Correlations and signatures of criticality in neural population models

NIPS Workshop 2015:
Modelling and inference for dynamics on complex interaction networks: joining up machine learning and statistical physics

Jakob Macke

caesar
center of advanced european studies and research

research center caesar - an associate of the Max Planck Society

Bernstein Center for Computational Neuroscience Tübingen

Are neural networks poised at a *thermodynamic* critical point?

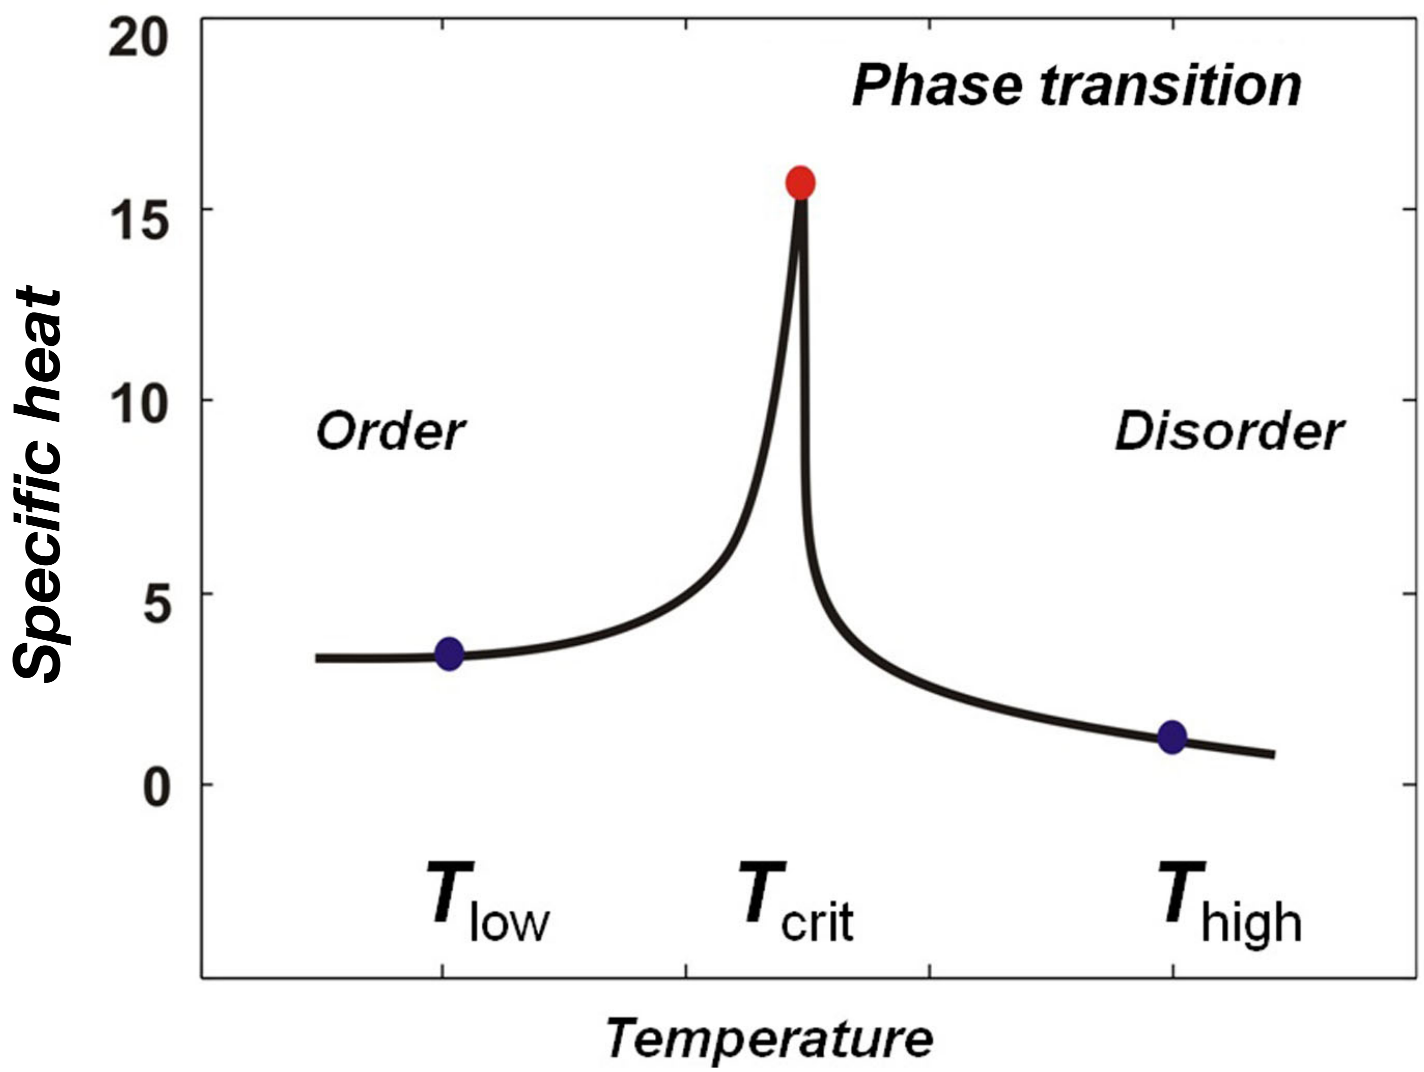

figure adapted from Beggs & Timme, 2012

Tkacik et al 2009, Mora & Bialek 2011, Stephens et al 2012, Yu et al 2013, Mora et al 2015, Tkacik et al 2015

Signatures of criticality in a recording of retinal ganglion cells: Are population codes ‘critical’?

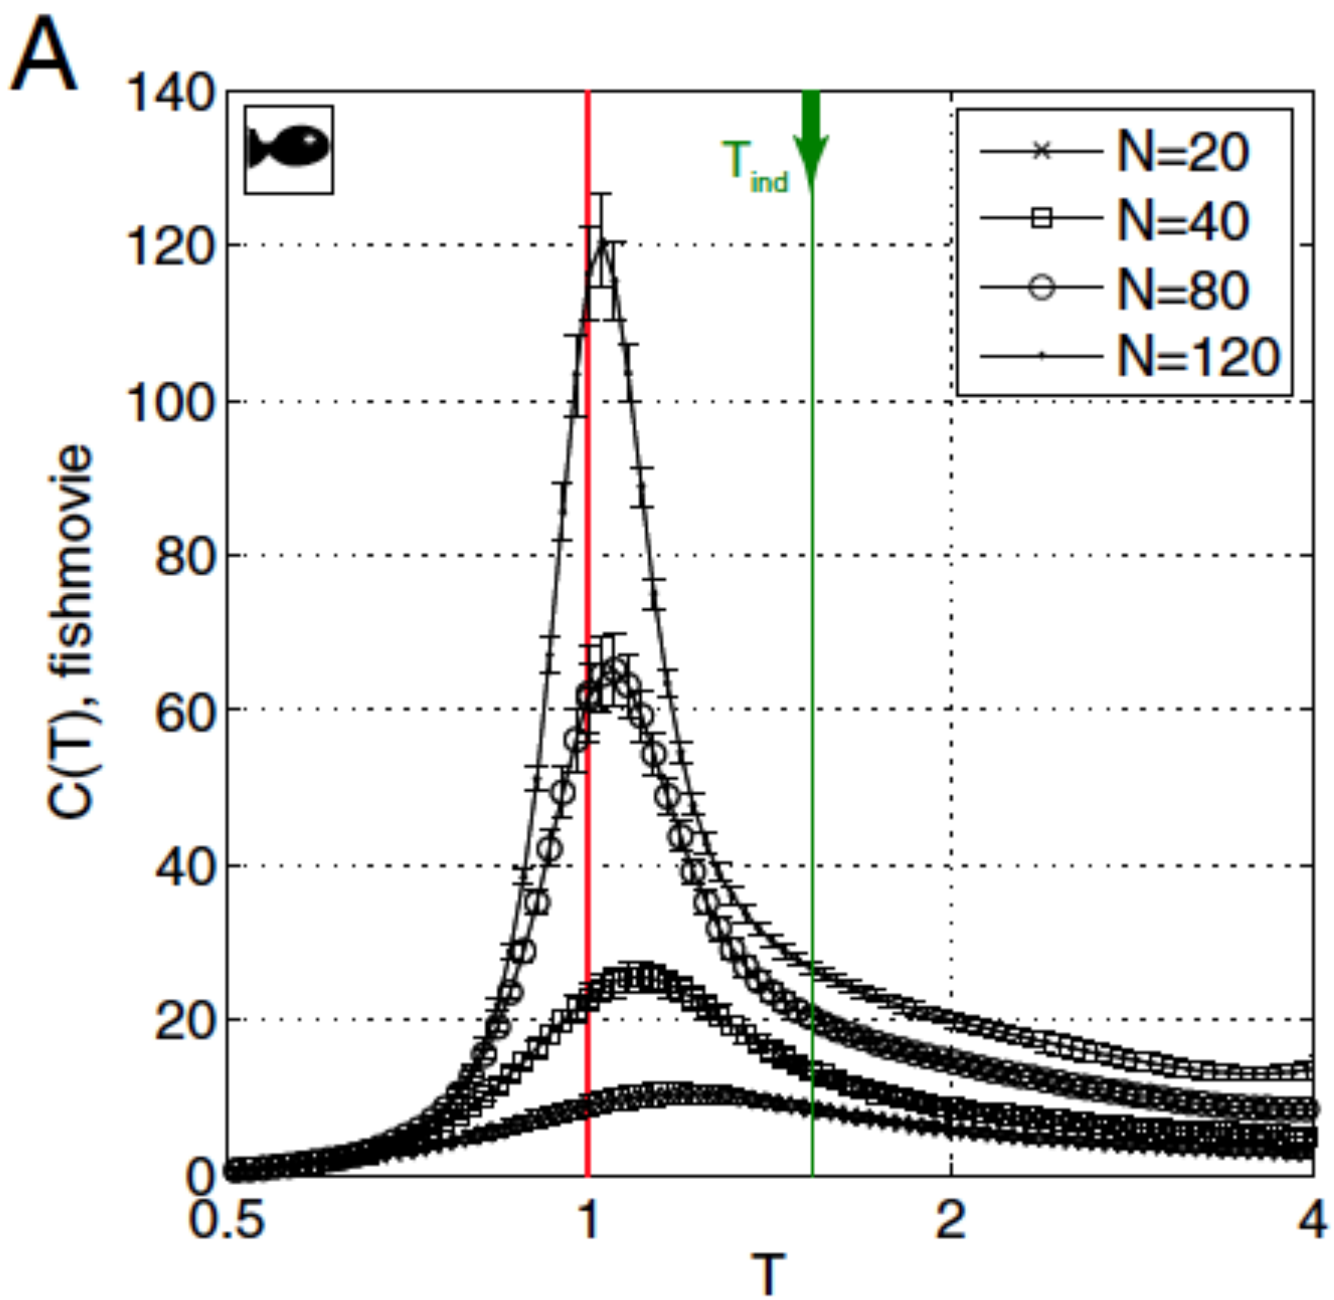

Tkacik, Mora, Marre, Amodei, Palmer, Berry, Bialek, PNAS 2015
Thermodynamics and signatures of criticality in a network of neurons

Tkacik et al 2009, Mora & Bialek 2011, Stephens et al 2012, Yu et al 2013, Mora et al 2015

Signatures of criticality in a recording of retinal ganglion cells: Are population codes 'critical'?

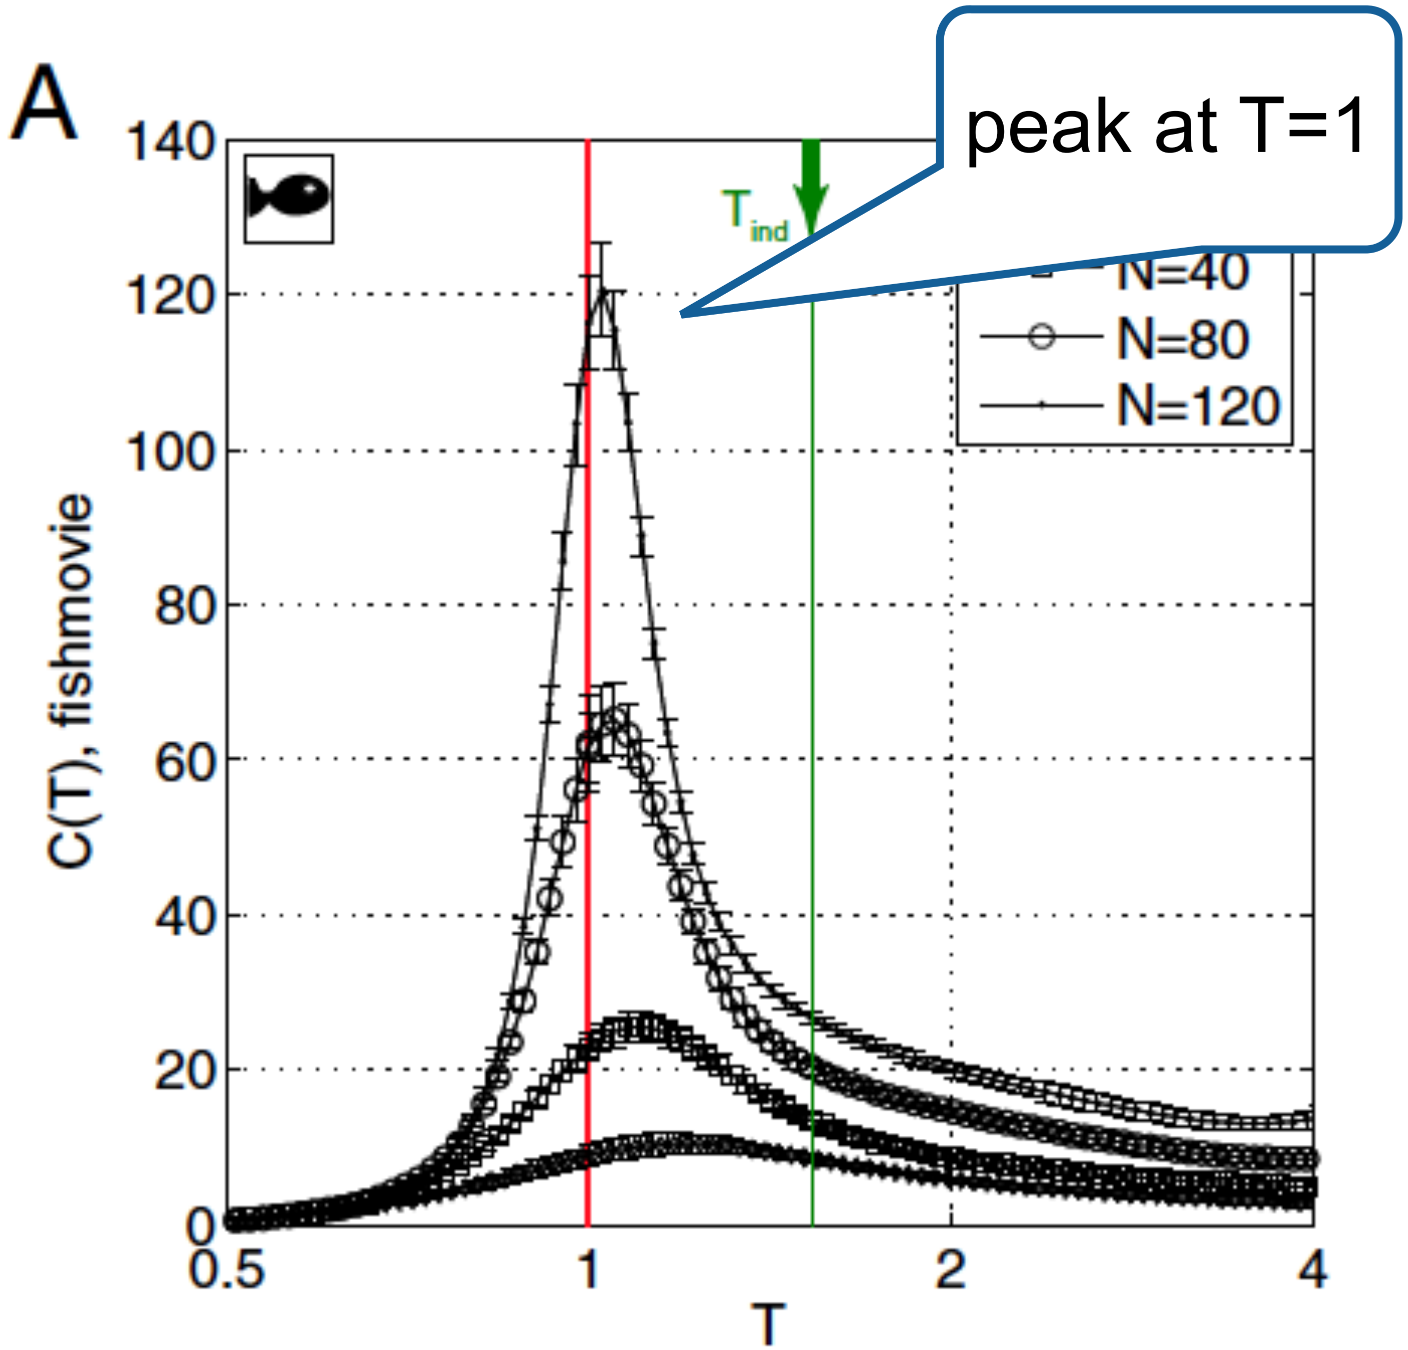

Tkacik, Mora, Marre, Amodei, Palmer, Berry, Bialek, PNAS 2015
Thermodynamics and signatures of criticality in a network of neurons

Tkacik et al 2009, Mora & Bialek 2011, Stephens et al 2012, Yu et al 2013, Mora et al 2015

Signatures of criticality in a recording of retinal ganglion cells: Are population codes ‘critical’?

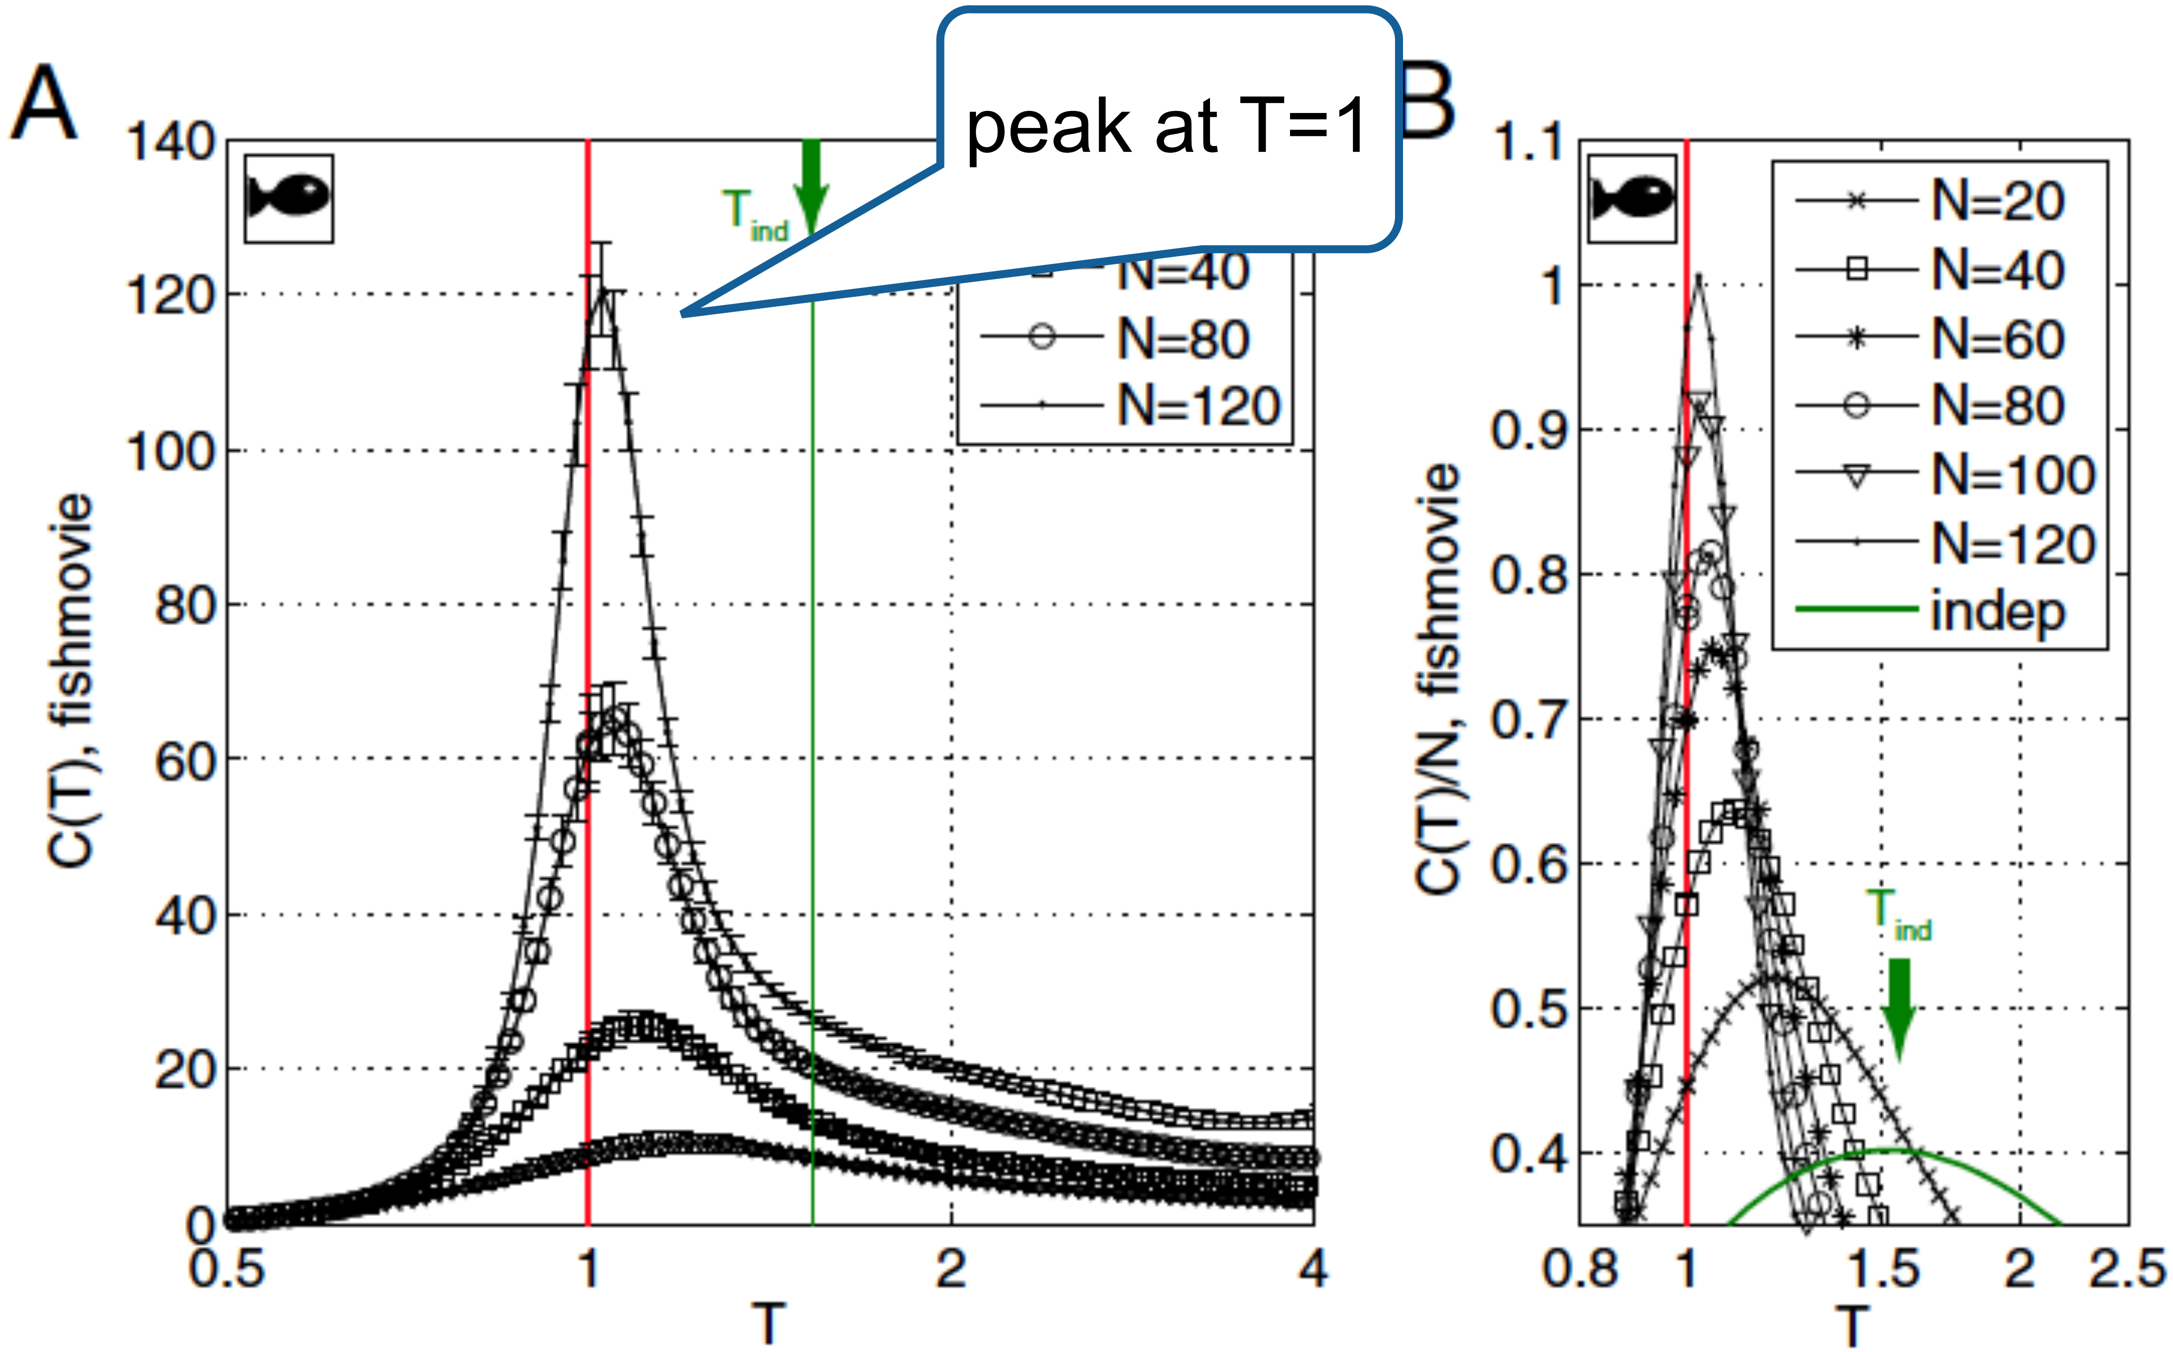

Tkacik, Mora, Marre, Amodei, Palmer, Berry, Bialek, PNAS 2015
Thermodynamics and signatures of criticality in a network of neurons

Tkacik et al 2009, Mora & Bialek 2011, Stephens et al 2012, Yu et al 2013, Mora et al 2015

Signatures of criticality in a recording of retinal ganglion cells: Are population codes 'critical'?

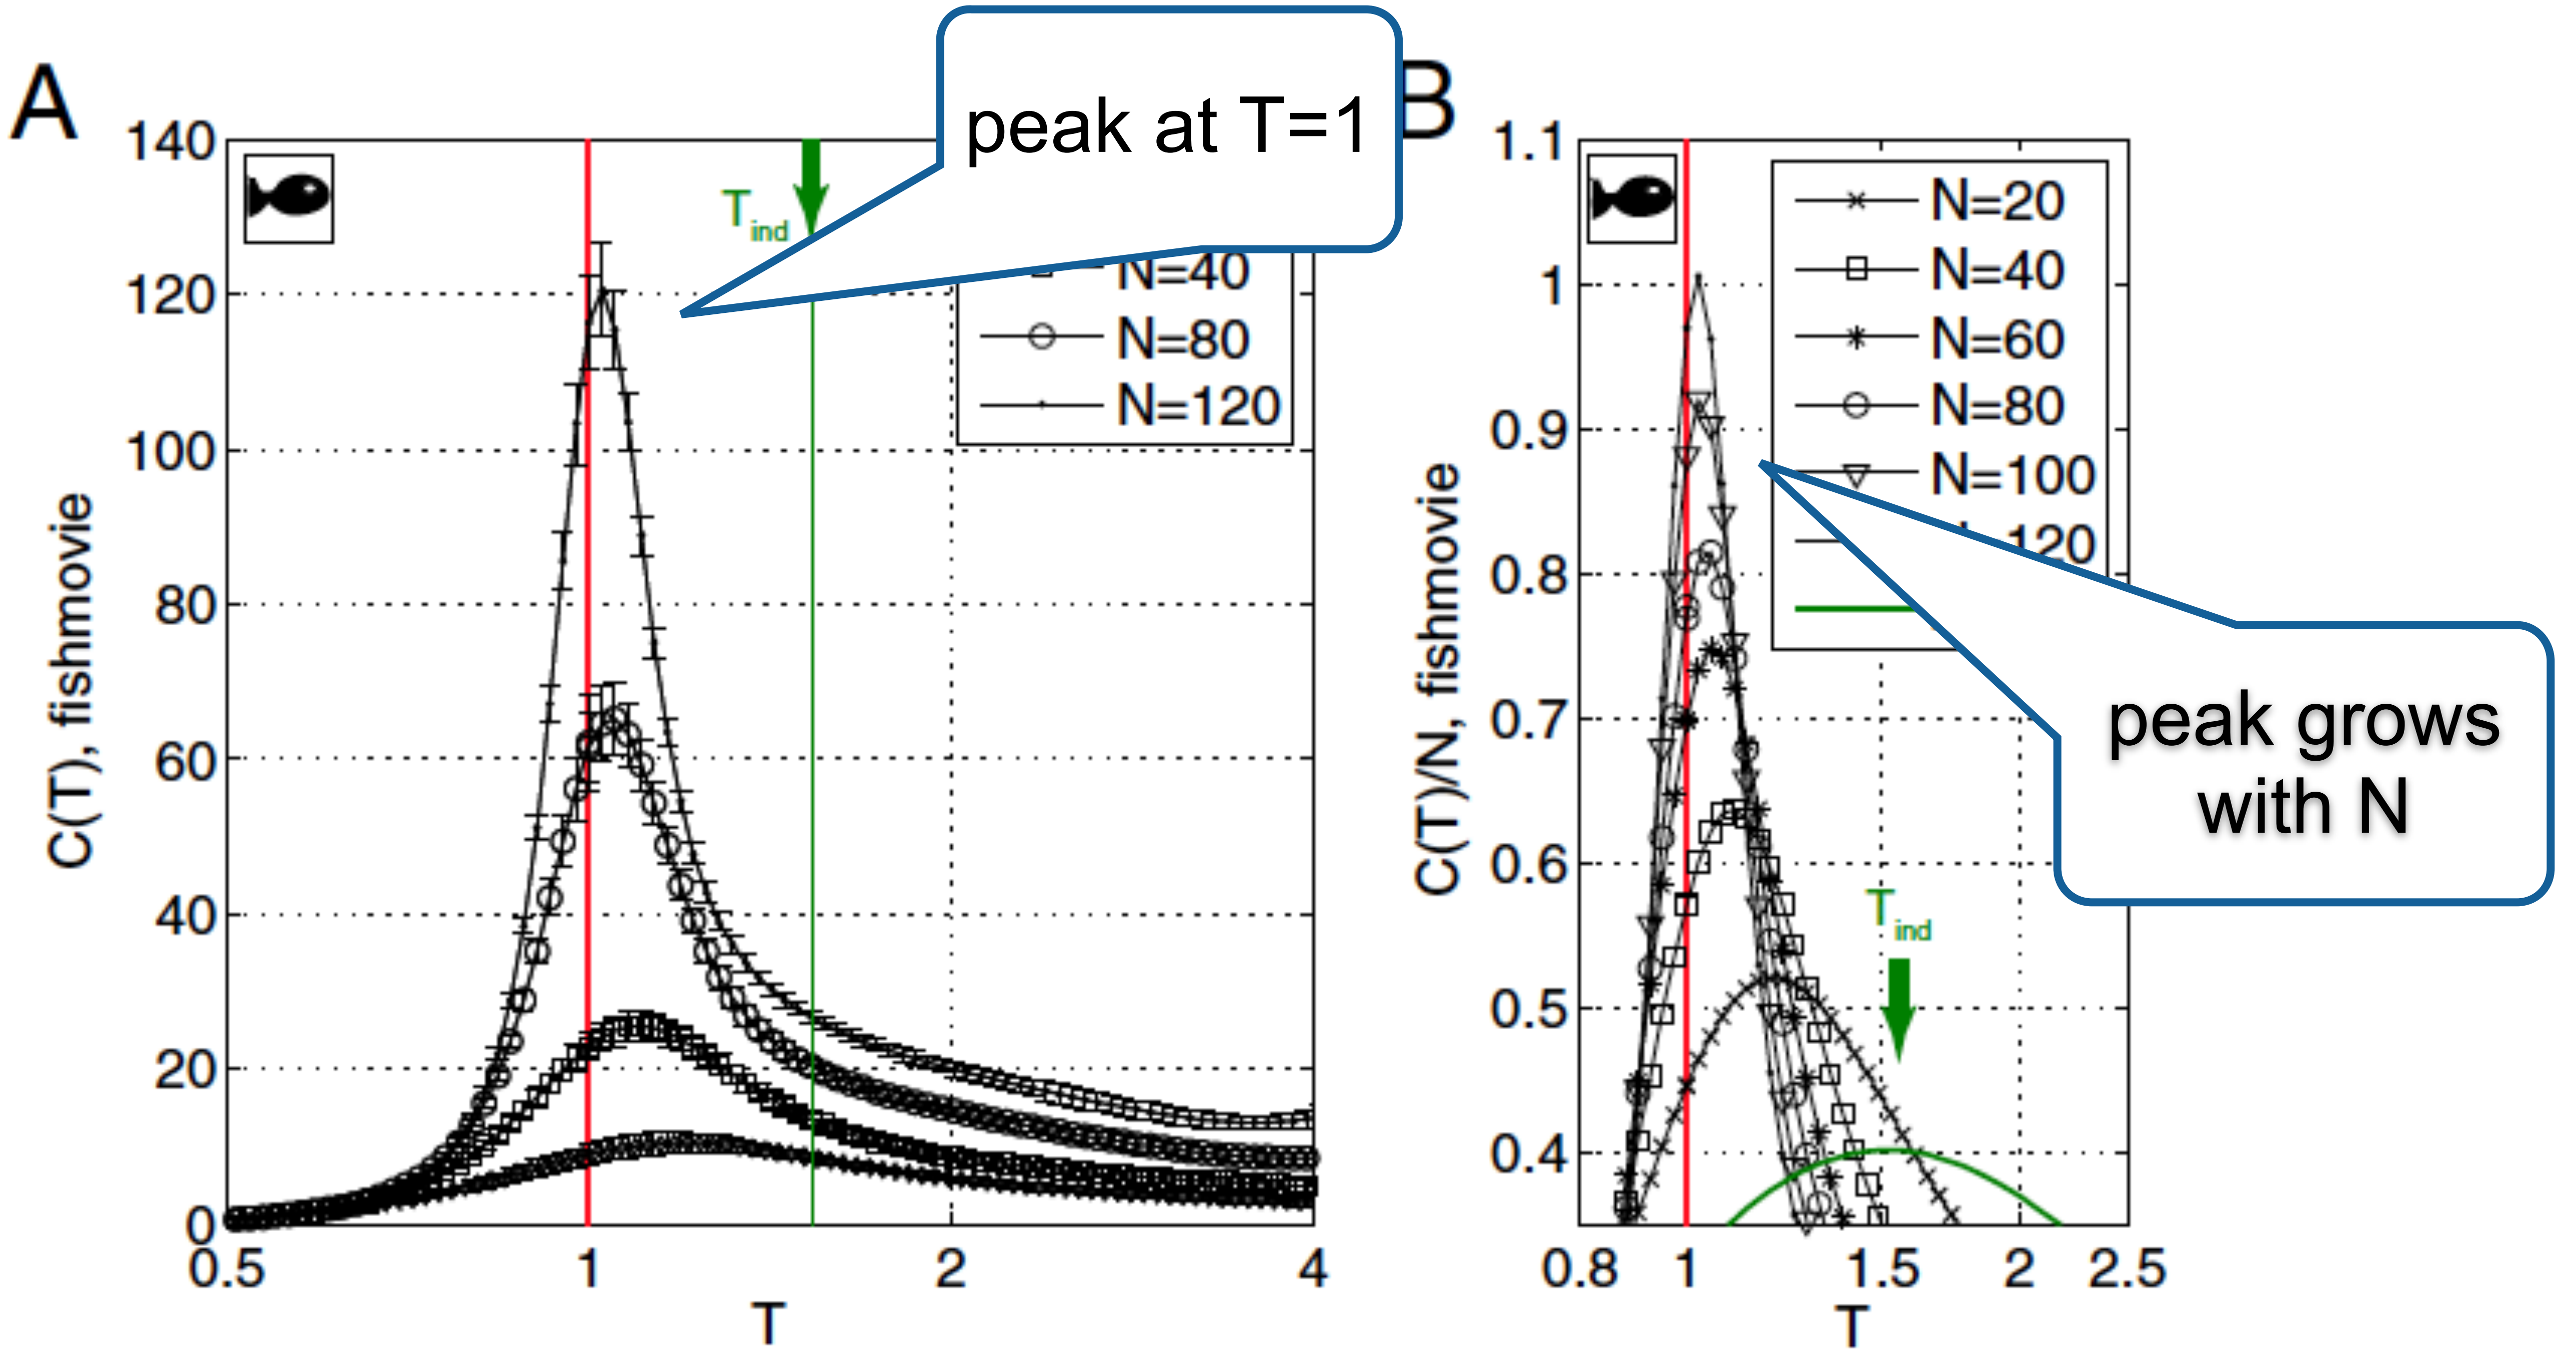

Tkacik, Mora, Marre, Amodei, Palmer, Berry, Bialek, PNAS 2015
Thermodynamics and signatures of criticality in a network of neurons

Tkacik et al 2009, Mora & Bialek 2011, Stephens et al 2012, Yu et al 2013, Mora et al 2015

We consider the distribution P(x) of binary 'spike-words' x.

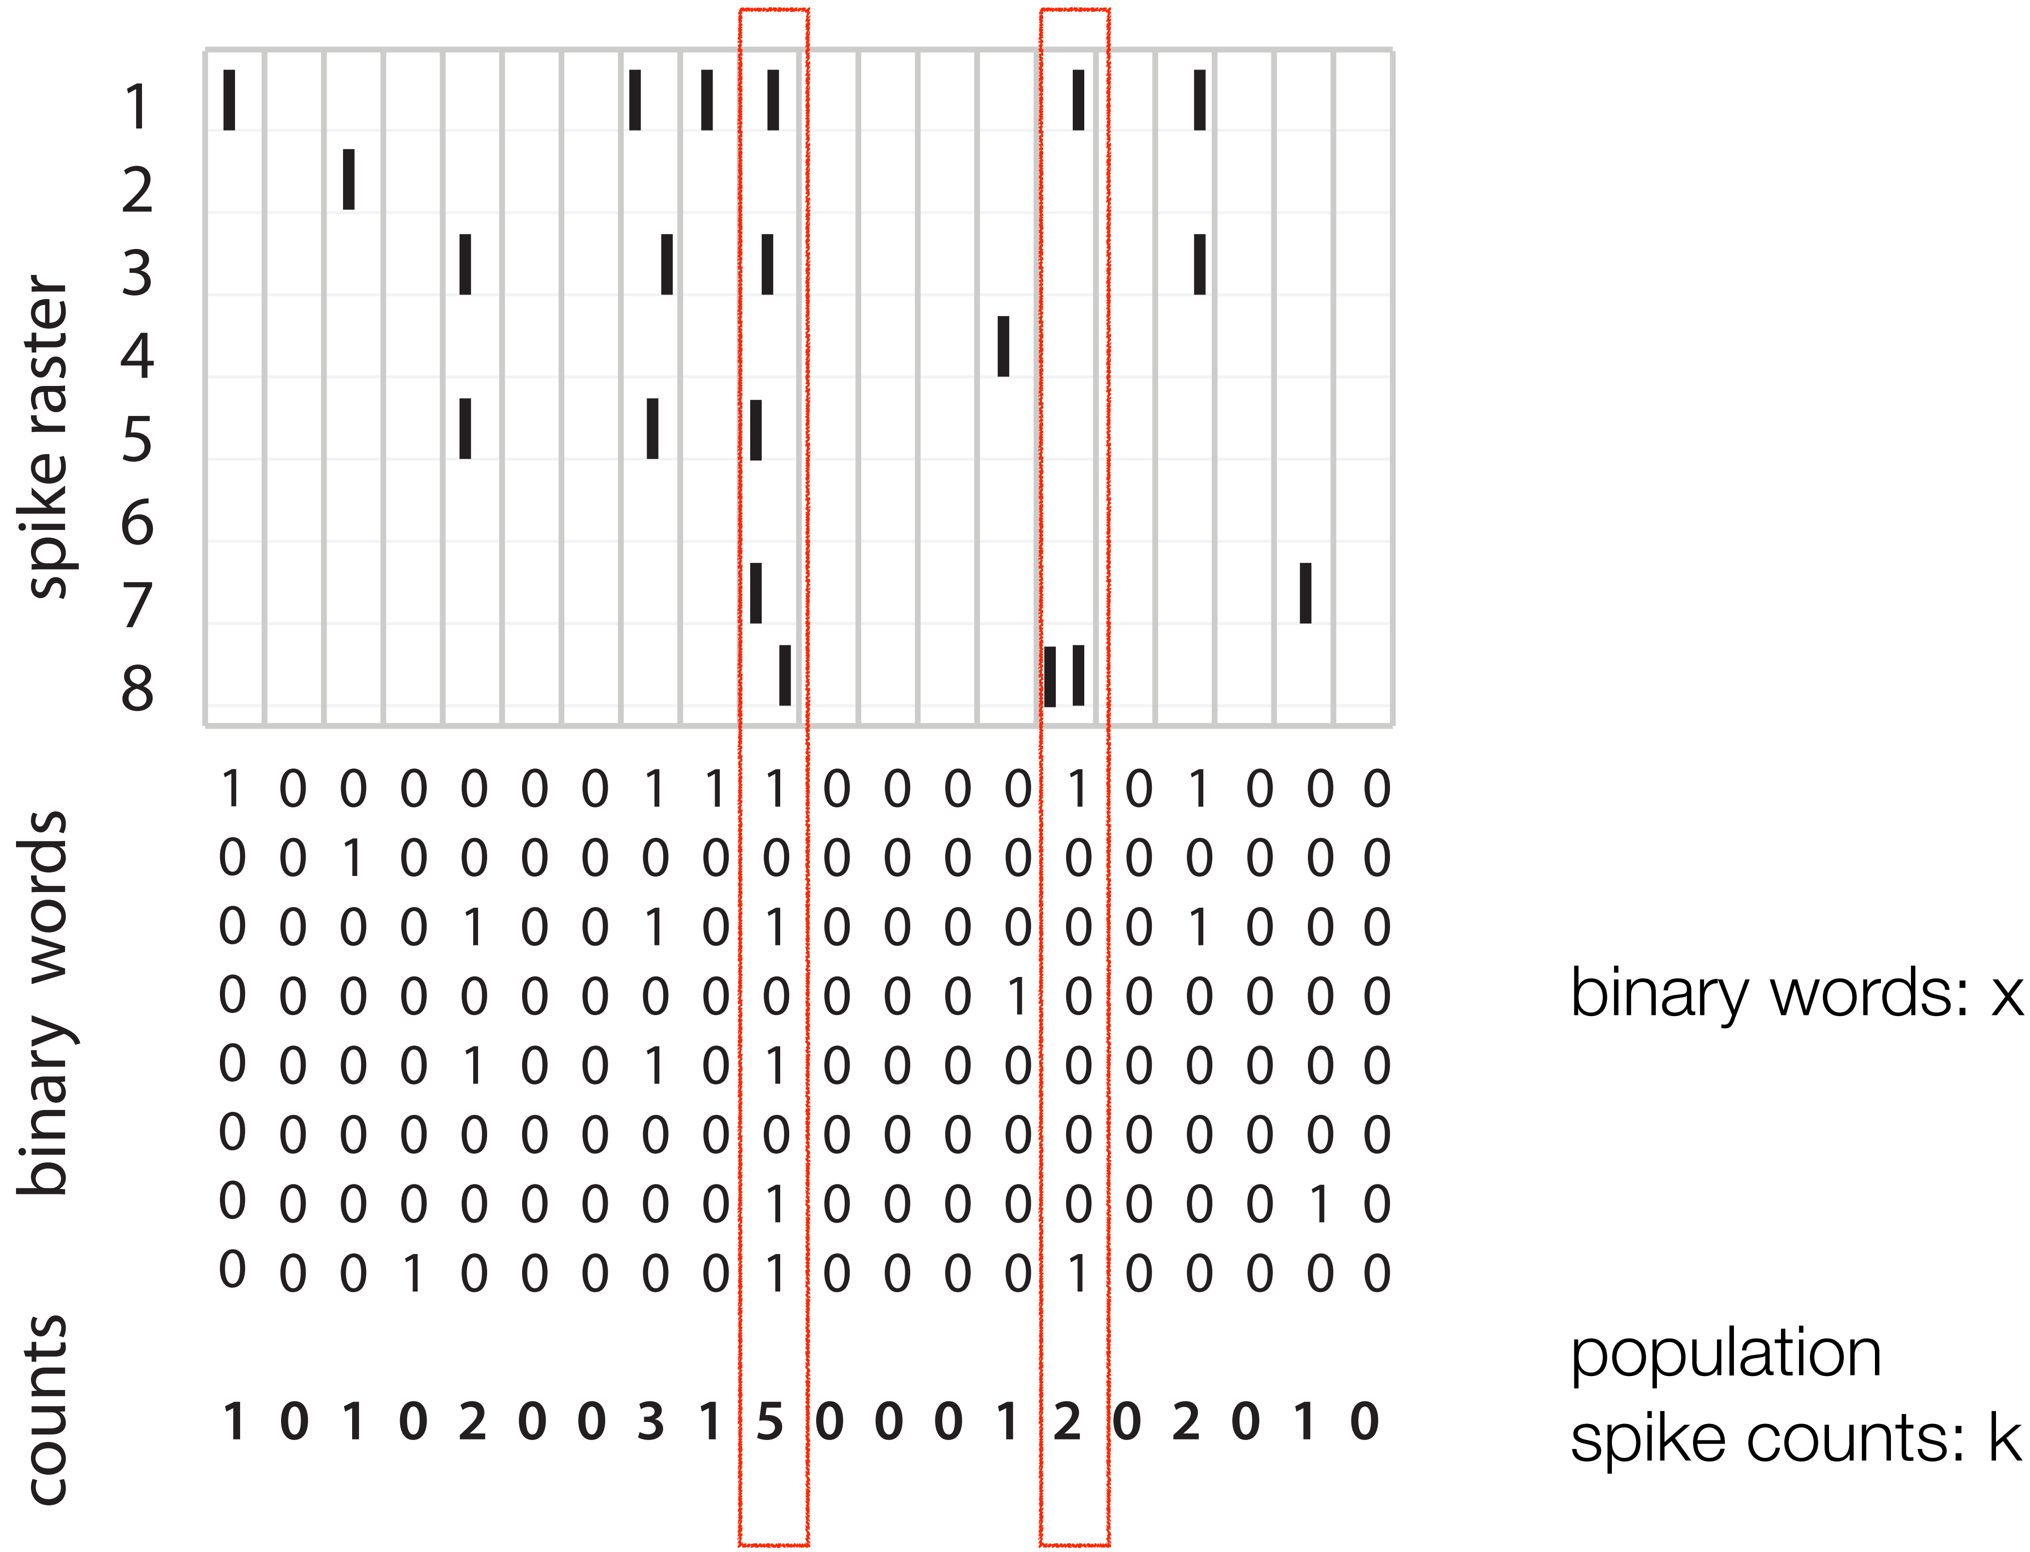

‘K-pairwise model’: An Ising model with spike-count constraints.

$$P_{full}(x) = \frac{1}{Z} \exp\left(h^\top x + x^\top J x + V_{\|x\|}\right)$$

Tkacik et al, Plos CB 2014

‘K-pairwise model’: An Ising model with spike-count constraints.

$$P_{full}(x) = \frac{1}{Z} \exp\left(h^\top x + x^\top J x + V_{||x||}\right)$$

N params to capture firing rates

Tkacik et al, Plos CB 2014

‘K-pairwise model’: An Ising model with spike-count constraints.

$$P_{full}(x) = \frac{1}{Z} \exp\left(h^\top x + x^\top J x + V_{||x||}\right)$$

N params to capture firing rates

N(N-1)/2 params to capture correlations

Tkacik et al, Plos CB 2014

‘K-pairwise model’: An Ising model with spike-count constraints.

$$P_{full}(x) = \frac{1}{Z} \exp\left(h^\top x + x^\top J x + V_{||x||}\right)$$

N params to capture firing rates

N(N-1)/2 params to capture correlations

N params to capture spike-count distribution

Tkacik et al, Plos CB 2014

‘K-pairwise model’: An Ising model with spike-count constraints.

$$P_{full}(x) = \frac{1}{Z} \exp\left(h^\top x + x^\top J x + V_{\|x\|}\right)$$

N params to capture firing rates

N(N-1)/2 params to capture correlations

N params to capture spike-count distribution

Pairwise model fails to capture spike-count distribution for large N

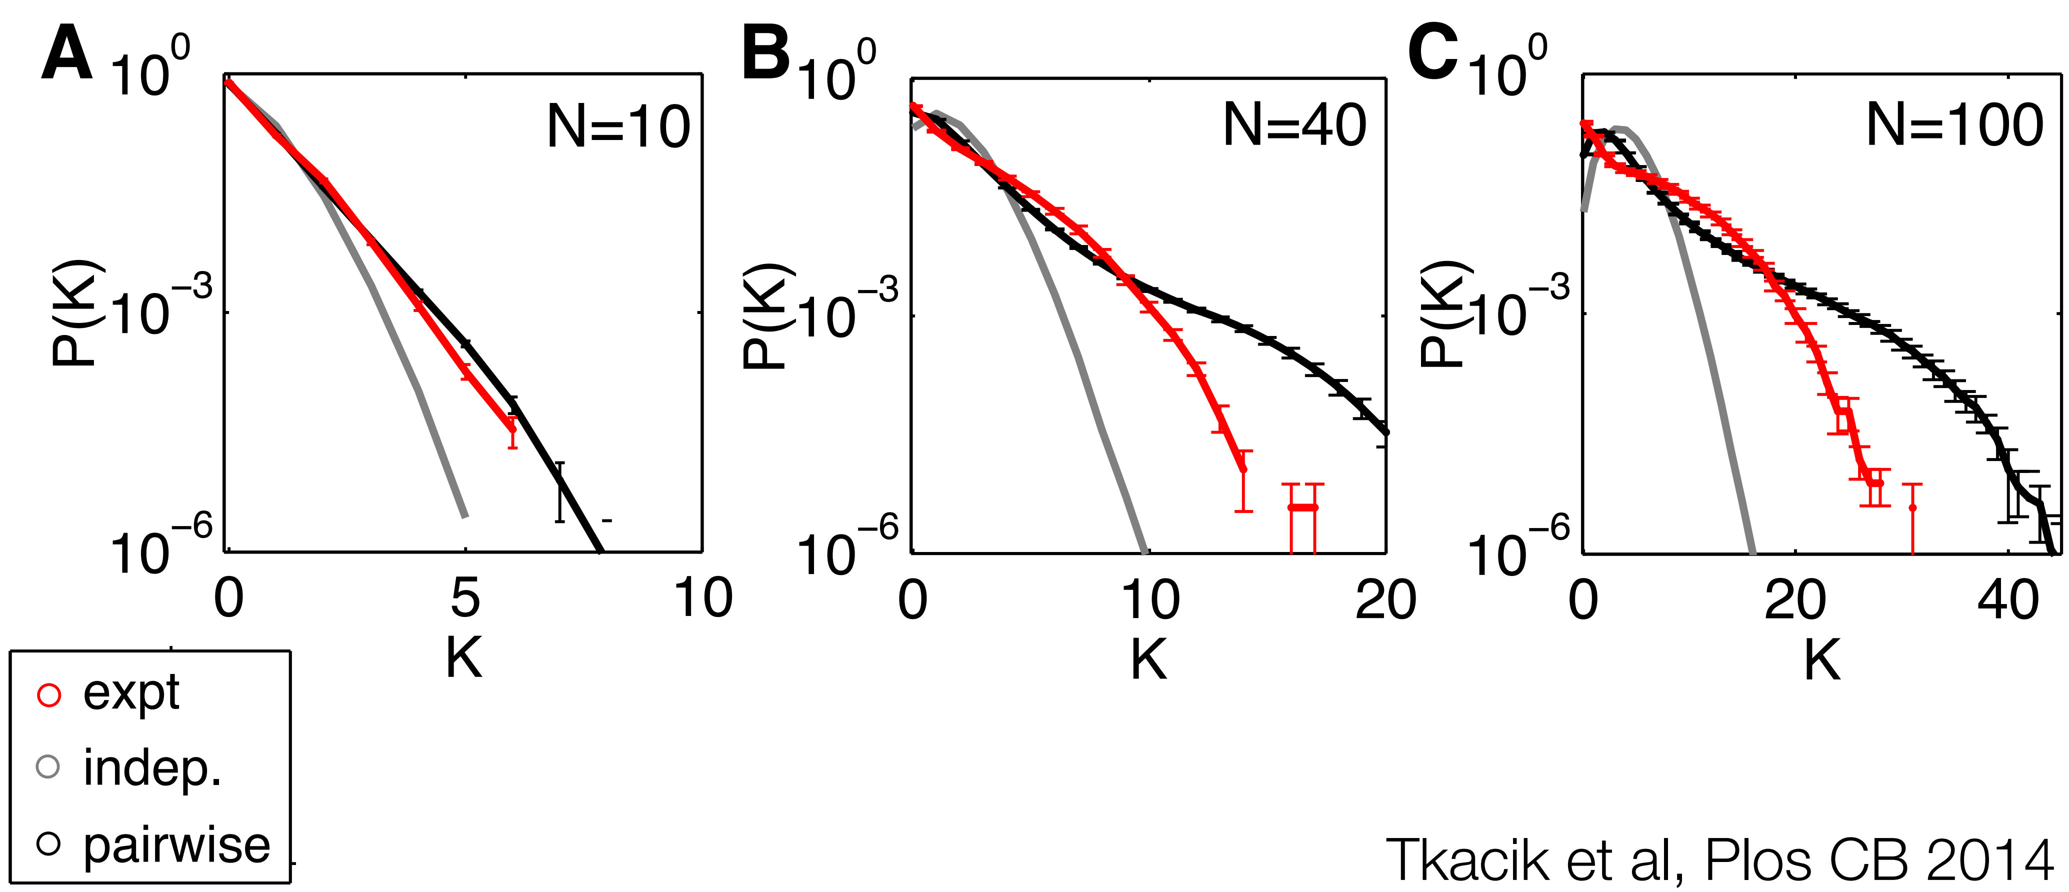

Tkacik et al, Plos CB 2014

Maximum entropy models can be used to translate concepts from thermodynamics to neural data.

Tkacik, Bialek et al

Maximum entropy models can be used to translate concepts from thermodynamics to neural data.

Probability: $$P_T(x) = \frac{1}{Z_T} \exp\left(-E(x)/T\right)$$

Tkacik, Bialek et al

Maximum entropy models can be used to translate concepts from thermodynamics to neural data.

Probability:

$$P_T(x) = \frac{1}{Z_T} \exp\left(-E(x)/T\right)$$

Heat capacity:

$$C_T = \frac{\partial \langle E(x) \rangle}{\partial T} = \frac{\mathrm{Var}\langle E(x) \rangle}{T^2}$$

Tkacik, Bialek et al

Maximum entropy models can be used to translate concepts from thermodynamics to neural data.

Probability:

$$P_T(x) = \frac{1}{Z_T} \exp\left(-E(x)/T\right)$$

Heat capacity:

$$C_T = \frac{\partial \langle E(x) \rangle}{\partial T} = \frac{\mathrm{Var}\langle E(x) \rangle}{T^2}$$

Thermo-dynamics

Tkacik, Bialek et al

Maximum entropy models can be used to translate concepts from thermodynamics to neural data.

Probability:

$$P_T(x) = \frac{1}{Z_T} \exp\left(-E(x)/T\right)$$

Heat capacity:

$$C_T = \frac{\partial \langle E(x) \rangle}{\partial T} = \frac{\mathrm{Var}\langle E(x) \rangle}{T^2}$$

Thermo-dynamics

Rewrite energy:

$$E(x) = -T \log P_T(x) - T \log Z_T$$

$$E(x) = -\log P(x) - \log Z$$

Tkacik, Bialek et al

Maximum entropy models can be used to translate concepts from thermodynamics to neural data.

Probability:

$$P_T(x) = \frac{1}{Z_T} \exp\left(-E(x)/T\right)$$

Heat capacity:

$$C_T = \frac{\partial \langle E(x) \rangle}{\partial T} = \frac{\mathrm{Var}\langle E(x) \rangle}{T^2}$$

Thermo-
dynamics

Rewrite energy:

$$E(x) = -T \log P_T(x) - T \log Z_T$$

$$E(x) = -\log P(x) - \log Z$$

Calculating P(x) at T:

$$P_T(x) = \frac{1}{\hat{Z}_T} \exp\left(\log P(x)/T\right) \propto P(x)^{1/T}$$

Tkacik, Bialek et al

Maximum entropy models can be used to translate concepts from thermodynamics to neural data.

Probability:

$$P_T(x) = \frac{1}{Z_T} \exp\left(-E(x)/T\right)$$

Heat capacity:

$$C_T = \frac{\partial \langle E(x) \rangle}{\partial T} = \frac{\mathrm{Var}\langle E(x) \rangle}{T^2}$$

Thermo-dynamics

Rewrite energy:

$$E(x) = -T \log P_T(x) - T \log Z_T$$

$$E(x) = -\log P(x) - \log Z$$

Calculating P(x) at T:

$$P_T(x) = \frac{1}{\hat{Z}_T} \exp\left(\log P(x)/T\right) \propto P(x)^{1/T}$$

Substitute into heat capacity:

$$C_T = \mathrm{Var}\langle \log P_T(x) \rangle$$

Tkacik, Bialek et al

Maximum entropy models can be used to translate concepts from thermodynamics to neural data.

Probability:

$$P_T(x) = \frac{1}{Z_T} \exp\left(-E(x)/T\right)$$

Heat capacity:

$$C_T = \frac{\partial \langle E(x) \rangle}{\partial T} = \frac{\mathrm{Var}\langle E(x) \rangle}{T^2}$$

Thermo-dynamics

Rewrite energy:

$$E(x) = -T \log P_T(x) - T \log Z_T$$

$$E(x) = -\log P(x) - \log Z$$

Calculating P(x) at T:

$$P_T(x) = \frac{1}{\hat{Z}_T} \exp\left(\log P(x)/T\right) \propto P(x)^{1/T}$$

Substitute into heat capacity: $C_T = \mathrm{Var}\langle \log P_T(x) \rangle$

Specific heat:

$$c_T = \frac{1}{N} C_T$$

Tkacik, Bialek et al

Maximum entropy models can be used to translate concepts from thermodynamics to neural data.

Probability:

$$P_T(x) = \frac{1}{Z_T} \exp\left(-E(x)/T\right)$$

Heat capacity:

$$C_T = \frac{\partial \langle E(x) \rangle}{\partial T} = \frac{\mathrm{Var}\langle E(x) \rangle}{T^2}$$

Thermo-dynamics

Rewrite energy:

$$E(x) = -T \log P_T(x) - T \log Z_T$$

$$E(x) = -\log P(x) - \log Z$$

Calculating P(x) at T:

$$P_T(x) = \frac{1}{\hat{Z}_T} \exp\left(\log P(x)/T\right) \propto P(x)^{1/T}$$

Substitute into heat capacity:

$$C_T = \mathrm{Var}\langle \log P_T(x) \rangle$$

Specific heat:

$$c_T = \frac{1}{N} C_T$$

Quantities measurable from statistical model

Tkacik, Bialek et al

Signatures of criticality in a recording of retinal ganglion cells: Are population codes ‘critical’?

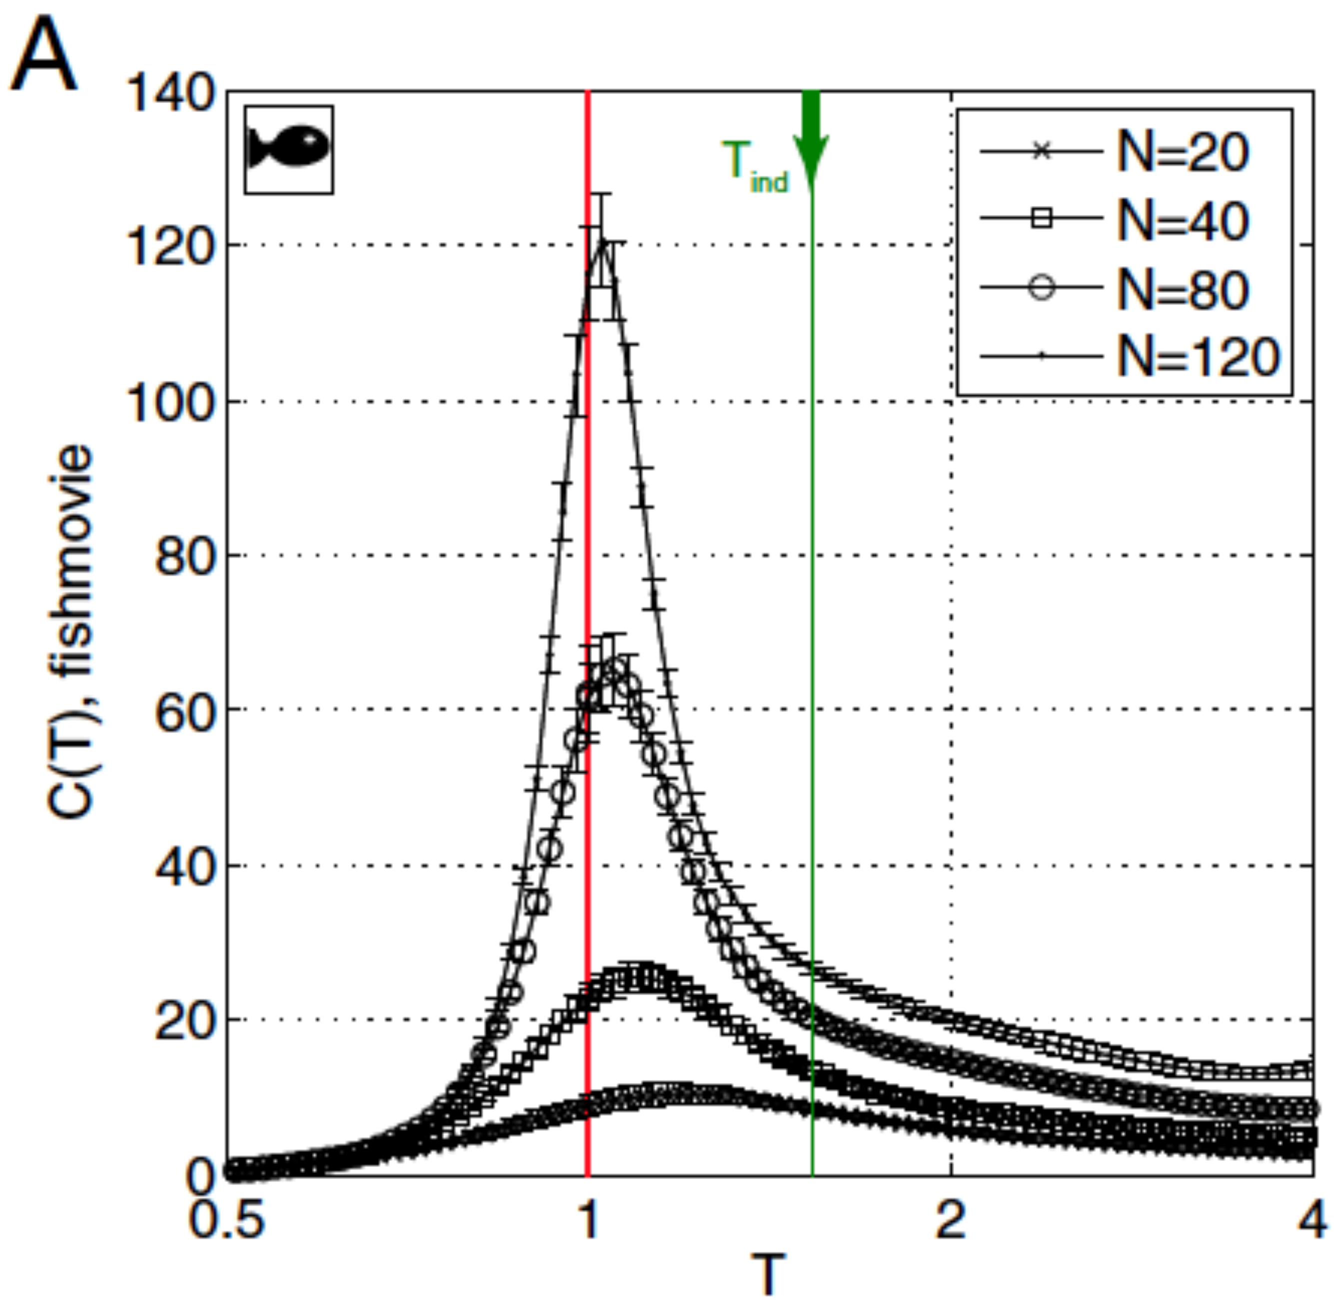

Tkacik, Mora, Marre, Amodei, Palmer, Berry, Bialek, PNAS 2015
Thermodynamics and signatures of criticality in a network of neurons

Tkacik et al 2009, Mora & Bialek 2011, Stephens et al 2012, Yu et al 2013, Mora et al 2015

Signatures of criticality in a recording of retinal ganglion cells: Are population codes ‘critical’?

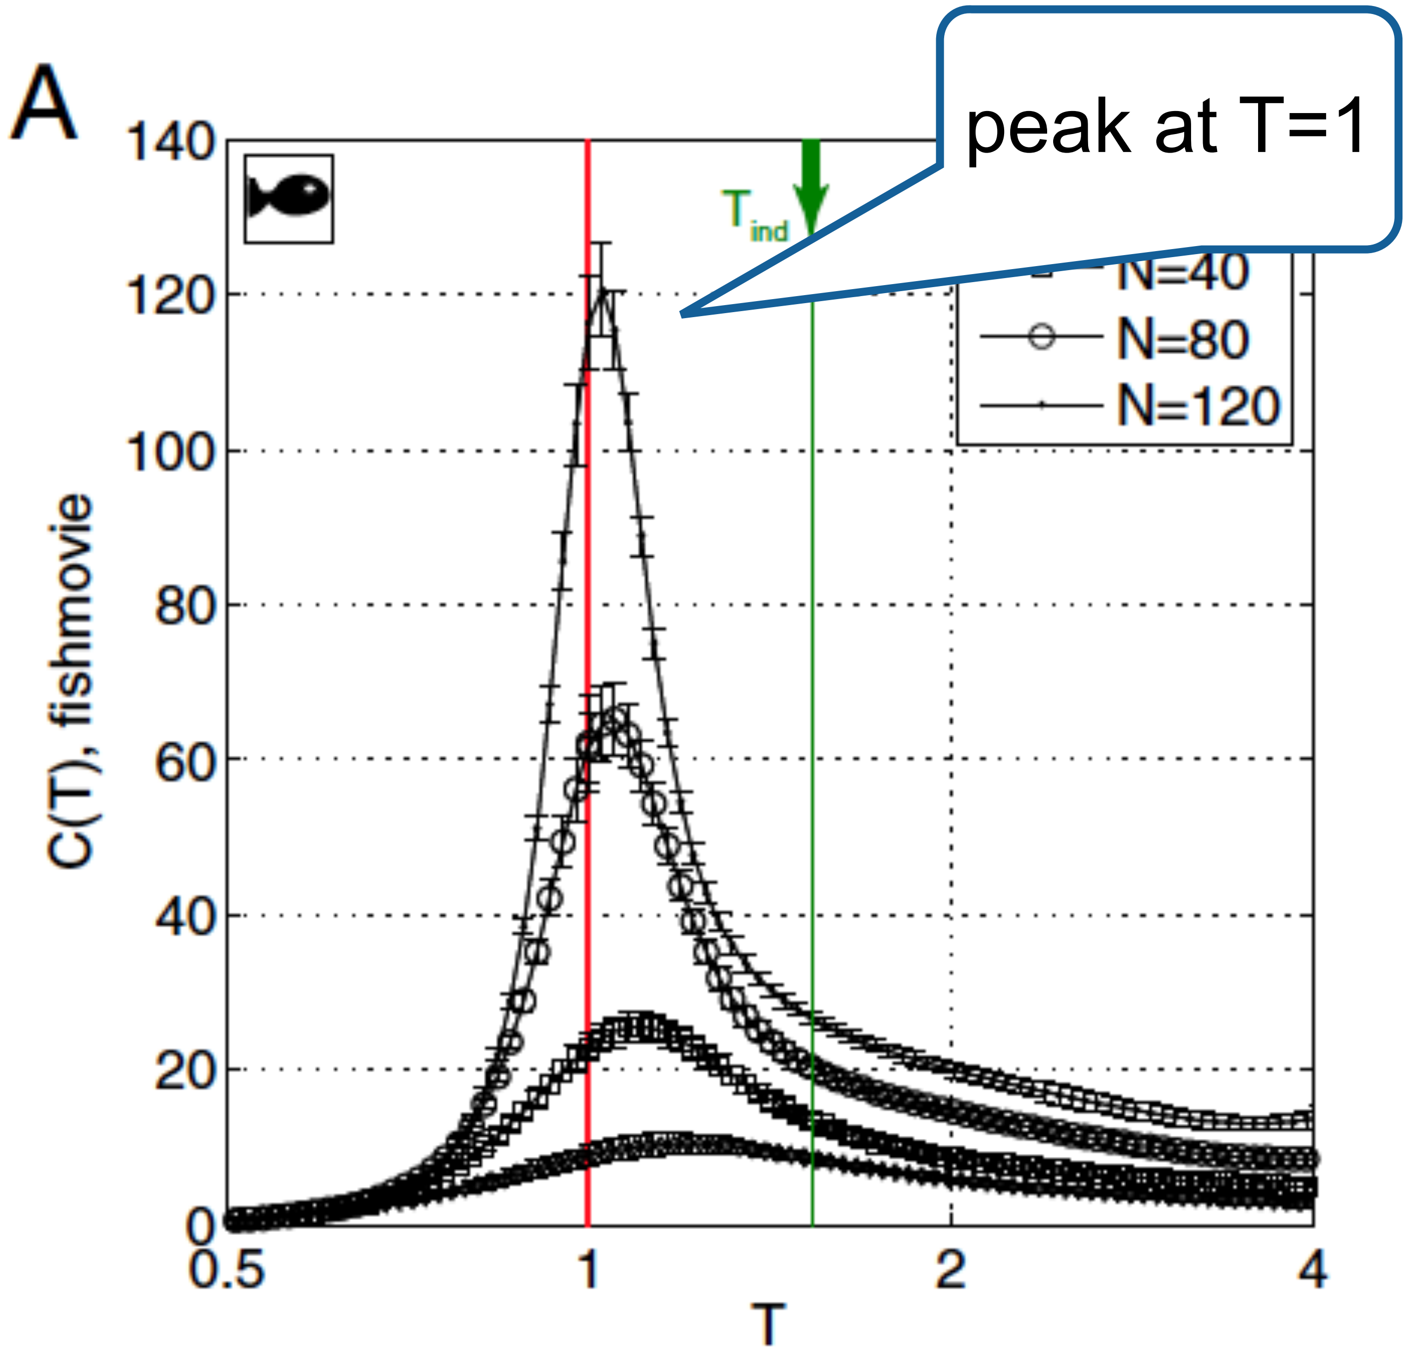

Tkacik, Mora, Marre, Amodei, Palmer, Berry, Bialek, PNAS 2015
Thermodynamics and signatures of criticality in a network of neurons

Tkacik et al 2009, Mora & Bialek 2011, Stephens et al 2012, Yu et al 2013, Mora et al 2015

Signatures of criticality in a recording of retinal ganglion cells: Are population codes ‘critical’?

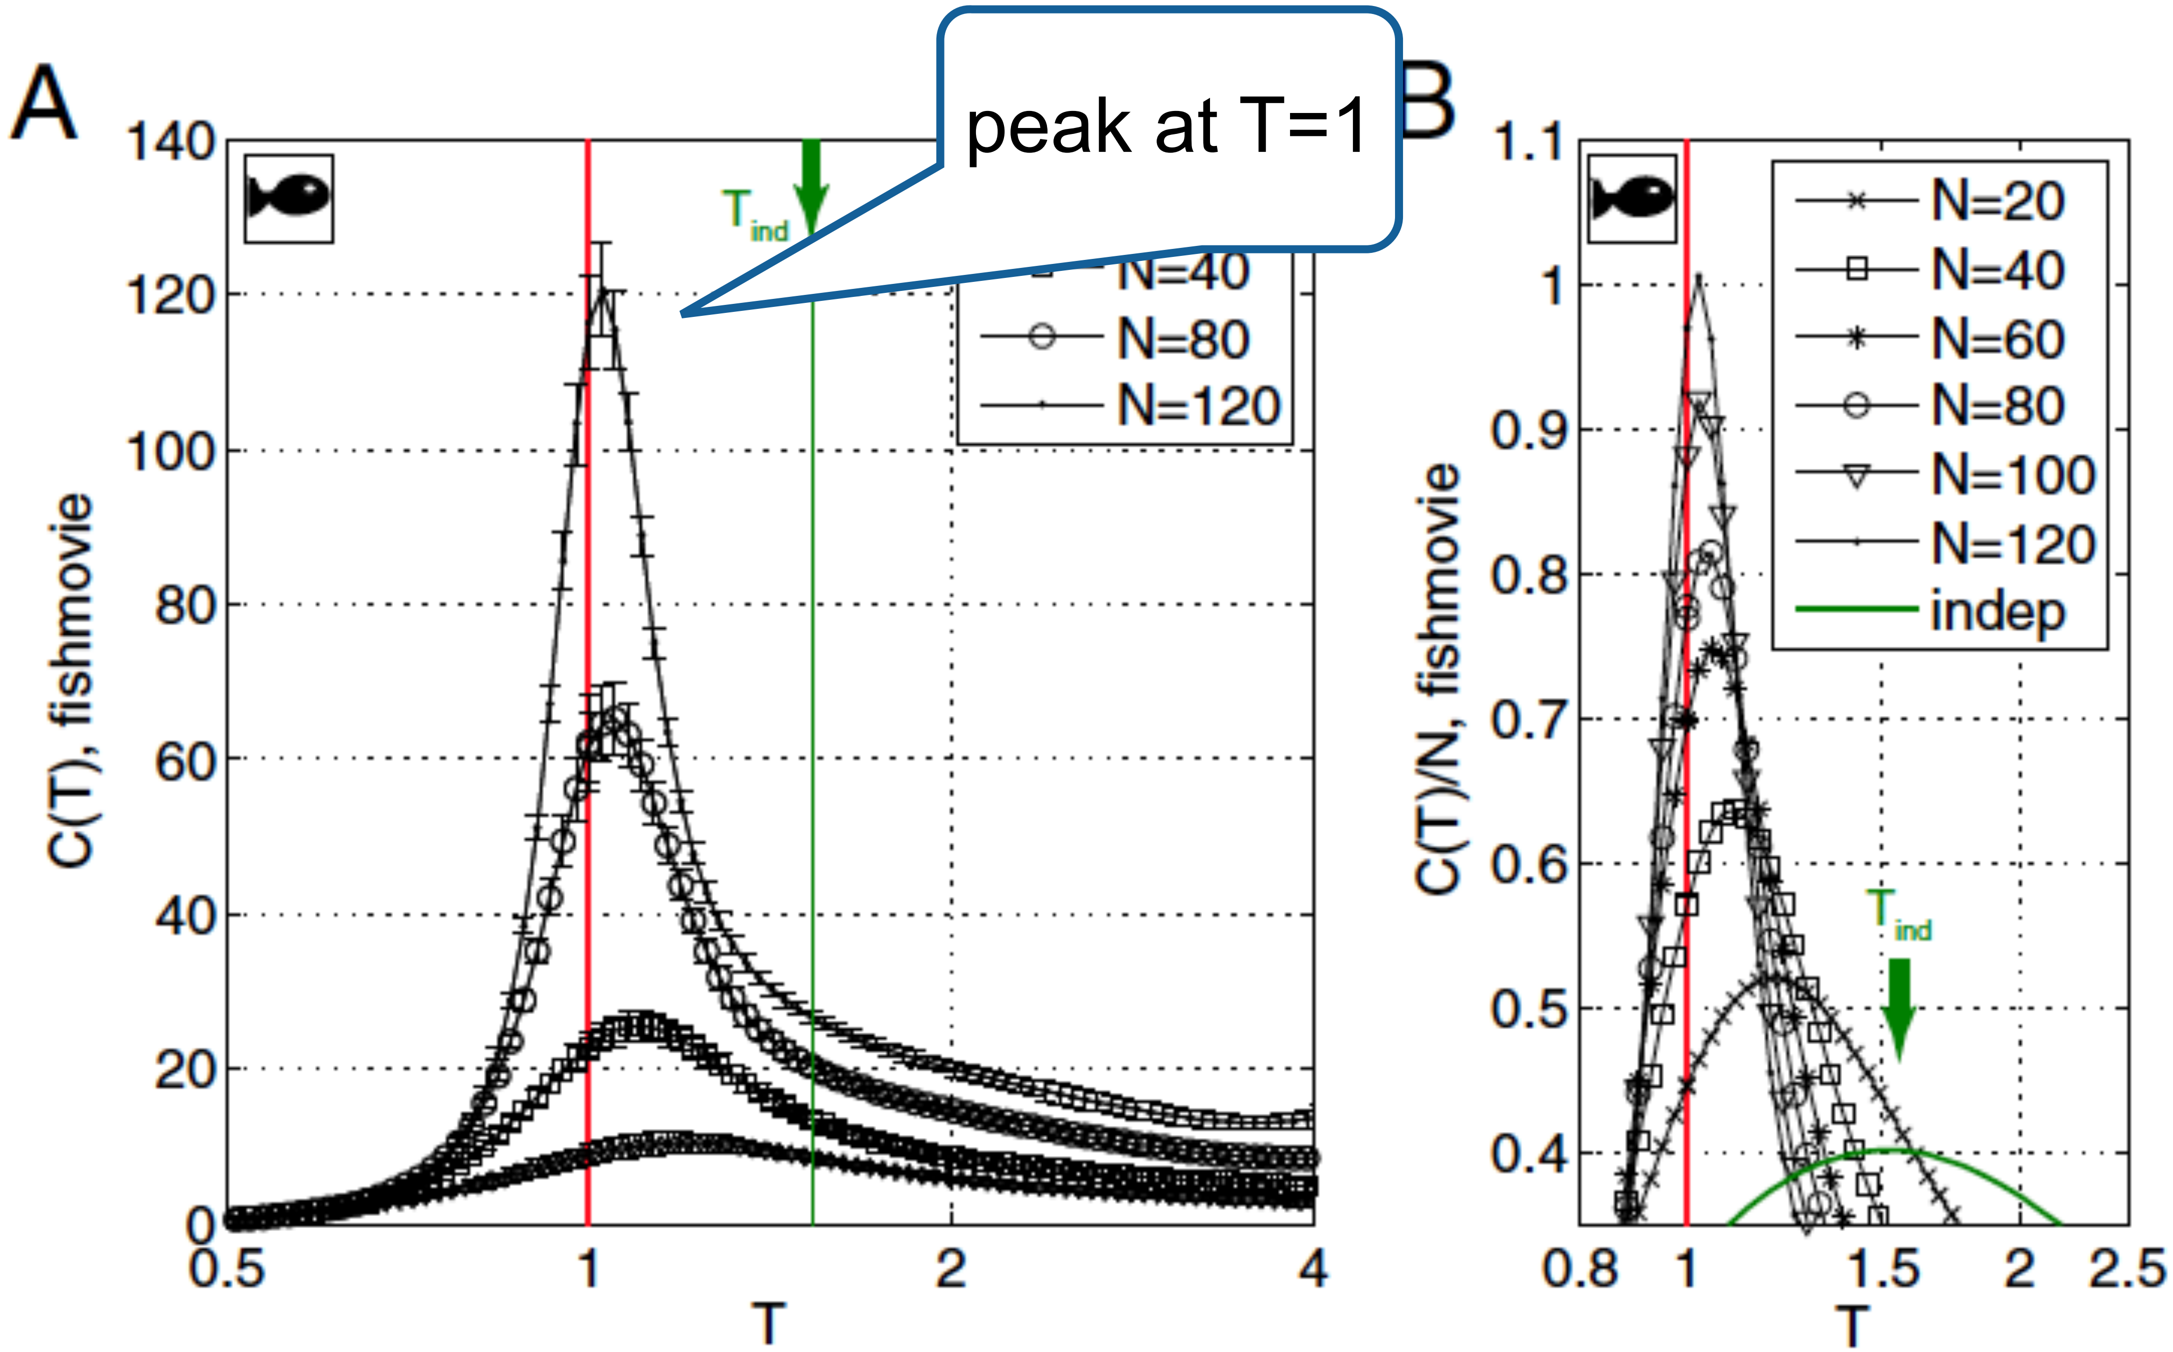

Tkacik, Mora, Marre, Amodei, Palmer, Berry, Bialek, PNAS 2015
Thermodynamics and signatures of criticality in a network of neurons

Tkacik et al 2009, Mora & Bialek 2011, Stephens et al 2012, Yu et al 2013, Mora et al 2015

Signatures of criticality in a recording of retinal ganglion cells: Are population codes 'critical'?

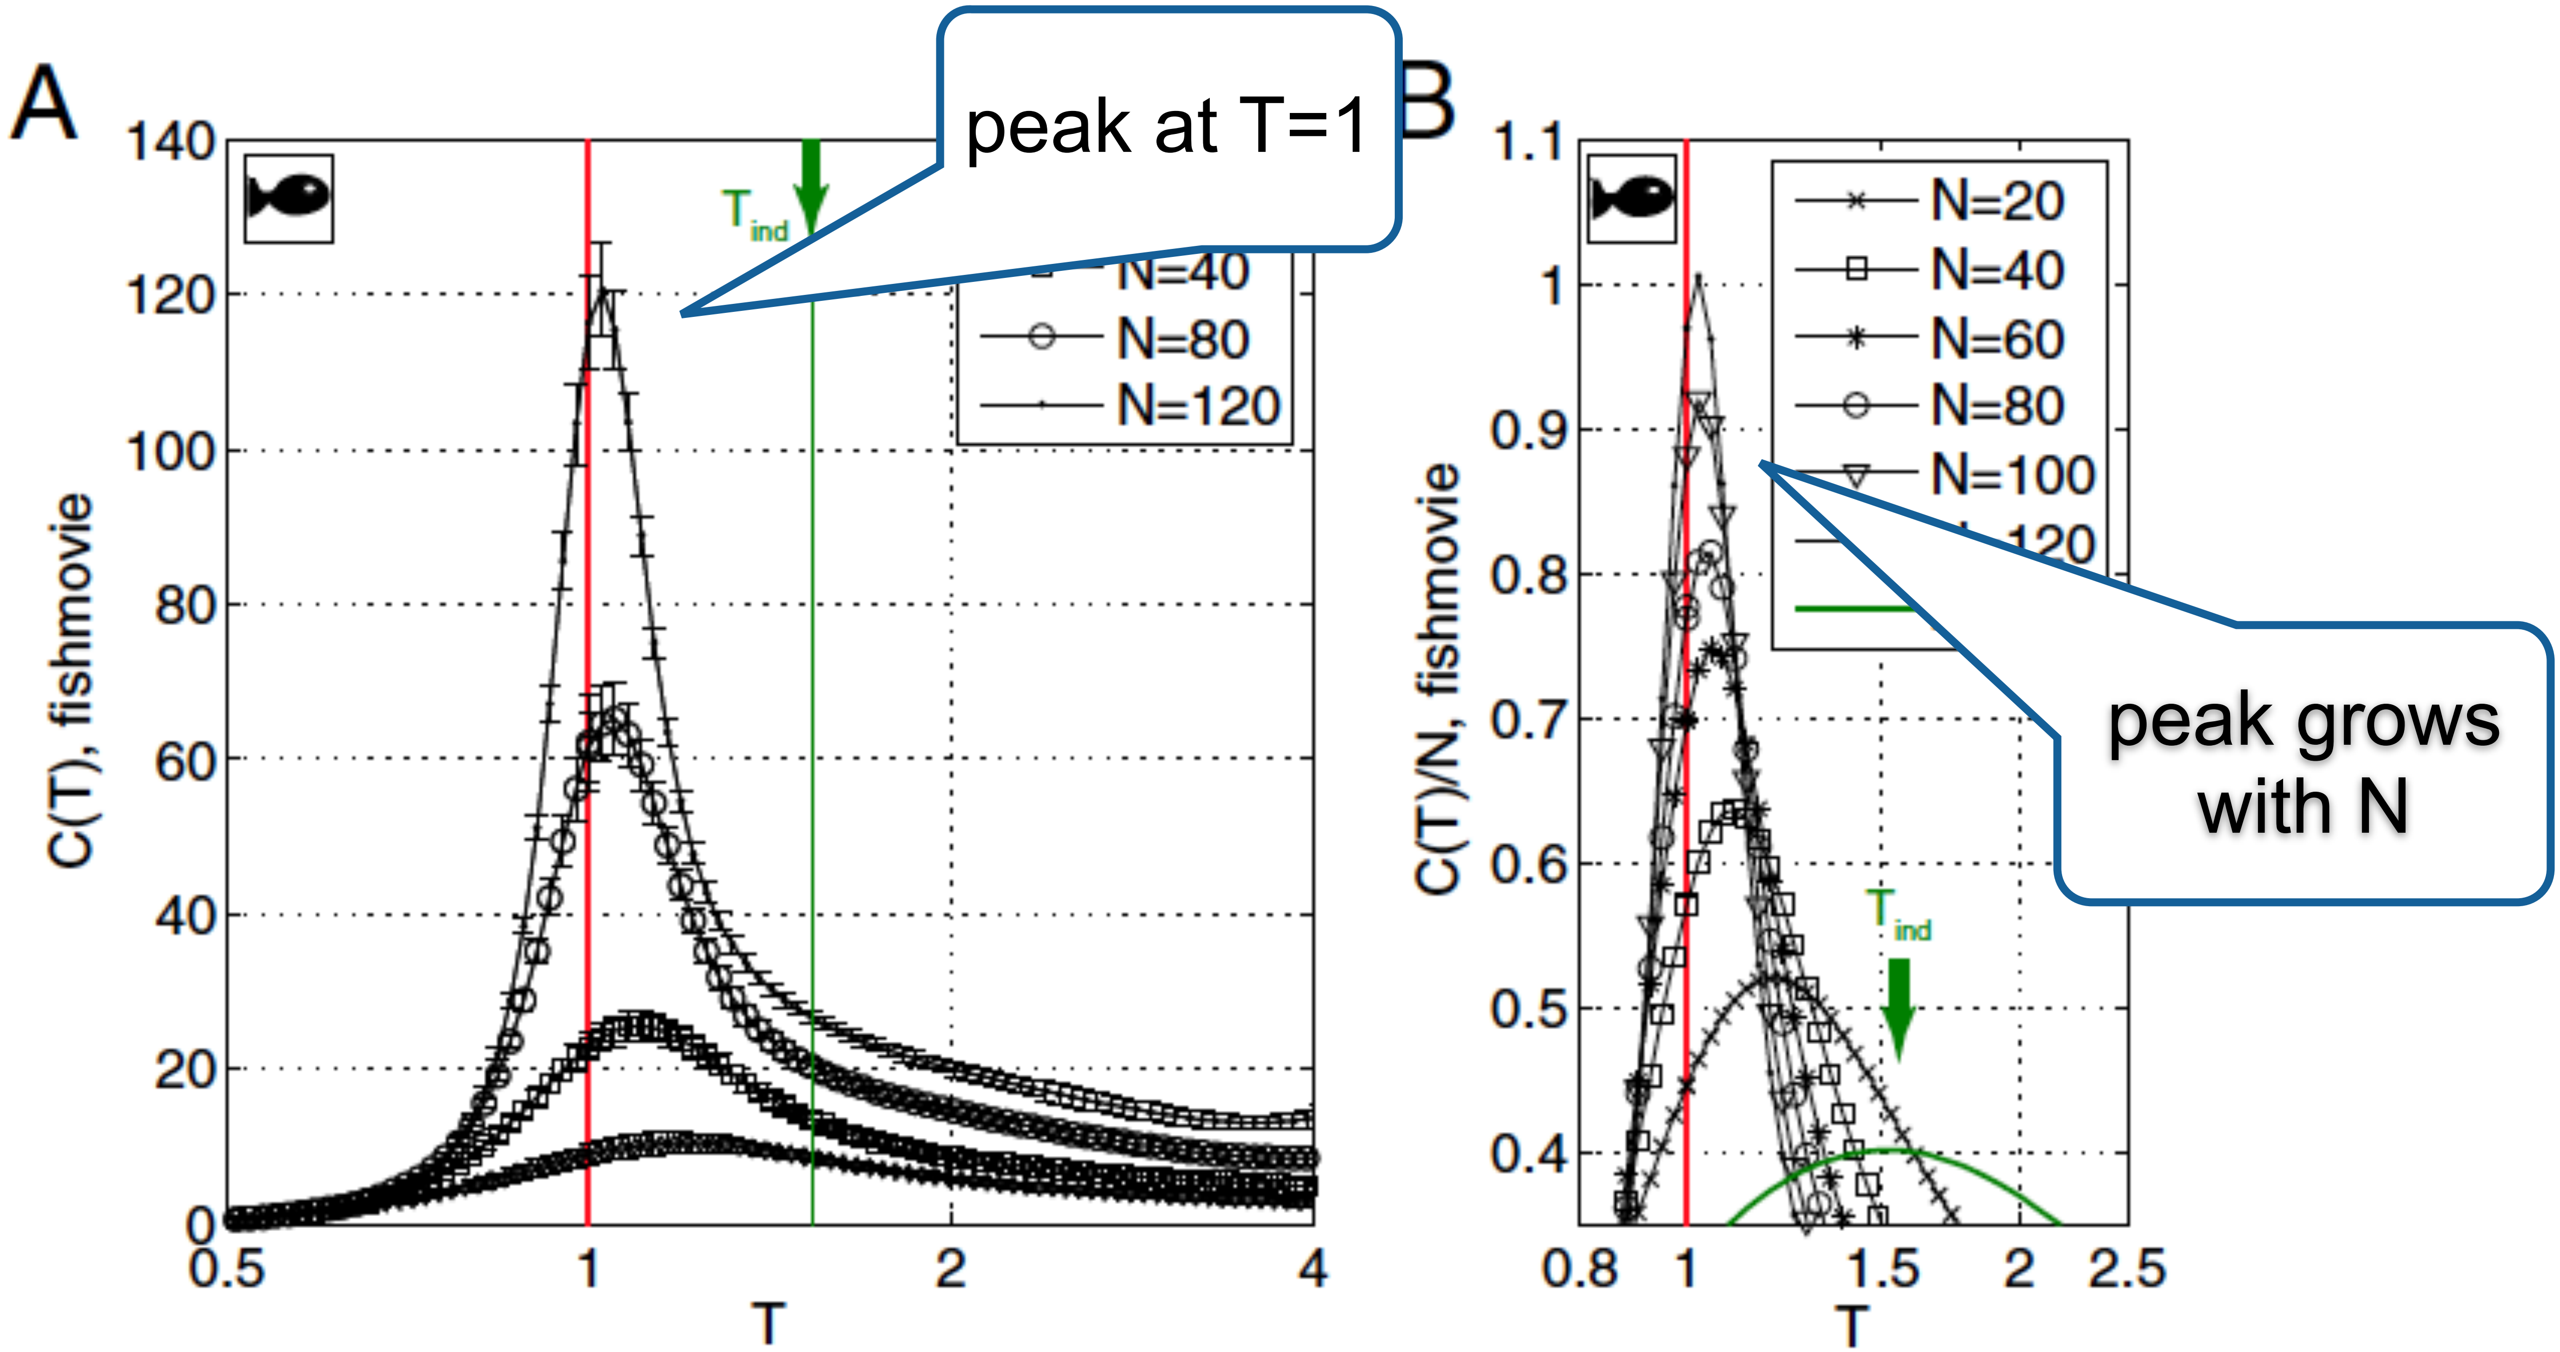

Tkacik, Mora, Marre, Amodei, Palmer, Berry, Bialek, PNAS 2015
Thermodynamics and signatures of criticality in a network of neurons

Tkacik et al 2009, Mora & Bialek 2011, Stephens et al 2012, Yu et al 2013, Mora et al 2015

- What neural mechanisms can explain this observation?

- What neural mechanisms can explain this observation?
- Does it require finely tuned parameters (e.g adaptation)?

- What neural mechanisms can explain this observation?
- Does it require finely tuned parameters (e.g adaptation)?
- What is a minimal model reproducing this behaviour?

- What neural mechanisms can explain this observation?
- Does it require finely tuned parameters (e.g adaptation)?
- What is a minimal model reproducing this behaviour?

Macke, Opper, Bethge 2011
Schwab & Metha et al 2014
Aitchison & Latham 2014

The simplest possible model of a patch of retina

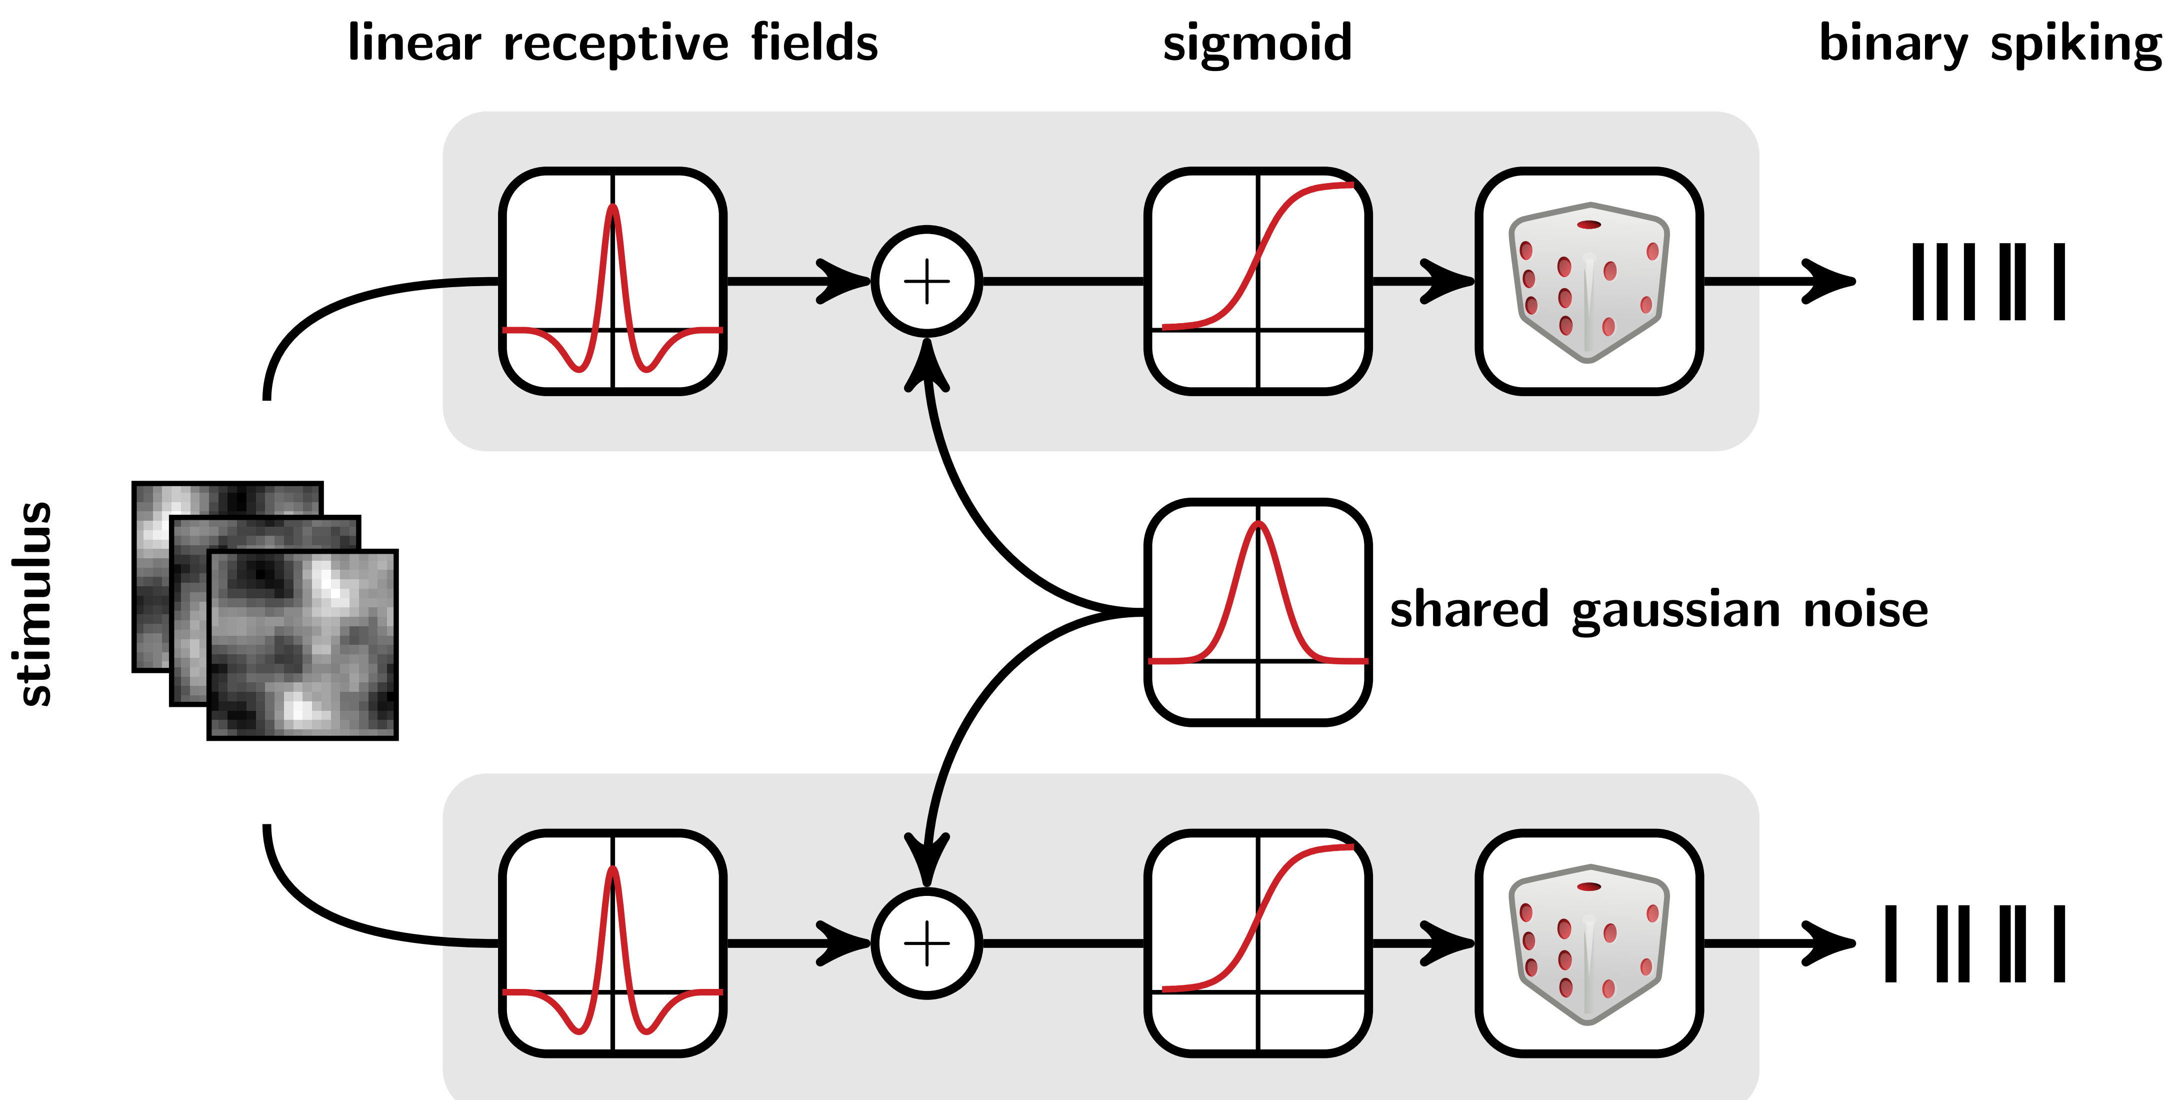

316 neurons

collaboration with C Behrens, P Berens, M Bethge

The simplest possible model of a patch of retina

Receptive field centres, N=316

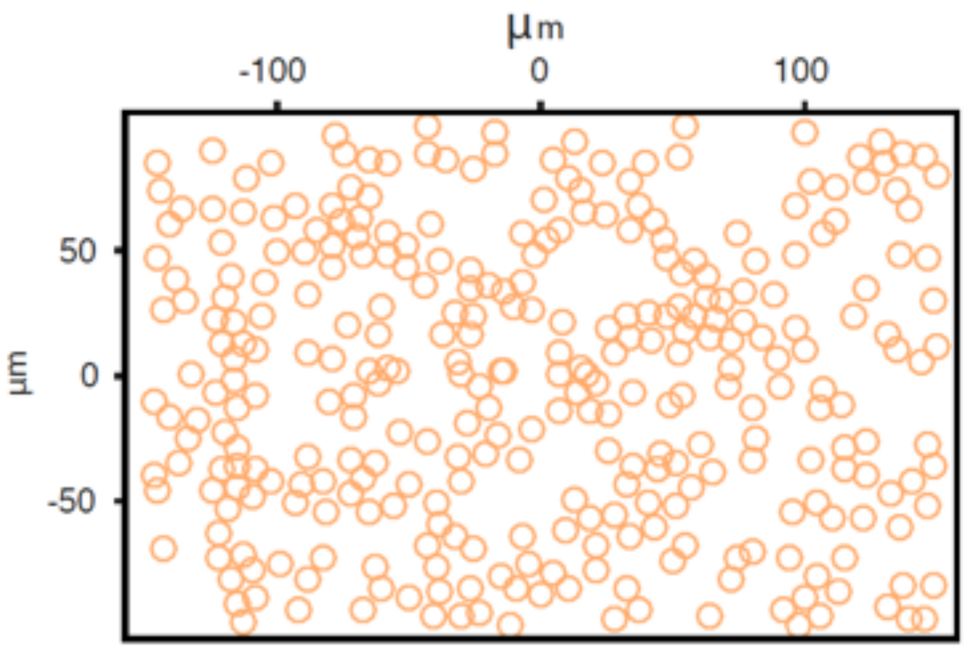

The simplest possible model of a patch of retina

Receptive field centres, N=316

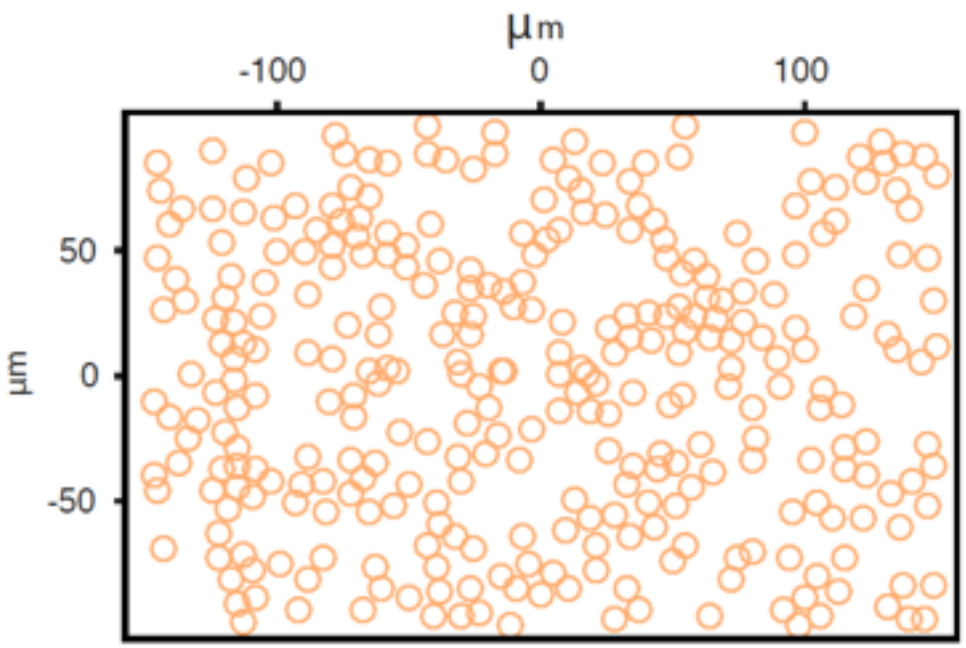

Simulated spike train

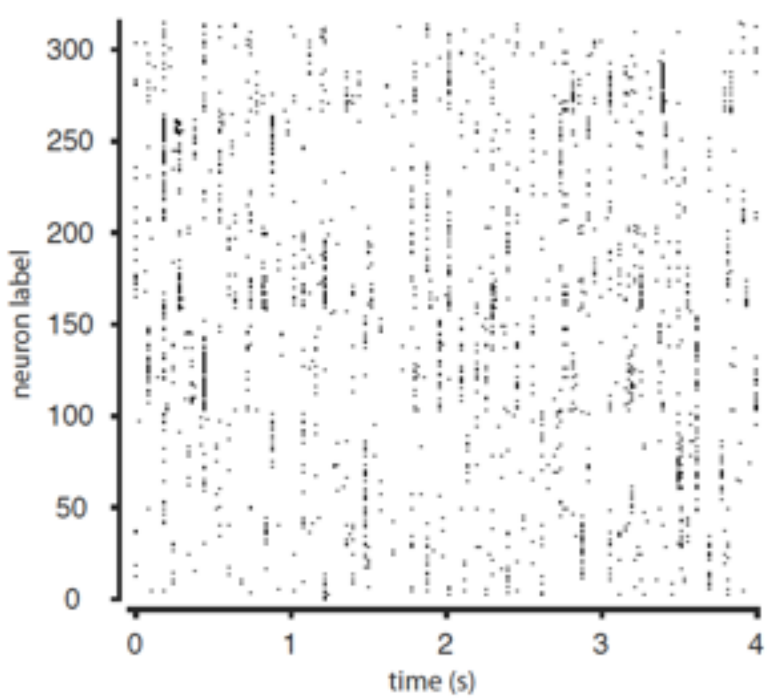

The simplest possible model of a patch of retina

Receptive field centres, N=316

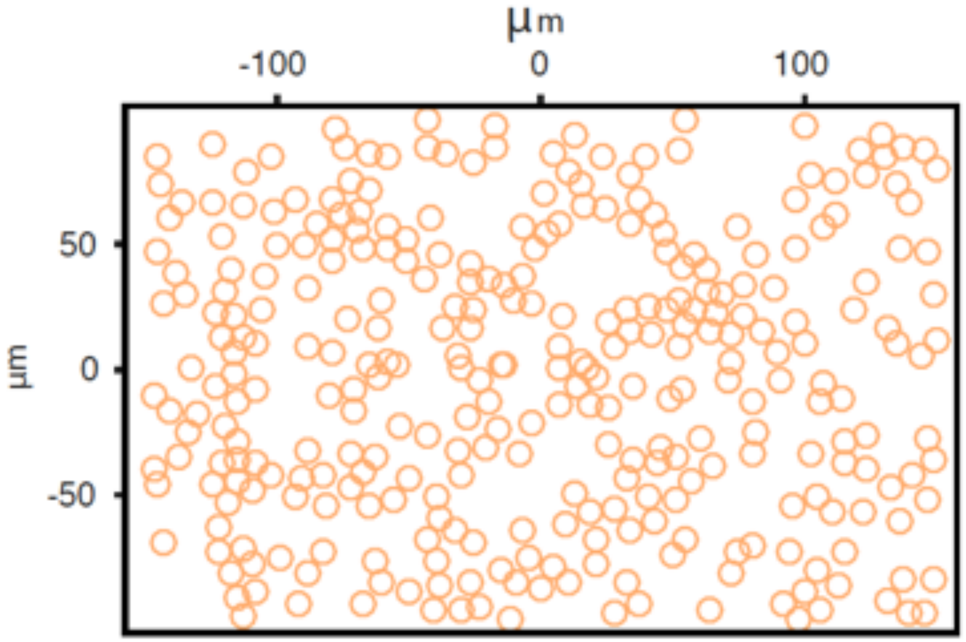

Histograms of firing rates and correlations

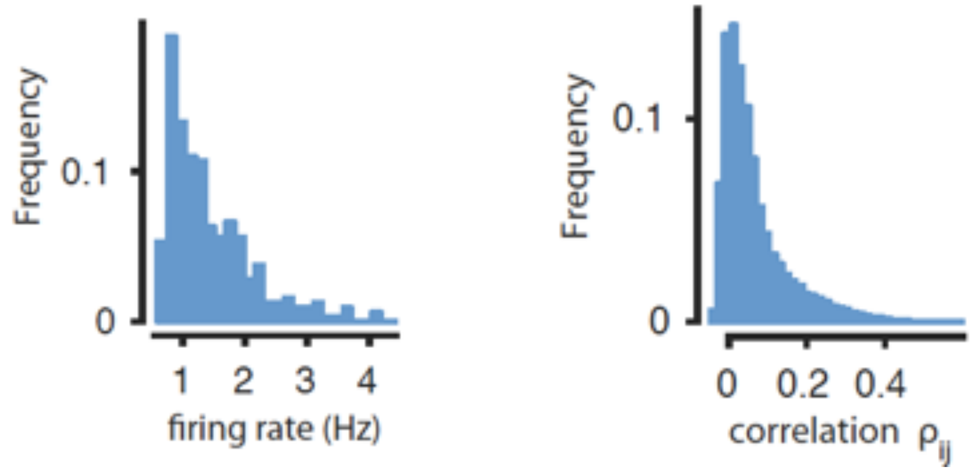

Simulated spike train

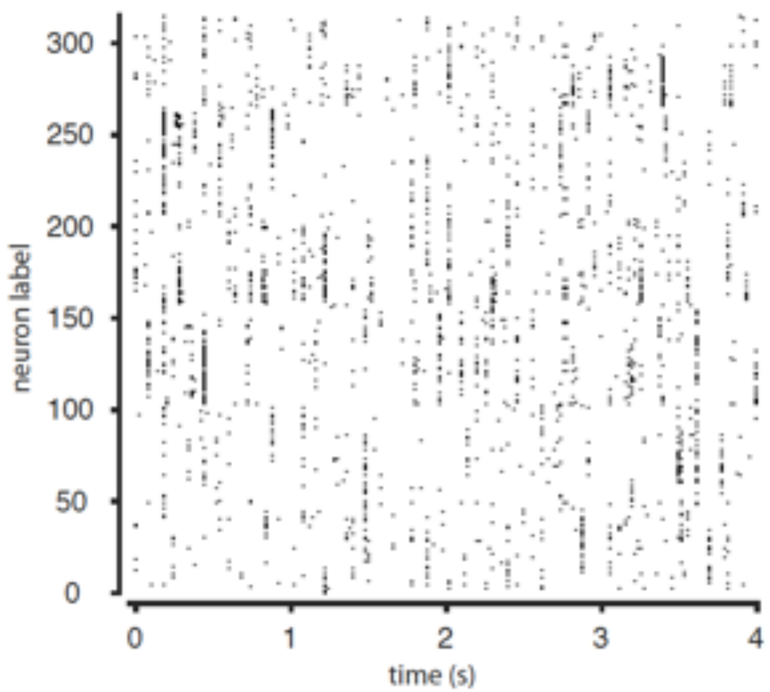

The simplest possible model of a patch of retina

Receptive field centres, N=316

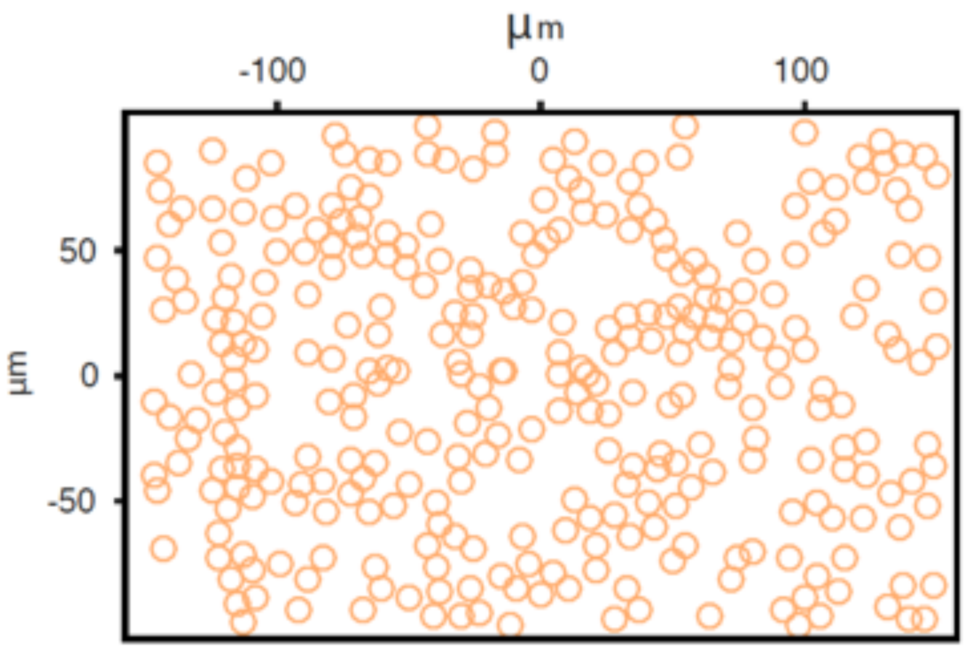

Histograms of firing rates and correlations

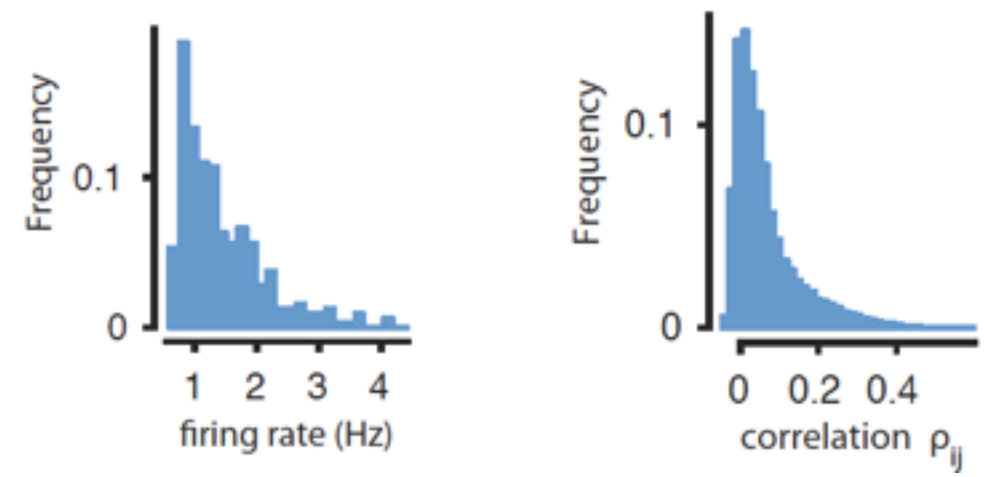

Simulated spike train

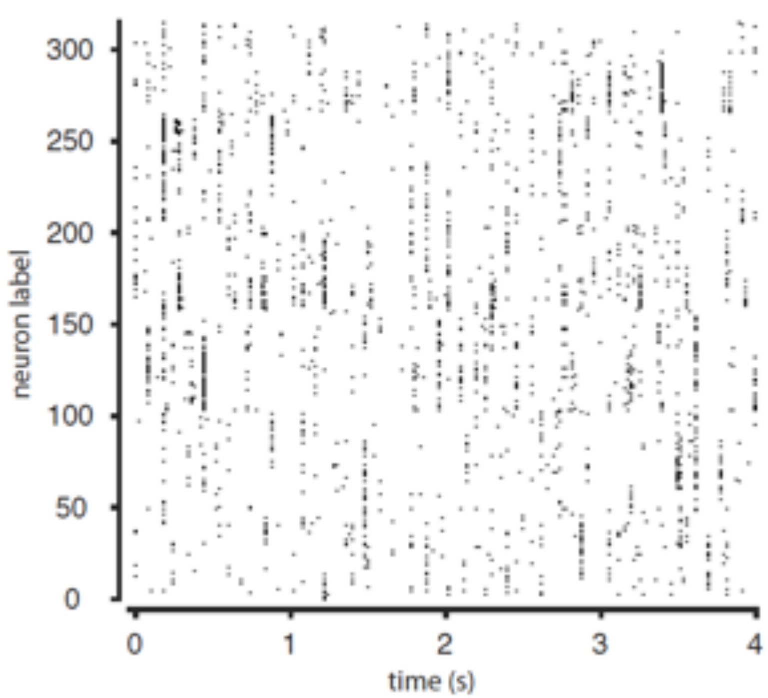

Population spike count distribution P(K)

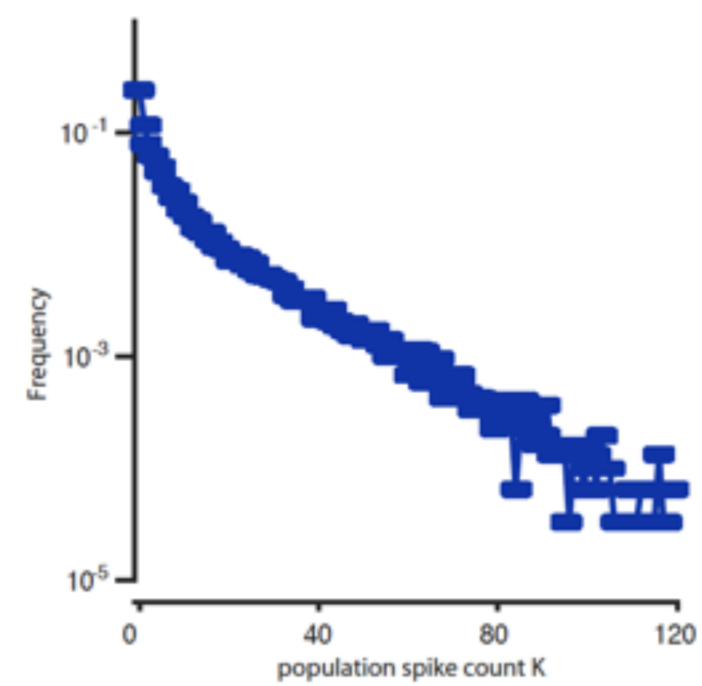

We infer model parameters of the K-pairwise model by maximising a penalised log-likelihood.

(Approximate) Log-likelihood:

$$L \approx \sum_{m=1}^{M} \lambda^\top F(x_m) - M \log \left(\sum_{x \in S} \exp\left(\lambda^\top F(x)\right) \right)$$

- Estimating correlations is most difficult, so do ‘pairwise Gibbs sampling’:

We infer model parameters of the K-pairwise model by maximising a penalised log-likelihood.

(Approximate) Log-likelihood:

$$L \approx \sum_{m=1}^{M} \lambda^{\top} F(x_m) - M \log \left(\sum_{x \in S} \exp \left(\lambda^{\top} F(x) \right) \right)$$

- Estimating correlations is most difficult, so do 'pairwise Gibbs sampling':

$$p(x_i x_j | x_1, x_2, ..., x_{i-1}, x_{i+1}, ...x_{j-1}, x_{j+1}, ...x_N, \lambda)$$

- This allows us to use Rao-Blackwell for the pairwise terms.

We infer model parameters of the K-pairwise model by maximising a penalised log-likelihood.

(Approximate) Log-likelihood: $$L \approx \sum_{m=1}^{M} \lambda^\top F(x_m) - M \log\left(\sum_{x \in S} \exp\left(\lambda^\top F(x)\right)\right)$$

- Estimating correlations is most difficult, so do 'pairwise Gibbs sampling':

$$p(x_i x_j | x_1, x_2, ..., x_{i-1}, x_{i+1}, ... x_{j-1}, x_{j+1}, ... x_N, \lambda)$$

- This allows us to use Rao-Blackwell for the pairwise terms.

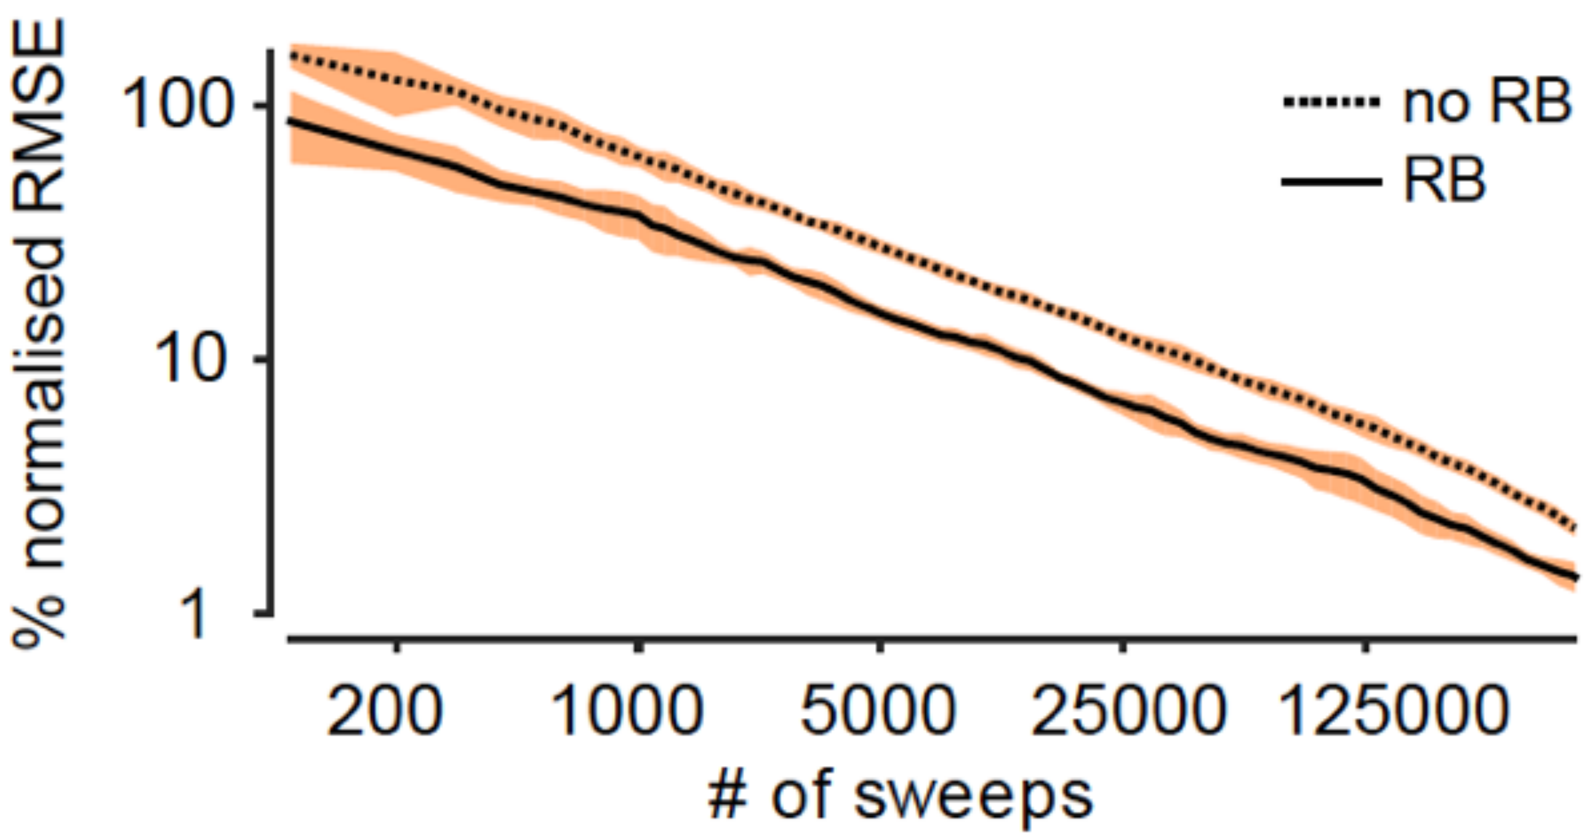

If you use use smart parameter-updates, you do not need to store the entire MCMC sample.

$$L \approx \sum_{m=1}^{M} \lambda^\top F(x_m) - M \log \left(\sum_{x \in S} \exp\left(\lambda^\top F(x)\right) \right)$$

Change in log-likelihood:

$$\delta = \hat{\lambda} - \lambda$$

$$\Delta L(\delta) = L(\hat{\lambda}) - L(\lambda)$$

$$\Delta L(\delta) \approx M \delta^\top \langle F(x) \rangle_{data} - M \log \left(\sum_{x \in S} \exp\left(\delta^\top F(x)\right) \right)$$

$$\text{For } x \in \{0, 1\} : \exp(\alpha x) = x \exp(\alpha) + 1 - x$$

$$\text{Assume that only } \delta_i F_i(x) \neq 0 :$$

$$\exp\left(\delta^\top F(x)\right) = F_i(x) \exp(\delta) + 1 - F_i(x)$$

As long as we only update one h or J term (or all of the spike count terms) at each iteration, we only need to store the feature-means, not the entire MCMC sample.

Signatures of criticality arise in a simple simulation of the retina, without any tuning or adaptation.

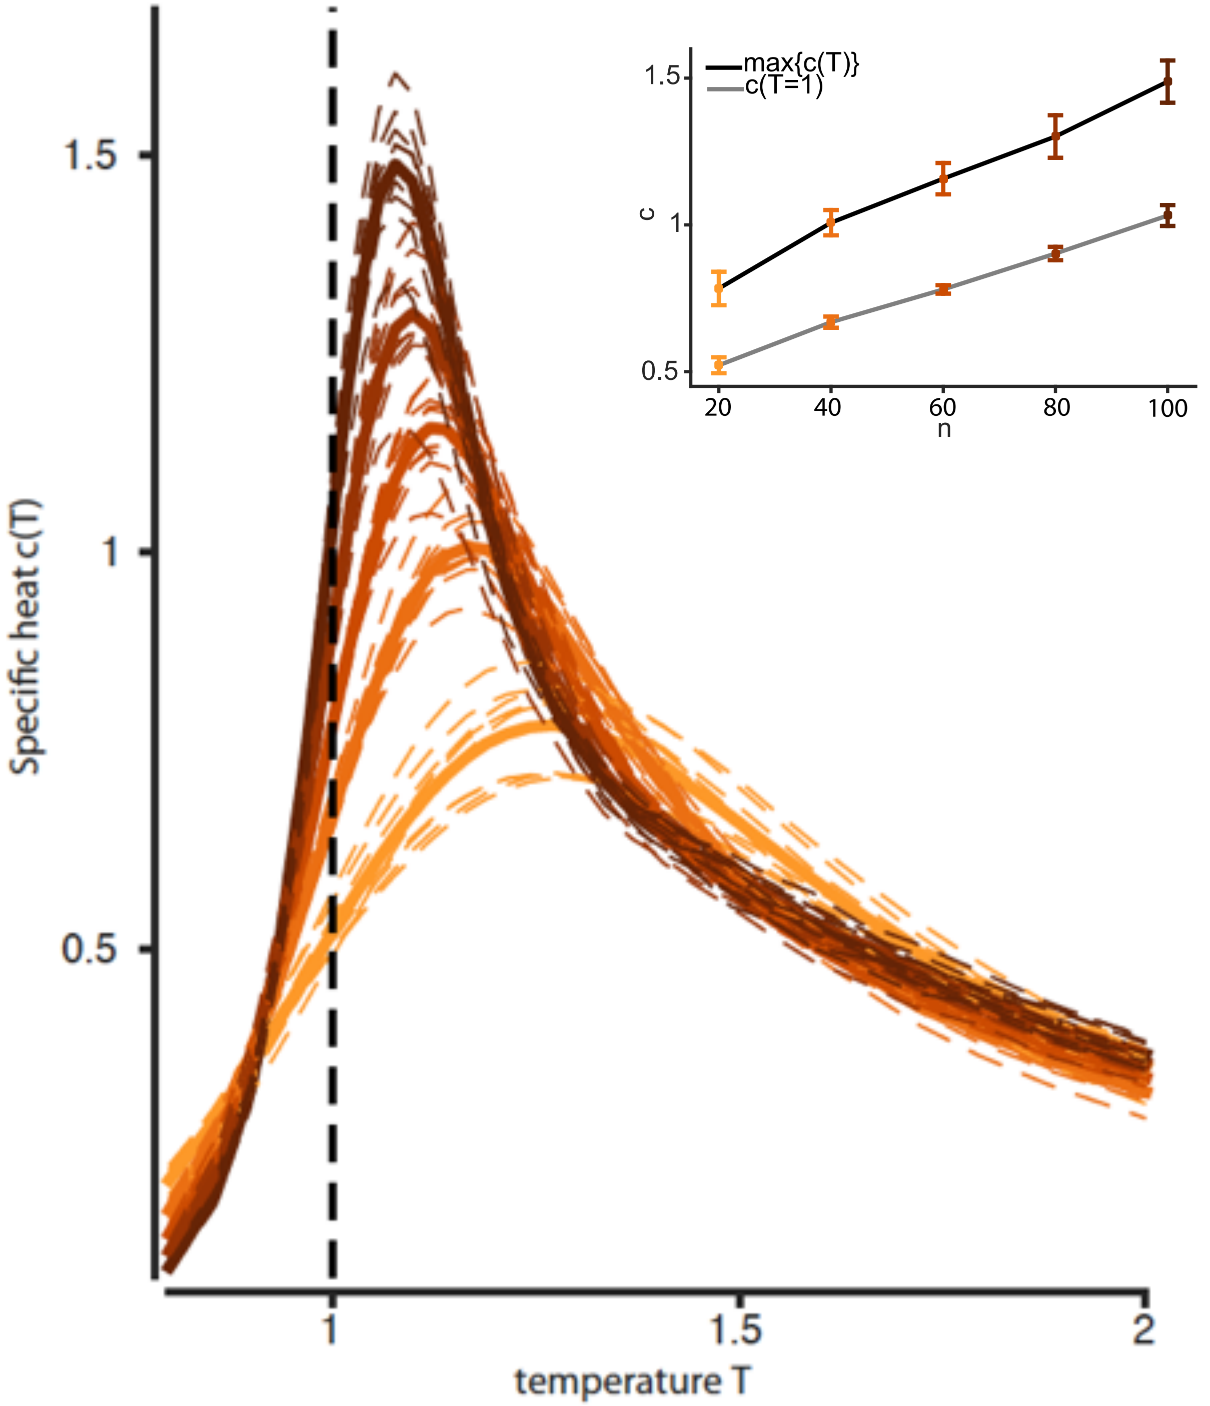

Signatures of criticality arise in a simple simulation of the retina, without any tuning or adaptation.

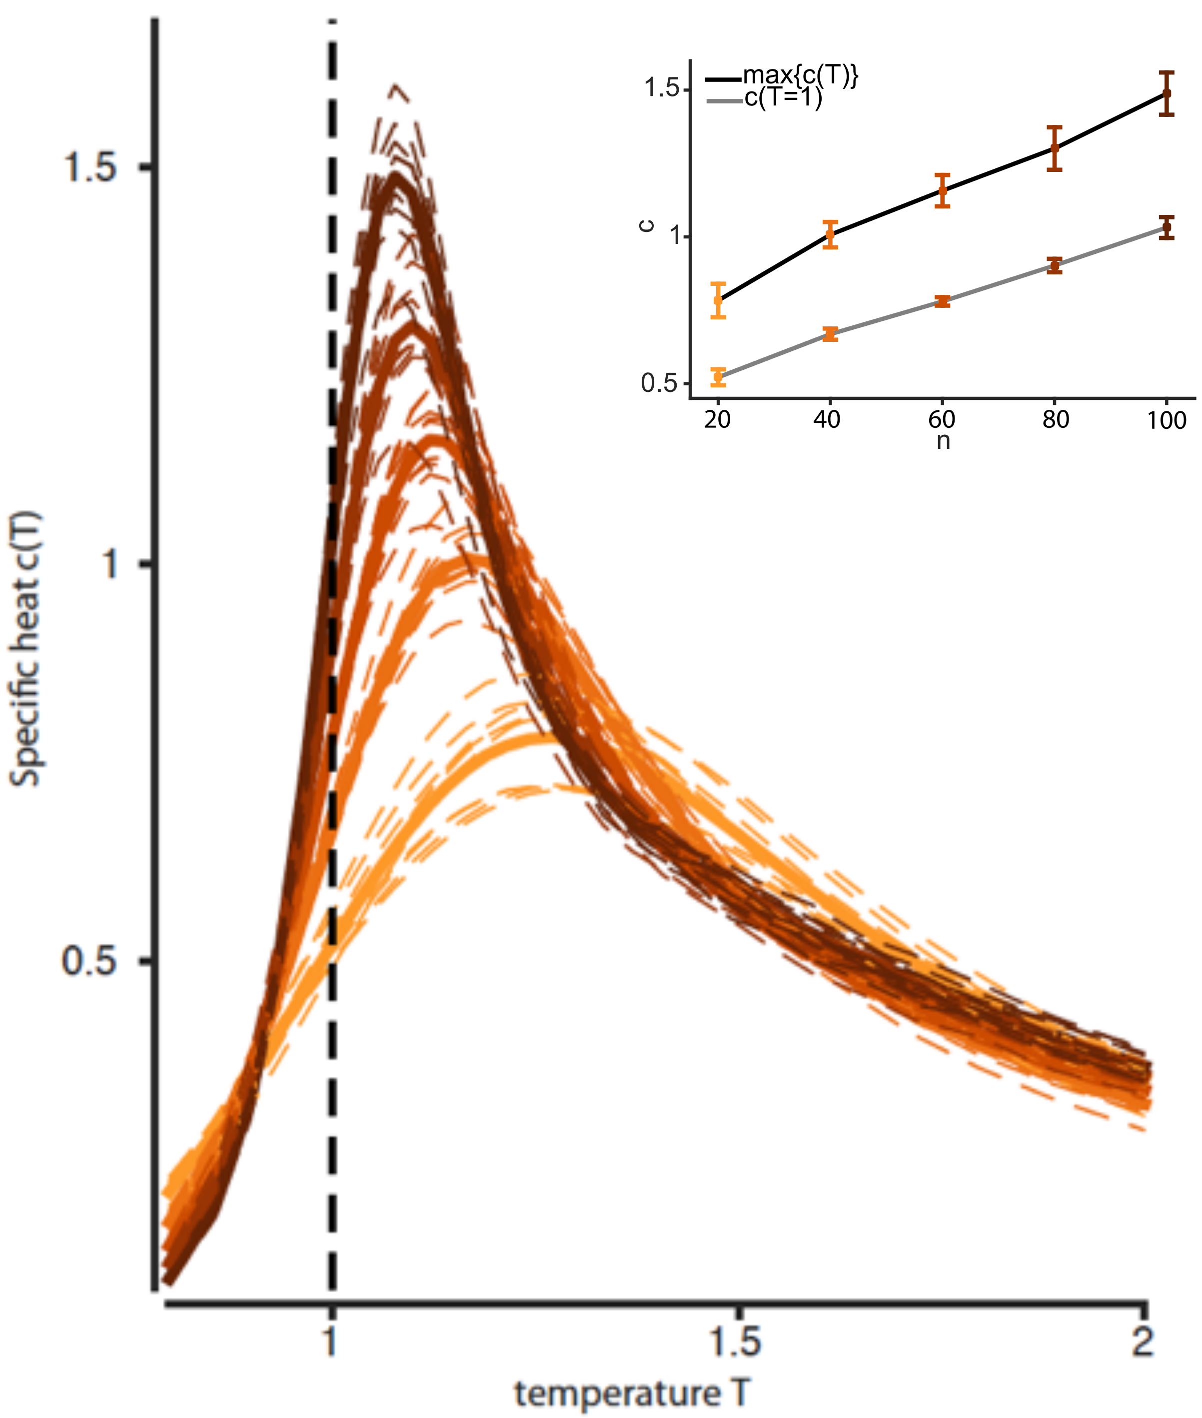

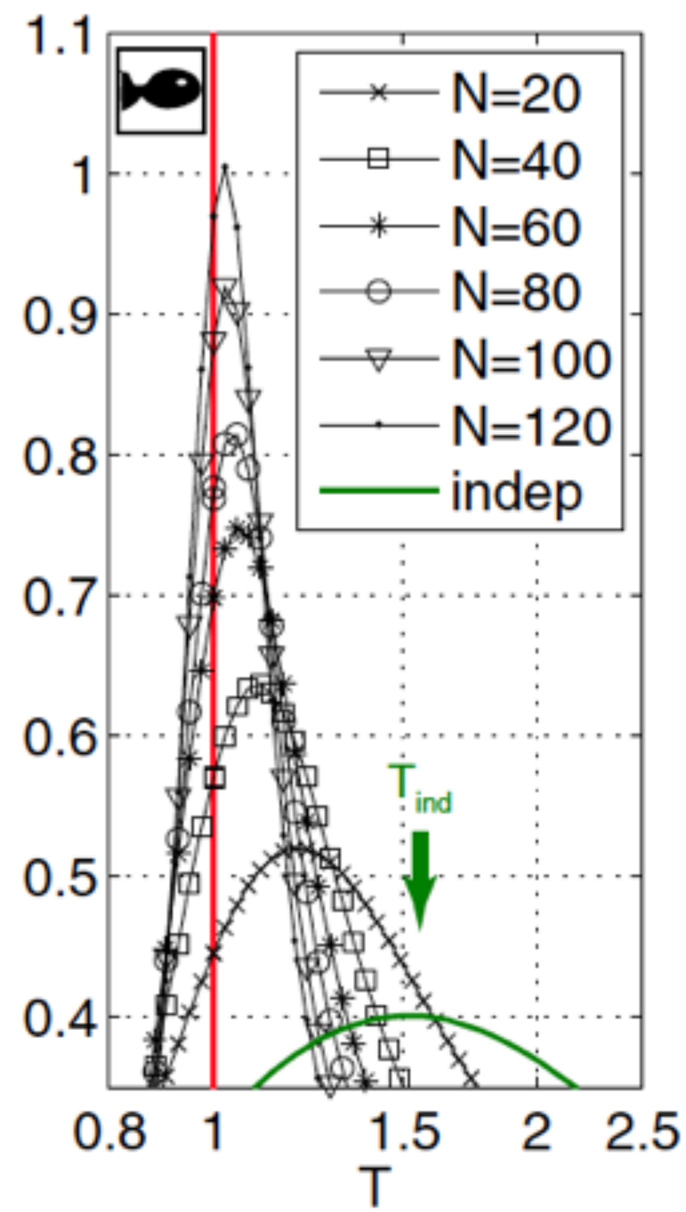

White noise | Natural movies | Full field flicker

Tkacik et al 2015

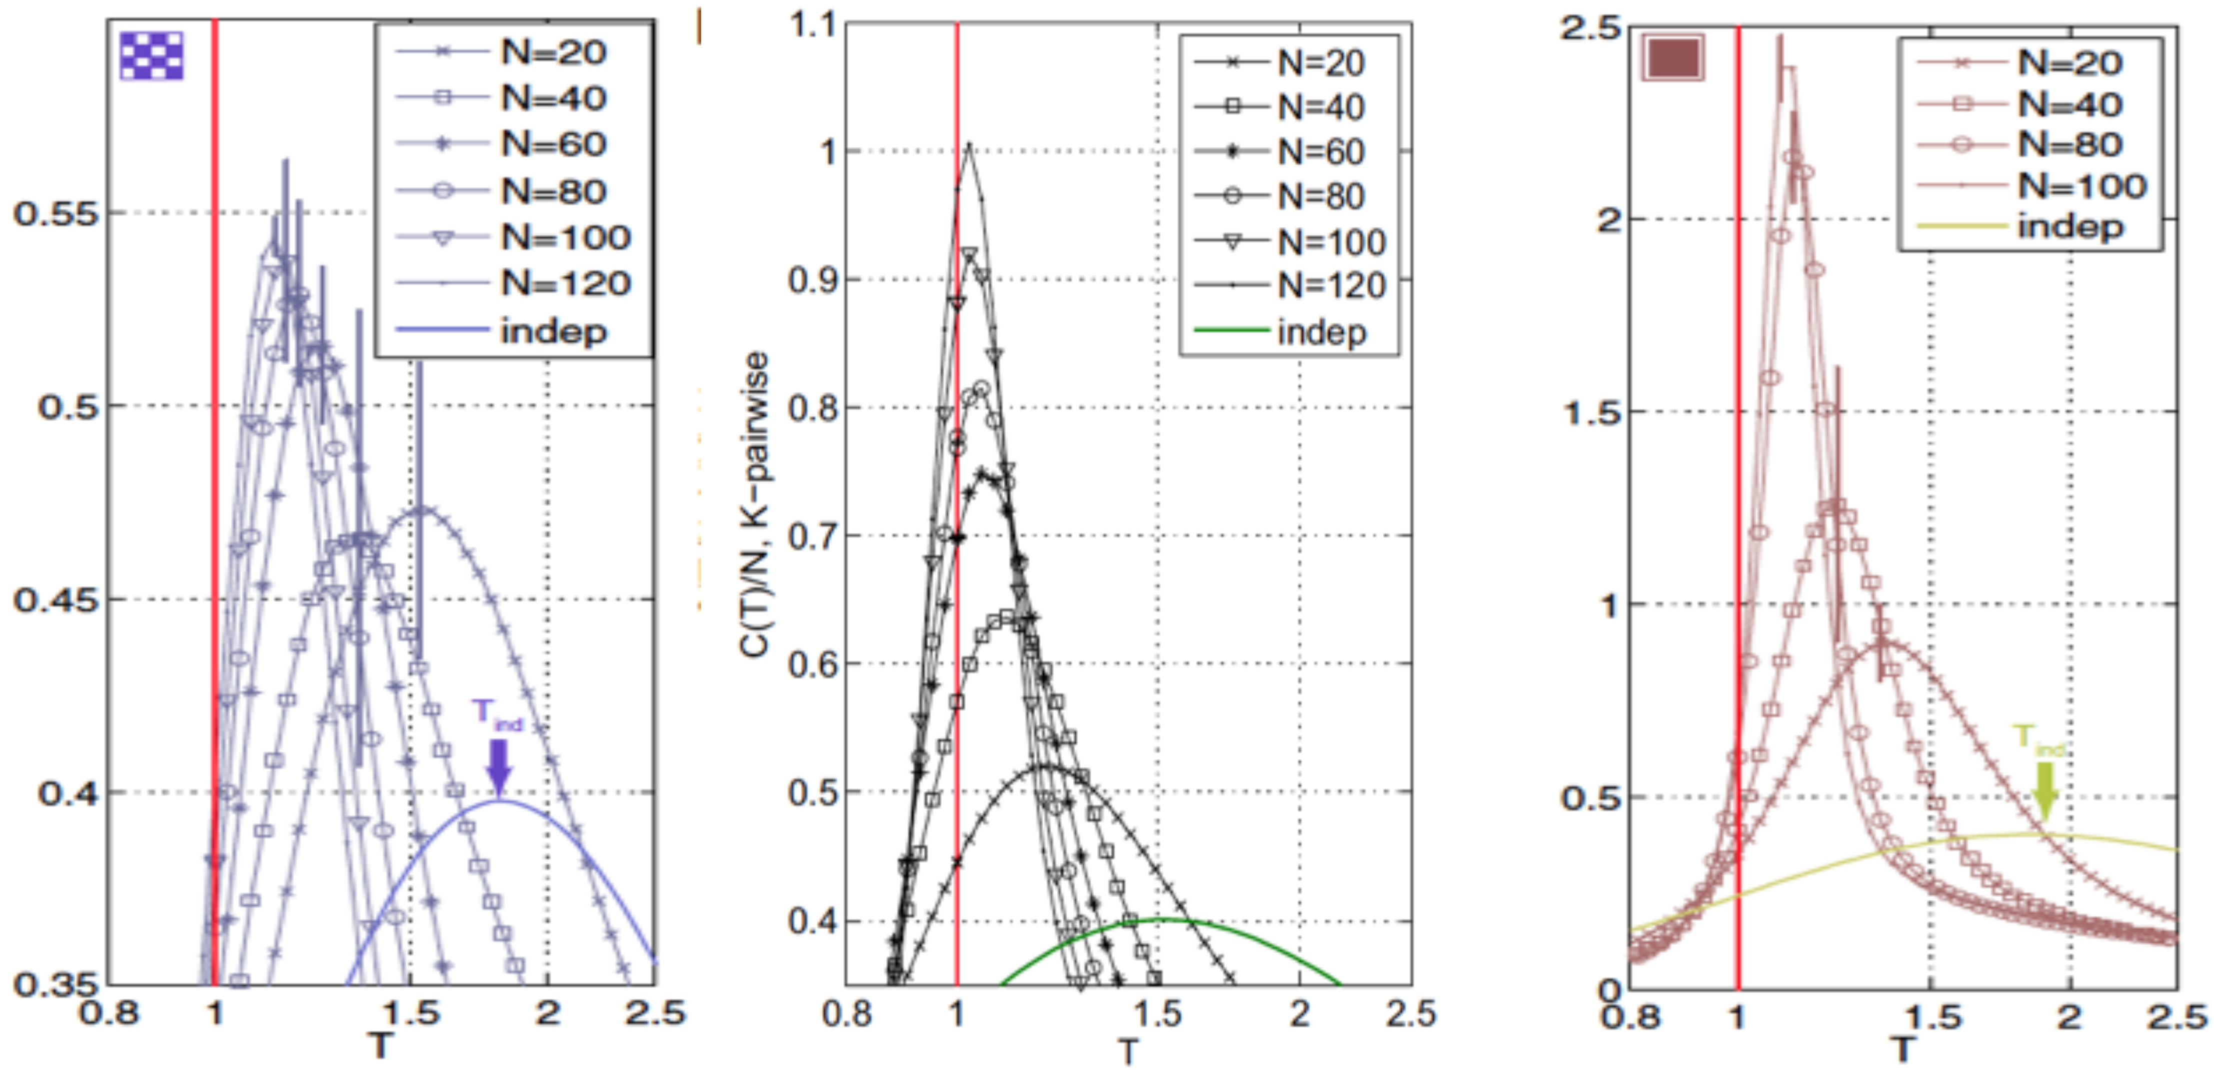

Our model

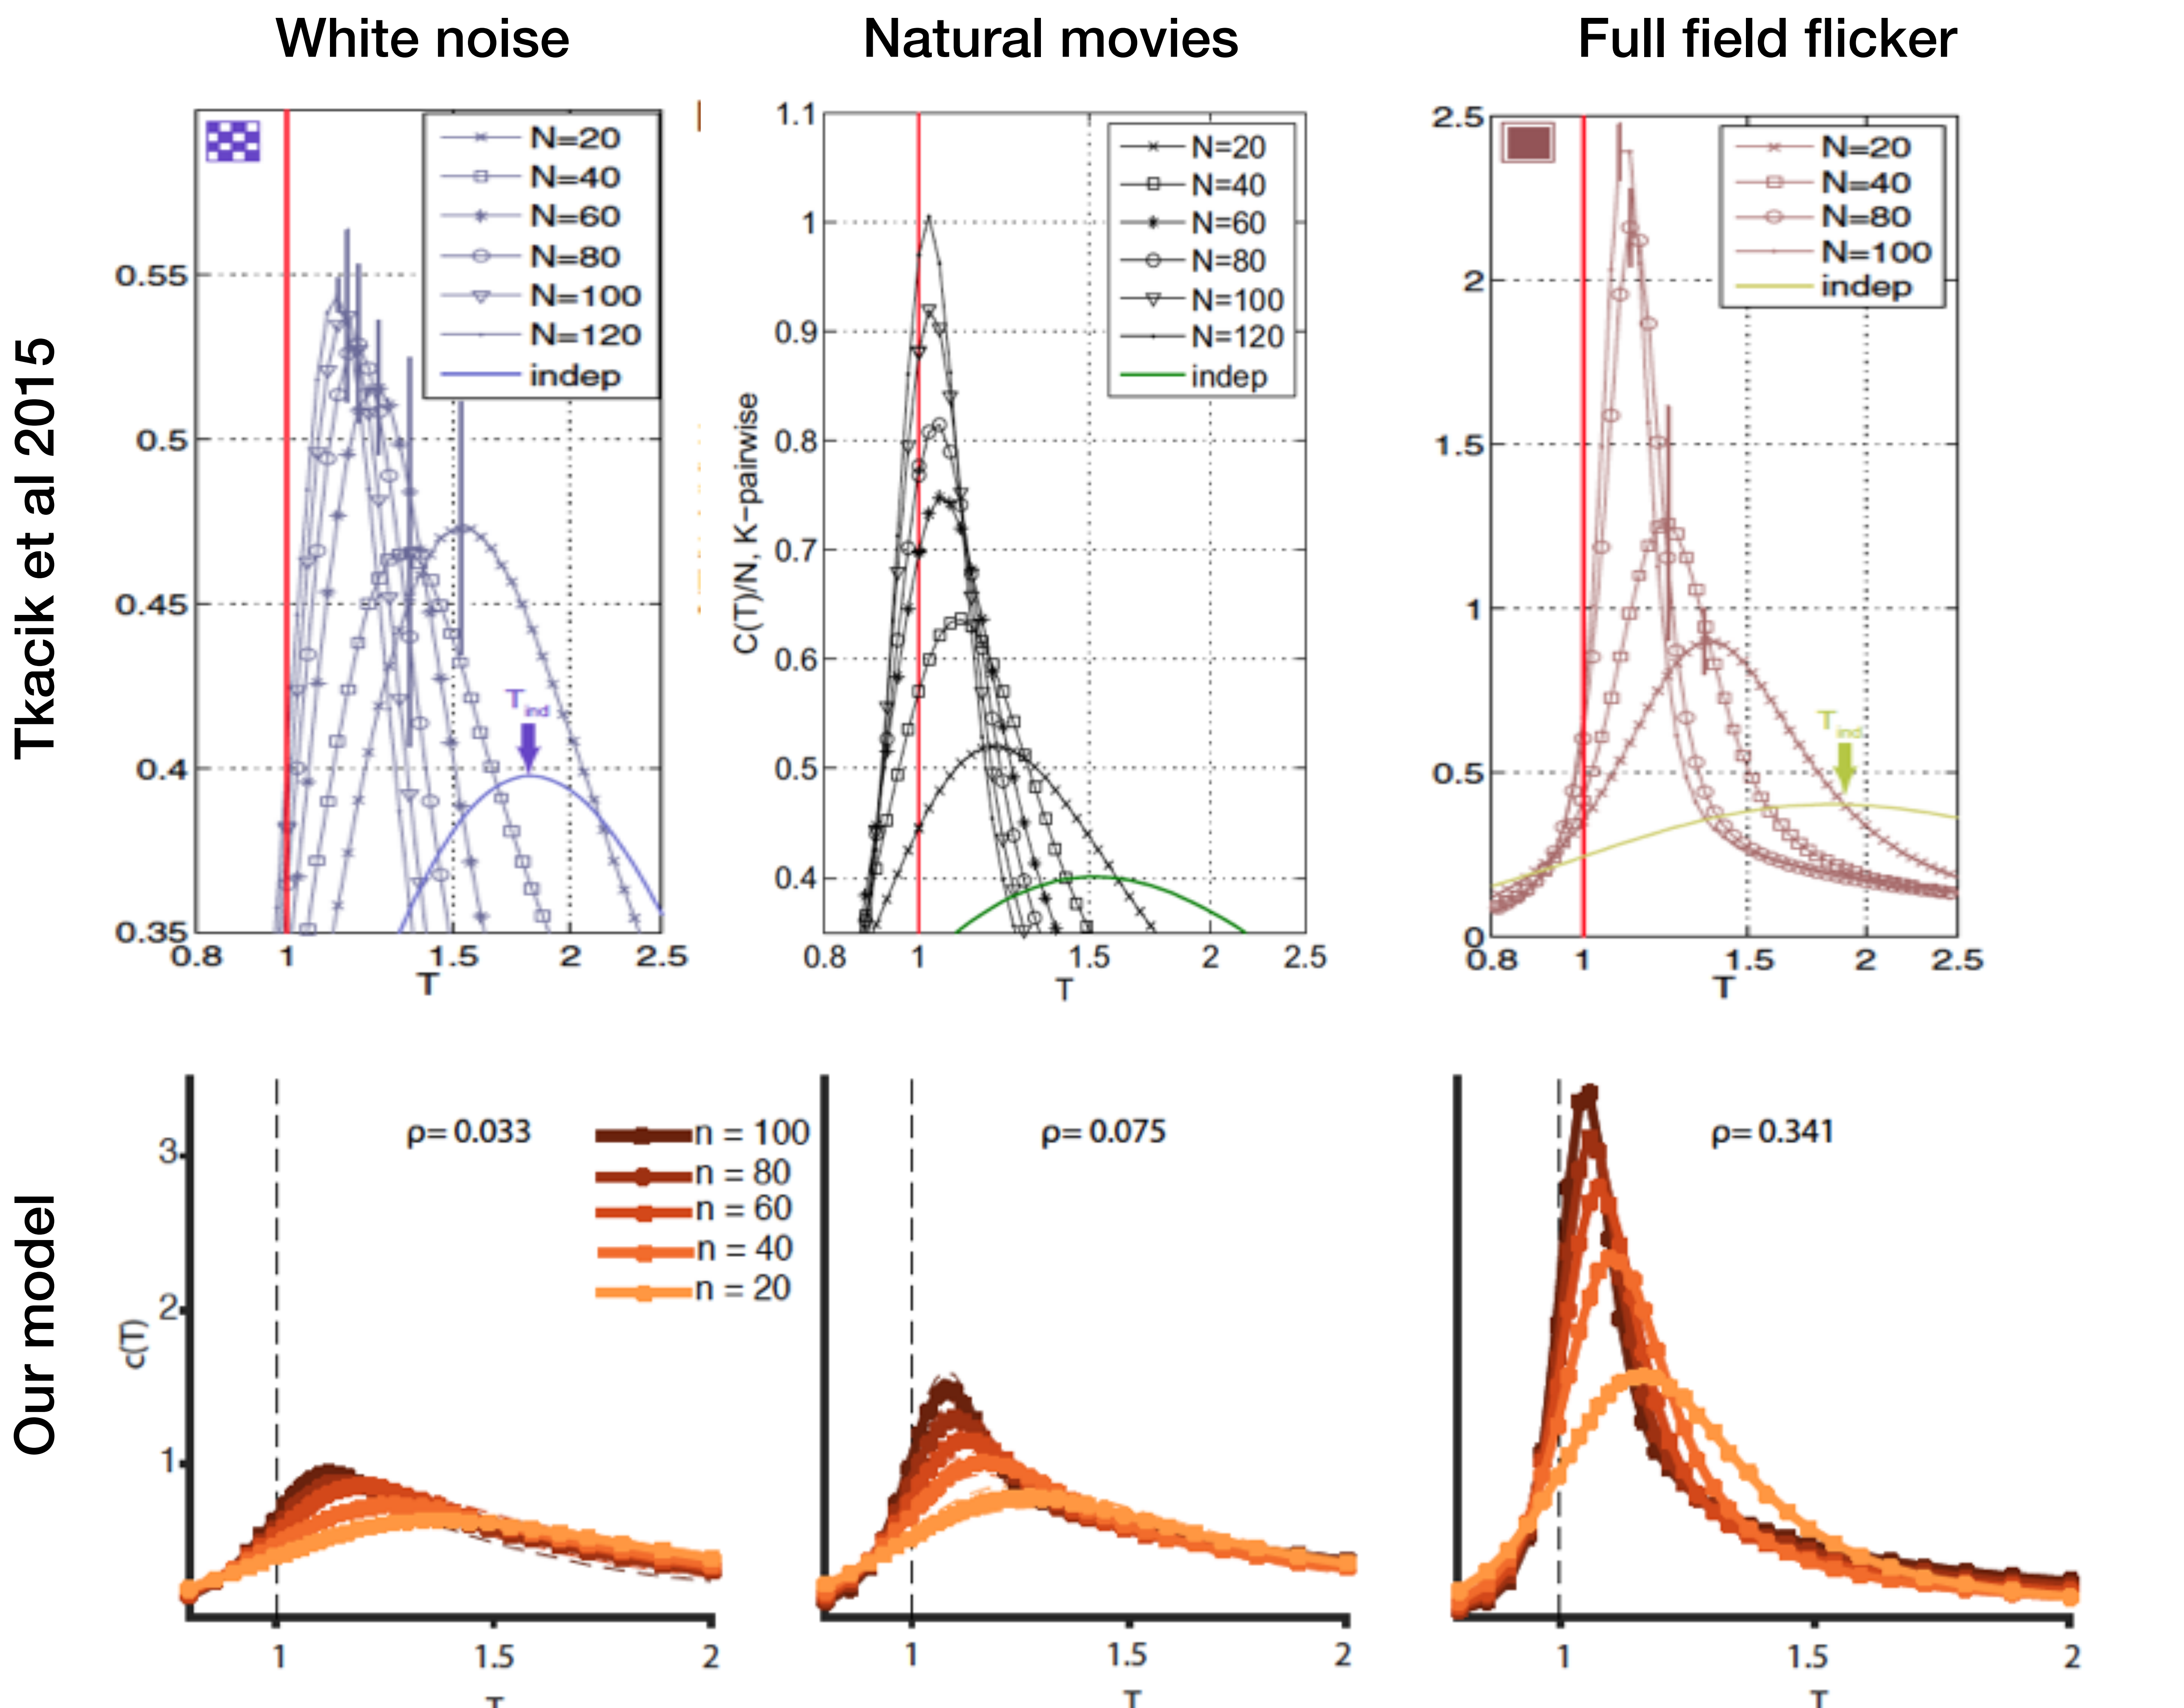
White noise
Natural movies
Full field flicker
Tkacik et al 2015
Our model
N=20
N=40
N=60
N=80
N=100
N=120
indep
T_ind
C(T)/N, K-pairwise
T
ρ= 0.033
ρ= 0.075
ρ= 0.341
n = 100
n = 80
n = 60
n = 40
n = 20
c(T)

Analysis: Subsampling a homogeneous population

Construct populations of size N by subsampling a large homogeneous population

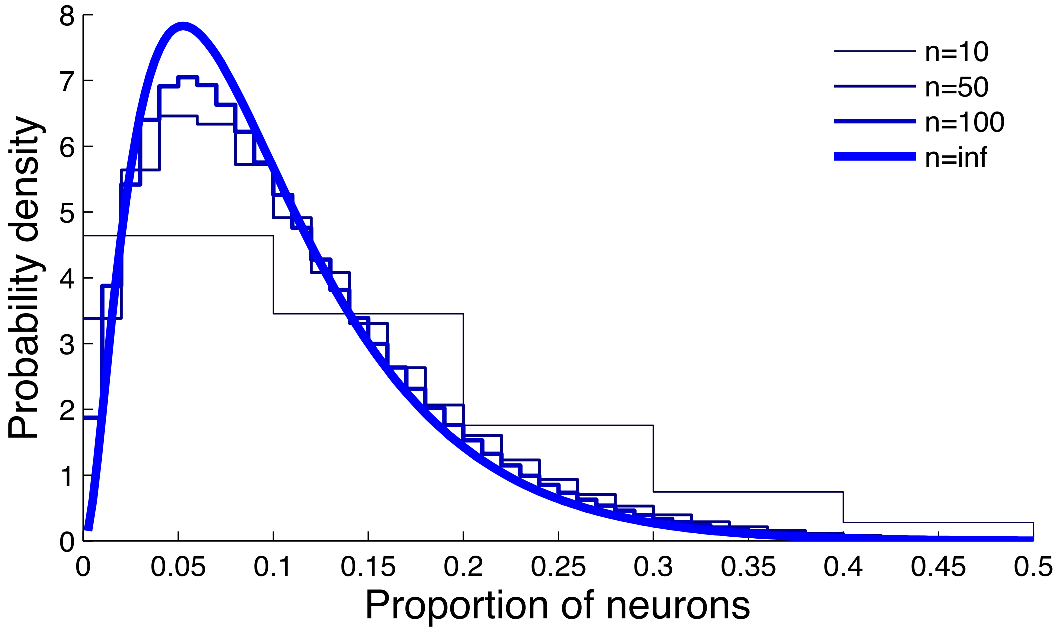

Amari et al 2003, Macke et al 2011, Tkacik et al 2013

Analysis: Subsampling a homogeneous population

Construct populations of size N by subsampling a large homogeneous population

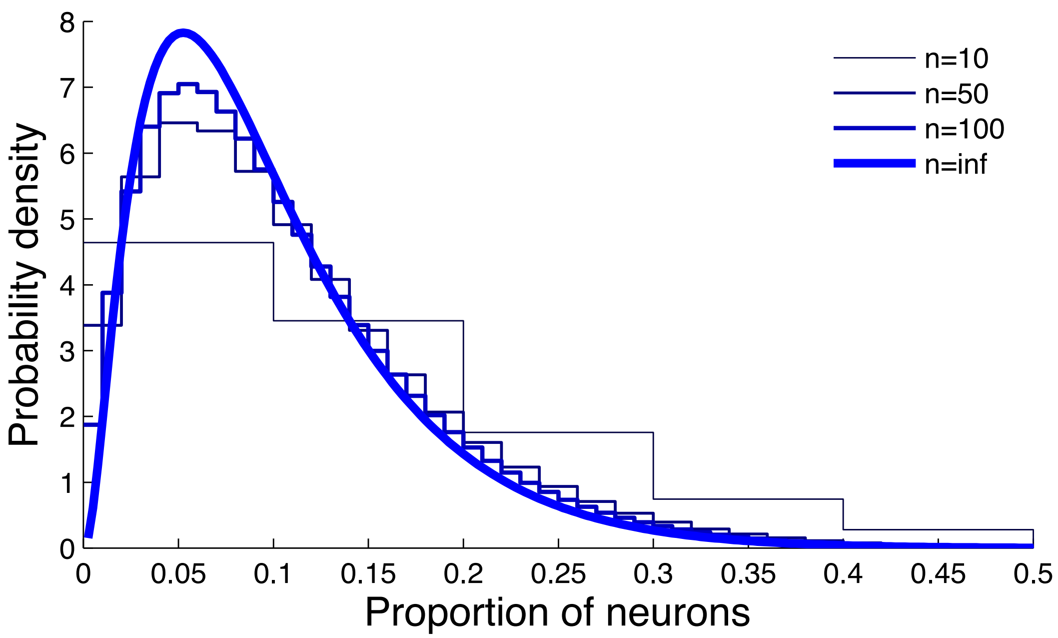

Beta-binomial population model

$$P(K = k) = \binom{N}{k} \frac{\mathrm{Beta}(\alpha + k, \beta + N - k)}{\mathrm{Beta}(\alpha, \beta)}$$

Amari et al 2003, Macke et al 2011, Tkacik et al 2013

Analysis: Subsampling a homogeneous population

Construct populations of size N
by subsampling
a large homogeneous population

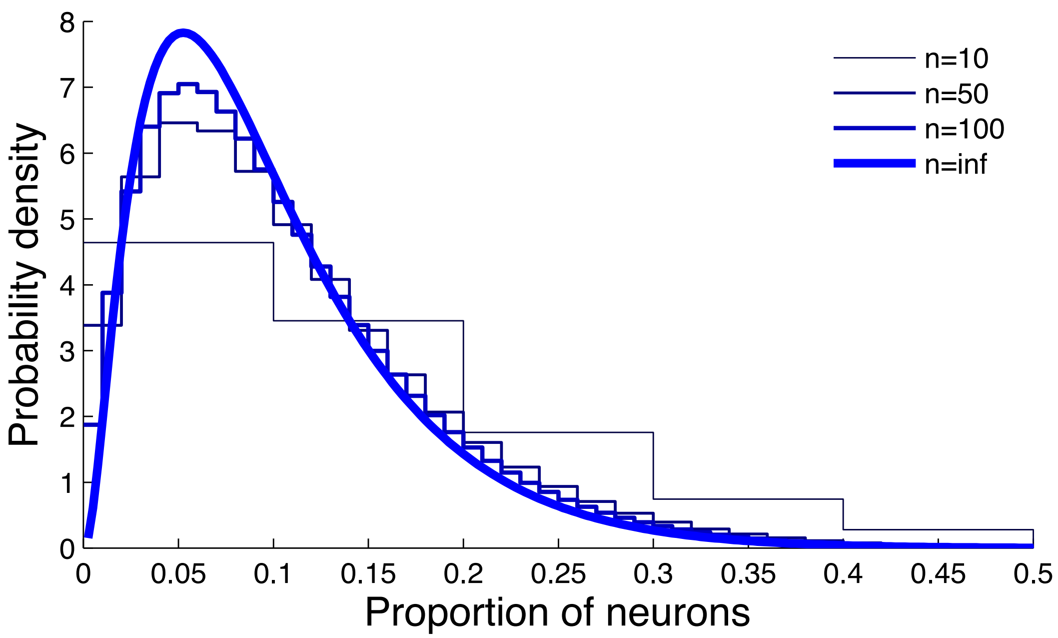

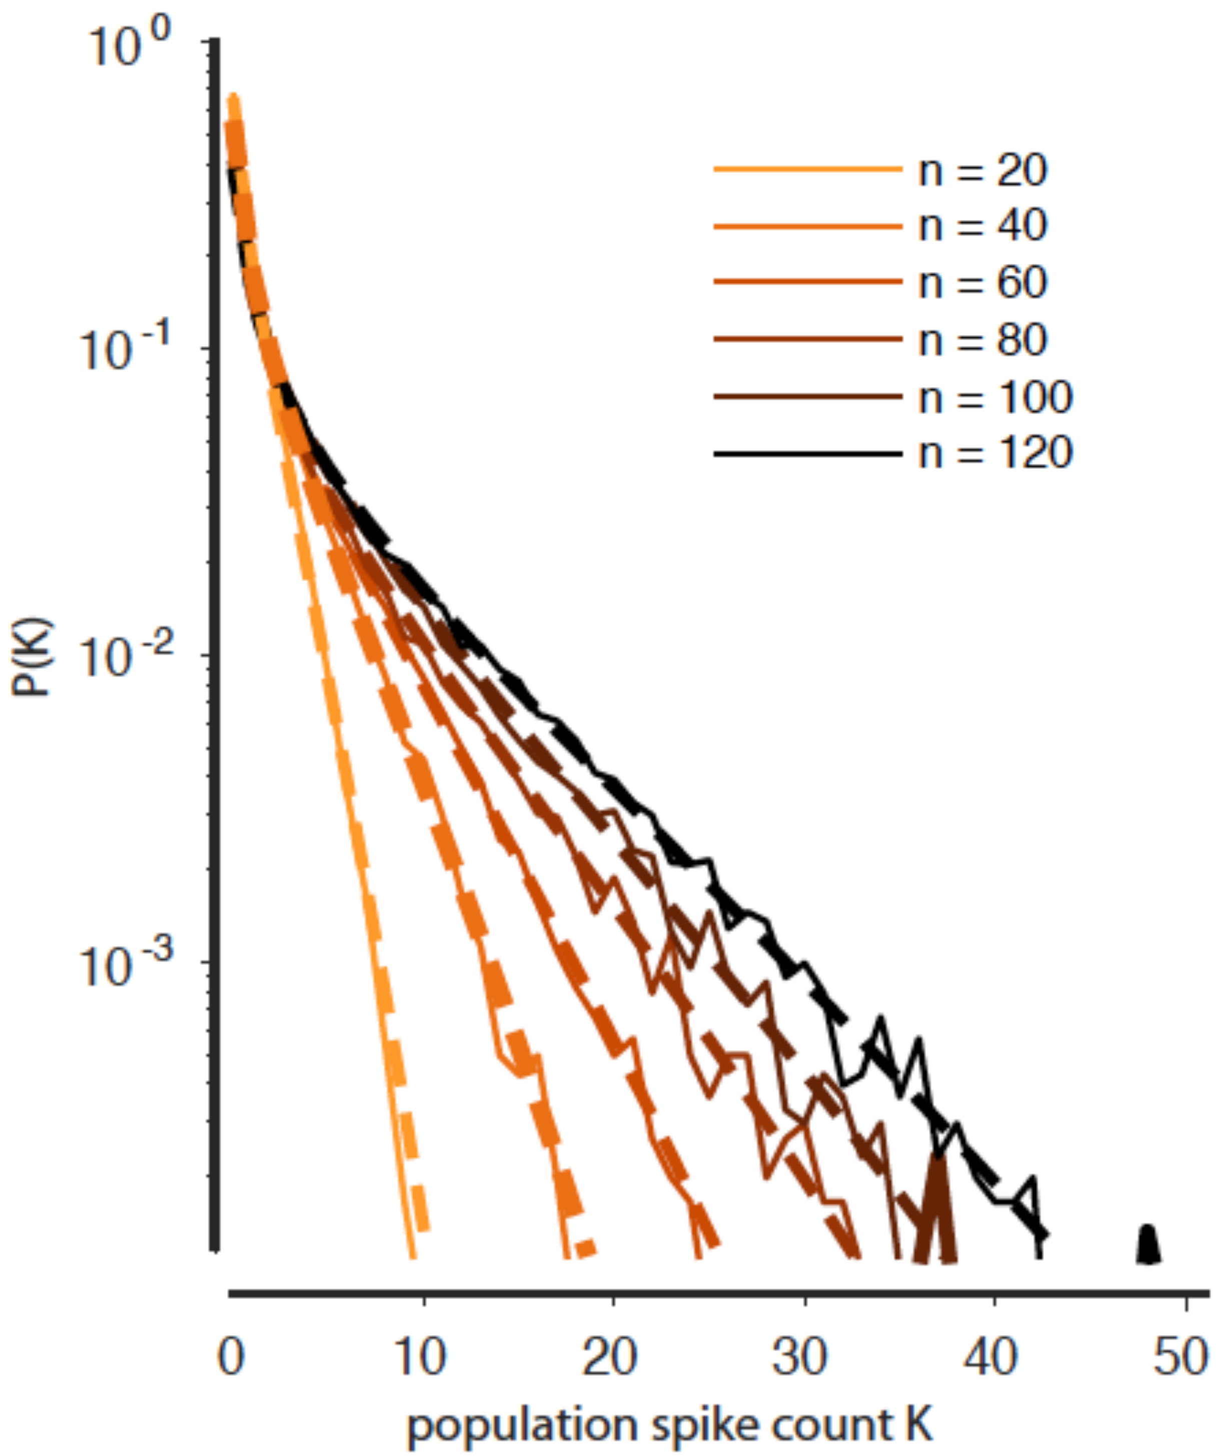

Beta-binomial population model

$$P(K = k) = \binom{N}{k} \frac{\mathrm{Beta}(\alpha + k, \beta + N - k)}{\mathrm{Beta}(\alpha, \beta)}$$

Amari et al 2003, Macke et al 2011, Tkacik et al 2013

In the beta-binomial model, all limits of interest can be calculated in closed form.

Specific heat diverges linearly, and peak always moves to T=1.

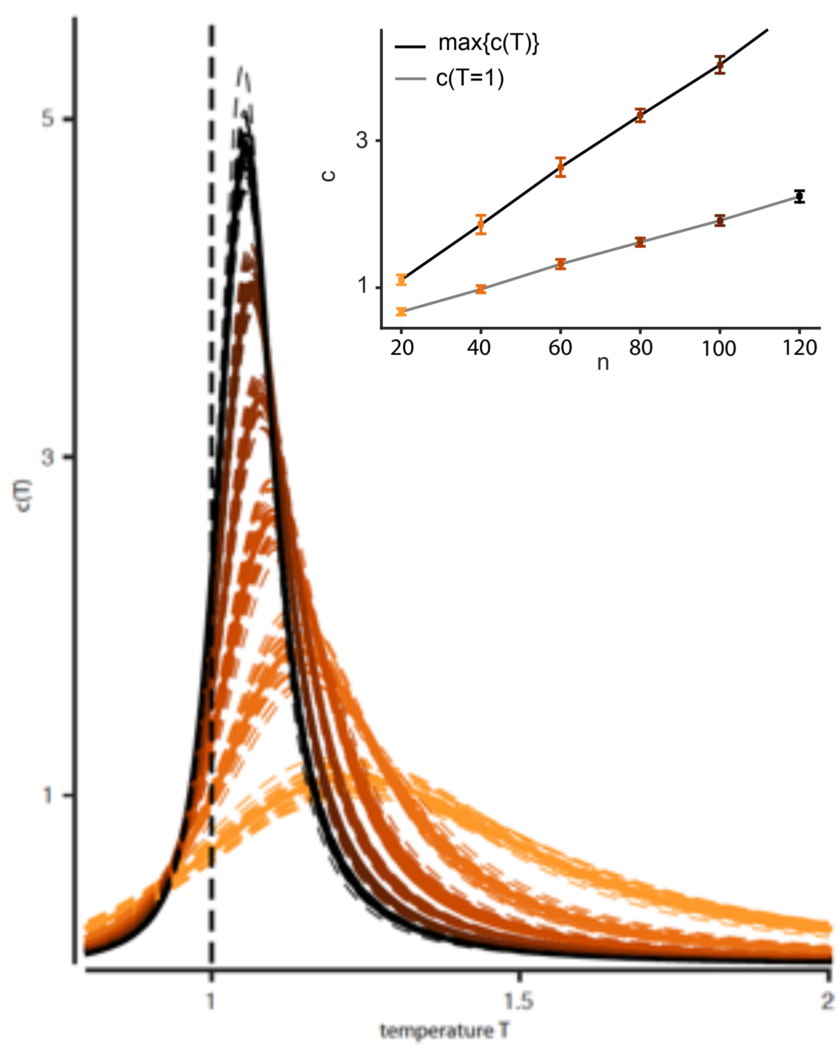

also see Schwab & Metha et al 2014, Aitchison & Latham 2014

In the beta-binomial model, all limits of interest can be calculated in closed form.

Specific heat diverges linearly, and peak always moves to T=1.

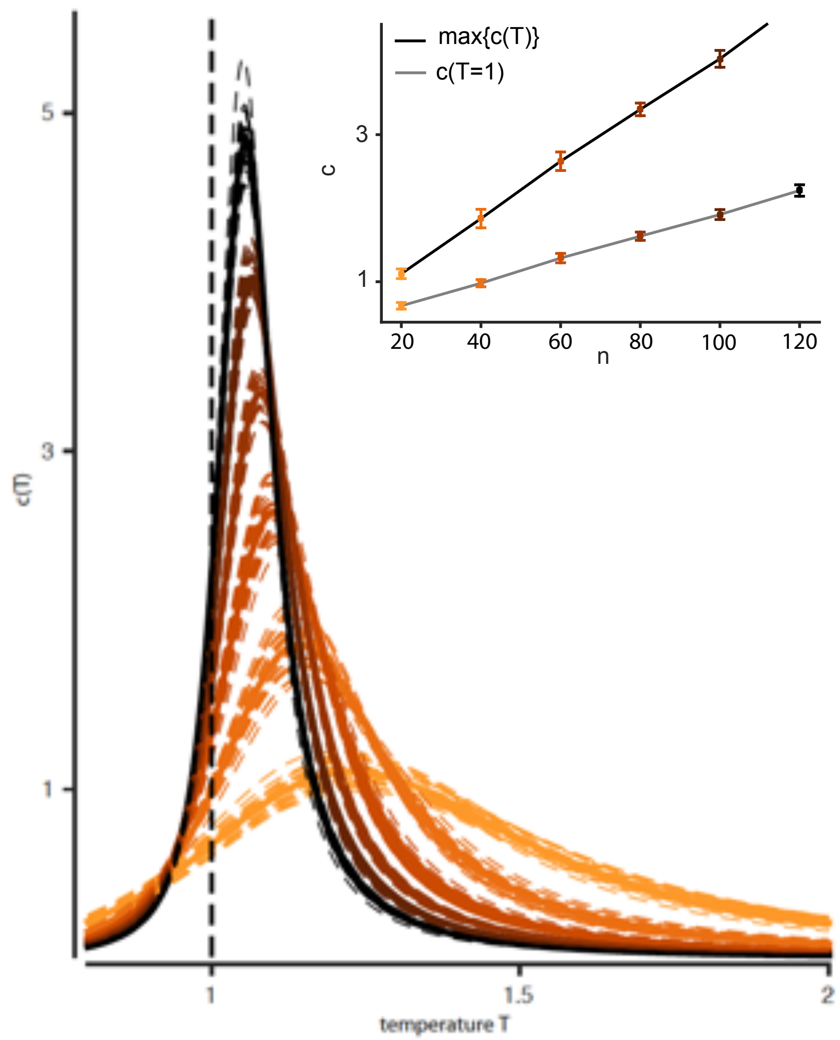

$$c = N\left(\frac{\alpha(\alpha+1)\psi_1(\alpha+1)+\beta(\beta+1)\psi_1(\beta+1)}{(\alpha+\beta)(\alpha+\beta+1)} - \psi_1(\alpha+\beta+1) + \frac{\alpha\beta\left(\psi_0(\alpha+1)-\psi_0(\beta+1)\right)^2}{(\alpha+\beta)^2(\alpha+\beta+1)}\right)$$

also see Schwab & Metha et al 2014, Aitchison & Latham 2014

In the beta-binomial model, all limits of interest can be calculated in closed form.

Specific heat diverges linearly, and peak always moves to T=1.

Rate of divergence: More correlation, more criticality

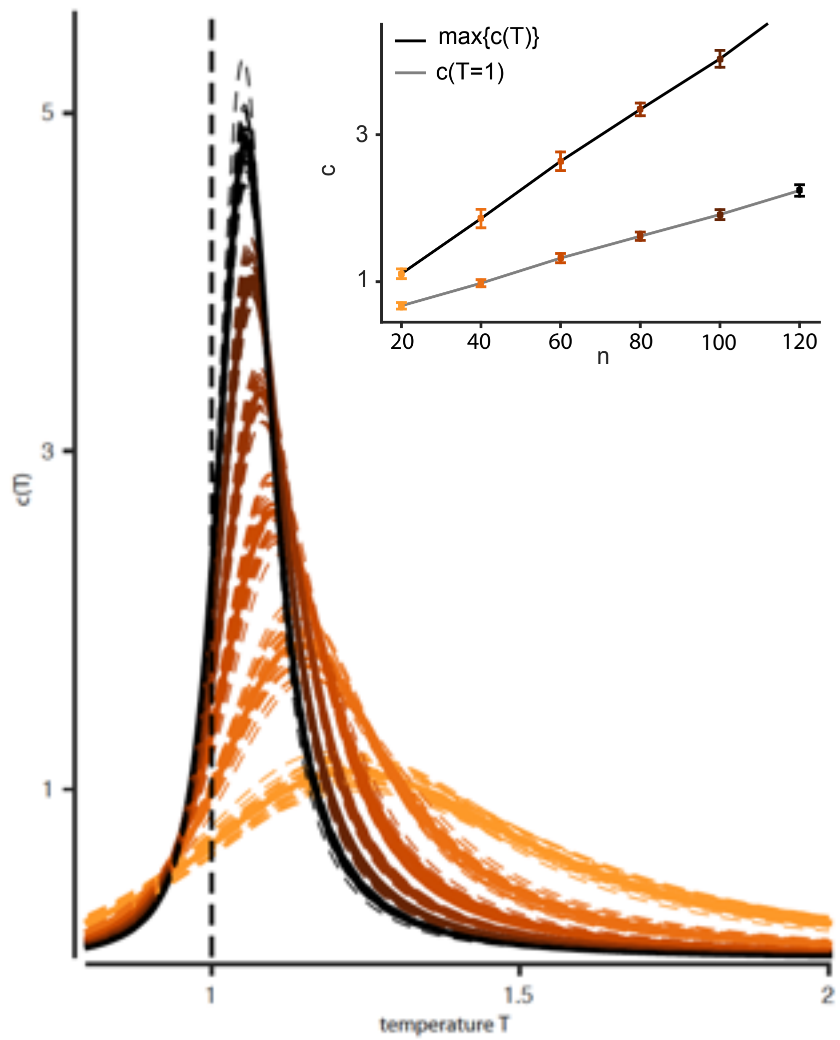

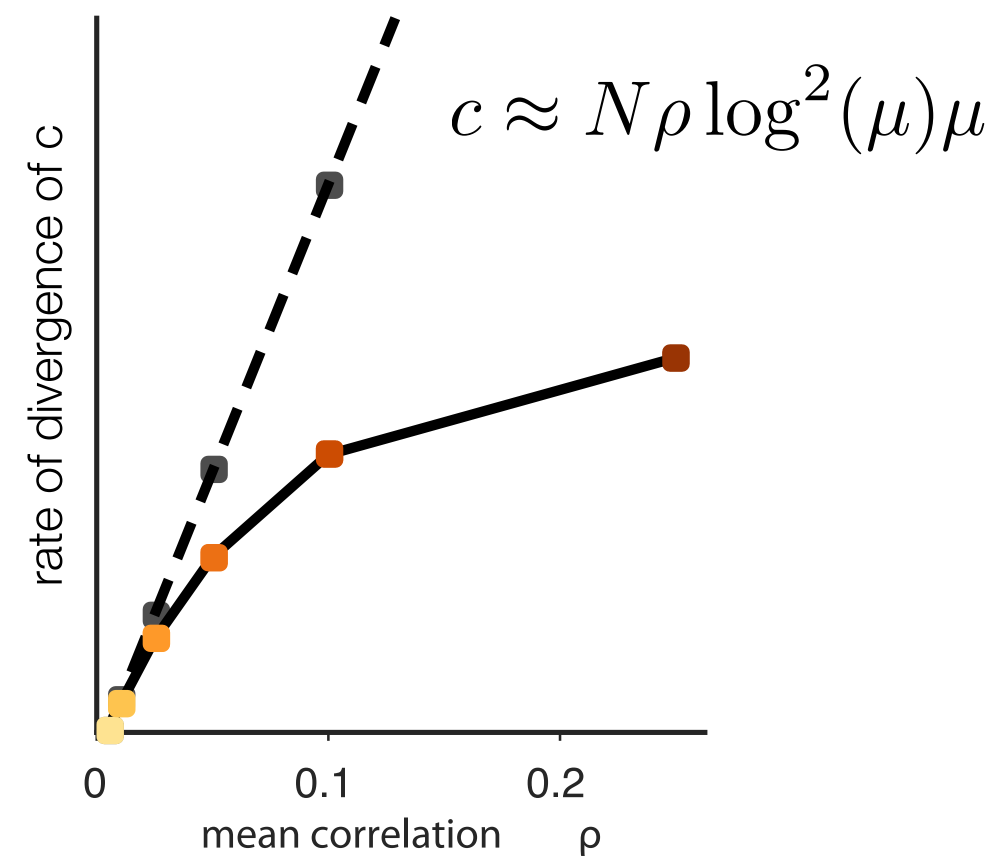

$$c = N\left(\frac{\alpha(\alpha+1)\psi_1(\alpha+1)+\beta(\beta+1)\psi_1(\beta+1)}{(\alpha+\beta)(\alpha+\beta+1)} - \psi_1(\alpha+\beta+1) + \frac{\alpha\beta\left(\psi_0(\alpha+1)-\psi_0(\beta+1)\right)^2}{(\alpha+\beta)^2(\alpha+\beta+1)}\right)$$

also see Schwab & Metha et al 2014, Aitchison & Latham 2014

One source of ‘criticality inducing correlations’: Random subsampling of a correlated population

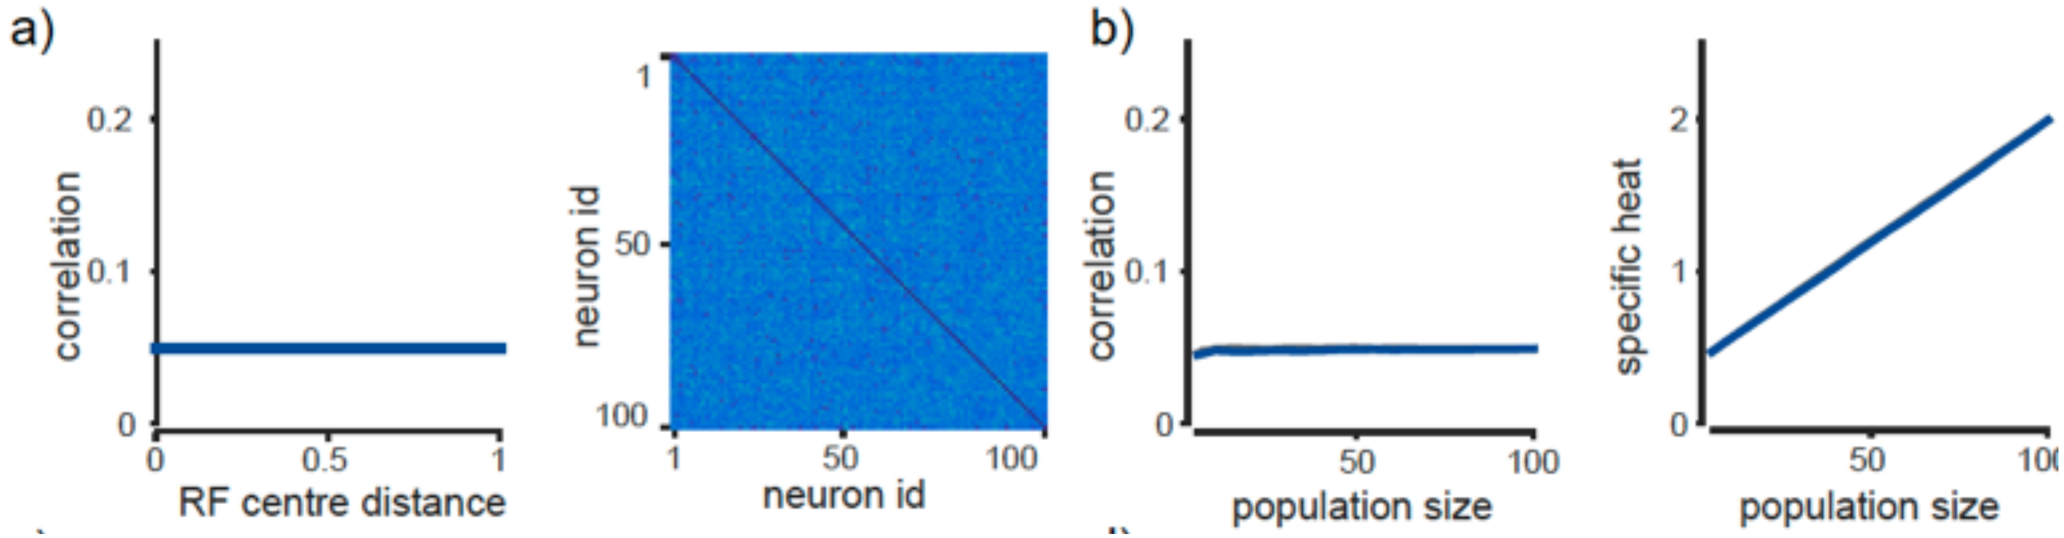

One source of ‘criticality inducing correlations’: Random subsampling of a correlated population

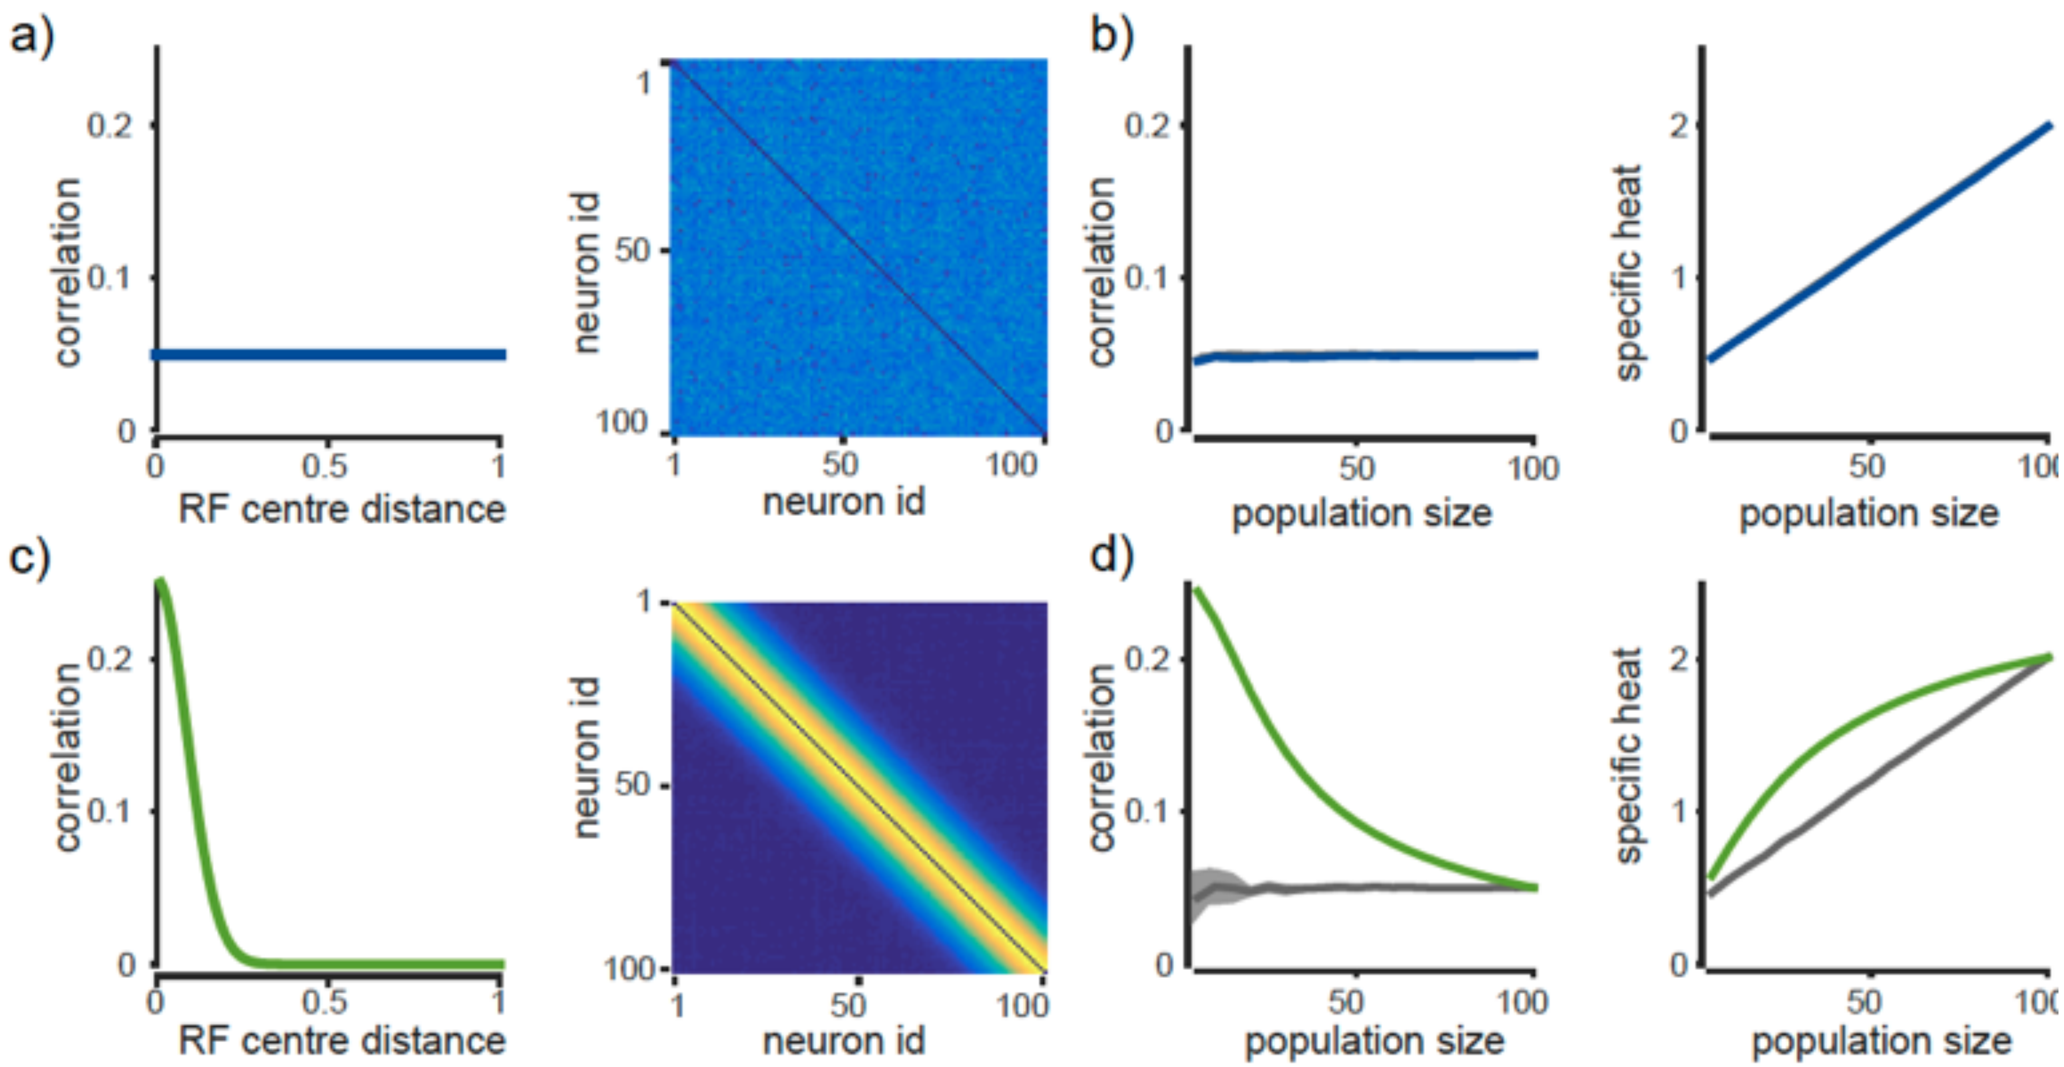

One source of 'criticality inducing correlations': Random subsampling of a correlated population

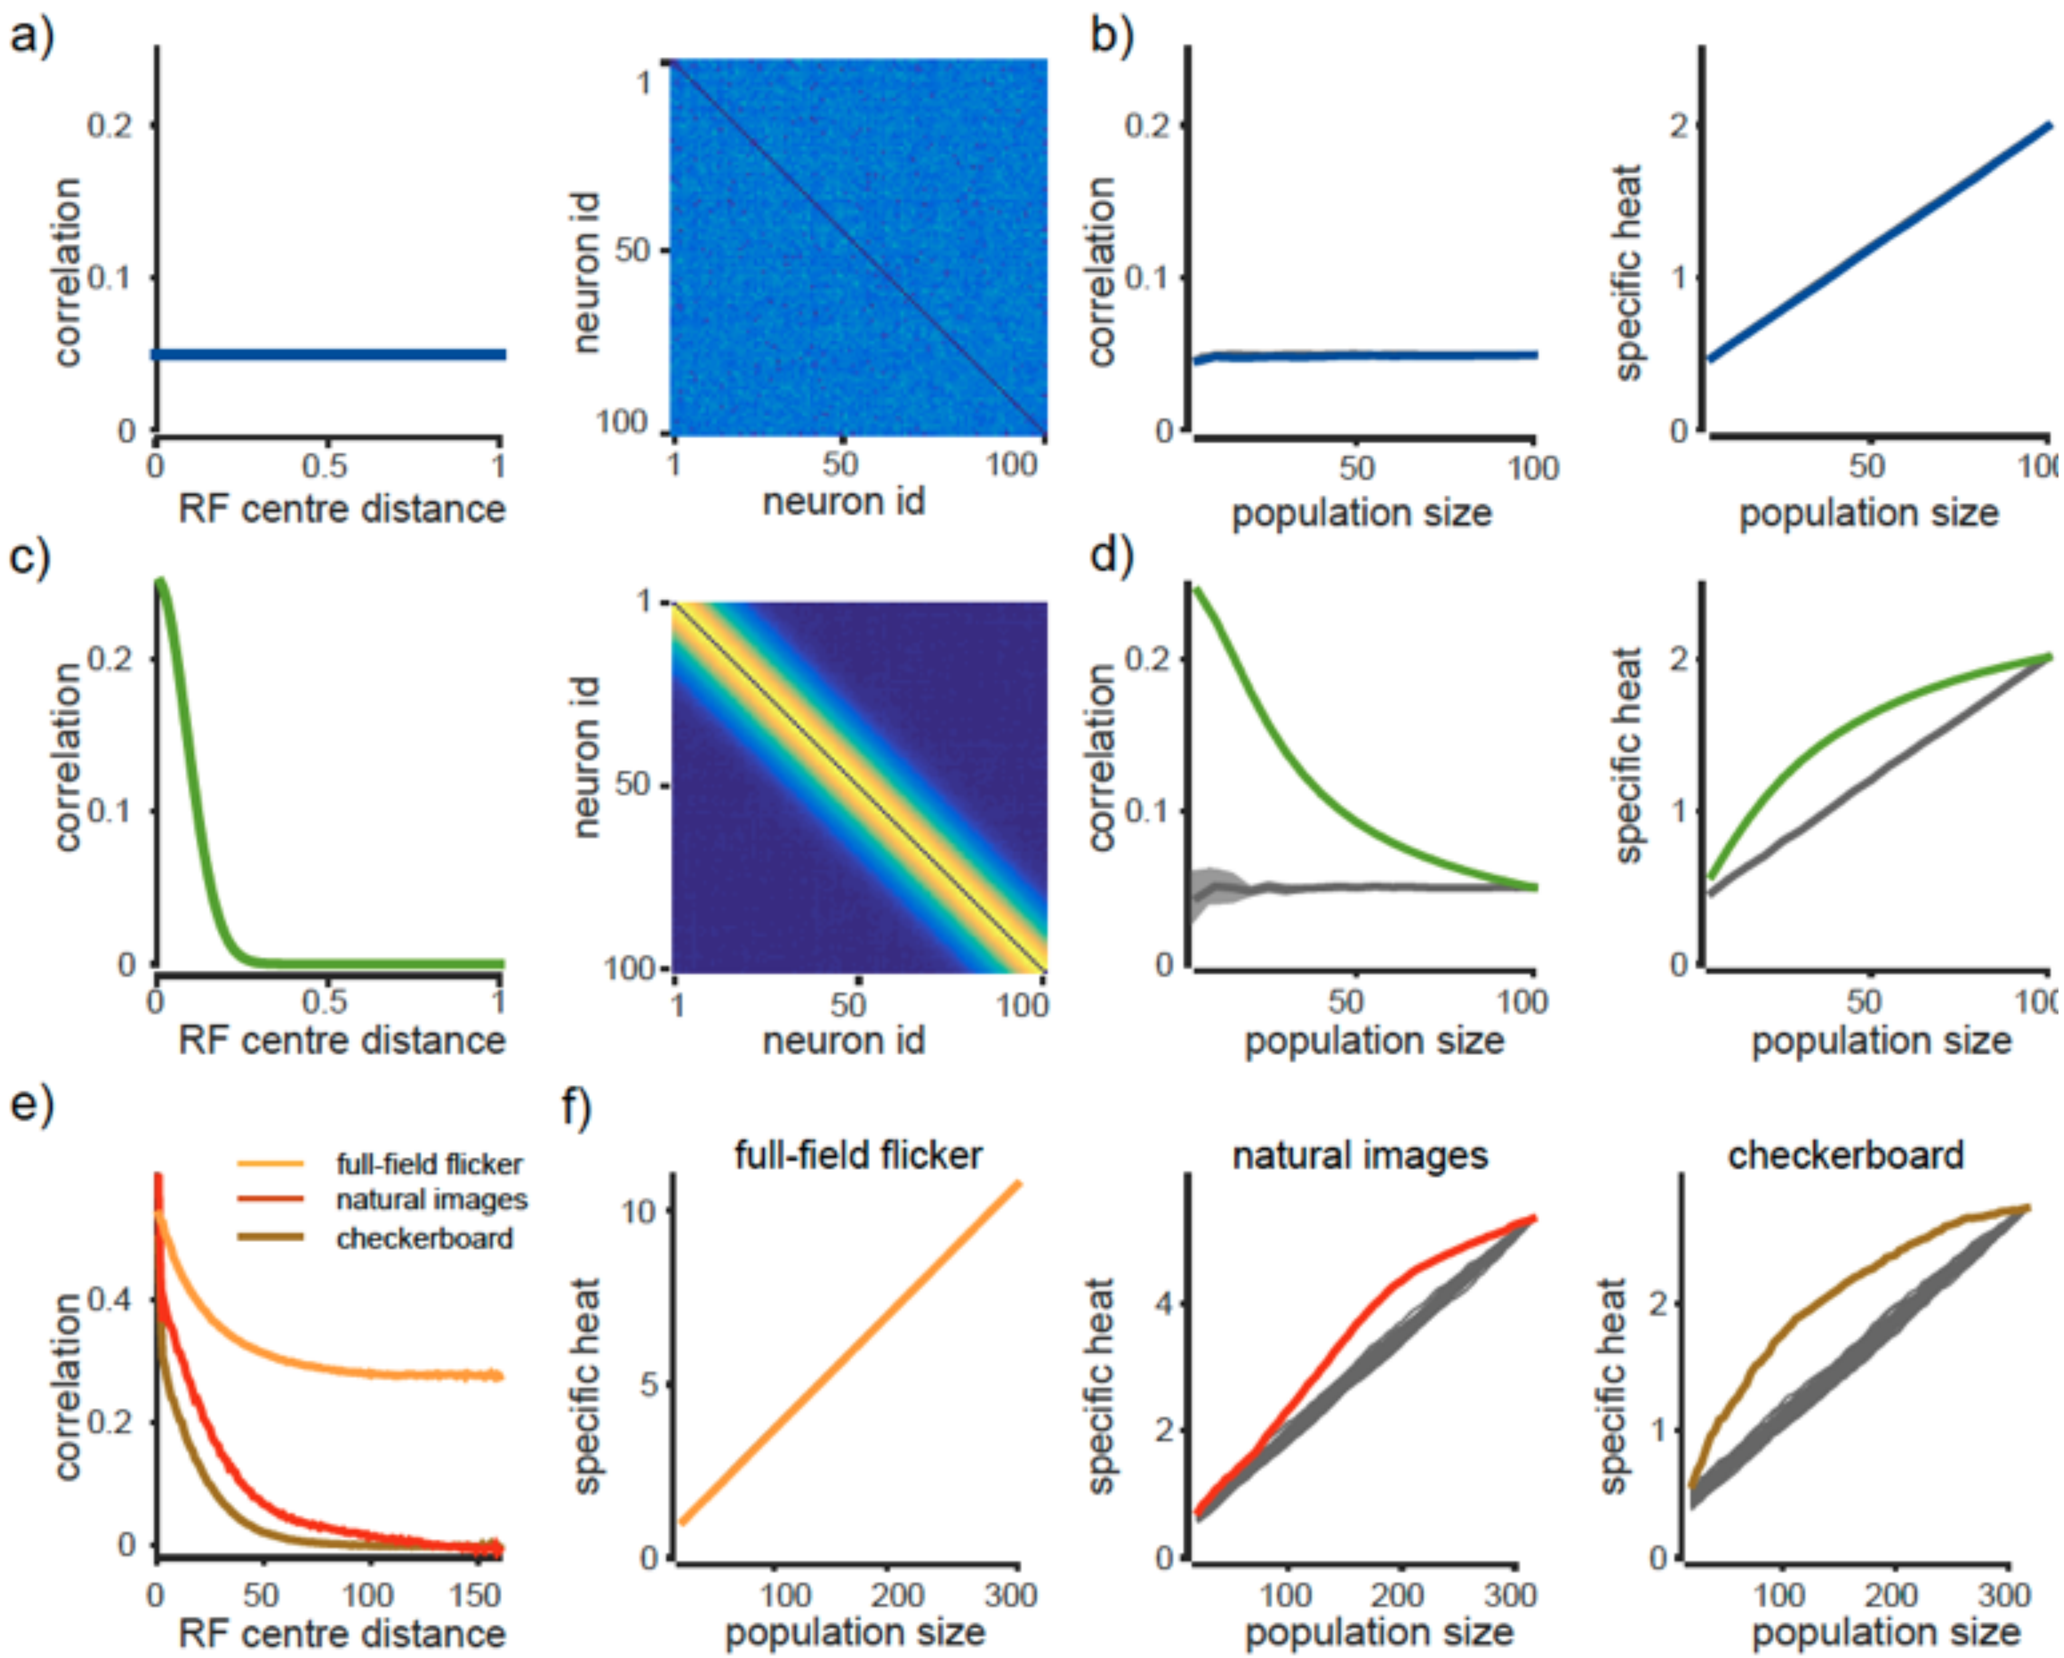

- Neural data analysis:
  - “Infinite-range” correlations are ubiquitous in local neural populations

- Neural data analysis:
  - “Infinite-range” correlations are ubiquitous in local neural populations
  - Shared inputs and latent modulators

- Neural data analysis:
  - “Infinite-range” correlations are ubiquitous in local neural populations
  - Shared inputs and latent modulators
  - Random subsampling gives rise to spurious infinite-range correlations

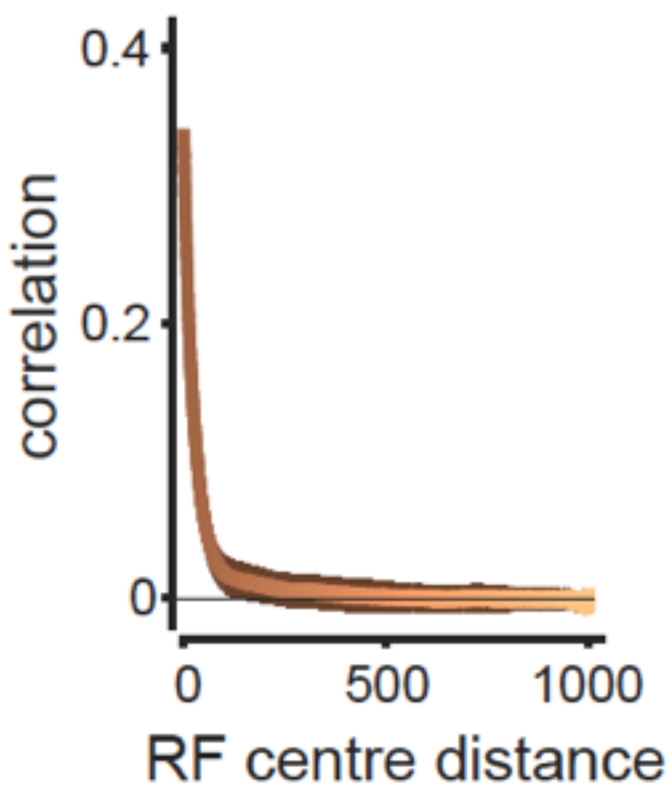

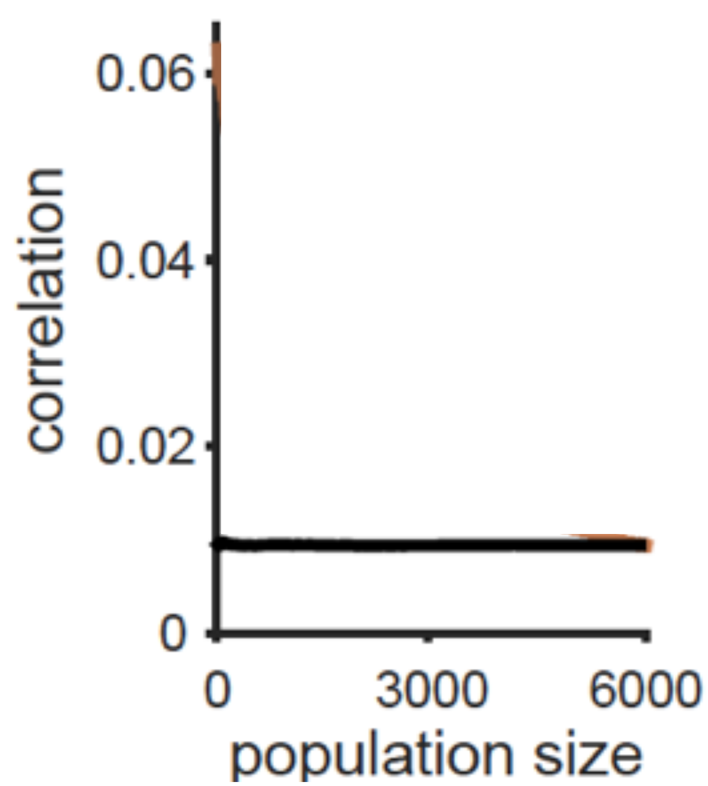

- Neural data analysis:
  - “Infinite-range” correlations are ubiquitous in local neural populations
  - Shared inputs and latent modulators
  - Random subsampling gives rise to spurious infinite-range correlations
- (Equilibrium) statistical physics:

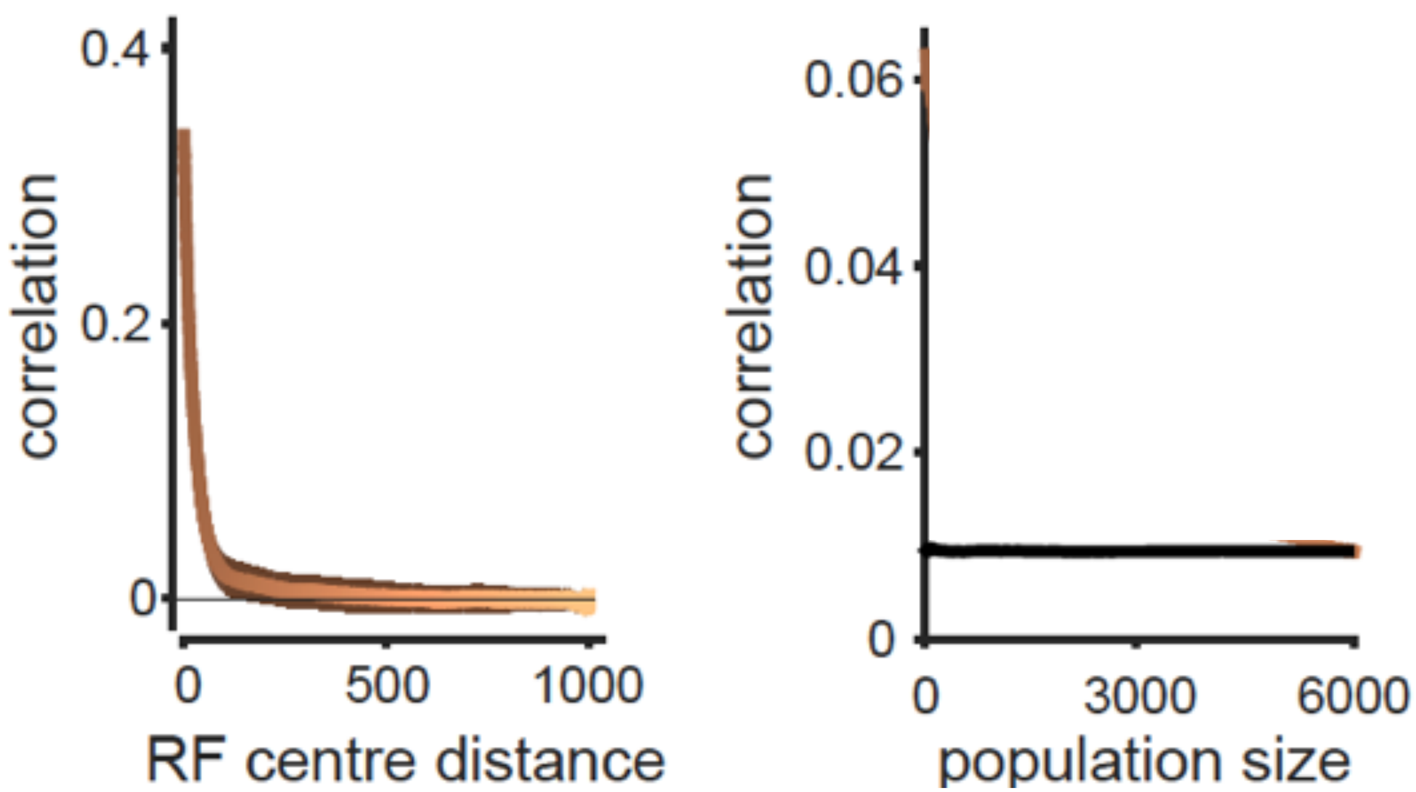

- Neural data analysis:
  - “Infinite-range” correlations are ubiquitous in local neural populations
  - Shared inputs and latent modulators
  - Random subsampling gives rise to spurious infinite-range correlations
- (Equilibrium) statistical physics:
  - “Infinite-range” correlations are unusual

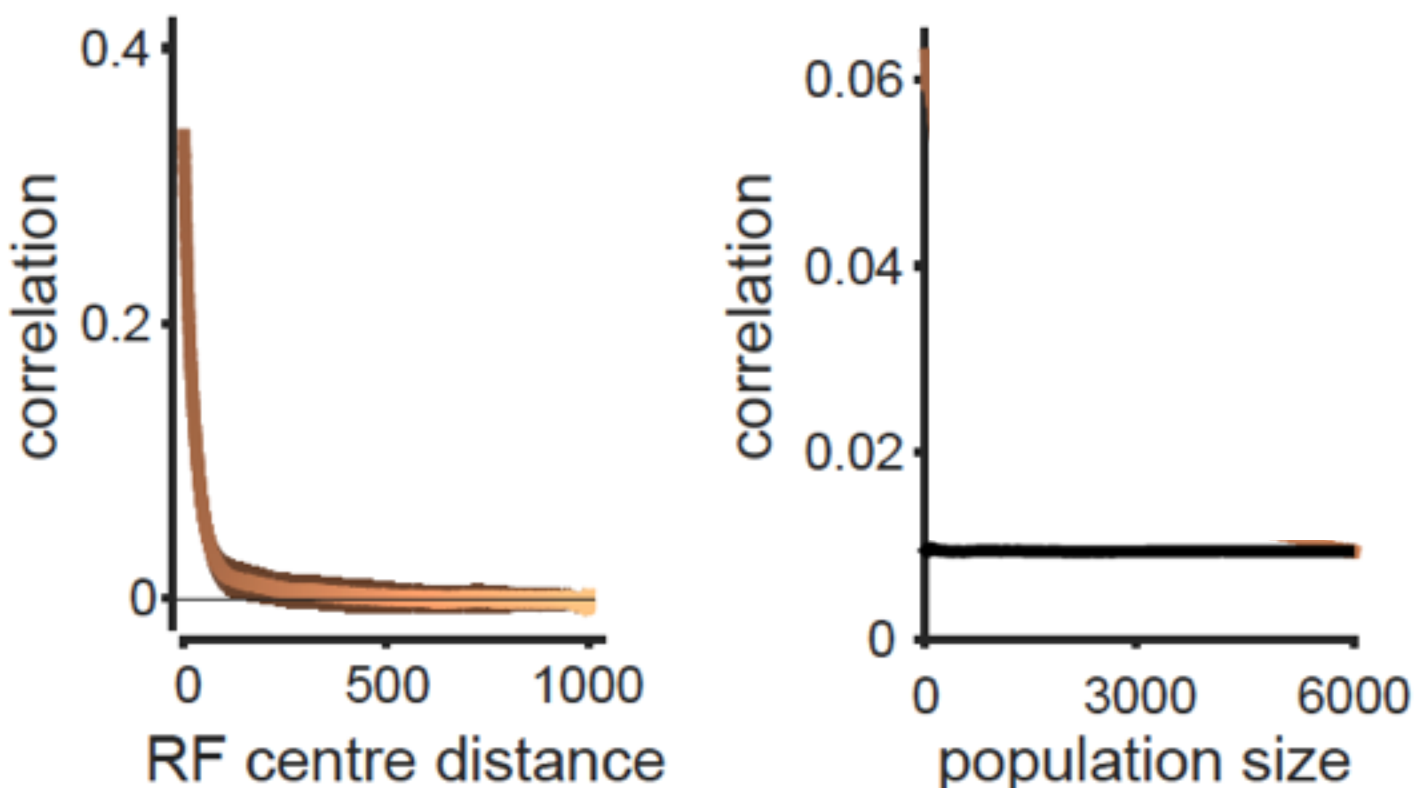

- Neural data analysis:
  - “Infinite-range” correlations are ubiquitous in local neural populations
  - Shared inputs and latent modulators
  - Random subsampling gives rise to spurious infinite-range correlations
- (Equilibrium) statistical physics:
  - “Infinite-range” correlations are unusual
  - No external, shared input

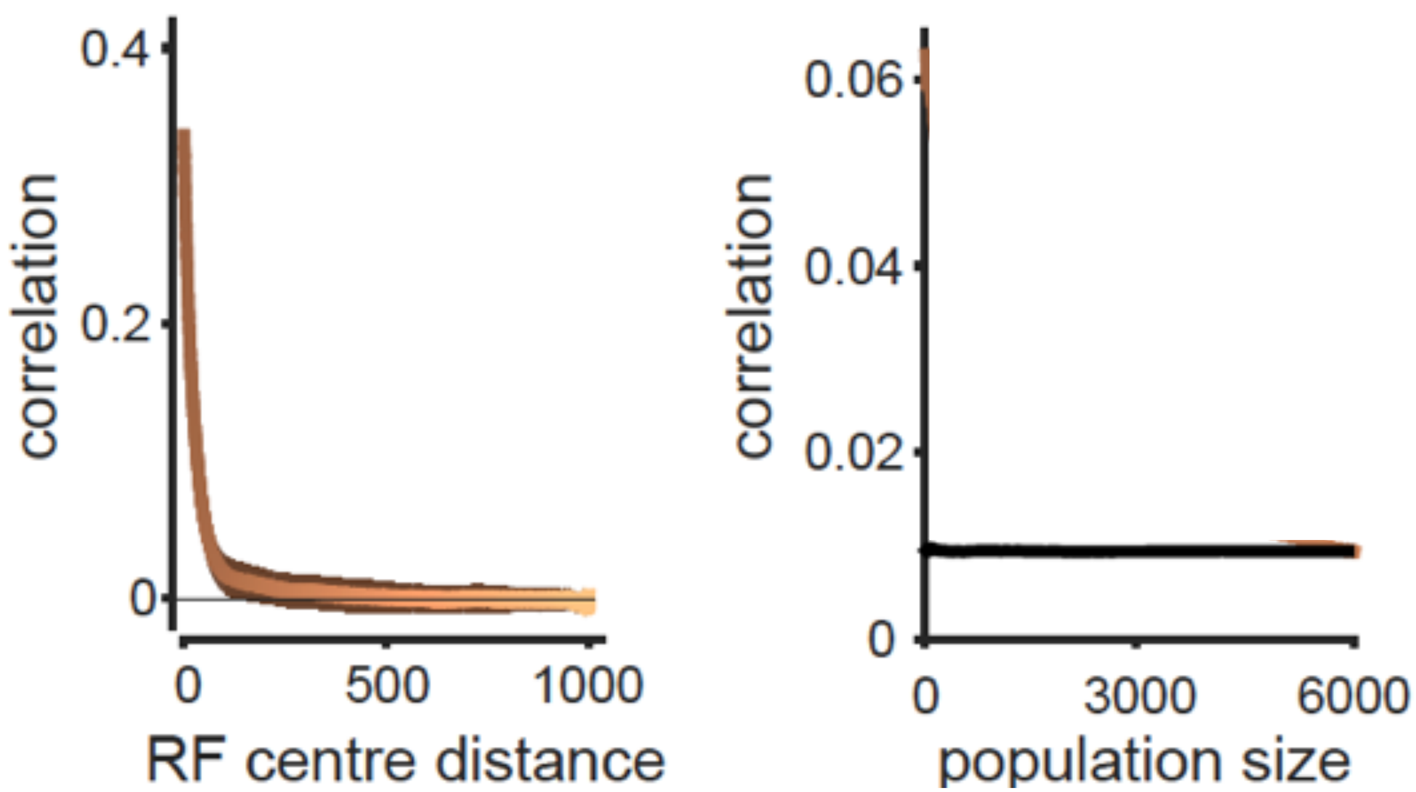

- Neural data analysis:
  - “Infinite-range” correlations are ubiquitous in local neural populations
  - Shared inputs and latent modulators
  - Random subsampling gives rise to spurious infinite-range correlations
- (Equilibrium) statistical physics:
  - “Infinite-range” correlations are unusual
  - No external, shared input
  - N is varied by considering systems of different size

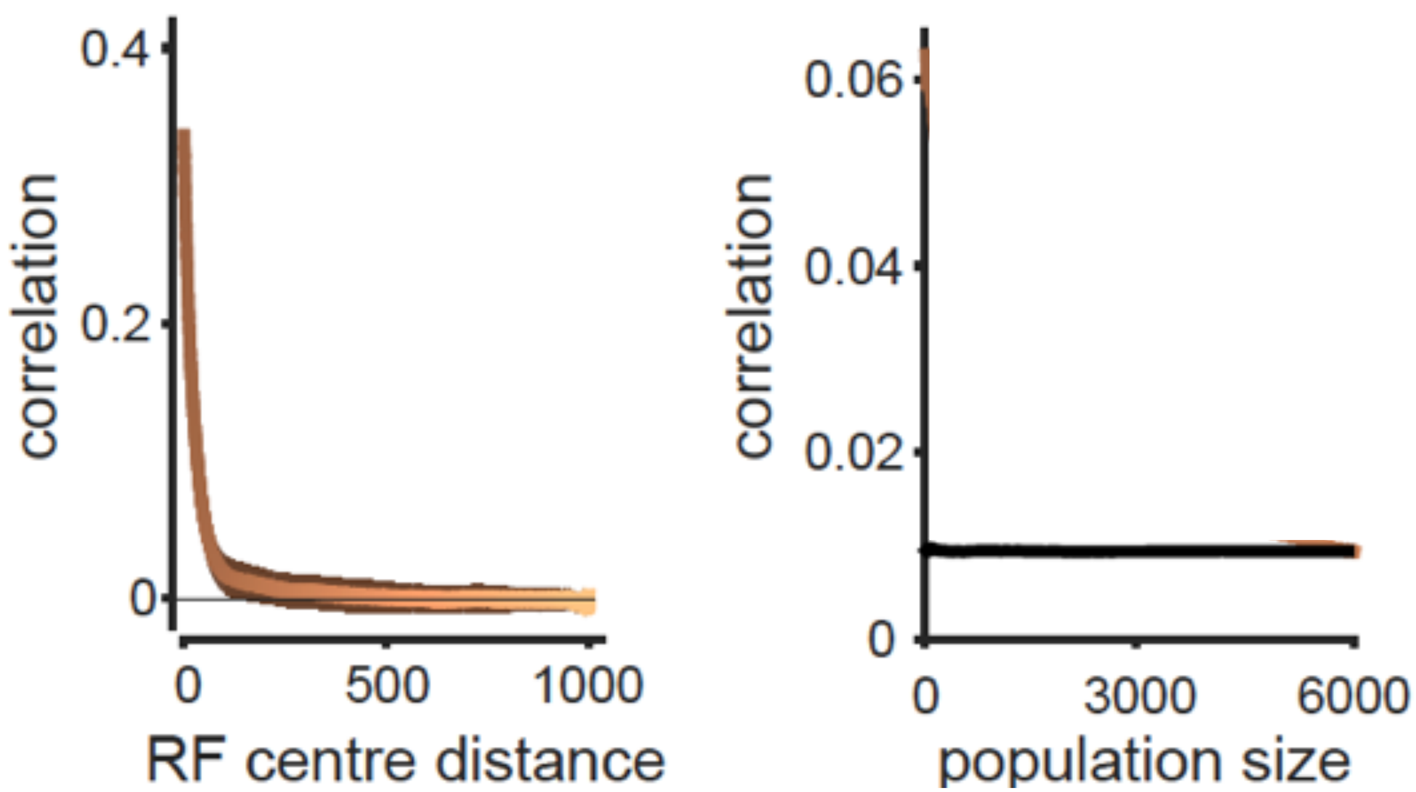

- We concentrated on the specific heat, and did not consider other signatures of criticality, such as Zipf-Scaling (Tkacik et al 2015, Aitchison et al 2014, Schwab et al 2014).

- We concentrated on the specific heat, and did not consider other signatures of criticality, such as Zipf-Scaling (Tkacik et al 2015, Aitchison et al 2014, Schwab et al 2014).
- Statistics of neural data can look unusual when viewed through the lens of equilibrium statistical physics, even when the underlying mechanisms are not.

- We concentrated on the specific heat, and did not consider other signatures of criticality, such as Zipf-Scaling (Tkacik et al 2015, Aitchison et al 2014, Schwab et al 2014).
- Statistics of neural data can look unusual when viewed through the lens of equilibrium statistical physics, even when the underlying mechanisms are not.
- We predict that signatures of criticality will also be found in other systems and brain areas.

- We concentrated on the specific heat, and did not consider other signatures of criticality, such as Zipf-Scaling (Tkacik et al 2015, Aitchison et al 2014, Schwab et al 2014).
- Statistics of neural data can look unusual when viewed through the lens of equilibrium statistical physics, even when the underlying mechanisms are not.
- We predict that signatures of criticality will also be found in other systems and brain areas.
- Is thermodynamic criticality an organising principle of neural codes?

Thanks to ...

Marcel Nonnenmacher

Christian Behrens
(Uni Tübingen)

Philipp Berens
(Uni Tübingen)

Matthias Bethge
(Uni Tübingen)

caesar

center of advanced
european studies
and research

research center caesar -
an associate of the
Max Planck Society

Bernstein Center for
Computational Neuroscience
Tübingen

IMPRS for Brain & Behavior

International Max Planck Research School

... and you all for your attention.