CS340: Machine Learning

- URL: www.ugrad.cs.ubc.ca/~cs340

Instructors

This week only

Nando de Freitas

Rest of class:

Kevin Murphy

TAs

- TAs:
  - Hao (Victor) Ren
  - Erik Zawadzki
- Discussion section (optional, but recommended - the TAs will go over homework problems, etc.)
  - T1A, 3:00 - 4:00pm Thursdays, DMP101
  - T1B, 8:30 - 9:30am Tuesdays, DMP201
- Office hours
  - Wed 3-4pm, CS 187

Textbook

- Required textbook (to arrive in UBC bookstore Friday Sep 8th) "Introduction to machine learning", Ethem Alpaydin

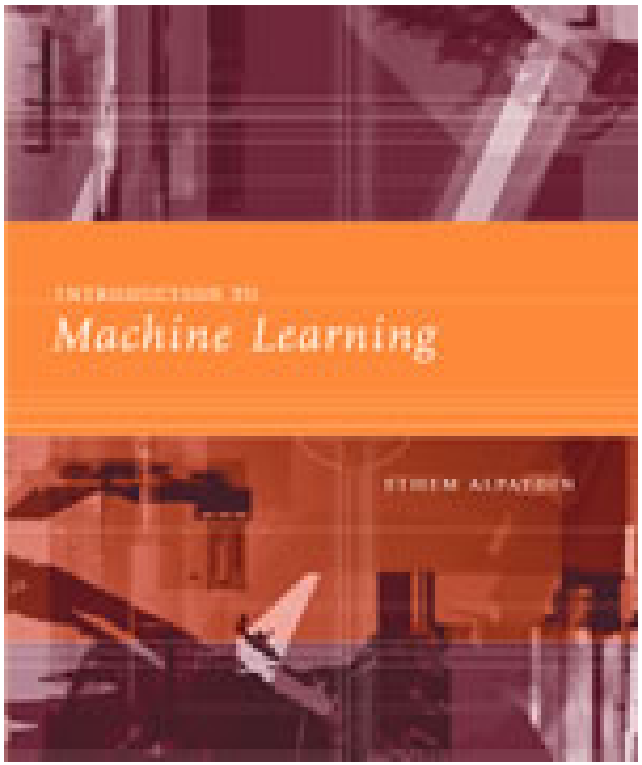

Other recommended books (more advanced)

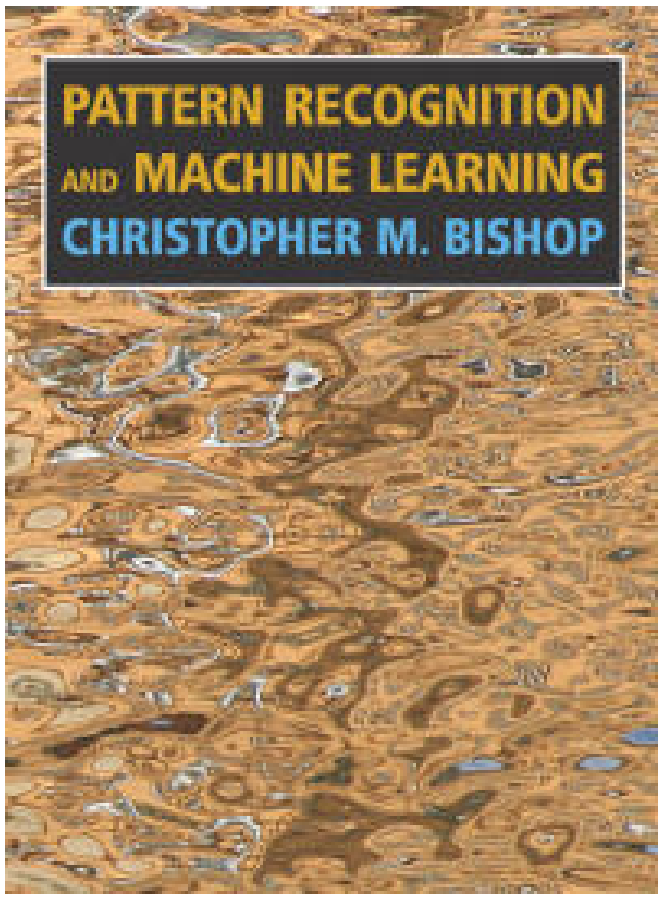

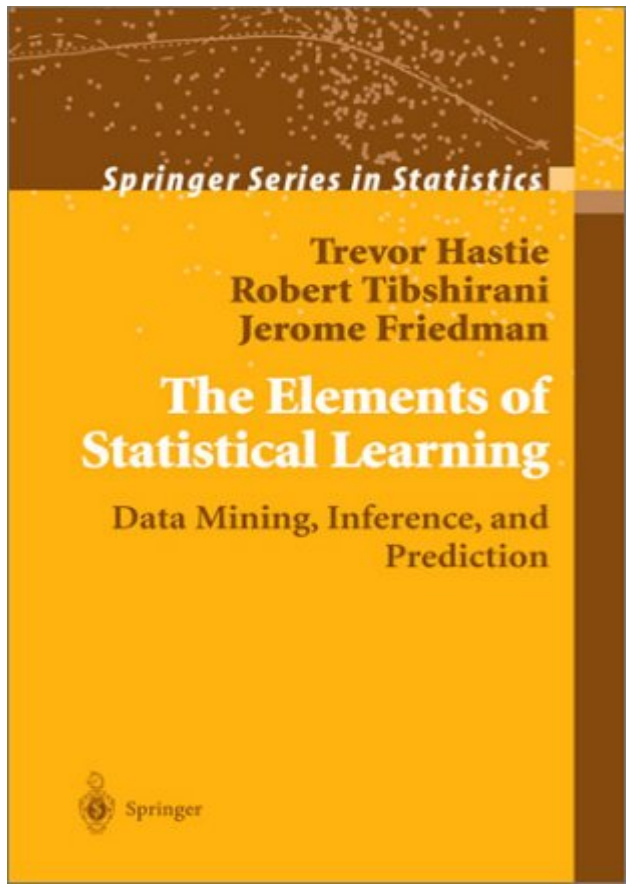

Reading

- Please read the sections of the book listed on the web page before class.
- Additional reading material will be put online; some optional, some required.
- Please keep up to date with reading!
- Lecture notes will be made available online after the class.

Grading

- Grading
  - Midterm: 30%
  - Final: 45%
  - Weekly Assignments: 25%
- Collaboration policy:
  - You can collaborate on homeworks if you write the name of your collaborators on what you hand in; however, you must understand everything you write, and be able to do it on your own (eg. in the exam!)
- Sickness policy:
  - If you cannot do an assignment or an exam, you must come see me in person; a doctor's note (or equivalent) will be required.

Pre-requisites

- You should know (or be prepared to learn)
  - Basic multivariate calculus e.g.,

$$\frac{\partial}{\partial x_j}\vec{x}^T\vec{x} = 2x_j$$

  - Basic linear algebra e.g.,

$$A\vec{u}_i = \lambda_i\vec{u}_i$$

  - Basic probability/ statistics e.g.

$$Cov(X,Y) = E\left[(X - EX)(Y - EY)\right] = E[XY] - E[X]E[Y]$$

  - Basic data structures and algorithms (e.g., trees, lists, sorting, dynamic programming, etc)

Matlab

- Everyone should have access to matlab on their CS account. If not, you can ask the TAs for a CS guest account.
- The TAs will hold a matlab tutorial session in Dmp 101.
- Various matlab tutorials on the class web-page. Best one is "Matlab for psychologists"
- The first homework is due in class on Monday 18th, and consists of some simple Matlab exercises.

What is machine learning?

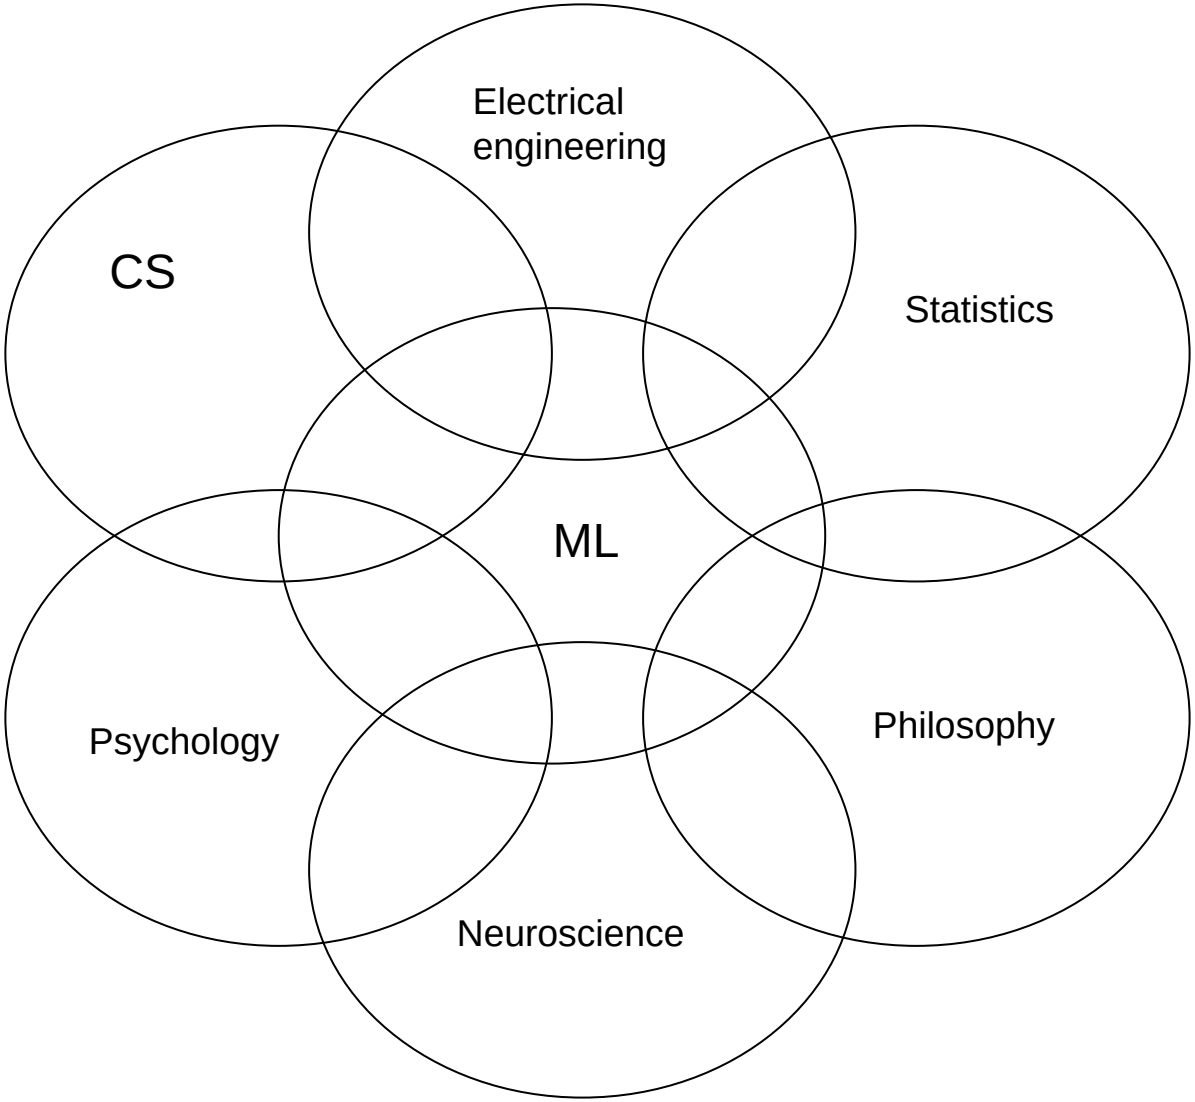

Machine Learning

Learning is the process of automatically constructing abstractions of the real world from a set of observations and past experiences

h:

d:

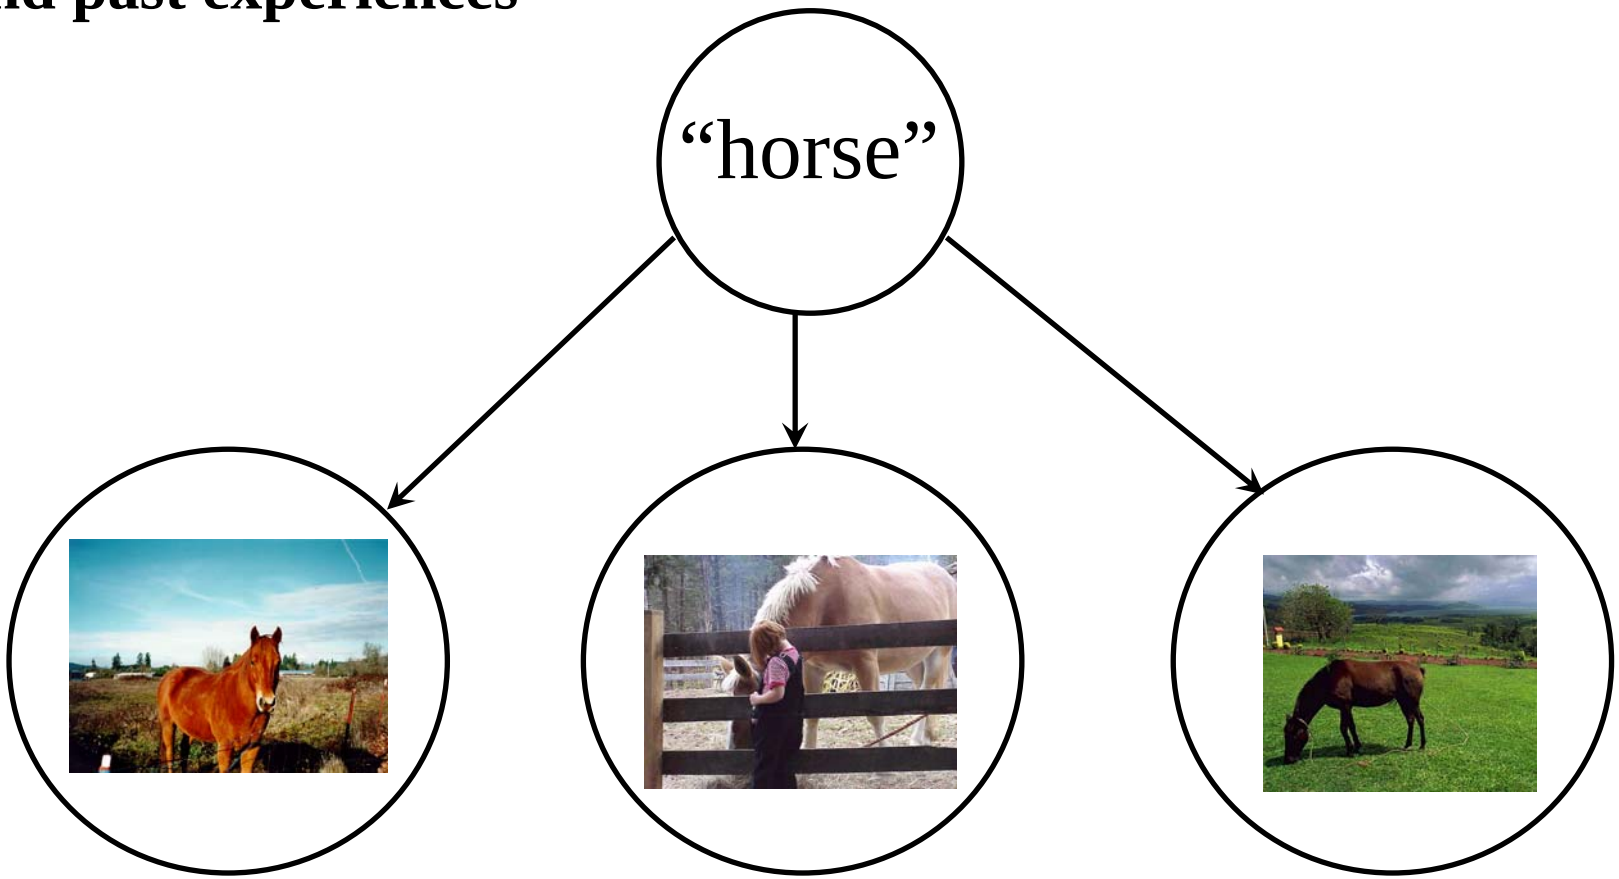

Learning concepts and words

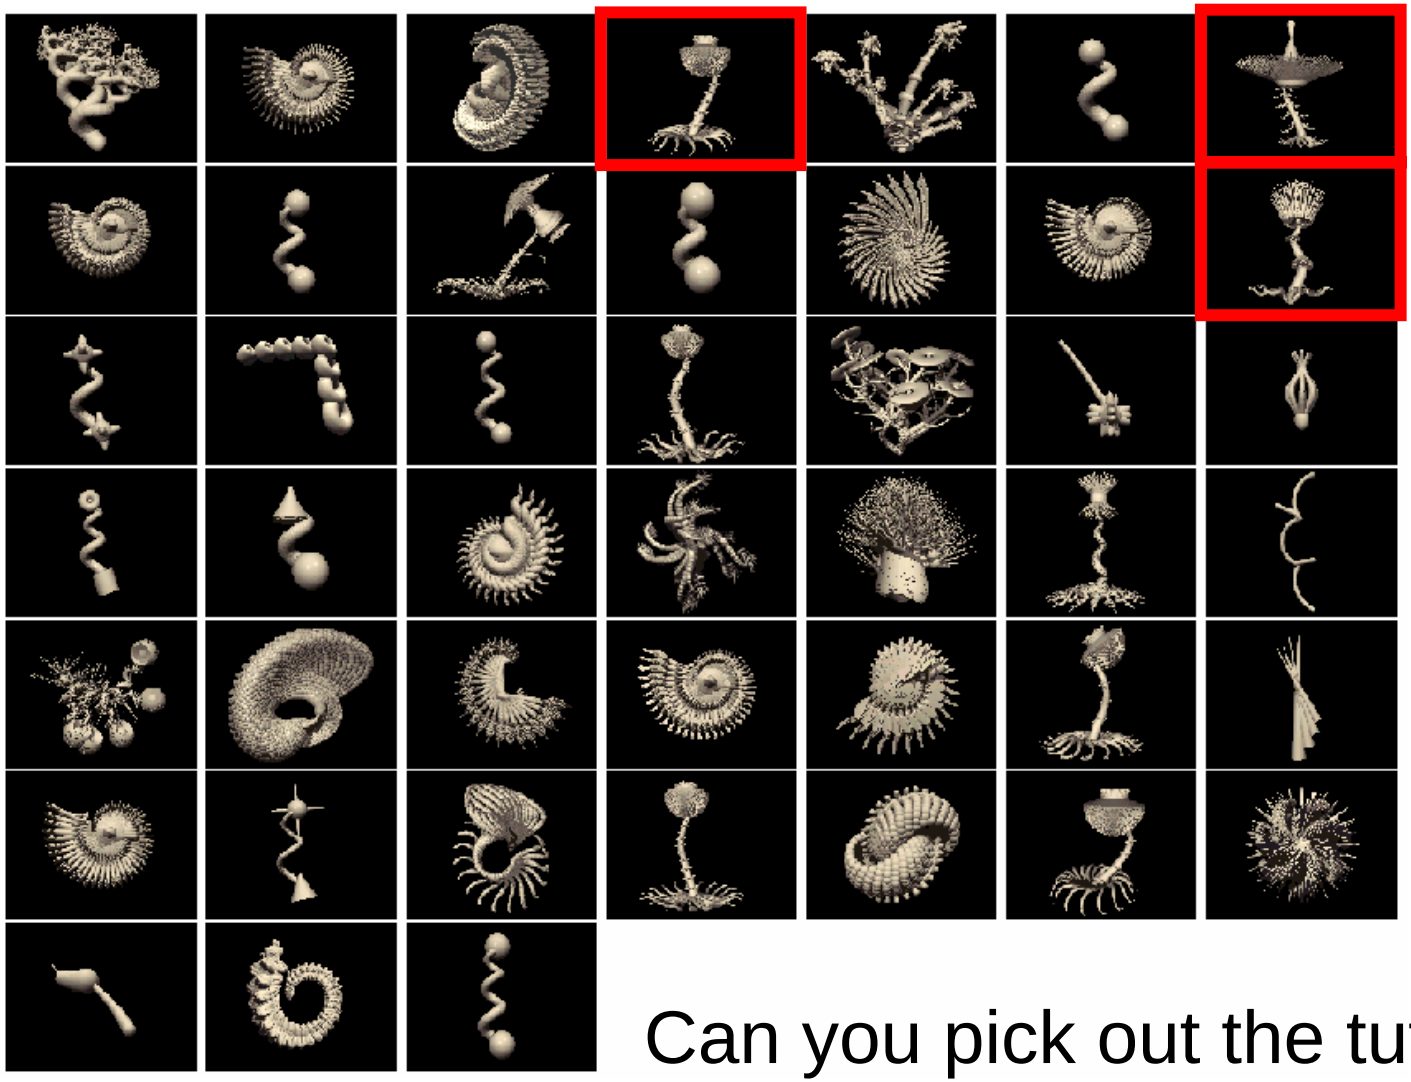

Can you pick out the tufas?

Information theory perspective

Data compression and transmission over a noisy channel provide some insight into the process of learning

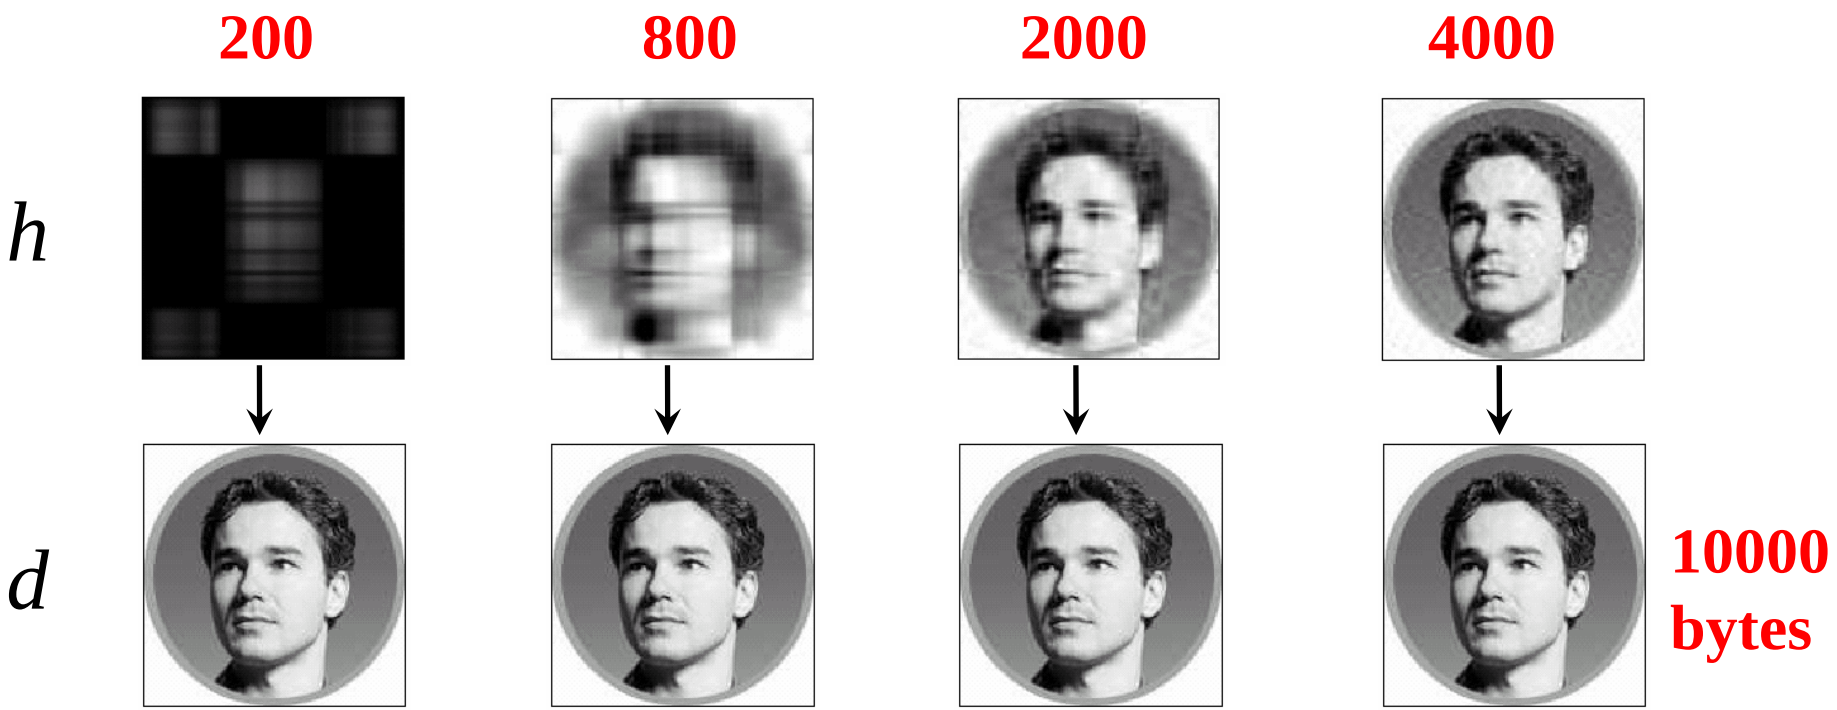

Which compressions capture the essence of the image?
Which one is best to recognize the same subject in a different photo?

Why “Learn” ?

- Machine learning is programming computers to optimize a performance criterion using example data or past experience.
- There is no need to “learn” to calculate payroll
- Learning is used when:
  - Human expertise does not exist (navigating on Mars),
  - Humans are unable to explain their expertise (speech recognition)
  - Solution changes in time (routing on a computer network)
  - Solution needs to be adapted to particular cases (user biometrics)

Perception-action cycle

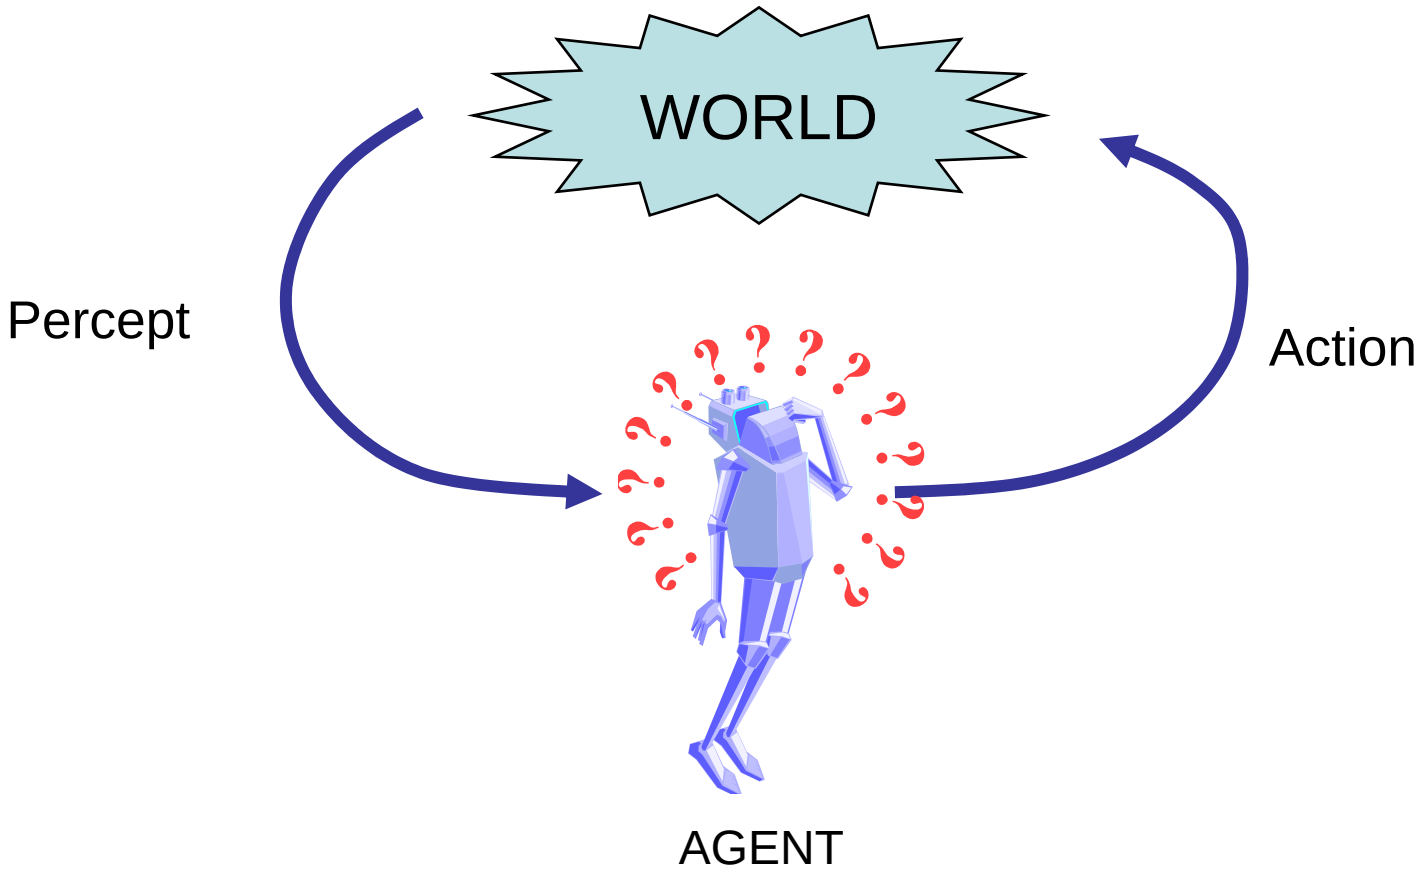

AI = designing intelligent agents
ML = designing agents that learn to be intelligent

Agents


AIBO® Entertainment Robot
Official U.S. Resources and Online Destinations
ERS-7
Entertainment Robot AIBO
ERS-7 with:
Wireless LAN
AIBO MIND software
Energy Station
AIBOne
Pink Ball
AIBO Cards (15)
WLAN Manager CD
Battery & AC Adapter
3rd Generation
Pre-order Now!

More agents

Electrolux Trilobite robot vacuum

Friendly Robotics lawn mower

Roomba from iRobot

Non-physical agents (chess)

May 11th, 1997

Computer won world champion of chess

(Deep Blue) (Garry Kasparov)

(Reuters = Kyodo News)

Non-physical agents (web-bots)

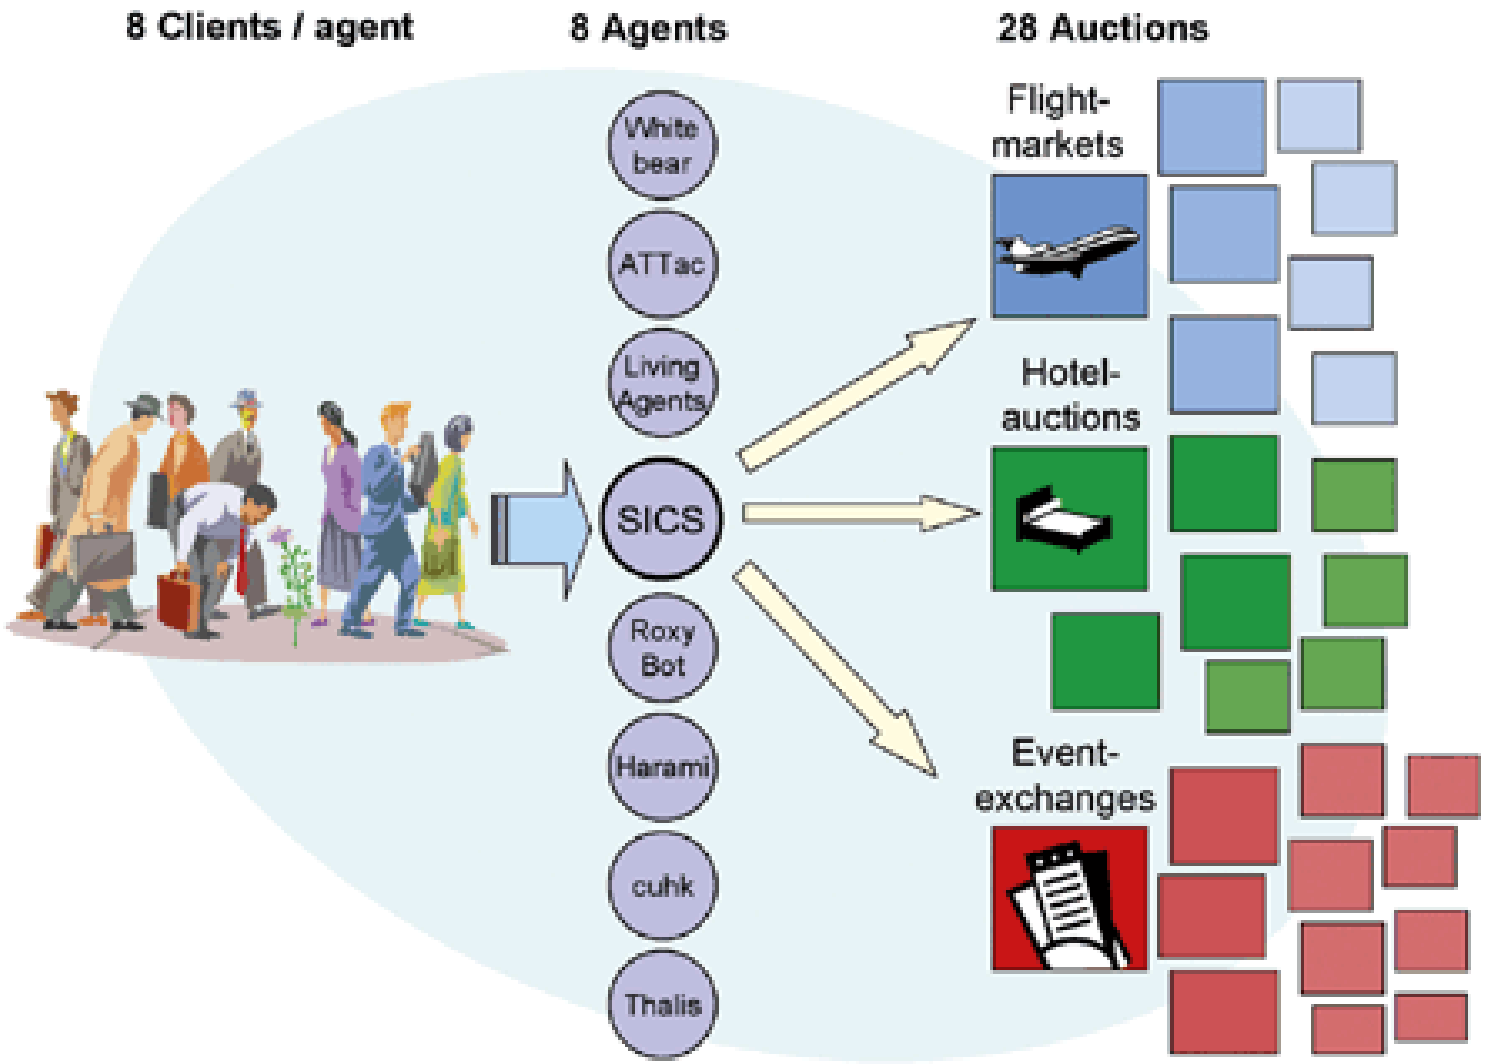

Multiple agents (robocup)

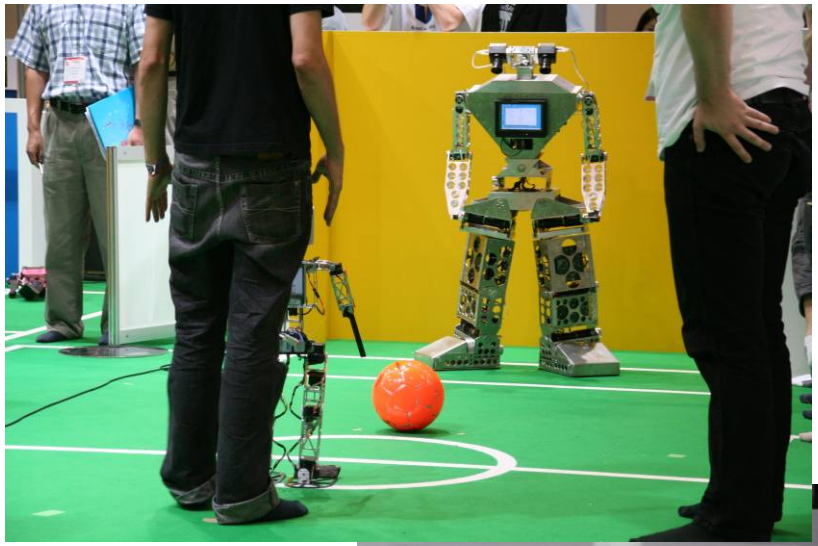

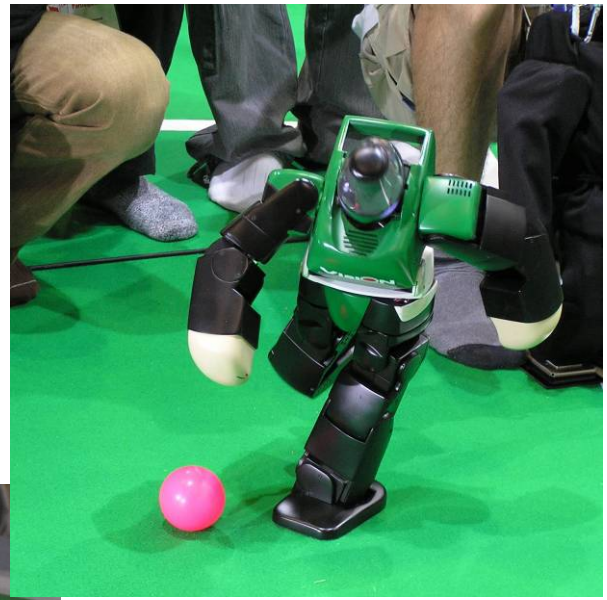

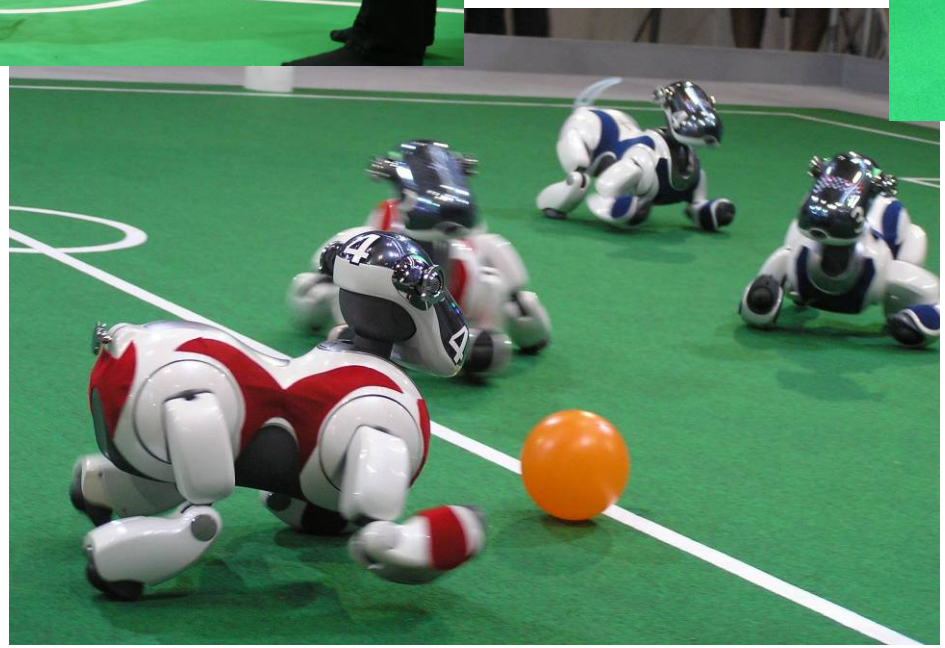

Perception

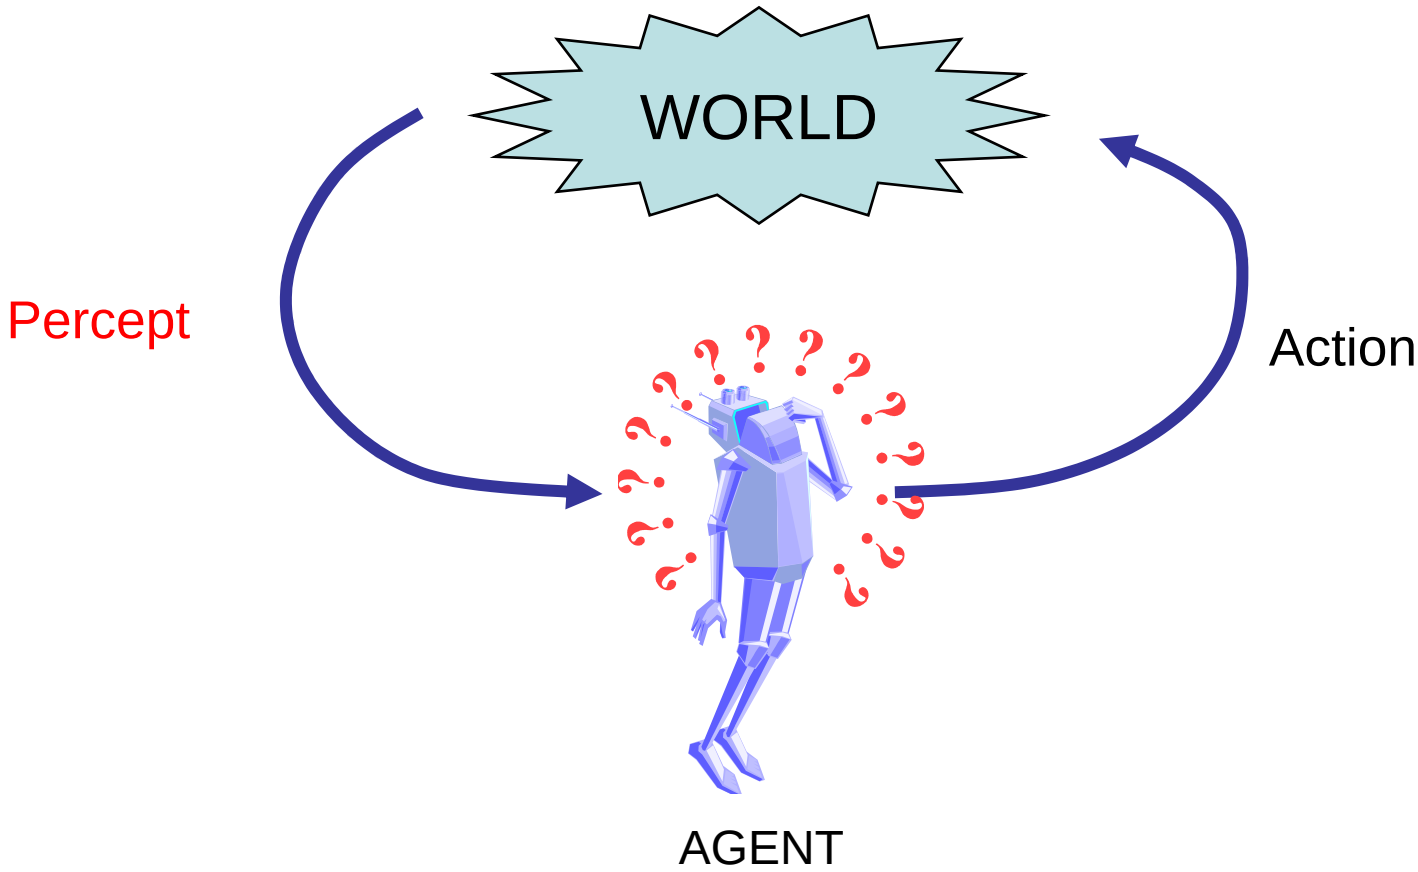

Bayesian inference perspective

Posterior probability

Observation model

Prior probability

$$p(h \mid d) = \frac{p(d \mid h)\, p(h)}{\sum_{h' \in H} p(d \mid h')\, p(h')}$$

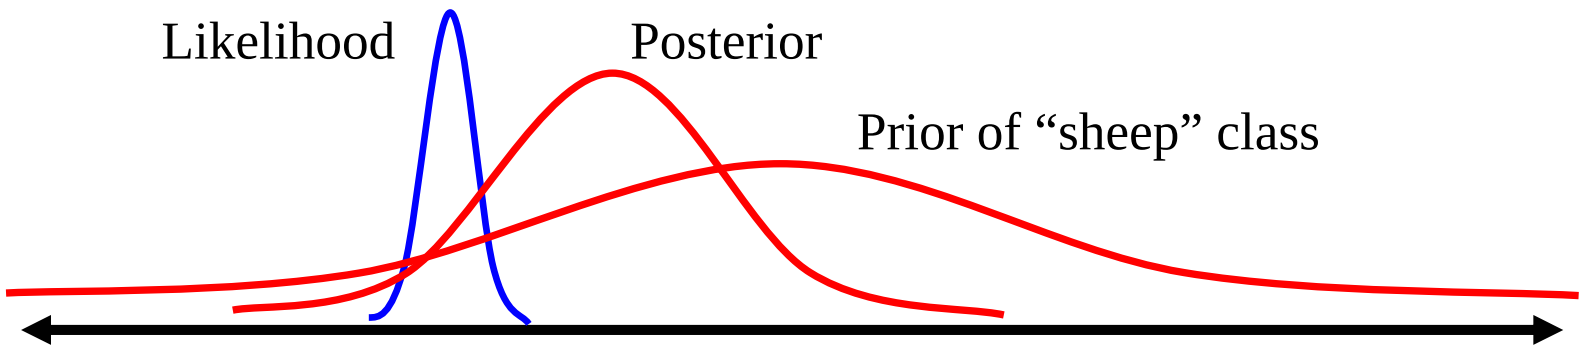

“sheep”

Vision = inverse graphics

- p(world | image) α p(image |world) x p(world)

Final beliefs

Likelihood of data

Initial beliefs

Inverse probability theory

(Bayes rule)

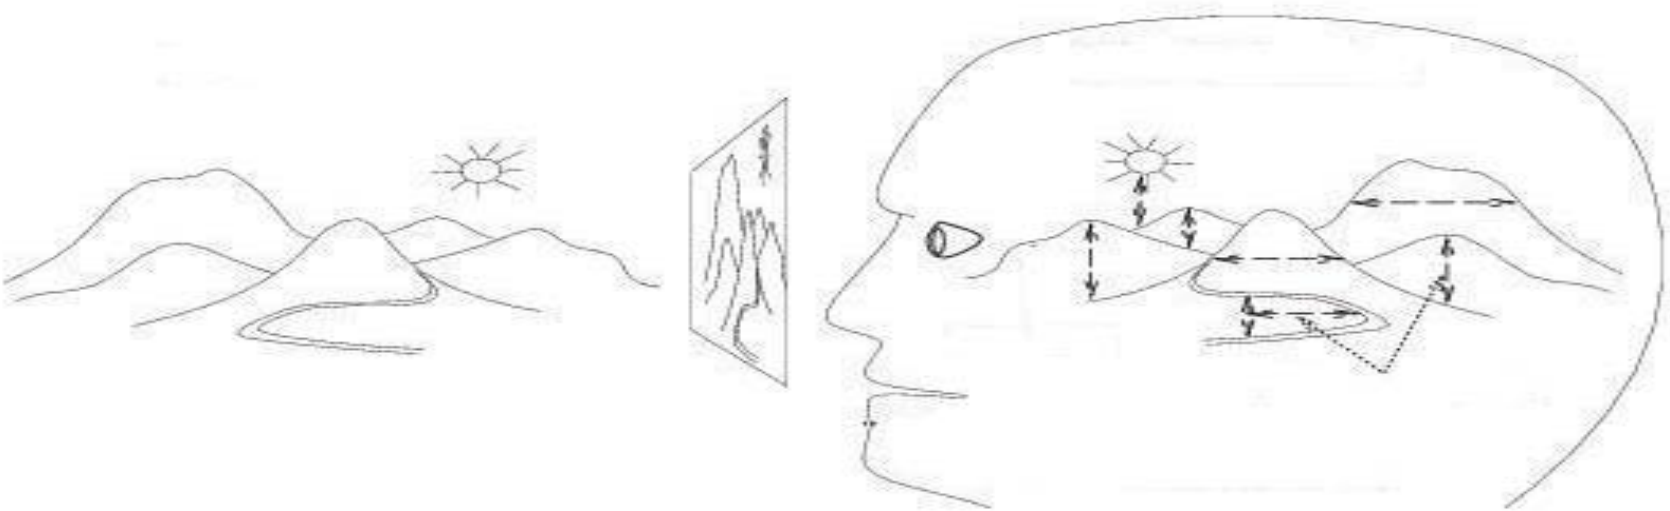

World

Image

Beliefs about world

People as Bayesian reasoners

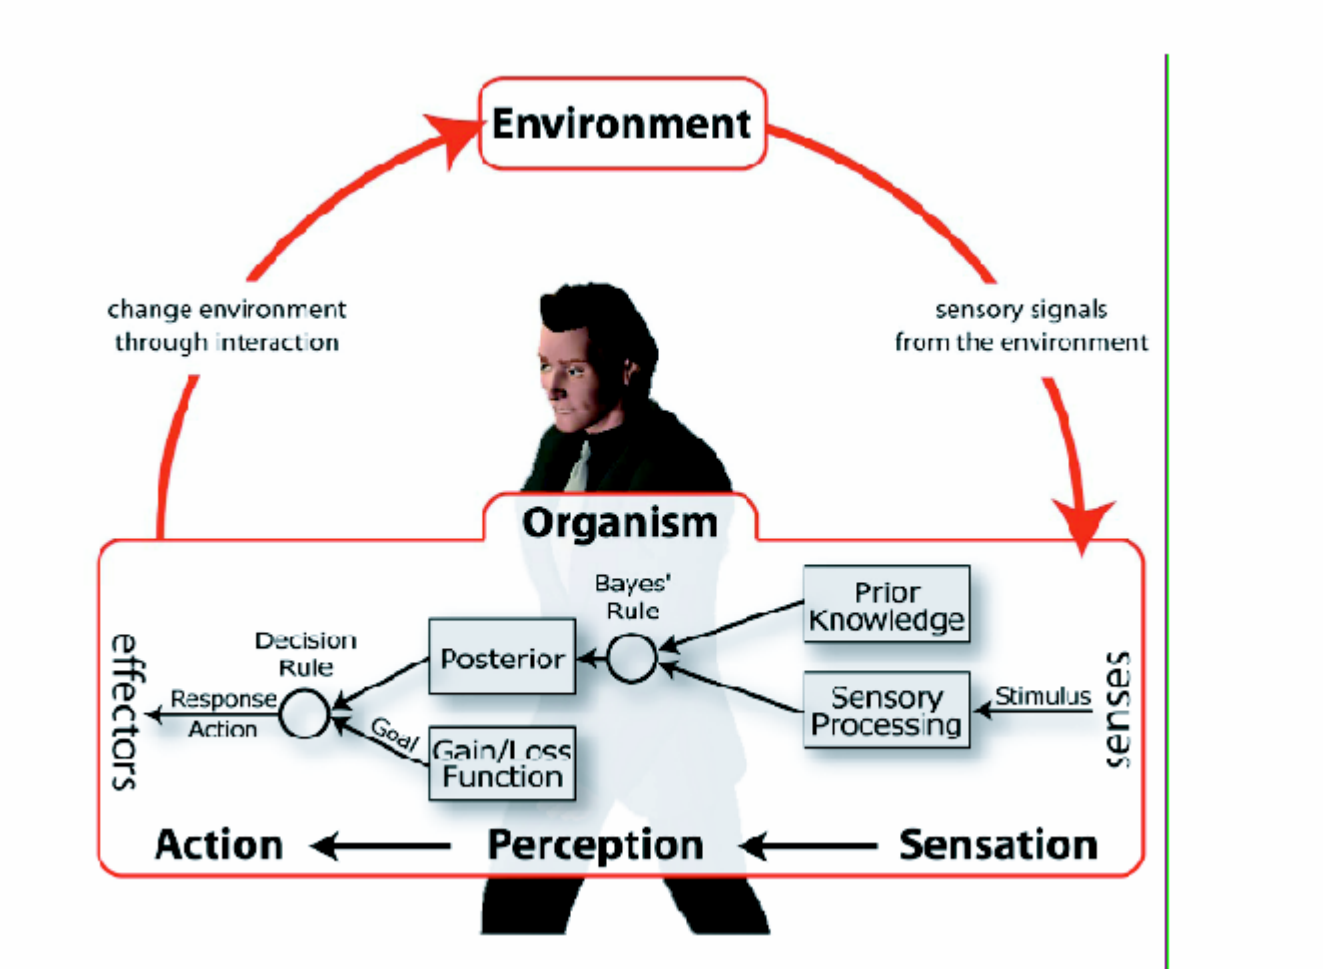

Speech recognition

- P(words | sound) α P(sound | words) P(words)

(Bayes rule)

Final beliefs

Likelihood of data
eg mixture of Gaussians

Language model
eg Markov model

Hidden Markov Model (HMM)

"Recognize speech"

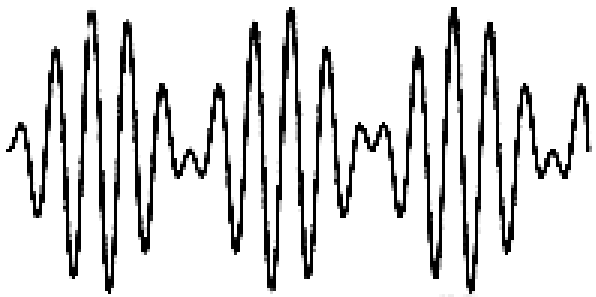

"Wreck a nice beach"

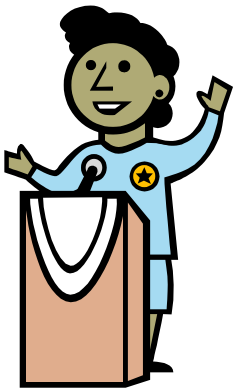

Natural language understanding

- P(meaning | words) α P(words | meaning) P(meaning)
- We do not yet know good ways to represent "meaning" (this is called the knowledge representation problem in AI)
- Current approaches involve "shallow parsing", where the meaning of a sentence can be represented by fields in a database eg
  - "Microsoft acquired AOL for $1M yesterday"
  - "Yahoo failed to avoid a hostile takeover from Google"
-

Buyer	Buyee	When	Price
MS	AOL	Yesterday	$1M
Google	Yahoo	?	?

Decision making under uncertainty

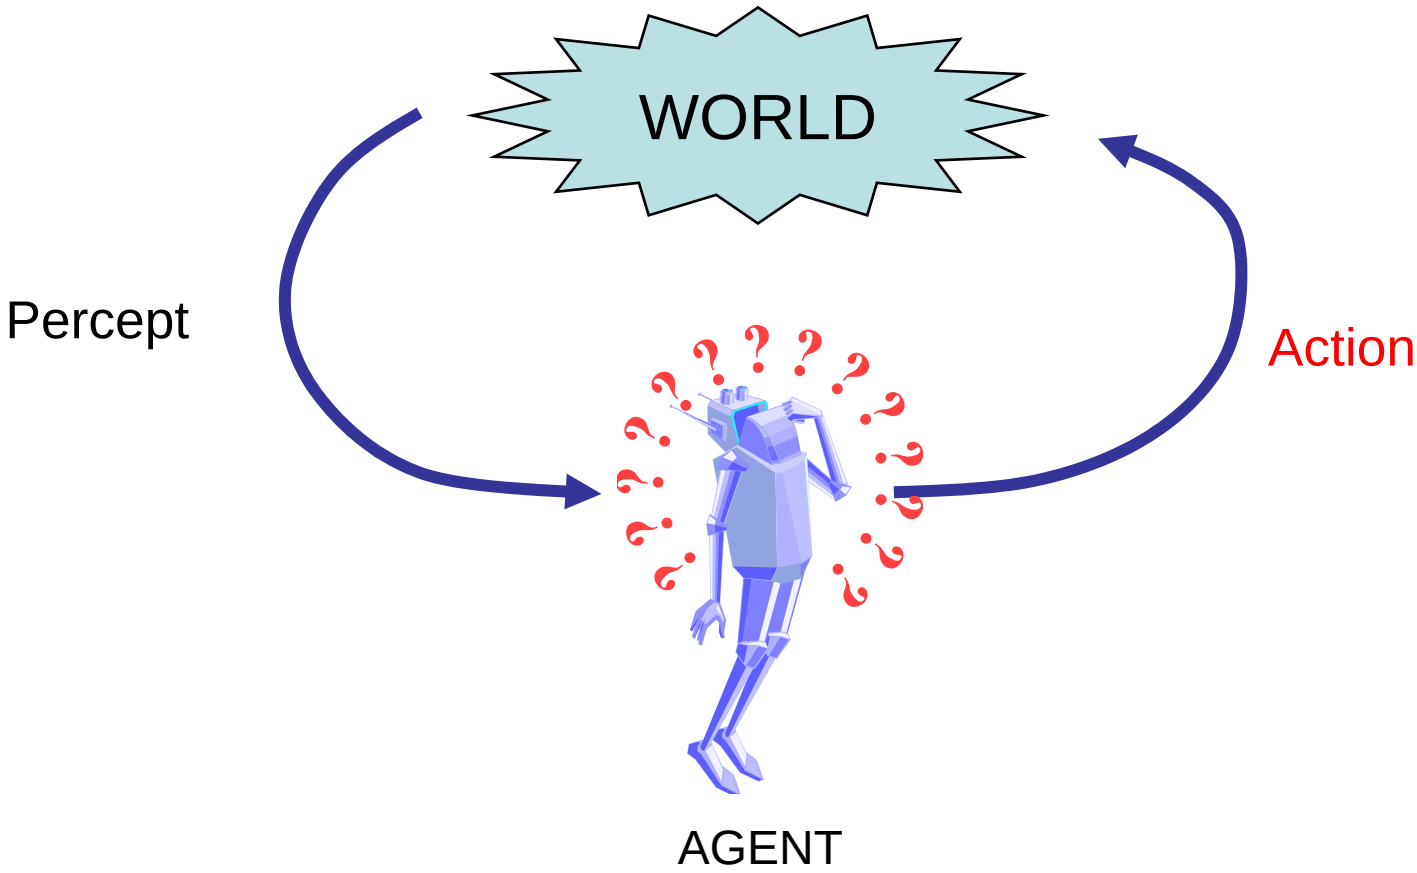

Decision theory perspective

Utilitarian view: We need models to make the right decisions under uncertainty. Inference and decision making are intertwined

Population model

$$p(\mathbf{x} = healthy) = 0.9$$

$$p(\mathbf{x} = cancer) = 0.1$$

Reward model

	$\mathbf{a} = no\ treatment$	$\mathbf{a} = treatment$
$\mathbf{x} = healthy$	0	-30
$\mathbf{x} = cancer$	-100	-20

We choose the action that maximizes the expected utility:

$$EU(\mathbf{a}) = \sum_{\mathbf{x} \in \{healthy, cancer\}} r(\mathbf{x}, \mathbf{a})\, p(\mathbf{x})$$

$$EU(\mathbf{a} = treatment) = \text{-27.2}$$

$$EU(\mathbf{a} = no\ treatment) = \text{-10}$$

Mobile robot navigation

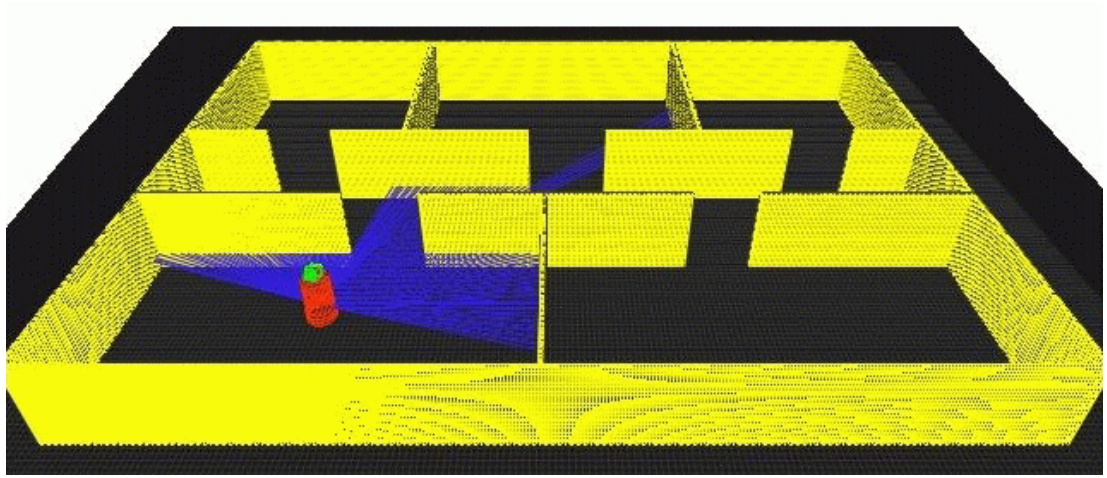

Learning how to fly

Learning how to make money

- In full 10-player games Poki is better than a typical low-limit casino player and wins consistently; however, not as good as most experts
- New programs being developed for the 2-player game are quite a bit better, and we believe they will very soon surpass all human players