The Ultimate Software
Machine Learning and Intelligence

Pedro Domingos
University of Washington

Where Does Knowledge Come From?

Evolution

Experience

Culture

Where Does Knowledge Come From?

Evolution

Experience

Culture

Computers

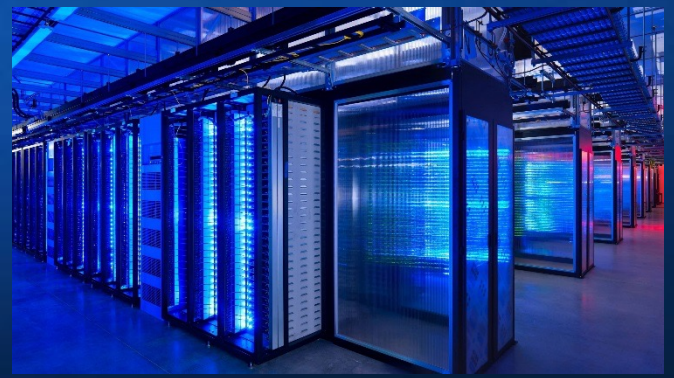

Most of the knowledge in the world in the future is going to be extracted by machines and will reside in machines.

– Yann LeCun, Director of AI Research, Facebook

Traditional Programming

Machine Learning

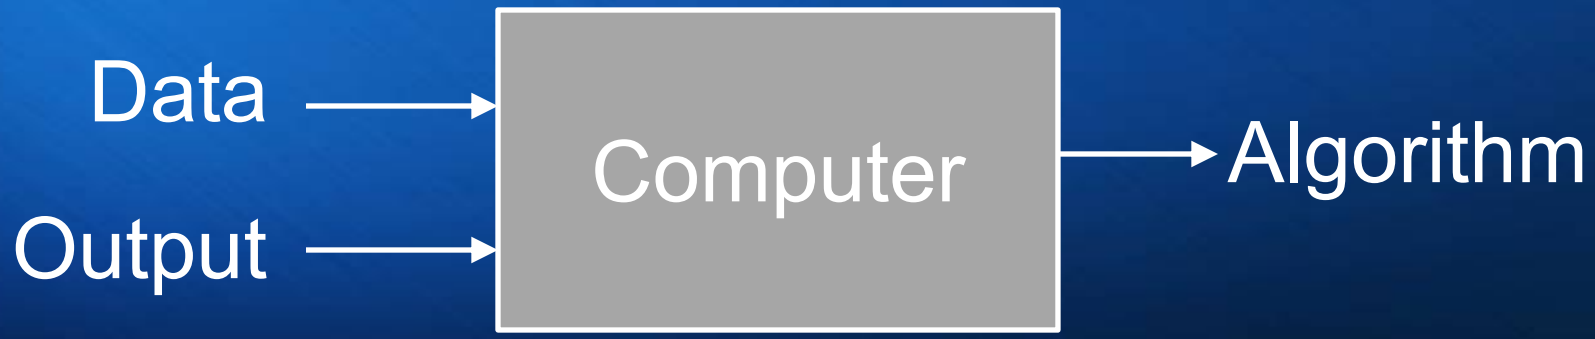

Traditional Programming

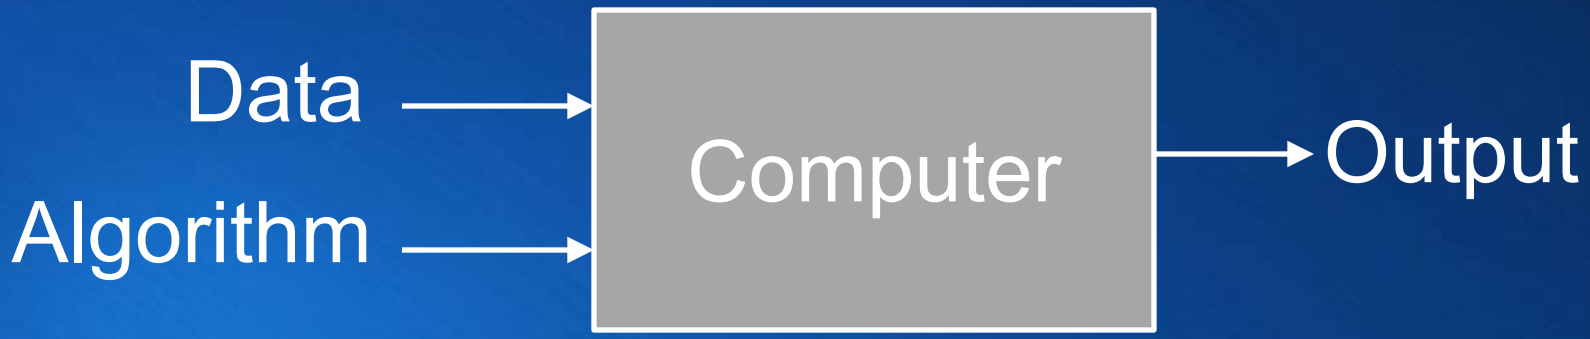

Machine Learning

So How Do Computers Discover New Knowledge?

1. Fill in gaps in existing knowledge
2. Emulate the brain
3. Simulate evolution
4. Systematically reduce uncertainty
5. Notice similarities between old and new

The Five Tribes of Machine Learning

Tribe	Origins	Master Algorithm
Symbolists	Logic, philosophy	Inverse deduction
Connectionists	Neuroscience	Backpropagation
Evolutionaries	Evolutionary biology	Genetic programming
Bayesians	Statistics	Probabilistic inference
Analogizers	Psychology	Kernel machines

Symbolists

Tom Mitchell

Steve Muggleton

Ross Quinlan

Inverse Deduction

Addition	Subtraction
2	2
+ 2	+ ?
= ?	= 4

Inverse Deduction

Deduction

Socrates is human

\+ Humans are mortal .

= ?

Induction

Socrates is human

\+ ?

= Socrates is mortal

Spot the Biologist in this Picture

Connectionists

Yann LeCun

Geoff Hinton

Yoshua Bengio

A Neuron

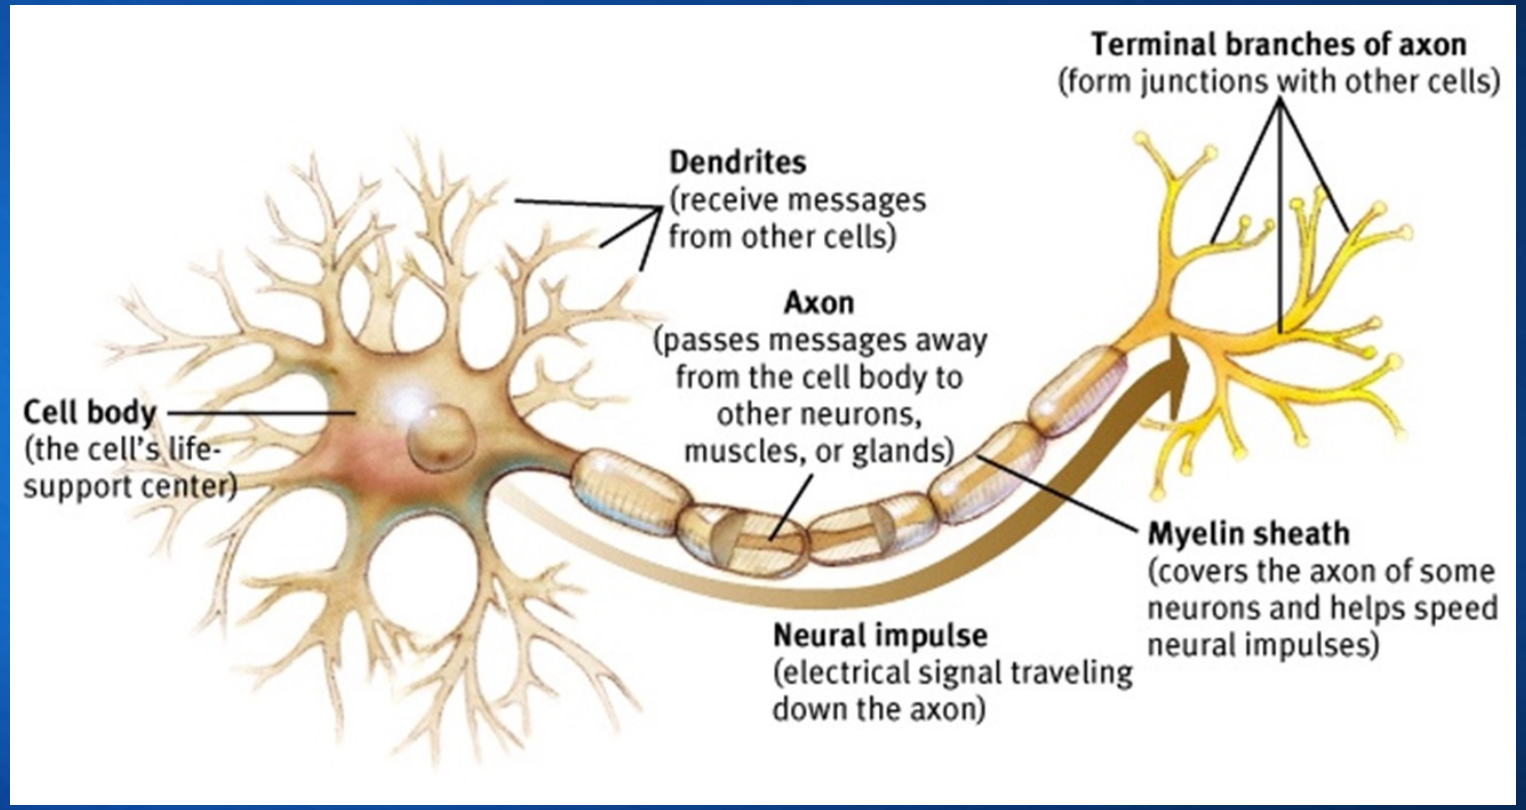
Terminal branches of axon
(form junctions with other cells)
Dendrites
(receive messages from other cells)
Axon
(passes messages away from the cell body to other neurons, muscles, or glands)
Cell body
(the cell's life-support center)
Myelin sheath
(covers the axon of some neurons and helps speed neural impulses)
Neural impulse
(electrical signal traveling down the axon)

An Artificial Neuron

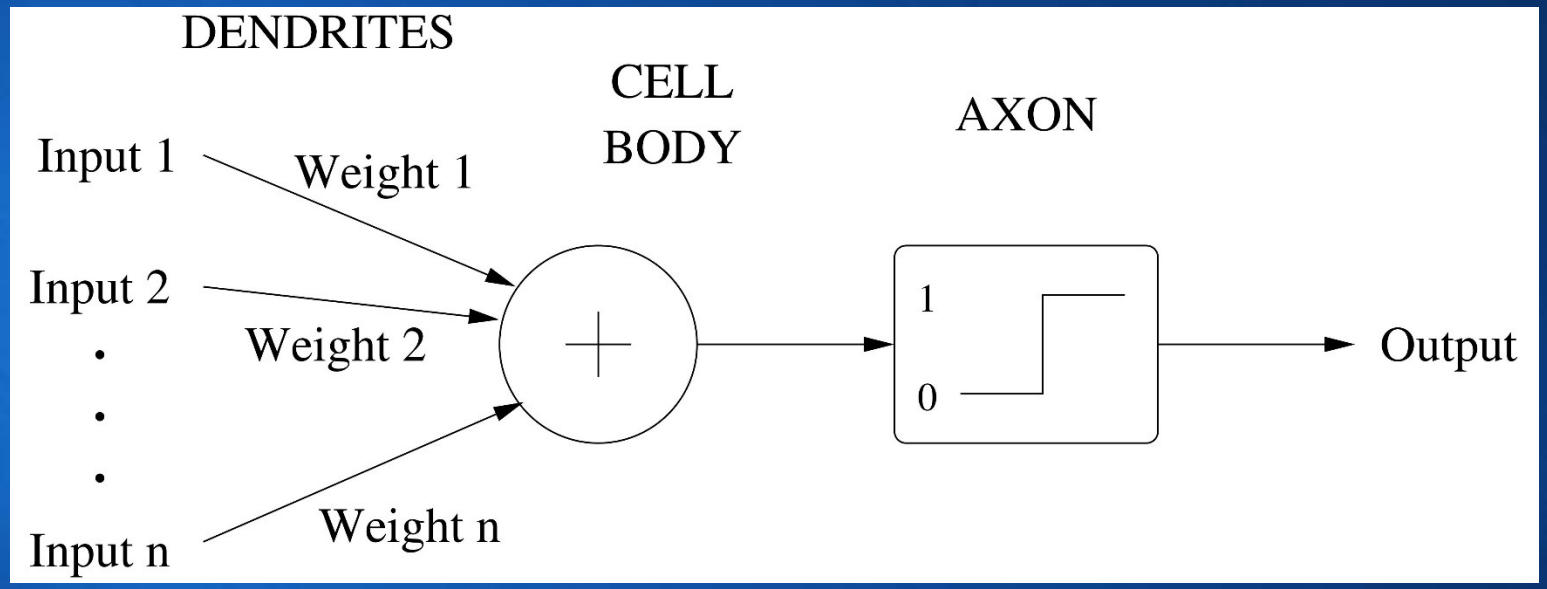

Backpropagation

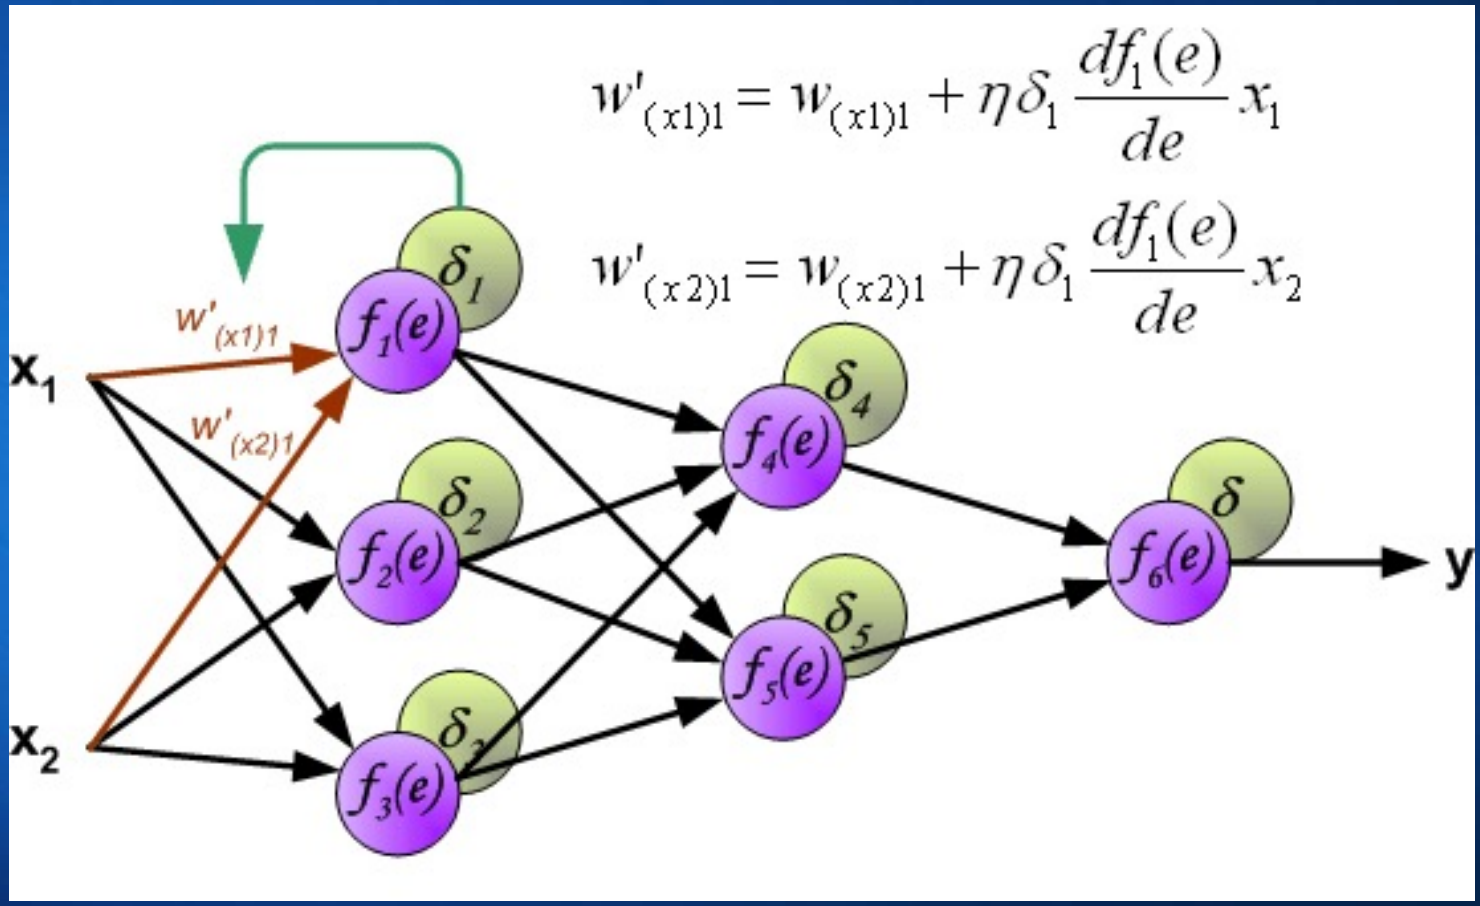

The Google Cat Network

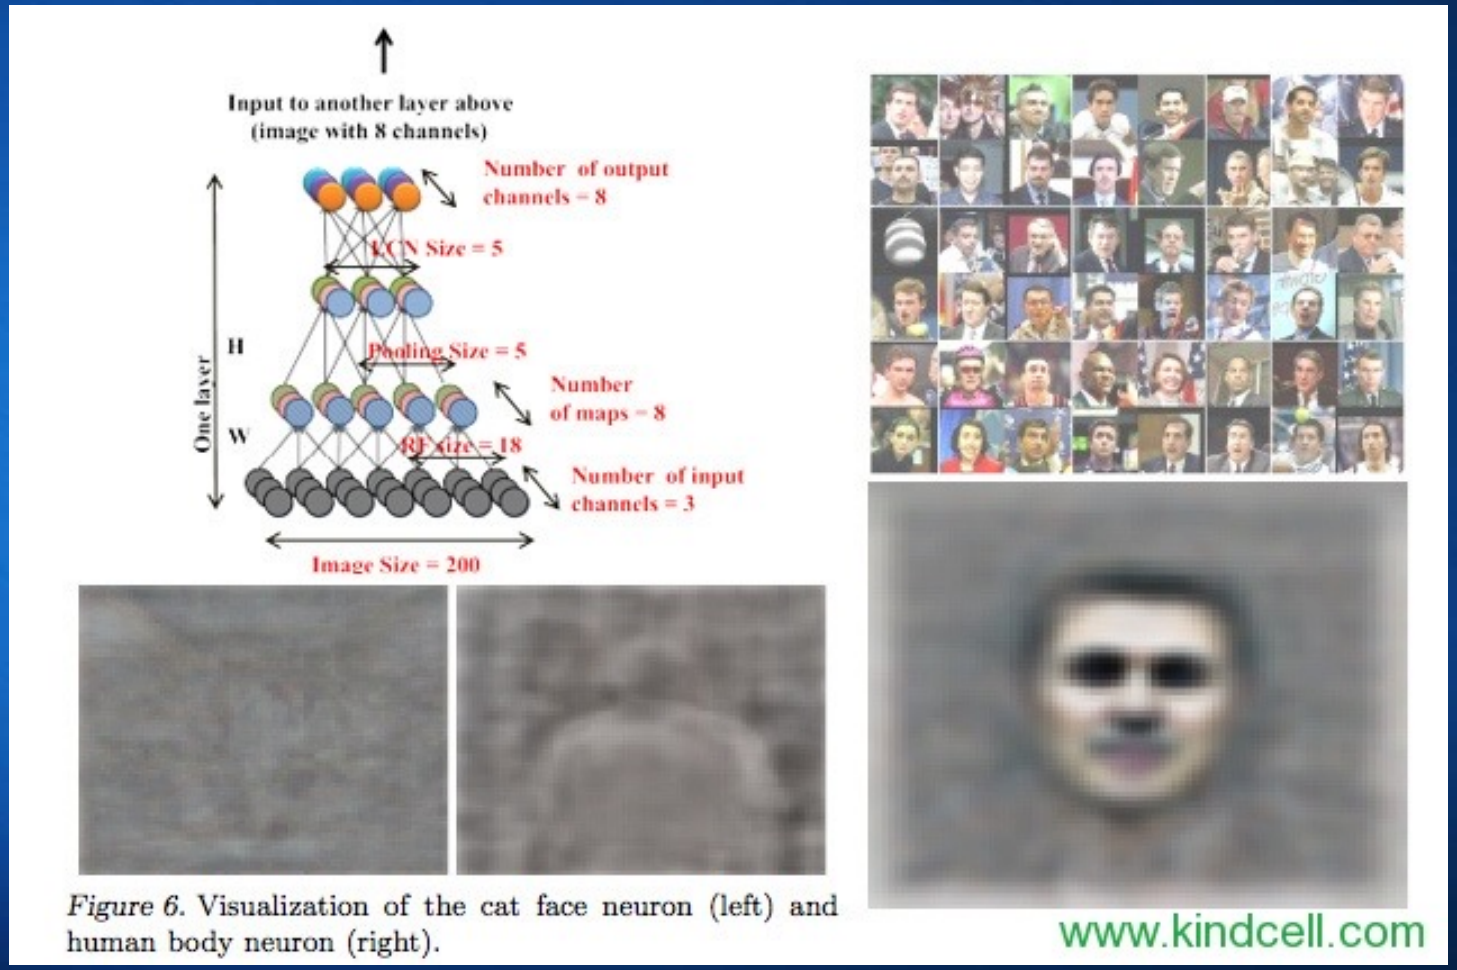

Figure 6. Visualization of the cat face neuron (left) and human body neuron (right).

Evolutionaries

John Koza

John Holland

Hod Lipson

Genetic Algorithms

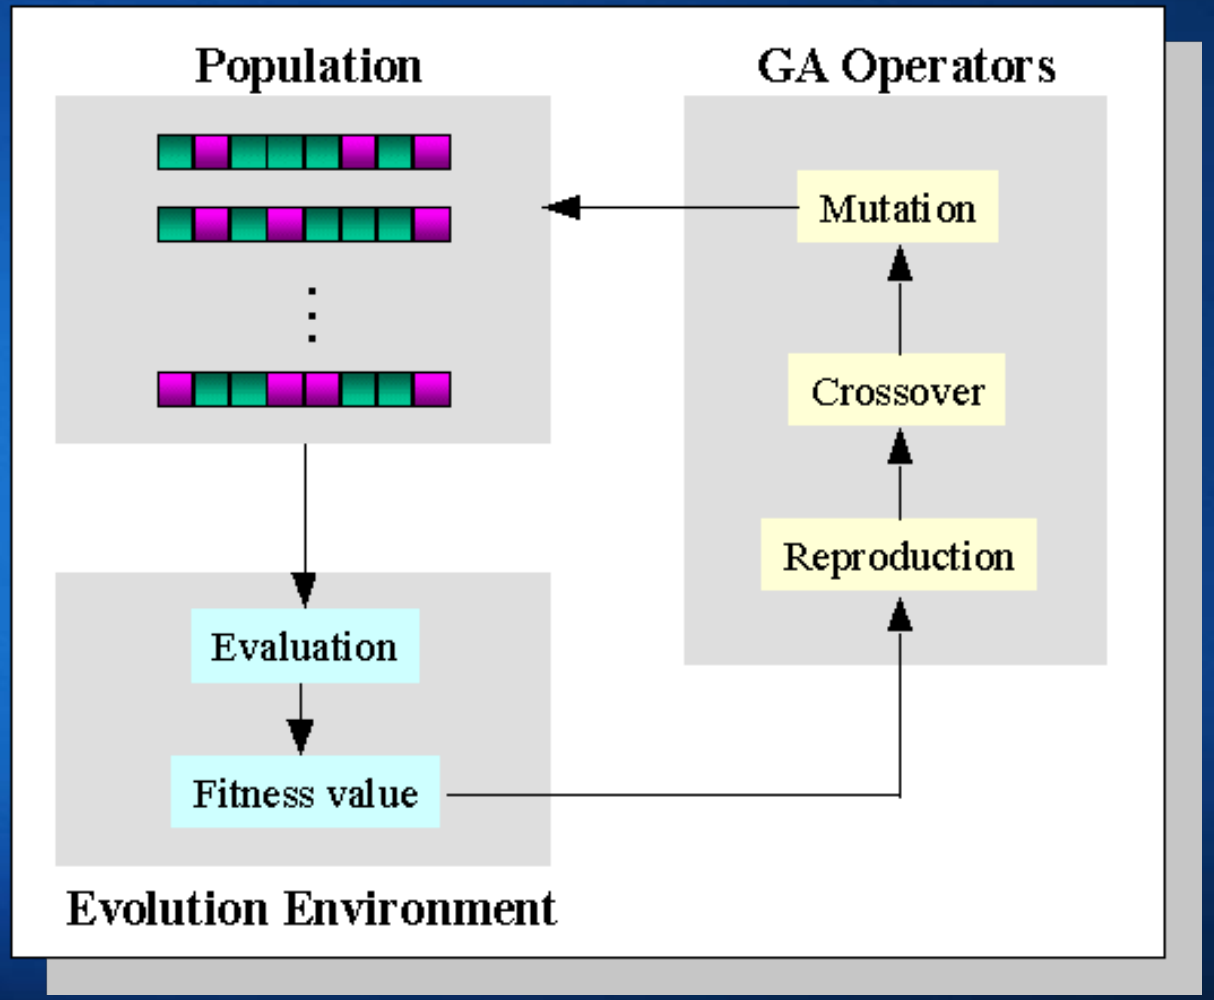

Genetic Programming

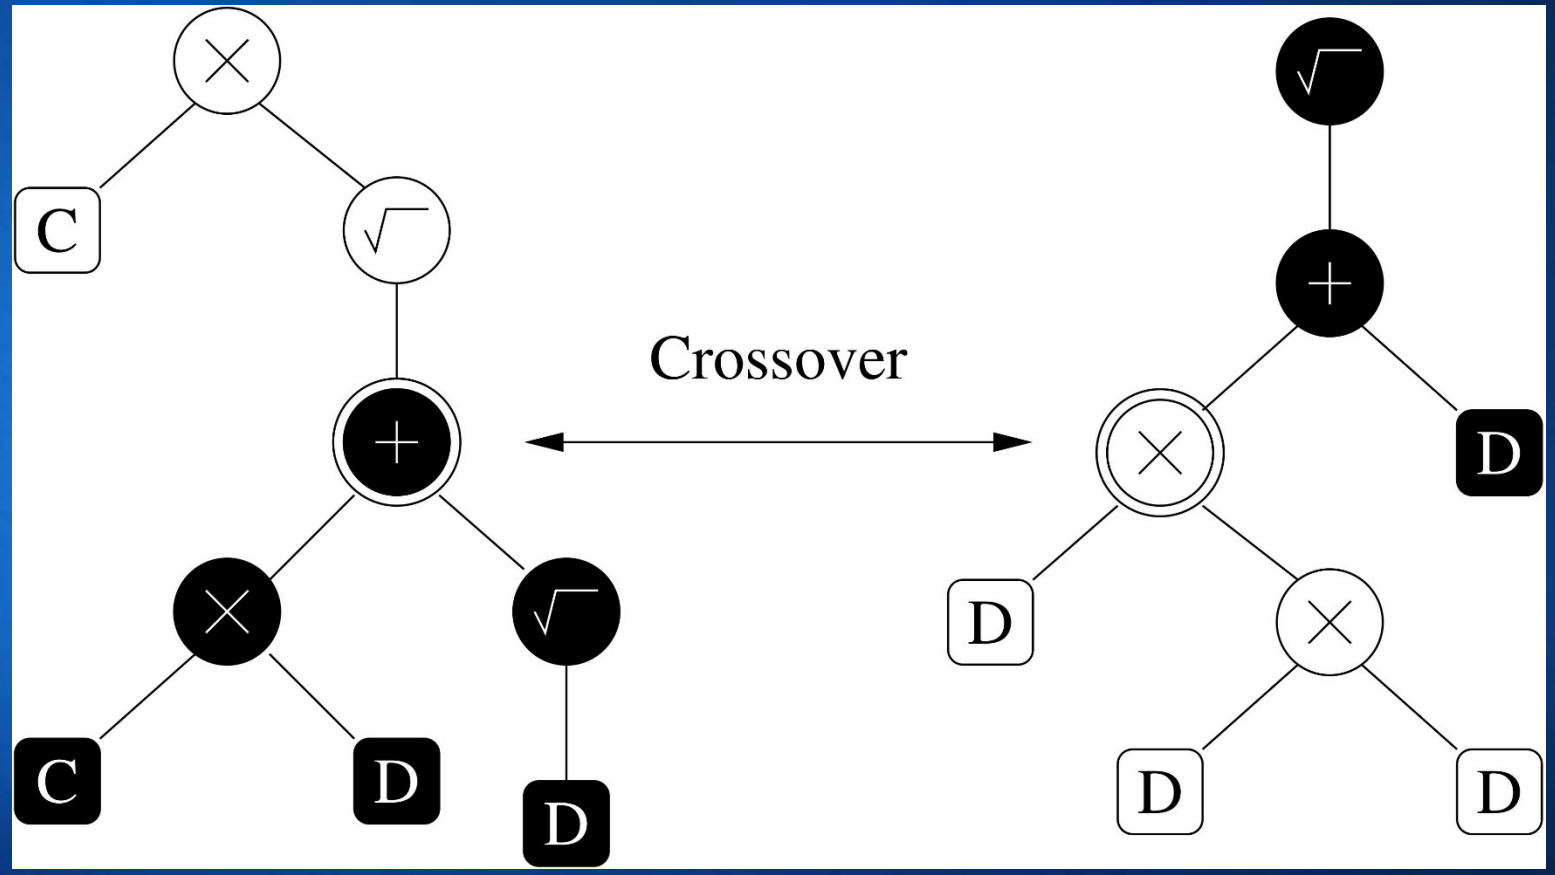

Evolving Robots

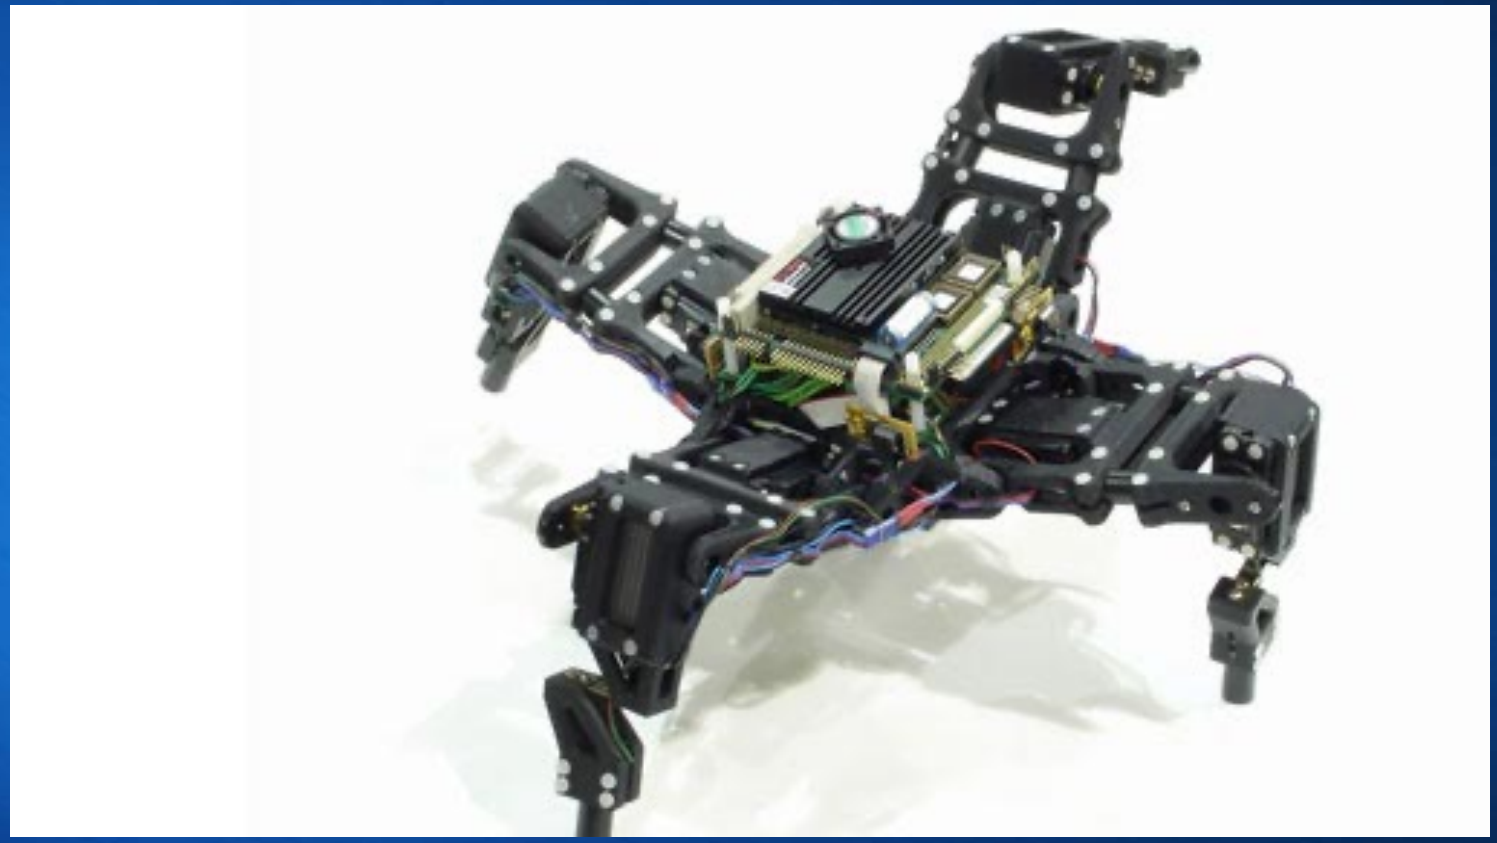

Bayesians

David Heckerman

Judea Pearl

Michael Jordan

Probabilistic Inference

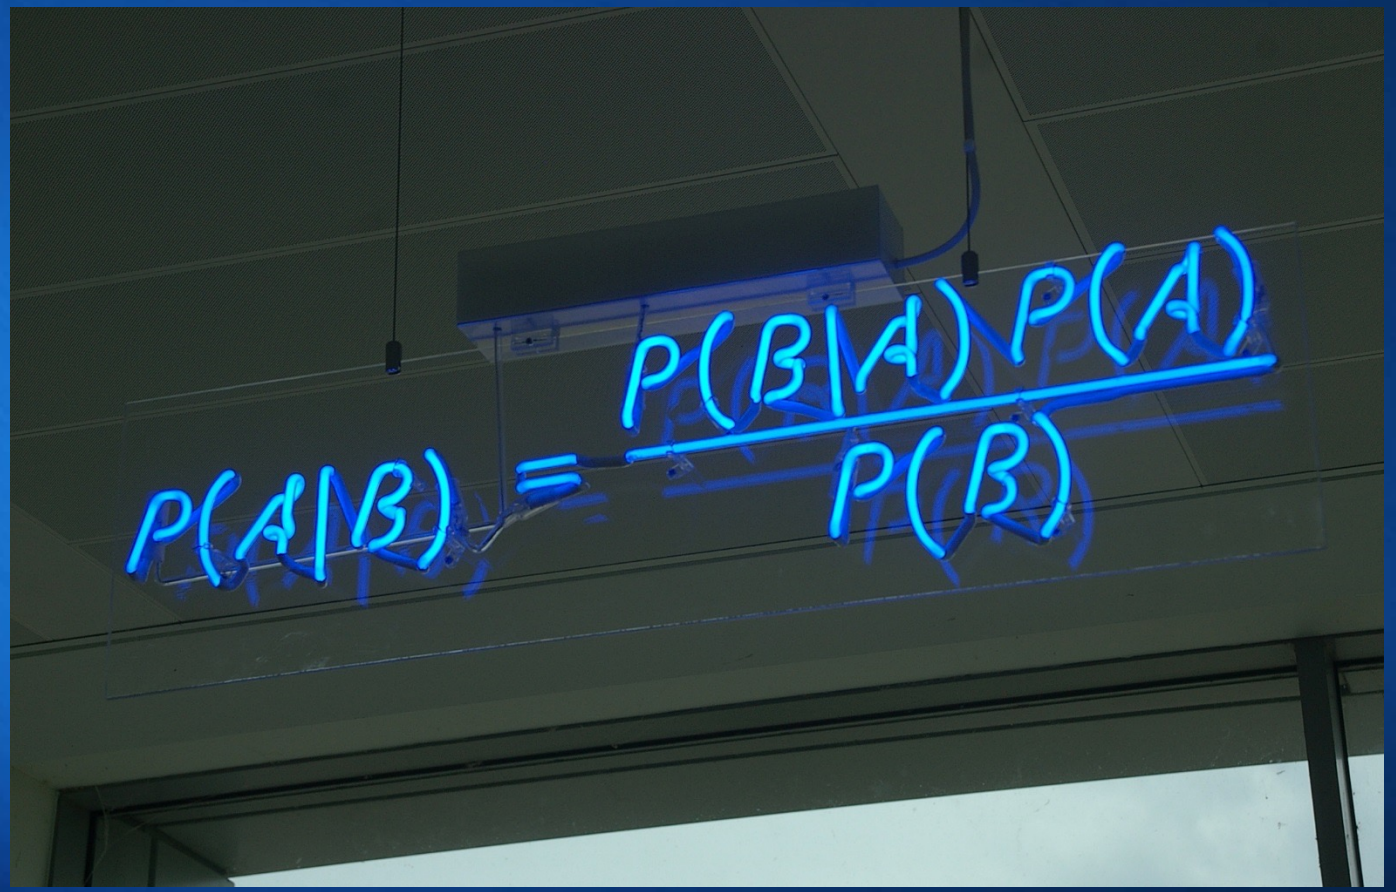

Probabilistic Inference

Likelihood

How probable is the evidence given that our hypothesis is true?

Prior

How probable was our hypothesis *before* observing the evidence?

$$P(H \mid e) = \frac{P(e \mid H)\, P(H)}{P(e)}$$

Posterior

How probable is our hypothesis given the observed evidence? (Not directly computable)

Marginal

How probable is the new evidence under all possible hypotheses?

$P(e) = \sum P(e \mid H_i)\, P(H_i)$

Spam Filters

Analogizers

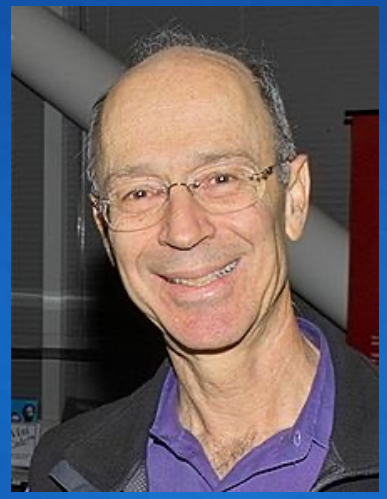

Peter Hart

Vladimir Vapnik

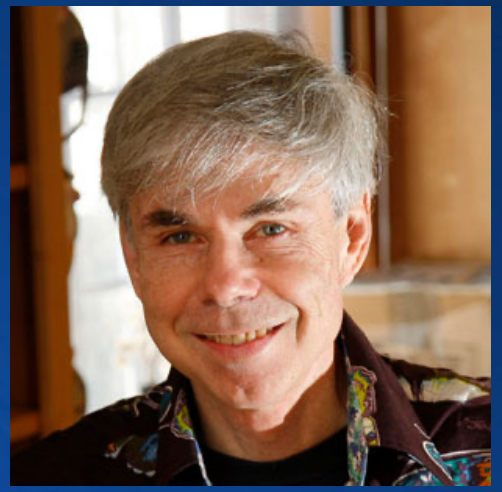

Douglas Hofstadter

Nearest Neighbor

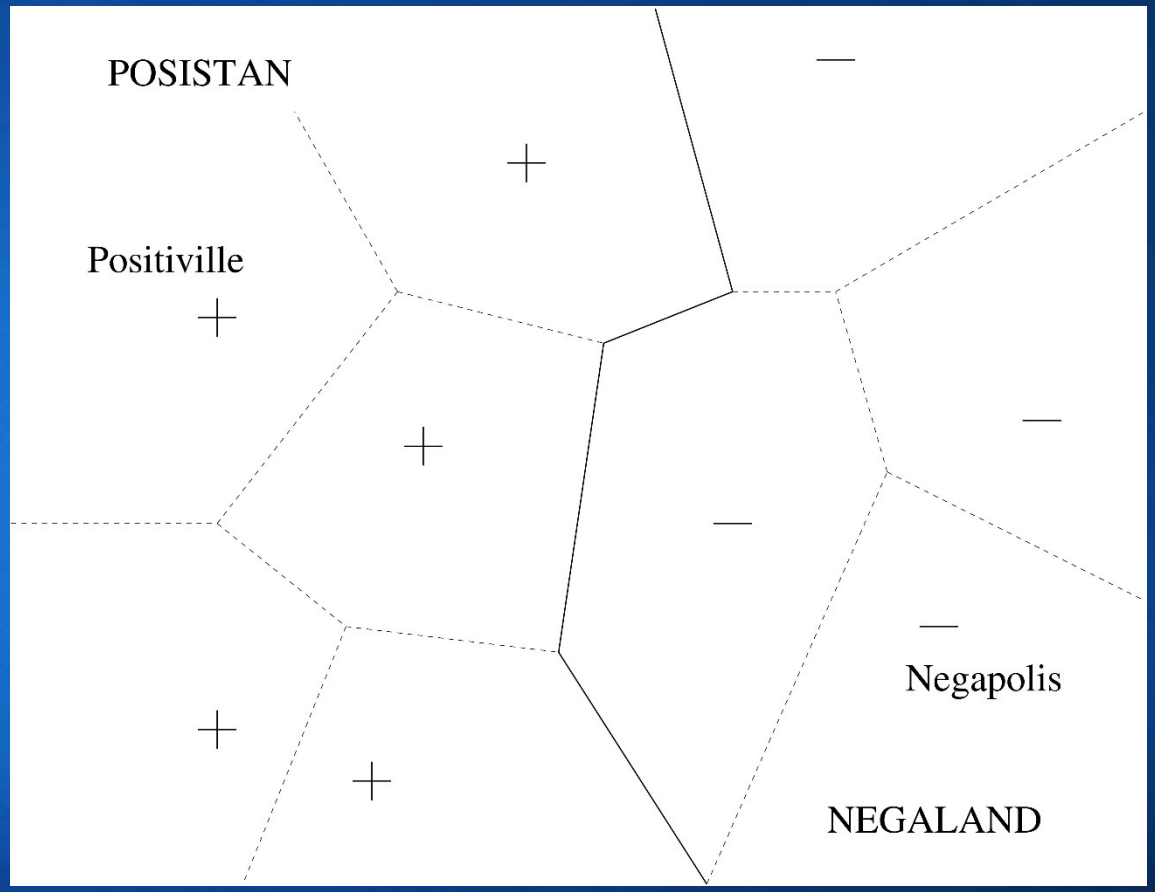

Kernel Machines

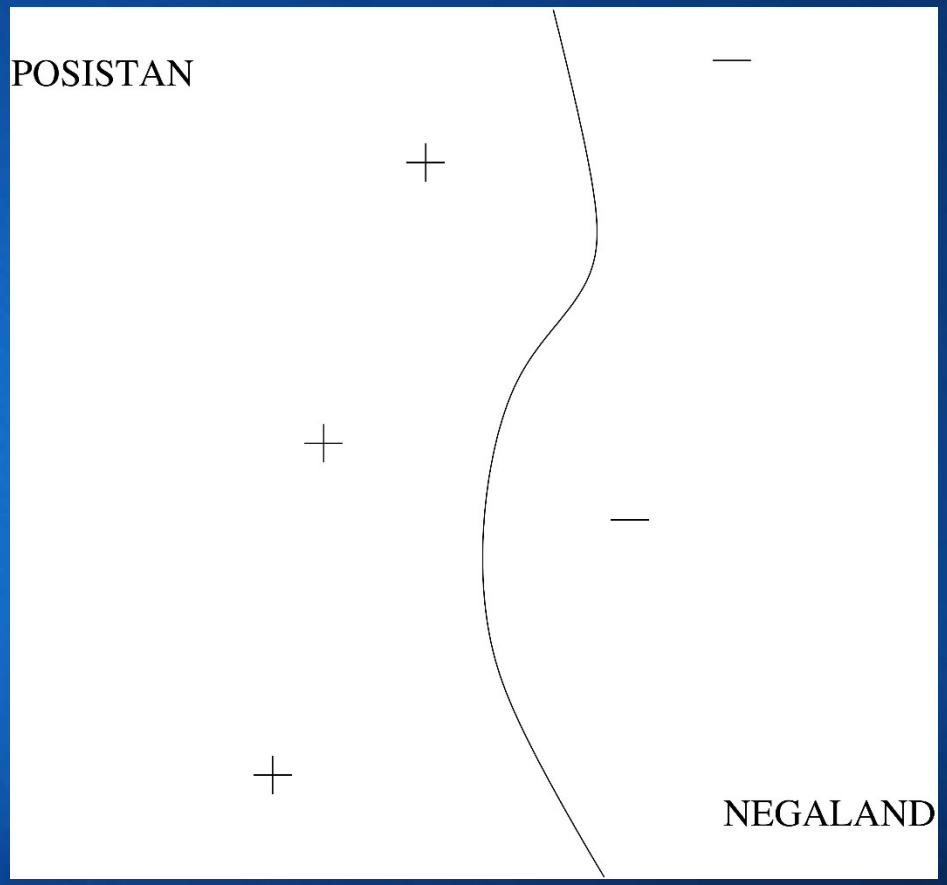

Recommender Systems

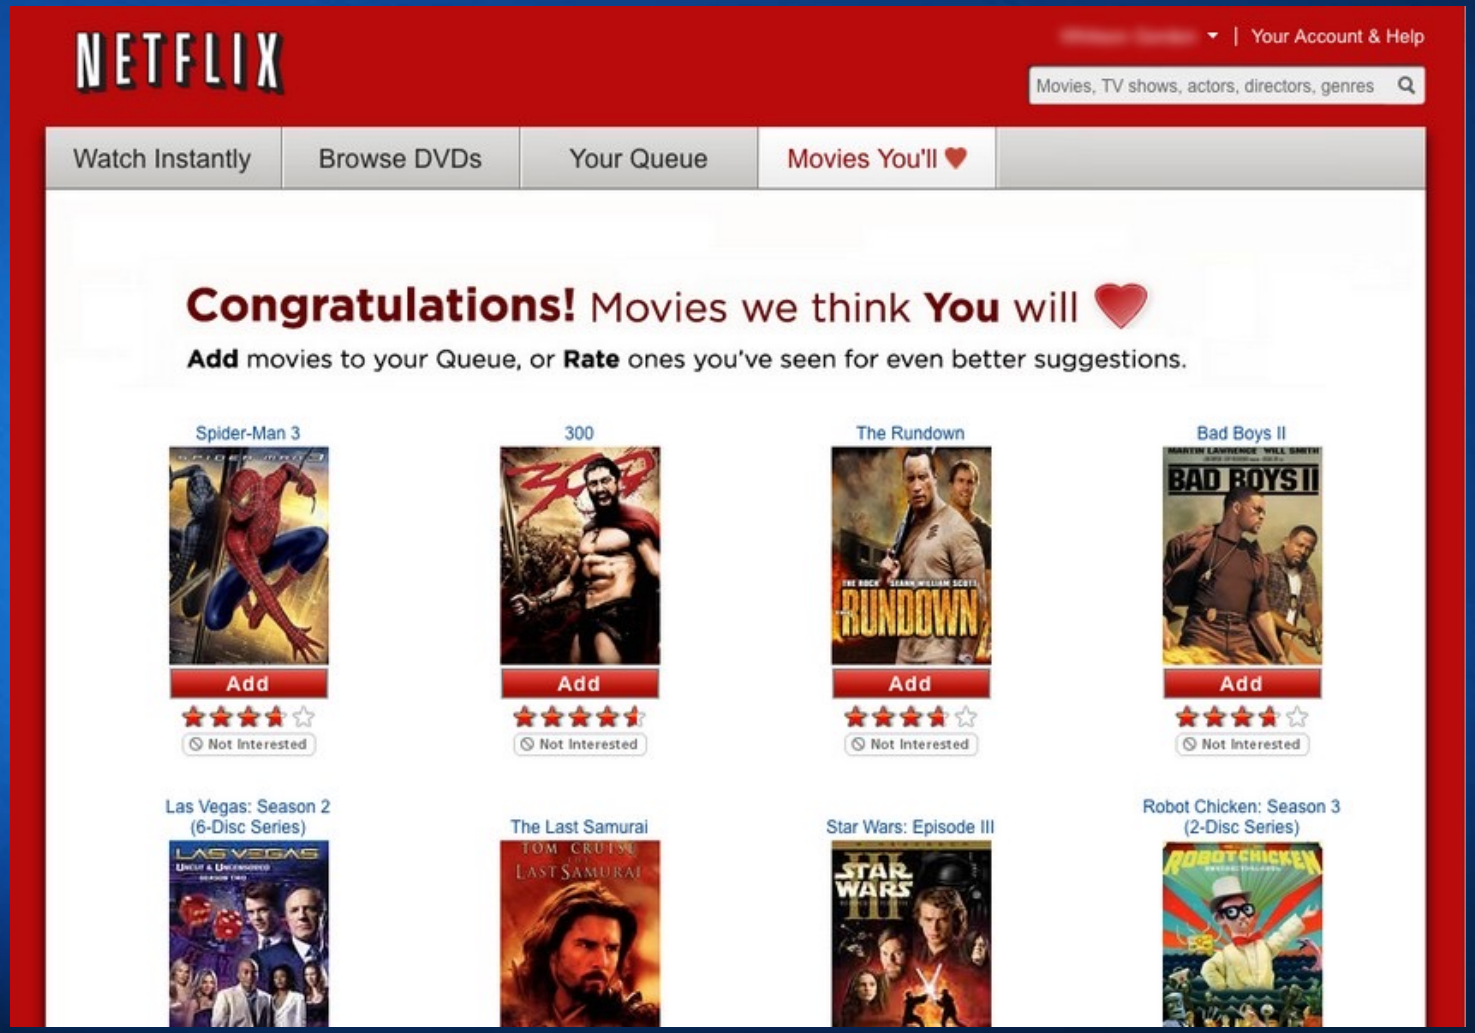

The Big Picture

Tribe	Problem	Solution
Symbolists	Knowledge composition	Inverse deduction
Connectionists	Credit assignment	Backpropagation
Evolutionaries	Structure discovery	Genetic programming
Bayesians	Uncertainty	Probabilistic inference
Analogizers	Similarity	Kernel machines

The Big Picture

Tribe	Problem	Solution
Symbolists	Knowledge composition	Inverse deduction
Connectionists	Credit assignment	Backpropagation
Evolutionaries	Structure discovery	Genetic programming
Bayesians	Uncertainty	Probabilistic inference
Analogizers	Similarity	Kernel machines

But what we really need is
a single algorithm that solves all five!

Putting the Pieces Together

- Representation
  - Probabilistic logic (e.g., Markov logic networks)
  - Weighted formulas → Distribution over states
- Evaluation
  - Posterior probability
  - User-defined objective function
- Optimization
  - Formula discovery: Genetic programming
  - Weight learning: Backpropagation

Toward a Universal Learner

- Much remains to be done . . .
- We need your ideas

What a Universal Learner Will Enable

Home Robots

World Wide Brains

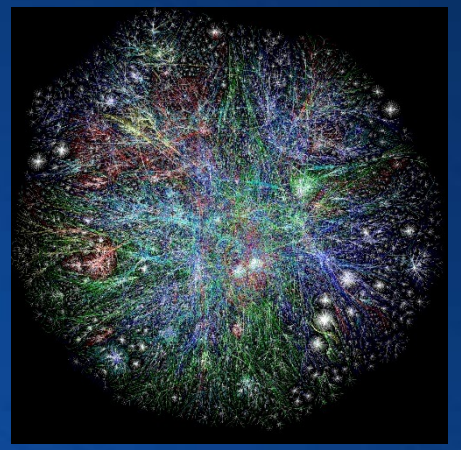

Cancer Cures

360° Recommenders

THE MASTER ALGORITHM

HOW THE QUEST FOR THE ULTIMATE LEARNING MACHINE WILL REMAKE OUR WORLD

PEDRO DOMINGOS