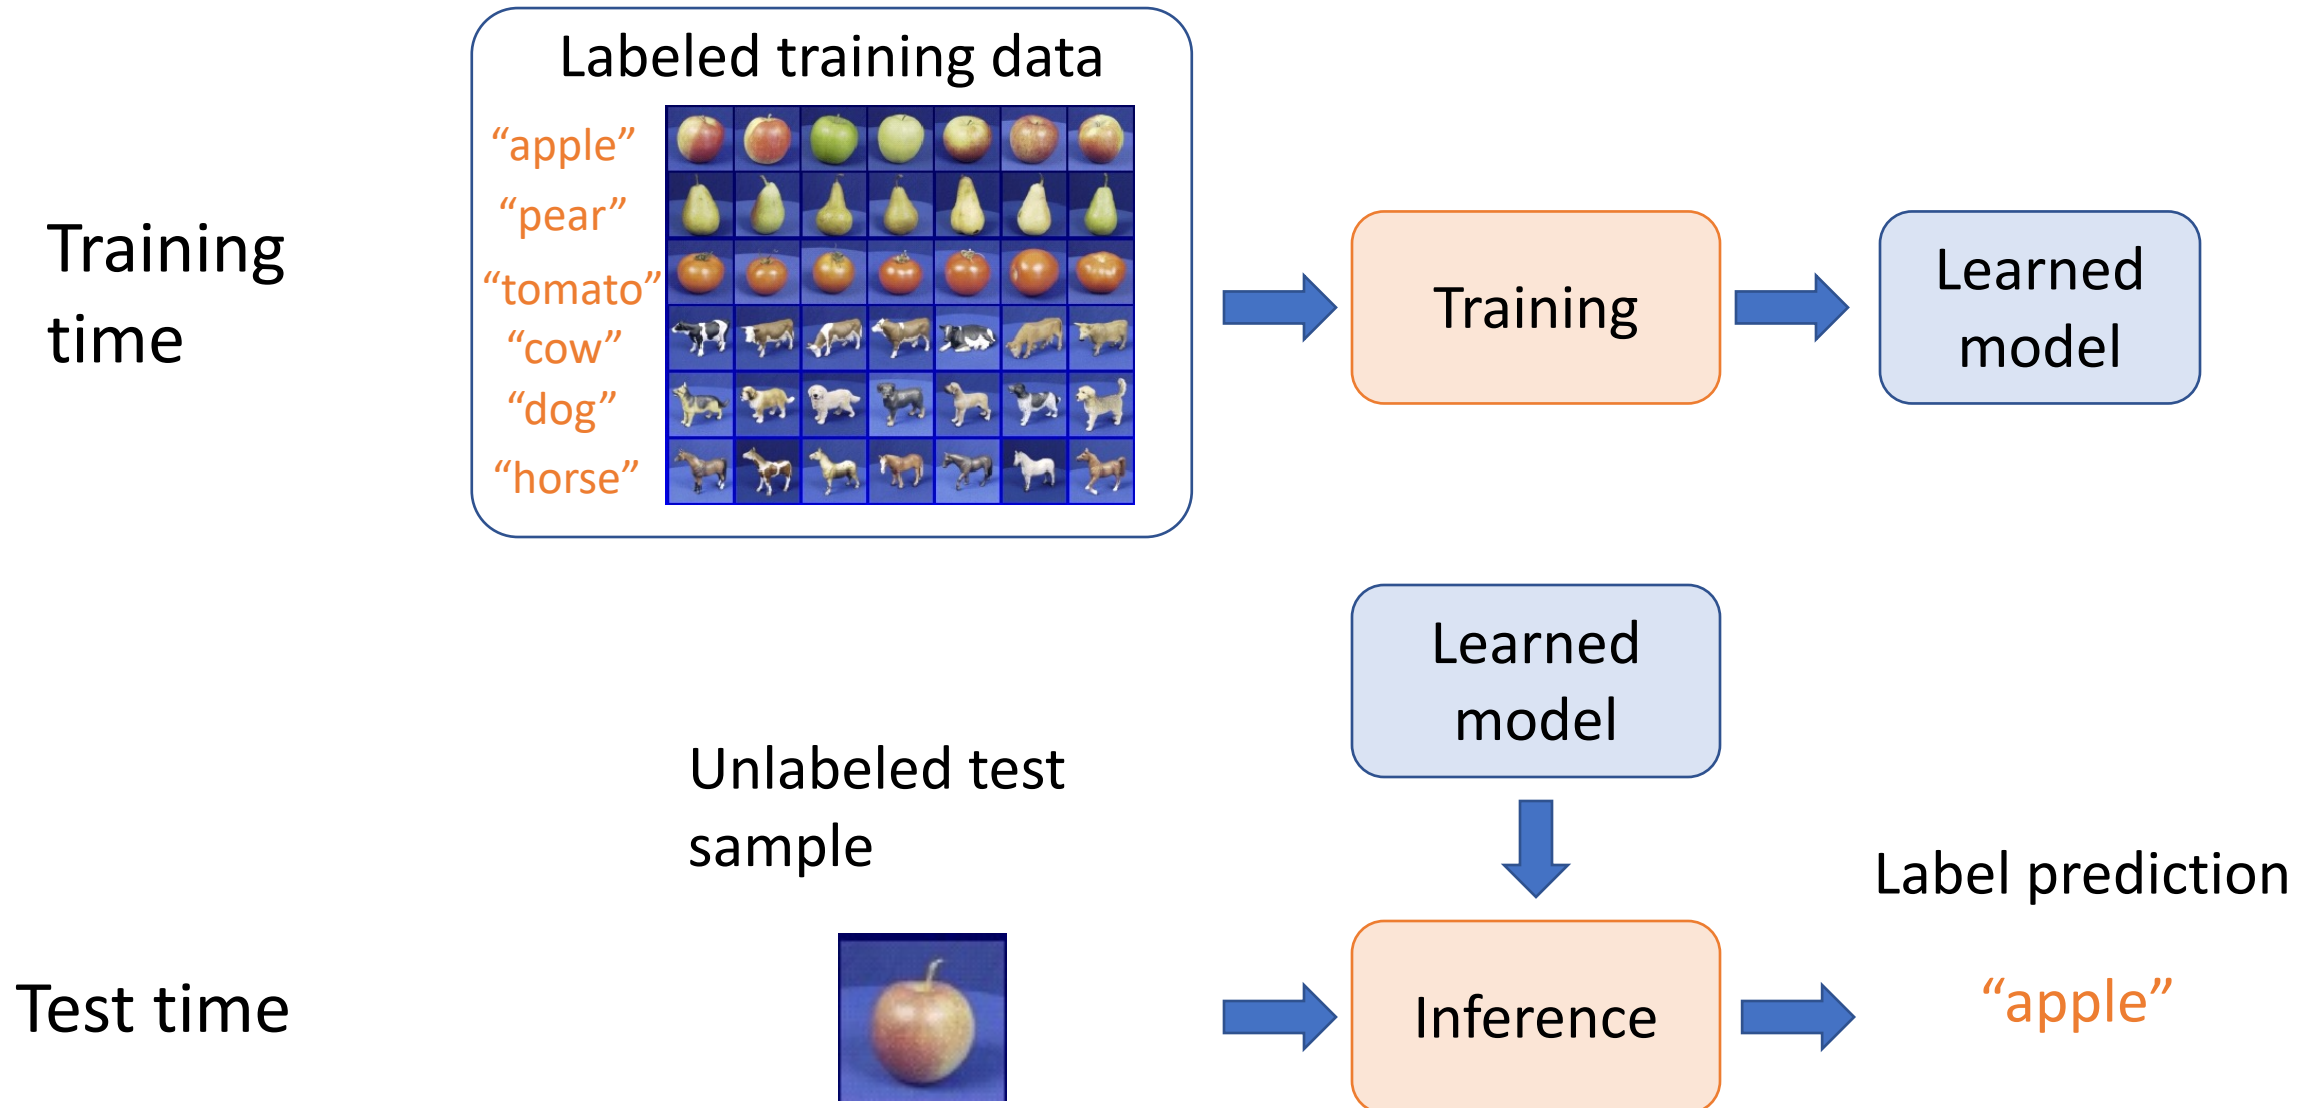


Keypoint prediction



Pose Prediction: Rotation that aligns object to a canonical pose

The basic ML framework (for supervised learning)



The basic ML framework (for supervised learning)

$$y = f(x)$$

↑
output

↑
prediction
function

↙
input



- **Training (or learning):** given a *training set* of labeled examples $\{(x_1, y_1), \dots, (x_N, y_N)\}$, instantiate a predictor f
- **Testing (or inference):** apply f to a new *test example* x and output the predicted value $y = f(x)$
- Rather than hand-defining how 2D projections of apples are different from pears, f will learn this from the data.


Is an image classifier all you need?

- Image Classification
- Object Detection
- Instance Segmentation
- Semantic Segmentation
- Image Captioning
- Depth Prediction
- Keypoint Prediction
- Pose Prediction
- ...

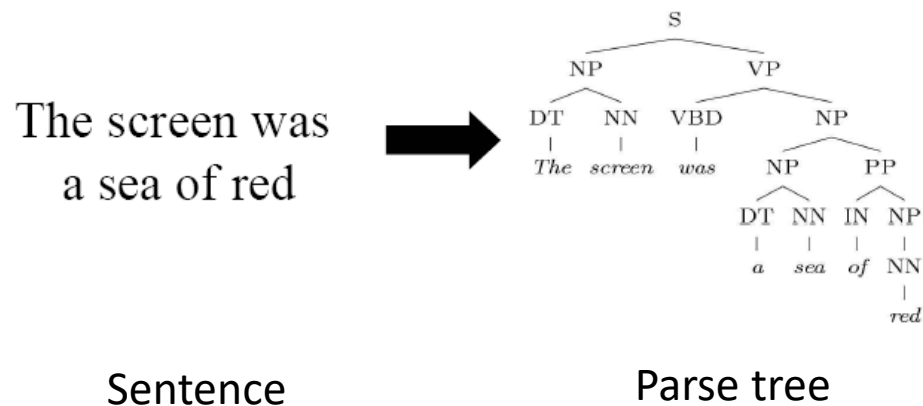


Taxonomy of learning problems

- **Type of output**
 - Classification
 - Regression
 - $y = f(x)$. y is an arbitrary scalar and not a class label.
 - Structured prediction
 - $y = f(x)$. y is a structured object.



Depth Prediction: how far is each pixel in the image



Several computer vision problems have structure in the output space, but often solving a classification problem with some simple post-processing (or even without) ends up being sufficient.