

ICS 435

Machine Learning

Fundamentals

Susanne Still
University of Hawaii at Manoa

Introduction

Learning and Adaptation

- Most intelligent systems show signs of learning.
- Most biological, “alive” systems utilize adaptation.

Challenges:

- Understand the principles of learning
- Build learning machines

Machine learning is crucial in robotics and AI:

- Often easier to build a learning system than to hand-code a program that works.
- Example: a walking robot on the moon. Many DOFs; changing environment.
- Typical tasks that require learning:
 - Speech
 - Handwriting and object recognition
 - Intelligent user interfaces
 - Motor behavior
- A true AI requires learning on many levels!

Machine learning is crucial for data analysis:

- Huge and very complex data sets *too large to analyze by hand*; for example:
 - CERN
 - data from complex systems, e.g. ecological
- High frequency task, *too fast to analyze by hand*; e.g. stock price prediction from trading data.
- Human can solve task, but can not explain how; e.g. character recognition.
- No human expert; e.g. DNA analysis.

Application Areas:

- Physics
- Bioinformatics
- Computer Vision
- Robotics
- Graphics
- Speech
- Financial analysis
- E-commerce
- Medicine
- Computer games
- Multimedia

Examples:

- ➔ particle physics
- ➔ microarray data
- ➔ object recognition
- ➔ decision making
- ➔ realistic simulations
- ➔ recognition, identification
- ➔ option pricing
- ➔ data mining
- ➔ diagnostics, drug design
- ➔ adaptive opponents
- ➔ retrieval across databases

Machine learning is one of the 21 century's core technologies.

Different kinds of learning

- **Supervised** -- there is a teacher signal
- **Unsupervised** -- no teacher signal
- Example: cancer tissue (UC Irvine ML repository)

diameter	perimeter	texture	...	outcome	time	training example
13.71	20.83	90.2	...	Recurrence	77	1
13	21.82	87.5	...	Normal	119	2
12.46	24.04	83.97	...	Normal	76	3

- **Supervised learning:** Binary classification. Predicting Yes/No label; here: cancer recurrent?
- **Supervised learning:** Regression. Predicting a continuous variable; here: time to recurrence.

- *Based on the role of the learner:*
active learning vs. **passive learning**
 - ▶ Most animal learning is active!
 - ▶ But most results in learning theory address passive learning. Complications for active learning. E. g. data are not i.i.d.
- *Based on the problem definition:*
Behavioral learning.
 - ▶ **Reinforcement learning:** a reward signal is present (a behavioral goal is specified)
 - ▶ Alternative theories for behavioral learning, e.g. based on information theory and physics.

Challenges:

- Understand the principles of learning
- Build learning machines



ICS 435

- Overview of some history and core ideas
- Supervised learning:
 - Introduction to statistical learning theory
 - Support Vector learning
 - Feed forward neural nets
 - Boltzmann machines, RBMs, deep learning
- Use of information theory in the objective function of the Boltzmann machine... leads us to explore:
- Physical limits of information processing...
 - ... those *directly* motivate information theoretic objective functions!
 - Introduction to information processing far from thermodynamic equilibrium, information theory, etc. (selected topics).
- ...the limits also lead straight to unsupervised machine learning.
- Unsupervised learning
 - Cluster analysis
 - Recurrent neural nets
- Optional: Behavioral Learning.

Learning machines

Some historical context

Machines that learn?



- Animals learn - how?
- What is learning?
- What is inference?
- What is a machine?
- Machines that can think?
- Machines that can do math?

Machines that compute

- Devices that help perform simple calculations go back to ancient civilizations!
- Mesopotamia (ca. 2700BC):
Abacus.
Pebbles (latin: calculi)
- Predicting eclipses,
“Antikythera mechanism”
(ca. 200BC, 30 gears),
“analog computation”.



- South America (ca.2600BC):
Quipu



Abacus vs pen and paper



a competition, 1503


“abacist vs algorist”

abacus faster, but no
written record

(introduction of
arabic numerals in
Europe)

Adding machine

- People make mistakes, arithmetic is tedious, so, let's build a machine!
- Pascal (1642)



Analog computing machines

- Mechanical device for tide prediction based on Fourier analysis
- Lord Kelvin (1872)
10 components
- 62 components,
Germany (1935)
(in use until 1968!)



“Modern” digital computer

- Based on boolean logic (George Boole 1847)



- Implemented basically via switches

Truth table for AND,

A	B	Z
0	0	0
0	1	0
1	0	0
1	1	1


$$Z=A.B$$

Switch realisation of AND:



Continuity occurs only when both A AND B are closed.

- Charles Babbage envisioned mechanical computers, but never completed one. (25,000 metal parts, 15 tons, precision issues.)



- Programmable

- Ada Lovelace wrote first “program” for his machine.

Probably the first built...

(Konrad Zuse 1938-41)

Z1-Z3, built in his parent's apartment in Berlin...

- 30,000 metal parts (reliability problems!)

Destroyed in
1944 by WW2
bombing



... the Z4 (1950)


First computer sold (that actually worked)
sold to ETH Zurich, Switzerland



- 2500 electro-mechanical relays
- **256 byte memory** (64 32-bit floating point words)
- **one multiply every 3 seconds** (modern laptops billions of times faster)


Another historical example: american Electronic Numerical Integrator And Computer (ENIAC, 1945)

- electronic (1000x faster than electromechanical; **356 *multiply/sec***)
- > 18,000 vacuum tubes
- programmed by plugboard and switches




UNIVAC

- first commercial computer produced in the US
- The first UNIVAC was accepted by the US Census Bureau on March 31, 1951
- 5,200 vacuum tubes, 29,000 pounds (13 metric tons), consumed 125 kW




- ***about 1,900 operations per second***


From vacuum tube computers, to transistor computers, integrated circuit computers, **but...**




tubes (4 logic gates, 1950s)



8 logic gates (1960s)



ICs hundreds (1970s)



ICs millions (1990s)

... basically

all digital computers are the same

The Turing-Church thesis (very informal)
any “computable function” that one computer can do,
any other digital computer can also do;
they are all are equivalent to Turing’s machine.
(speeds may vary enormously)


Origins of “modern computer science” (1930s)

- Foundations were laid by some breakthroughs starting in the 1930s with the work of Alan Turing, Alonzo Church, Kurt Goedel, John v. Neumann, Emil L. Post, and others.
- ***Alan Turing (1936) introduced abstract model of a general purpose computer.***

Ambition: understand how living systems think; understand intelligence!

- Motivation driving the great minds that started this “new discipline”, which is now Computer Science.
- Goes back to Alan Turing, John von Neumann, Norbert Wiener (cybernetics), and others, starting in the 1930s
- *Physicists and mathematicians working together with neuroscientists*
- First models of neurons. (Pitts, McCulloch)
=> neural nets => machine learning

Cybernetics



- Information processing, control and (self-) regulation in the animal and the machine
- 10 Macy Conferences (1946-53); eclectic group of interdisciplinary participants.

Visionary field spun/influenced a many areas, for example:

- Control theory and *Information theory*
- Computational neuroscience
- *Neural networks (machine learning)*
- Neuromorphic engineering
- Theory of self-organization of living systems

Information theory

- Basis for communication (wired and wireless), e.g. transmitting TV signals, internet, phone.
- At the core lies the question:

What is information?
How to measure it?

- Original intellectual motivation: move from an energy based description of the world to an information based one.

Claude Shannon

- While Wiener attacked the hard problem of continuous information processing,
- Shannon made progress by considering discrete symbols: he showed that to fulfill simple assumptions, information is best measured by $\log(1/p)$
- Average
$$-\sum_x p(x) \log[p(x)]$$
 is directly related to Gibbs entropy (see E.T. Jaynes)
- Measures uncertainty, and the maximum (average) amount of information that can be gained by measuring the outcome of the random variable x

- **Information** is uncertainty reduction

$$I[X, Y] = H[X] - H[X|Y]$$

$$H[X] = - \sum_x p(x) \log[p(x)]$$

$$H[X|Y] = - \sum_{x,y} p(x, y) \log[p(x|y)]$$

- **Channel capacity** is the maximally transmittable (rate of) information $\max_{p(x)} I[x, y]$
maximum over all information sources
(x is input, y is output of the information channel)

- A **continuous signal** has infinite information rate.
- But infinite resolution is irrelevant for most applications, some level of distortion is tolerable.
- The achievable rate of a continuous information source, if transmitted to finite resolution, i.e. for fixed average distortion is:

$$R(D) := \min_{p(s|x)} I[s, x]$$

$$\text{s.t. } \langle d(s, x) \rangle_{p(s, x)} = D$$

- Represent original signal, x , by encoded signal, s .
Given: distortion function $d(s, x)$; information source $p(x)$
- (units: convert between information in bits per symbol, and rate in bits per second: multiply by a constant - symbols per second)

Computational neuroscience to neural networks

- Computational neuroscience produces mathematical models of neurons of varying degree of complexity
- One of the first was the McCulloch-Pitts model, pioneered by Pitts in the early 1940s
- Simple model leads to simplest “learning machine”, the “Perceptron” (F. Rosenblatt 1957)
- From perceptron to multi-layer neural nets to the “Neocognitron”, to “deep learning”...

...we will hear more about this later...

Neuromorphic engineering

- Interesting observation (Mead, late 1980s)
Transistors in the sub-threshold regime:
current depends approximately exponentially
on gate voltage. Similar current-voltage curve
in ion channels (building blocks of cell
membranes and neurons in the brain).
- Allows for a biology-inspired approach to
computing, adaptive and able to learn
- *analog* VLSI

Neuromorphic engineering

- Achievements include:

- Silicon retina
(Mahowald 1988)

- Silicon cochlea
(Lyon and Mead 1988)

- Silicon neuron (Mahowald 1991)

- Silicon synapse (Dorio et al 1996)

- Cognitive systems (Indiveri et al. 2013-present)

- Applications include:

- Prosthetics


- Low power devices

Back to the main track...

- **Origins of neural networks** come from mathematical descriptions of neurons (brain cells)
- *Physicists and mathematicians working together with neuroscientists*
- From neuron models to neural nets to machine learning...

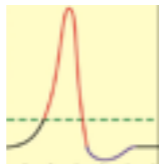
From brains to learning machines



- **Neurons:** central nervous system (CNS) has ca. 10^{11} . They are:
- separate entities (Ramon y Cajal, Nobel Price 1906)
- connected by synapses (Sherrington and Adrian, Nobel Price, 1932)
- many different types of neurons with different functionality



Optic tectum (sparrow), drawing by Cajal

Neurons communicate


- Information processing within a neuron: electrical. Action potential (“spike”). 
- Information processing between neurons: (mostly) chemical (exception: gap junctions)



http://ffden-2.phys.uaf.edu/212_fall2003.web.dir/Keith_Palchikoff/biological%20neuron.JPG

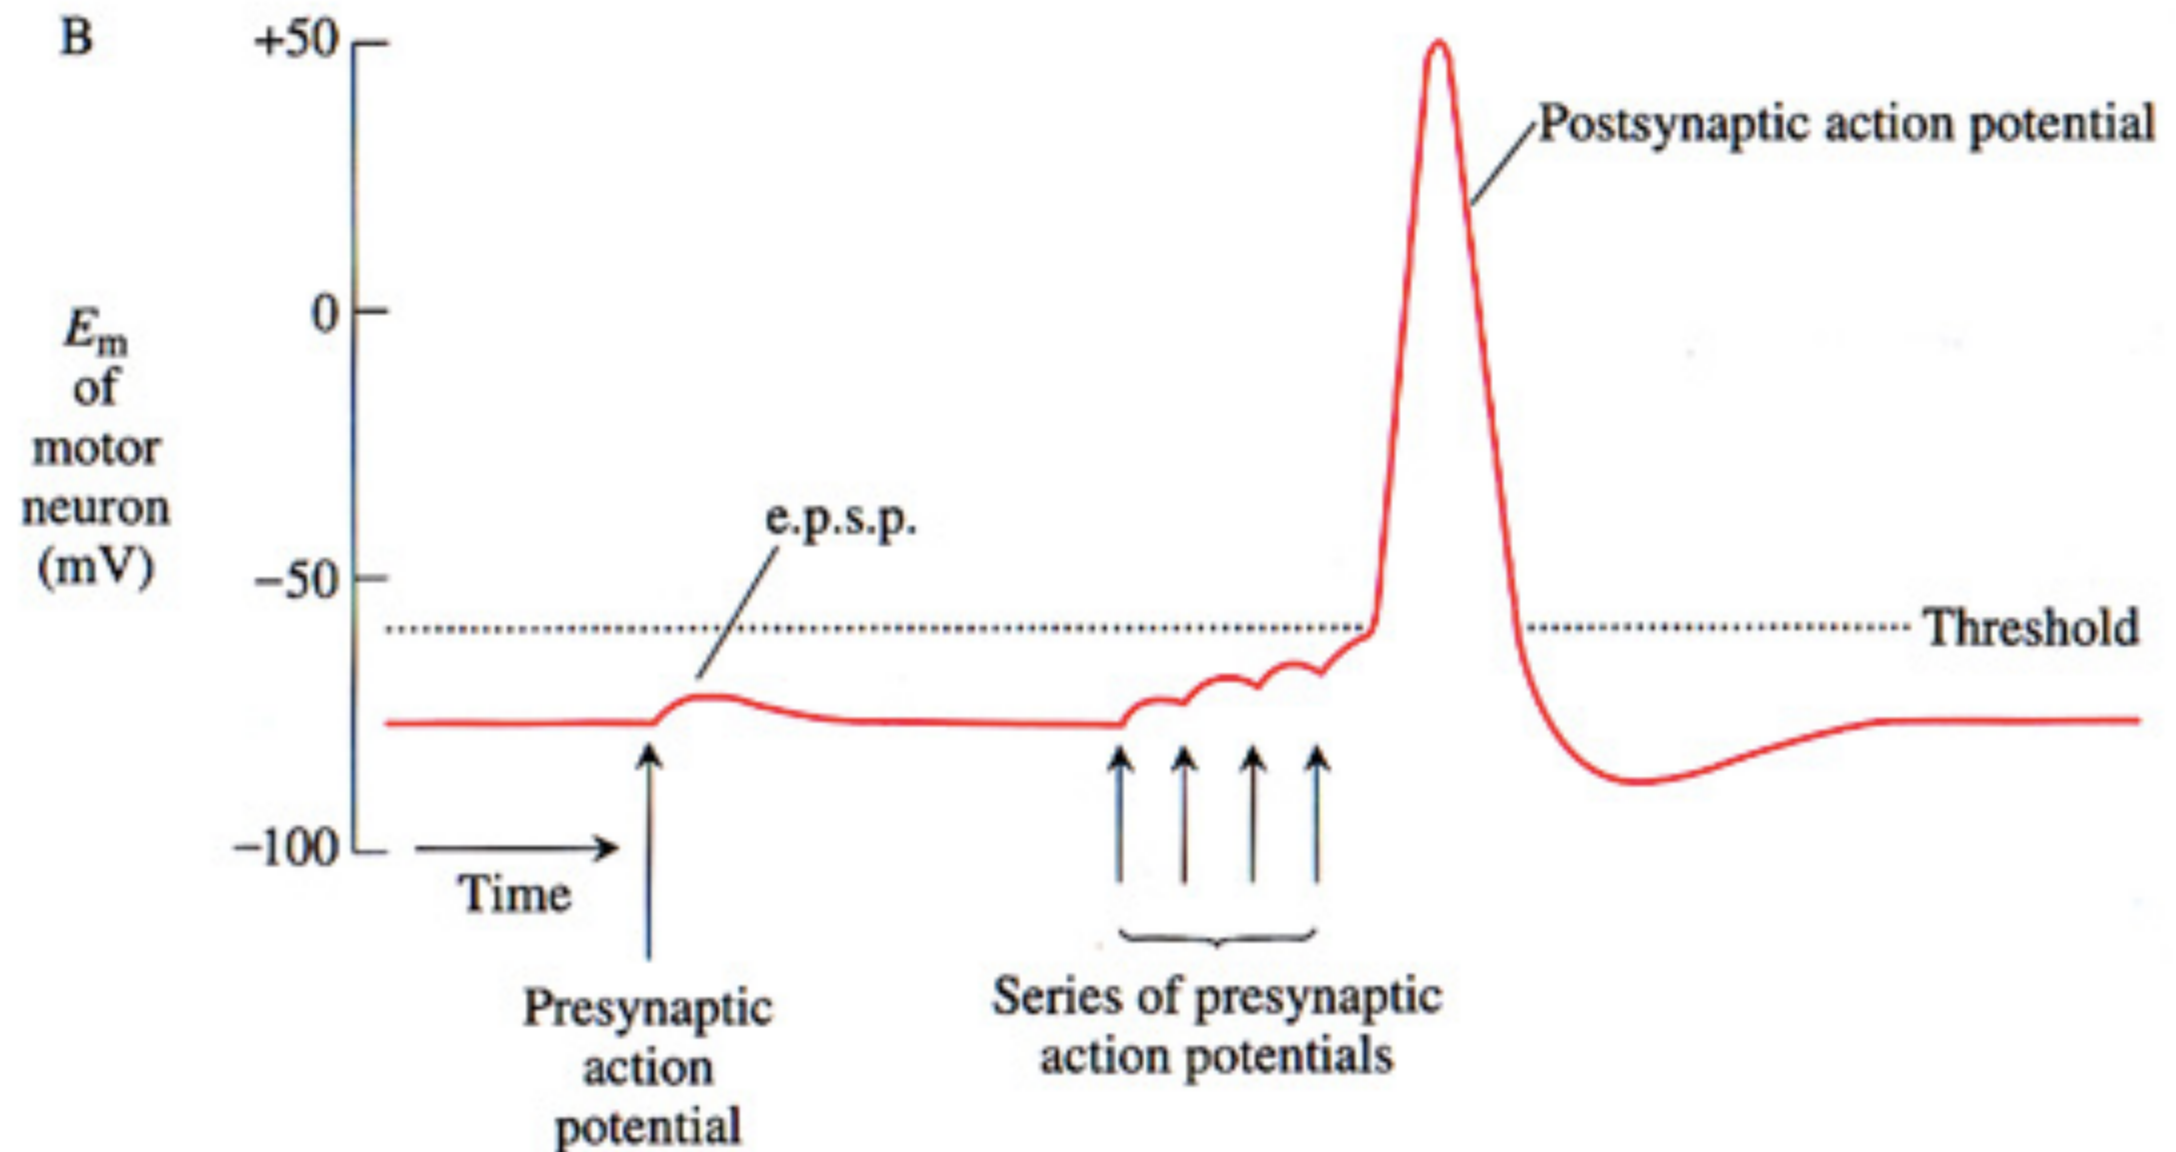
http://www.brain.riken.go.jp/english/g_braaw/images/g5/synapse.gif

- Pre-synaptic action potential depolarizes synaptic terminal
- Voltage-sensitive calcium ion channels open => calcium enters
- Release of neurotransmitter; diffusion through synaptic cleft.
- Neurotransmitter binds to post-synaptic receptors
- Ions enter post-synaptic cell => EPSP

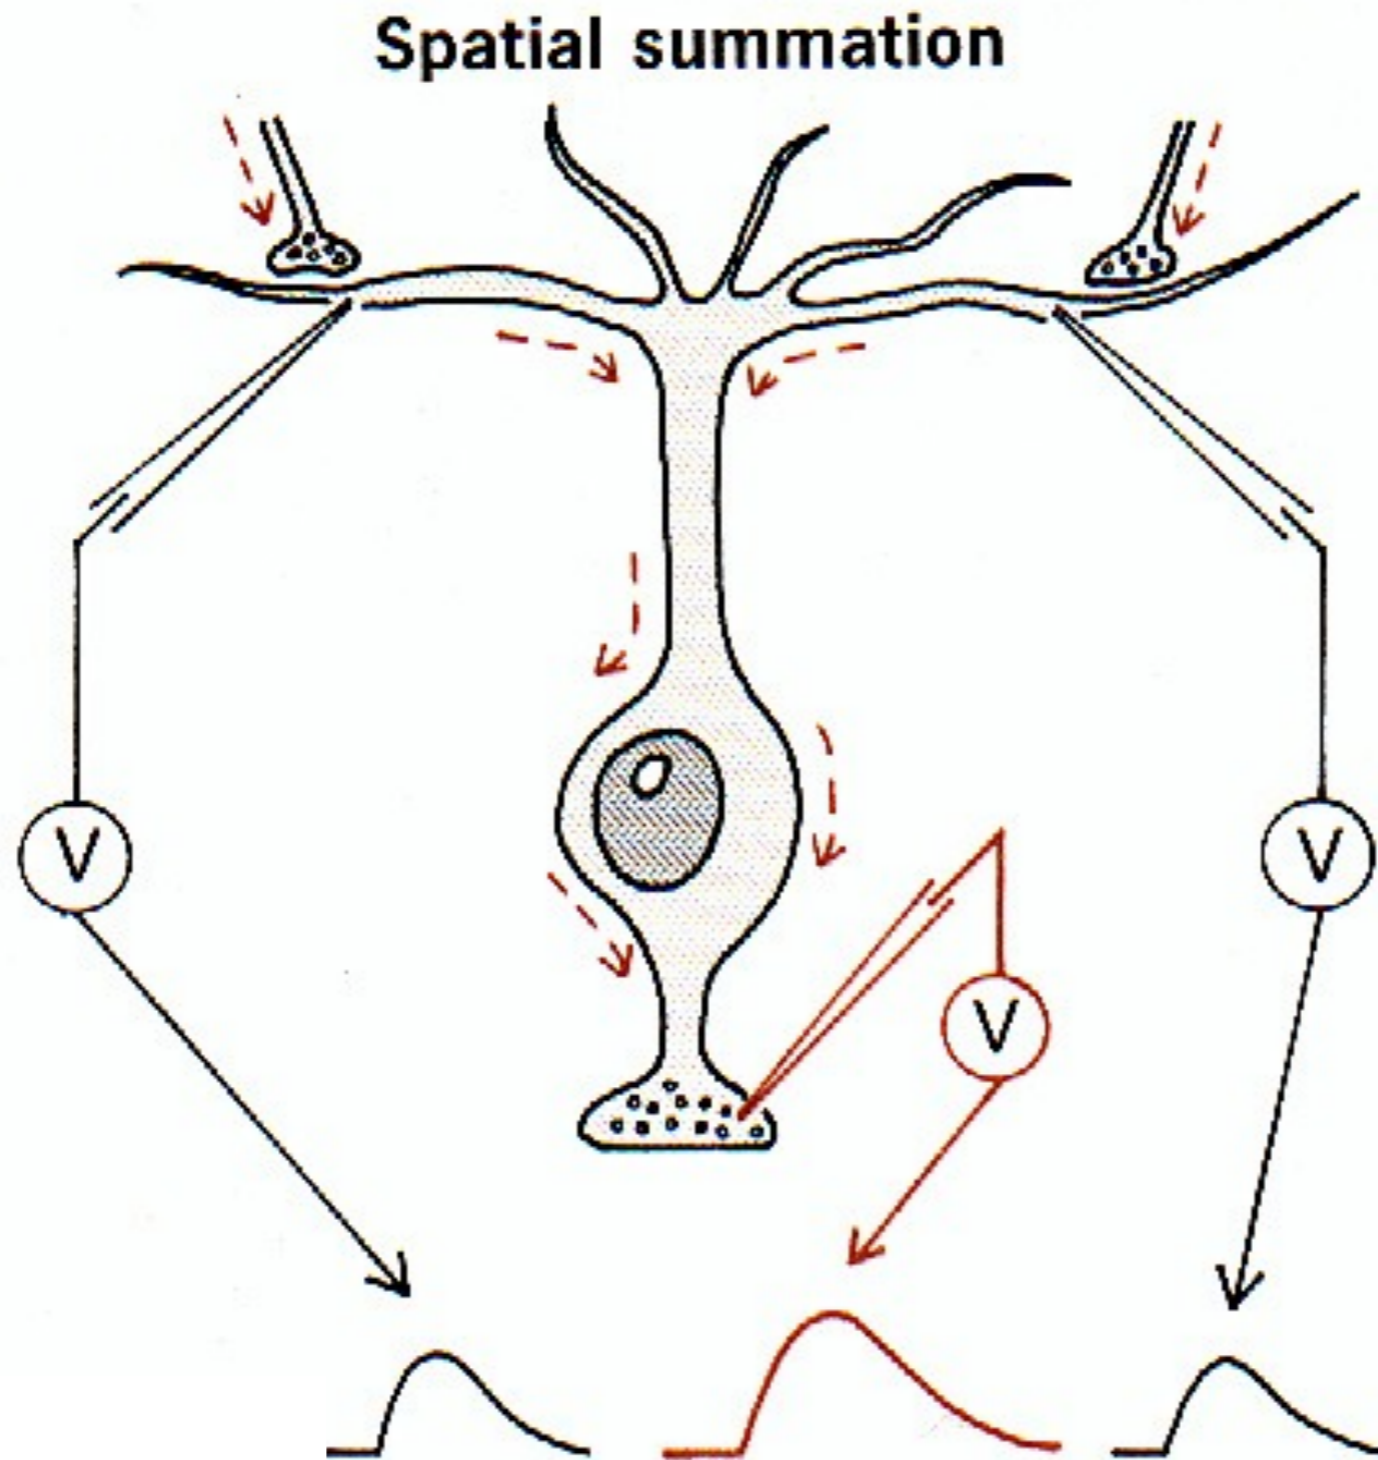


(Temporal Integration of EPSPs)

(B) Responses of the postsynaptic motor neuron to action potentials in the presynaptic sensory neuron. At the upward arrows, action potentials are triggered in the presynaptic neuron by an electrical stimulus.

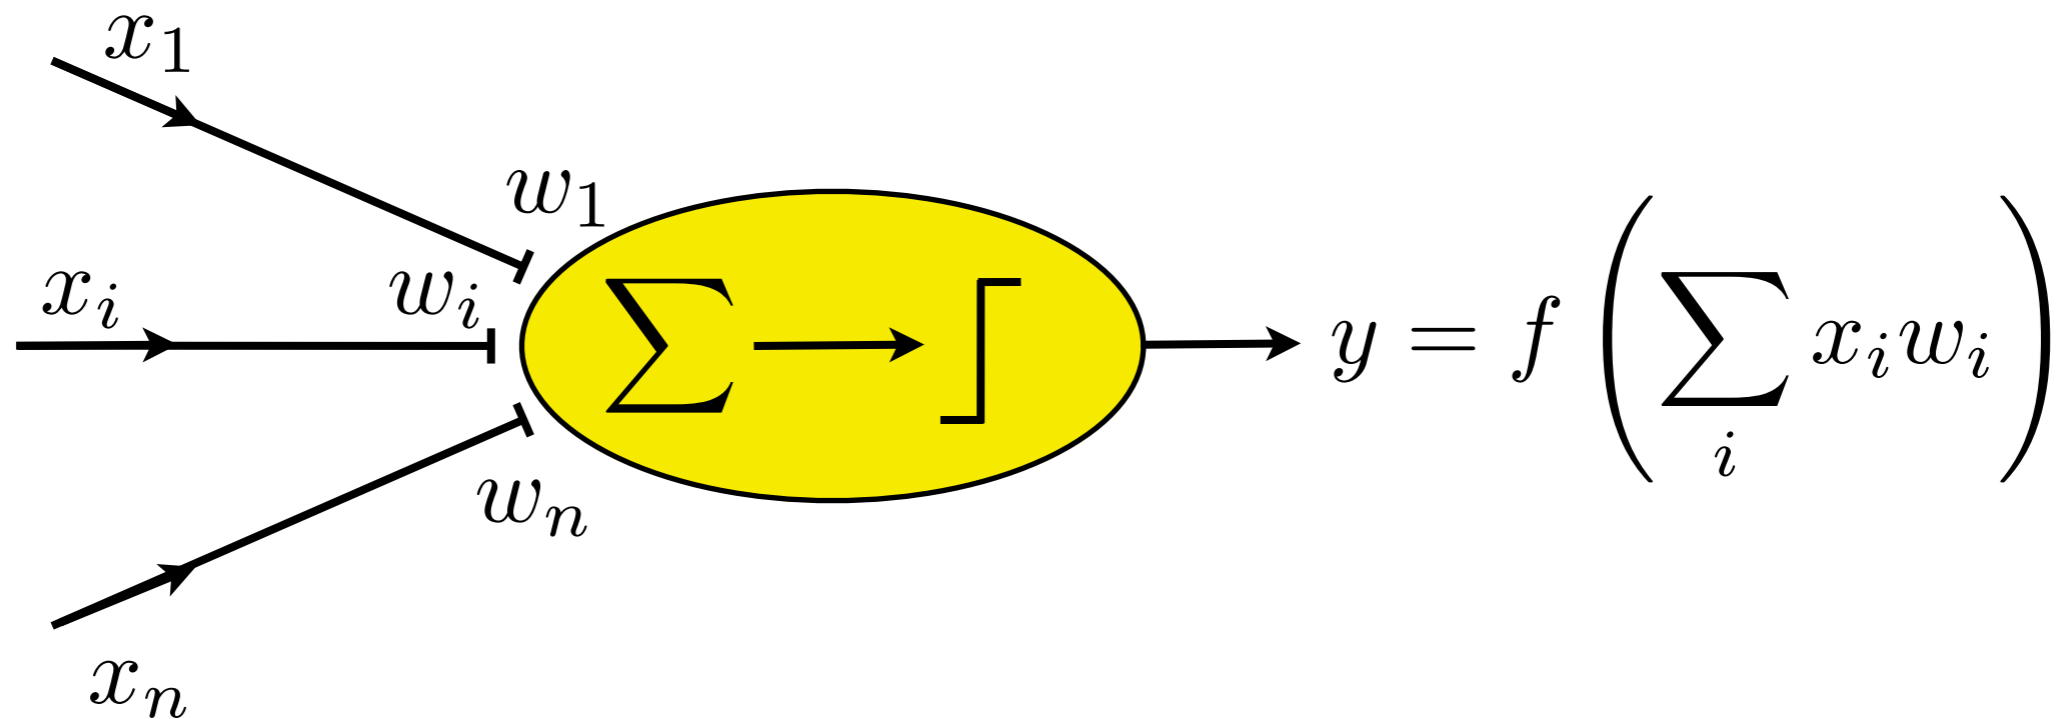


Spatial Integration of EPSPs



Mathematical simplification


- Neurons are either on or off: represented by binary value.



- Inputs form n neurons: x_i
- then added $\sum_i x_i w_i$
- get multiplied by weights at “synapses”: $x_i w_i$
- if above threshold, then neuron is “on”.
- Remark: add x_0 to the input vector, in order to write the bias as $b = x_0 w_0$ to get more compact form $\mathbf{w}\mathbf{x}$, rather than $\mathbf{w}\mathbf{x} + b$


Transfer function

- Step function:



- Used by Pitts & McCulloch (1942), and in Rosenblatt's Perceptron (1957)


- Sigmoid (used in many “modern” feed forward neural nets):



- Sigmoidal function is differentiable (good for deriving gradient decent learning rule; Backpropagation algorithm, Werbos 1974)

Neural networks (1940s/50s)

- Pitts and McCulloch (1942/3): “Formal” (mathematical) neurons thought of as processing units
- *Networks* can emulate any logical function.
- D. Hebb: Connections between neurons can change. “Learning rule”: strengthen proportional to correlation between activity of pre- and post synaptic neuron.



Perceptron (Rosenblatt, 1957)

- Probably the first “learning machine”. Artificial neuron that solved a classification task (supervised learning):
- Given: N input vectors \mathbf{x} together with labels l (these labels are the “teaching” signal).
- Goal: for any given input, the output of the classifier should be the same as the label.
- Simplify: assume labels are binary, either -1 or $+1$.
➔ *Binary classification.*



Perceptron learning


- *adjust weights according to the correctness of the output* until all input data in training set are classified correctly.
- **error measure:** compare output of the neuron (y) to desired output (= label, l): both are either -1 or 1, so:
 - $y \cdot l = 1$: correct classification, then: $l - y = 0$
 - $y \cdot l = -1$: incorrect classification, then: $l - y = 2l$
- adjust weight vector \mathbf{w} for each *misclassified* input \mathbf{x} , by adding $l \cdot \mathbf{x}$. This turns \mathbf{w} towards/away from \mathbf{x} if $l=1/-1$
- can do $\mathbf{w} \leftarrow \mathbf{w} + c(l-y)\mathbf{x}$ for all training examples \mathbf{x}
- c : step size parameter called “learning rate”; $0 < c < 1$

Representational power

- the decision boundary of the perceptron is a line:

$$\mathbf{w}\mathbf{x} + b$$

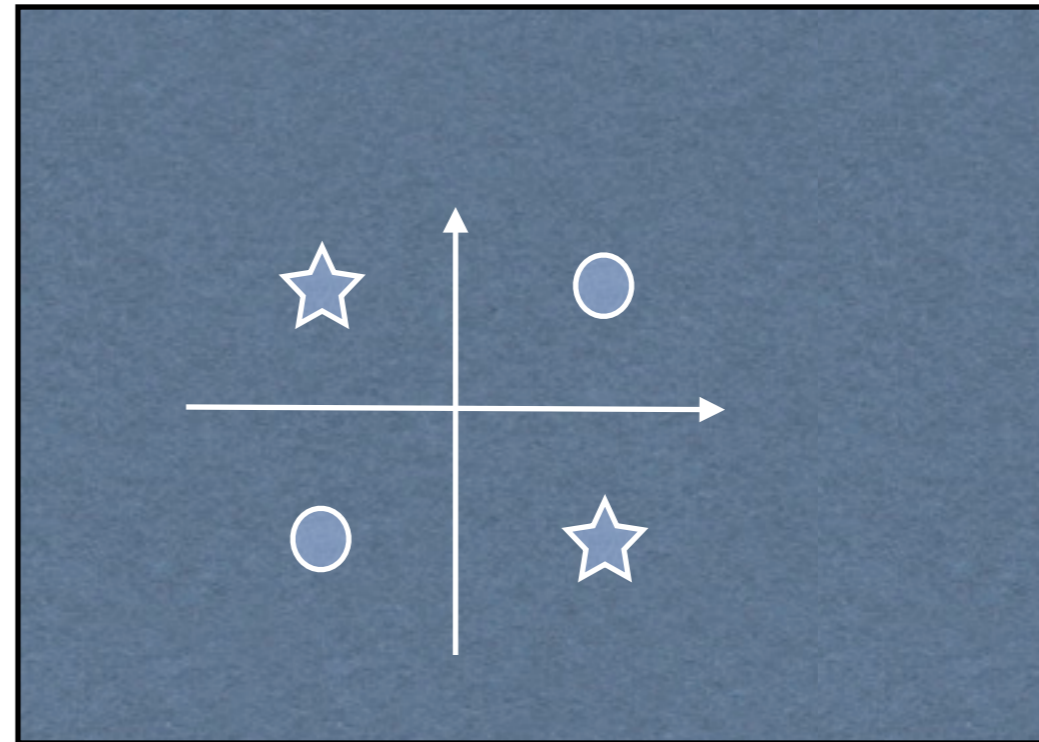
- the perceptron learning algorithm *converges* if input data are **linearly separable**.



- Novikoff (1962): Perceptron Convergence Theorem; first *margin-based error bound* (in a way the “dawn” of statistical learning theory)


XOR problem

- Perceptron algorithm can not classify input data which can not be separated by a line.
- For example XOR




XOR problem

- Perceptron algorithm can not classify input data which can not be separated by a line.
- For example XOR




XOR problem

- Perceptron algorithm can not classify input data which can not be separated by a line.
- For example XOR




XOR problem

- Perceptron algorithm can not classify input data which can not be separated by a line.
- For example XOR




XOR problem

- Perceptron algorithm can not classify input data which can not be separated by a line.
- For example XOR




XOR problem

- Perceptron algorithm can not classify input data which can not be separated by a line.
- For example XOR



XOR problem


- Perceptron algorithm can not classify input data which can not be separated by a line.
- For example XOR



- **A *single artificial neuron* can express the boolean functions AND, OR, and NOT, but not XOR.**

Minsky and Papert (1969)

- Shortcomings of perceptron were pointed out in a fierce manner by proponents of **symbolic AI** in what became a very influential book: “Perceptrons”.
- Their intuition was that **neural nets would not be useful in practical applications.**
- This put a stop to neural network funding, and research was significantly diminished.




Come-back of neural nets and beginnings of machine learning

- Backpropagation algorithm (Werbos, 1974).
- Neocognitron (Fukushima, late 1970s).
- Backprop re-discovered, used to train multi-layer networks in the 1980s: first big successes in pattern recognition.
 - ➔ Training convolutional NNs, eventually “deep learning”.
- 1960s-1980s: Vapnik with Chervonenkis and others looked carefully at the classification problem and developed statistical learning theory, addressing the problem of *model complexity control*. Foundation for:
 - ➔ Support Vector machines (SVM, 1963)
 - ➔ Kernel machines (1990s)

Backpropagation

- Use differentiable transfer function
- Gradient descent over all units in network: use chain rule to compute gradient.
- Forward pass: compute outputs of all units starting from input layer, ending with output layer
- Backward pass: compute gradient starting from output layer.
- Update the weights.
- Stochastic gradient descent - helps with computational cost and thus speed.



Kunihiko Fukushima

- Built the first electronic retina out of discrete components (1970, with Y. Yamaguchi, M. Yasuda, S. Nagata)
- Neocognitron, the “grandfather” of convolutional neural nets (1979/80)

Neocognitron

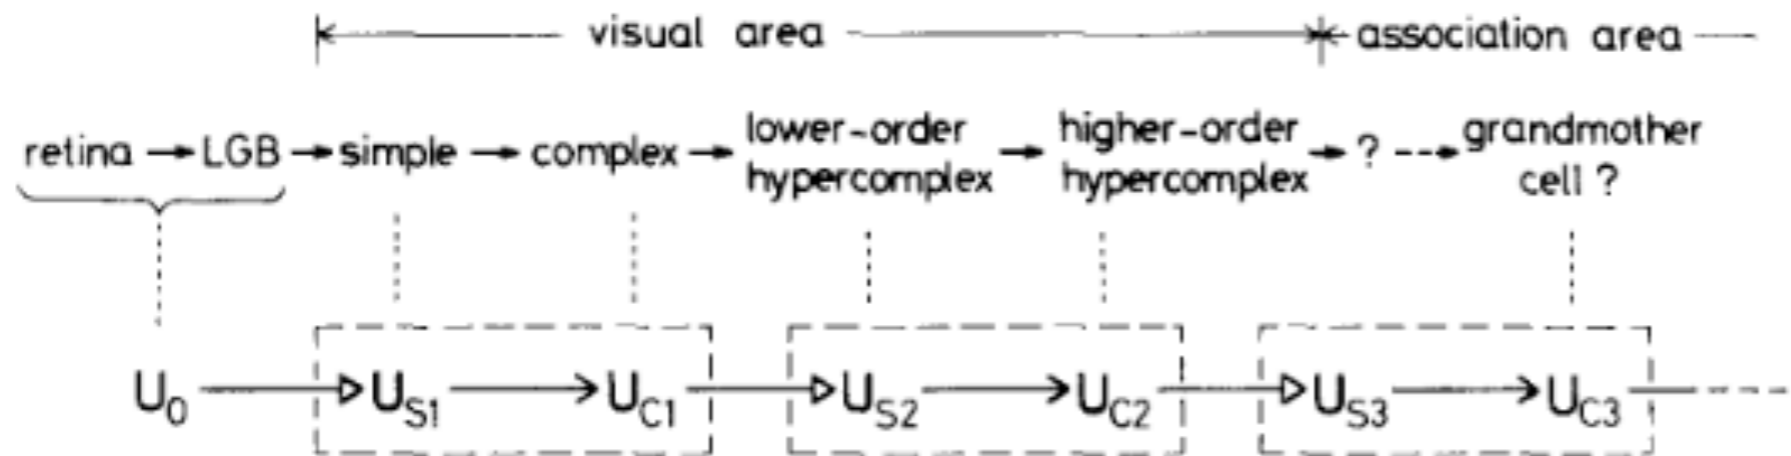


Fig. 1. Correspondence between the hierarchy model by Hubel and Wiesel, and the neural network of the neocognitron




Fig. 2. Schematic diagram illustrating the interconnections between layers in the neocognitron

ca. 1980s-mid 90s: Neural nets

- Recurrent NNs as Ising models (Hopfield)
- Boltzmann machine (Hinton, Sejnowski)
- Convolutional network trained by backprop alg. for *hand written digit recognition* (LeCun and others).

1990s: SVMs, Kernels and Bayes

- Support vector machine alg. (Vapnik&Chervonenkis 1963)
- Nonlinear classifiers via Kernel trick (Vladimir Vapnik with Isabelle Guyon, Bernard Moser, 1992)
- Soft margin SVM implements *empirical risk minimization* (Vapnik and Corinna Cortes, 1995)
- More Kernel machines (Kernel PCA etc., Bernhard Schoelkopf, A. Smolla, and many others)
- Bayesian Inference and probabilistic foundations of ML (MacKay, Neal, Hinton, Bishop, Jordan,...)
- Increasing amount of information theory applied to machine learning, and to neuroscience (Bialek and others)

New Millenium

- Proliferation of ML methods!
- Many new application areas
- “Big data”
- NNs big comeback in “deep learning”
- New hardware:
 - GPUs
 - Quantum computers... ?!
 - Quantum Machine Learning

ML techniques

- This is just a selection...
- NIPS conference proceedings (Neural Information Processing Systems)
- Journal of Machine learning research
- Other conferences, such as UAI (Uncertainty in Artificial Intelligence), ICML(Int. Conf. Machine Learning)
- Obscure “old” books

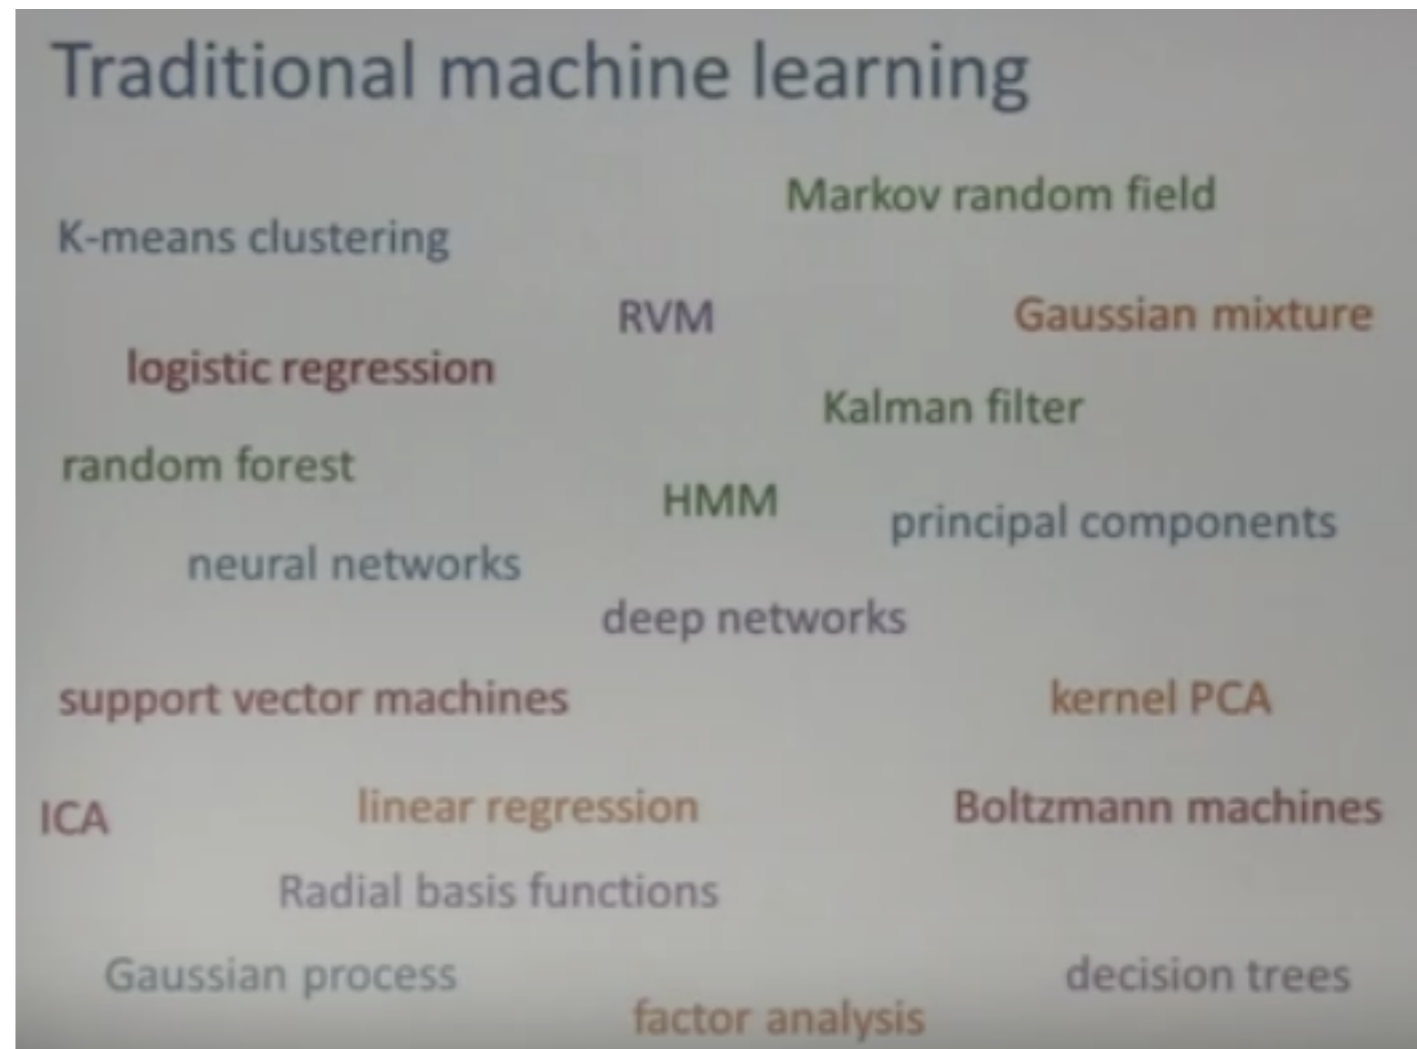


Image stolen from Chris Bishop's lecture at MLSS 2013, MPI Tuebingen

Ack! Too much “stuff”

- Are there some simple principles?
- *Is there physical reality behind all of this?*
 - ▶ Living systems are physical systems. Can we find physical principles to guide us through the jungle of learning and data processing algorithms?
- Study physical limits of information processing...
- Hope to get to building principles!

Homework

- Implement perceptron learning algorithm: N Inputs \mathbf{x}_j & labels l_j
 - ▶ Initialize weight vector \mathbf{w}
 - ▶ While there exist misclassified examples:
 - ▶ Compute output $y_j = \theta(\mathbf{w}\mathbf{x}_j)$
 - ▶ For each example, update the weights: $\mathbf{w} += c(l_j - y_j)\mathbf{x}_j$
- Play around with the parameter (learning rate) and the input data, and verify for yourself what the Perceptron can and can not do
 - ▶ Make a movie of Perceptron converging, and one of Perceptron failing on the XOR.
- What else do you notice?
 - ▶ Is every solution the same? If not, are some “better” than others in some sense?

BREAK

Organizational items

- Ask questions!
 - 3-5 Homework projects (no Midterm): Oral presentation. 40% of grade.
 - Final project: Oral presentation, and written report. 40% of grade.
 - Final exam (multiple choice): 20 % of grade.
 - Form study groups and work in groups for HW and projects.
- *Student feedback welcome -- you can shape the course of the lectures by expressing your interests!*

how to contact me

- best: you can talk to me any time!
don't be shy
- second best: email
sstill @ hawaii . edu

Take a look at my website!

www2.hawaii.edu/~sstill/

where you can find all information,
including office hours, publications, open projects