343H: Honors AI

Lecture 18: Decision Networks and VOI

3/27/2014

Kristen Grauman

UT Austin

Slides courtesy of Dan Klein, UC Berkeley

Unless otherwise noted

Recall: Inference in Ghostbusters

- A ghost is in the grid somewhere
- Sensor readings tell how close a square is to the ghost
  - On the ghost: red
  - 1 or 2 away: orange
  - 3 or 4 away: yellow
  - 5+ away: green

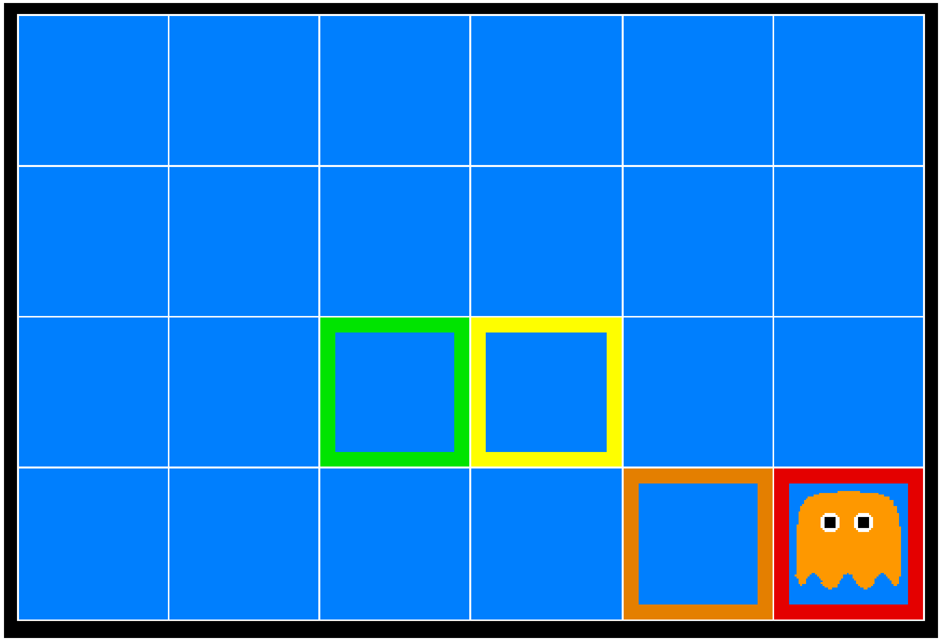

- Sensors are noisy, but we know P(Color | Distance)

P(red | 3)	P(orange | 3)	P(yellow | 3)	P(green | 3)
0.05	0.15	0.5	0.3

Inference in Ghostbusters

Inference in Ghostbusters

- Need to decide when and what to sense!

Decision Networks

- MEU: choose the action which maximizes the expected utility given the evidence

- Can directly operationalize this with decision networks

- New node types:

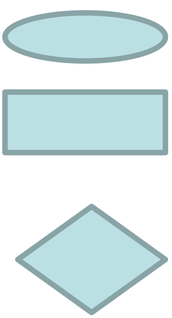

  - Chance nodes (just like BNs)
  - Actions (cannot have parents, act as observed evidence)
  - Utility node (depends on action and chance nodes)

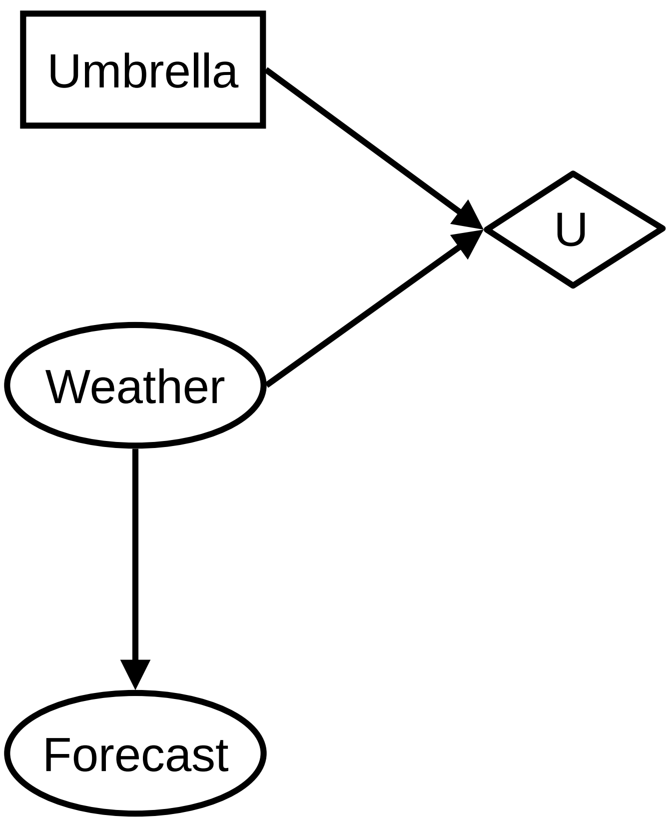

Decision Networks

- Action selection:
  - Instantiate all evidence
  - Set action node(s) each possible way
  - Calculate posterior for all parents of utility node, given the evidence
  - Calculate expected utility for each action
  - Choose maximizing action

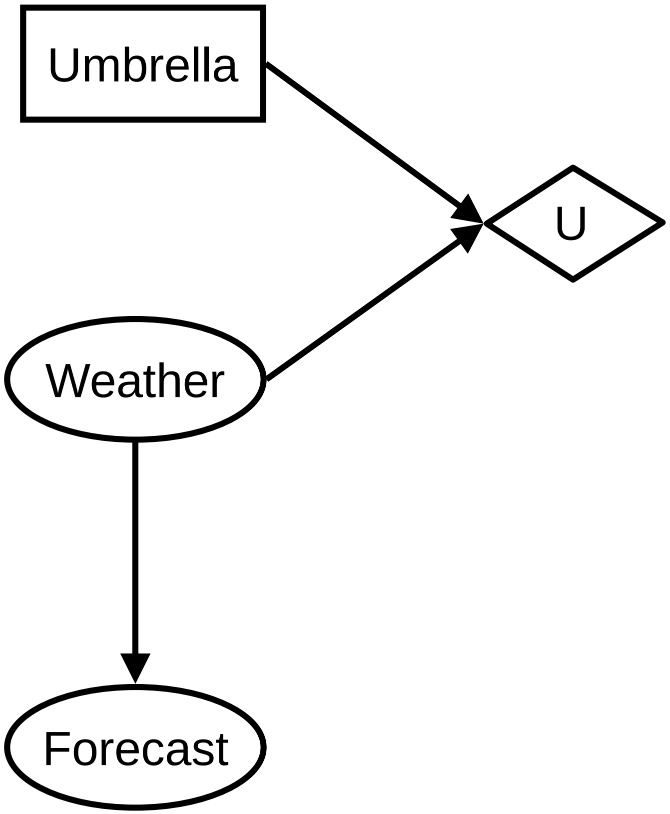

Example: Decision Networks

Umbrella = leave

$$\text{EU}(\text{leave}) = \sum_w P(w)U(\text{leave}, w)$$

$$= 0.7 \cdot 100 + 0.3 \cdot 0 = 70$$

Umbrella = take

$$\text{EU}(\text{take}) = \sum_w P(w)U(\text{take}, w)$$

$$= 0.7 \cdot 20 + 0.3 \cdot 70 = 35$$

Optimal decision = leave

$$\text{MEU}(\emptyset) = \max_a \text{EU}(a) = 70$$

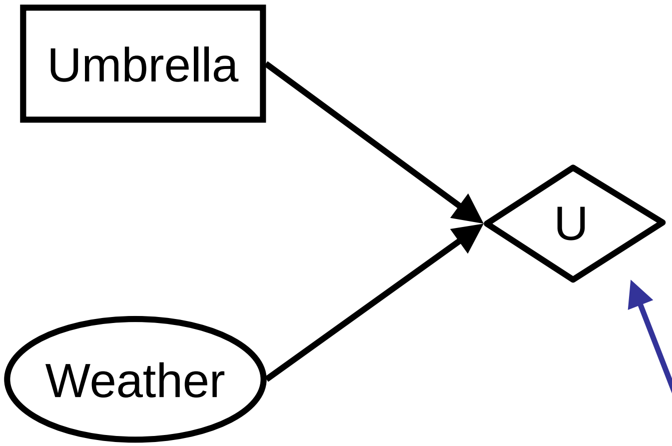

W	P(W)
sun	0.7
rain	0.3

A	W	U(A,W)
leave	sun	100
leave	rain	0
take	sun	20
take	rain	70

Decisions as Outcome Trees

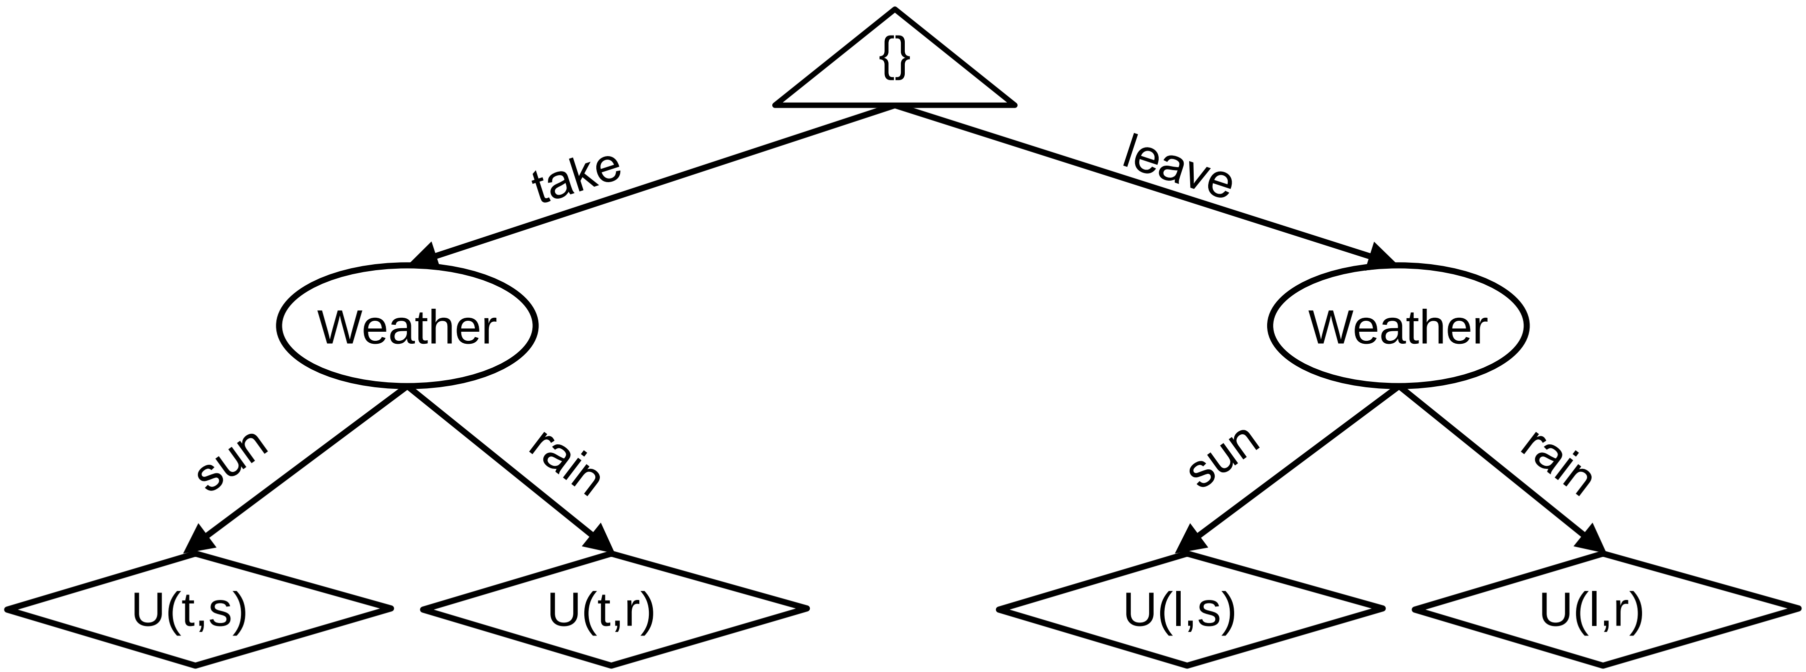

- Almost exactly like expectimax / MDPs
- What’s changed?

Example: Decision Networks

Umbrella = leave

$$\text{EU}(\text{leave}|\text{bad}) = \sum_w P(w|\text{bad})U(\text{leave}, w)$$

$$= 0.34 \cdot 100 + 0.66 \cdot 0 = 34$$

Umbrella = take

$$\text{EU}(\text{take}|\text{bad}) = \sum_w P(w|\text{bad})U(\text{take}, w)$$

$$= 0.34 \cdot 20 + 0.66 \cdot 70 = 53$$

Optimal decision = take

$$\text{MEU}(F = \text{bad}) = \max_a \text{EU}(a|\text{bad}) = 53$$

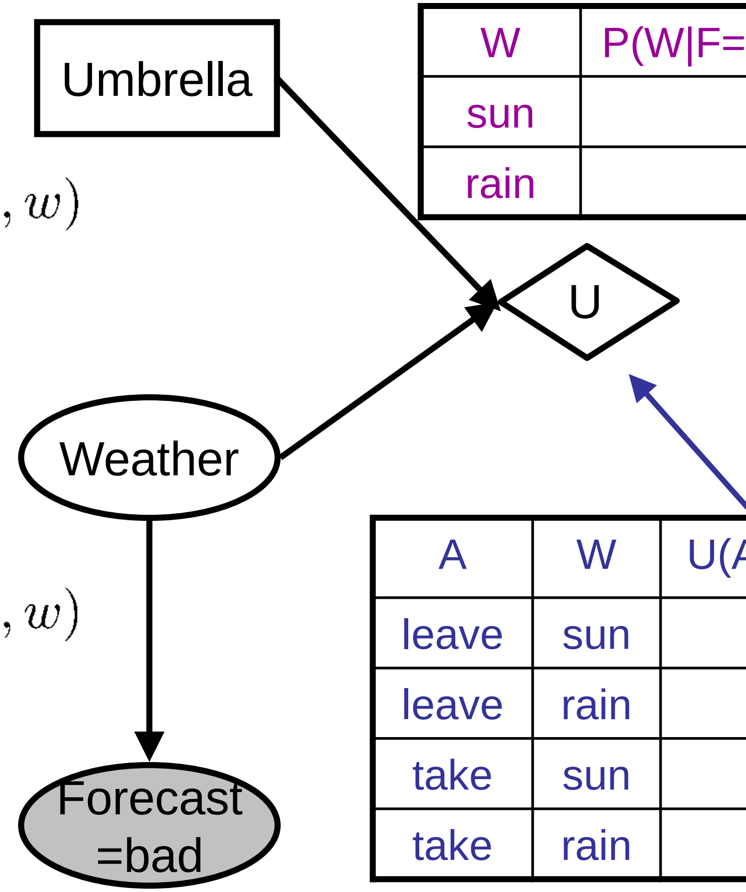

W	P(W|F=bad)
sun	0.34
rain	0.66

A	W	U(A,W)
leave	sun	100
leave	rain	0
take	sun	20
take	rain	70

Decisions as Outcome Trees

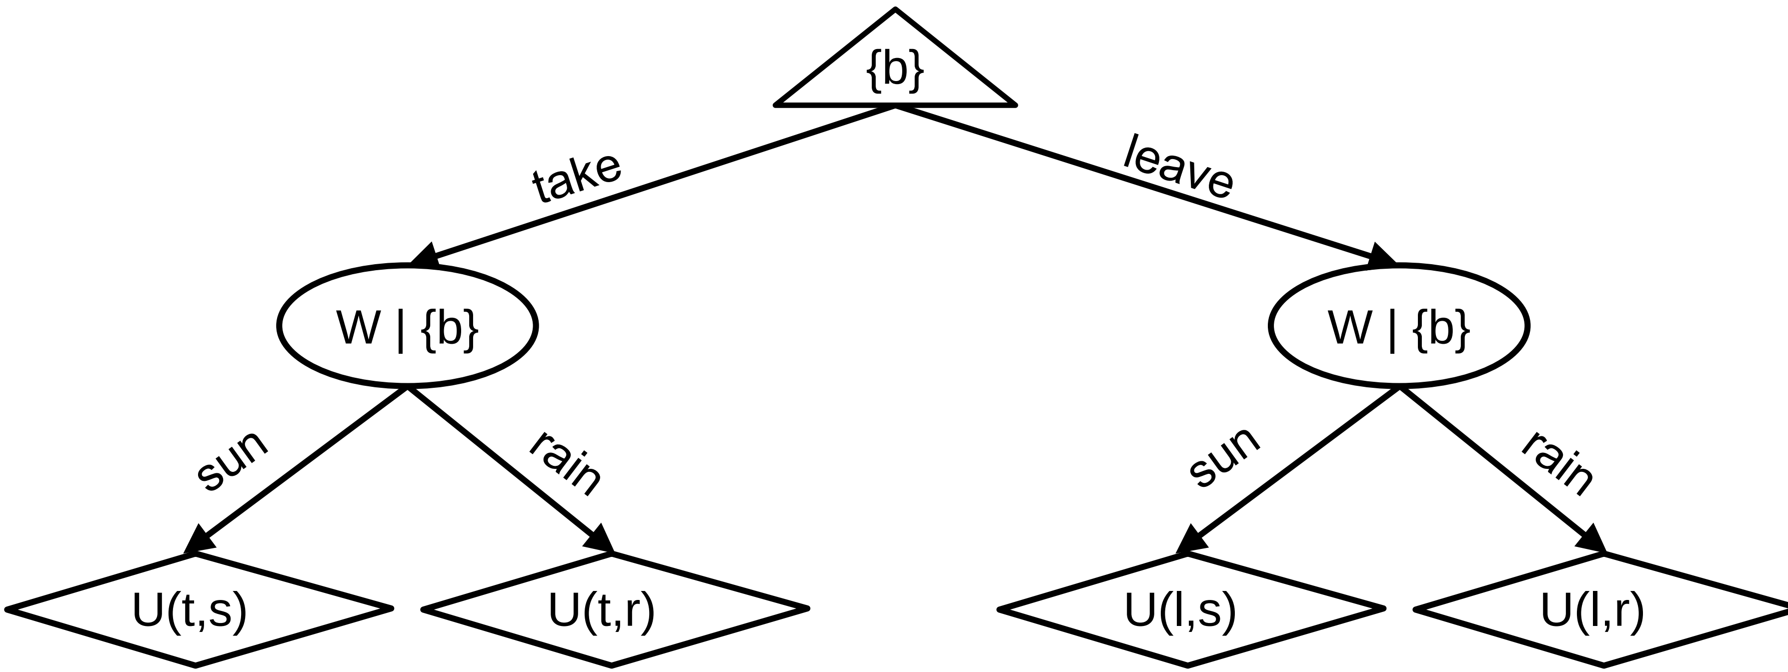

Ghostbusters decision network

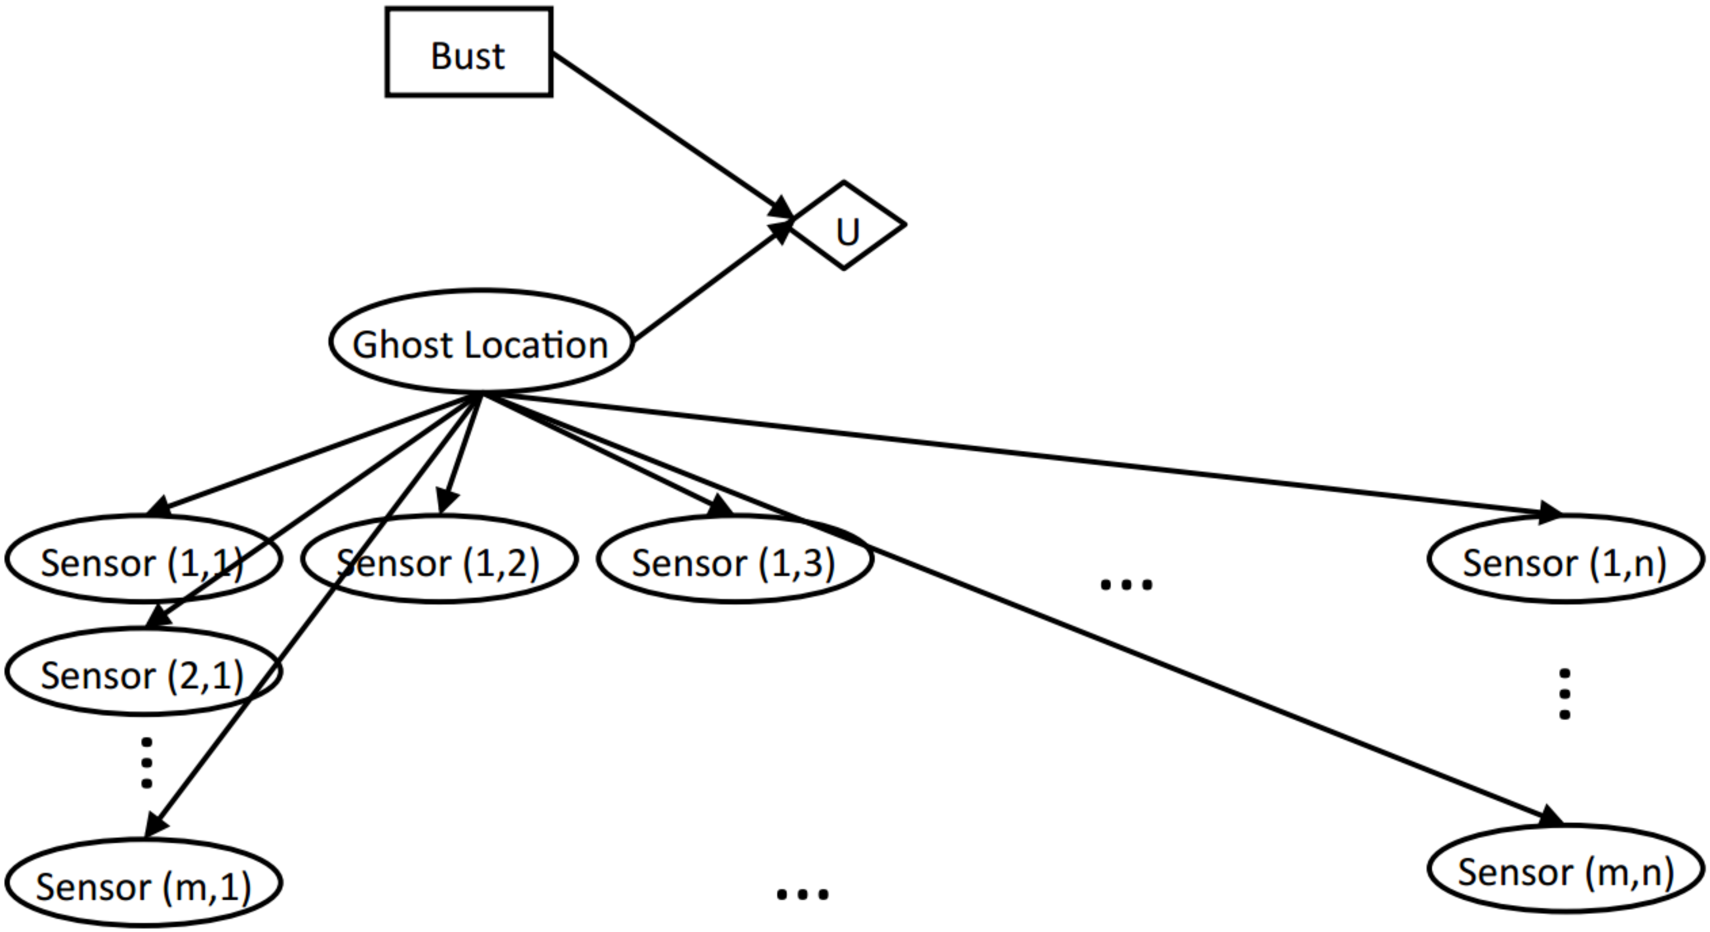

Value of Information

- Idea: compute value of acquiring evidence
  - Can be done directly from decision network

- Example: buying oil drilling rights
  - Two blocks A and B, exactly one has oil, worth k
  - You can drill in one location
  - Prior probabilities 0.5 each, & mutually exclusive
  - Drilling in either A or B has EU = k/2, MEU = k/2

- Question: what's the value of information of O?
  - Value of knowing which of A or B has oil
  - Value is expected gain in MEU from new info
  - Survey may say "oil in a" or "oil in b," prob 0.5 each
  - If we know OilLoc, MEU is k (either way)
  - Gain in MEU from knowing OilLoc?
  - VPI(OilLoc) = k/2
  - Fair price of information: k/2

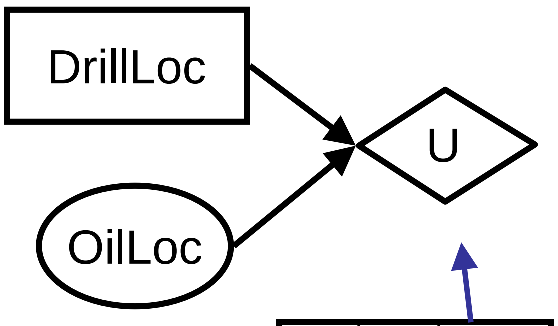

O	P
a	1/2
b	1/2

D	O	U
a	a	k
a	b	0
b	a	0
b	b	k

VPI Example: Weather

MEU with no evidence

$$\mathrm{MEU}(\emptyset) = \max_a \mathrm{EU}(a) = 70$$

MEU if forecast is bad

$$\mathrm{MEU}(F = \mathrm{bad}) = \max_a \mathrm{EU}(a|\mathrm{bad}) = 53$$

MEU if forecast is good

$$\mathrm{MEU}(F = \mathrm{good}) = \max_a \mathrm{EU}(a|\mathrm{good}) = 95$$

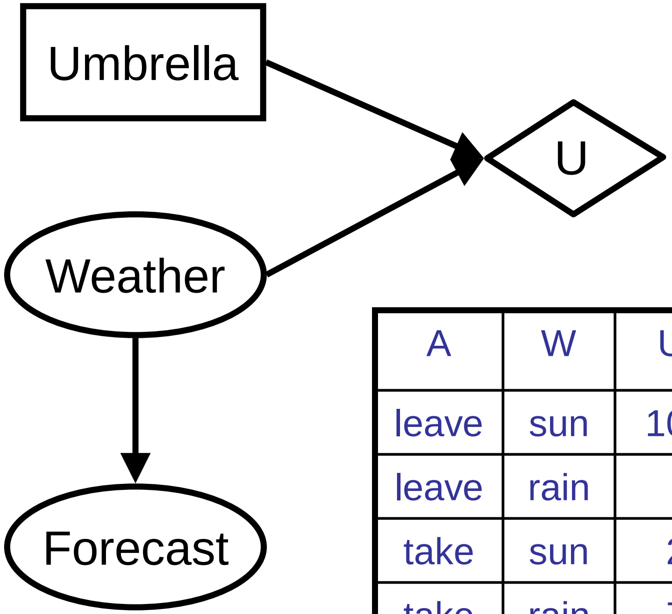

A	W	U
leave	sun	100
leave	rain	0
take	sun	20
take	rain	70

VPI Example: Weather

MEU with no evidence

$$\text{MEU}(\varnothing) = \max_a \text{EU}(a) = 70$$

MEU if forecast is bad

$$\text{MEU}(F = \text{bad}) = \max_a \text{EU}(a|\text{bad}) = 53$$

MEU if forecast is good

$$\text{MEU}(F = \text{good}) = \max_a \text{EU}(a|\text{good}) = 95$$

Forecast distribution

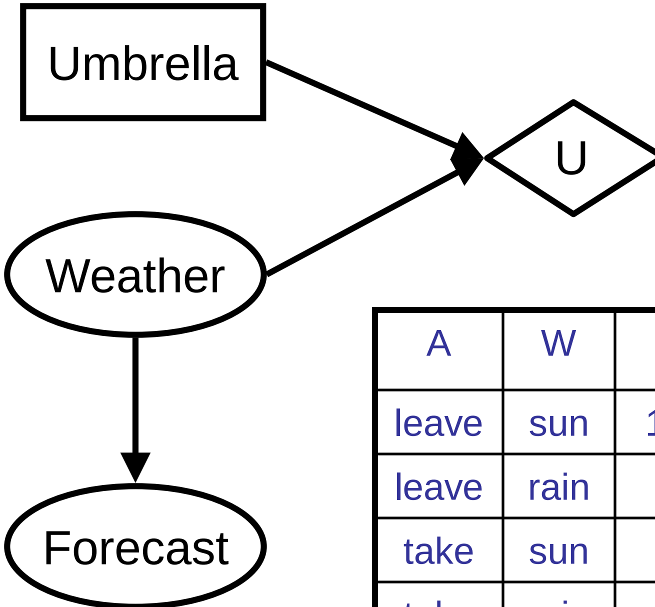

A	W	U
leave	sun	100
leave	rain	0
take	sun	20
take	rain	70

F	P(F)
good	0.59
bad	0.41

$$0.59 \cdot (95) + 0.41 \cdot (53) - 70$$

$$77.8 - 70 = 7.8$$

$$\text{VPI}(E|e') = \left(\sum_{e'} P(e'|e)\text{MEU}(e, e') \right) - \text{MEU}(e)$$

Value of Information

- Assume we have evidence E=e. Value if we act now:

$$\text{MEU}(e) = \max_a \sum_s P(s|e)\, U(s,a)$$

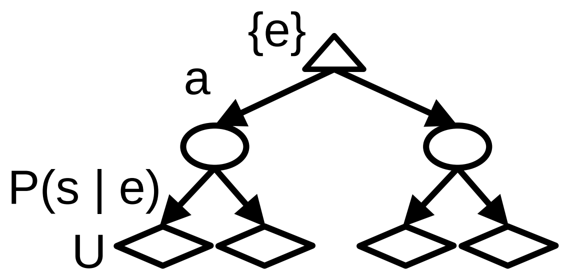

- Assume we see that E' = e'. Value if we act then:

$$\text{MEU}(e,e') = \max_a \sum_s P(s|e,e')\, U(s,a)$$

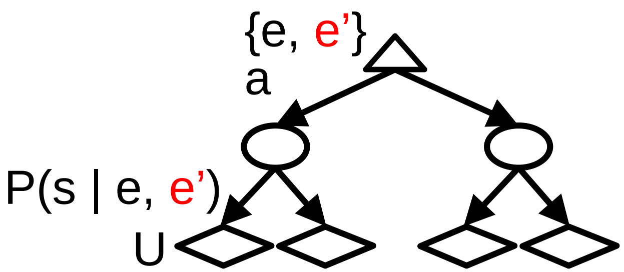

- BUT E' is a random variable whose value is unknown, so we don't know what e' will be

- Expected value if E' is revealed and then we act:

$$\text{MEU}(e,E') = \sum_{e'} P(e'|e)\text{MEU}(e,e')$$

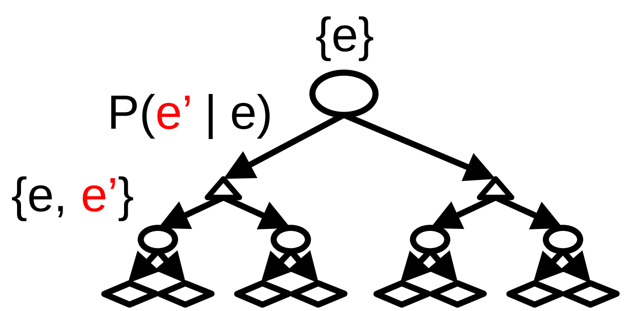

- Value of information: how much MEU goes up by revealing E' first then acting, over acting now:

$$\text{VPI}(E'|e) = \text{MEU}(e,E') - \text{MEU}(e)$$

VPI Properties

- Nonnegative

$$\forall E', e : \text{VPI}(E'|e) \geq 0$$

- Nonadditive – consider, e.g., observing E_j twice

$$\text{VPI}(E_j, E_k|e) \neq \text{VPI}(E_j|e) + \text{VPI}(E_k|e)$$

- Order-independent

$$\begin{aligned}\text{VPI}(E_j, E_k|e) &= \text{VPI}(E_j|e) + \text{VPI}(E_k|e, E_j) \\ &= \text{VPI}(E_k|e) + \text{VPI}(E_j|e, E_k)\end{aligned}$$

Quick VPI Questions

- The soup of the day is either clam chowder or split pea, but you wouldn't order either one. What's the value of knowing which it is?
- There are two kinds of plastic forks at a picnic. One kind is slightly sturdier. What's the value of knowing which?
- You're playing the lottery. The prize will be $0 or $100. You can play any number between 1 and 100 (chance of winning is 1%). What is the value of knowing the winning number?

Value of imperfect information?

- No such thing
- Information corresponds to the observation of a node in the decision network
- If data is “noisy”, that just means we don’t observe the original variable, but another variable which is a noisy version of the original one.

VPI Question

- VPI(OilLoc)?
- VPI(ScoutingReport)?
- VPI(Scout)?
- VPI(Scout | ScoutingReport)?

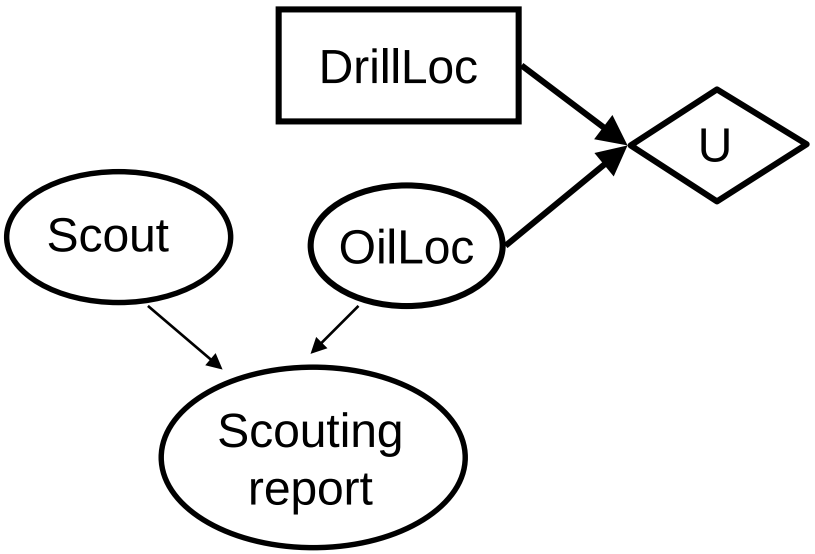

If Parents(U) $\perp\!\!\!\perp$ Z | CurrentEvidence

Then VPI(Z | CurrentEvidence) = 0

Another VPI example

Training an object recognition system: The standard pipeline

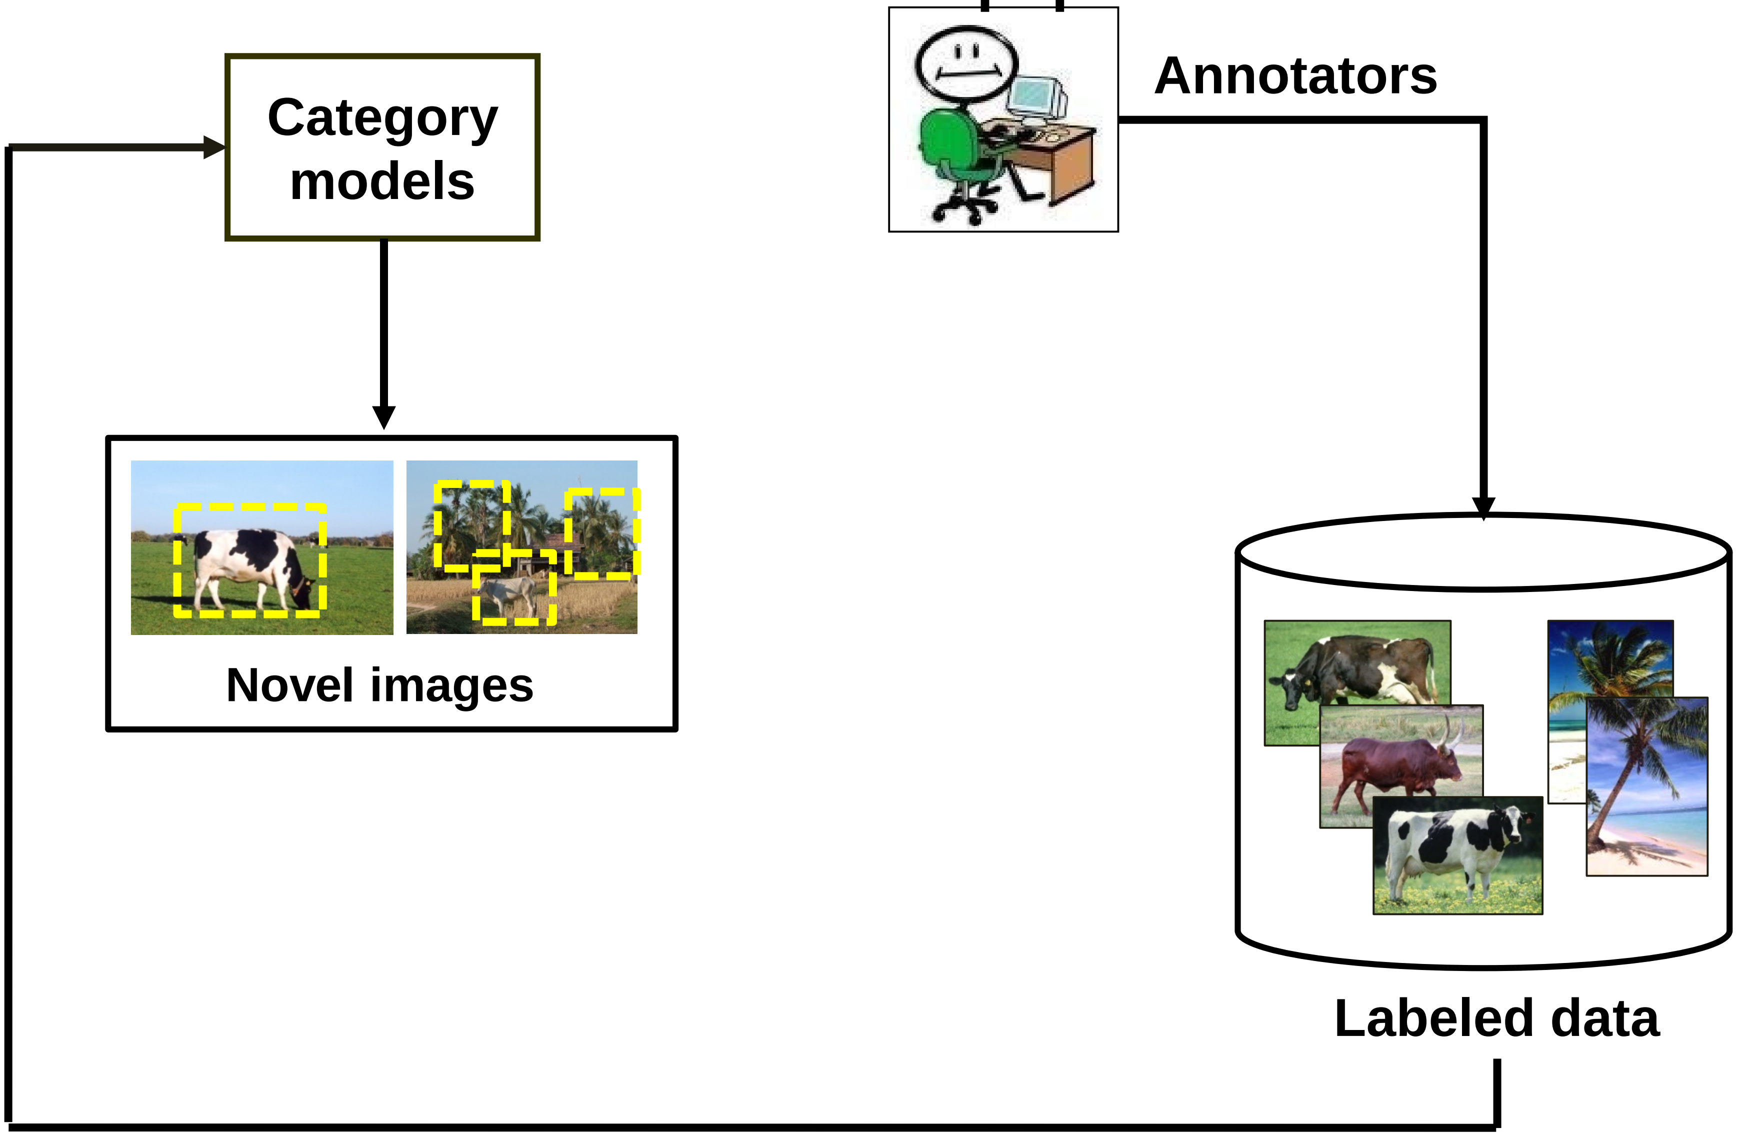

Kristen Grauman

The active visual learning pipeline

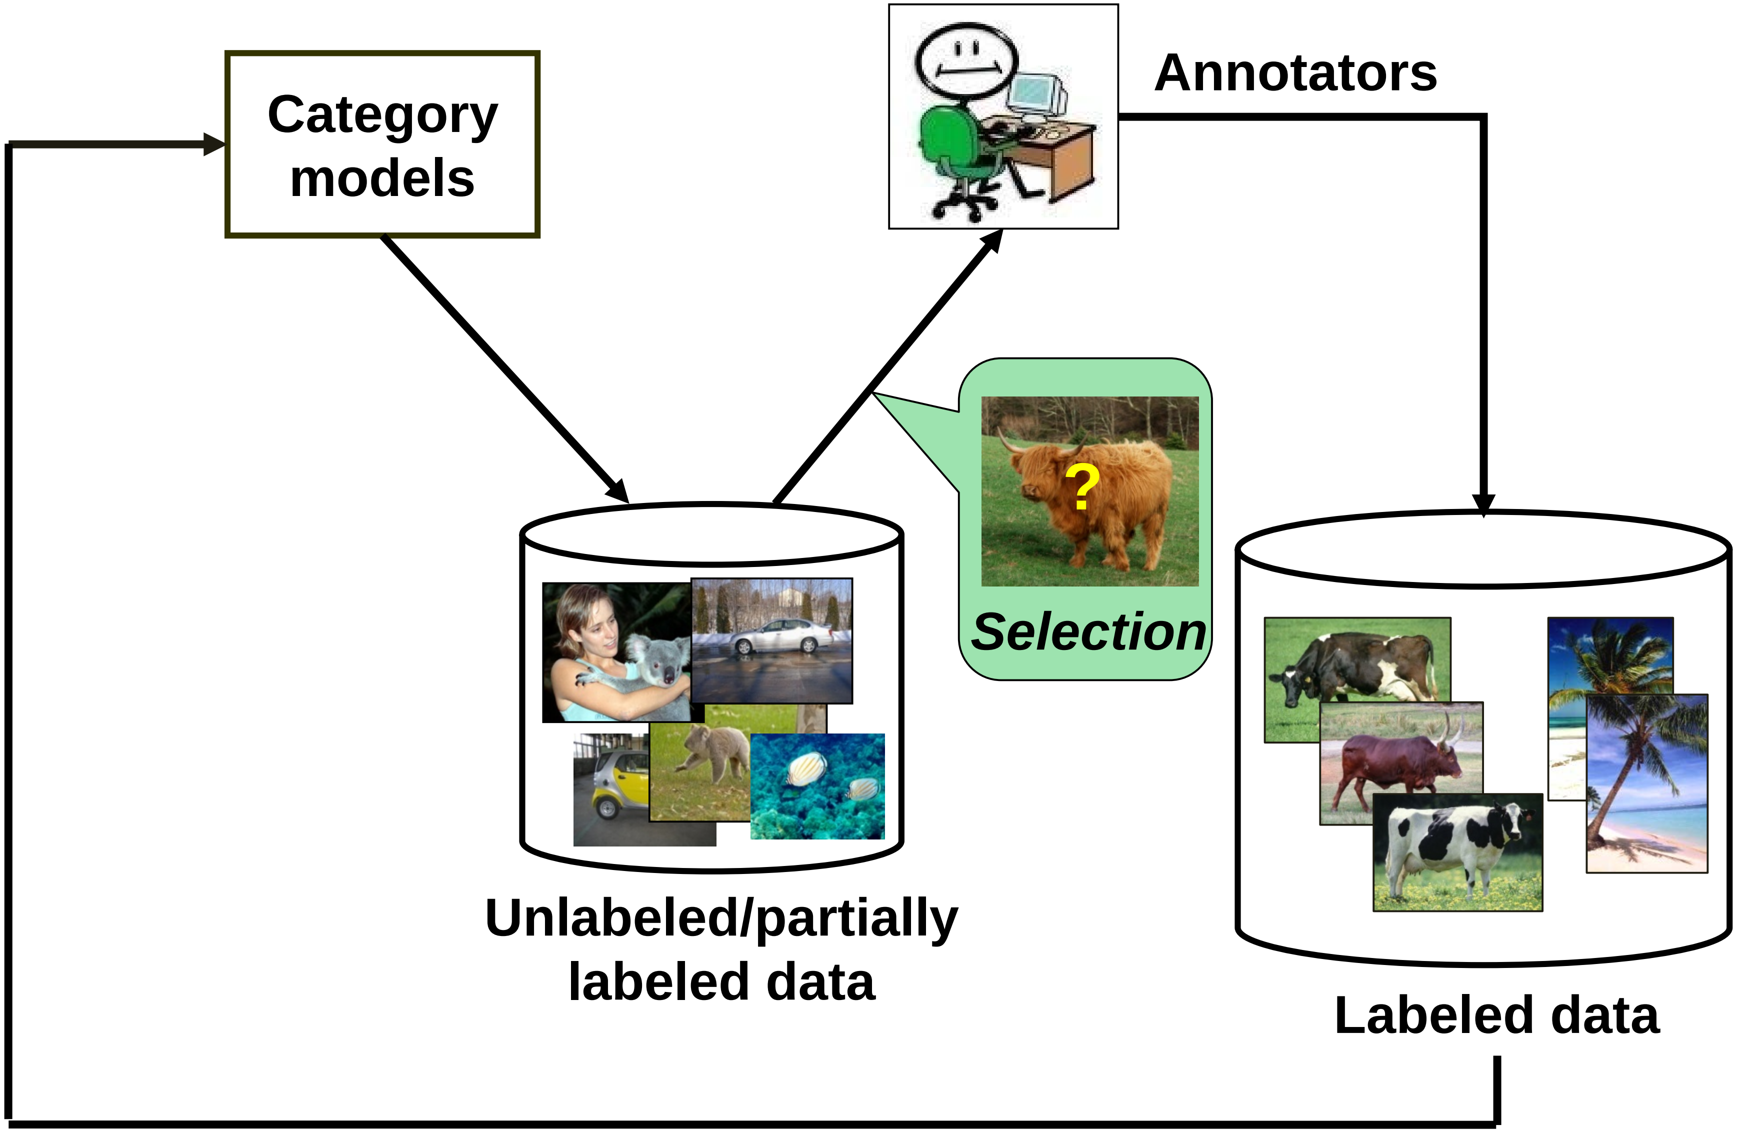

Kristen Grauman

Active selection

- **Traditional active learning** reduces supervision by obtaining labels for the most informative or uncertain examples first.

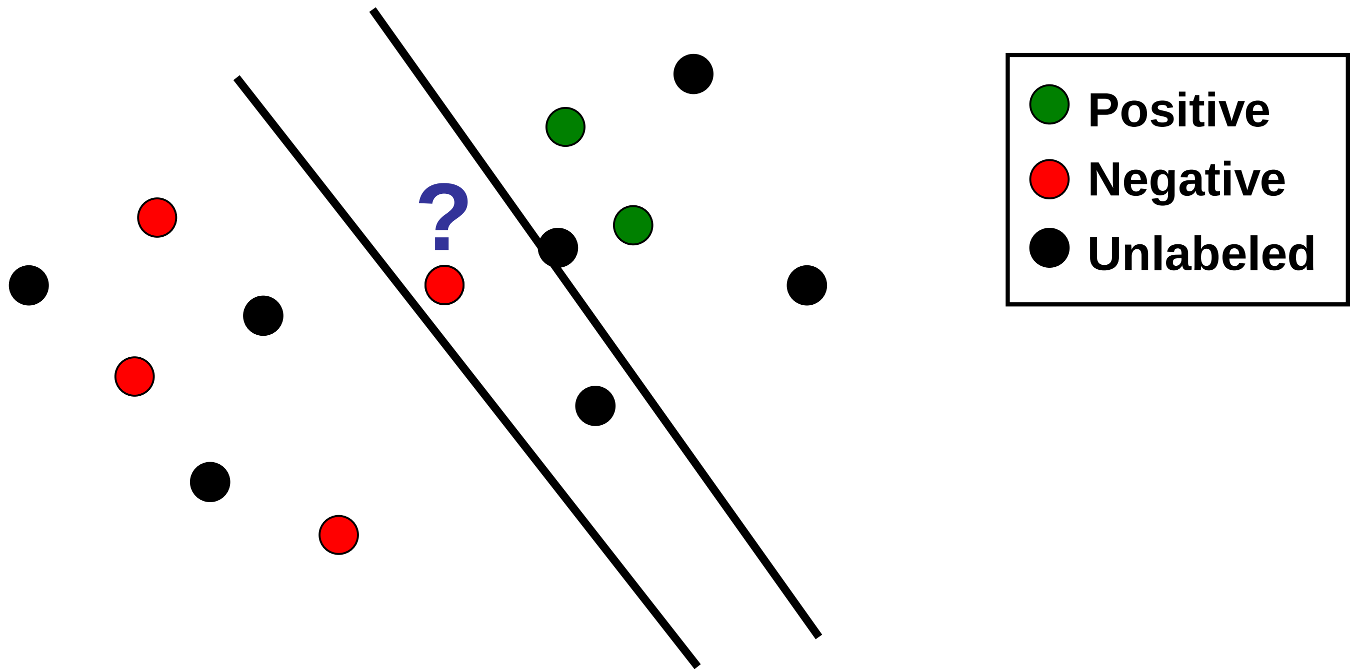

[Mackay 1992, Freund et al. 1997, Tong & Koller 2001, Lindenbaum et al. 2004, Kapoor et al. 2007,…]

Kristen Grauman

Problem: Active selection and recognition

Less expensive to obtain

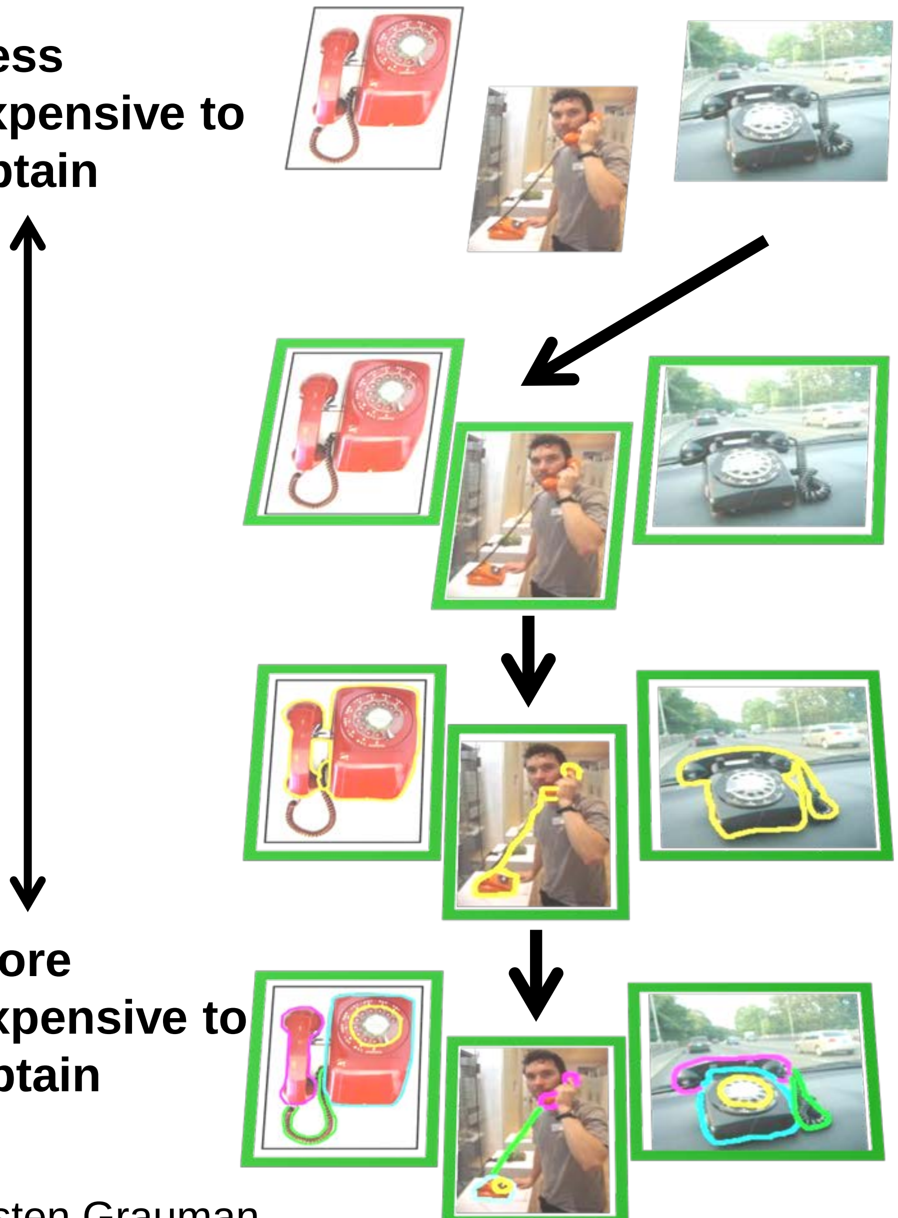

More expensive to obtain

- **Multiple levels** of annotation are possible
- **Variable cost** depending on level *and* example

Kristen Grauman

Idea: Cost-sensitive multi-level active learning

- Compute decision-theoretic active selection criterion that weighs both:
  - which *example* to annotate, and
  - what *kind* of annotation to request for it

  as compared to
  - the *predicted effort* the request would require

[Vijayanarasimhan & Grauman, NIPS 2008, CVPR 2009]

Idea: Cost-sensitive multi-level active learning

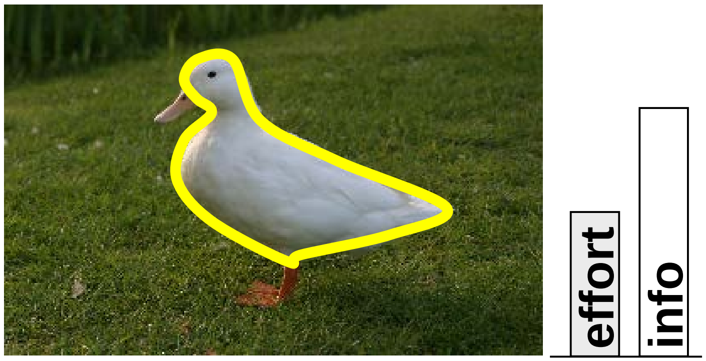

Most regions are understood, but this region is unclear.

This looks expensive to annotate, and it does not seem informative.

…

This looks expensive to annotate, but it seems very informative.

This looks easy to annotate, but its content is already understood.

…

Kristen Grauman

Multi-level active queries

- Predict which query will be most informative, given the cost of obtaining the annotation.
- Three levels (types) to choose from:

1. What object is this region?

2. Does the image contain object X?

3. Segment the image, name all objects.

Kristen Grauman

Decision-theoretic multi-level criterion

$$\underbrace{\text{VALUE}(O,Q)}_{\text{Value of asking given question about given data object}} = \underbrace{\text{RISK}(\mathcal{X}_L, \mathcal{X}_U)}_{\text{Current misclassification risk}} - \underbrace{\widehat{\text{RISK}}(\mathcal{X}_L \cup O_A, \mathcal{X}_U \setminus O)}_{\text{Estimated risk if candidate request were answered}} - \underbrace{\text{COST}(O,Q)}_{\text{Cost of getting the answer}}$$

Estimate risk of incorporating the candidate before obtaining true answer A by computing expected value:

$$\widehat{\text{RISK}}(\mathcal{X}_L \cup O_A, \mathcal{X}_U \setminus O) = \sum_{\ell \in \mathbb{L}} \text{RISK}(\mathcal{X}_L \cup O_\ell, \mathcal{X}_U \setminus O) \;\; p(\ell | O)$$

where $\mathbb{L}$ is set of all possible answers.

$$\text{VPI}(E'|e) = \text{MEU}(e, E') - \text{MEU}(e)$$

$$\text{MEU}(e, E') = \sum_{e'} P(e'|e)\text{MEU}(e, e')$$

Kristen Grauman

Decision-theoretic multi-level criterion

Estimate risk of incorporating the candidate before obtaining true answer A by computing expected value:

$$\widehat{\text{RISK}}(\mathcal{X}_L \cup O_A, \mathcal{X}_U \setminus O) = \sum_{\ell \in \mathbb{L}} \text{RISK}(\mathcal{X}_L \cup O_\ell, \mathcal{X}_U \setminus O) \;\; p(\ell | O)$$

where $\mathbb{L}$ is set of all possible answers.

3.

How many terms are in the expected value?

Kristen Grauman

Decision-theoretic multi-level criterion

Estimate risk of incorporating the candidate before obtaining true answer A by computing expected value:

$$\widehat{\text{RISK}}(\mathcal{X}_L \cup O_A, \mathcal{X}_U \setminus O) = \sum_{\ell \in \mathbb{L}} \text{RISK}(\mathcal{X}_L \cup O_\ell, \mathcal{X}_U \setminus O) \;\; p(\ell | O)$$

where $\mathbb{L}$ is set of all possible answers.

3.

Compute expectation via Gibbs sampling:

- Start with a random setting of the labels.
- For S iterations:
  - o Temporarily fix labels on M-1 regions; train.
  - o Sample remaining region's label.
  - o Cycle that label into the fixed set.

Kristen Grauman

Decision-theoretic multi-level criterion

Estimate risk of incorporating the candidate before obtaining true answer A by computing expected value:

$$\widehat{\text{RISK}}(\mathcal{X}_L \cup O_A, \mathcal{X}_U \setminus O) = \sum_{\ell \in \mathbb{L}} \text{RISK}(\mathcal{X}_L \cup O_\ell, \mathcal{X}_U \setminus O) \;\; p(\ell | O)$$

where $\mathbb{L}$ is set of all possible answers.

3.

For *M* regions $O = \{o_1, \ldots, o_M\}$

$$\approx \frac{1}{S} \sum_{k=1}^{S} \text{RISK}\left(\mathcal{X}_L \cup \{o_1^{(a_1)_k}, \ldots, o_M^{(a_M)_k}\}, \mathcal{X}_U \setminus O\right)$$

Kristen Grauman

Decision-theoretic multi-level criterion

$$\textsc{Value}(O,Q) = \underbrace{\textsc{Risk}(\mathcal{X}_L, \mathcal{X}_U)}_{\text{Current misclassification risk}} - \underbrace{\widehat{\textsc{Risk}}(\mathcal{X}_L \cup O_A, \mathcal{X}_U \setminus O)}_{\text{Estimated risk if candidate request were answered}} - \underbrace{\textsc{Cost}(O,Q)}_{\text{Cost of getting the answer}}$$

Cost of the answer: domain knowledge, or directly predict.

Kristen Grauman

Recap: Actively seeking annotations

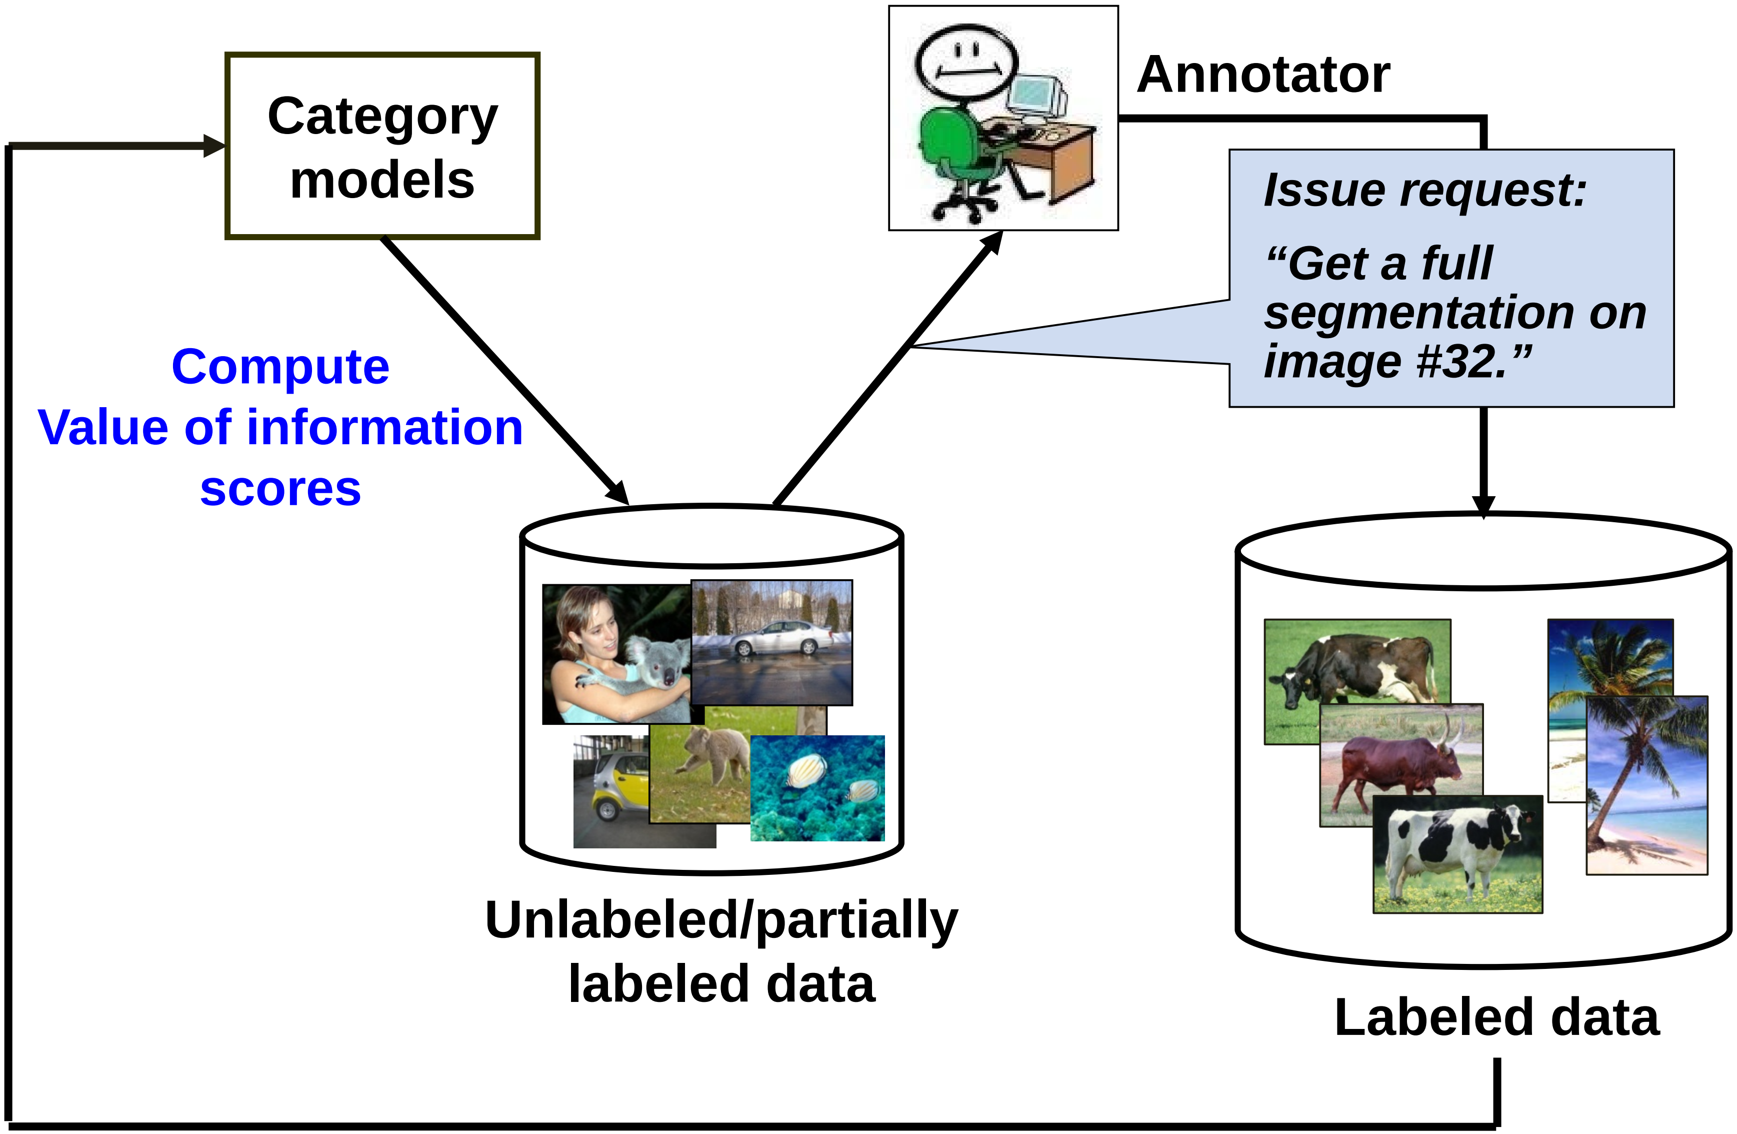

Kristen Grauman

Multi-level active learning curves

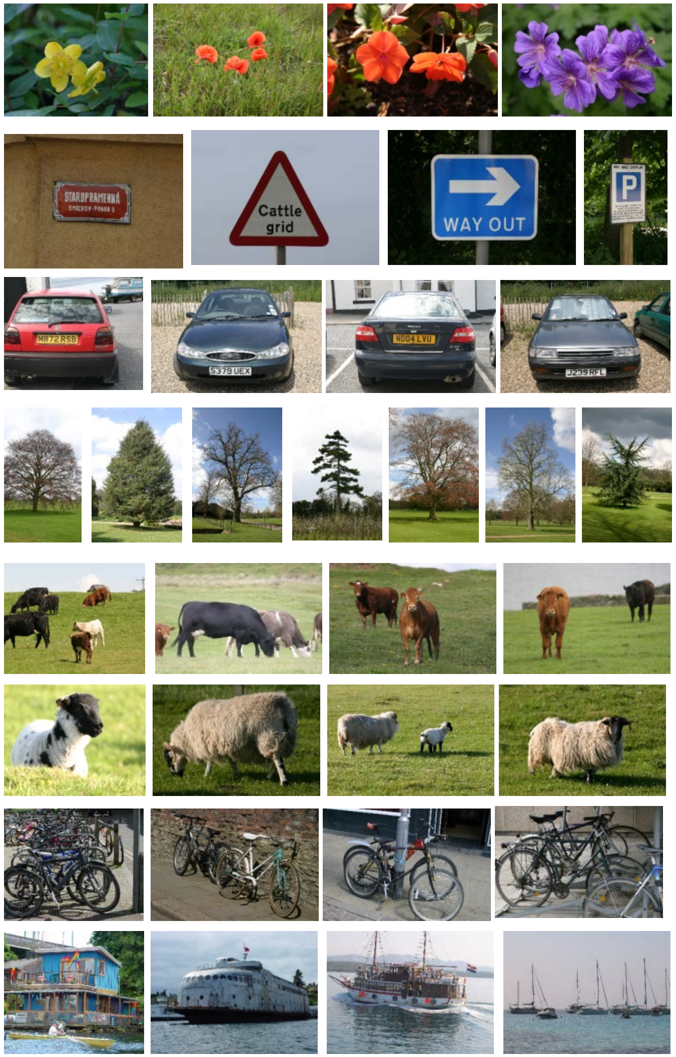

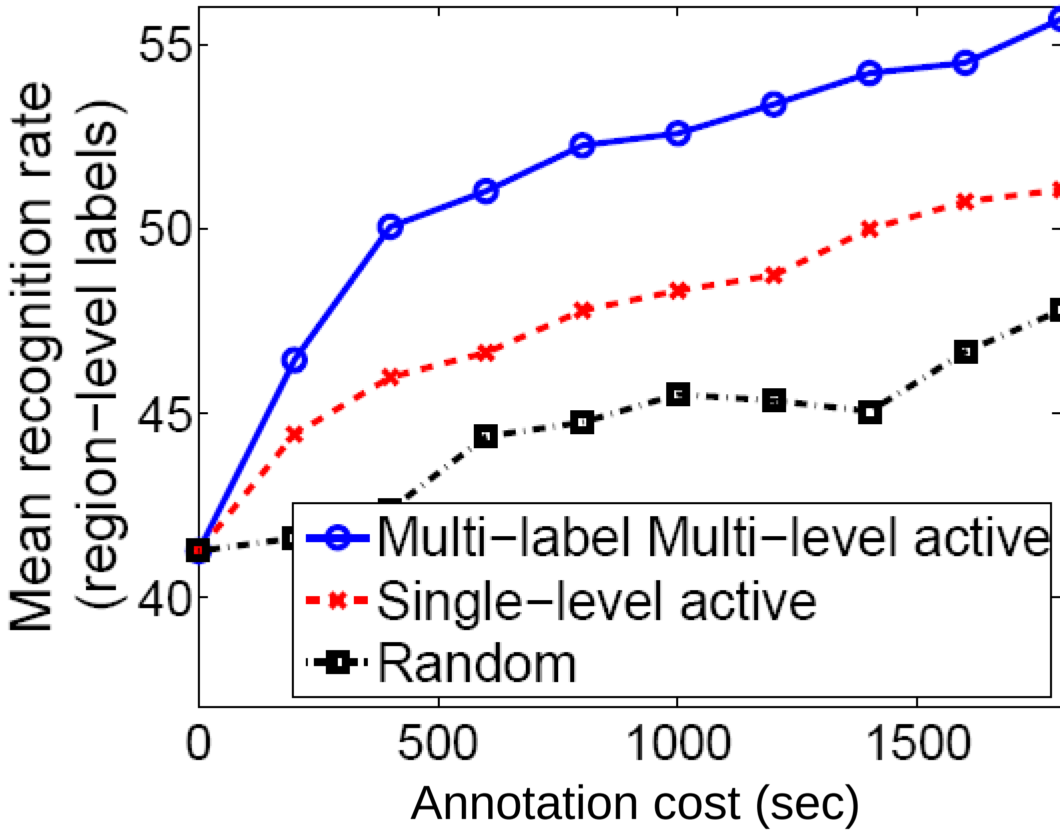

Region features: texture and color

Kristen Grauman

Recap

- Decision networks:
  - What action will maximize expected utility?
  - Connection to expectimax
- Value of information:
  - How much are we willing to pay for a sensing action to gather information?