Carnegie Mellon University
Language Technologies Institute

Artificial Intelligence

What is it?

Anatole Gershman

What is Intelligence?

What is Intelligence?

Merriam-Webster

A) **The ability to learn or understand or to deal with new or trying situa**

B) **The ability to apply knowledge to manipulate one's environment**

Carnegie Mellon University
Language Technologies Institute

What is Intelligence?

Merriam-Webster

A) The ability to learn or understand or to deal with new or trying situa
B) The ability to apply knowledge to manipulate one's environment

AlphaGo

AlphaGo Zero is by far the world best chess and go pl

Is it intelligent?

Intelligent Machines

This course is about how to make machines intelligent

Intelligent Machines

This course is about how to make machines intelligent

But first, a few words myself

A few words about myself

Carnegie Mellon University
Language Technologies Institute

Born in Moscow

Ph.D. Yale University

Bell Laboratories
Schlumberger
accenture
Carnegie Mellon University
School of Computer Science

40 years in Industry and Academi

Why am I interested in AI?

Carnegie Mellon University
Language Technologies Institute

How the mind emerges from the brain?

What makes us smart?

Building smart machines may help us answer these questions

Two views of Artificial Intelligence

General Intelligence:

How the mind emerges from the brain

Task-specific Intelligence:

How to automate tasks that require intelligence

Which tasks require intelligence?

In the early days of computing, text formatting was considered a task that required intelligence

EMACS editor was developed at the MIT AI lab

In the 1970-80s, intelligence was required to execute rule-based tasks. Expert systems were developed to automate these tasks. They separated the rules from the execution “engine”.

Now, every business system does that without calling it “expert system” or AI.

Which tasks require intelligence?

In your opinion, what are the most intelligent apps today?

AI as automation of Cognitive Tasks

- **For businesses, technology has always been a means of automation**
- **Artificial Intelligence is not an exception**

- **Technology now enables automation of cognitive tasks that require perception, reasoning and decision making**

Cognitive Tasks

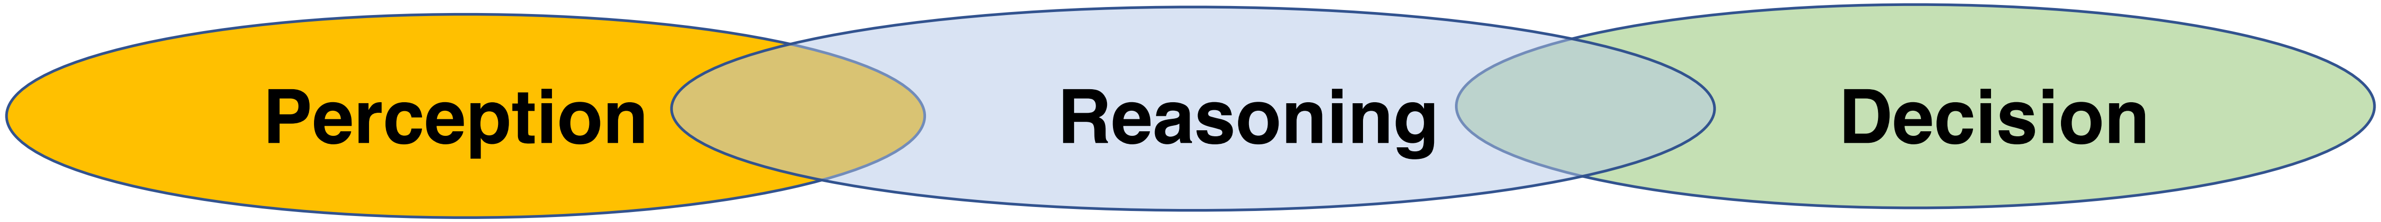

What do I see or hear? **What does it mean?** **What do I do?**

Cognitive Tasks

Carnegie Mellon University
Language Technologies Institute

Perception | **Reasoning** | **Decision**

What do I see or hear?	**What does it mean?**	**What do I do?**

Customer:
My printer does not print red

Hypotheses:
Out of ink .8
Nozzle clogged .4

Robot:
Did you check the red ink level?

Cognitive Tasks

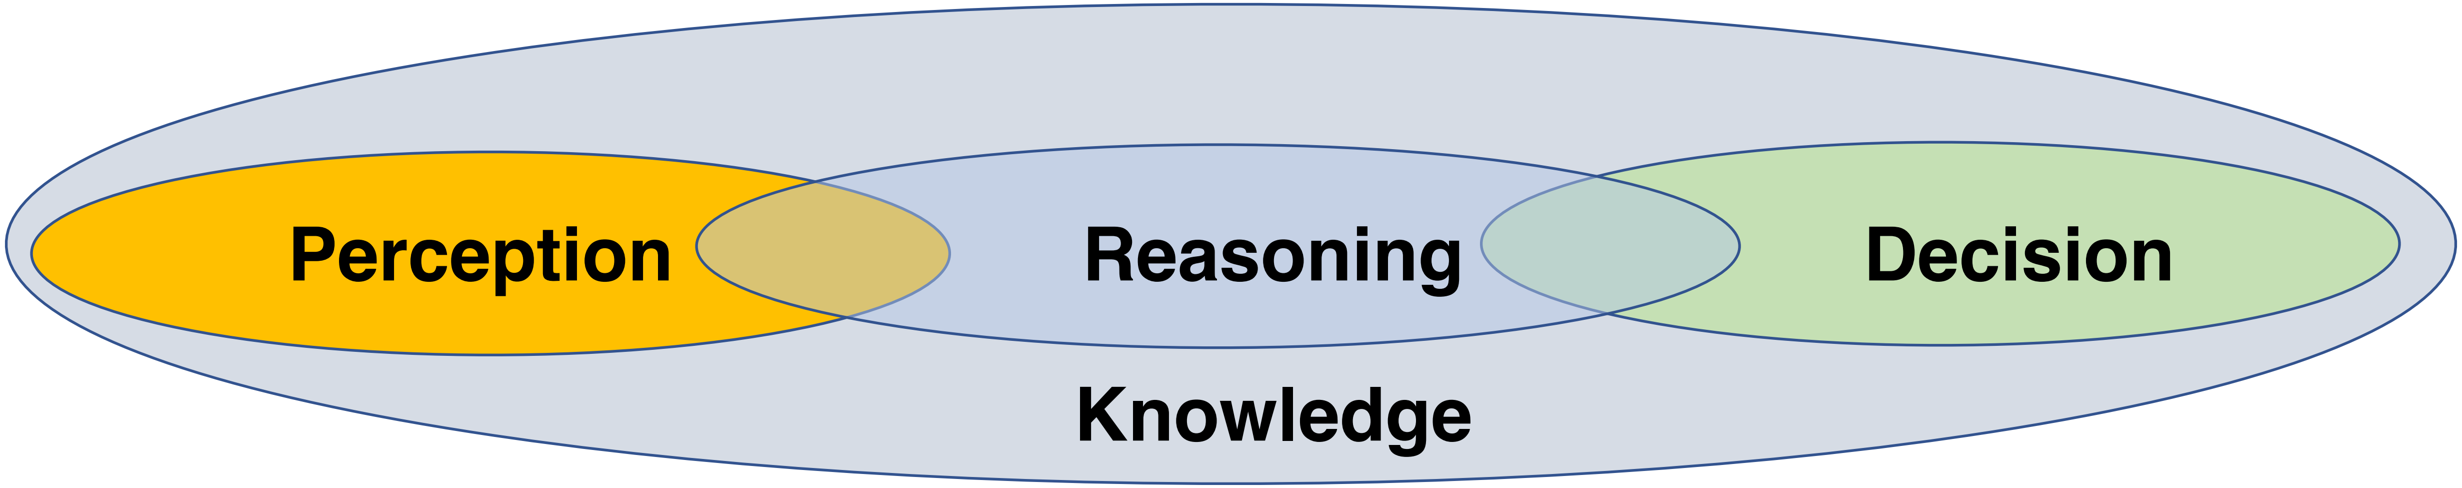

What do I see or hear? **What does it mean?** **What do I do?**

<u>Customer</u>:
My printer does not print red

<u>Hypotheses</u>:
Out of ink .8
Nozzle clogged .4

<u>Robot</u>:
Did you check the red ink level?

Life Insurance Underwriting Example

Carnegie Mellon University
Language Technologies Institute

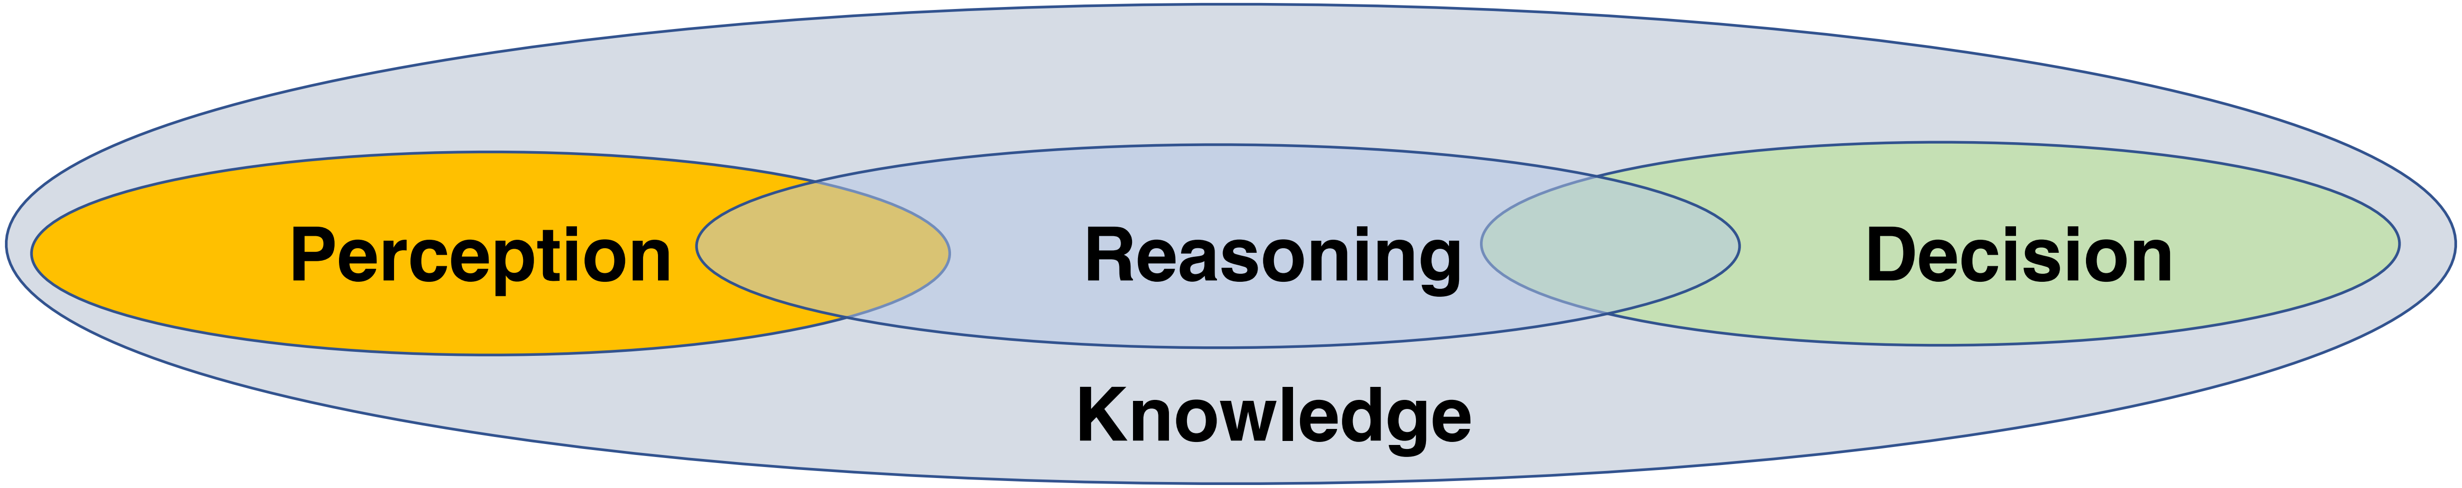

What do I see or hear?

What does it mean?

What do I do?

Application Medical Data: Moderately high blood sugar

Excess Mortality +75%

Approve, Standard Rate

Perception

- **Images** – recognize objects and events on a picture
- **Audio** – recognize sounds, speech to text
- **Video** – recognize object movement and interactions
- **Text** – extract mentions of entities, relations and eve

Perception

Image Recognition

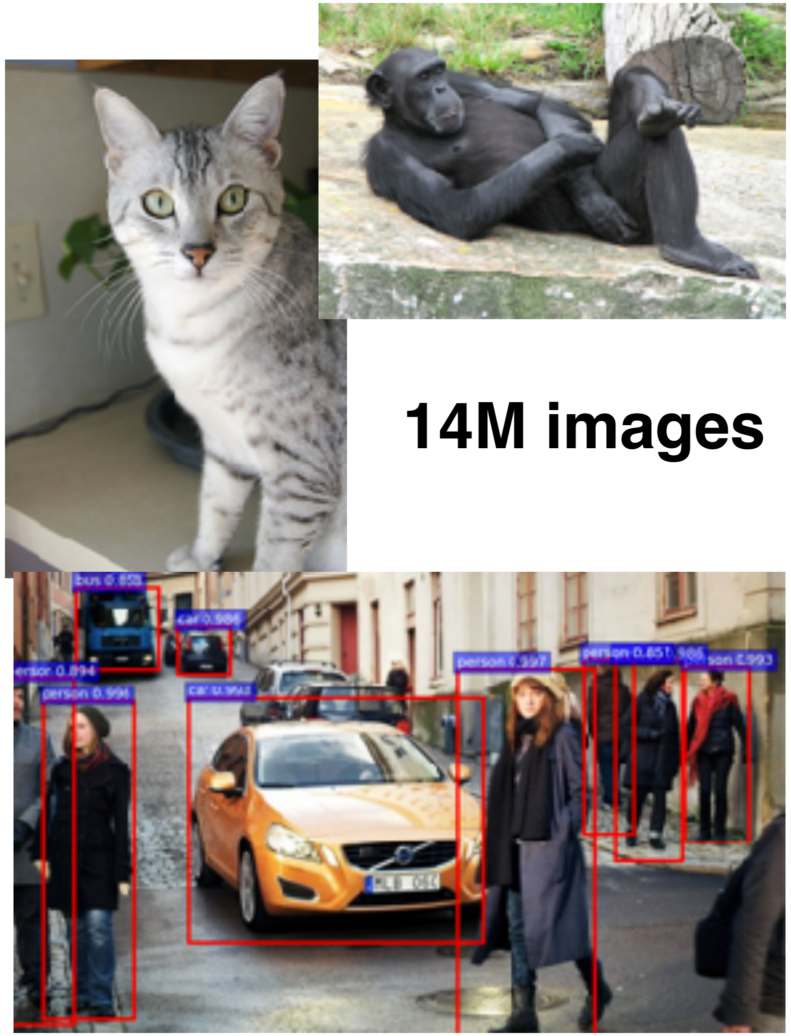

14M images

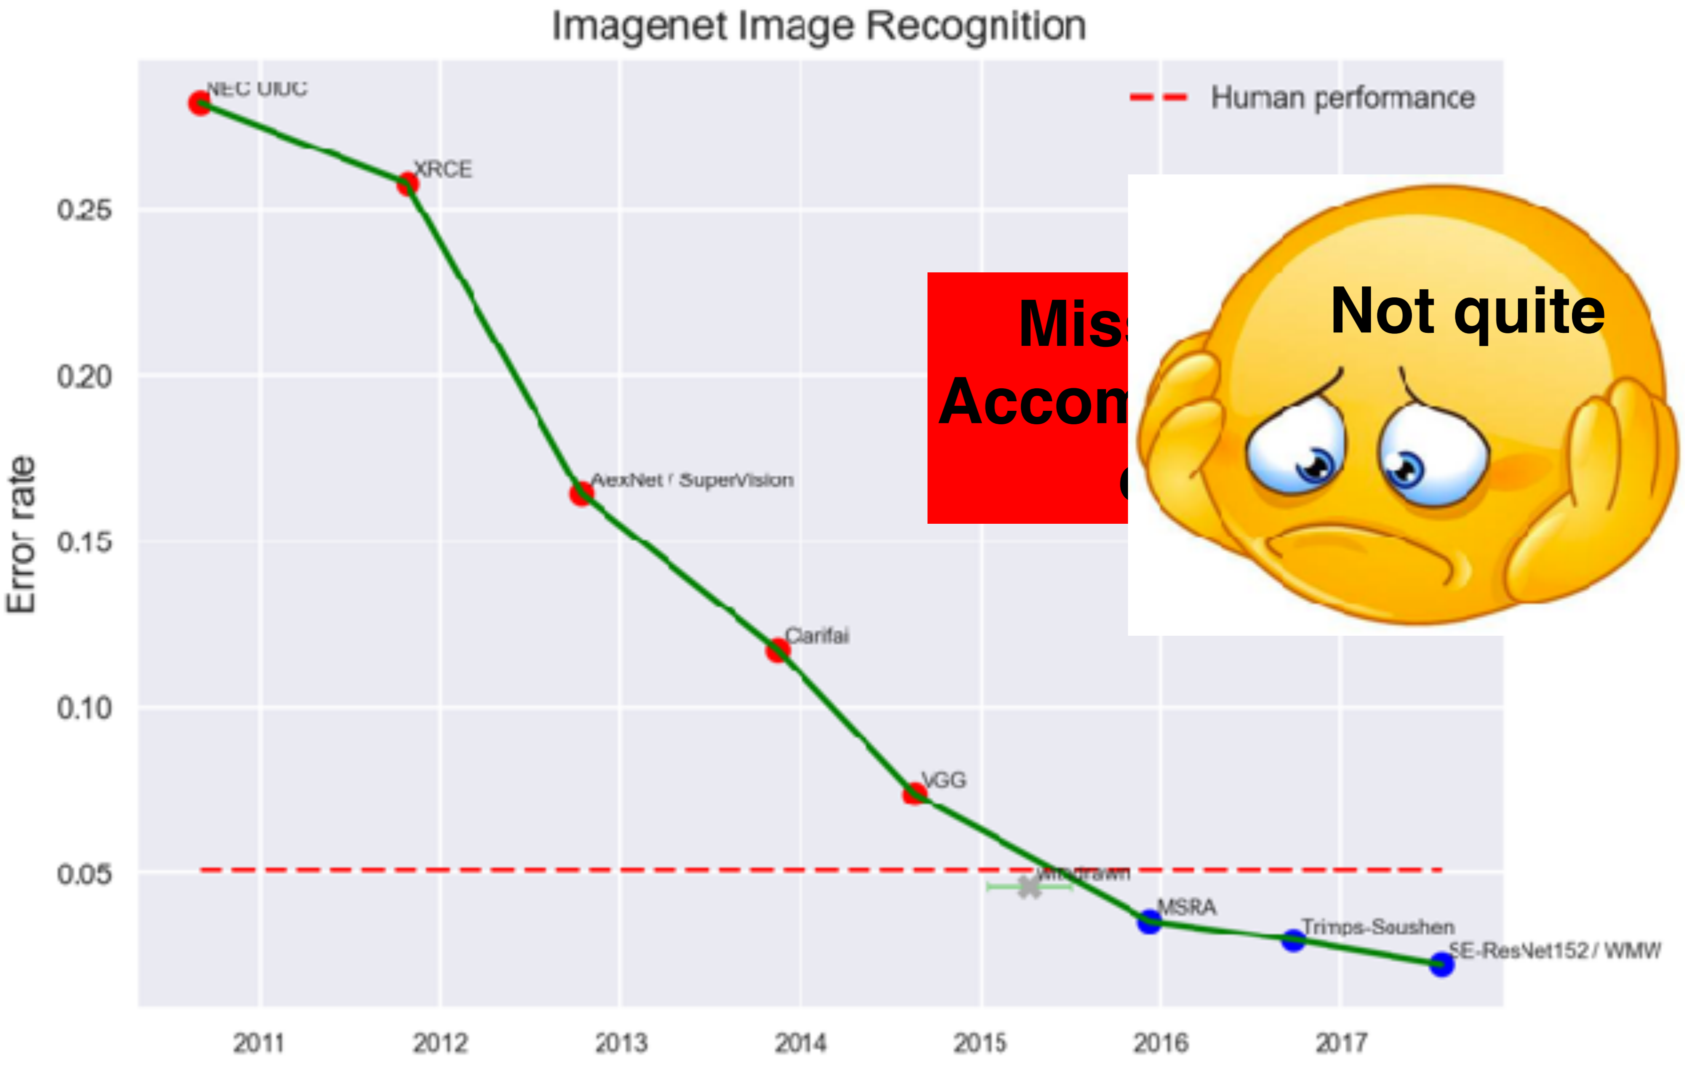

Perception

Carnegie Mellon University
Language Technologies Institute

Image Recognition remains h

- Accuracy in natural settings is not nearly as good
- Training requires thousands of pictures – a person can learn from one
- Easy to fool and manipulate

From Piekniewski blog:

Cat 98% conf.

No idea – best guess
"fire screen" 10%

From OpenAI blog:

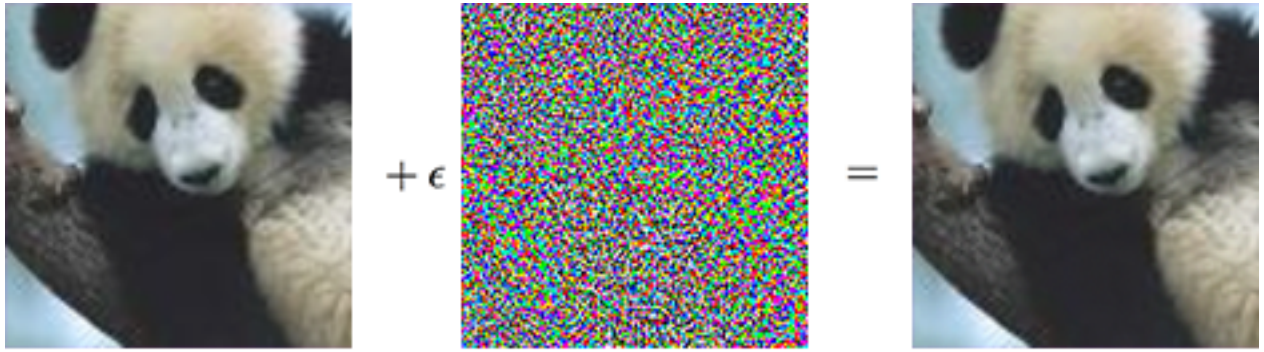

Panda 57.7% conf. + a little noise Gibbon 98.3% conf.

Perception

Carnegie Mellon University
Language Technologies Institute

What is the problem? Why is perception hard?

Are there cars in this picture?

- ➢ The lack of common sense
- ➢ Good perception requires reasoning

Perception

Speech Recognition

Carnegie Mellon University
Language Technologies Institute

“Switchboard” data set of 40 phone calls

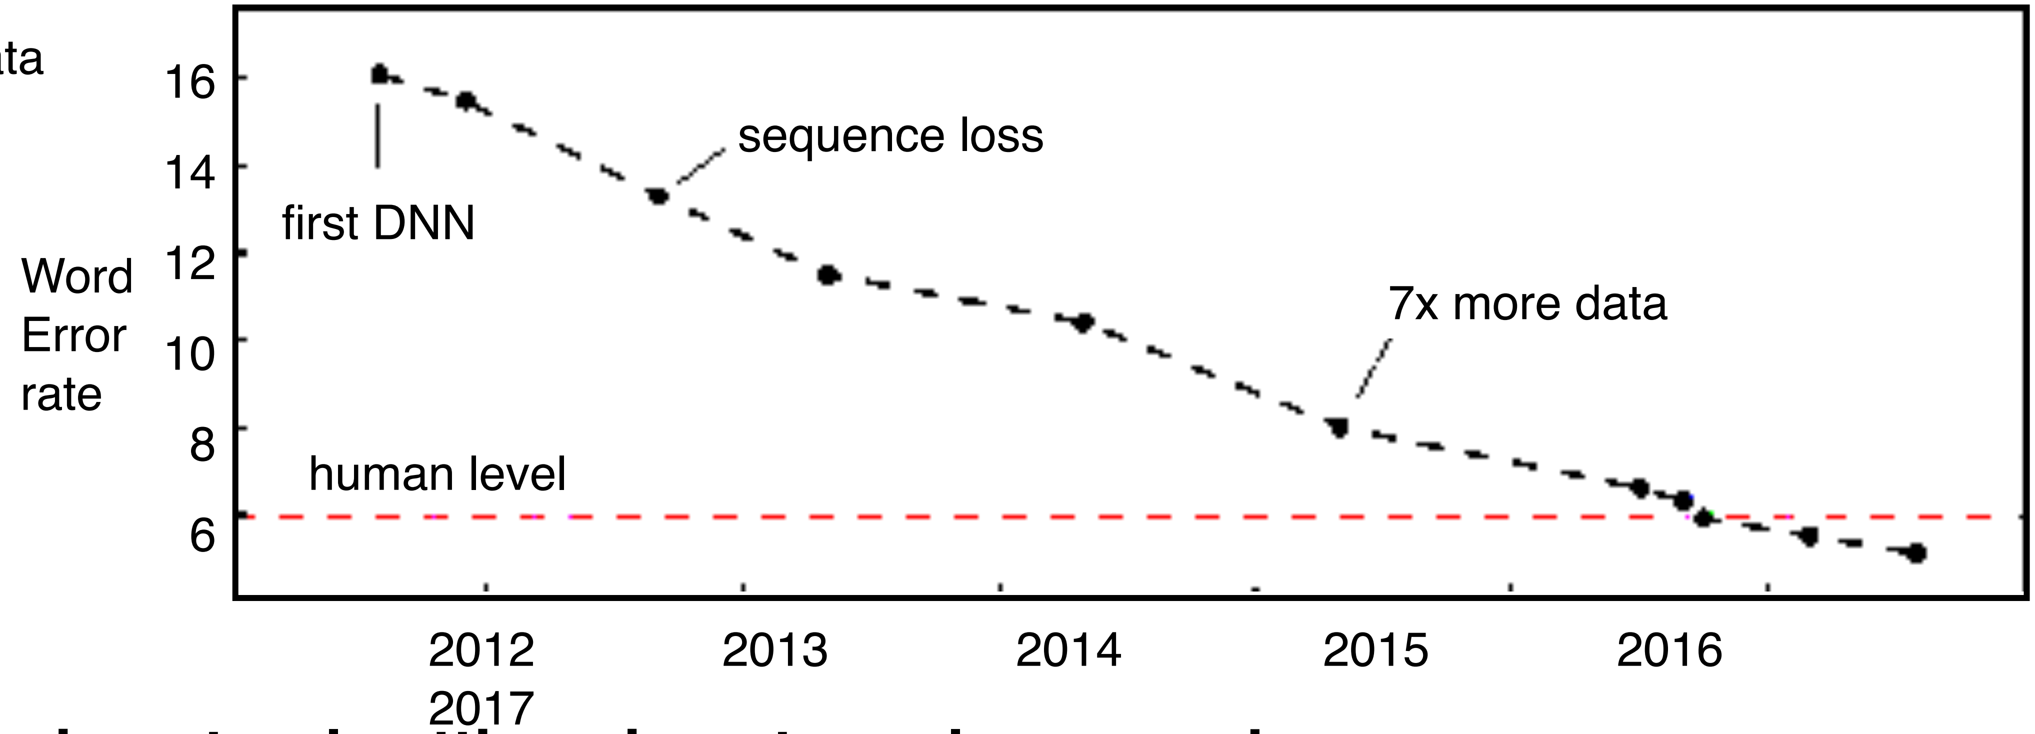

- ❑ **Accuracy in natural settings is not nearly as good**
- ❑ **Easy to fool and manipulate**

I can play “slightly doctored” Beethoven to Alexa and it will order a case of be

Perception

Text Processing

- ➢ Indexing and Search
- ➢ Information Extraction
- ➢ Machine Translation

- Information Extraction is by far most important for cognitive task automation and remains the most challenging
- There are no good universal benchmarks for Information Extraction
- Statistical methods alone cannot solve the Information Extraction problem

Perception

Text Processing

I saw the Grand Canyon flying to New York

Who is flying?

According to the most popular NLP tools: the Grand Canyon is flying

Google Translate: 그랜드 캐년이 뉴욕으로 날아가는 것을 봤습니다

To achieve robustness and high accuracy, statistics is not enough, perception needs reasoning and knowledge

Reasoning

Carnegie Mellon University
Language Technologies Institute

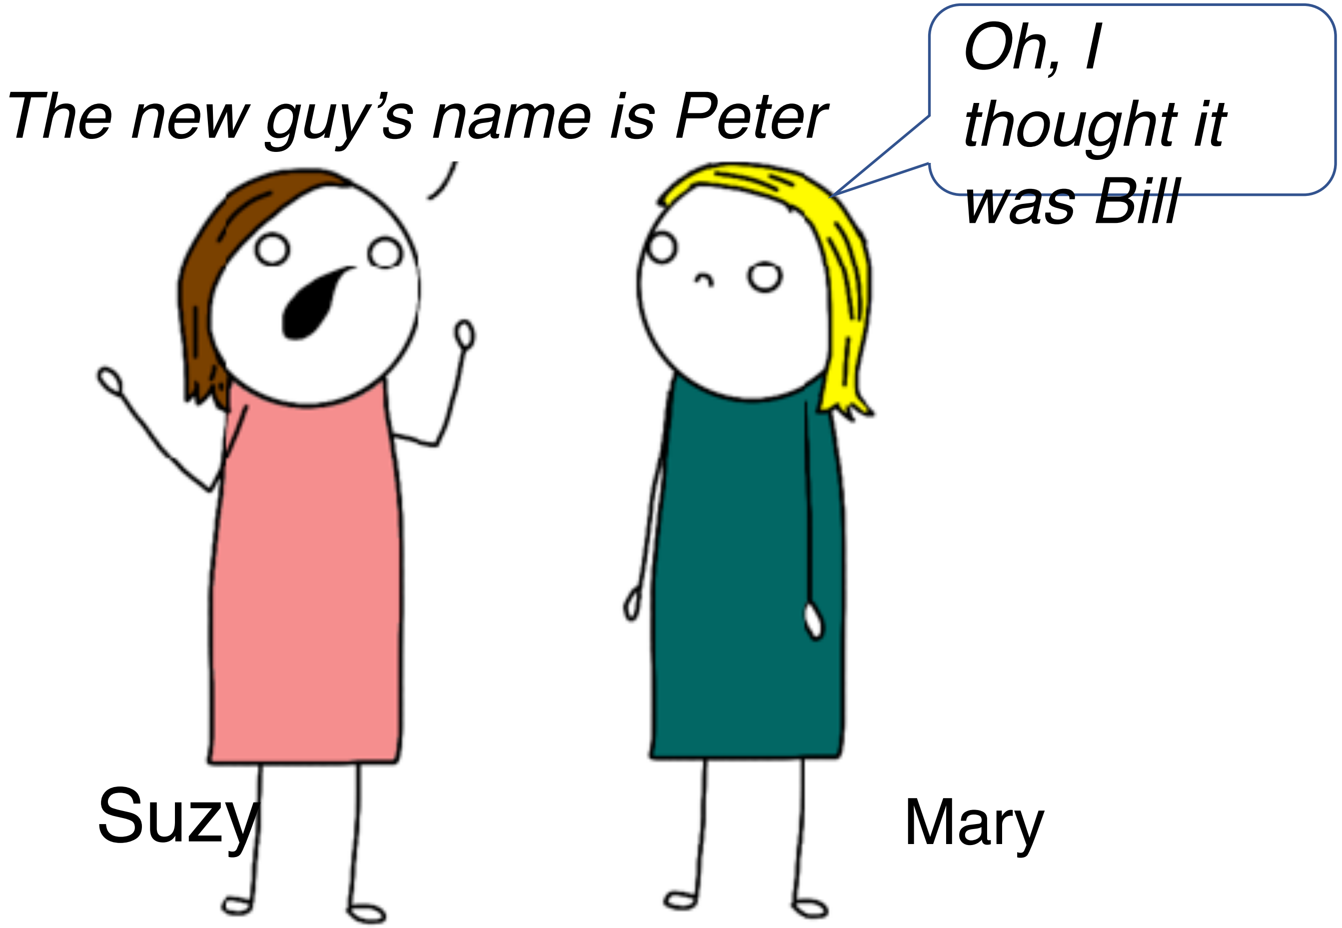

What should Mary believe now?

Reasoning

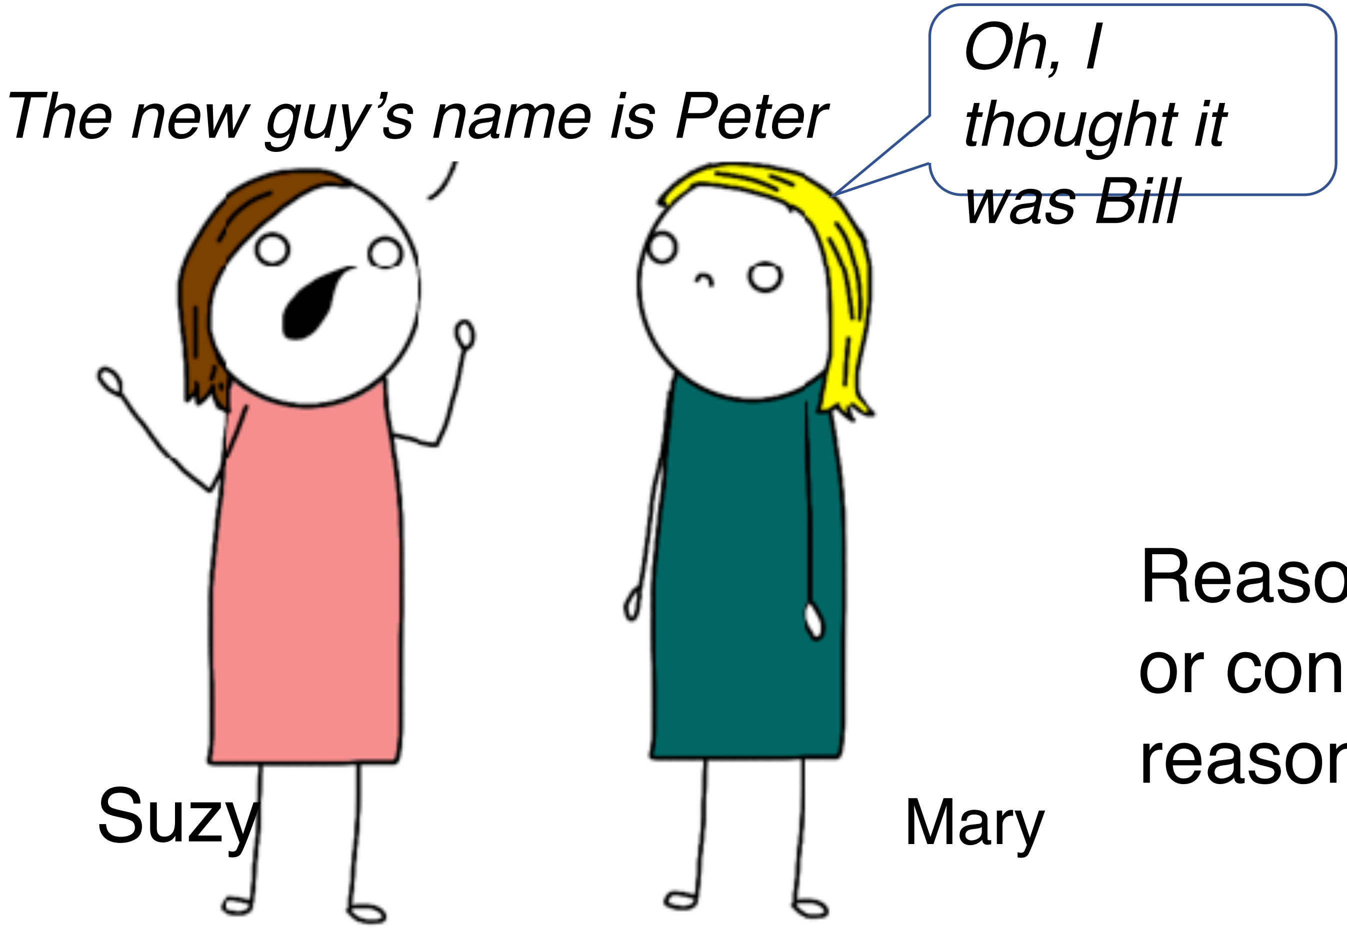

What should Mary believe now?

Merriam-Webster

Reasoning: the drawing of inferences or conclusions through the use of reason

Reasoning

Carnegie Mellon University
Language Technologies Institute

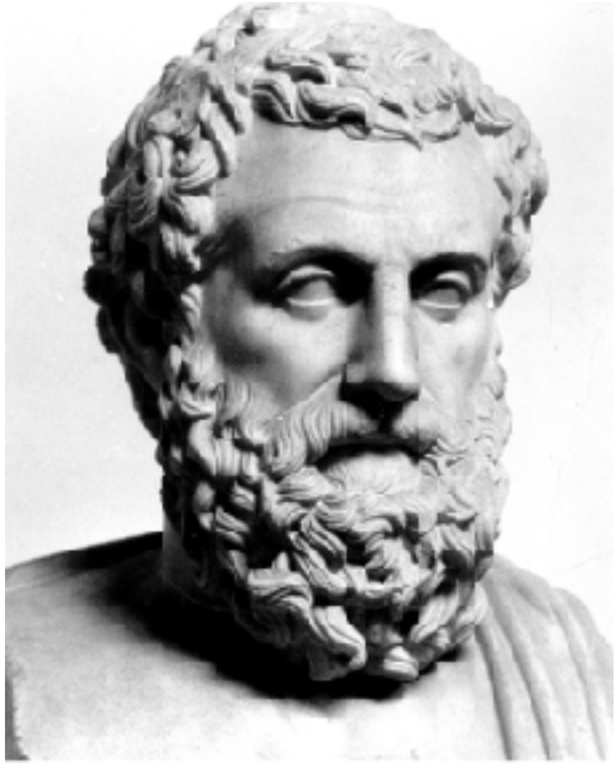

Aristotle
384 - 322 BC

At least since Aristotle, reason has been equated with logic
Logic deals with propositions:

P - name is Peter
B - name is Bill
S - Suzy says that the name is Peter
R - Suzy is trustworthy (reliable)

Given:

$$P \bigwedge B = False$$
$$S \bigwedge R \rightarrow P$$
$$S = True$$
$$R = True$$

We can prove that

$$P = True$$
$$B = False$$

Reasoning

Carnegie Mellon University
Language Technologies Institute

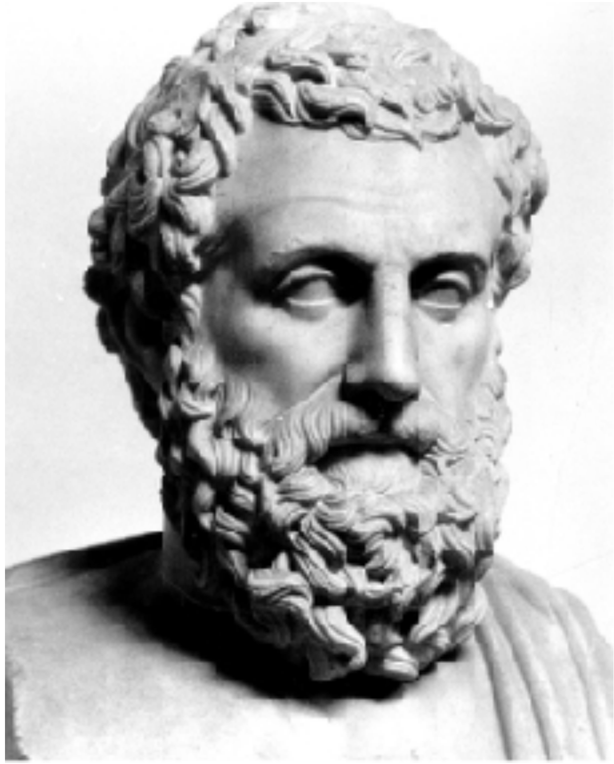

Aristotle
384 - 322 BC

At least since Aristotle, reason has been equated with logic
All this math is a fancy way of saying:

If Person X says Y and if X is trustworthy then Y is True

Since Suzy is trustworthy then Mary should now believe that the neighbor's name is Peter, not Bill

But we can prove it mathematically and such proofs can be automated

Reasoning

Expert Systems

In the 1970-80s many “Expert Systems” were built to conduct reasoning based on the rules of logic.

They had thousands of rules and were quite complex

Logical rule engines were what the neural nets are today

It is common now to think that they failed, but this is not true

Most business systems today are their descendants

Problems with Logic

1. **Real rules are often not deterministic**
   Diabetes increases the likelihood of death
2. **We are often uncertain about facts**
   We are 70% confident that the patient has diabetes

What if Suzy was 80% trustworthy and Mary was only 60% confident that the neighbor's name was Bill?

Reasoning

What do we mean when we say:

Mary is 60% confident that the neighbor's name is Bill?

Suppose, we assign a "plausibility" or "confidence" number to every hypothesis

If two hypotheses are equally plausible, their plausibility numbers should be the same

Small changes in a hypothesis should lead to small changes in its plausibility

True hypotheses should have higher plausibility numbers than false hypotheses

It can be shown mathematically that under these assumptions, plausibility numbers must be proportional to probabilities and obey probability rules

Reasoning

Probability as Logic

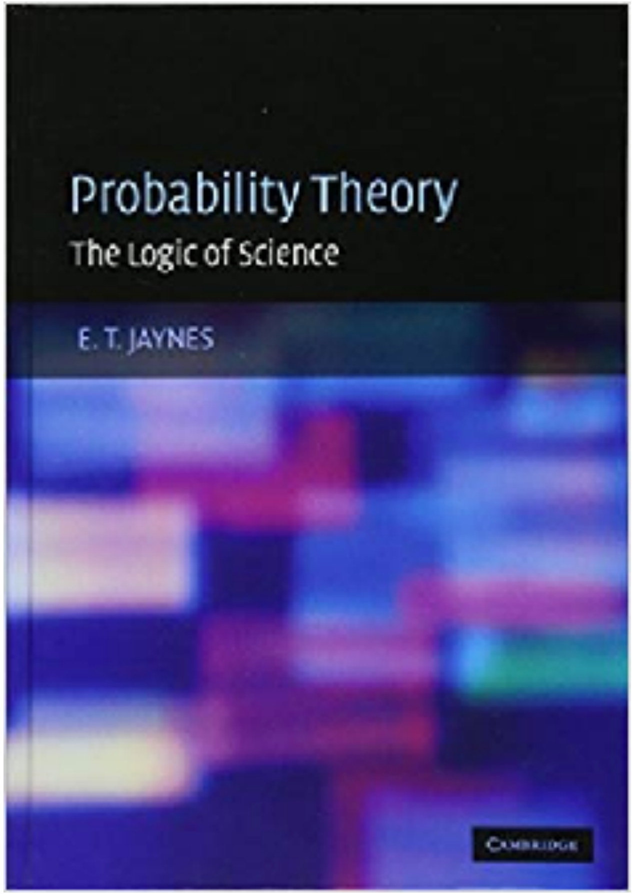

E.T. Jaynes, 2003

If you are interested, I highly recommend this book

Probability theory is nothing, but common sense reduced to calculations

Laplace, 1819

Reasoning

In terms of probabilities,
"Mary is 60% confident that the neighbor's name is Bill"
can be expressed as:

$$P(\text{name} = \text{Bill}) = .6$$

But what is the probability of any other name, say Peter?

Clearly, we cannot specify a complete probability distribution over all possible names

Carnegie Mellon University
Language Technologies Institute

Open Sets

The set of all names is open

We can get around this problem by assuming that Mary said: "I am 60% confident it is Bill, but it could be any one of 99 other names". We can further assume that all other names are equally likely.

$$P(name = X) = \begin{cases} .6, & if\ x = Bill \\ .4 * \frac{1}{99}, & if\ x \neq Bill \end{cases}$$

{Bill: .6, other-99: .4}

$$P(name = Peter)\, .4 * \frac{1}{99} = .00404$$

Reasoning

What do we mean when we say: Suzy is 80% trustworthy?

80% of the time, Suzy accurately reports what she observes

But what happens the other 20% of the time?

She draws her statements from some other distribution

In our example, she might draw a name at random from 100 names

P(Suzy says the name is Peter | Suzy is unreliable) $= .01$

Reasoning

Carnegie Mellon University
Language Technologies Institute

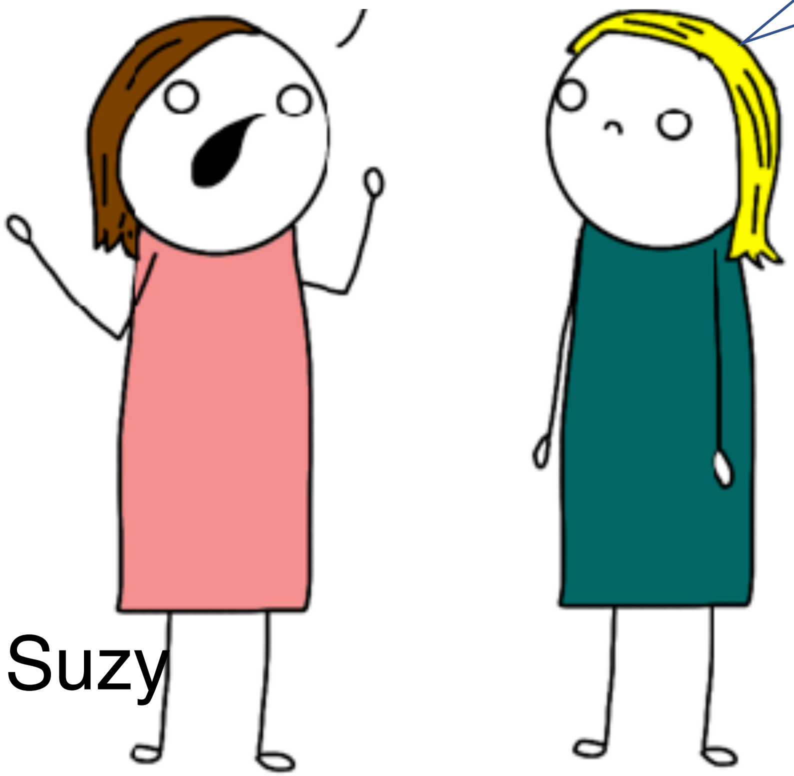

Mary was 60% confident that the name was Bill
Suzy is 80% reliable
We can now apply probability calculus (Bayes rule) to update Mary's beliefs:

Peter: .62
Bill: .23
Other: .15

What should Mary believe now?

Reasoning

The Model

$$S = \begin{cases} N, \; if \; R = True \\ \sim X, \; if \; R = False \end{cases}$$

We want to know

$$P(N \mid S = Peter)$$

$$P(N = Bill \mid S = Peter)$$

$$P(N = Peter \mid S = Peter)$$

$$P(N = other \mid S = Peter)$$

Neighbor's name

{Bill: .6, other-99: .4}

Suzy is reliable

{True: .8, False: .2}

Suzy's guess distr.

{other-100: 1.0}

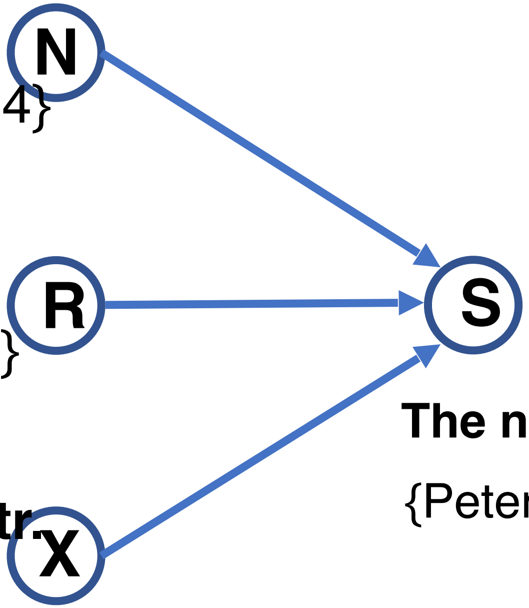

The name Suzy says

{Peter: 1.0}

Reasoning

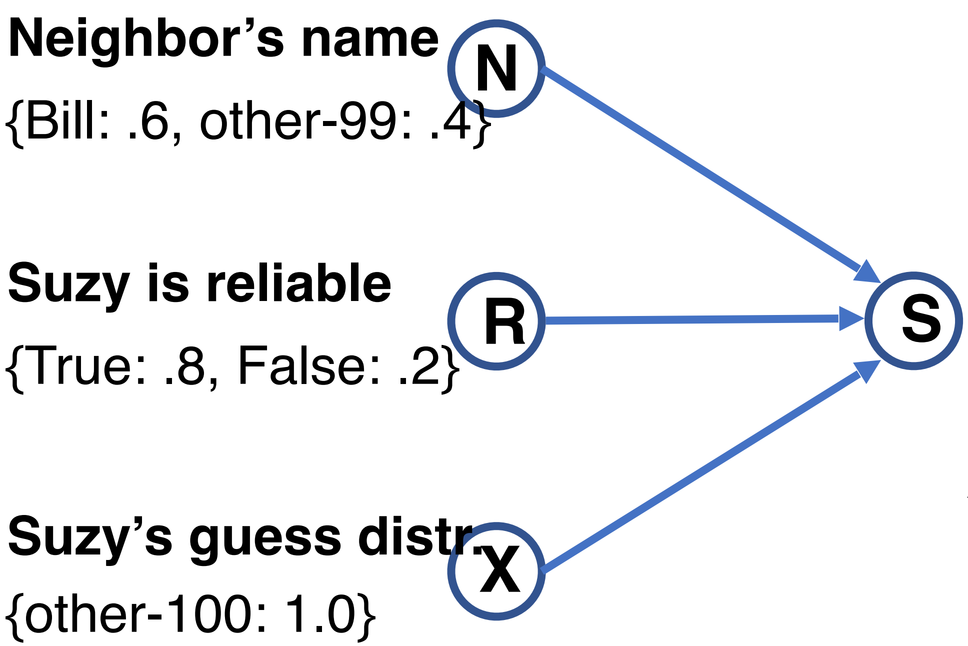

The Model

$$S = \begin{cases} N, \ if \ R = True \\ \sim X, \ if \ R = False \end{cases}$$

We want to know

$$P(N \mid S = Peter)$$

$$P(N = Bill \mid S = Peter)$$

$$P(N = Peter \mid S = Peter)$$

$$P(N = other \mid S = Peter)$$

Rules of Probability

$P(X,Y) = P(X \mid Y) * P(Y)$ or $P(X \mid Y) \propto P(X, Y)$

$P(X,Y) = P(X) * P(Y)$ When X and Y are independent

$$P(X) = P(X,Y) + P(X, \ not \ Y)$$

Reasoning

The Model

$$S = \begin{cases} N, \; if \; R = True \\ \sim X, \; if \; R = False \end{cases}$$

What is

$$P(N \mid S = Peter)$$

Neighbor's name
{Bill: .6, other-99: .4}

Suzy is reliable
{True: .8, False: .2}

Suzy's guess distr.
{other-100: 1.0}

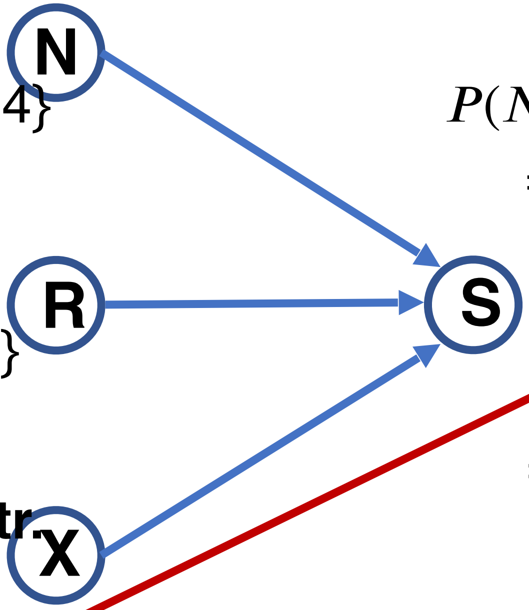

$$P(N \mid S = Peter) \propto P(N,\; S = Peter) =$$

$$= P(N,\; S = Peter,\; R = True) + P(N,\; S = Peter, R = False) =$$

$$= P\big(S = Peter \mid N,\; R = True\big) * P(N) * P(R = True) +$$

$$+ P(S = Peter \mid R = False) * P(N) * P(R = False) =$$

$$= P(N) * [P(S = Peter \mid N, R = True) * P(R = True) +$$

$$+ P(S = Peter \mid R = False) * P(R = False)] =$$

$$= P(N) * \big[P(S = Peter \mid N, R = True\big) * .8 + .01 * .2]$$

When R = False, Suzy ignores N

R and N are independent

Carnegie Mellon University
Language Technologies Institute

Reasoning

Neighbor's name

{Bill: .6, other-99: .4}

Suzy is reliable

{True: .8, False: .2}

Suzy's guess distr.

{other-100: 1.0}

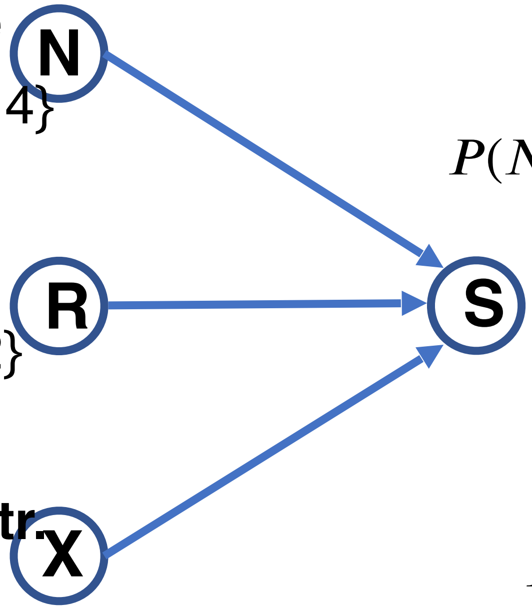

The Model

$$S = \begin{cases} N, \ if \ R = True \\ \sim X, \ if \ R = False \end{cases}$$

What is

$$P(N = Bill \mid S = Peter)$$

$$P(N = Bill) * \left[P(S = Peter \mid N = Bill, R = True\right) * .8 + \ .002] =$$

$$= .6 \ * (0 \ * \ .8 + \ .002) = \ .0012$$

This is 0

$$P(N = Bill \mid S = Peter) \propto .0012$$

Reasoning

Neighbor's name
{Bill: .6, other-99: .4}

Suzy is reliable
{True: .8, False: .2}

Suzy's guess distr
{other-100: 1.0}

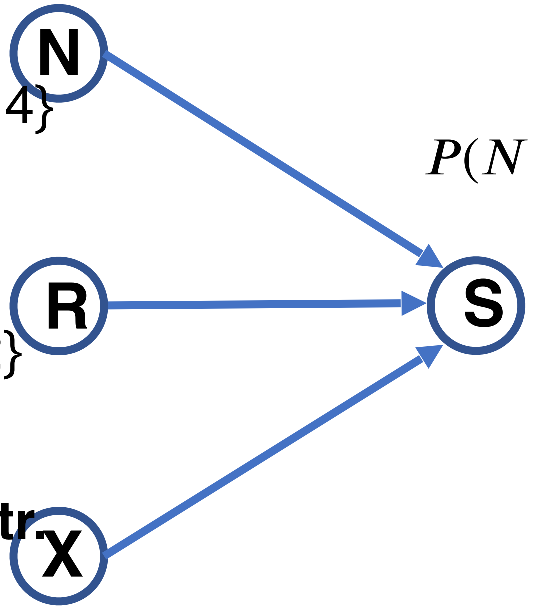

The Model

$$S = \begin{cases} N, \ if \ R = True \\ \sim X, \ if \ R = False \end{cases}$$

What is

$$P(N = Peter \mid S = Peter)$$

$$P(N = Peter) * \left[P(S = Peter \mid N = Peter, R = True)\right) * .8 + \ .002] =$$

This is 1

$$= .4 * 1/99 \ * (1 \ * \ .8 + \ .002) = \ .0032404$$

$$P(N = Peter \mid S = Peter) \propto \ .0032404$$

Reasoning

Neighbor's name
{Bill: .6, other-99: .4}

Suzy is reliable
{True: .8, False: .2}

Suzy's guess distr.
{other-100: 1.0}

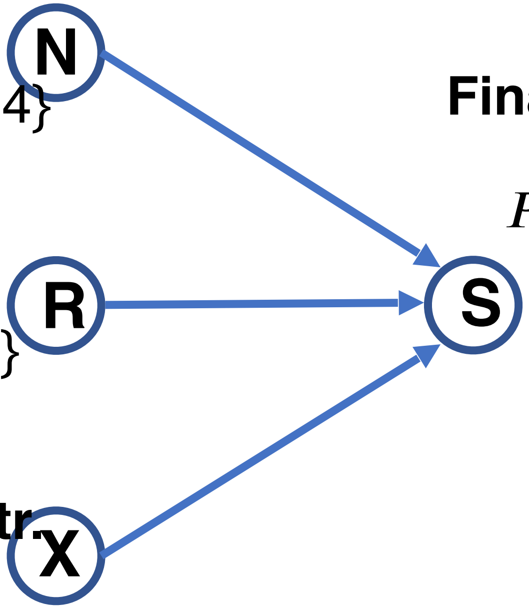

The Model

$$S = \begin{cases} N, \; if \; R = True \\ \sim X, \; if \; R = False \end{cases}$$

What is

$$P(N = Z \mid S = Peter)$$

Finally, Z is neither Bill nor Peter

$$P(N = Z) * \left[P(S = Peter \mid N = Z, R = True \right) * .8 + \; .002] =$$

This is 0

$$= .4 * 98/99 * (0 * \; .8 + \; .002) = \; .0007919$$

$$P(N = not \; Bill \; or \; Peter \mid S = Peter) \propto \; .0007919$$

Carnegie Mellon University
Language Technologies Institute

Reasoning

Neighbor’s name
{Bill: .6, other-99: .4}

Suzy is reliable
{True: .8, False: .2}

Suzy’s guess distr
{other-100: 1.0}

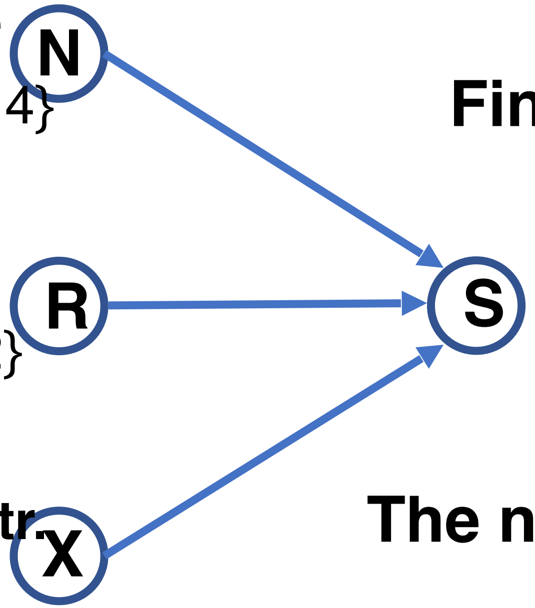

The Model

$$S = \begin{cases} N, \ if \ R = True \\ \sim X, \ if \ R = False \end{cases}$$

What is

$$P(N \mid S = Peter)$$

Finally, we can compute the posteriors

$$P(N = Bill \mid S = Peter) \propto .0012$$

$$P(N = Peter \mid S = Peter) \propto .0032$$

$$P(N = not \ Bill \ or \ Peter \mid S = Peter) \propto .0008$$

The normalizing factor is .0052 and the posteriors are:

Peter: .62
Bill: .23
Other: .15

Reasoning

Carnegie Mellon University
Language Technologies Institute

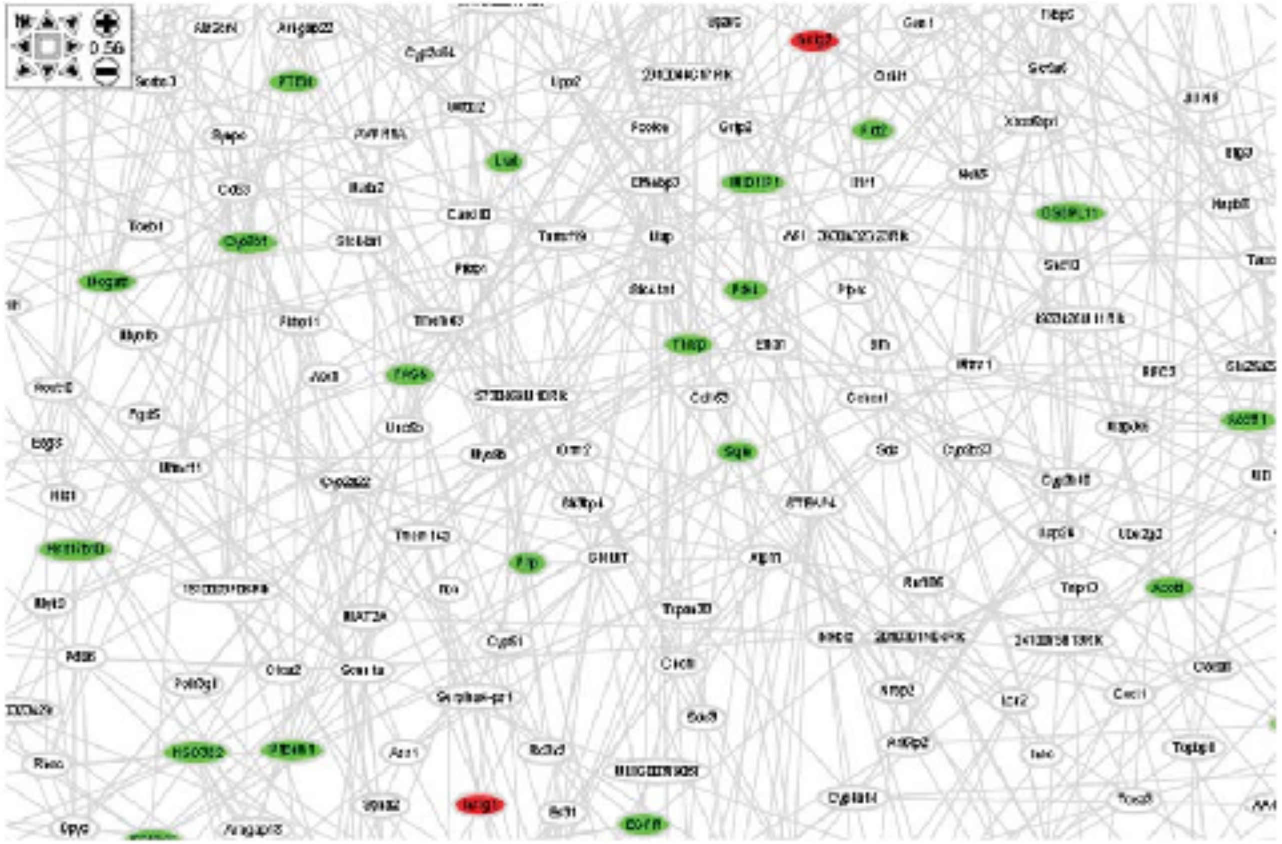

The network representing gene expression signatures

- **The real models can be very complex**
- **We use approximate methods to make large-scale inference practical**
- **There are many probabilistic reasoning tools**
- **We will discuss them later in the course**

No Model, no Reasoning!

Reasoning

The "it's all in the data" fallacy

Date	Wet Grass
1-10-2018	yes
1-21-2018	no
2-15-2018	yes
2-17-2018	no
3-05-2018	yes
3-12-2018	yes
…	…

Will the grass be wet tomorrow?

- **Without a model, the data is "dumb"**
- **A qualitative or partially-specified model can be fine-tuned to fit the data – this is called Machine Learning**
- **All ML classifiers assume a models whose parameters are learned from the data**
- **Models are based on the assumptions about causality**

Only a model gives meaning to the data

Cognitive Tasks

Carnegie Mellon University
Language Technologies Institute

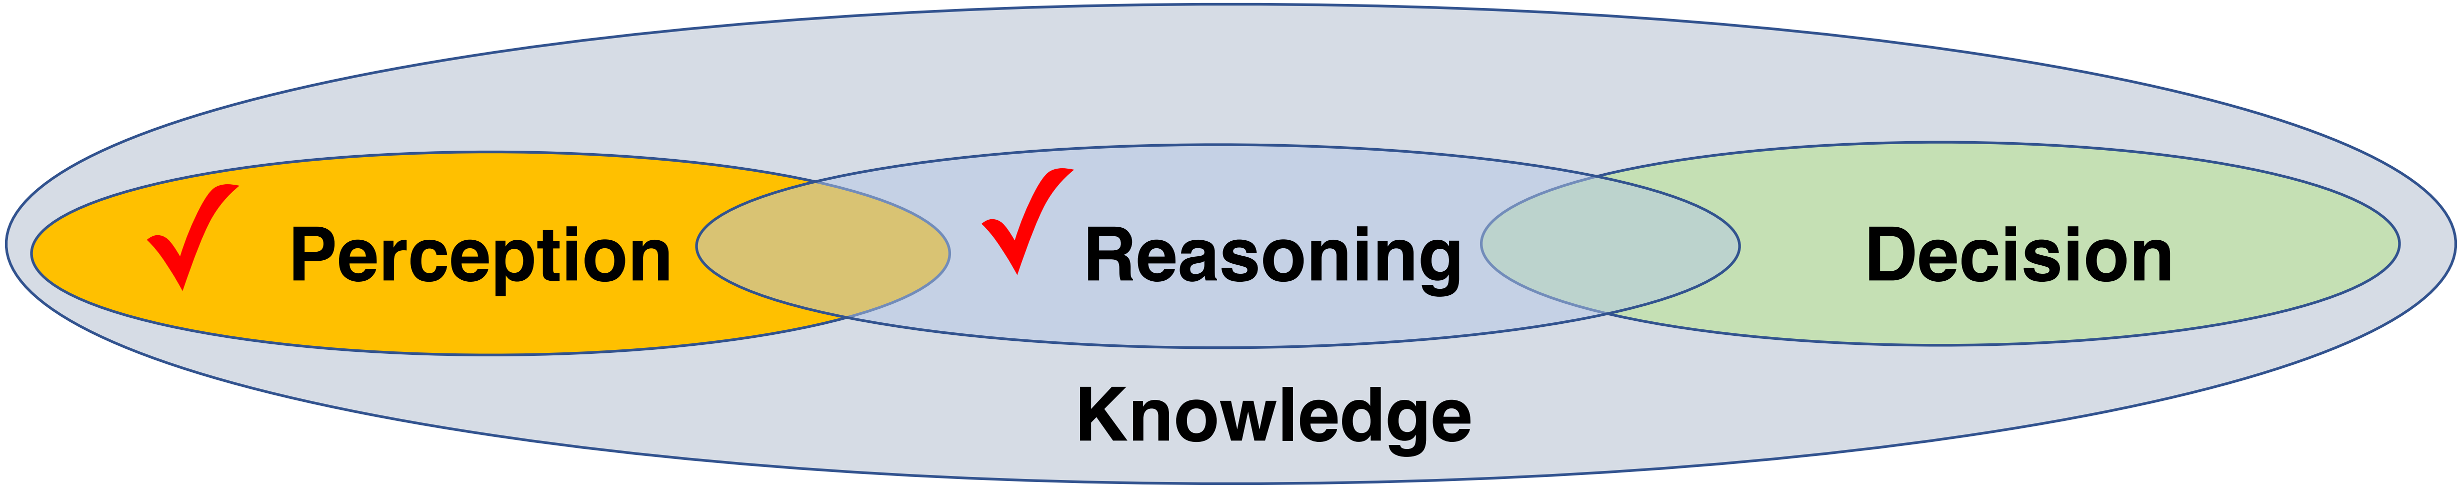

What do I see or hear? What does it mean? What do I do?

- **Perception and integration of perceived information require reasoning**
- **Reasoning is application of logic to a model – reasoning is impossible without model**
- **Probability theory is an extension of logic**