Inductive Reasoning

Logical Reasoning and Human Nature

- Historically, many researchers believed that logical reasoning is an essential part of human nature
  - Aristotle
  - Boole (1854). A book on logical calculus
    "An investigation of the laws of thought"
- Rational behavior = logical thinking (deductive reasoning)

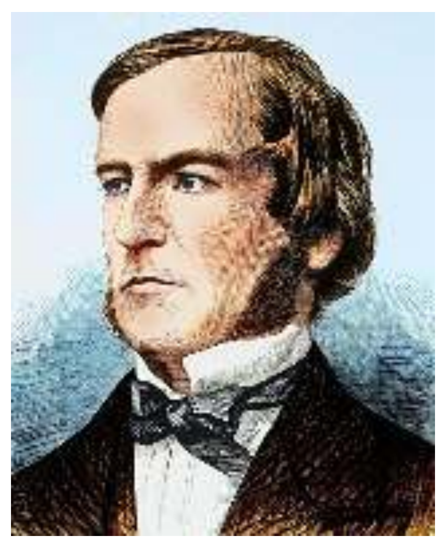

Inductive vs. Deductive Reasoning

- Deductive reasoning:
  - Conclusion follows logically from premises
- Inductive reasoning:
  - Conclusion is likely based on premises
  - Involves a degree of uncertainty
- Reasoning in real-world is often based on induction

- In computer science:
  - Logical reasoning is often associated with symbolic AI
  - Inductive reasoning is often associated with statistical machine learning

Deductively Valid?

- Premise: All cars have wheels
- Premise: All wheels are round
- Conclusion: All cars have round wheels

- Premise: I have a diamond
- Premise: Most diamonds are shiny
- Conclusion: My diamond is shiny

- Premise: John is 93
- Conclusion: John will not do a double back flip today

Inductive Reasoning

- Reason from observable information to unobservable and uncertain information

What makes people smart?

- Memory? No.
- Deductive inference? No.
- Intuitions and inductions.

Everyday Inductive Leaps

- How can people learn so much about the world from such limited evidence?
  - Learning concepts from examples – “horse”






Real World Inductive Inferences

- Medical diagnosis:
  - Symptoms, test outcomes (observable) → Diseases (unobservable)
- Scientific reasoning:
  - Experimental data (observable) → Hypotheses (unobservable)
- Law:
  - Facts **(证据)** (observable) → Guilt **(罪行)** (unobservable, uncertain)

Reasoning under Uncertainty

- Bayes rule tells us how to optimally reason with uncertainty
  - Bayes estimation
- Allows us to say how we believe something to be true based on prior beliefs and new available evidence

Thomas Bayes (1702-1761)

Bayes Rule

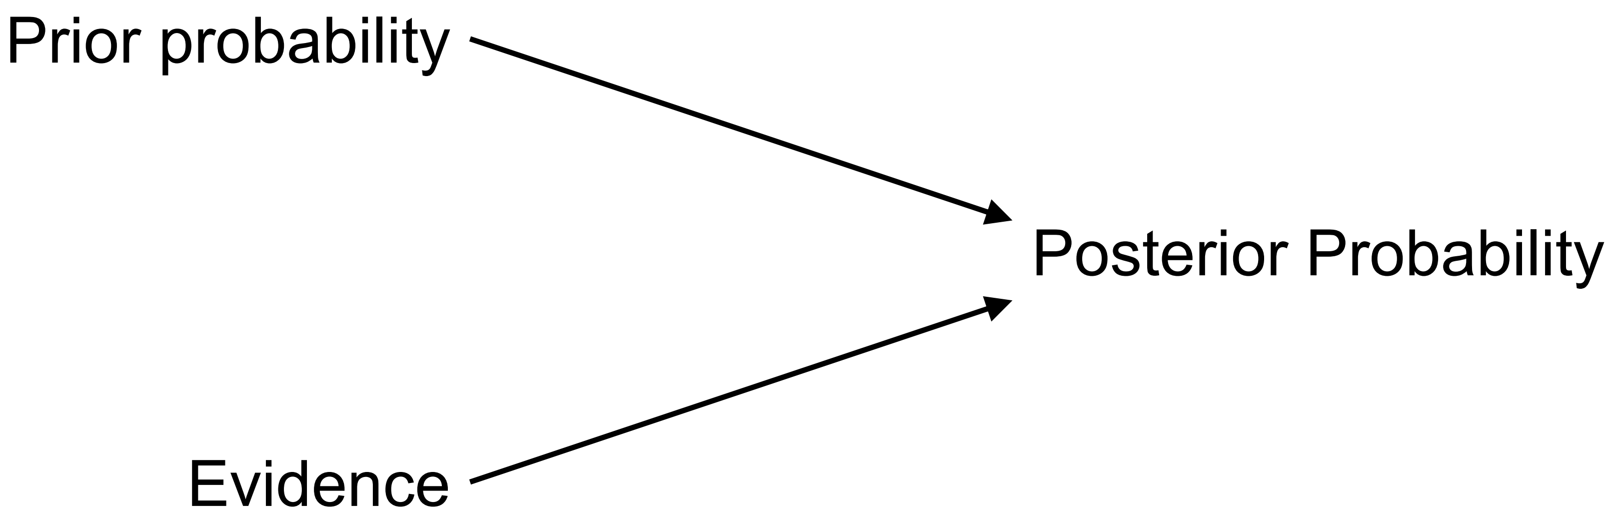

Bayes rule tells us how the available evidence should alter our belief in something being true

Do People Reason like Bayes Rule?

- Problems understanding **conditional probability**
  - Doctors need to calculate the probability of disease given the observed symptoms: P(disease | symptoms)
  - Sometimes P(symptoms | disease) is used incorrectly when reasoning about the likelihood of a disease

- Why is this wrong?

The Base Rate is Important

- To get P(disease | symptom), you need to know about P(symptom | disease) and also the **base rate** -- prevalence of the disease before you have seen patient
- More intuitive example:
  - What is the probability of being tall given you are player in the NBA?
  - What is the probability of being a player in the NBA given that you are tall?

P (NBA player | tall) ≠ P (tall | NBA player)

Reasoning with Base Rates

- Suppose there is a disease that affects 1 out of 100 people
- There is a diagnostic test with the following properties:
  - If the person has the disease, the test will be positive 98% of the time
  - If the person does not have the disease, the test will be positive 1% of the time
- A person tests positive, what is the probability that this person has the disease?
  - Frequent answer =
  - Correct answer ≈

Are we really that bad in judging probabilities?

- According to some researchers (e.g., Gigerenzer), it matters how the information is presented and processed

- Processing frequencies is more intuitive than probabilities

A Counting Heuristic (in tree form)

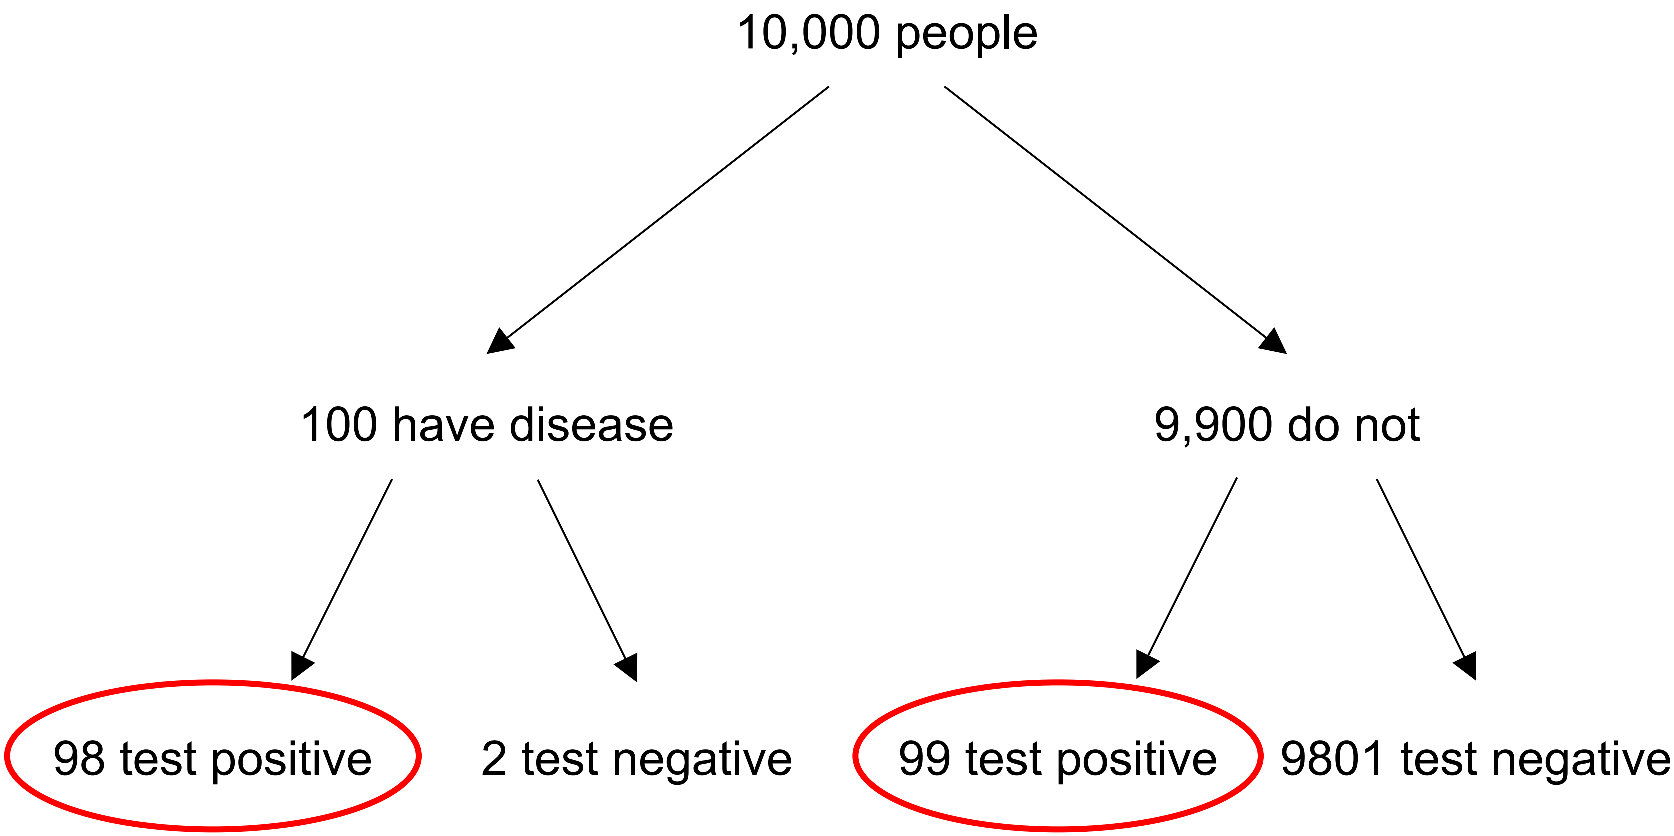

P(disease | test positive) = 98 / (98 + 99) ≈ .50

The Same Thing in Words ...

- Let’s take 10,000 people
- On average, 100 out of 10,000 actually have the disease and 98 of those will test positive (98% true positive rate)
- Among the 9,900 who do not have the disease, the test will falsely identify 1% as having it. 1% of 9,900 = 99
- On average, out of 10,000 people:
  98 test positive and they have the disease
  99 test positive and they do not have the disease
- Therefore, a positive test outcome implies a 98/(98+99) ≈ 50% chance of having the disease

Change the Example

- What now if the disease affects only 1 out of 10,000 people?

- Assume same diagnosticity of test (98% true positive rate, 1% false positive rate)

- A person tests positive, what now is the probability that this person has the disease?

A Counting Heuristic (in tree form)

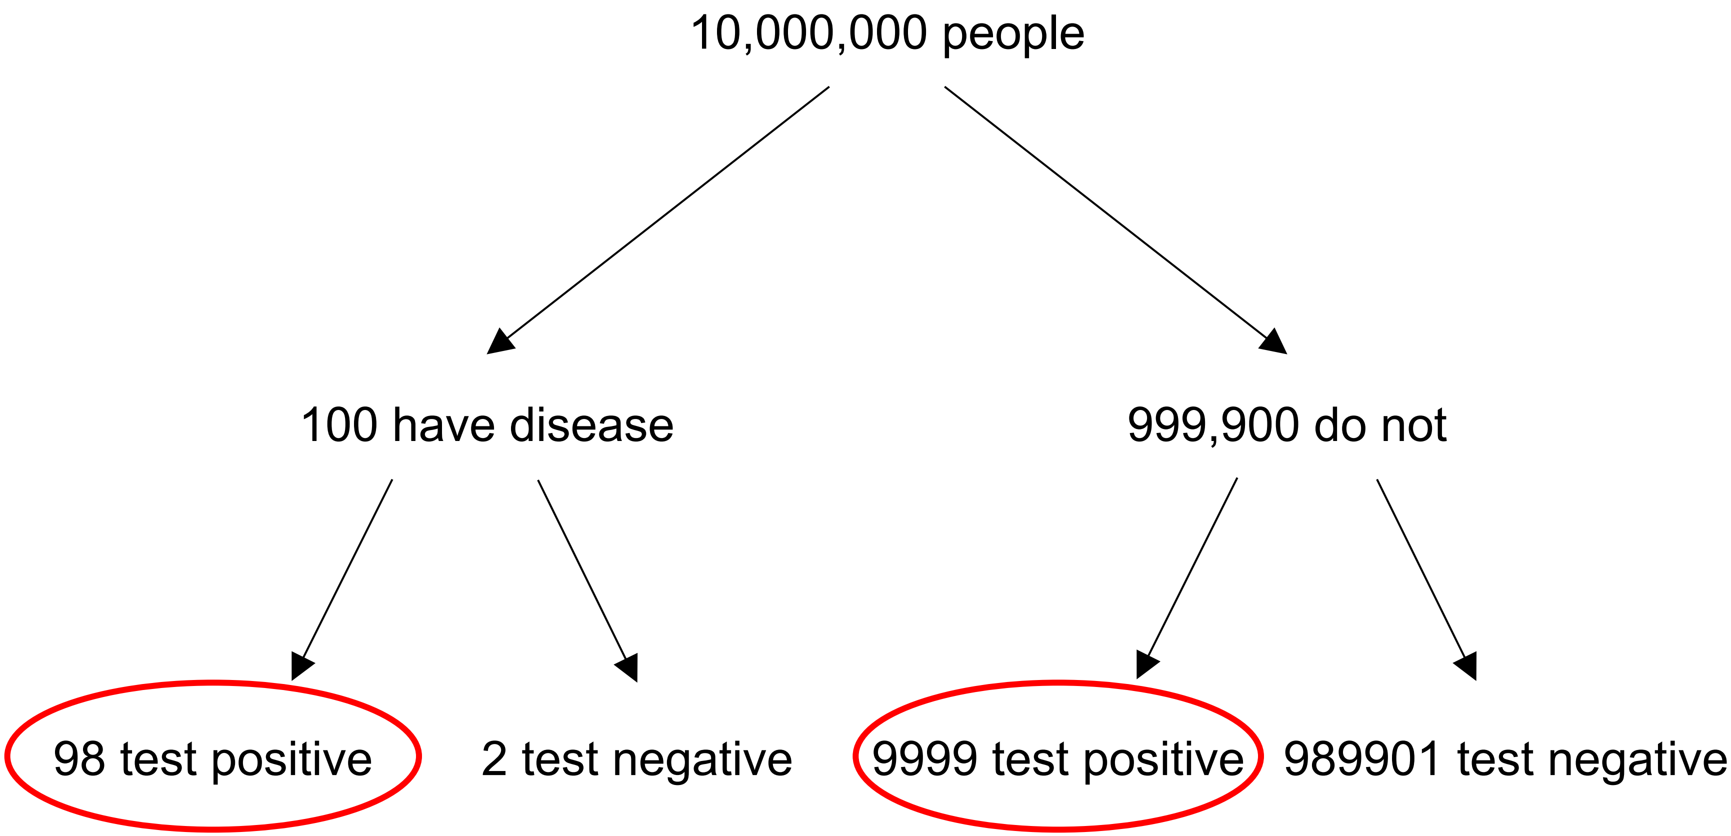

P (disease | test positive) = 98 / (98 + 9999) = .0097

(smaller than 1%)

Bayes Rule

- The previous example essentially is a simple way to apply Bayes rule:

$$P(disease \mid positive) = \frac{P\left(positive \mid disease\right) P(disease)}{P\left(positive \mid disease\right) P(disease) + P\left(positive \mid not\ disease\right) P(not\ disease)}$$

$$P\left(positive \mid disease\right) = .98$$

$$P\left(positive \mid not\ disease\right) = .01 \quad \Rightarrow \quad P\left(disease \mid positive\right) = .0097$$

$$P\left(disease\right) = .0001$$

Normative Model

- Bayes rule tells you how you should reason with probabilities -- it is a prescriptive (i.e., **normative (标准模型)**) model
- But do people reason like Bayes?

  In certain circumstances, we observe **base rate neglect (基率忽视)**

The **base rate fallacy**, also called **base rate neglect** or **base rate bias**, is an error that occurs when the conditional probability of some hypothesis H given some evidence E is assessed without taking into account the prior probability ("base rate") of H and the total probability of evidence E