CS 4649/7649
Robot Intelligence: Planning

Probability Primer

Sungmoon Joo

School of Interactive Computing
College of Computing
Georgia Institute of Technology

*Slides based in part on Dr. Mike Stilman and Dr. Andrew Moore's slides

Administrative– Final Project

- CS7649
- project topic decision, grouping: Due Oct. 23 → update Wiki (group, description)
- project proposal: Due Oct. 30, 2-3page (motivation, technical gap, approach, expected result)
- project final report: Due Dec. 4, 23:59pm, conference-style paper (format is on the course web)
- project presentation: Dec. 11, 11:30am - 2:20pm

*there may be meetings between project teams and the instructor to see if projects are progressing as scheduled.

- CS4649
- project reviewer assignment: Oct. 28
- proposal review report: Due Nov. 6
- project review report(for the assigned project): Due Dec. 11, 11:30am
- project presentation review*(for all presentation): Due Dec. 11, 2:20pm

*presentation review sheets will be provided

Two Sources of Error

Your models are Incorrect!

- **Sensing & State Estimation → Uncertainty**
  - Sensors have noise
  - You don't know exactly what the state is (e.g. mapping, localization,…)
- **Action Execution → Uncertainty**
  - Your actuators do not do what you tell them to
  - The system responds differently than you expect
    : Frictions, etc.

Need to formalize uncertainty!

Random Variable & Probability

(Scalar) Random Variable
 - a variable but rather a function that maps events (or outcome) to real numbers, each with an associated probability

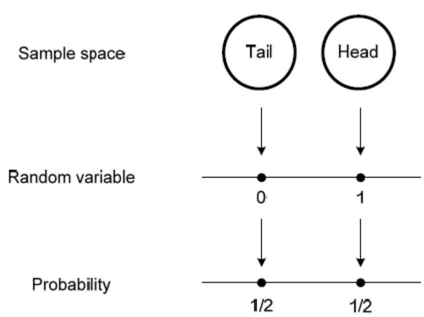

$x : \Omega \rightarrow \mathbb{R}$

(i) $x(Head) = 1$, and $x(Tail) = 0$

(ii) $p(x = 1) = 0.5$, and $p(x = 0) = 0.5$.

* A random vector is a vector-valued random variable whose components are scalar random variables.

Markov Processes: Formalizing Uncertainty

A random process is a collection of continuous-valued random variables indexed by a continuous-valued parameter

$$x(t_i), t_0 \le t_i \le t_f.$$

A Markov Process is a random process which can be completely specified by giving the density functions $p(x(t)|x(\tau))$ and $p(x(\tau))$ for all $t, \tau \in (t_0, t_f)$

We will, mostly, work on discrete version !
- Markov sequence or Markov chain

A random sequence is a collection of random variables (scalar or vector) indexed by a discrete-valued parameter such as $x_0, x_1, \ldots, x_N$

A random sequence is called Markov Sequence if

$$p(x_{k+1}|x_0, \ldots, x_k) = p(x_{k+1}|x_k), \forall k \in 0, 1, \ldots, N$$

Markov Decision Process (MDP)

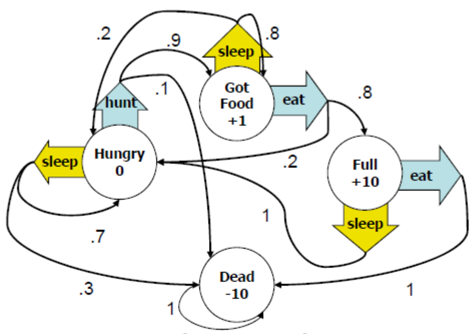

States $\Sigma = \{s_1, \ldots, s_n\}$ **Actions** $A = \{a_1, \ldots, a_m\}$

Rewards $R = \{r_1, \ldots, r_n\}$

Transition Model

$$P_{ij}^l = P(\text{next} = s_j | \text{ current} = s_i \text{ and action} = a_l)$$

* $p_{ij} = P(x_{k+1} = j | x_k = i)$

Value Iteration

The following hold for the discounted problem

(a) The value iteration algorithm converges to the optimal cost J^*

$$J_s^{k+1} = \min_l [r_s + \lambda \sum_{j=1}^{n} p_{s,j}^l J_j^k], 0 < \lambda < 1, s = 1, \ldots, n$$

starting from arbitrary initial conditions

(b) The optimal cost J_s^* of the discounted problem satisfy Bellman's equation, and they are the unique solution of Bellman's equation

$$J_s^* = \min_l [r_s + \lambda \sum_{j=1}^{n} p_{s,j}^l J_j^*], 0 < \lambda < 1, s = 1, \ldots, n$$

(c) $|J_s^k - J_s^*| \leq \alpha \lambda^k, \alpha = \text{a constant}$

"Dynamic Programming and Optimal Control" by D.P. Bertsekas

Where are we headed?

Look familiar?

Creates an interesting looking data structure! (aka Dynamic Programming)

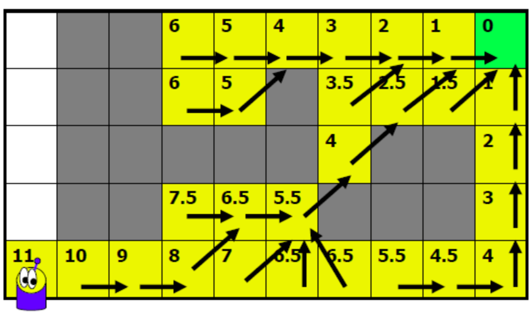

Where are we headed?

'Cost-to-Go' of each state → Policy

Creates an interesting looking data structure! (aka Dynamic Programming)

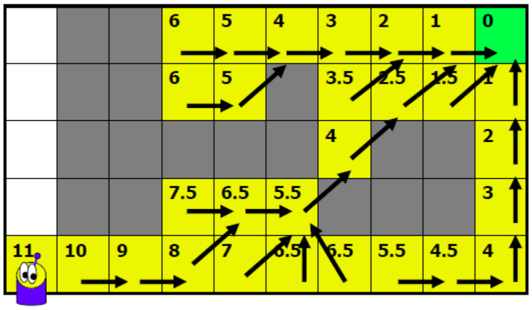

Where are we headed?

'Value' of being at each state

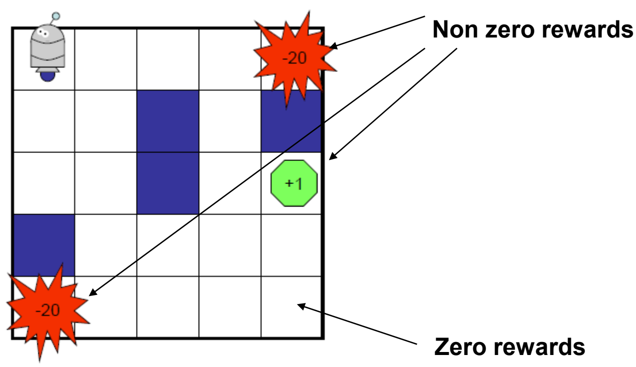

Where are we headed?

'Value' of being at each state → Policy

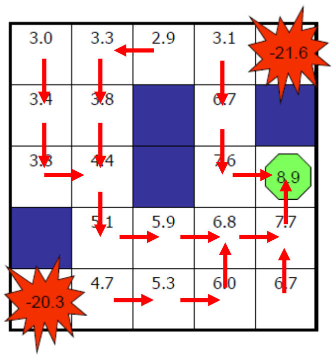

After some iterations

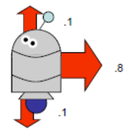

Where are we headed?

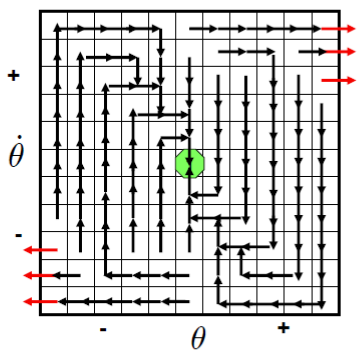

Relevance of Probability

- **Directly related to uncertainty in the world:**

– Things don't always go the way you plan
– Yet, some regularity can be found!

- **Currently, AI has taken on concepts from the 18th century:**

– Probability Theory
– Bayesian Methods

Philosophy

Points of View:

- Probability describes regularities in events
- Probability encodes regularities for which we don't know the cause

Applications:

- Probability distributions describe the dynamics of a process
- Probability distributions are more accurate than guesswork
- An agent's beliefs can be represented as a probability distribution

→ **Bayesian Reasoning**

Random Variables

• A random variable is the result of a chance event, that you can measure or count (continuous or discrete R.V.)

• A is a boolean-valued random variable if:

– A denotes an event

– There is some degree of certainty as to whether A occurs

Examples of events
- When you toss a coin, the outcome is heads
- You will learn something about probability from this lecture
- You wake up tomorrow with a headache

Probability

P(**A**) = the fraction of possible worlds in which **A** is **true**

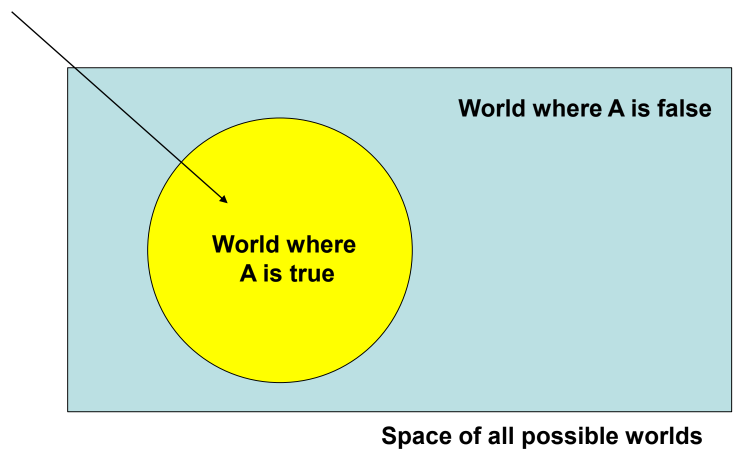

The Axioms of Probability

- $0 \leq P(A) \leq 1$
- P(True) = 1
- P(False) = 0
- P(A or B) = P(A) + P(B) – P(A and B)

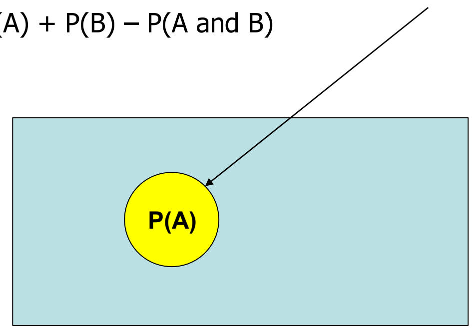

The Axioms of Probability

- $0 \leq P(A) \leq 1$
- P(True) = 1
- P(False) = 0
- P(A or B) = P(A) + P(B) – P(A and B)

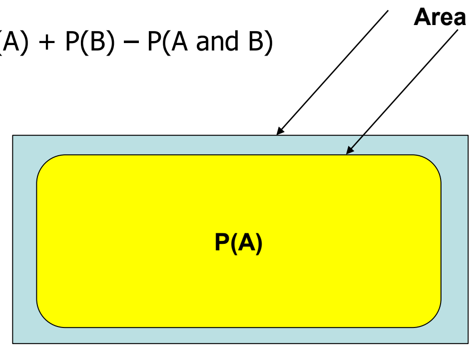

The Axioms of Probability

- $0 \le P(A) \le 1$
- P(True) = 1
- P(False) = 0
- P(A or B) = P(A) + P(B) – P(A and B)

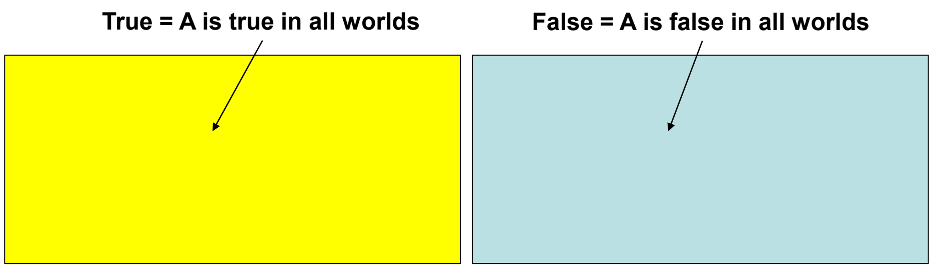

The Axioms of Probability

- $0 \le P(A) \le 1$
- P(True) = 1
- P(False) = 0
- P(A or B) = P(A) + P(B) – P(A and B)

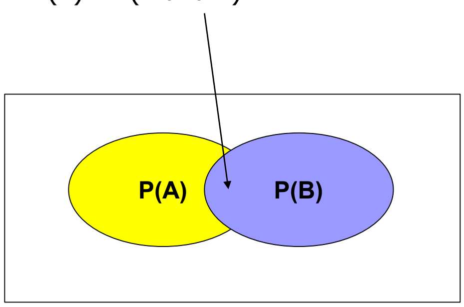

Axioms of Probability vs. Others

- **There have been attempts to handle uncertainty:**

– Fuzzy Logic
– Three-valued Logic (True/False/Unknown)
– …

- **But, axioms of probability are the ONLY system such that:**

– "If you gamble using them you can't be unfairly exploited by an opponent using some other system."

- de Finetti 1931

Probability Theorems: 1

(1) $0 \leq P(A) \leq 1$ (2) P(True) = 1 (3) P(False) = 0
(4) P(A or B) = P(A) + P(B) – P(A and B)

Axioms

Prove: P (not A) = P (¬ A) = 1 – P(A)
Intuition:

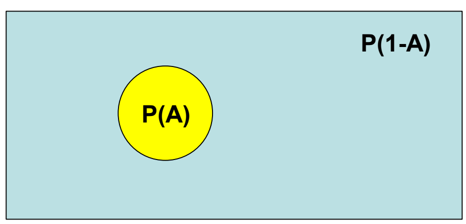

Probability Theorems: 1

(1) $0 \leq P(A) \leq 1$ (2) P(True) = 1 (3) P(False) = 0
(4) P(A or B) = P(A) + P(B) – P(A and B)

Axioms

Prove: **P (not A) = P (¬ A) = 1 – P(A)**
Proof:

P(A or not A) = P(True) = 1 ⟵ Ax. 2
P(A and not A) = P(False) = 0 ⟵ Ax. 3
1 = P(A) + P(not A) – 0 ⟵ **Ax. 4**

P(not A) = 1 – P(A)

Probability Theorems: 2

(1) $0 \leq P(A) \leq 1$ (2) P(True) = 1 (3) P(False) = 0
(4) P(A or B) = P(A) + P(B) – P(A and B)

Axioms

Prove: **P (A)** = P (A and B) + P(A and not B)
= P (A ∧ B) + P(A ∧ ¬B)

Proof:

P (A) = P (A ∧ True) ⟵ Truth Table
= P (A ∧ (B ∨ ¬ B)) ⟵ Truth Table
= P ((A ∧ B) ∨ (A ∧ ¬ B)) ⟵ Distribute
= P(A ∧ B) + P(A ∧ ¬B) ⟵ **Ax. 4**
- P((A ∧ B) ∧ (A ∧ ¬ B))
= P(A ∧ B) + P(A ∧ ¬ B) ⟵ **Ax. 3**

Multi-Valued Random Variables

A is a **Random Variable with arity k** if it can take on exactly one value from the set:

$\{v_1, v_2, \ldots, v_k\}$

$$P(A = v_i \wedge A = v_j) = 0 \text{ if } i \neq j$$

$$P(A = v_1 \vee A = v_2 \vee \ldots \vee A = v_k) = 1$$

Probability Theorems: 3

(1) $0 \leq P(A) \leq 1$ (2) $P(True) = 1$ (3) $P(False) = 0$
(4) $P(A \text{ or } B) = P(A) + P(B) - P(A \text{ and } B)$
(5) $P(A = v_i \wedge A = v_j) = 0$ if $i \neq j$
(6) $P(A = v_1 \vee A = v_2 \vee \ldots \vee A = v_k) = 1$

Axioms

Prove: $P(A = v_1 \vee A = v_2 \vee \ldots \vee A = v_i) = \sum_{j=1}^{i} P(A = v_j)$

Expand with Ax. 4, remove P(A ∧ B) with Ax. 3, 5

Also: $\sum_{j=1}^{k} P(A = v_j) = 1$ **Ax. 6**

Probability Theorems: 4

(1) $0 \leq P(A) \leq 1$ (2) $P(True) = 1$ (3) $P(False) = 0$
(4) $P(A \text{ or } B) = P(A) + P(B) - P(A \text{ and } B)$
(5) $P(A = v_i \wedge A = v_j) = 0$ if $i \neq j$
(6) $P(A = v_1 \vee A = v_2 \vee \ldots \vee A = v_k) = 1$

Axioms

Prove: $P(B \wedge (A = v_1 \vee A = v_2 \vee \ldots \vee A = v_i)) = \sum_{j=1}^{i} P(B \wedge A = v_j)$

Distribute (and) over (or) and follow Theorem 3

Leads to: $P(B) = \sum_{j=1}^{k} P(B \wedge A = v_j)$

Interim Summary

Axioms

(1) $0 \leq P(A) \leq 1$ (2) $P(True) = 1$ (3) $P(False) = 0$

(4) $P(A \text{ or } B) = P(A) + P(B) - P(A \text{ and } B)$

(5) $P(A = v_i \wedge A = v_j) = 0$ if $i \neq j$

(6) $P(A = v_1 \vee A = v_2 \vee \ldots \vee A = v_k) = 1$

Theorems

$P(\text{not } A) = P(\neg A) = 1 - P(A)$

$P(A) = P(A \wedge B) + P(A \wedge \neg B)$

$P(A = v_1 \vee A = v_2 \vee \ldots \vee A = v_i) = \sum_{j=1}^{i} P(A = v_j)$

$P(B \wedge (A = v_1 \vee A = v_2 \vee \ldots \vee A = v_i)) = \sum_{j=1}^{i} P(B \wedge A = v_j)$

Visualizing Probability

P(¬A) + P(A) = 1

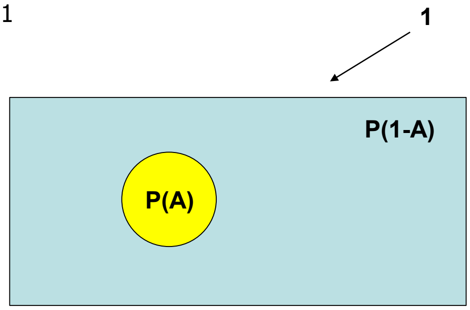

S. Joo (sungmoon.joo@cc.gatech.edu) 10/23/2014 29

Visualizing Probability

P(B) = P(B ∧ A) + P(B ∧ ¬ A)

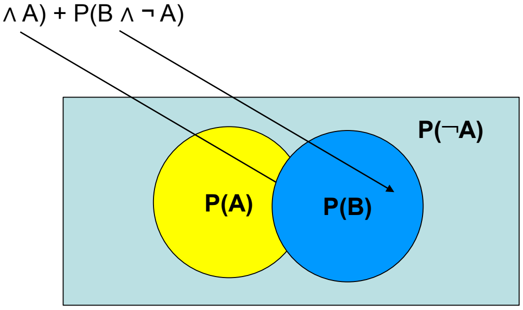

S. Joo (sungmoon.joo@cc.gatech.edu) 10/23/2014 30

Visualizing Probability

$$\sum_{j=1}^{k} P(A = v_j) = 1$$

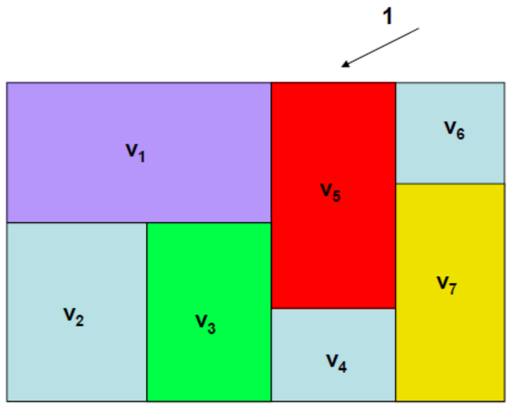

Conditional Probability

P (A | B) = The fraction of possible worlds in which B = true that also have A = true

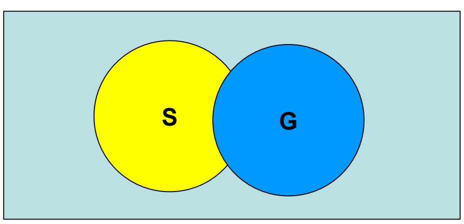

P(S | G) = Fraction of goal reaching worlds in which your sensor works properly.

$$= \frac{\text{\# worlds with goal and sensor}}{\text{\# worlds with goal}}$$

$$= \frac{\text{area of (S and G) region}}{\text{area of G region}}$$

$$\mathbf{= \frac{P(S \wedge G)}{P(G)}}$$

G = Robot reaches the goal
S = Sensor works properly

Conditional Probability

- Conditional Probability (Definition)

$$P(A|B) = \frac{P(A \wedge B)}{P(B)}$$

- Corollary

$$P(A \wedge B) = P(A|B)P(B)$$

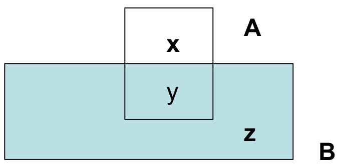

$$P(B|A) = \frac{y}{x+y}$$

Probabilistic Inference

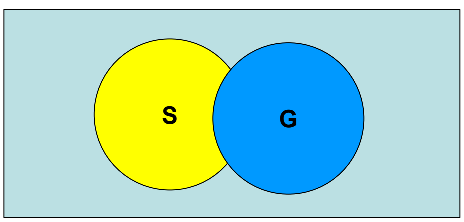

G = Robot reaches the goal
S = Sensor works properly

P(G) = 1/20
P(S) = 1/10
P(S | G) = 3/4

P(S ∧ G) = P(S | G) P(G) = 3/4 × 1/20 = 3/80

P(G | S) = P(S ∧ G) / P (S) = (3/80) / (1/10) = **3/8**

Bayes Rule

$$P(A|B) = \frac{P(A \wedge B)}{P(B)} = \frac{P(B|A)P(A)}{P(B)}$$

$$P(B|A) = \frac{P(B \wedge A)}{P(A)}$$

Alternative Formulations:

$$P(A|B) = \frac{P(B|A)P(A)}{P(B|A)P(A) + P(B|\neg A)P(\neg A)}$$

$$P(A|B \wedge X) = \frac{P(B|A \wedge X)P(A \wedge X)}{P(B \wedge X)}$$

$$P(A = v_i|B) = \frac{P(B|A = v_i)P(A = v_i)}{\sum_{j=1}^{k} P(B|A = v_i)P(A = v_j)}$$

Thomas Bayes (1763)

Probability Distributions

Boolean

A	P(A=v_i)
T	.3
F	.7
Total	1

Multi-Valued

A	P(A=v_i)
v_1	.3
V_2	.2
V_3	.15
V_4	.15
V_5	.2
Total	1

Continuous

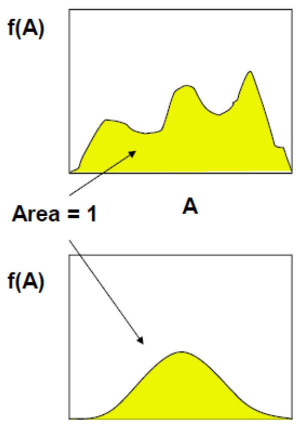

$$E(A) = \sum_i v_i P(v_i) \equiv \int_{-\infty}^{\infty} v f(v) dv$$

$$Var(A) = E(A^2) - (E(A))^2$$

S.D. = Sqrt (Var)

How do we answer questions?

Suppose we have a world with random variable A, B, and C

- Q1: What is the probability that A, B are true but C is false?
- Q2: Given that B is true, what is the probability of A being true?
- Q3: Given that B and C are true, what is the probability of A being true?

Joint Distributions

Analogous to a truth table:

List all combinations of values for each of the random variables.

Assign a probability to each of the combinations (rows).

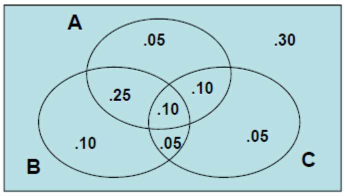

A	B	C	Prob
0	0	0	.30
0	0	1	.05
0	1	0	.10
0	1	1	.05
1	0	0	.05
1	0	1	.10
1	1	0	.25
1	1	1	.10

Joint Distributions

- Q1: What is the probability that A, B are true but C is false?

A	B	C	Prob
0	0	0	.30
0	0	1	.05
0	1	0	.10
0	1	1	.05
1	0	0	.05
1	0	1	.10
1	1	0	.25
1	1	1	.10

$$P(A \wedge B \wedge \neg C) = P(A \wedge B) - P(A \wedge B \wedge C)$$
$$= 0.25 + 0.1 - 0.1$$
$$= 0.25$$

$$P(A) = P(A \wedge B) + P(A \wedge \neg B)$$

$$P(E) = \sum_{rows\ matching\ E} P(row)$$

$$P(E_1|E_2) = \sum_{rows\ matching\ E_1\ and E_2} P(row)$$

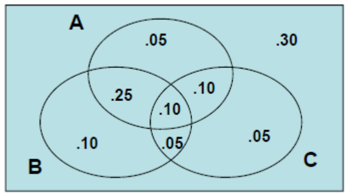

Joint Distributions

- Q2: Given that B is true, what is the probability of A being true?

Using a joint distribution:

A	B	C	Prob
0	0	0	.30
0	0	1	.05
0	1	0	.10
0	1	1	.05
1	0	0	.05
1	0	1	.10
1	1	0	.25
1	1	1	.10

$$P(E_1|E_2) =$$
$$= \frac{P(E_1 \wedge E_2)}{P(E_2)}$$
$$= \frac{\sum_{rowsmatchingE_1 andE_2} P(row)}{\sum_{rowsmatchingE_2} P(row)}$$

$$P(A|B) = \frac{P(A \wedge B)}{P(B)}$$
$$= \frac{0.25 + 0.1}{0.1 + 0.05 + 0.25 + 0.1}$$
$$= \frac{0.35}{0.5}$$

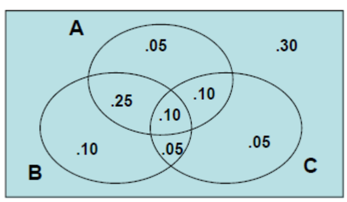

Joint Distributions

- Q3: Given that B and C are true, what is the probability of A being true?

Using a joint distribution:

$$P(E_1|E_2) =$$
$$= \frac{P(E_1 \wedge E_2)}{P(E_2)}$$
$$= \frac{\sum_{rowsmatchingE_1andE_2} P(row)}{\sum_{rowsmatchingE_2} P(row)}$$

A	B	C	Prob
0	0	0	.30
0	0	1	.05
0	1	0	.10
0	1	1	.05
1	0	0	.05
1	0	1	.10
1	1	0	.25
1	1	1	.10

$$P(A|B \wedge C) = \frac{P(A \wedge B \wedge C)}{P(B \wedge C)}$$
$$= \frac{0.1}{0.05 + 0.1}$$
$$= \frac{0.1}{0.15}$$