Temporal Dynamics of Learning, San Diego, February 2009

Levels, Time and Models

Tony Bell
Redwood Center for Theoretical Neuroscience
University of California at Berkeley

I want to connect these two things -

Theory of Probabilistic Learning ⟷ Empirically observed structure of living matter

in order to understand what learning *really* is.

Observation:
Parts of objects get summaries of the activities of parts of other objects:

ie: Objects summarise and broadcast:

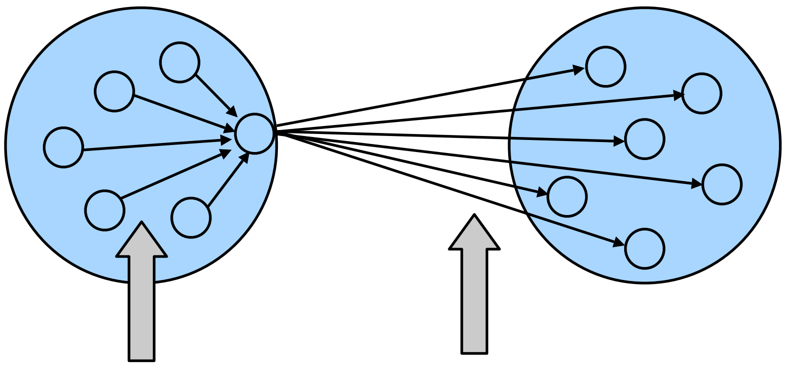

Emergence and s*ubmergence* of information are simple logical consequences of modular structure.

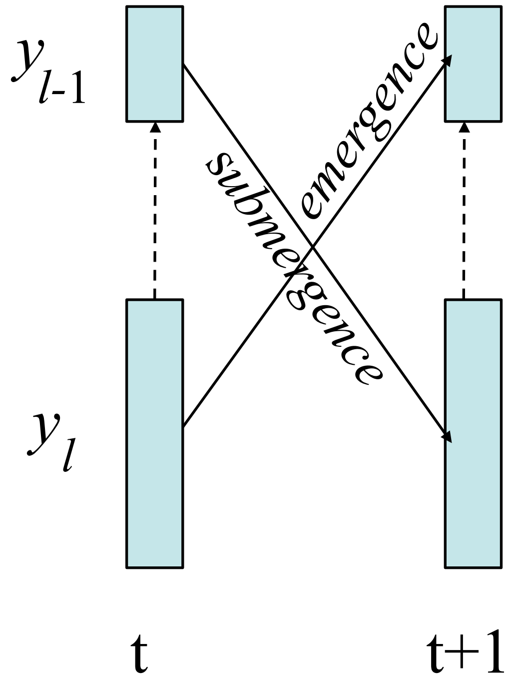

ie: Objects summarise and broadcast:

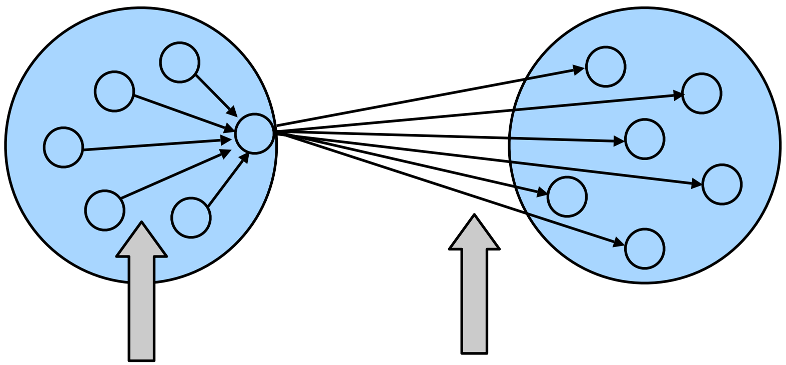

Emergence and s*ubmergence* of information are simple logical consequences of modular structure.

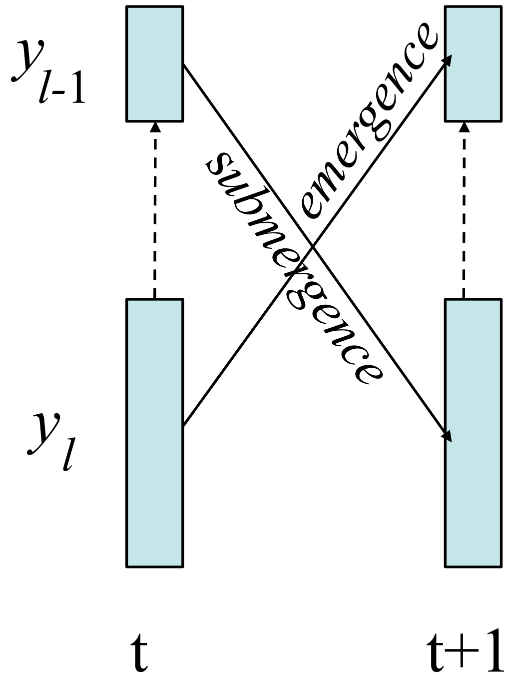

INTERACTIONS ARE NOT HORIZONTAL

Neurons don't talk to neurons, *they talk to synapses*.

This is what I mean by submergence.

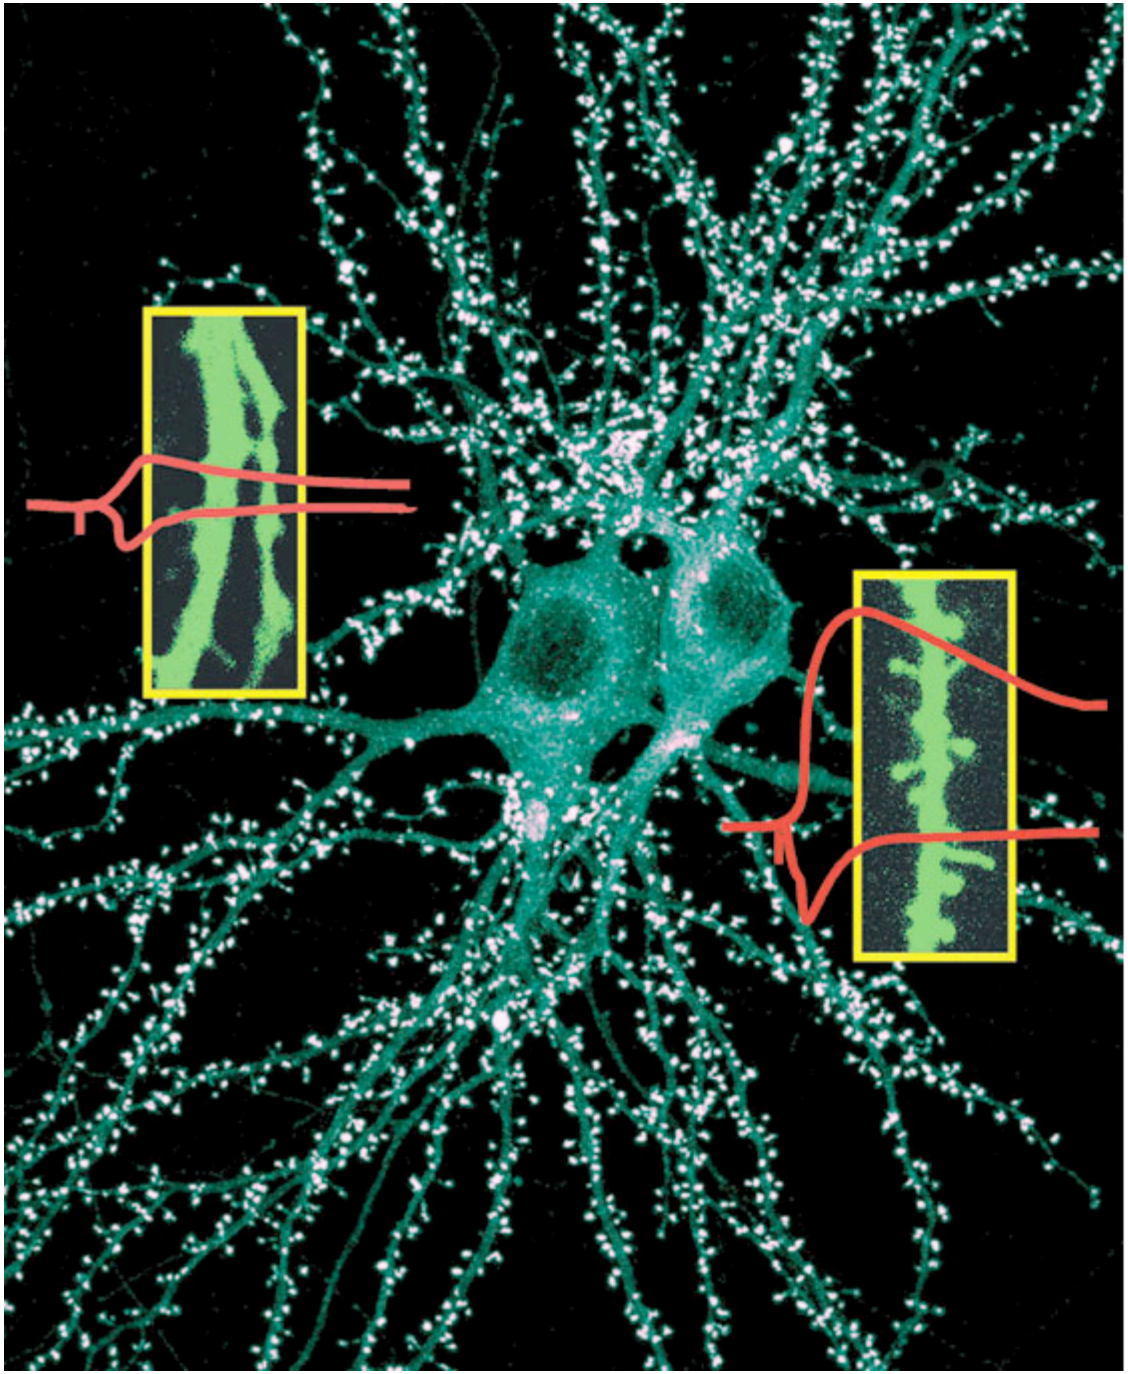

Figure: A hippocampal neuron with synapses stained for post-synaptic proteins Shank and Homer (white puncta). Overexpression of dominant negative form of Homer (Homer1a) causes loss of dendritic spines and suppression of postsynaptic responses. *Picture by Carlo Sala.*

Examples: Submergence from neurons to synapses via spikes.
Submergence from synapses to macromolecules via EPSP's.
(In both cases into massively overcomplete storage spaces.)

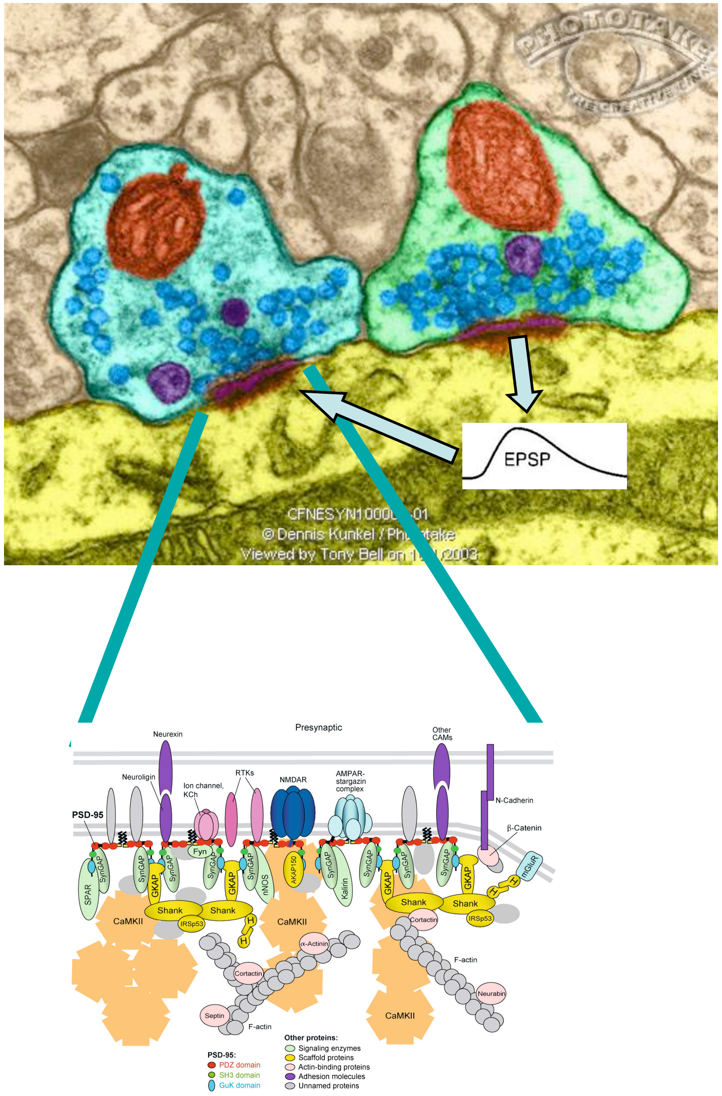

Sheng M, Hoogenraad CC. 2007.
Annu. Rev. Biochem. 76:823–47

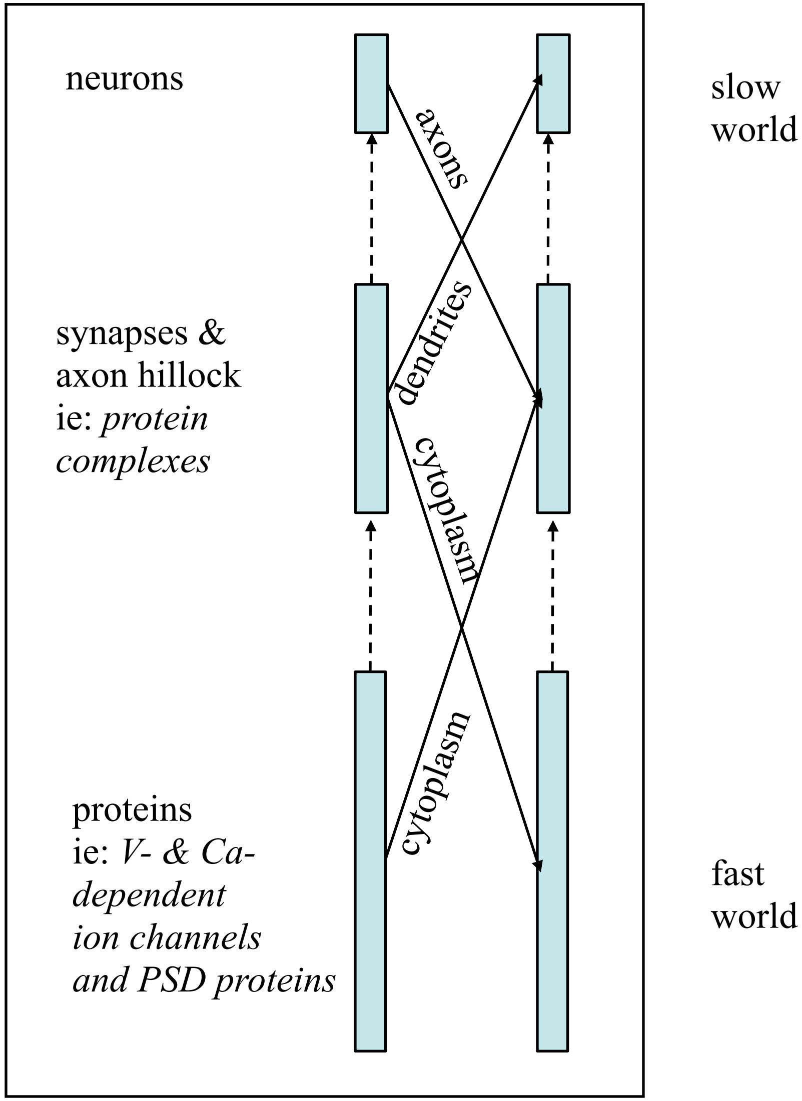

And it continues: calcium plays a role inside the synapse like that played by voltage in the cell as a whole.

calcium microdomain

20nm

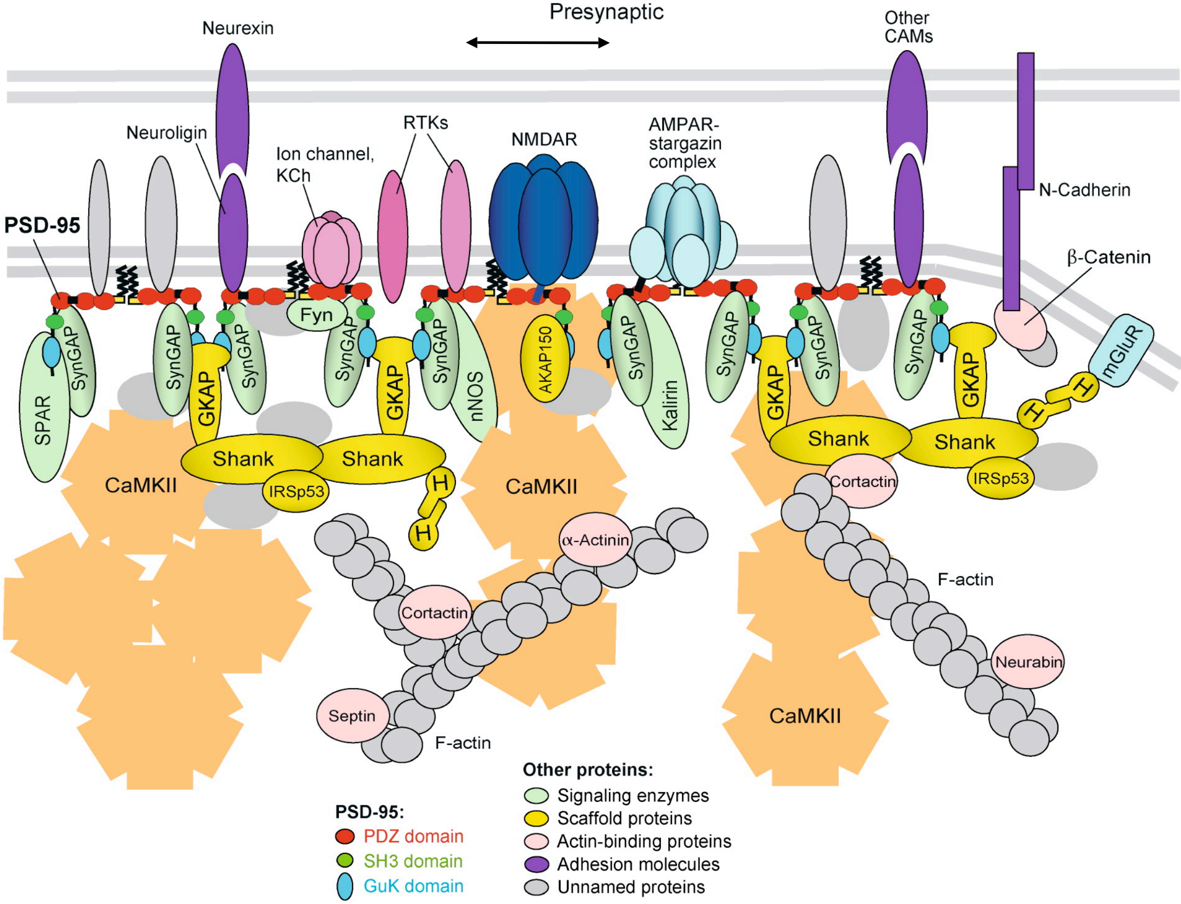

Sheng M, Hoogenraad CC. 2007.
Annu. Rev. Biochem. 76:823–47

(image from Morgan Sheng)

Some synaptic plasticity researchers are coming to terms with this:

Cell PRESS

Neuron
Perspective

Ubiquitous Plasticity and Memory Storage

Sang Jeong Kim[1,*] and David J. Linden[2,*]
[1]Department of Physiology, Seoul National University College of Medicine, 110-799 Seoul, Korea
[2]Department of Neuroscience, Johns Hopkins University School of Medicine, Baltimore, MD 21205, USA
*Correspondence: dlinden@jhmi.edu (D.J.L.), sangjkim@snu.ac.kr (S.J.K.)
DOI 10.1016/j.neuron.2007.10.030

To date, most hypotheses of memory storage in the mammalian brain have focused upon long-term synaptic potentiation and depression (LTP and LTD) of fast glutamatergic excitatory postsynaptic currents (EPSCs). In recent years, it has become clear that many additional electrophysiological components of neurons, from electrical synapses to glutamate transporters to voltage-sensitive ion channels, can also undergo use-dependent long-term plasticity. Models of memory storage that incorporate this full range of demonstrated electrophysiological plasticity are better able to account for both the storage of memory in neuronal networks and the complexities of memory storage, indexing, and recall as measured behaviorally.

Going further down: the energy landscapes of macromolecules are functionally coupled to that of water.

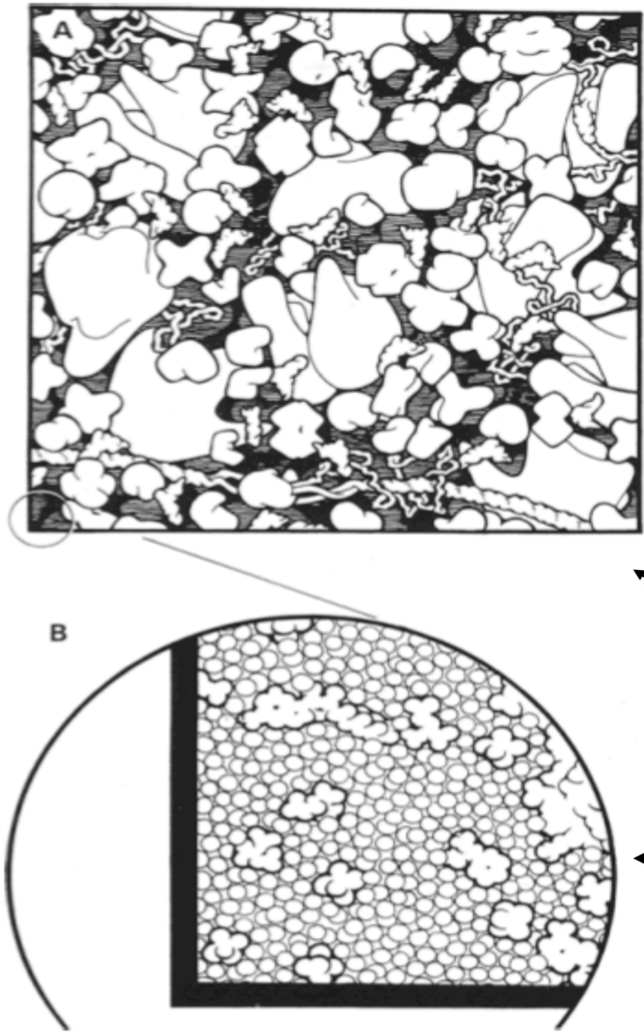

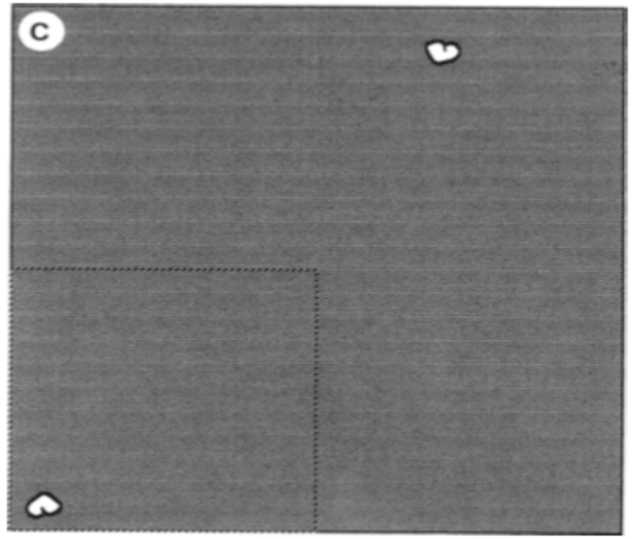

100 nm x 100nm
proteins 10nm apart

small solute molecules in water

Hoppert & Mayer 99

Water forms clusters:
(keyword: “water clusters”)

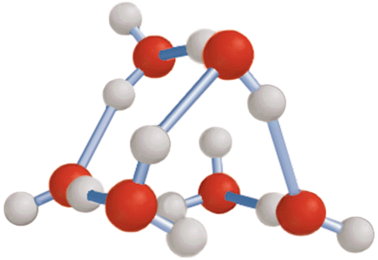

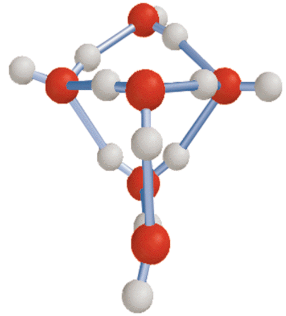

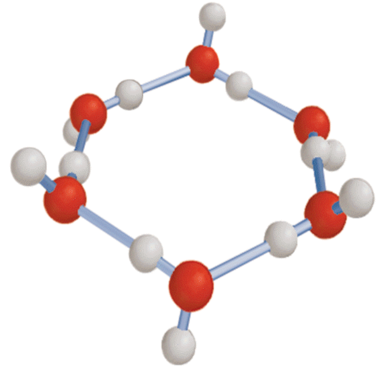

based on dynamic switching of hydrogen bonds

There is elasticity: the molecules are 'pulled out of shape':

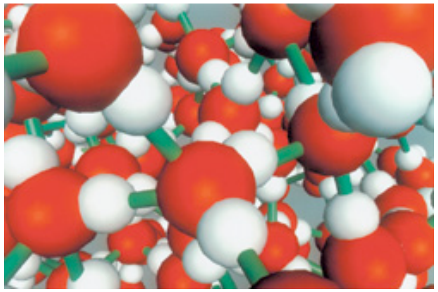

→

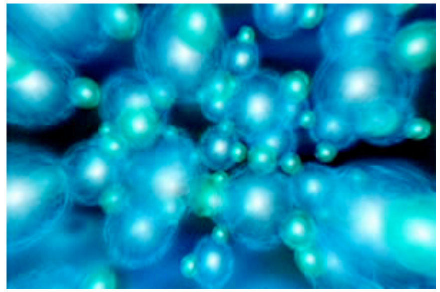

and there are domains of all sizes:

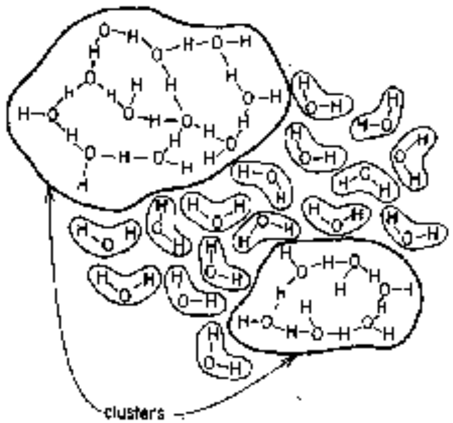

These last for several picoseconds and are implicated in the stabilisation of and signalling between macromolecules

“The many roles water plays in biomolecular processes, and particularly the coupling between its motions and the dynamics of proteins and nucleic acids, are currently widely acknowledged.”
- Y.Levy and J.N.Onuchic *Annu. Rev. Biophys. Biomol. Struct*. 2006. 35:389-415

Above the level of the cell: oscillating assemblies

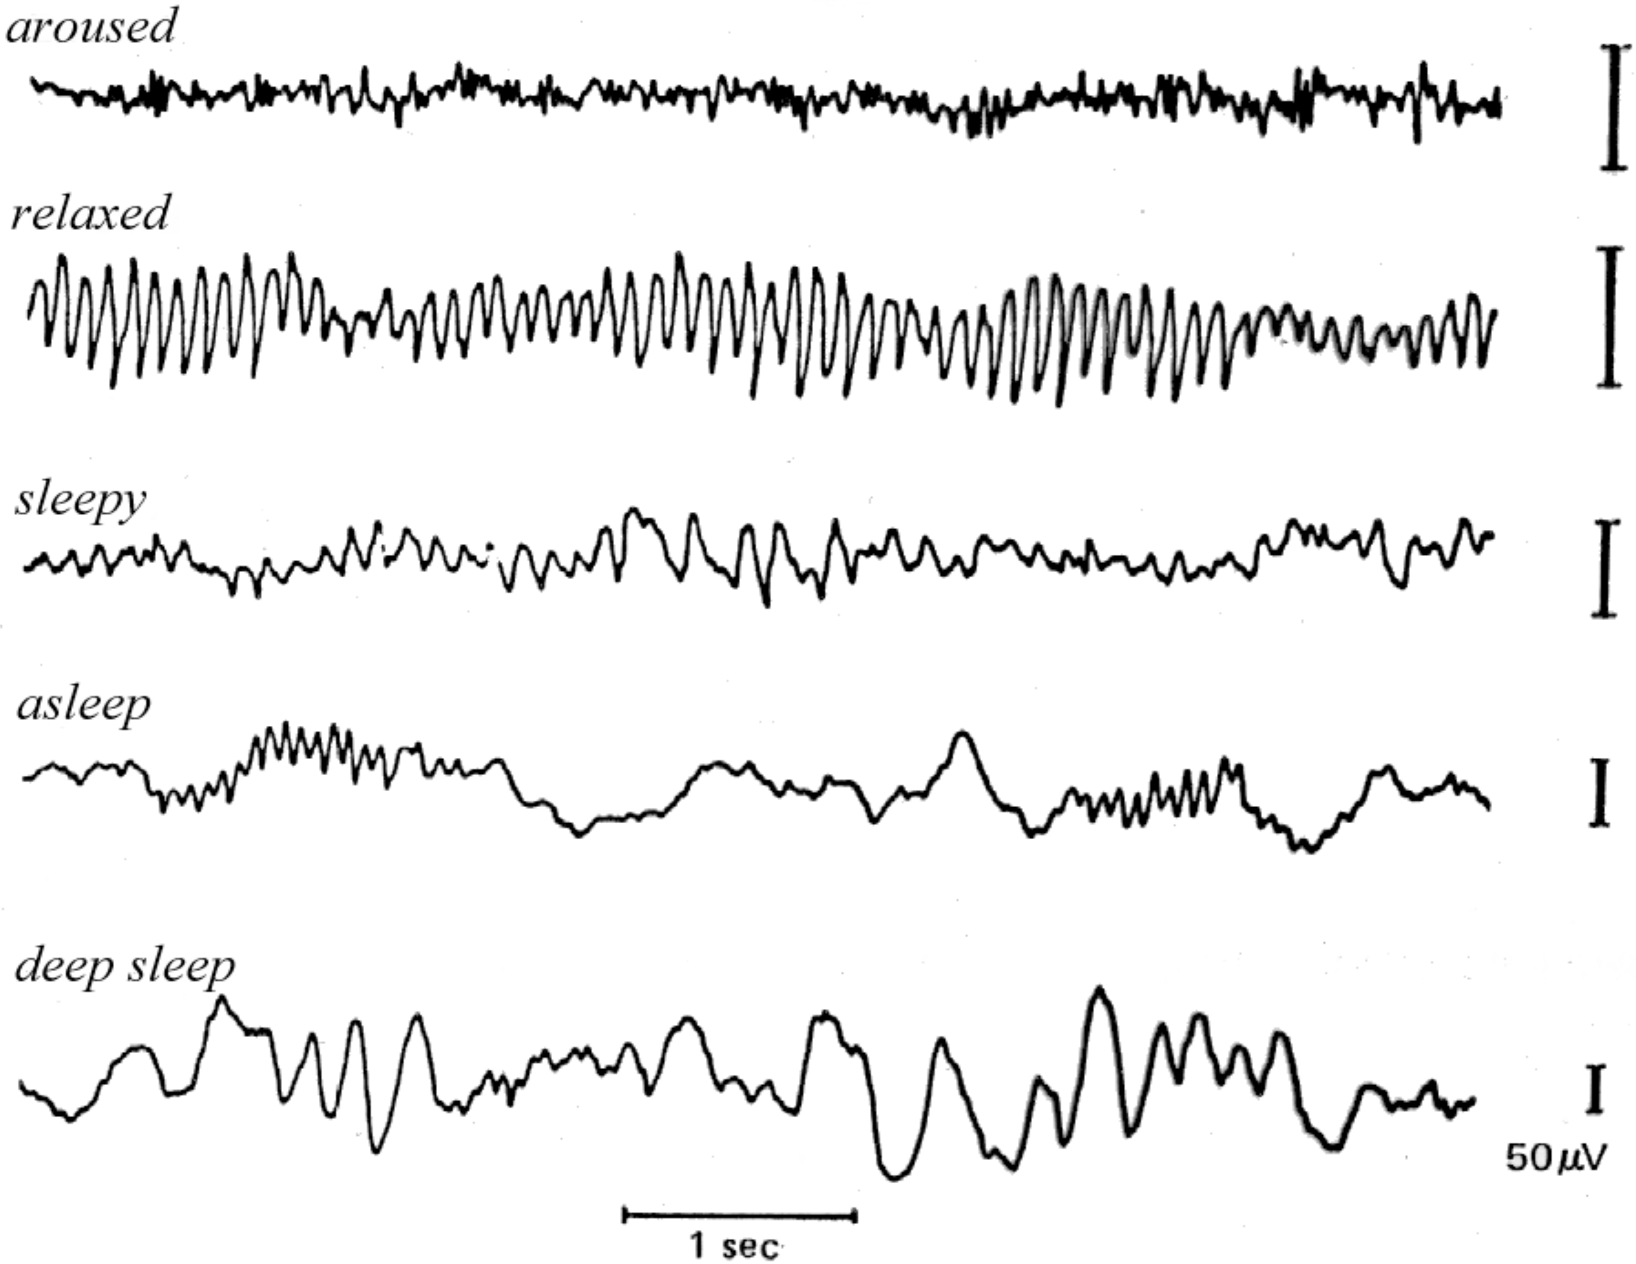

Emergence through *synchrony*, submergence through *phase amplitude coupling*.

ie: *large-scale cell assemblies* map into an overcomplete space: *small assemblies* (Lakatos, Schroeder, Canolty)
ie: *small-scale cell assemblies* map into an overcomplete space: *neurons* (Fries, Koepsell)

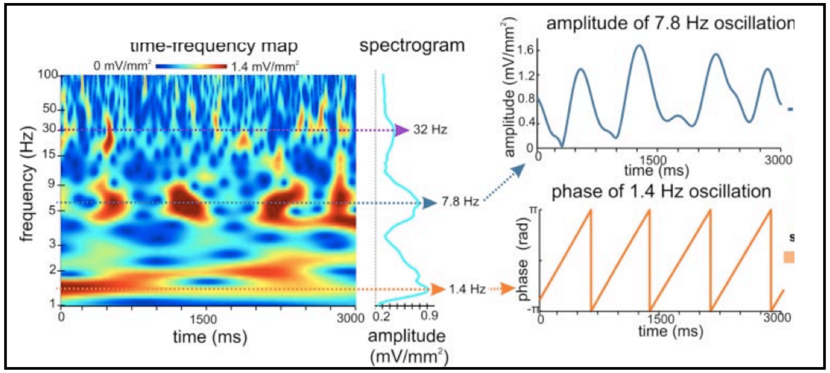

delta to theta coupling *(Lakatos et al)*

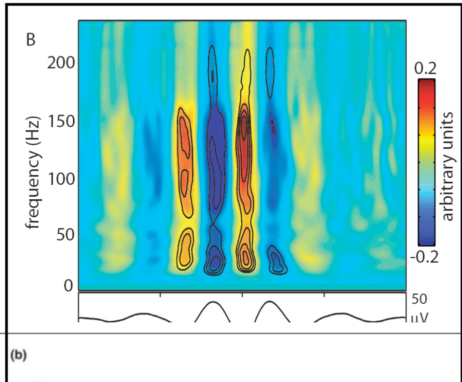

theta to hi-gamma coupling *(Canolty et al)*

gamma to spike coupling *(Fries et al)*

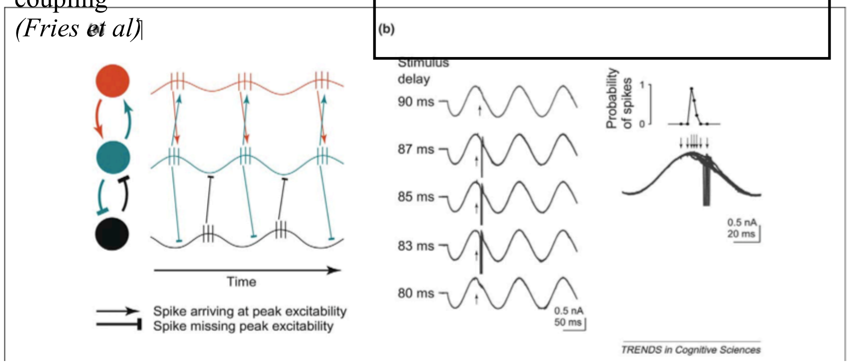

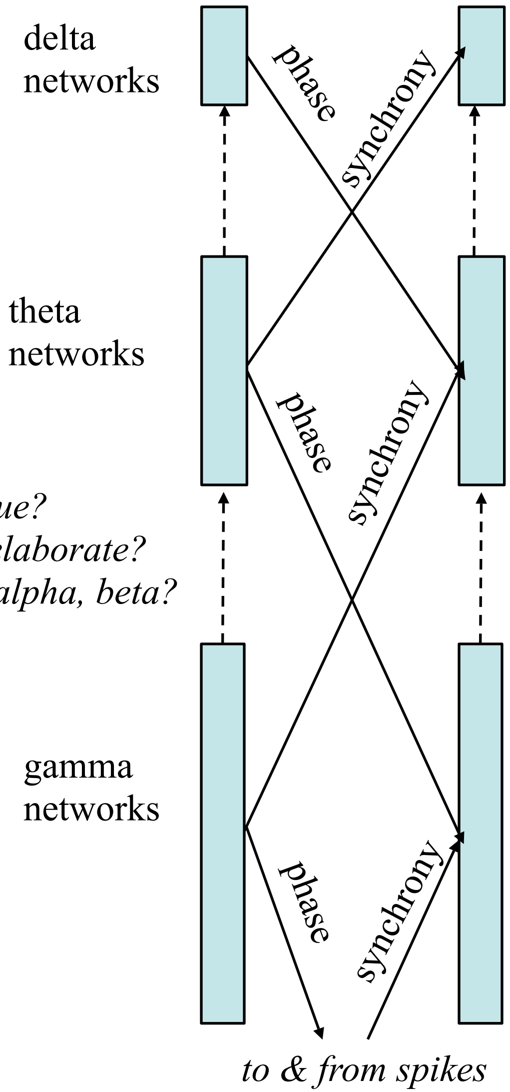

is this true?
can we elaborate?
what of alpha, beta?

Global low-frequency oscillations *index into* more local higher frequency oscillations (that is: they set up space-time patterns that allow only certain more local patterns to occur).

The *retrieval* is then the result of the more local computation,which is passed back up to more global networks.

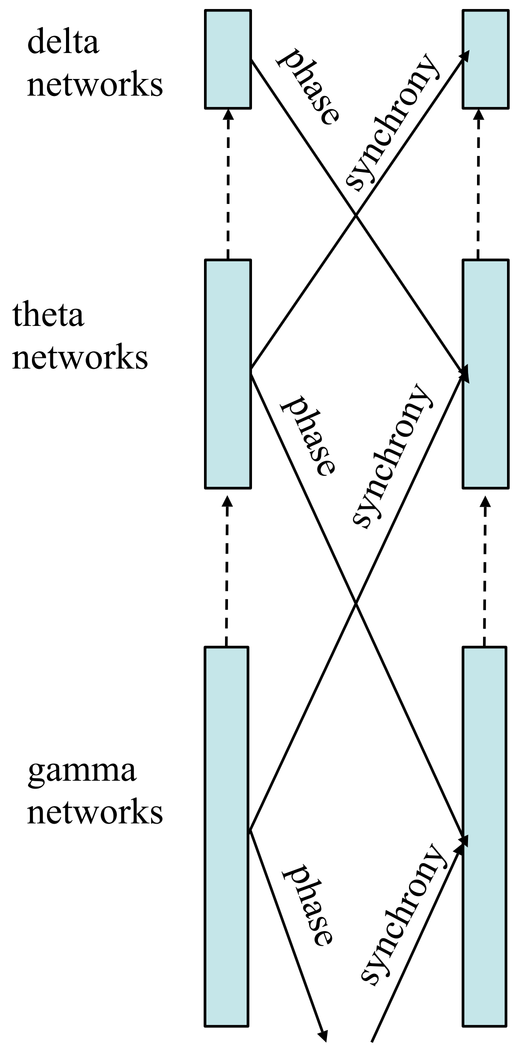

CLAIM:

The adaptive power of living systems comes from the gating of information flows across levels, and no horizontal models have this power.

Horizontal models are fine for fleshing out the rungs of the ladder, but nothing will ever move up and down the ladder.

(How are memories stored and retrieved?)
(How are actions or decisions generated?)
(What is the mechanism of learning?)
(What happens when a thought occurs to you?)

Think about these questions in the light of this claim.

Emergence:

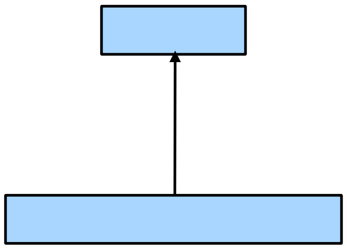

an undercomplete mapping into a lower dimensional space for the purpose of sending messages in a more macroscopic network

Submergence:

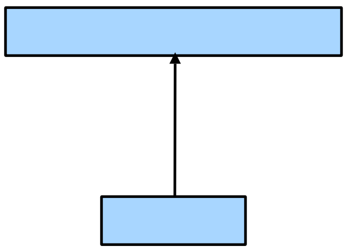

an overcomplete mapping into a higher dimensional space for the purpose of sending messages in a more microscopic network

This picture of the storage and retrieval roles of the microscopic is a fundamental challenge to computational and 'effective physical' theories.

Horizontal interactions only occur in anomalies like the computer, and other ***models****, in which emergence from the microscopic is screened out:*

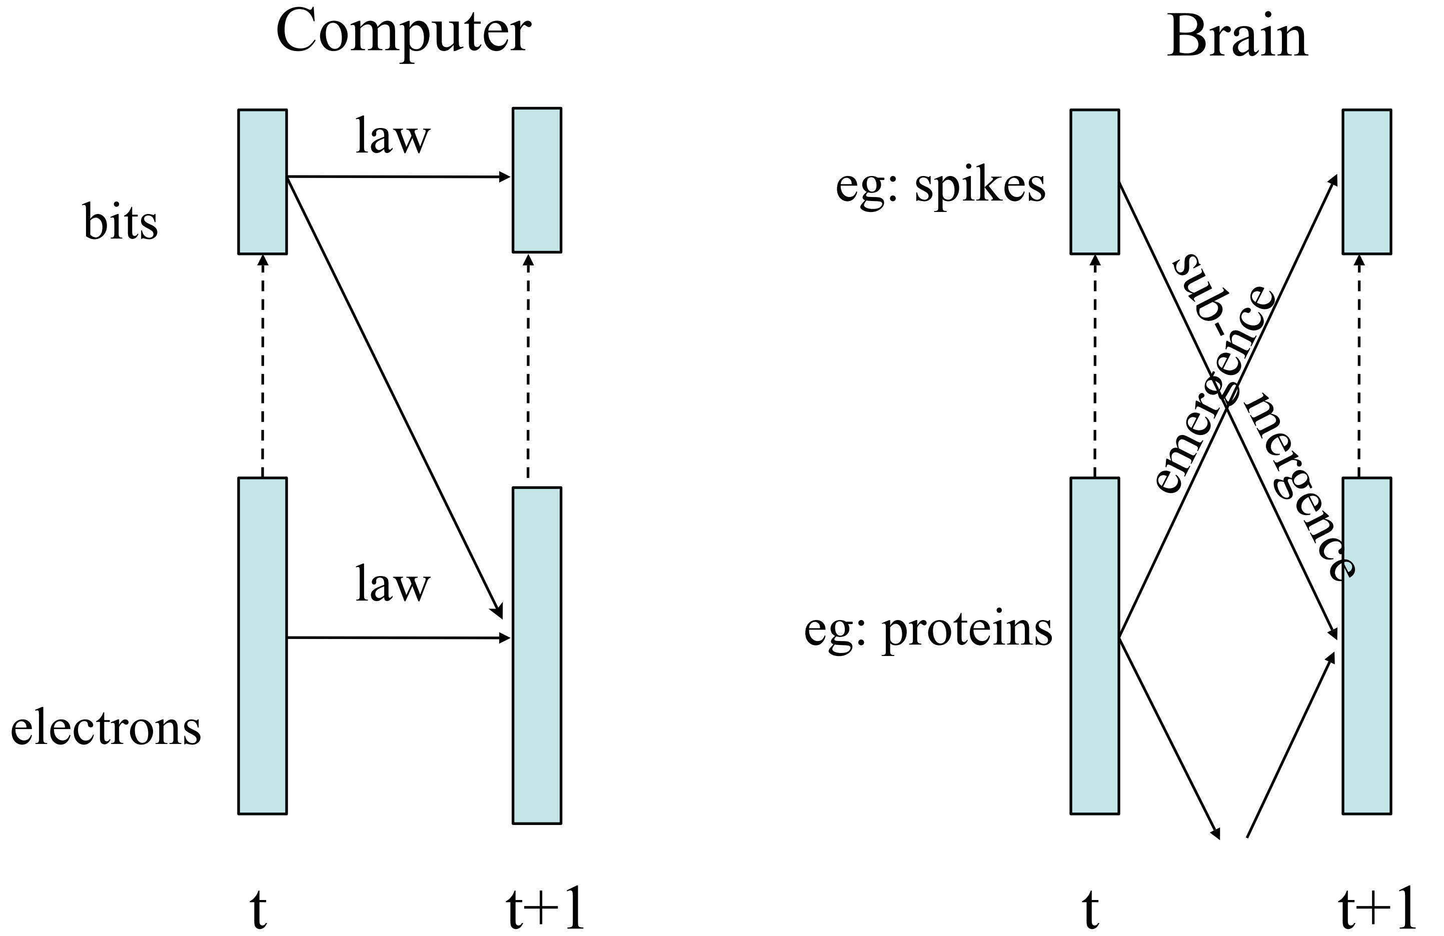

(minimal) deterministic causal relation. deterministic structural relation

Levels determined by Minimum Adequate Description criterion

Living systems are:
all the way down and
all the way up:

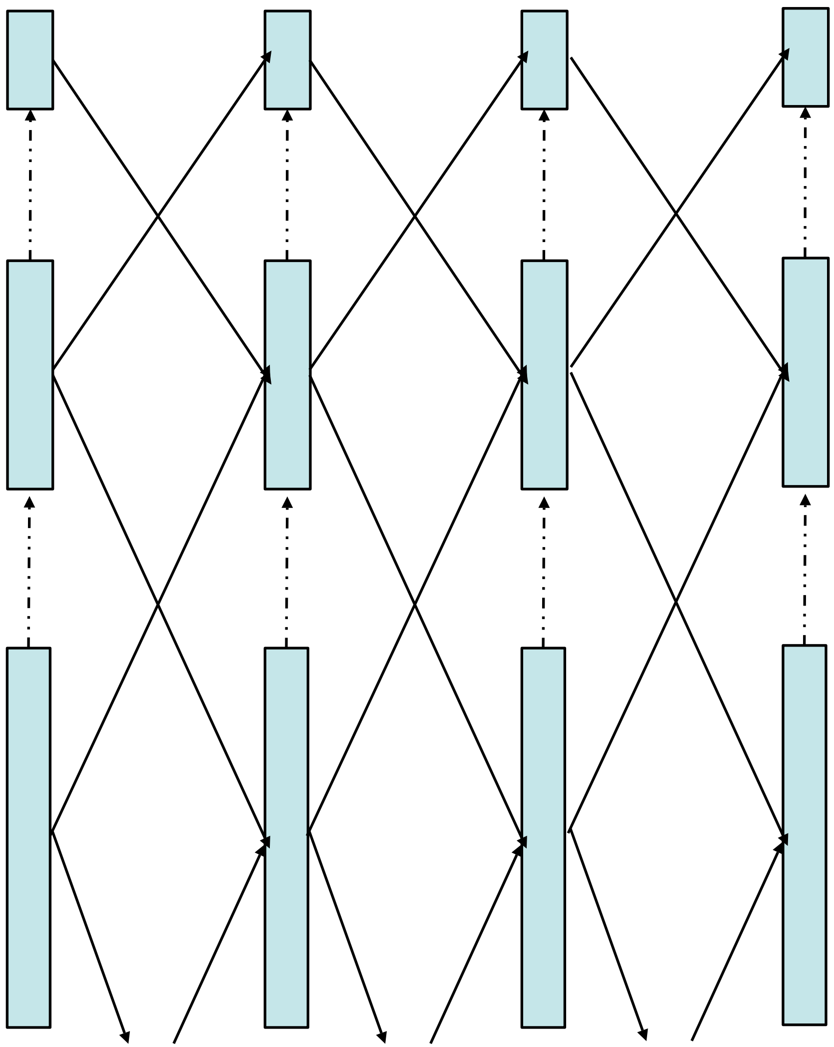

with each interaction taking the form of the outside of an object interacting with the inside to create the local space-time-level state:

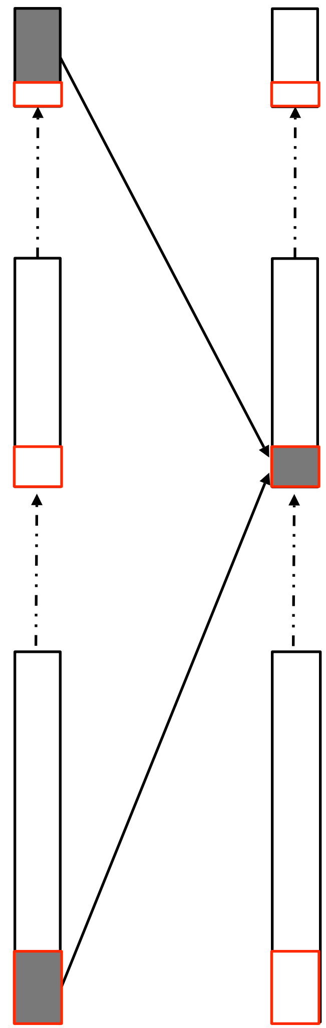

Q. How to get to grips with this theoretically?
A. Through linking causality to probabilistic models.

Remember deterministic Infomax
(in which the transformation IS the model):

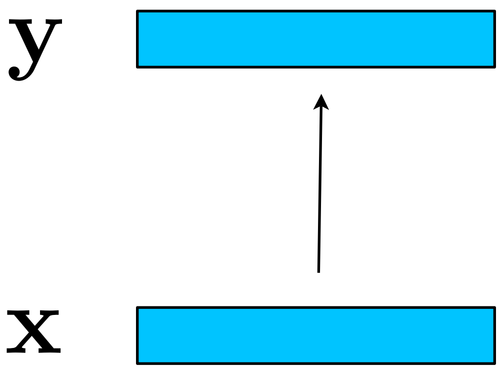

model on outputs (= unit p.d.f.)

Jacobian of transformation

model on inputs

$$q(\mathbf{y}) \left| \frac{\partial \mathbf{y}}{\partial \mathbf{x}} \right| = q(\mathbf{x})$$

What if we were to generalise this to:

“everywhere there is causality, we put a piece of probabilistic modeling, and then we combine them all together to make a big probabilistic model over all time, space and levels”.

Q. What is a piece of causality?
A. It is an object going " How likely was that?"

ie: given my internal information, how likely are these external things that are affecting me?

For this we need the overcomplete, conditional density estimation generalisation of Infomax,

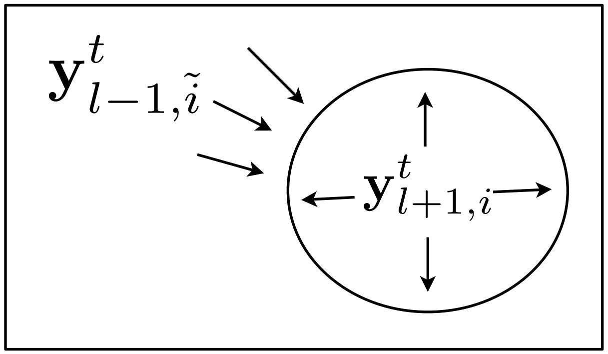

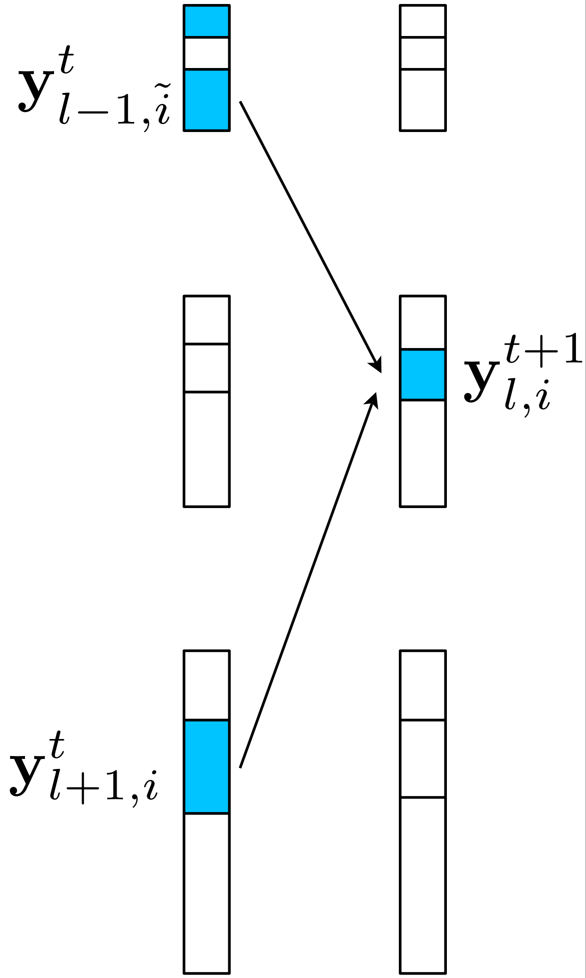

which looks like this:

$$q(\mathbf{y}^{t+1}_{l,i} | \mathbf{y}^t_{l+1,i}) \left| \frac{\partial \mathbf{y}^{t+1}_{l,i}}{\partial \mathbf{y}^t_{l-1,\tilde{i}}}^T \frac{\partial \mathbf{y}^{t+1}_{l,i}}{\partial \mathbf{y}^t_{l-1,\tilde{i}}} \right|^{1/2} = \underbrace{q(\mathbf{y}^t_{l-1,\tilde{i}} | \mathbf{y}^t_{l+1,i})}_{\text{How likely was that?}}$$

How do these combine? - think about a system in time.....

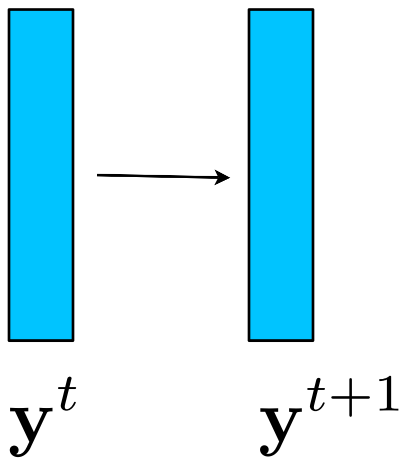

1 $$\frac{d\,\mathbf{W}}{d\,t} \propto \partial_{\mathbf{W}} \log q(\mathbf{y}^t)$$ Gradient Learning

2 $$q(\mathbf{y}^{t+1}) \left| \frac{\partial \mathbf{y}^{t+1}}{\partial \mathbf{y}^t} \right| = q(\mathbf{y}^t)$$ Transformation of densities

$$\Big\downarrow \; \partial_{\mathbf{W}} \log(\cdot)$$

$$\partial_{\mathbf{W}} \log q(\mathbf{y}^{t+1}) - \partial_{\mathbf{W}} \log q(\mathbf{y}^t) = -\partial_{\mathbf{W}} \log \left| \frac{\partial \mathbf{y}^{t+1}}{\partial \mathbf{y}^t} \right|$$

This is just a discrete time log likelihood derivative!

So: $$\boxed{\frac{d^2\mathbf{W}}{dt^2} \propto -\partial_{\mathbf{W}} \log \left| \frac{\partial \mathbf{y}^{t+1}}{\partial \mathbf{y}^t} \right|}$$

Weight acceleration is *model-free*!

20

Implications

Single level models can't capture inter-level information flows and thus can't explain memory, thought, perception or action. Horizontal theories are essential, but *wrong*!

There is *no cutoff level* but rather submergence all the way down and emergence all the way up.

Agent-centred theories with concepts like *reward* and *fitness* are wrong because behaviours are just meso-scale emergences (this has been faced in Evolutionary Theory by Multi-Level Selection Theorists).

Uncertainty in a model should never be confused with "noise-in-the-system", the latter being a completely undefined concept (which should be abolished).

Prospects

Progress towards mathematically rigorous, concrete, implemented, example models of adaptive emergence and submergence has been frustratingly slow, though with some flashes of insight.

But if we can work this out, it will be big.