Information Flow and Phase Transitions in Complex Dynamical Systems

Lionel Barnett

l.c.barnett@sussex.ac.uk

Sackler Centre for Consciousness Science (SCCS)
Department of Informatics
University of Sussex, UK

www.sussex.ac.uk/sackler

CIDNET14
Max Planck Institute for the Physics of Complex Systems
Dresden, Germany
June 2014

A murmuration of starlings flocking over the West Pier in Brighton, UK

Phase transitions in complex dynamical systems

- Phase transitions are ubiquitous in complex systems featuring large ensembles of dynamically interacting elements
- In many cases, these may be classed as *order-disorder* transitions, for example:
- Natural systems:
  - Physical systems (states of matter, magnetization, etc.)
  - Neural systems
  - Financial markets
  - Ecosystems
- Model systems:
  - Physics models (statistical ensemble models, spin-system models, etc.)
  - Particle swarm (flocking/schooling) models
  - Agent-based models (economics, games, etc.)
  - Phase synchronization models (e.g. Kuramoto oscillators)
  - Cellular automata
  - Random boolean networks

Phase transitions in complex dynamical systems

- In many natural systems, *dis*order is associated with “healthy” dynamics, while *order* is associated with “pathological” dynamics; a disorder $\longrightarrow$ order transition signals a calamitous event

Financial markets - crashes associated with “herding” behaviour

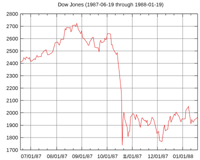

Neural systems - onset of epileptic seizures associated with runaway synchronisation

Ecosystems - catastrophic loss of biodiversity

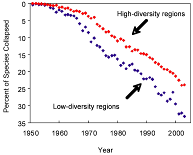

Phase transitions in complex dynamical systems

- Can we **predict** if a system is moving towards a phase transition?

- More specifically: Can we predict an imminent disorder $\rightarrow$ order phase transition using information flow dynamics?

- Perhaps . . . proof of concept:

  L. Barnett, J. T. Lizier, M. Harré, A. K. Seth and T. Bossomaier, *Information flow in a kinetic Ising model peaks in the disordered phase*, Phys. Rev. Lett. **111(17)**, 2013.

What is a phase transition (thermodynamics)?

- Probability distribution $P(\boldsymbol{s})$ over *microstates* $\boldsymbol{s}$ of system (ensemble)
  - $\boldsymbol{s} = (s_1, s_2, \ldots, s_N)$ where s_i is state of ith element, N is system size
  - E.g. Boltzmann-Gibbs distribution in statistical mechanics: $P(\boldsymbol{s}) = \frac{1}{Z} e^{-E(\boldsymbol{s})/kT}$
  - $E(\boldsymbol{s})$ is the *energy* associated with microstate $\boldsymbol{s}$ (specifies interactions between elements), T is temperature, k is Boltzmann's constant, $Z \equiv \sum_{\boldsymbol{s}} e^{-E(\boldsymbol{s})/kT}$ is the *partition function*
- Thermodynamic limit: system size $N \to \infty$
- $P(\boldsymbol{s})$ depends on *control parameters* (temperature, pressure, volume, ...)
- *Order parameter* (particle density, magnetization, ...) - may be scalar or vector
- *Susceptibility* measures statistical fluctuations of the order parameter
- *Correlation* may be measured between states of system elements at various distance scales
  - Correlation frequently decays exponentially with distance, defining a *correlation length*
- Various other thermodynamic properties are associated with the system (entropy, free energy, internal energy, specific heat capacity, ...)

Characterizing order/disorder phase transitions

- In the thermodynamic limit ...
- Order parameter is zero beyond a *critical point* in the trajectory of the system in the (phase) space of control parameters - this defines the disordered regime
  - Types of phase transitions (1st order, 2nd order, etc.) are characterised by the discontinuity of the order parameter or its derivatives at the critical point
- Susceptibility, specific heat and correlation length diverge at the critical point

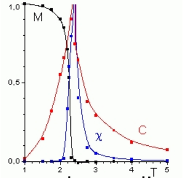

- In general, thermodynamic properties have universal power-law scaling properties at a critical point (critical exponents)
  - Everything goes fractal
- Spontaneous *symmetry breaking* and *ergodicity breaking* (may) occur

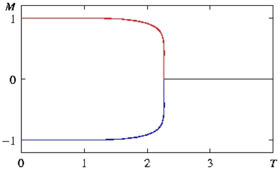

Dynamical ensemble systems

- No dynamics implied yet ... just $P(\boldsymbol{s})$
- Stationary* stochastic process $\boldsymbol{S}(t)$ describes how microstates evolve over time
  - Stationary $\implies P\big(\boldsymbol{S}(t) = \boldsymbol{s}\big) = P(\boldsymbol{s})$ for all microstates $\boldsymbol{s}$
- Frequently modelled as Markov process: transition probability $P(\boldsymbol{s} \to \boldsymbol{s}')$
  - Reversible $\implies$ *detailed balance*: $\dfrac{P(\boldsymbol{s} \to \boldsymbol{s}')}{P(\boldsymbol{s}' \to \boldsymbol{s})} = \dfrac{P(\boldsymbol{s}')}{P(\boldsymbol{s})}$
- Useful for *sampling* microstates from $P(\boldsymbol{s})$ in simulation: Markov Monte Carlo methods
- Examples (Boltzmann-Gibbs distribution):
  - $P(\boldsymbol{s} \to \boldsymbol{s}') = \begin{cases} e^{-\Delta E/kT} & \Delta E > 0 \\ 1 & \Delta E \le 0 \end{cases}$ - Metropolis dynamics
  - $P(\boldsymbol{s} \to \boldsymbol{s}') = \left(1 + e^{\Delta E/kT}\right)^{-1}$ - Glauber (heat-bath) dynamics
  - ... where $\Delta E \equiv E(\boldsymbol{s}') - E(\boldsymbol{s})$ is the energy difference between states

Below the thermodynamic limit: $N < \infty$

- Phase transition is "blurred"
  - The order parameter (and its derivatives) are *not* discontinuous at the critical point
  - Nor do susceptibility, correlation length, etc. diverge
- However, susceptibility does *peak* near the critical point; this furnishes a "working definition" for identifying a critical point in practice
  - As N increases, the susceptibility peak becomes sharper and approaches the true (thermodynamic limit) critical point
- *Critical slowing down*: for a dynamical system near a critical point, fluctuations in the dynamics generally become larger and larger (this is a thorny issue for simulation...)
  - Monte Carlo Markov simulations take longer and longer to approach stationarity
  - Statistical estimates become noisier
- Finite-system fluctuations break symmetry and ergodicity breaking!
  - Especially just on the ordered side of a critical point (again, a thorny issue for simulation studies)

Below the thermodynamic limit: $N < \infty$

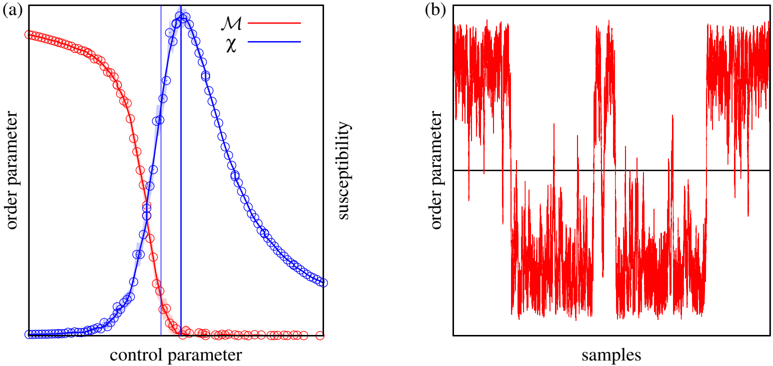

(a) "Blurring" of phase transition for finite systems: susceptibility doesn't diverge and peak (bold vertical line) is shifted from true (thermodynamic limit) critical point (light vertical line). (b) Symmetry breaking and ergodicity breaking fail: even though strictly in the ordered regime, the order parameter "flips" repeatedly. (2d lattice Ising model of size $N = 32 \times 32$ with Glauber dynamics.)

The 2d lattice Ising model

- The “fruit fly” of phase transitions - almost everything we know about phase transitions stems from analysis of this model
- 2-spin model: $s_i = \pm 1$, where i specifies a site on a 2d square lattice (usually wrapped to torus)

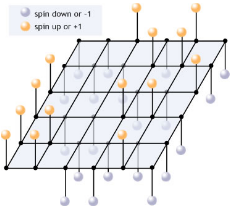

Boltzmann-Gibbs model with interaction energy (isotropic, zero external field)

$$E(\boldsymbol{s}) = -\sum_{<i,j>} s_i s_j$$

where $<i,j>$ denotes summation over lattice neighbours

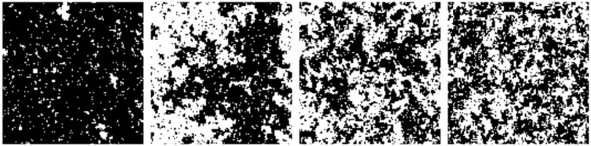

increasing temperature $T \longrightarrow$

The 2d lattice Ising model

- Has a 2nd order phase transition: model exactly solved (for zero external field) by Lars Onsager in 1944:

critical temperature	T_c	$=$	$\dfrac{2}{\log\left(1+\sqrt{2}\right)}$
magnetization	$\mathcal{M}$	$=$	$\pm\left(1-\sinh^{-4}2\beta\right)^{\frac{1}{8}}$ for $T < T_c$, $\quad \mathcal{M}=0$ for $T \ge T_c$
free energy	$-2\beta\mathcal{F}$	$=$	$\log\left(2\cosh^2 2\beta\right)+\dfrac{2}{\pi}\displaystyle\int_0^{\pi/2}\log\left(1+\sqrt{1-\kappa^2\sin^2\theta}\right)d\theta$
internal energy	$-\mathcal{U}$	$=$	$\coth 2\beta\left[1+\dfrac{2}{\pi}\left(\kappa\,\sinh 2\beta-1\right)\displaystyle\int_0^{\pi/2}\frac{d\theta}{\sqrt{1-\kappa^2\sin^2\theta}}\right]$

Thermodynamic limits for the 2d lattice Ising model: here $\beta \equiv 1/T$, $\kappa \equiv 2\tanh 2\beta\,\mathrm{sech}\,2\beta$.

Information theory: mutual information

- Entropy: X a random variable, $p_X(x)$ the density of X.

$$H(X) \equiv -\int p_X(x) \log p_X(x)\, dx$$

  - Intuition: "the uncertainty of X"

- Mutual information: X, Y jointly distributed

$$I(X:Y) \equiv H(X) + H(Y) - H(X,Y)$$

  - Intuition: "the degree to which X disambiguates Y (and vice-versa)"
  - $I(X:Y) = 0 \iff X, Y$ are *independent*

- Conditional versions: X, Y, Z jointly distributed

$$\begin{aligned} H(X|Z) &\equiv H(X,Z) - H(Z) \\ I(X:Y|Z) &\equiv H(X|Z) + H(Y|Z) - H(X,Y|Z) \end{aligned}$$

Information theory: multivariate information measures

- Multivariate system of random variables $\boldsymbol{X} = (X_1, \ldots, X_N)$ with adjacency relation (network)
- Mean pairwise (bivariate) mutual information

$$I_{pw}(\boldsymbol{X}) \equiv \frac{1}{\nu} \sum_{<i,j>} I(X_i, X_j)$$

  - $\nu =$ number of unique adjacent pairs
  - Intuition: "the degree to which adjacent variables are statistically independent"
  - $I_{pw}(\boldsymbol{X}) = 0 \iff$ all adjacent pairs X_i, X_j are independent
  - Note: no conditioning, so doesn't take into account *indirect* connections between variables
- Multi-information ("global")

$$I_{gl}(\boldsymbol{X}) \equiv \sum_i H(X_i) - H(\boldsymbol{X})$$

  - Intuition: "the degree to which the system $\boldsymbol{X}$ is statistically uncoupled"
  - $I_{gl}(\boldsymbol{X}) = 0 \iff X_1, \ldots, X_N$ are *mutually* independent
- For a highly ordered system entropies are low, so our information measures take low values. But for a highly *dis*ordered system variables behave independently, so measures are *again* low; thus we may expect them to peak at some intermediate order
  - We wouldn't be surprised if they peak at an order/disorder phase transition ...
  - *cf. integration/segregation balance* theory in neuroscience (Tononi)

Information theory: information flow (transfer)

- Transfer entropy (Schreiber) - $X(t), Y(t)$ jointly stationary* stochastic processes:

$$T(Y \to X) \equiv I\big(X(t) : Y^{(\ell)}(t) \,|\, X^{(\ell)}(t)\big)$$

  - $X^{(\ell)}(t) \equiv X(t-1), \ldots, X(t-\ell)$ denotes the *history* (past) of $X(t)$ of length ℓ
  - Intuition (i): "the degree to which the past of Y disambiguates current X, *given past* X"
  - Intuition (ii): "the quantity of information transferred from Y to X per unit time"
  - $T(Y \to X) = 0 \iff X$, conditional on its own past, is independent of the past of Y
  - Conditioning on the past of X is essential to take into account *shared history* between the processes

- Conditional version:

$$T(Y \to X \,|\, Z) \equiv I\big(X(t) : Y^{(\ell)}(t) \,|\, X^{(\ell)}(t), Z^{(\ell)}(t)\big)$$

Information theory: multivariate information flow measures

- Mean pairwise (bivariate) transfer entropy

$$T_{pw}(\boldsymbol{X}) \equiv \frac{1}{\nu} \sum_{<i,j>} T(X_j \to X_i)$$

  - Intuition: "the degree to which adjacent variables are dynamically independent"
  - Again, doesn't take into account indirect connections between variables
  - *Note*: some given number of lags ℓ is assumed

- Global transfer entropy

$$T_{gl}(\boldsymbol{X}) \equiv \frac{1}{N} \sum_{i} T(\boldsymbol{X} \to X_i)$$

  - Intuition (i): "the degree to which the system is dynamically uncoupled"
  - Intuition (ii): "density of gross information transfer per unit time"
    - *cf. causal density*: Seth *et al.* Proc. Natl. Acad. Sci. U.S.A. **103** 10799, 2006
  - $T_{gl}(\boldsymbol{X}) = 0 \iff$ each X_i, conditional on its own past, is independent of past $\boldsymbol{X}$; that is of the history of the entire system

- As for the "static" mutual information measures, these measures may also be expected to peak at an intermediate order
  - Again, we wouldn't be surprised if they peak at an order/disorder phase transition ...
  - ... but should we be surprised if they *don't*?

Information measures for the 2d lattice Ising model

- Analytic results
  - The mutual information measures in the thermodynamic limit may be calculated entirely analytically, using the Onsager results. We have

$$I_{pw} = -2\sum_{\sigma=\pm1} p_\sigma \log p_\sigma + \sum_{\sigma,\sigma'=\pm1} p_{\sigma\sigma'} \log p_{\sigma\sigma'}$$

$$\frac{1}{N} I_{gl} = -\sum_{\sigma} p_\sigma \log p_\sigma - \frac{1}{T}(\mathcal{U} - \mathcal{F})$$

    with $p_\sigma = \frac{1}{2}(1+\sigma\mathcal{M})$, $p_{\sigma\sigma'} = \frac{1}{4}[1 + (\sigma+\sigma')\mathcal{M} - \frac{1}{2}\sigma\sigma'\mathcal{U}]$ for $\sigma, \sigma' = \pm1$.

  - The finite-lag* transfer entropy measures for Glauber dynamics may be expressed in terms of ensemble averages that must be estimated in simulation:

$$NT_{pw} = -q \sum_{\sigma=\pm1} \log \frac{q}{p_\sigma} + \sum_{\sigma'=\pm1} q_{\sigma'} \sum_{\sigma=\pm1} \log \frac{q_{\sigma'}}{p_{\sigma\sigma'}}$$

$$NT_{gl} = -q \sum_{\sigma=\pm1} \log \frac{q}{p_\sigma} + r$$

    with $q = \frac{1}{2}\langle P_i(\boldsymbol{S})\rangle$, $q_{\sigma'} = \frac{1}{4}(\langle P_i(\boldsymbol{S})\rangle + \sigma' \langle S_j P_i(\boldsymbol{S})\rangle)$ and $r = \langle P_i(\boldsymbol{S}) \log P_i(\boldsymbol{S})\rangle$ for an arbitrary site i, where $P_i(\boldsymbol{S})$ denotes the Glauber spin-flip probability at site i for a random state $\boldsymbol{S}$ sampled from the Boltzmann-Gibbs distribution

*It is easy to show that in the thermodynamic limit, if the number $\ell < \infty$ of historical lags is fixed, then the result is the same as for 1 lag; however, taking the "long time range" infinite lags limit *before* proceeding to the thermodynamic limit has so far proved intractable. In this sense the measures as calculated are essentially short-range.

Information measures for the 2d lattice Ising model

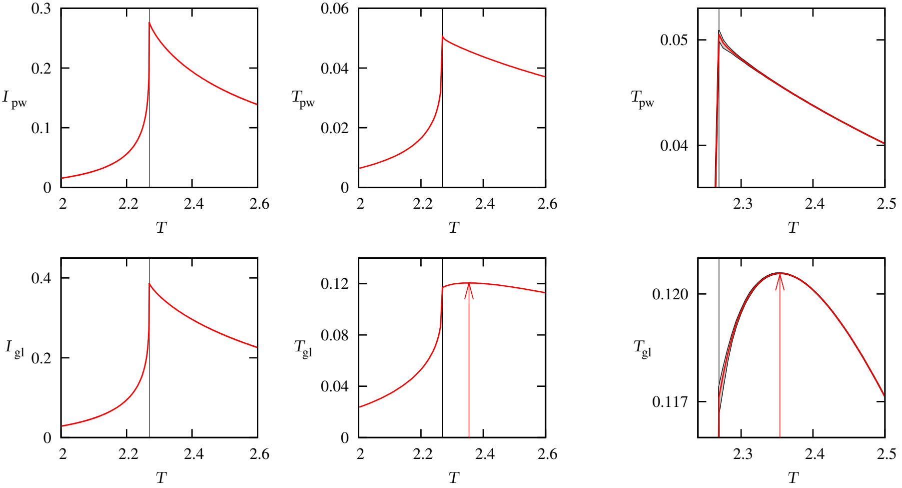

Information measures for the 2d lattice Ising model. Mutual information measures were calculated analytically. Transfer entropy measures were calculated for Glauber dynamics from large-scale simulations on a 512×512 toroidal lattice. From L. Barnett *et al.*, Phys. Rev. Lett. **111(17)**, 2013. The critical temperature is $T_c \approx 2.269$, while the T_{gl} peak is at ≈ 2.354.

Analysis

- Unlike the other measures, T_{gl} peaks *strictly in the disordered phase*
  - Tested with large-scale simulation to a high degree of statistical significance
  - Surprising: for the Ising model just about anything that peaks does so *at* the critical temperature
- From the paper:

> The pairwise measures incorporate putative statistical dependencies intermediated by the joint distribution of the remaining system elements with the spin pair in question. The global measures do not suffer from this effect; nonetheless, the static global measure I_{gl} also peaks at the phase transition, while the dynamic measure T_{gl} peaks strictly in the disordered regime. We conclude that a postcritical peak is not simply a consequence of accounting for common influences, nor is it a consequence alone of incorporating past-conditional dependencies; both factors are required.

- Intuitive explanation? Admission: we don't have a really good one ... from the paper:

> Preliminary analysis implicates a subtle interplay between differing contributions to T_{gl} from sites within and on the boundaries of same-spin domains, and the change in distribution of domain sizes as the temperature increases and domains disintegrate.

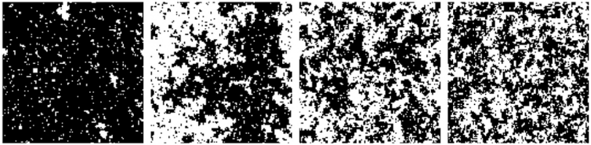

increasing temperature $T \longrightarrow$

Why is this significant?

- *Scenario*: We watch a "healthy" (disordered?) complex dynamical system over time, collating information statistics over a sliding window (in "shortish" time windows the process appears near-stationary). We observe both I_{pw} and T_{gl} to be slowly increasing. Then T_{gl} starts to fall, but I_{pw} continues to increase; we anticipate an imminent phase transition

- *Interpretation*: some (unidentified) control parameter is changing slowly, pushing the sytem towards a critical point

- *Note*: for real-world systems it is likely that we don't know what the control parameters are, nor may we have a sutitable candidate for an order parameter ... however, it has been demonstrated for a range of model systems that I_{pw} peaks at an order/disorder phase transition (e.g. Ising model, Vicsek particle swarm model, RBNs)

  And even if the I_{pw} peak does not coincide precisely with the phase transition, a shift in peak of T_{gl} with respect to I_{pw} may still furnish a useful indicator

  - In fact, a peak in I_{pw} may be your best bet in practice for *identifying* a phase transition

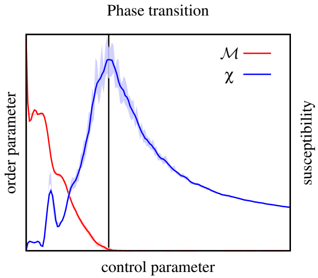

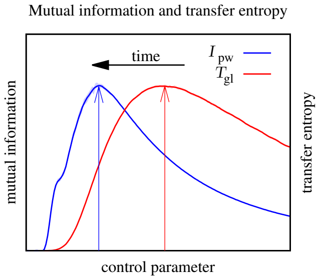

Practical issues

- Just how *universal* is this result? That is, to what class of systems/phase transitions is it likely to hold?
    - Current research by our group is clarifying this question
    - Strong dependence on interaction (network) topology
- Given stationary time-series data, I_{pw} is straightforward (if computationally costly for large N) to estimate in sample. However, it seems as if T_{gl}—even to 1 lag—will be problematic, as the *joint* distribution of states of all N system elements is implicated
    - For real-world data we don't have the luxury of the analytic result as for the Ising model
- If, however, (as in lattice spin models) interactions between elements scale $\ll N$ then the problem becomes more tractable
    - Example: particle swarm models where only nearby particles interact
- Another possibility is to take a lower-dimensional "proxy" (e.g. first few principal components) for the full system $\boldsymbol{X}$ in the $T(\boldsymbol{X} \to X_i)$ terms in T_{gl}
- Information-theoretic quantities are notoriously tricky to estimate in sample. One approach is to consider *linear models*, perhaps under Gaussian assumptions; thus mutual information corresponds to (generalised) *correlation*, while transfer entropy corresponds to *Granger causality* [Barnett *et al.*, Phys. Rev. Lett. **103(23)**, 2009]
- *Statistical significance* may be difficult to establish - but essential to distinguish behaviour (especially peaks) of the measures

Acknowledgements

- I would like to thank my collaborators on this study:

  Prof. Terry Bossomaier, *Centre for Research in Complex Systems, Charles Sturt University, Panorama Avenue, Bathurst, New South Wales 2795, Australia*

  Dr. Joseph T. Lizier, *CSIRO Computational Informatics, P.O. Box 76, Epping, New South Wales 1710, Australia*

  Dr. Michael Harré, *School of Civil Engineering, The University of Sydney, Sydney, New South Wales 2006, Australia*

  Prof. Anil K. Seth, *Sackler Centre for Consciousness Science, School of Informatics, University of Sussex, Brighton BN1 9QJ, UK*

- I would also like to thank the Sackler Centre for financial support