

25/07/2017

WEAVING TOGETHER STATISTICAL MECHANICS, MACHINE LEARNING

SURYA GANGULI

AND NEUROSCIENCE

INTRODUCTION

High dimensional statistics \rightarrow assumption about simplicity of the data \leftrightarrow associated algorithms

Theoretical neuroscience \leftrightarrow machine learning successes, for which we know all of the ingredients
 \hookrightarrow could the method we use to understand these (still to be invented) paved the way of understanding the brain.

- OUTLINE:
1. Oldies but goodies
 2. High dimensional statistics: theory and experiments
 3. Deep learning: theory and practice


REF: STATISTICAL MECHANICS OF

LECTURE I OLDIES BUT GOODIES



Hodgkin Huxley \rightarrow quantitative description of action potential.

Hopfield \rightarrow an idea of memory

Hodgkin Huxley:




\rightarrow all information transported by this voltage.



if for any reason $V_{rest} \uparrow$
the channels open
 \hookrightarrow Na^+ rushes in

FAST POSITIVE FEEDBACK


⊕ POTASSIUM CHANNEL




if $V_{rest} \uparrow$, K^+ channels open slowly
 \hookrightarrow K^+ rushes out

SLOW NEGATIVE FEEDBACK

⊕ energy cost to being excitable to maintain the non equilibrium steady state




\rightarrow we have the minimal ingredients of an excitable system in biology:



→ The spike happening somewhere on the membrane will then propagate, through a wave along the membrane.

→ Increase of V happens at synapse:



↳ SPIKE FROM PRESYNAPTIC NEURONS

↳ Ca^{2+} rushes in,

↳ vesicles with neurotransmitter open and release

↳ neurotransmitter diffuses

↳ post synaptic neuron open neurotransmitter gated ion channels.

↳ complicated transmission, that actually contains many degrees of freedom that allow a very sensible fine tuning of the spikes → store memories for a lifetime?


→ Then to study the brain, we need to coarse grain... REF: Model Neuron: From Hodgkin Huxley to Hopfield (Alkon and Kepler 1990) -

+ time scale separation: spikes ~ 2ms.
cognition ~ 100 ms.


↳ INTEGRATE AND FIRE: "leaky integrator of membrane voltage"

$$\tau \frac{dV}{dt} = -V + I(t) \quad \rightarrow \text{if } V > V_{\text{threshold}} \text{ emit a spike, then go back to } V_{\text{rest}}$$


input current $I(t) < V_{\text{rest}}$



input $I(t) > V_{\text{threshold}}$



→ under constant input current response = constant rate of spikes



⇒ collective modelling: $I_i(t) = \sum_j W_{ij} \sum_{t_j} K(t - t_j)$


SYNAPTIC STRENGTH ← W_{ij}
time of spike of j ← t_j

Rk: can formulate mean field theory of integrate and fire networks → a bit complicated

↳ SPIKE RATES: one more step of coarsening.

r = rate = spikes per unit time:

Ex: compute analytically



→ origin of nonlinear activation function of artificial NN

$$\tau \frac{dr_i}{dt} = -r_i + \sum_j W_{ij} \psi(r_j^{\text{ext}}) + I_i^{\text{ext}}$$

Rk: fastest neuron 200 Hz

spontaneous random neuron 3 Hz

↳ SPIN NEURON

$$s_i(t+1) = \text{sgn} \left(\sum_j J_{ij} s_j(t) + h_i \right)$$

→ zero-temperature Glauber dyn of spin system —

Rk: Other separation of time scales with the learning, for which we have on much longer time scales plasticity = changes on the J_{ij} / w_{ij} .


HOPEFIELD MODEL

What is a memory? → pattern completion: associate name with picture —
stimulus with action —

↳ Store on N neurons, P memories:

memories: ξ_i^μ $\left\{ \begin{array}{l} i=1 \dots N \\ \mu=1 \dots P \end{array} \right.$
 $w(0,1)$ ✓
 RANBOLIID THEORIES

→ we wish that if initializing in ξ_i^μ , system will converge to ξ_i^μ cleaning up the corrupted memory.
 → image of basins of attraction
 → how to implement that?



↳ Hopfield realized that one needs a dynamics with decreasing energy to rule out the oscillations → ok Glauber

↳ How to have energy minimums $E = - \sum_{ij} J_{ij} s_i s_j$ close to ξ_i^μ

→ idea that neurons that fire together, wire together (Hebb):

$$\Delta J_{ij} = \xi_i^\mu \xi_j^\mu \text{ each time see a pattern}$$

$$\text{after seeing all the patterns: } J_{ij} = \frac{1}{N} \sum_{\mu=1}^P \xi_i^\mu \xi_j^\mu$$

$$E = - \frac{1}{N} \sum_{ij} \sum_{\mu} \xi_i^\mu \xi_j^\mu s_i s_j = - \frac{1}{N} \sum_{\mu=1}^P \left(\sum_i \xi_i^\mu s_i \right)^2 = -N \sum_{\mu=1}^P (m^\mu)^2$$

$$\text{with } E = -N \sum_{\mu=1}^P (m^\mu)^2$$

$$m^\mu = \frac{1}{N} \sum_i \xi_i^\mu s_i$$

magnetisation of μ^{th} memory = overlap

↳ looking good

↳ what about interference between memories? Is $m = (m^1, \dots, m^P) = (1, 0, 0, \dots)$ stable? ORDER PARAMETER

↳ replica calculation: 0 RETRIEVAL SPIN GLASS → $\alpha = P/N$ "loading!"

lowest free energy

also superposition of memories

↳ yet not too problematic, given an initialization close to a pattern —

Rk: Correlated patterns? pseudo inverse rule?

→ Is ξ_i^μ a fixed point of $T=0$ dynamics? sufficient constraint to be in retrieval phase!

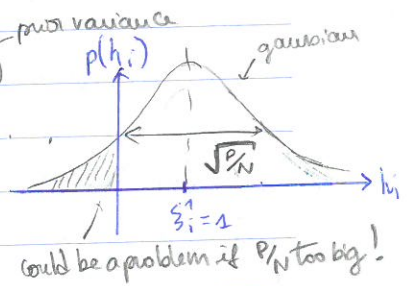
first pattern: $\xi_i^1 = \text{sgn} \left(\frac{1}{N} \sum_j \sum_{\mu=1}^P \xi_i^\mu \xi_j^\mu \xi_j^1 \right)$

$$= \text{sgn} \left(\underbrace{\frac{1}{N} \sum_j \xi_i^1 \xi_j^1}_{\xi_i^1} + \underbrace{\frac{1}{N} \sum_{\mu=2}^P \sum_j \xi_i^\mu \xi_j^\mu \xi_j^1}_{\text{interference of competing memories}} \right)$$

interference of competing memories

CLT $\sim W(0, \alpha)$


factor $\frac{1}{N^2} \times P \times N \times O(1)$
variance $\frac{1}{N^2} \times P \times N \times O(1)$
nb terms



$\Rightarrow \xi_i^1 = \text{sgn} \left(\underbrace{\xi_i^1}_{h_i^1} + \eta_i^1 \right)$ with $\eta_i \sim W(0, \frac{P}{N})$

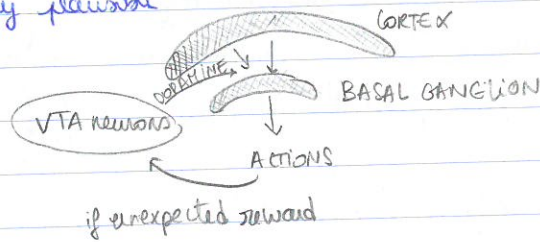
5/07/2017

ABOUT LEARNING (UNSUPERVISED / SUPERVISED)




Different learning tasks:

- unsupervised learning
 - supervised learning $x \rightarrow y \equiv y^*$
 - reinforcement learning not told y^* , but tell you whether good/bad direction
 - more biologically plausible
- THE TRAINING OF LIFE:



UNSUPERVISED LEARNING: what is Hebbian learning doing?



a single neuron $v = \vec{w} \cdot \vec{u}$ linear

Imagine streaming input pattern \vec{u}^α $\alpha = 1 \dots n$

And having weights changing according to Hebb rule: $\Delta \vec{w} = \lambda \vec{u}^\alpha v^\alpha$

biological assumptions to the behavior of a neuron.

small lambda (slow plasticity)

\hookrightarrow after n inputs: $\vec{w}(n) = \vec{w}(0) + \sum_{\alpha=1}^n \lambda \vec{u}^\alpha v^\alpha \rightarrow \tau \frac{d\vec{w}}{dt} = \langle \vec{u} v \rangle_P(\vec{u})$

\hookrightarrow plug-in $v = \vec{w}^T \vec{u} \rightarrow \tau \frac{dw_i}{dt} = \left\langle \sum_k w_k u_k u_i \right\rangle = \sum_k w_k \underbrace{\langle u_k u_i \rangle}_{Q_{ki}} \Rightarrow \tau \frac{d\vec{w}}{dt} = \vec{Q} \vec{w}$

learning driven by second order correlation for linear neurons!
 \hookrightarrow of $P(u)$