bias^2

$$= \cancel{\mathbb{E}_D}\left[\left(f(x_j) - \mathbb{E}_D\left[g(x_j;\hat{\theta}_D)\right]\right)^2\right] +$$

$$+ \mathbb{E}_D\left[\left(g(x_j;\hat{\theta}_D) - \mathbb{E}_D\left[g(x_j;\hat{\theta}_D)\right]\right)^2\right] + \text{cross terms} \to 0$$

because mean is equal to real mean.

variance

how similar is the model we learn every time to the average model you learn

does not depend on the data on the training

$$\mathbb{E}[(C)] = \text{noise} + \text{bias}^2 + \text{variance}$$

noise ↳ σ_ε^2

how different is the true model f from the average model you learn.

blue graph is $C_{in}(\{y\},\{g(\vec{x},\theta)\}) = \sum_i (y_i - g(\vec{x}_i,\theta))^2$ → which gives you $\hat{\theta}_D$

red one is the Cost in $\mathbb{E}_{D,\varepsilon}[\text{Cost}(y, g(x,\hat{\theta}_D))]$

So what you care about is minimising Cost and you need to remember about noise, bias^2 + variance.

The exact definition of noise, bias^2 and variance only holds for Gaussian noise models but you get similar definitions for non-quadratic cost functions.

(Friday start here)

Eleni lecture 4

So for optimization

$$\hat{C} = \sum \frac{Q_e}{C_e} + \lambda \sum C_e^\gamma$$

$\gamma < 1$  $\gamma > 1$ 

Now lets focus on how we get this and look at Development. (so growth)

Lets look at the Murray model.

But make it look like energy.

$$E = \sum \frac{Q_e^2}{C_e} + C_0 \sum C_e^{\gamma}$$ (same as previous eqn with notation changes)

If the network is developing $\frac{dC_e}{dt} \neq 0$ (growing)

$$\frac{dE}{dt} = -\sum \frac{Q_e^2}{C_e^2}\frac{dC_e}{dt} + C_0\gamma \sum C_e^{\gamma-1}\frac{dC_e}{dt} + \sum \frac{2Q_e}{C_e}\left(\nabla_C Q_e \cdot \frac{dC}{dt}\right)$$

(this term much smaller)

We do not consider topological network rearrangements, links are allowed to appear and disappear but not do this

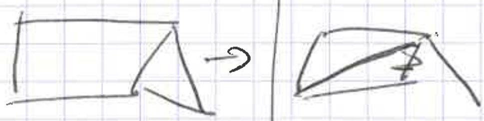

These are called adaptive networks → a whole field

$$= \sum\left(-\frac{Q_e^2}{C_e^2} + C_0\gamma C_e^{\gamma-1}\right)\frac{dC_e}{dt}$$

This term comes from the fact that Q_e is a function of C_e.

Life can choose an adaptation term that $\frac{dC_e}{dt}$ looks like:

adaptation term → steady state $Q_e \sim C_e^{\frac{\gamma+1}{2}}$ motivates this

$$\frac{dC_e}{dt} = +\frac{Q_e^2}{\gamma C_e^2} + C_0\gamma C_e^{\gamma-1}$$

(use it / loose it)

then we can write $\frac{dE}{dt}$ in this form. (because it the same thing as in the bracket above)

$$\frac{dE}{dt} = -\sum(\quad)^2$$

Sometimes (determined by the conductances) the system decides to use them a lot, then they grow, and they are used even more and so on

→ so an instability

Conversely, if you don't use it, it decreases.

General form of the adaptation term.

$$\frac{dC_e}{dt} = \beta Q_e^{g} - \alpha C_e \quad , \quad g = \frac{2}{\gamma+1}$$

in complex flow systems, check [illegible] [illegible] related to vessel [illegible] (Poiseuille etc)

this describes the architecture, the type of flow etc.

we transferred the C_e exponent to g (took it out).

Remember Q_e is a function of C_e (it can be any function),

Go back to Poiseuille flow

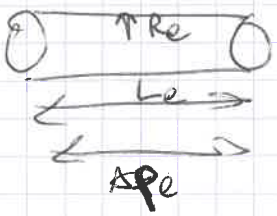

$$Q_e = \frac{\pi R_e^4}{8 L \mu} \Delta P_e$$

under control of pertains to this particular vessel.

$\Rightarrow C_e$ — conductance

The vessels adapt by changing their diameter to increased shear stress. $P_i > P_o$

Shear stress

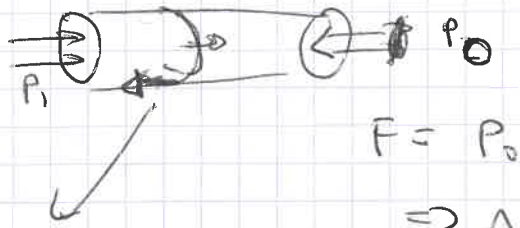

$F = P_i \pi R^2$

$F = P_o \cdot \pi R^2$

$\Rightarrow \Delta F = \Delta P \pi R^2$

to balance the flow

the wall exerts friction.

$\Delta F = \tau 2\pi R L$

↑ shear stress

Force balance: $\Delta P \pi R^2 = \tau 2\pi R L \Rightarrow$ shear stress $\boxed{\tau = \frac{4\mu Q}{\pi R^3}}$

What people have measured (a way to test that vessels adapt to shear stress):

$$(*)\ \frac{dR_e}{dt} = \Theta_1 \left(|\tau_e|^\sigma - \tau_d^\sigma \right)$$

Θ_1: constant parameter; $|\tau_e|$: actual value; τ_d: desired value.

expect that if a fitting parameter

$R_e \sim C_e^{1/4}$

$$\frac{dC_e}{dt} = \Theta'' \left(\frac{|Q_e|^\sigma}{C_e^{3\sigma/4} - 1} - \tau_d^\sigma C_e \right)$$

This agrees with theory. Vessels optimize cost and dissipation by optimizing (changing) shear stress.

(*) This is an exponential eqn. It means that the vessel mechanically would like to be receiving a specific shear stress. Because if its less you would recycle that vessel, waste of material. If its too much it can burst. Vessels adapt, but it makes the decision to grow or decrease based on local flow, local direction, local <u>shear stress</u>. This allows it to find a local minimum that affects the total system, because its hydrodynamically coupled. The vessel does <u>not</u> measure global flow, only local.

Question of timescales (these models assume all change is instantaneous), the other question is about spatial distance → how far away do you feel a perturbation → complicated question.

(If you measure the pressure in your ankle can you figure out you had a stroke? → not right now with the current models).

Lets look at some limits:

$g \to 0$, $\gamma \to \infty$: steady state $C_e = \beta/\alpha$
large pipes are expensive — here you have loads of links and small links are used.

$g \gg 1$ $\gamma \to 0$; $\frac{dC_e}{dt} = V - \alpha C_e \Rightarrow$ exponential approach to $C_e = V/\alpha$
large pipes cheap — you get trees, you start deleting links below a certain threshold — every link used, much more economic design.

These optimisation eqn are also time optimisation eqns → they are solutions to MFPT eqns on graphs, and for random walks.

Topology and graph theory
(some generally applicable tools),

This also maps into other networks,

There is an exact mathematical analogy between spring networks and what we've seen before, components in 1D but it generalizes.

Flow	Electrical	Mechanical
pressure p_i	voltage V_i	displacement x_i
current Q_{ij}	electrical current I_{ij}	spring force F_{ij} defined as the tension elongation Δx (strain)
pressure drop Δp_{ij}	voltage drop ΔV_{ij}	
net current q_i	net current q_i	external forces or masses f_i
conductance C_{ij}	conductance C_{ij} (Conductor)	spring constant k_{ij}
pressure drop flow $C\Delta p = Q$	Ohm's law $R = U/I$	$F = k\Delta x$
current conservation $\sum Q_i = q$	current conservation	conservation of force $\sum F_i = 0$
dissipation $\frac{Q^2}{C} = \Delta p^2 \cdot C$		$k\Delta x^2 = E$ energy stored in a mechanical network.

So we can take this code and count the energy stored in a mechanical network etc,

So a current dipol becomes a mechanical dipole,

Engineers knew this in the 1960s.

(We still need capacitors,

They could not solve mechanical systems. They tested mechanical systems on electrical prototypes (because it was cheaper),

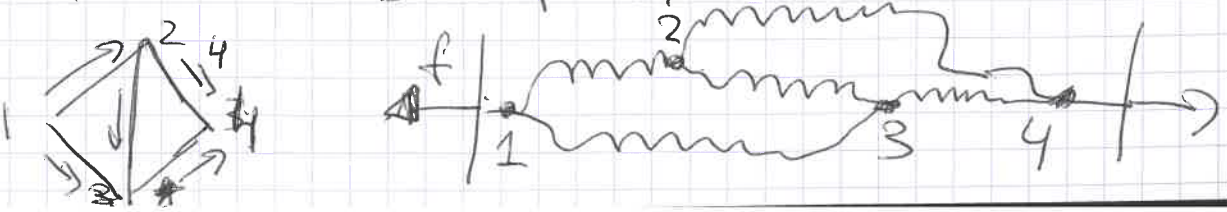

Random walks on graphs

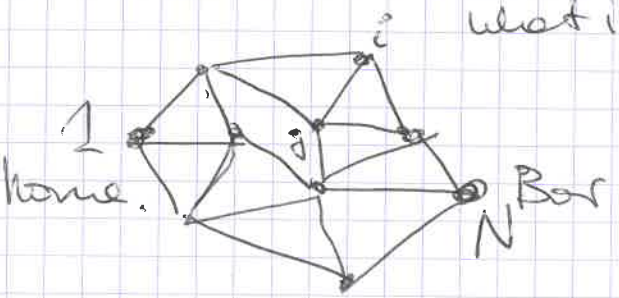

what is the prob that a drunken sailor (expects home before getting home

P_i = prob that a random walker is present on node i at timestep k.

At each iteration of the random walk the walker will move to a neighboring node with prob $W_{ij} = \frac{a_{ij}}{\sum_i a_{ij}}$

Each link has a weight a_{ij} (how attractive the street is), considered by the neighbors of that node.

(*) $\vec{P}^{(k+1)} = \hat{D}^{-1} \hat{A} \vec{P}^{[k]}$

$\vec{P}^{[k]}$ prob of walker at each node

$\hat{D} = \text{diag}\left\{ \sum_i a_{ij} \right\}$ sum all the links that are connected to node i

weighted adjacency matrix a_{ij}.

$(\hat{D}^{-1}\hat{A})$ will be related to the discrete Laplacian → see next page

Boundary condition: the walker stops at home, the prob that you will be home at the end is 1, $P_1^\infty = 1$.

If you start at the bar, $P_N^\infty = 0$ you will never make it home.

So now we add boundary conditions.

If you know the P_i of all your neighbors, and you know P_j^∞ then your P_i^∞:

$P_i^\infty = \sum_j W_{ij} P_j^\infty$

is the weighted sum of all your neighbors.

$\Rightarrow \vec{P}^\infty = \hat{D}^{-1} \hat{A} \vec{P}^\infty$ ← so this is the steady state of (*)

$$\hat{D}\vec{P}^{\infty} = \hat{A}\vec{P}^{\infty} \Rightarrow \underbrace{\left(\hat{D}-\hat{A}\right)}_{\hat{L}\leftarrow \text{discrete Laplacian}}\vec{P}^{\infty} = 0$$

P_i^{∞} ← is the prob. that you end up at node i before you get to N. So this is solved by the backward eqn. But A is symmetric so $W^T = W$, (Nodes are bi directional).

$$\hat{L}\vec{P}^{\infty} = 0$$

Travel time:

How long on average does it take to go from node i to to node 1 (expected # of steps).

If $t(v_i, 1) = t_i$

this from node i to node 1

If I knew t_i that would be equal to the travel time from my neighbors + 1 (for one additional step).

$$\vec{t} = \underbrace{\hat{D}^{-1}\hat{A}}_{\hat{W}\leftarrow \text{one of } \hat{W} \text{ if } (N-1)\times(N-1) \text{ because we eliminate node 1.}}\vec{t} + \vec{1}$$

this does not apply for node 1.

$$\Rightarrow \underbrace{\left(\hat{D}-\hat{A}\right)}_{\hat{L}}\vec{t} = \underbrace{\hat{D}\vec{1}}_{-\vec{q}} \Rightarrow \boxed{\hat{L}\vec{t} = -\vec{q}}$$

so this is an eqn. we've seen. But now we have $\vec{q}$.

Before we saw $L\vec{p} = -\vec{q}$.

But there is a subtlety now $\vec{q}(\vec{D})$ $\vec{q}$ depends on $\vec{D}$.

Before we had current BC, Now we fix the t_{ij} so in a sense this would be proportional

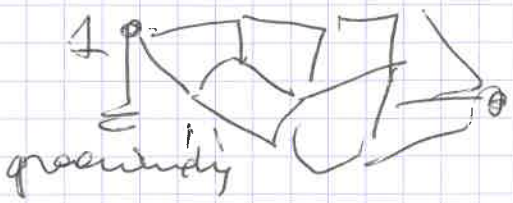

But now the exponential depends on the neighbours! $q_i = -\sum a_{ij}$

And this could be happening in the real!
Before things were strongly coupled.

Last thing

The novelty here in $1 \to i$ is not the same as $i \to 1$ (it's easier to get out of a small village and go to SF than to SF to a small village).

This translates into the need to symmetrize $\xi_{ij} = t(v_i, v_j) + t(v_j, v_i)$

both and both this (symmetrized too).
it's the effective resistance of a graph.

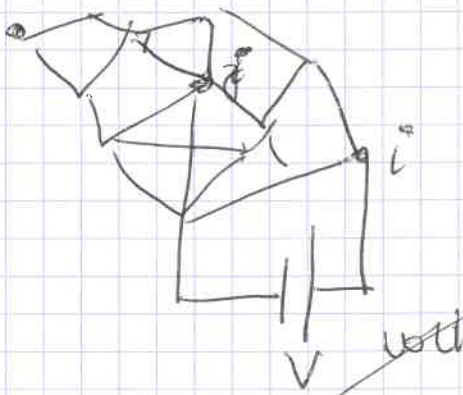

$R_{eff} = \frac{\Delta P}{Q}$ (or $\frac{\Delta V}{a}$)

$\xi_{ij} = R_{eff}$

voltage drop.

it's the same thing as the effective resistance

it effectively gives you a METRIC
a distance measure!

Networks in all systems obey the same laws.

David lecture 2

gradient descent

$E(\theta) = \sum_{i=1}^{n} e_i(x_i, \theta)$ ← optimization lecture.

$\vec{v}_t = \eta_t^0 \nabla_\theta E(\theta_t)$ step size.

update rule.

$\theta_{t+1} = \theta_t - v_t$

actually works very well often.

If your steps are too small you [illegible] they not go down before your step size is big.

If it's too small it will take forever.

Newton's method

$E(\theta + v) \simeq E(\theta) + \nabla_\theta E(\theta) v + \frac{1}{2} v^T H v$

$0 = \nabla_\theta E(\theta) + H v_{opt}$

$v_t = H^{-1}(\theta_t) \nabla_\theta E(\theta_t)$

$\theta_{t+1} = \theta_t - v_t$

maximal η is $\frac{2}{\lambda_{max}}$ (above that you explode)

a theme for getting an idea for the step size

[illegible] you replace the Hessian by

you should use.

If you use different step sizes in different directions you will converge faster but it will not generalize as well.

You won't perform as well on the test data although you will fit the training data better.

Stochastic gradient descent

$\nabla_\theta E \to \sum_{i \in B_t} \nabla_\theta e_i(x_i, \theta)$

you take a random subset of your data in each minibatch B_t. In each minibatch it's a different subset.

You can optimize the hyperparameters and use different ones on top of that.

Stochastic GD (with momentum)

Whatever you pick you need v_t you have some memory of what happened at the previous step.

$v_t = \gamma v_{t-1} + \eta_t \nabla(\quad)$ so this is inertia.

How do you pick γ? In practice you pick a value, you use it, you test it, see how it works.

Next topic:

Bayesian inference.

Introduce a few concepts:

- likelihood function: $p(x|\theta)$

  x is fixed, is a function of θ

- prior distribution $p(\theta)$

- posterior distribution $p(\theta|x) = \frac{p(x|\theta)p(\theta)}{\int d\theta' p(x|\theta')p(\theta')}$

⇒ now we are outputting a distribution of the parameters.

You can then average over this distribution of parameters → proper Bayesian.

This is often not done because it's computationally.

So usually you use some frequency statistics,

$\langle\theta\rangle = \int d\theta\, \theta\, p(\theta|x)$ Bayes estimate

$\hat{\theta}_{MAP} = \arg\max_\theta p(\theta|x)$ maximum a posterior estimate.

→ low variance high bias.

If you want to work without a prior you can.

$\hat{\theta}_{ML} = \arg\max_\theta p(x|\theta)$ ML estimate.

↑ high variance low bias regime.

Example: linear regression the value.

$$y = x^T w^* + \varepsilon$$

likelihood $p(y|x, \theta) = N(y | \mu(x), \sigma^2)$

where $\mu = x^T w$ ← mean the value.

$\theta = \{w, \sigma^2\}$

$$\hat{\theta} = \arg\max_\theta \log p(y | x, \theta) = -\frac{1}{2\sigma^2} \sum_{i=1}^{M} (y_i - x_i^T w)^2 + \text{const}(\sigma)$$

Choose the prior:

$$p(\theta) = \prod_i \sqrt{\frac{\lambda}{2\pi}} e^{-\lambda \theta_i^2}$$

$$p(\theta) = \prod_i \frac{\lambda}{2} e^{-\lambda |\theta_i|}$$

Gaussian prior.

~~Laplacian~~ exponential prior that results in an L_1 penalty.

$$\hat{\theta}_{MAP} = \arg\max \log p(x|\theta) p(\theta) =$$

Ridge regression $= \sum_{i=1}^{n} (y_i - x_i^T w)^2 + \tilde{\lambda} \|w\|_2^2$

tries to get penalty L_2 norm sense small wij, discourages sparsity

$L_1 \rightarrow$ Lasso regression.

↓

It pulls you away from fitting the noise of the data

⇒ regularization.

There is an equivalence between using a prior and using regularization (or penalty parameter dispersion).

Let's say you want to learn the Hamiltonian of an Ising model. (1D)

$$H = -J \sum_{j=1}^{L} s_j s_{j+1}$$

Learn J matrix. Our model does not know things are 1D, so we are going to run over all pairs.

$H_{model} = \sum_{jk} J_{jk} S_j S_k$ so N^2 values

we are going to deal with linear regression and then put on different activation functions.

Ising does well, and it is the only one that gets the right model. Often do badly. And it would not have to be the case that we get the right model.

Here we get the right model because it is within the class of models we are working.

But it does not have to be the case.

Next logistic regression (this is working on discrete cases).

→ discussion.

logistic regression two dimensions.

The language of logistic regression is carried to the language of neural networks.

logistic regression:

Our task is to figure out whether something is class 1 (dogs) or 0 (cats)

$y_i = 0$

$y_i = 1$

5 features (color) at the data & effort.

signal → $s = x^T w + b_0$

outputs zero or one:

PERCEPTRON $\sigma(s) = \text{sign}(s)$ ← this is what a perceptron does. sharp decision.

logistic regression does a soft version of this:

$\sigma(s) = \frac{1}{1+e^{-s}}$ LOGISTIC REGRESSION.

There a various ways of justifying $\frac{1}{1+e^{-s}}$ as hypothesis, but we will not do that,

$$p(y_i=1 \mid x_i, b) = \frac{1}{1+e^{-(x_i^T w + b)}}$$

↑ prob of class 1 (dogs).

Write down likelihood model:

$$p(D|\theta) = \prod_{i=1}^{n} [\sigma(x_i^T, w)]^{y_i} [1 - \sigma(x_i^T w)]^{1-y_i}$$

Maximize the log-likelihood as the usual,

$$+ \log p(D|\theta) = \Sigma + y_i \log \sigma(x_i^T w) + (1-y_i) \log (1-\sigma(x_i^T w))$$

this is called then "cross entropy"

- To learn the model do gradient descent.
  You can compute the gradient but you get a transcendental eqn, so you cannot set it to zero and get parameters, analytically. Can do it numerically.

Generally this 'nets not just binary outputs but these classes and can turn it into a neural network → next time.

Bill Bialek lecture (seminar)

RG - inspired to networks of real neurons.
The world is complicated. You cannot take into account everything that has been seen.
Why are you writing down simple models?
Two reasons: - the world really is simpler than it might be at the microscopic models → nice predictions.

that is probably not the most obviously simple models work for biological systems: already molecules are complicated

- simplicity emerges as we try to describe macroscopic phenomena like life
→ phenomenology of neurons.

RG

Let's think of systems that are in thermal equilibrium

If I have local variables $\{\sigma_i\}_{i=1,2,\dots N}$ N variables.

$$P(\{\sigma_i\}) = \frac{1}{Z} e^{-\beta E(\{\sigma_i\})}$$

What does it mean to say that I am looking for simplicity to emerge as I zoom out macroscopic phenomena.

original variables σ_i are on a scale a

Let's imagine defining some coarse-grained variables on scale ℓ $\tilde{\sigma}_i(\ell)$ (N variables)

(imagine defocusing your microscope).

What is the prob. distribution on scale ℓ:

$$P_\ell(\{\tilde{\sigma}_i^\ell\}) = ?$$

$$E = g_1 \hat{O}(\{\sigma_i\}) + g_2 \hat{O}_2(\{\sigma_i\}) + \dots$$

I can expand the energy in some way.

I can think of the probability distribution that I end up with also as a Boltzmann distribution under the energy as an expansion in the new $\tilde{\sigma}_i^\ell$ basis. If still have N variables,

$$P_\ell(\{\tilde{\sigma}_i^\ell\}) = ? \sim e^{-\beta E(\ell)(\{\tilde{\sigma}_i^\ell\})}$$

$$E(\ell) = \tilde{g}_1^\ell \hat{O}(\{\tilde{\sigma}_i^\ell\}) + \tilde{g}_2^\ell \hat{O}(\{\sigma_i\}) + \dots$$

but now I need less coefficients.

I still have the same # of variables, but less coefficients because most go to zero. That is what it means that the model is simpler.

$(g_1, g_2 \ldots) \to (\tilde{g}_1, \tilde{g}_2, \ldots)$ a lot

$\frac{d\tilde{g}_1}{dl} > 0$ relevant variables

$\frac{d\tilde{g}_n}{dl} < 0$ irrelevant

} keep iterating until the RG flow to a fixed point)

if most terms were relevant and a few irrelevant we would be in trouble. Luckily it the other way round: most terms are irrelevant.

Exercise 1D Ising model

↑ ✕ ↓ | ↑ ✕ ↑ | ✕ ↓ average over every other spin.

$2n$ $2n+1$ $2n+2$

$E = -J \sum_n \sigma_n \sigma_{n+1}$ $E_{\square} = -J(\sigma_{2n}\sigma_{2n+1} + \sigma_{2n+1}\sigma_{2n+2})$

$\sum_{\{\sigma_{2n+1}\}} e^{-\beta E(\{\sigma_n\})} \Rightarrow e^{-\beta E'}$ ← it looks like $2n$ and $2n+2$ now interact

sum over odd spins

But $2n+2$ does not interact

$E' = -J \sum_n \sigma_{2n}\sigma_{2n+2}$ with $2n-2$, so spin still interacts with NN. So I still have a Boltzmann distribution. But with a different energy.

The absolute value of J' is always smaller than J.

So as you zoom out the spins become less and less interacting and then the spins never order.

In 2D ↑ ✕ ↑ / ✕ ↑ ✕ / ↑ ✕ ↓ things get complicated because even if it's NN then you have interacting along the diagonal.

In 3D you get more terms and gradually all these terms go away.

At the end for 2 days 2 survive:

Coarse grain:

- decimation (eg. every other spin)
- block spins

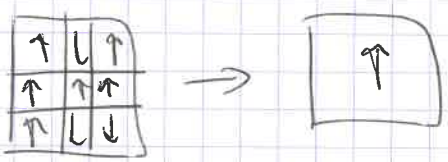

real space methods via numerical simulation you would do this

- Fourier transform

$$\sigma(k_i) = \sum_{k} e^{ik\vec{x}_i} \sigma(\vec{k})$$

I also prop to produce coarse grained variables

$$\tilde{\sigma}_i = z_i \sum_{|k|<\frac{1}{a}} e^{ikx_i} \sigma(k)$$

momentum shell methods

pen & paper you would do this

The goal is not to make these constructions but to follow the probability distribution of the variables that are left.

Start with these techniques go work on the methods you care about.

What is the problem?

real space $\overset{\text{needs}}{\longleftrightarrow}$ neighbourhood

I need distance, what is a distance?

Real distance, probably not all the neurons in a given neighborhood matter.

If I want to do momentum shell? what is momentum.

momentum shell $\hookrightarrow$ momentum?

We need to have candidates for what neighborhood & momentum mean?

These may not be the best solutions, they are just candidates.

version of this already exist in disordered systems but problems here are severe.

I do know how to calculate the covariance matrix
$\langle \sigma_i \sigma_j \rangle - \langle \sigma_i \rangle \langle \sigma_j \rangle = C_{ij}$

If your variables things live on a lattice and a ~~generally~~ positive

Pearson correlation coefficient → nearest neighbour of j $\longleftrightarrow$ $\text{argmax}_{j \ne i}\ C_{ij}/\sqrt{C_{ii} C_{jj}}$

How would I change my length scale of averaging by a factor of 2;

find max C_{ij} $(i \neq j)$ $(i, j^*(i))$ "NN" in this sense of i

$$\sigma_i + \sigma_{j^*(i)} = \tilde{\sigma}_i$$

repeat . . . $\{\sigma_i\}_{i=1 \ldots N} \longrightarrow \{\tilde{\sigma}_i^{(1)}\}_{i=1 \ldots N/2}$

This statement makes sense for ~~all~~ one spin.
Now you need to do this for all ~~of~~ them (pair them up). Here proposing to do this in a greedy way.

This is a coarse group strategy.
It depends on what my ~~the~~ observables are.

For momentum shell you want translational invariance

$$C_{ij}\, u_j^\mu = \lambda_\mu\, u_i^\mu \leftarrow e^{i k_\mu x_i}$$
$$\lambda_\mu \Leftrightarrow \int d^3r\ C(\vec{r})\, e^{i\vec{k}\vec{r}}$$
$$G(k)$$

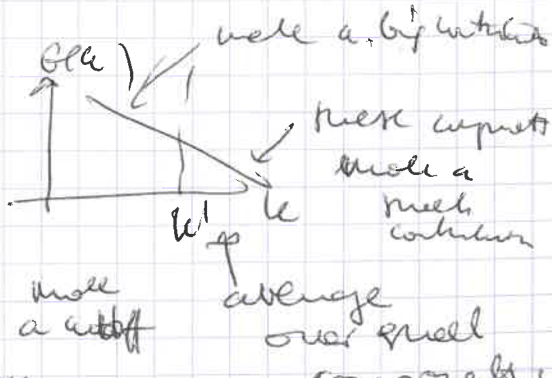

this sounds like PCA —— make a cutoff

the question in PCA is where is the cut-off.

In most system the ~~[illegible]~~ spectrum is continuous.
So RG tells you you should ask what changes when you change the cut-off.

The proposal:
what corresponds to neighborhood is the most correlated component
what corresponds to momentum are the principal components.
(correct: and this is for one true point)
implement → how do you follow flow?

"flow" $(g_1, g_2, \ldots) \to ?$
what it is for the Ising model,
$P_\ell(\tilde{\sigma}_1^\ell)$

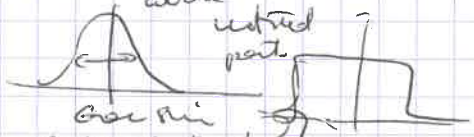

What happens when you change ℓ

People who have been doing Monte Carlo simulations have thought about this:

$$\langle \sigma_i^{(\ell)}(t)\, \sigma_i^{(\ell)}(t+\tau) \rangle_c = G_\ell(\tau) \underset{\uparrow}{=} C\left(\tau/\tau_c(\ell)\right)$$

It flows are becoming simpler when you vary ℓ.
Dynamical scaling

$$\tau_\ell \sim \ell^z$$

These are the examples of the kinds of things you can check for.
We are going to follow these coarse graining procedures
We are going to do it for real data for a mouse's
fluorescent protein that lights up when calcium is
high (action potential).

In most regions of the mouse brain nothing happens unless the mouse behaves (let the mouse run on a styrofoam ball),
→ virtual reality for a mouse.

continuous $x_i(t)$ variables

or $\sigma_i = \{0, 1\}$ discrete variable.

Real space version

compute correlation matrix, most most correlated pair cells,

normalize so that mean nonzero signal = 1

continue until N cells become N/2

iterate

Momentum Flow! ← non spiky silence. ← probability that the coarse grained cluster size, distribution is zero,

$$P_k(x) = P_0(k)\,\delta(x) + [1 - P_0(k)]\,Q_k(x)$$

$$\int dx\, Q_k(x) = 1$$ ↖ probability that the coarse grained distribution is not zero.

experimentally you see

$$P(k) = \exp(-\alpha k^{\beta})$$

$$\beta \simeq 0.87 \pm 0.014 \pm 0.015$$

(0.014: within one mouse; 0.015: within different mice)

If you look at the continuous part of the distribution (the non-silent one).

The distribution of the coarse grained distribution is approaching a fixed distribution.

Dynamic scaling: if you rescale by the correlation time you see a collapse.

correlation time τ

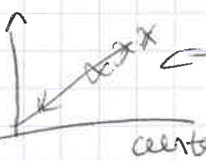

← power law, $\tau_c \propto k^{\tilde{z}}$

$\tilde{z} \sim$ 8 ... [illegible]

cluster size

Now momentum space
→ distribution of coarse grained variables approaches a fixed non-Gaussian form.

and you can get dynamic scaling $\varepsilon_c \sim k^2$
and the exponents agree.