Statistical physics of disordered systems

Course information

5 weeks, 4 classes per week

(Roughly) every other class consists of problem solving in groups

No homework, but if you miss a ‘problem solving’ class it becomes homework

Graded by attendance and participation in problem solving

Professor: Jaron Kent-Dobias, Room 313 IFT-Unesp, `jaron@ictp-saifr.org`

What is a *non*-disordered system?

Energy depends on configurations $\boldsymbol{s}$ and their momenta $\boldsymbol{p}$, giving Hamiltonian $H(\boldsymbol{s}, \boldsymbol{p})$

Hamiltonian implies dynamics of configurations which evolve in time like $\boldsymbol{s}(t)$, $\boldsymbol{p}(t)$

We observe macroscopic properties of the system over some small region (or coupled to an external bath) over some relatively large timescale,

$$\langle A\rangle(\boldsymbol{s}_0, \boldsymbol{p}_0) = \frac{1}{t_f - t_i}\int_{t_i}^{t_f} dt\, A(\boldsymbol{s}(t), \boldsymbol{p}(t))$$

Ergodic hypothesis tells us that averages of macroscopic properties over macroscopic timescales is the same as their average at appropriate temperature T, with

$$\langle A\rangle = \frac{1}{Z'}\int d\boldsymbol{s}\, d\boldsymbol{p}\, e^{-H(\boldsymbol{s},\boldsymbol{p})/T} H(\boldsymbol{s}, \boldsymbol{p}) = \frac{1}{Z}\int d\boldsymbol{s}\, e^{-H(\boldsymbol{s})/T} H(\boldsymbol{s})$$

At the end, predictions depend only few model parameters and temperature

What is a disordered system?

Configurations are split into fast dynamic degrees of freedom $\boldsymbol{s}$ and slow frozen degrees of freedom $\boldsymbol{J}$, Hamiltonian is $H(\boldsymbol{s}, \boldsymbol{J})$

The frozen configurations can result from the system itself having some degrees of freedom 'get stuck', or they can be set from the outset

Everything from the previous slide follows for the fast degrees of freedom: measuring macroscopic properties over macroscopic timescales gives Boltzmann:

$$\langle A(\boldsymbol{J}) \rangle = \frac{1}{Z} \int d\boldsymbol{s}\, e^{-H(\boldsymbol{s},\boldsymbol{J})/T} H(\boldsymbol{s}, \boldsymbol{J})$$

Now, results depend in principle on *extremely many* degrees of freedom $\boldsymbol{J}$ which vary from system to system

Examples of disordered systems

Spin glasses

Substances made from magnetic impurities (iron or manganese) embedded as an alloy in a conductor (copper or gold)

When cooled, the position of the impurities are fixed and spread randomly: frozen degrees of freedom

The magnetic moments (spins) of the impurities fluctuate and interact with each other with an exchange interaction mediated by the conductor that depends very sensitively on the distance between them: fast degrees of freedom

At low temperature, the spins undergo a phase transition into a amorphous state

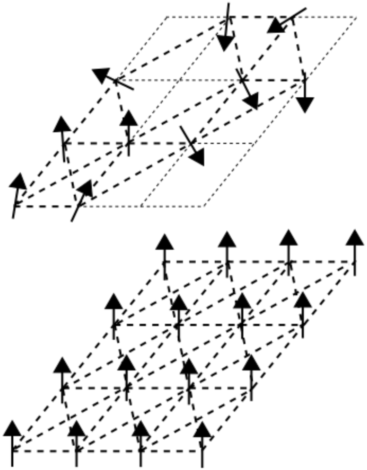

Examples of disordered systems

Spin glasses

At the spin glass transition, the magnetic susceptibility has a cusp, like antiferromagnetism (but no sign of antiferromagnetism)

But, the location of the cusp depends on the frequency of the measurement!

There is also not the usual sign of a phase transition in the specific heat

. . . and, the structure of the cusp at low temperatures depends on whether the sample was cooled in the presence of a magnetic field (FC for field cooled) or not (ZFC for zero field cooled)

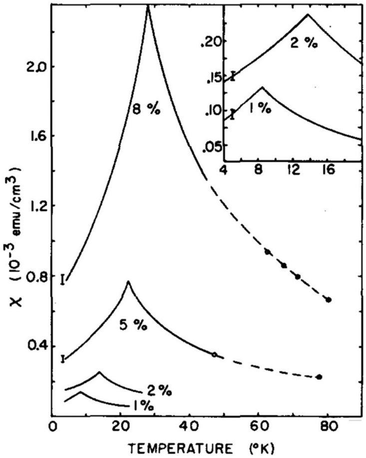

FIG. 1. The magnetic susceptibility of spin glasses: the results for Au-Fe alloys in the region of low Fe concentration. From Cannella, Mydosh, and Budnick, 1971.

Examples of disordered systems

Spin glasses

At the spin glass transition, the magnetic susceptibility has a cusp, like antiferromagnetism (but no sign of antiferromagnetism)

But, the location of the cusp depends on the frequency of the measurement!

There is also not the usual sign of a phase transition in the specific heat

. . . and, the structure of the cusp at low temperatures depends on whether the sample was cooled in the presence of a magnetic field (FC for field cooled) or not (ZFC for zero field cooled)

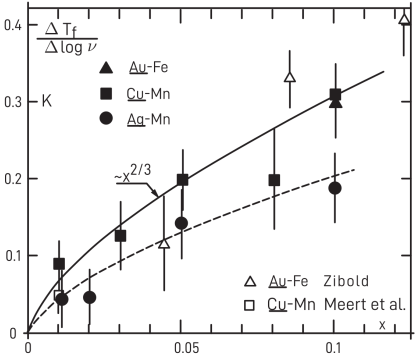

FIG. 3. The absolute variation of T_f (equivalent to what we call T_g in this review) per decade of time, reported as a function of x, the percentage fraction of Mn, for CuMn ($1\% < x < 10\%$). This can be represented by an $x^{2/3}$ law. From Tholence, 1981.

Examples of disordered systems

Spin glasses

At the spin glass transition, the magnetic susceptibility has a cusp, like antiferromagnetism (but no sign of antiferromagnetism)

But, the location of the cusp depends on the frequency of the measurement!

There is also not the usual sign of a phase transition in the specific heat

...and, the structure of the cusp at low temperatures depends on whether the sample was cooled in the presence of a magnetic field (FC for field cooled) or not (ZFC for zero field cooled)

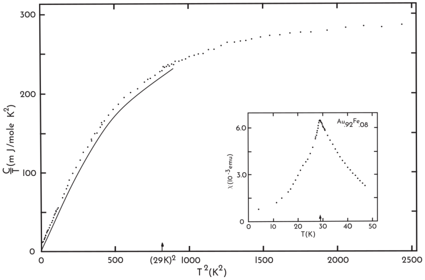

FIG. 4. Specific heat of $Au_{0.92}Fe_{0.08}$ in the temperature region 3–50 K shown as a plot of C/T vs T^2. The solid curve is the calculated nonmagnetic contribution to the specific heat of the alloy between 0 and 30 K. Inset: susceptibility results of the same sample, which were provided by S. A. Werner. From Wenger and Keesom, 1975, 1976.

Examples of disordered systems

Spin glasses

At the spin glass transition, the magnetic susceptibility has a cusp, like antiferromagnetism (but no sign of antiferromagnetism)

But, the location of the cusp depends on the frequency of the measurement!

There is also not the usual sign of a phase transition in the specific heat

. . . and, the structure of the cusp at low temperatures depends on whether the sample was cooled in the presence of a magnetic field (FC for field cooled) or not (ZFC for zero field cooled)

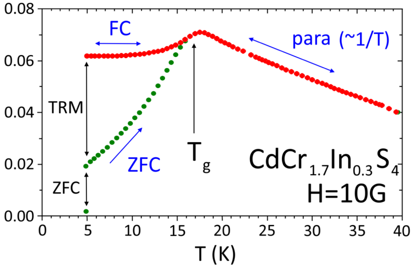

FIG. 5. The field-cooled (FC, upper), zero-field-cooled (ZFC, lower), and thermoremanent (TRM, middle) magnetizations against temperature for $CdCr_{1.7}In_{0.3}S_4$ in an external magnetic field $H = 10$ G. From Dupuis, 2002, and Vincent, 2024.

Examples of disordered systems

Spin glasses

At the spin glass transition, the magnetic susceptibility has a cusp, like antiferromagnetism (but no sign of antiferromagnetism)

But, the location of the cusp depends on the frequency of the measurement!

There is also not the usual sign of a phase transition in the specific heat

. . . and, the structure of the cusp at low temperatures depends on whether the sample was cooled in the presence of a magnetic field (FC for field cooled) or not (ZFC for zero field cooled)

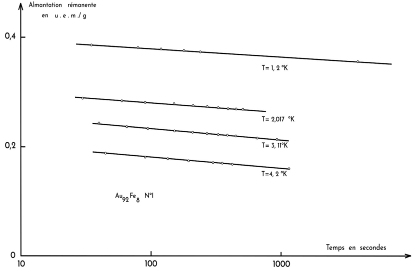

FIG. 6. Time decay of the TRM of $Au_{92}Fe_8$. From Tournier, 1965.

Examples of disordered systems

Spin glasses

Not only do properties of the low-temperature state depend on how it was cooled, they also change in time: the difference between the FC and ZFS procedures depends how long after cooling you measure the susceptibility

Time dependence is also reflected in the structure factor, which relaxes more slowly the longer you wait

The slow drift of properties of a cooled sample with waiting time after cooling is called *aging*, indicates something nonequilibrium is occurring

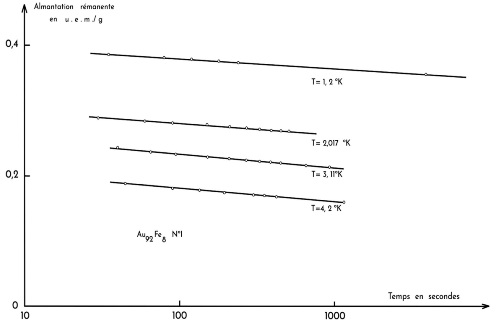

FIG. 6. Time decay of the TRM of $Au_{92}Fe_8$. From Tournier, 1965.

Examples of disordered systems

Spin glasses

Not only do properties of the low-temperature state depend on how it was cooled, they also change in time: the difference between the FC and ZFS procedures depends how long after cooling you measure the susceptibility

Time dependence is also reflected in the structure factor, which relaxes more slowly the longer you wait

The slow drift of properties of a cooled sample with waiting time after cooling is called *aging*, indicates something nonequilibrium is occurring

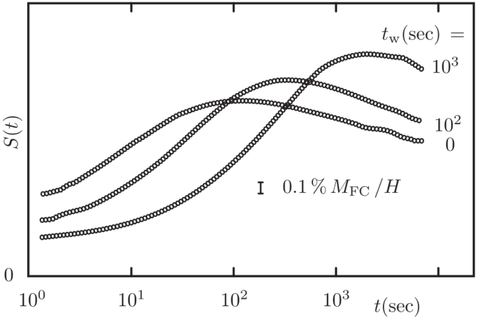

FIG. 8. Relaxation rates $S(t)$ of the remanent magnetization of CuMn 5 at. % ($T_g = 28$ K) at $T = 21$ K for different waiting times t_w. A relaxation rate of 0.1% of the FC susceptibility value is indicated. From Nordblad *et al.*, 1986.

Examples of disordered systems

Spin glasses

In equilibrium, the fluctuation–dissipation theorem relates the time-dependent correlation function $C(t', t)$ to the integrated response $\chi(t', t)$ by $\chi(t', t) = \beta C(t', t)$.

Spin glasses break the FDT, don't seem to recover it at large waiting times

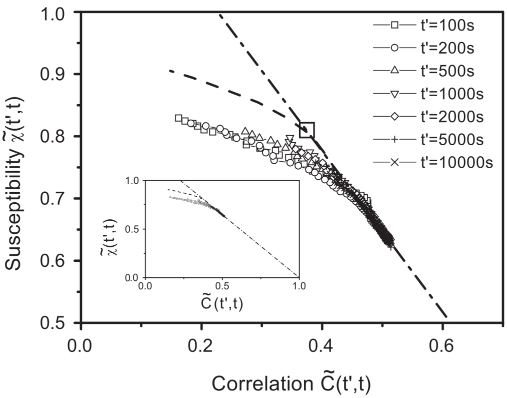

Examples of disordered systems

Spin glasses

In equilibrium, the fluctuation–dissipation theorem relates the time-dependent correlation function $C(t', t)$ to the integrated response $\chi(t', t)$ by $\chi(t', t) = \beta C(t', t)$.

Spin glasses break the FDT, don't seem to recover it at large waiting times

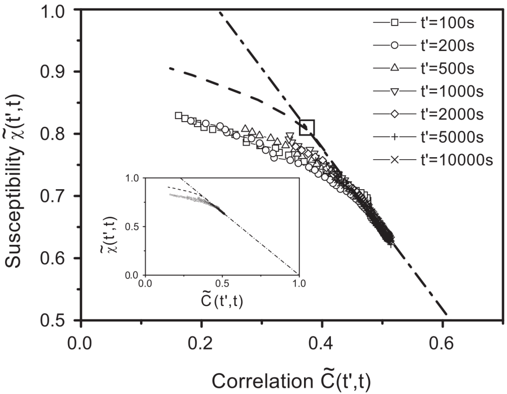

Examples of disordered systems

Spin glasses

Samples that underwent some time-dependent procedure at a certain temperature

1. *rejuvenate* at lower temperatures (behave as if they haven't aged)...
2. but *remember* that procedure when reheated (show the aged susceptibility)

Low-temperature spin glasses behave very differently from tradition non-disordered systems

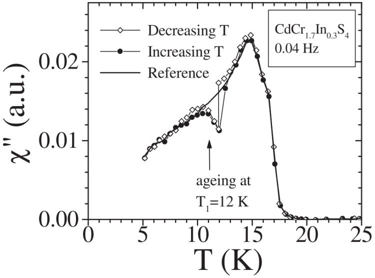

FIG. 45. The out-of-phase susceptibility χ'' of the $CdCr_{1.7}In_{0.3}S_4$ spin glass. The solid line was measured upon heating the sample at a constant rate of 0.1 K/ min (reference curve). The open diamonds indicate a subsequent measurement made during cooling at the same rate, except that the cooling procedure was stopped at 12 K over 7 h to allow for aging. Cooling then resumed down to 5 K: χ'' is not influenced and goes back to the reference curve (chaos). The solid circles indicate data taken after this cooling procedure at the previous constant rate, exhibiting memory of the aging stage at 12 K. From Jonason *et al.*, 2000.

This course

Spin glass theory

Most of this course will consist of the *mean field theory* of *spin glass models*

At the end, we will be able to at least qualitatively explain all of the behaviors outlined above

However, no one learns spin glass theory today to explain spin glasses: they are a "useless material" – 2021 Nobel Prize winner Giorgio Parisi

The ideas and methods originally developed for mean field spin glasses now find applications in diverse fields: glasses, neuroscience, machine learning, ecology, etc

Examples of disordered systems

Structural glasses

Glasses are made from certain substances (silica, rubber, etc) that fail to crystalize at low temperatures

Sometimes they need to be cooled quickly ("quenched") to prevent crystalization, but not always

Glasses are rigid, but have a microscopic structure that does not appear different from the liquid

"The deepest and most interesting unsolved problem in solid state theory is probably the theory of the nature of glass and the glass transition." – P.W. Anderson

Examples of disordered systems

Structural glasses

Glasses are made from certain substances (silica, rubber, etc) that fail to crystalize at low temperatures

Sometimes they need to be cooled quickly ("quenched") to prevent crystalization, but not always

Glasses are rigid, but have a microscopic structure that does not appear different from the liquid

"The deepest and most interesting unsolved problem in solid state theory is probably the theory of the nature of glass and the glass transition." – P.W. Anderson

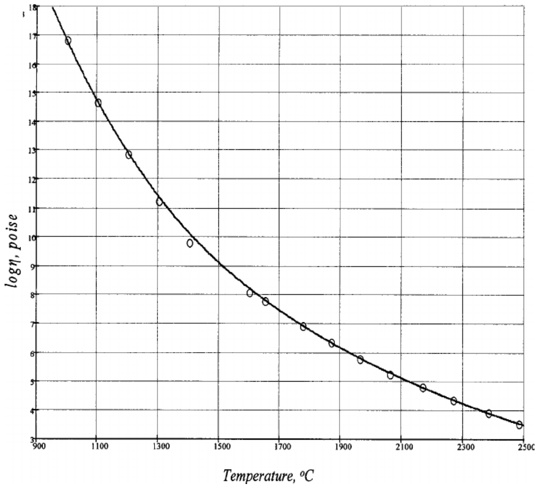

Examples of disordered systems

Structural glasses

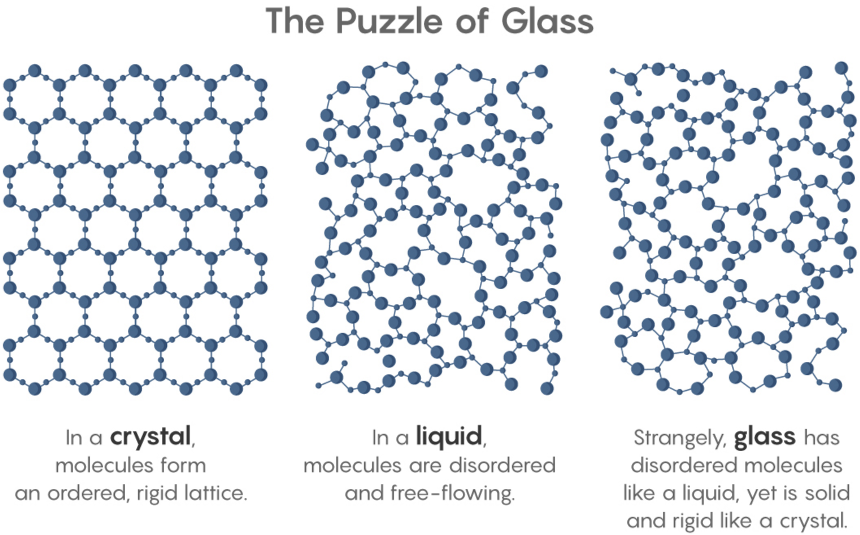

Examples of disordered systems

Structural glasses

Structure factor doesn't notably change at glass transition (376K in the figure at right)

The fast and slow variables are difficult to separate: every glass is different in the positions of the atoms, but their positions are also not completely frozen

Is it is phase transition or just really slow liquid? We don't really know

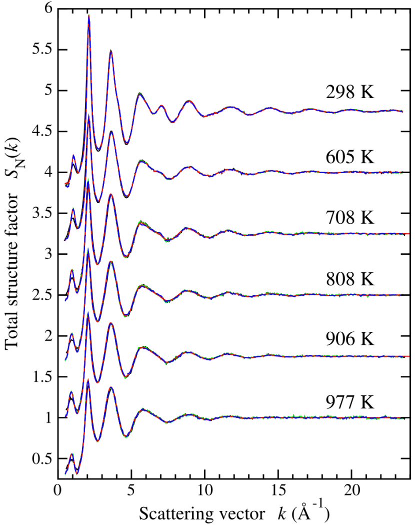

Fig. 1. The neutron total structure factor $S_N(k)$ for glassy and liquid $ZnCl_2$ where the results

Examples of disordered systems

Structural glasses

Typical 'phase diagram' for glass: crystalization is avoided, liquid remains but is metastable, eventually reaches a cooling-rate-dependent state that leaves the metastable equilibrium liquid

What would happen if you were able to wait and follow the metastable liquid? Perhaps a paradox (Kauzmann): the extrapolated entropy of the metastable liquid appears to go below that of the crystal at finite temperature

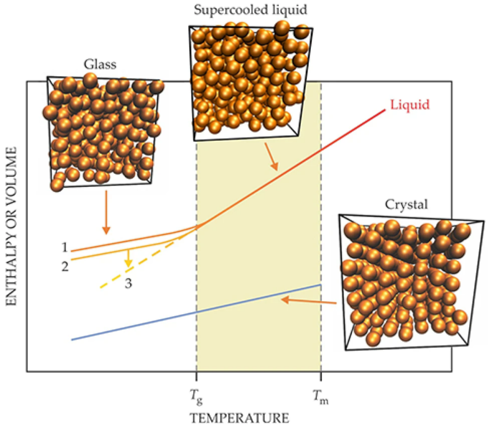

Examples of disordered systems

Structural glasses

Typical 'phase diagram' for glass: crystalization is avoided, liquid remains but is metastable, eventually reaches a cooling-rate-dependent state that leaves the metastable equilibrium liquid

What would happen if you were able to wait and follow the metastable liquid? Perhaps a paradox (Kauzmann): the extrapolated entropy of the metastable liquid appears to go below that of the crystal at finite temperature

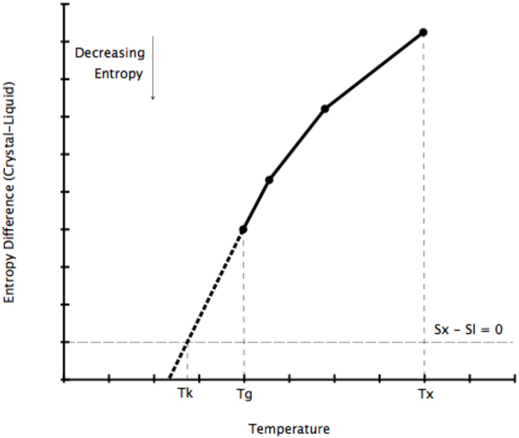

Examples of disordered systems

Structural glasses

Glasses show aging: properties depend on the amount of time since they are prepared or stressed

Shown both in bulk response and in the correlation functions (here for simulations of binary Lenard–Jones particles)

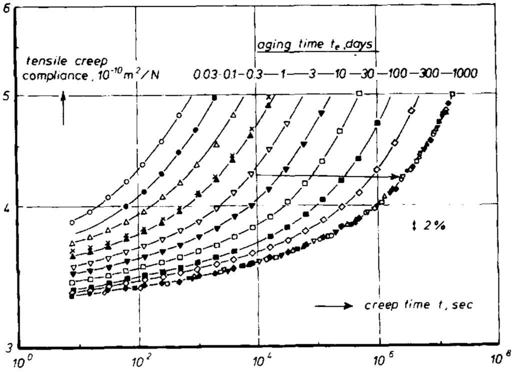

Fig. 3. Small-strain tensile creep curves of rigid PVC quenched from 90°C (i.e., about 10°C above T_g) to 40°C and further kept at 40 ± 0.1°C for a period of 4 years. The different curves were measured for various values of the time t_e elapsed after the quench. The master curve gives the result of a superposition by shifts which were almost horizontal; the shifting direction is indicated by the arrow. The crosses refer to another sample quenched in the same way, but measured for creep at a t_e of 1 day only.

Examples of disordered systems

Structural glasses

Glasses show aging: properties depend on the amount of time since they are prepared or stressed

Shown both in bulk response and in the correlation functions (here for simulations of binary Lenard–Jones particles)

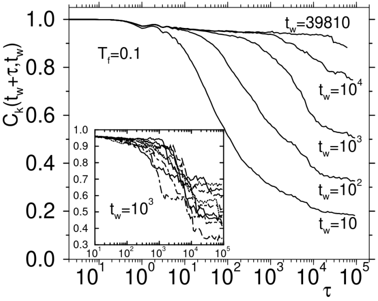

Fig. 10. Main figure: time dependence of the correlation functions $C_k(t_w + \tau, t_w)$, for the waiting times $t_w = 10$, 100, 1000, 10 000, and 39 810. $T_f = 0.1$, $k = 7.23$. Inset: the same correlation function for $t_w = 1000$ for the individual runs.

Examples of disordered systems

Machine learning

Machine learning models take huge amounts of data (all the text/images on the internet) and use them to train models with even larger numbers of parameters to perform simple tasks (write text, identify images, etc)

Despite not being thermal, they have the form of a disordered system: set of data is fixed from the outset (frozen DOF), weights are varied to make network perform well (fast DOF)

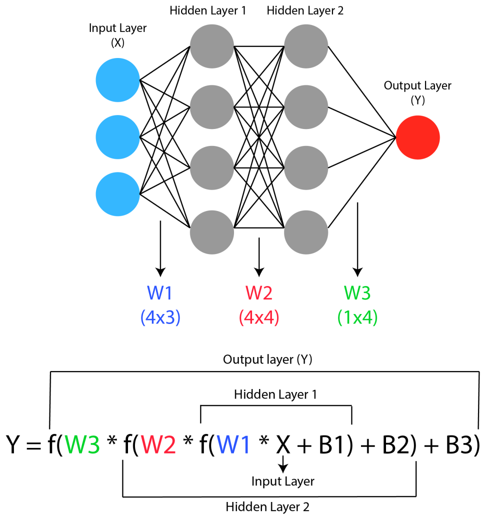

Examples of disordered systems

Machine learning

Neural network is iterative nonlinear function of input depending on weight matrices W_i and bias vectors B_i taking input X to label Y

Want to find the values of W_i and B_i such that $Y = F(X \mid \{W, B\})$ is "correct" for all (X, Y)

Trained by taking dataset $\{(X_i, Y_i)\}$ and minimizing loss

$$L(\{W, B\})) = \sum_{i=1}^{M} [F(X_i \mid \{W, B\}) - Y_i]^2$$

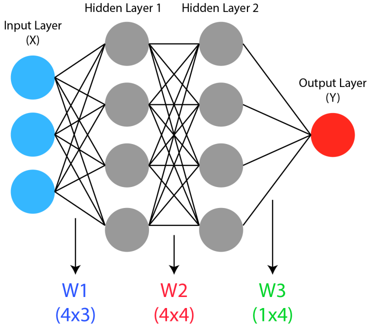

Output layer (Y)

Hidden Layer 1

$$Y = f(W3 * f(W2 * f(W1 * X + B1) + B2) + B3)$$

Input Layer

Hidden Layer 2

Examples of disordered systems

Machine learning

Neural networks are *not* glassy: if 'cooling' (training) them to find 'low-energy' (low loss function) was slow, they would be bad models!

Instead, massive overparameterization (more fit parameters than data points) makes loss easy to train to zero

Shockingly, massive overparameterization *works*, given the way people train networks (calculating the gradient of the loss with small subsets of the data works *better* than using all the data)

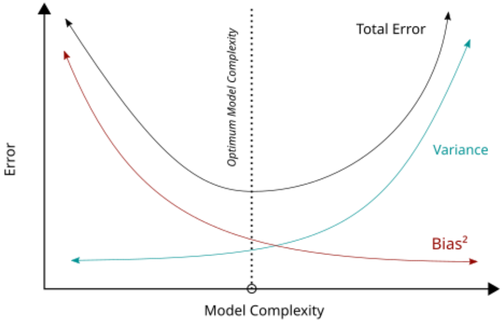

Examples of disordered systems

Machine learning

Neural networks are *not* glassy: if 'cooling' (training) them to find 'low-energy' (low loss function) was slow, they would be bad models!

Instead, massive overparameterization (more fit parameters than data points) makes loss easy to train to zero

Shockingly, massive overparameterization *works*, given the way people train networks (calculating the gradient of the loss with small subsets of the data works *better* than using all the data)

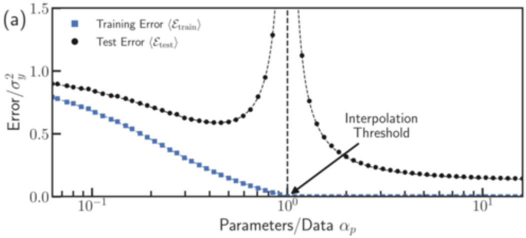

Linear least squares

The bad, the good, the ugly, and the weird

$M = 40, \ N = 2$

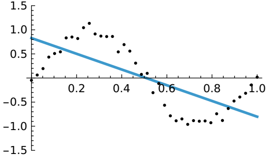

Underfit

χ^2 is large

Best fit has *high bias*

$M = 40, \ N = 7$

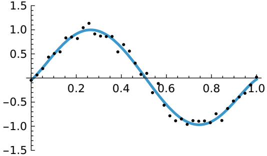

Good fit!

χ^2 is moderate

Best fit has *low variance* and *low bias*

$M = 40, \ N = 40$

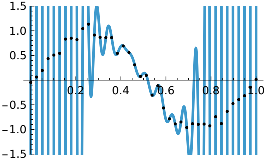

Overfit

χ^2 is zero

Best fit has *high variance*

$M = 40, \ N = 80$

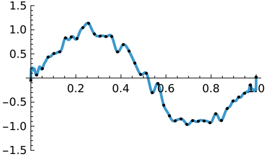

Good fit?

χ^2 is zero

Best fit has *low variance* and *low bias*

Examples of disordered systems

Machine learning

How does noisy out-of-equilibrium training drive performance? Can we describe the dynamics of training quantitatively?

How well do networks learn a task with a known answer? If a network can in principle *perfectly* learn a task, does it?

How does the structure of data and/or the structure of the network influence how the task is learned, and to what quality?

Are there different models of learning besides feed-forward neural networks that are energetically cheaper to train and evaluate?

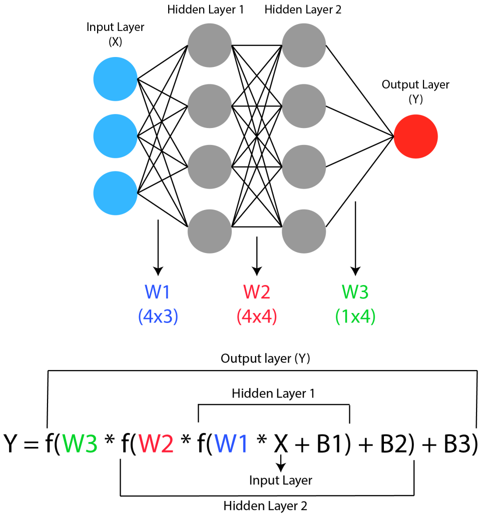

Examples of disordered systems

Ecosystems

Ecosystems contain many species that interact with each other

The set of species that an ecosystem could contain and their interactions with each other are slow (frozen) degrees of freedom, their abundances are fast degrees of freedom

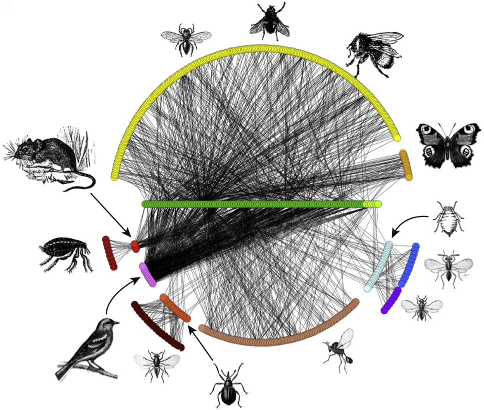

Credit: Quintessence Network DOI: 10.1016/j.tree.2015.12.003

Examples of disordered systems

Ecosystems

Commonly modelled through Lotka–Volterra differential equations for the abundances n_i for each of the species,

$$\frac{\partial n_i}{\partial t} = \mu n_i (N_i - n_i) + \sum_j J_{ij} n_j + \xi$$

which is just a logistic equation for each specie plus a linear interaction with all of the others

The matrix of interactions J is fixed and 'random'

Depending on the properties of J, the resulting evolution can qualitatively change

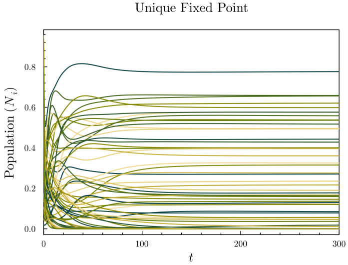

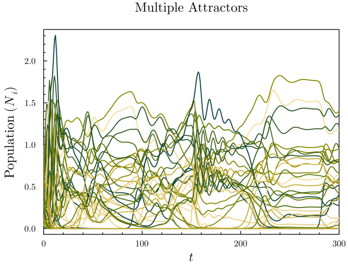

Examples of disordered systems

Ecosystems

Can we predict the stability of ecosystems from first principles?

What is the relationship between diversity and stability of large ecosystems? Should we expect 'typical' ecosystems to be very stable or nearly unstable?

Do 'healthy' and 'unhealthy' ecosystems have matrix structures J that are measurably different?

The Js are not really random, they come from evolution. What does this mean for their structure?

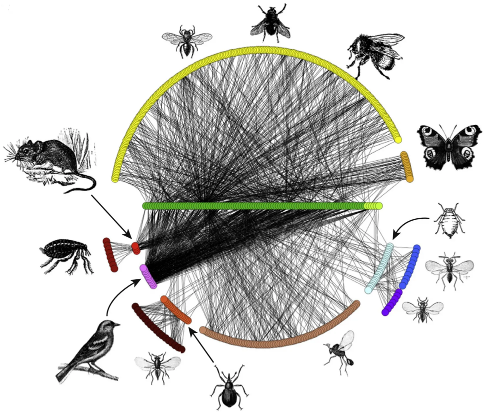

Credit: Quintessence Network DOI: 10.1016/j.tree.2015.12.003

Examples of disordered systems

Neuronal networks

Neurons are connected to each other with relatively stable network, excitations are fast

Also modelled with differential equations, with activity of neurons r_i evolving like

$$\frac{\partial r_i}{\partial t} = -\mu r_i + \sum_j \int dt' \, J_{ij}(t - t') r_j(t')$$

where $J_{ij}(t - t')$ is a *kernel* that takes into account the recent activity of the neighboring neurons

Similar to recurrent neural networks in machine learning

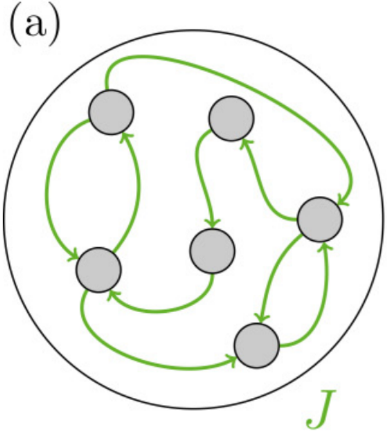

This course

The methods we will introduce in this course are

		Applicable to. . .				
	Method	Spin glasses	Glasses	'AI'	Ecology	Neurons
Statics	Replica method	✓	✓/×	✓	×	×
Dynamics	DMFT, equilibrium	✓	✓	×	×	×
	DMFT, nonequilibrium	✓	✓	✓	✓	✓
Structure	Complexity	✓	✓	✓	✓/×	✓/×
	Franz–Parisi potential	✓	✓	✓	×	×