Data Reliability

Interpreting Sensor Data

Reminders:

- HW1 due now.
- HW2 out this evening.

Analogy

- System reliability challenge:
  - Building reliable systems from less reliable components
- Data reliability challenge:
  - Making reliable conclusions from less reliable (sensor) data

Making Conclusions from Probabilistic Data

- Cyber-physical systems obtain data about their environment via sensors
- Sensors (or data sources in general) are often imperfect
- The challenge is: how to correctly compute the "state of the environment" given sensor readings?

Review: Things You Should Know About Probabilities

- Probability of multiple simultaneous events
  - What are the odds that it rains and my basement floods? Say P(rains) = 0.2. P(flood) = 0.1

Review: Things You Should Know About Probabilities

- Probability of multiple simultaneous events
  - What are the odds that it rains and my basement floods?
  - Answer: It is the odds that "it rains", times the odds that "my basement floods given that it rains":

$$P(\text{rain}, \text{flood}) = P(\text{rain})\, P(\text{flood}|\text{rain})$$

Note: $P(\text{flood}|\text{rain})$ is larger than $P(\text{flood})$

Review: Things You Should Know About Probabilities

- $P(A,B) = P(A|B).P(B)$

- Corollary: If events A and B are independent, the odds of them happening together is the product of their individual probabilities.
- $P(A,B) = P(A).P(B)$

Note: This is because $P(A|B) = P(A)$

Review: Probability versus Conditional Probability

- Probability: the odds that something, say X, happens, P(X).
- Conditional probability: the odds that X happens *given* that a certain condition, C, has occurred, P(X|C).
  - This condition may affect the odds.

Review: Probability versus Conditional Probability

- Probability: the odds that something, say X, happens, P(X).
- Conditional probability: the odds that X happens *given* that a certain condition, C, has occurred, P(X|C).
  - This condition may affect the odds.
- Traffic example:
  - P(Accident) may be low
  - P(Accident|Black ice) is a lot higher!

A Visual Interpretation

The universe of all possibilities

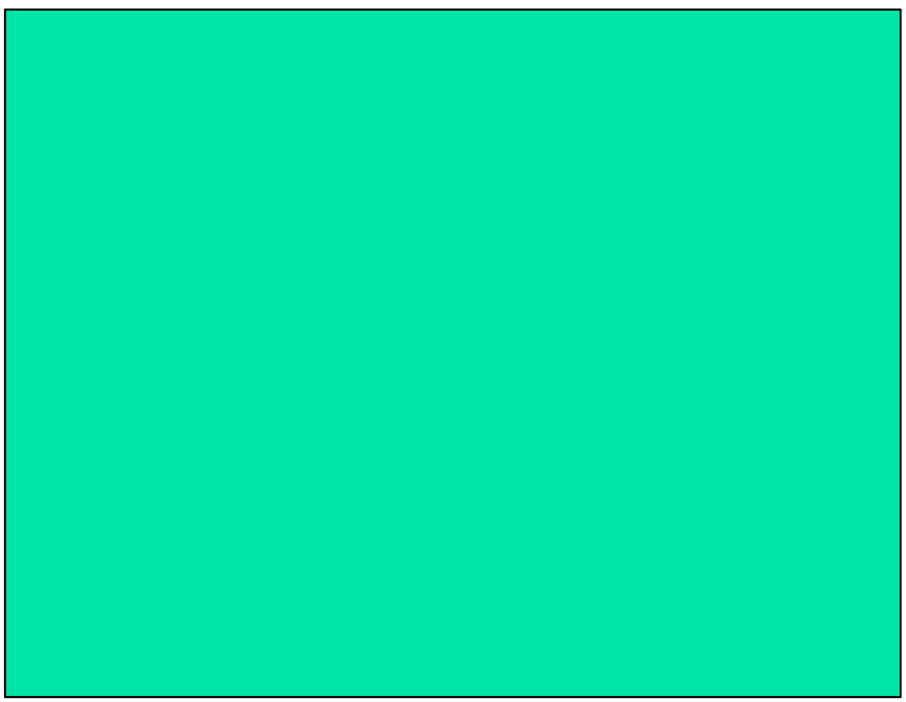

A Visual Interpretation

The universe of all possibilities

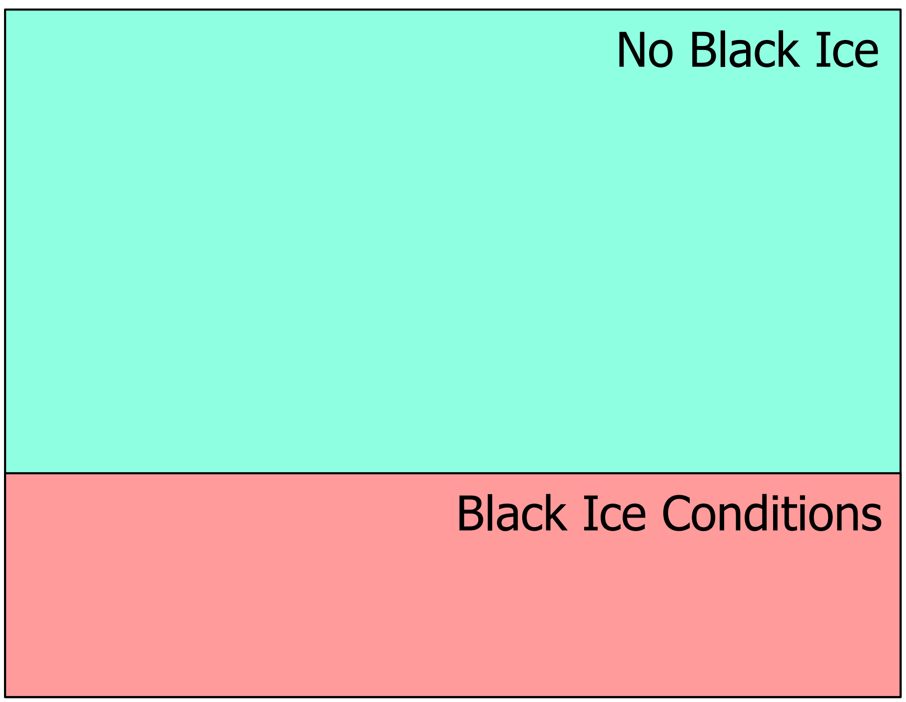

Consider a graphical visualization of probabilities where area represents probability

A Visual Interpretation

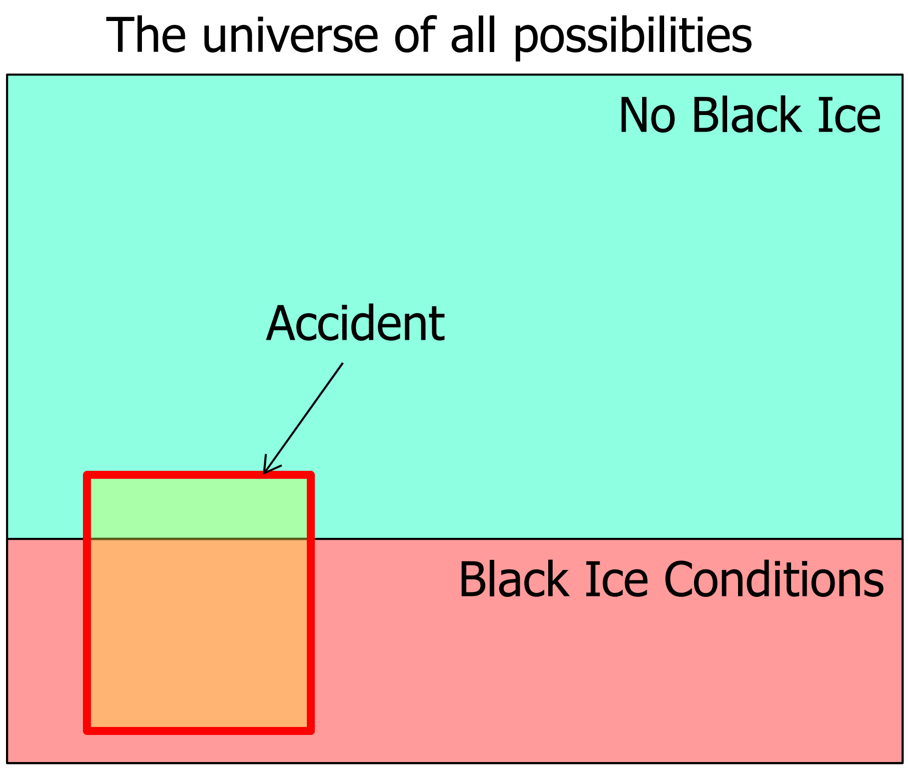

Consider a graphical visualization of probabilities where area represents probability

A Visual Interpretation

The universe of all possibilities

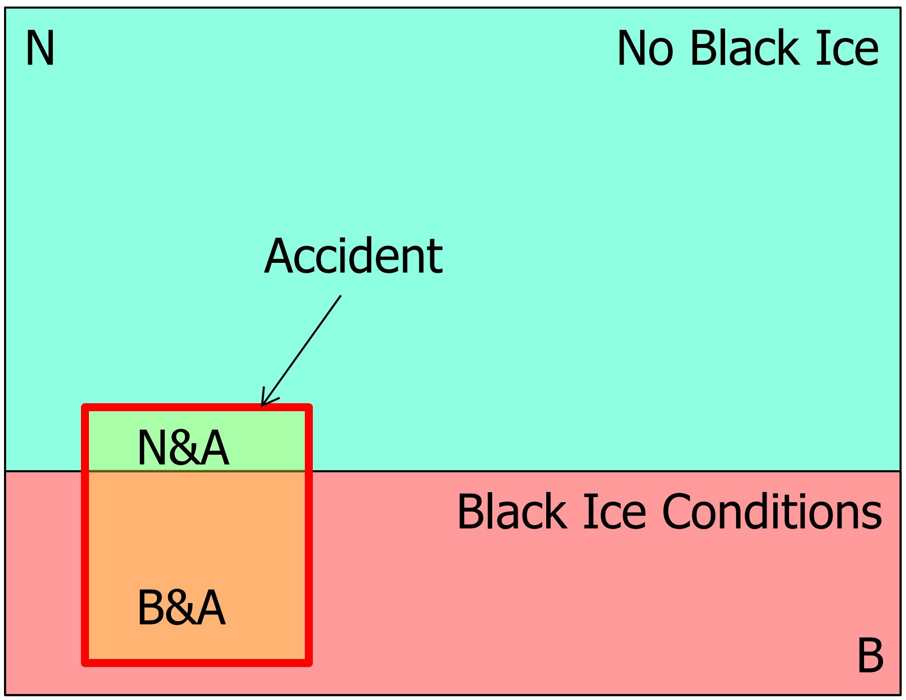

Consider a graphical visualization of probabilities where area represents probability

A Visual Interpretation

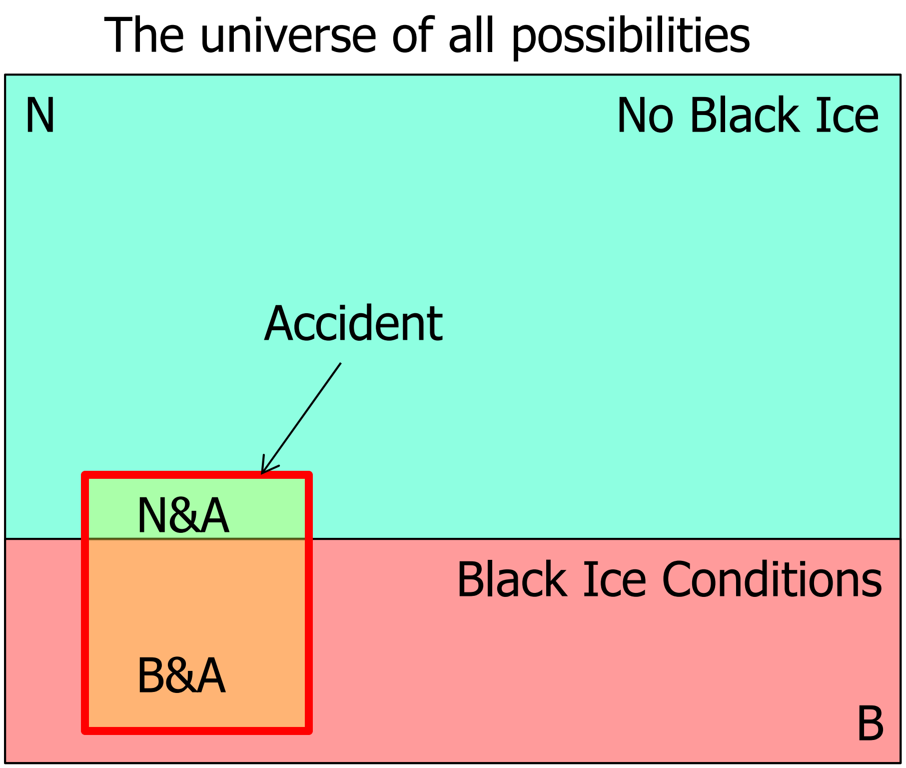

- P(Accident) = (N&A+B&A)/(N+B)

A Visual Interpretation

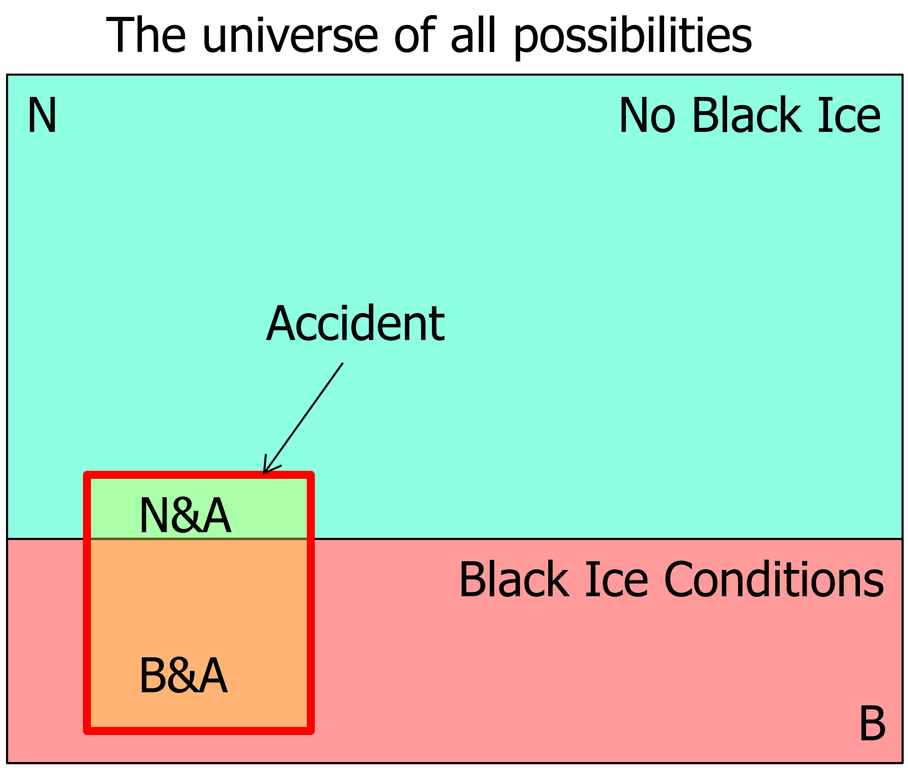

- P(Accident) = (N&A+B&A)/(N+B)
- P(Accident|Ice) = (B&A)/B
- P(Accident|No Ice) = (N&A)/N

A Visual Interpretation

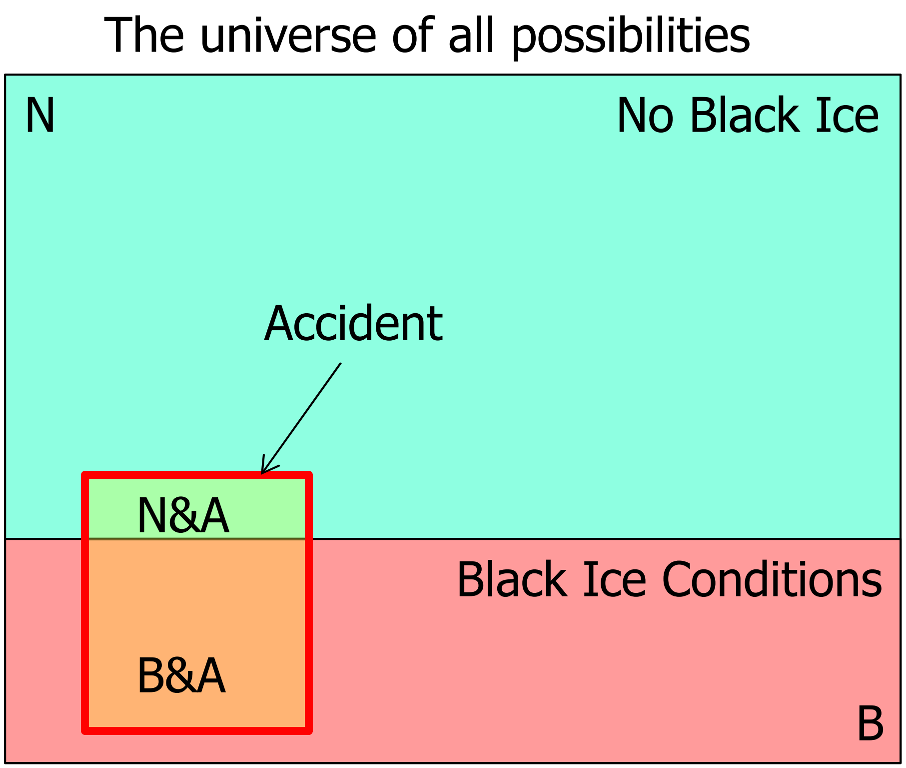

- P(Accident) = (N&A+B&A)/(N+B)
- P(Accident|Ice) = (B&A)/B
- P(Accident|No Ice) = (N&A)/N

Question: If P(Accident|Ice) = 0.1, P(Accident|No-Ice) = 0.02, P(Ice) = 0.2, what are the odds of accidents in general, P(Accident)?

A Visual Interpretation

The universe of all possibilities

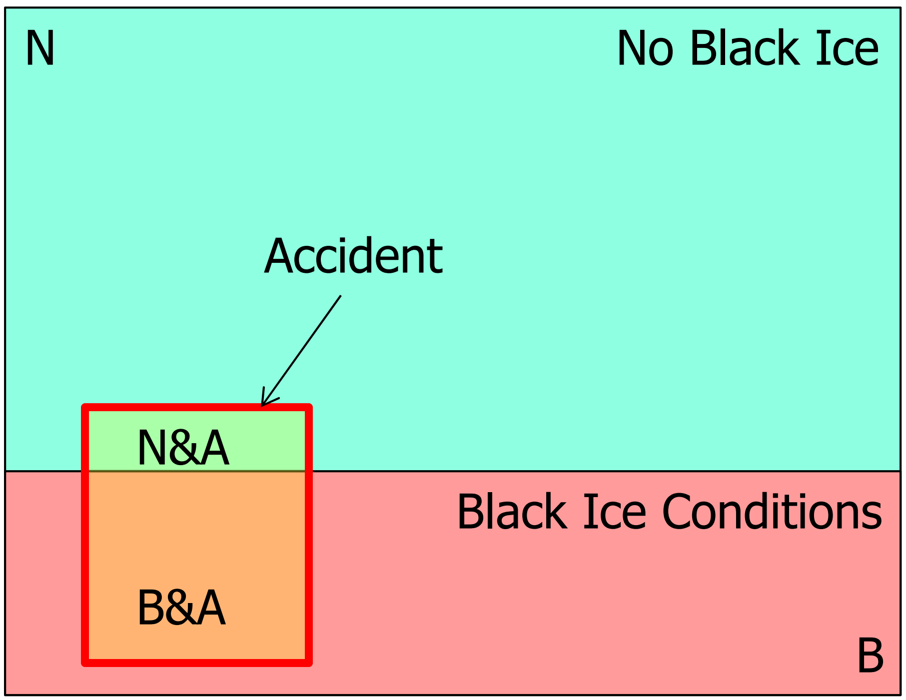

- (B&A)/B = 0.1
- (N&A)/N = 0.02
- B/(B+N) = 0.2
- (N&A+B&A)/(N+B)?

Question: If P(Accident|Ice) = 0.1, P(Accident|No-Ice) = 0.02, P(Ice) = 0.2, what are the odds of accidents in general, P(Accident)?

Review:
Total Probability Theorem

$P(\text{Something}) = P(\text{Something}|X)\,P(X) + P(\text{Something}|\overline{X})\,P(\overline{X})$

The universe of all possibilities

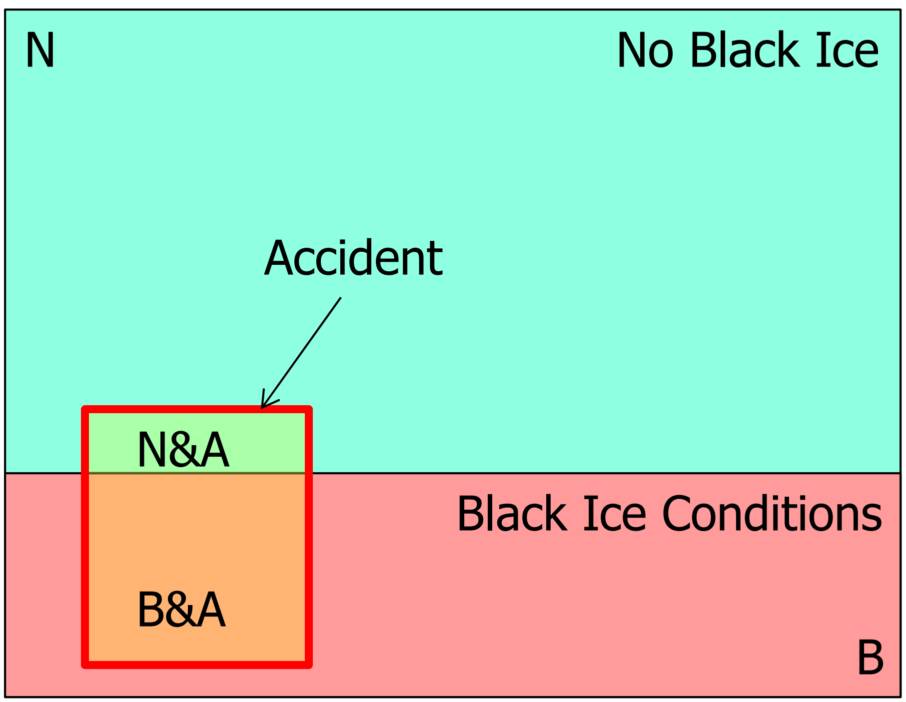

- P(Accident) =
  0.1 * 0.2
  + 0.02 * 0.8
  = 0.032

Question: If P(Accident|Ice) = 0.1, P(Accident|No-Ice) = 0.02, P(Ice) = 0.2, what are the odds of accidents in general, P(Accident)?

Example: Probability versus Conditional Probability

- A father is accused of murdering his abused child. The lawyer says that only 2% of men who abuse their children actually end up killing them, so the odds that this is a murder are very low.

- Is this argument statistically valid? If so, explain why (mathematically). If not, why not?

Example: Probability versus Conditional Probability

- A father is accused of murdering his abused child. The lawyer says that only 2% of men who abuse their children actually end up killing them, so the odds that this is a murder are very low.

- The relevant statistic is: *given that* an abused child is murdered, what are the odds that the father did it? (This happens to be 50%)

Example: Probability versus Conditional Probability

- A father is accused of murdering his abused child. The lawyer says that only 2% of men who abuse their children actually end up killing them, so the odds that this is a murder are very low.

Conditional probability

- The relevant statistic is: *given that* an abused child is murdered, what are the odds that the father did it? (This happens to be 50%)

A Visual Interpretation

The universe of all possibilities

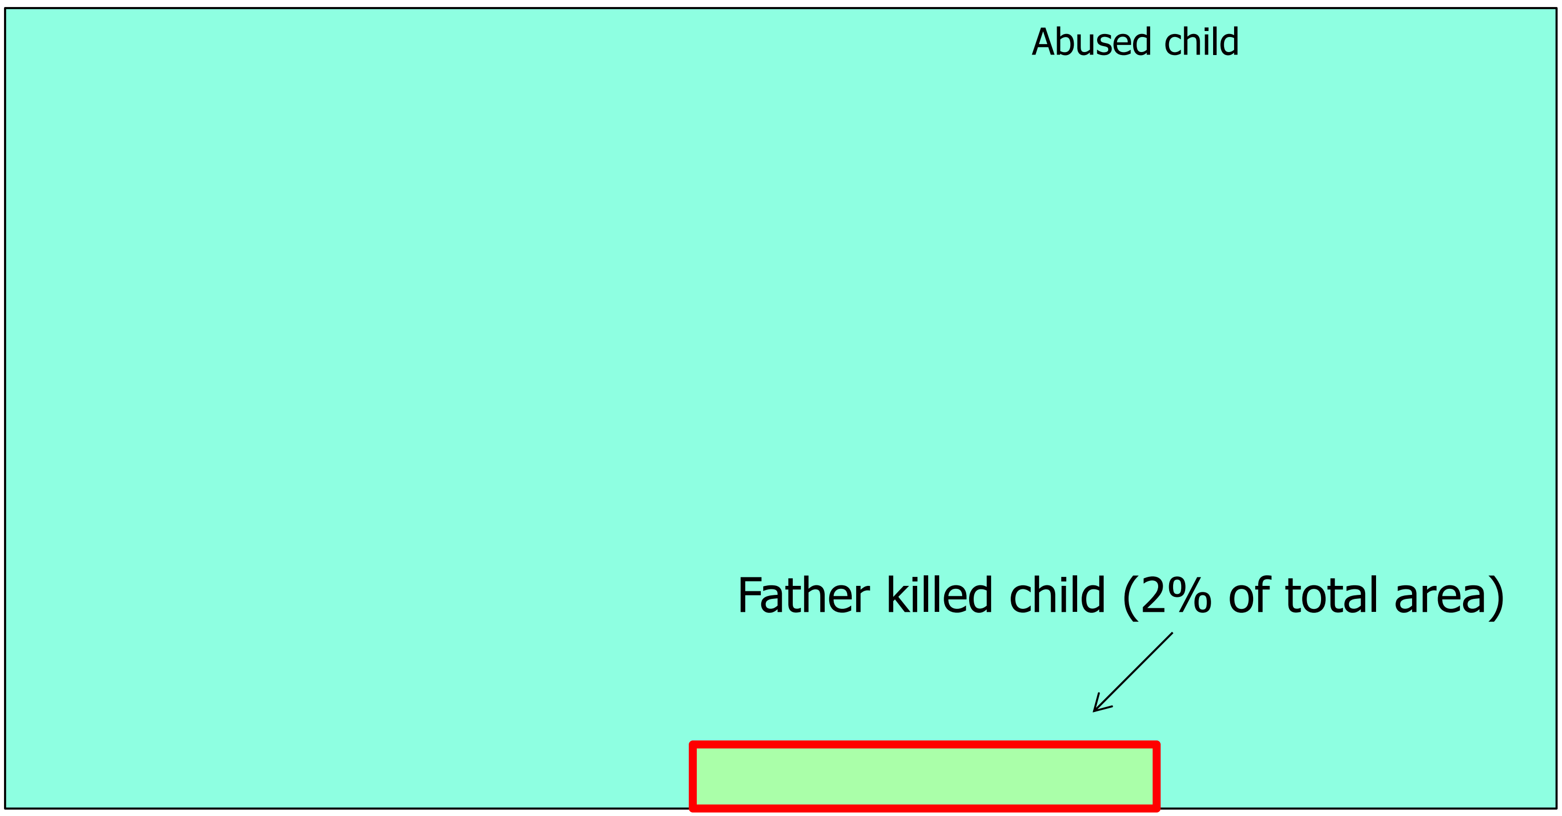

A Visual Interpretation

The universe of all possibilities

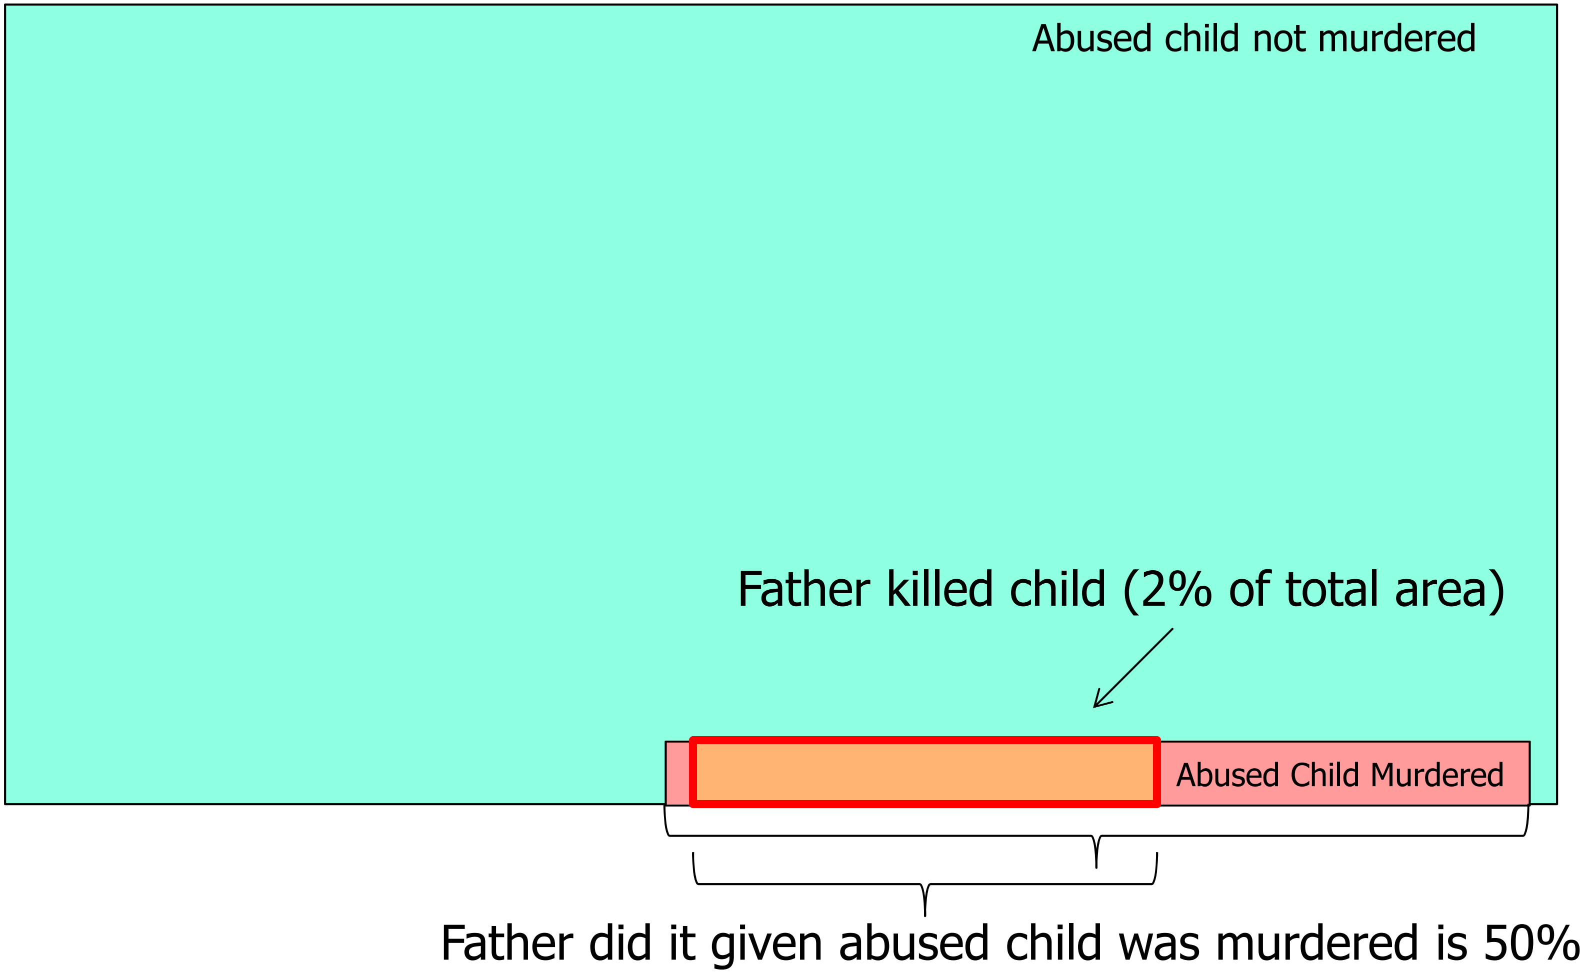

Example: Intrusion Detection

- A motion alarm is used to detect unauthorized access to a warehouse after hours. The motion sensor is mounted near the only entrance to the warehouse. If a burglar enters the building, there is a 99% chance that the burglar triggers the motion alarm.
- At 9pm, on September 16th, 2013, the alarm was set off. What are the odds that a burglar is in the building?

Example: Asteroid Collision with Earth

- If a major Asteroid collides with Earth in St. Louis, traffic on I-74 will be backed up.

- On Feb 10th, 2014, there was a big back-up on I-74. What are the odds that a major Asteroid collided with Earth?

Example: Asteroid Collision with Earth

- If a major Asteroid collides with Earth in St. Louis (A), traffic on I-74 will be backed up (B).
  - P (B|A) = P (Backup given Asteroid) = 1

- On Feb 10th, 2014, there was a big back-up on I-74 (B). What are the odds that a major Asteroid collided with Earth in St. Louis (A)?
  - P(A|B) = P (Asteroid given Backup) = ?

Factors to Consider

P (Asteroid given Backup) ?

Factors to Consider
P (Asteroid given Backup) ?

- How often do major asteroids hit earth?
  - P(A) = P (Asteroid) = ?
  - The less often it happens, the less likely it is that a traffic jam is attributed to an asteroid.
- How often traffic backs up on I-74
  - P(B) = P (Backup) = ?
  - The more often this happens the less likely it is to be an indicator of asteroid collision

Factors to Consider
P (Asteroid given Backup) ?

- How often do major asteroids hit earth?
  - P(A) = P (Asteroid) = ?
  - The less often it happens, the less likely it is that a traffic jam is attributed to an asteroid.
- How often traffic backs up on I-74
  - P(B) = P (Backup) = ?
  - The more often this happens the less likely it is to be an indicator of asteroid collision
- $P(A|B) = P(B|A). P(A)/P(B)$ ← Bayes Theorem

Factors to Consider

P (Asteroid given Backup) ?

- $P(A) = P(\text{Asteroid}) = 0.00001$
- $P(B) = P(\text{Backup}) = 0.01$
- $P(B|A) = 1$

- $P(A|B) = P(B|A).\ P(A)/P(B) = 0.001$

Review of Important Theorems

- Total Probability Theorem:
  $P(A) = P(A|C_1)\ P(C_1) + \ldots + P(A|C_n)\ P(C_n)$
  where $C_1, \ldots, C_n$ cover the space of all possibilities

- Bayes Theorem:
  $P(A|B) = P(B|A).\ P(A)/P(B)$

- Other: $P(A,B) = P(A|B)\ P(B)$

Intrusion Detection, Again

- A motion alarm is used to detect unauthorized access to a warehouse after hours. The motion sensor is mounted near the only entrance to the warehouse. If a burglar enters the building, there is a 99% chance that the burglar triggers the motion alarm.
- At 9pm, on September 16th, 2013, the alarm was set off. What are the odds that a burglar is in the building?

Intrusion Detection, Again

- A motion alarm is used to detect unauthorized access to a warehouse after hours. The motion sensor is mounted near the only entrance to the warehouse. If a burglar enters the building, there is a 99% chance that the burglar triggers the motion alarm.
- At 9pm, on September 16th, 2013, the alarm was set off. What are the odds that a burglar is in the building?
- Assume the alarm goes off about 3 days a year and burglaries happen about once a year

Intrusion Detection, Again

- A motion alarm is used to detect unauthorized access to a warehouse after hours. The motion sensor is mounted near the only entrance to the warehouse. If a burglar enters the building, there is a 99% chance that the burglar triggers the motion alarm.
- At 9pm, on September 16th, 2013, the alarm was set off. What are the odds that a burglar is in the building?
- Assume the alarm goes off about 3 days a year and burglaries happen about once a year

  - $P(A) = P(Alarm) = 3/365$
  - $P(B) = P(Burglar) = 1/365$
  - $P(A|B) = 0.99$

Intrusion Detection, Again

- A motion alarm is used to detect unauthorized access to a warehouse after hours. The motion sensor is mounted near the only entrance to the warehouse. If a burglar enters the building, there is a 99% chance that the burglar triggers the motion alarm.
- At 9pm, on September 16th, 2013, the alarm was set off. What are the odds that a burglar is in the building?
- Assume the alarm goes off about 3 days a year and burglaries happen about once a year

  - P(A) = P(Alarm) = 3/365
  - P(B) = P(Burglar) = 1/365
  - P(A|B) = 0.99
  - P(B|A) = P(A|B).P(B)/P(A) = 0.33 (i.e., if alarm sounds, there is a 33% chance it is a burglar)

Multisensor Data Fusion

A Second Sensor

- In the intrusion detection example, assume that there is a vibration sensor on the floor that detects footsteps. If a burglar enters the building, there is a 90% chance that the vibration sensor will fire. If the vibration sensor fires, what are the odds that there is a burglar? Assume that the vibration sensor fires 10 times a year

A Second Sensor

- In the intrusion detection example, assume that there is a vibration sensor on the floor that detects footsteps. If a burglar enters the building, there is a 90% chance that the vibration sensor will fire. If the vibration sensor fires, what are the odds that there is a burglar? Assume that the vibration sensor fires 10 times a year
  - P (Burg|vib) = P(Vib|Burg).P(Burg)/P(Vib)

    = 0.9 * (1/365) / (10/365) = 0.09

Two Sensor Example

- In the intrusion detection example, what are the odds of burglary if both sensors fire?
- P (Burg|A, Vib) = ?

Remember: If burglar enters, motion alarm fires 99% of the time and vibration alarm fires 90% of the time. Burglaries occur once a year, motion alarm fires 3 times a year, and vibration alarm fires 10 times a year.

Two Sensor Example

- In the intrusion detection example, what are the odds of burglary if both sensors fire?
- P (Burg|A, Vib) = ?
- $P(B|A,V) = P(A,V|B)\,P(B)/P(A,V)$

Remember: If burglar enters, motion alarm fires 99% of the time and vibration alarm fires 90% of the time. Burglaries occur once a year, motion alarm fires 3 times a year, and vibration alarm fires 10 times a year.

Remember: If burglar enters, motion alarm fires 99% of the time and vibration alarm fires 90% of the time. Burglaries occur once a year, motion alarm fires 3 times a year, and vibration alarm fires 10 times a year.

Two Sensor Example

- In the intrusion detection example, what are the odds of burglary if both sensors fire?
- $P(Burg|A, Vib) = ?$
- $P(B|A,V) = P(A,V|B)\,P(B)/P(A,V)$

Now what?

Is it OK to say $P(A,V|B) = P(A|B)P(V|B)$?

Is it OK to say $P(A,V) = P(A)P(V)$?

Independence versus Conditional Independence

- John and Sally follow Mike on Twitter.
- When Mike tweets something, John re-tweets it with a 50% probability. Sally re-tweets it with a 30% probability.
- Are John's and Sally's tweets independent?

Independence versus Conditional Independence

- John and Sally follow Mike on Twitter.
- When Mike tweets something, John re-tweets it with a 50% probability. Sally re-tweets it with a 30% probability.
- Are John's and Sally's tweets independent?
  - No. However, given that Mike says something, their decisions to re-tweet it are independent (conditional independence)

Remember: If burglar enters, motion alarm fires 99% of the time and vibration alarm fires 90% of the time. Burglaries occur once a year, motion alarm fires 3 times a year, and vibration alarm fires 10 times a year.

Two Sensor Example

- In the intrusion detection example, what are the odds of burglary if both sensors fire?
- P (Burg|A, Vib) = ?
- P (B|A,V) = P(A,V|B) P(B)/P(A,V)

Now what?

OK to say P(A,V|B) = P(A|B)P(V|B)

~~P(A,V) = P(A)P(V)~~?

Remember: If burglar enters, motion alarm fires 99% of the time and vibration alarm fires 90% of the time. Burglaries occur once a year, motion alarm fires 3 times a year, and vibration alarm fires 10 times a year.

Two Sensor Example

- In the intrusion detection example, what are the odds of burglary if both sensors fire?
- P (Burg|A, Vib) = ?
- $P(B|A,V) = P(A,V|B)\,P(B)/P(A,V)$ where

  $P(A,V) = P(A,V|B)\,P(B) + P(A,V|\bar{B})\,P(\bar{B})$

  and $P(A,V|B) = P(A|B)P(V|B)$

Remember: If burglar enters, motion alarm fires 99% of the time and vibration alarm fires 90% of the time. Burglaries occur once a year, motion alarm fires 3 times a year, and vibration alarm fires 10 times a year.

Two Sensor Example

- In the intrusion detection example, what are the odds of burglary if both sensors fire?
- P (Burg|A, Vib) = ?
- P (B|A,V) = P(A,V|B) P(B)/P(A,V) where

$$P(A,V) = P(A,V|B)\,P(B) + P(A,V|\bar{B})\,P(\bar{B})$$

and $P(A,V|B) = P(A|B)P(V|B)$

$$P(A,V|\bar{B}) = P(A|\bar{B})P(V|\bar{B})$$

Remember: If burglar enters, motion alarm fires 99% of the time and vibration alarm fires 90% of the time. Burglaries occur once a year, motion alarm fires 3 times a year, and vibration alarm fires 10 times a year.

Two Sensor Example

- P (Burg|A, Vib) Solution steps:
  - Find the probability of false alarms from:

    $P(A) = P(A|B)\,P(B) + P(A|\bar{B})\,P(\bar{B})$

    $P(V) = P(V|B)\,P(B) + P(V|\bar{B})\,P(\bar{B})$
  - Find the probability of both sensors firing:

    $P(A,V) = P(A,V|B)\,P(B) + P(A,V|\bar{B})\,P(\bar{B})$

    where $P(A,V|B) = P(A|B)P(V|B)$

    $P(A,V|\bar{B}) = P(A|\bar{B})P(V|\bar{B})$
  - $P(B|A,V) = P(A,V|B)\,P(B)/P(A,V) = 94.62\%$

Key Things to Remember

- Total Probability Theorem:
  - P(A) = P(A|C1) P(C1) + … + P(A|Cn) P(Cn)

  where C1, …, Cn cover the space of all possibilities
- Bayes Theorem: P(A|B) = P(B|A). P(A)/P(B)
- Other: P(A,B) = P(A|B) P(B)

Practice and Review: Conditional Probabilities

- A driver with a history of substance abuse is severely injured when his car spins out of control hitting a stop sign. Police and ambulance vehicles are dispatched. A police officer at the scene suggests to a paramedic that the driver must have been intoxicated. The paramedic answers: "Statistically, only 10% of individuals with substance abuse problems end up in severe car crashes, so it may well have been something else".
  - Is the argument valid?

Practice and Review: Conditional Probabilities

- A driver with a history of substance abuse is severely injured when his car spins out of control hitting a stop sign. Police and ambulance vehicles are dispatched. A police officer at the scene suggests to a paramedic that the driver must have been intoxicated. The paramedic answers: "Statistically, only 10% of individuals with substance abuse problems end up in severe car crashes, so it may well have been something else".
  - Is the argument valid?
  - No. The relevant statistic here is: Given that a person with a history of substance abuse ends up in a severe car crash, what are the odds that there were under influence? (These odds are very high.)

Practice and Review: Conditional Probabilities

- An obstacle sensor on a robot is used to detect nearby obstacles on the path. When an obstacle is present, it is detected with a 99.9% probability. When the sensor fires, what is the probability that an obstacle lies in the robot's path?

Practice and Review: Conditional Probabilities

- An obstacle sensor on a robot is used to detect nearby obstacles on the path. When an obstacle is present, it is detected with a 99.9% probability. When the sensor fires, what is the probability that an obstacle lies in the robot’s path?
  - Not clear; we do not know the rate of false positives (instances when the sensor fires in the absence of obstacles).

Conditional Independence

- Two obstacle sensors are mounted near each other on the front of a robot. The probability that one of them fires in a particular time interval is 5%. What is the probability that both fire simultaneously?

Conditional Independence

- Two obstacle sensors are mounted near each other on the front of a robot. The probability that one of them fires in a particular time interval is 5%. What is the probability that both fire simultaneously?
  - We cannot compute the numeric answer because the sensors are not independent. We need more information.

Conditional Independence

- Two obstacle sensors are mounted near each other on the front of a robot. *When an obstacle appears in front of the robot*, the probability that one of them fires is 75%. What is the probability that both fire at that time?

Conditional Independence

- Two obstacle sensors are mounted near each other on the front of a robot. *When an obstacle appears in front of the robot*, the probability that one of them fires is 75%. What is the probability that both fire at that time?
  - $(0.75)^2$

Multisensor Fusion

- A defense sensor (radar) is used to detect incoming missiles. The probability of a missile being fired in the vicinity of the sensor is 1%. When a missile is fired, the sensor detects it with 98% probability. The sensor also has a 20% false positive rate. Three sensors are deployed. When all three sensors indicate a missile, what is the probability that an actual missile is fired?

Three Sensor Example

- $P(M|S1,S2,S3) = 54.3\%$
- $P(M|S1,S2,S3) = P(S1,S2,S3|M)\,P(M)/P(S1,S2,S3)$, where:

$$P(S1,S2,S3) = P(S1,S2,S3|M)\,P(M) + P(S1,S2,S3|\text{no }M)\,P(\text{no }M)$$

and

$$P(S1,S2,S3|M) = P(S1|M)P(S2|M)P(S3|M)$$

$$P(S1,S2,S3|\text{no }M) = P(S1|\text{no }M)P(S2|\text{no }M)P(S3|\text{no }M)$$

Multisensor Fusion

- In the preceding example, how many sensors do you need in order to raise the probability of there being a target (when all sensors fire) to over 95%?

Robot Example:

- A robot has a camera that detects obstacles with probability 70%, a bump sensor that detects imminent collisions with a probability of 99.9% (when an obstacle is 1 inch away), and a cliff sensor that detects imminent falls off a cliff with a probability of 99.9% (when the cliff is 1 inch away). The robot has breaks that can stop it within 0.1 second. The mission is to deliver supplies from point A to point B, safely.
  - What are safety-critical requirements?
  - What are mission-critical (i.e., performance) requirements?
  - What is a safe state?
  - How to ensure well-formed dependencies?
  - What is a safe speed for the robot?
  - Is the algorithm that computes speed based on preferred arrival time and route safety-critical or mission-critical?