Learning and memory

Arvind Murugan
Assistant Professor, Physics
U Chicago
Nov 2017

Associative memory

How do you keep multiple memories from interfering?

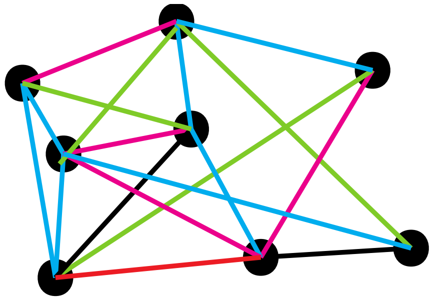

Learning and memory

Memory of examples vs learning from examples

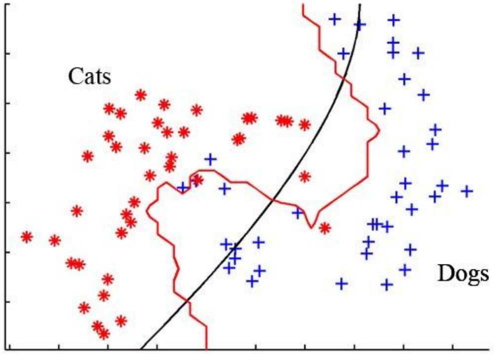

Not for today

Temporal dynamics in biology

Circadian clocks Metabolism

w/ Rust lab (U Chicago)

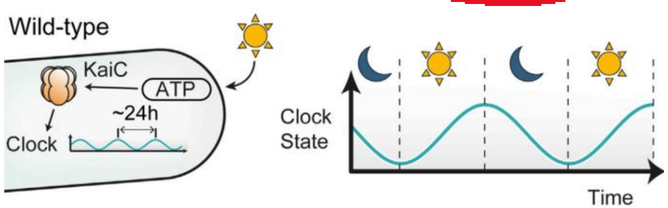

Kalman filter..

Temporal control of gene regulation, fate etc

w/ Tay lab (U Chicago) + others

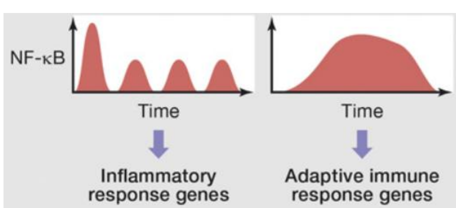

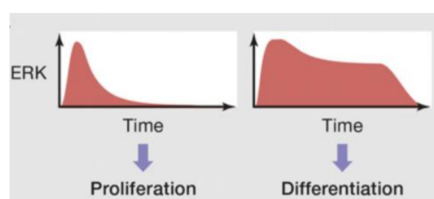

Purvis/Lahav 2014

Specificity, allostery/ cooperativity in time..

Associative memory in neural networks

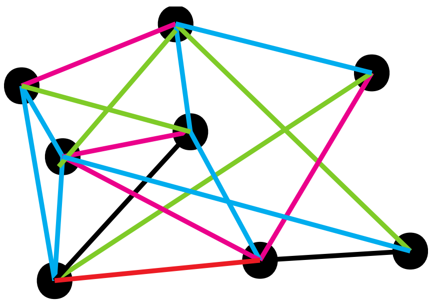

$$J_{ij} = J_{ij} + J_{ij} + J_{ij}$$

$$H = \sum_{ij} J_{ij} s^i s^j$$

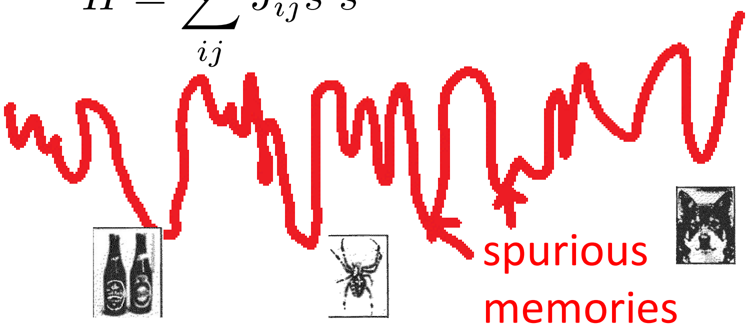

Above capacity

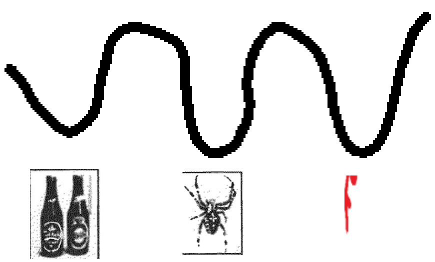

Under capacity
(Hopfield 1982)

Initial Cond.

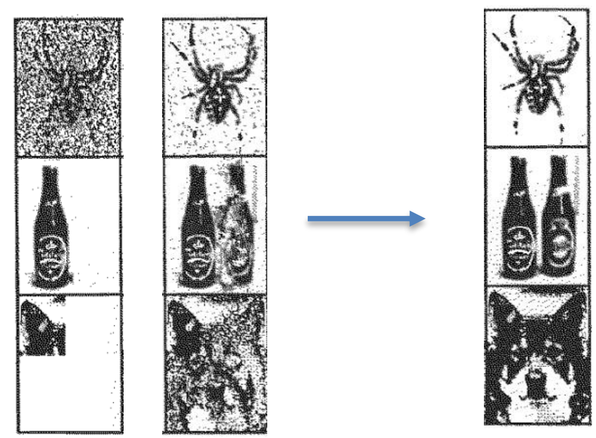

Retrieval by association

Associative memory in neural networks

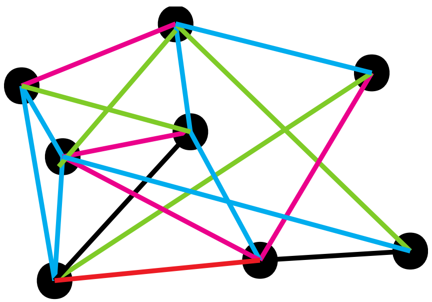

$$J_{ij} = J_{ij} + J_{ij} + J_{ij}$$

Hopfield 1982
Amit et al 1985

$$H = \sum_{ij} J_{ij} s^i s^j$$

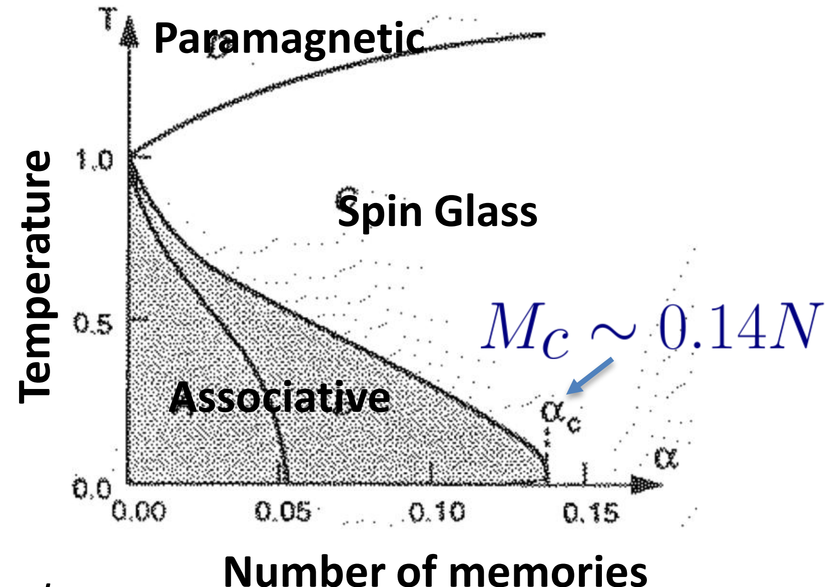

$$J_{ij} = x_i x_j, \quad \vec{x} = \text{spider}$$

$$J_{ij} = x_i x_j, \quad \vec{x} = \text{bottles}$$

$$J_{ij} = x_i x_j, \quad \vec{x} = \text{wolf}$$

$$\vdots$$

Final state

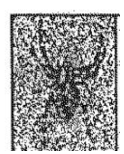

Success

Initial Cond.

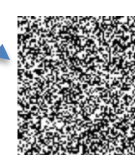

Failure

Associative memory in neural networks

1. Correlations between memories reduce capacity

Ideally: $\vec{x} \cdot \vec{x} = 0$

2. Complex learning rules can (slightly) increase capacity

Hopfield w/ linear Hebbian rule

$$J_{ij} = J_{ij} + J_{ij} + J_{ij}$$

$$M_C = 0.14N$$

E. Gardner

Optimal J_{ij}

$$M_C = 2N$$

Associative memory in neural networks

3. Range of interactions is important

Fully connected
(infinite dimensions)

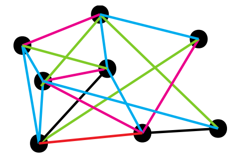

$$M_C = 0.14N$$

$$J_{ij} = x_i x_j, \quad \vec{x} = \text{spider}$$

Finite dimensions

$$M_C = O(1)$$

4. Nature of memories is important

Original model :
Each memory = point attractor

Place cell model (spatial memories):
Each memory = continuous attractor

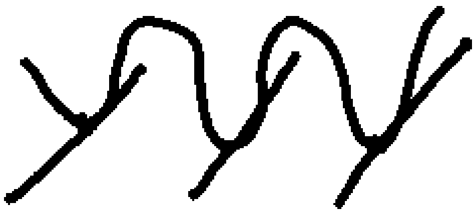

Associative memory in materials

Zorana Zeravcic

Stanislas Leibler

Michael Brenner

PNAS 2015

J. Stat. Phy. 2017
- W Zhong, D. Schwab

Menachem Stern

Nat. Comm. 2017,
PRX 2017
+ in progress

DNA Brick assembly

Yin lab,
Harvard Medical School

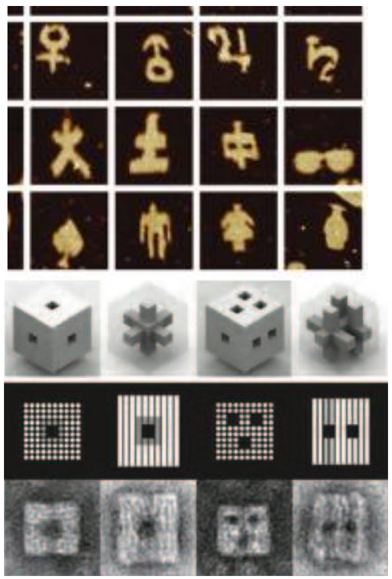

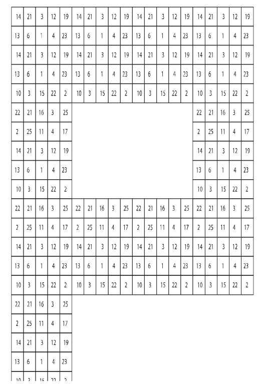

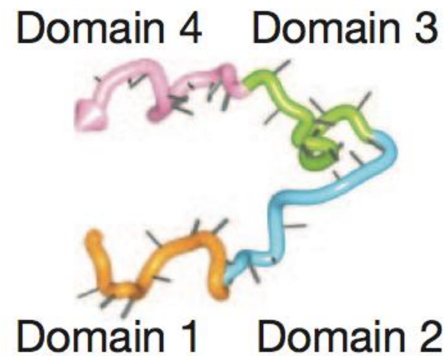

Exactly 1 partner for every binding site

Abstract view

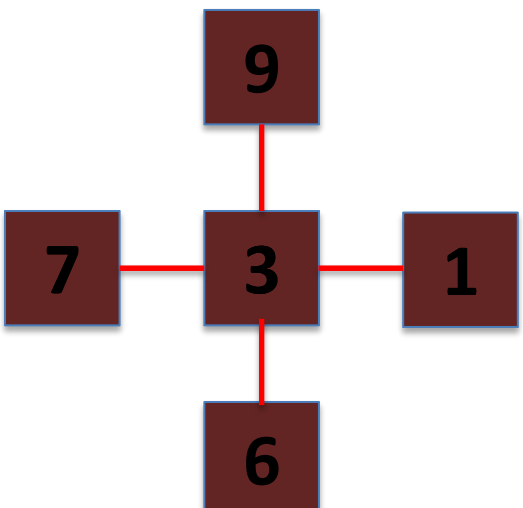

Assembly mixtures

- We ask:

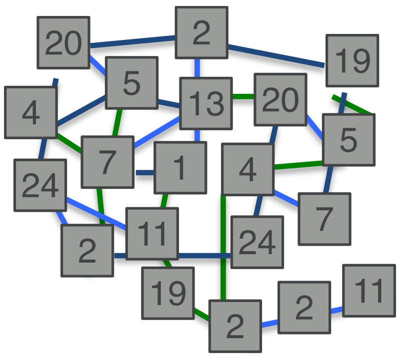

Cue A

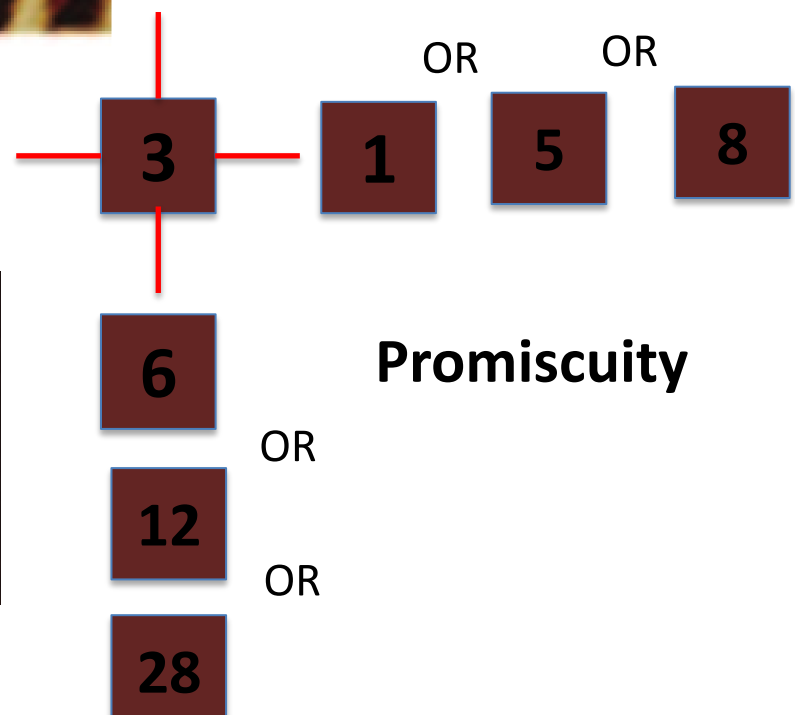

Assembly mixtures

- General model:

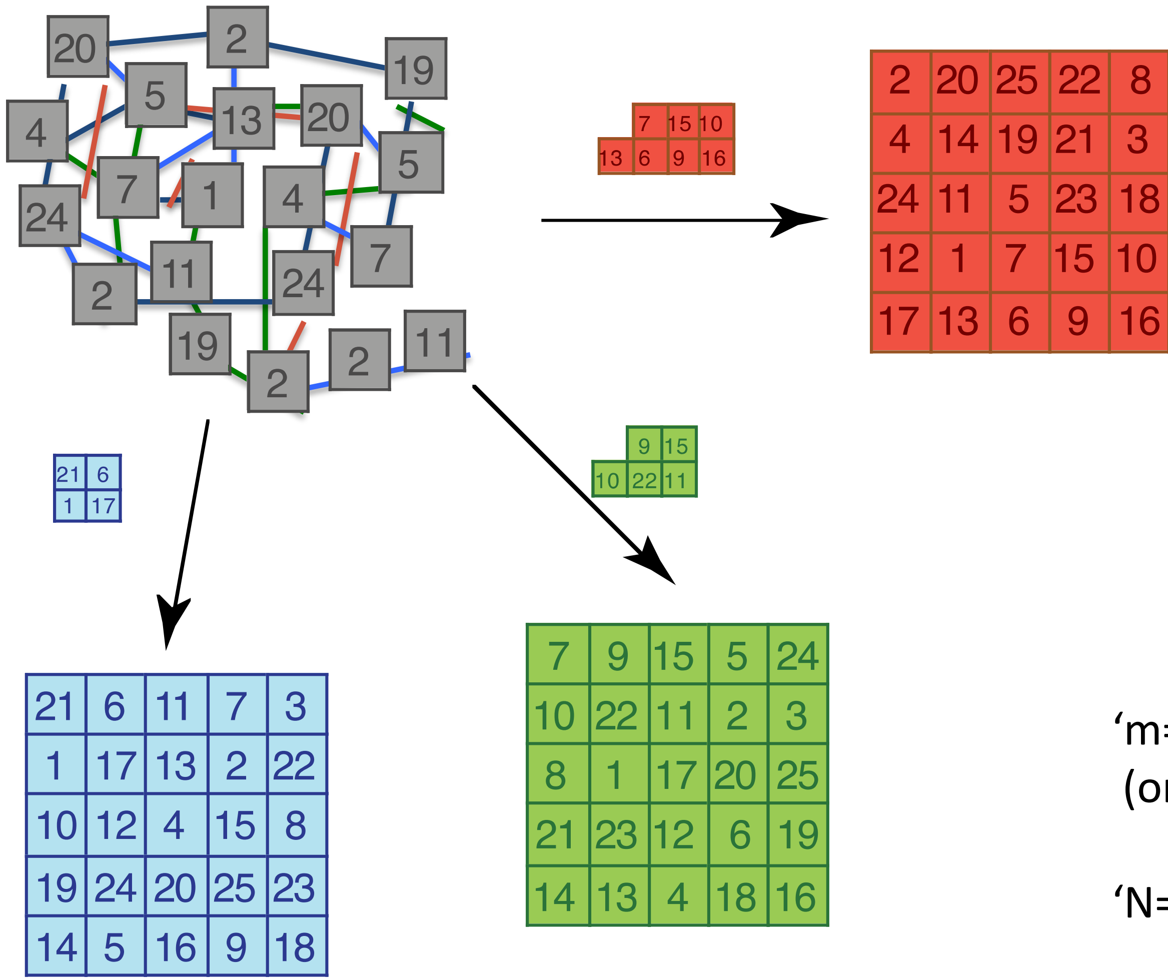

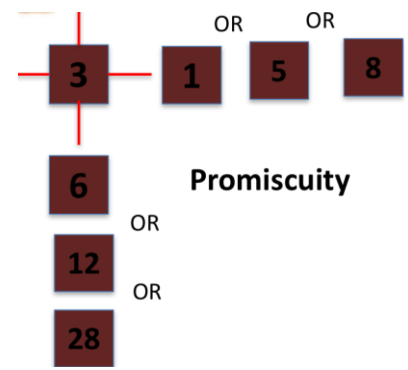

'm=3' stored structures
(or 'memories')

'N=25' species

Colors represent bonds

Monte-Carlo Simulations

m=5 stored memories

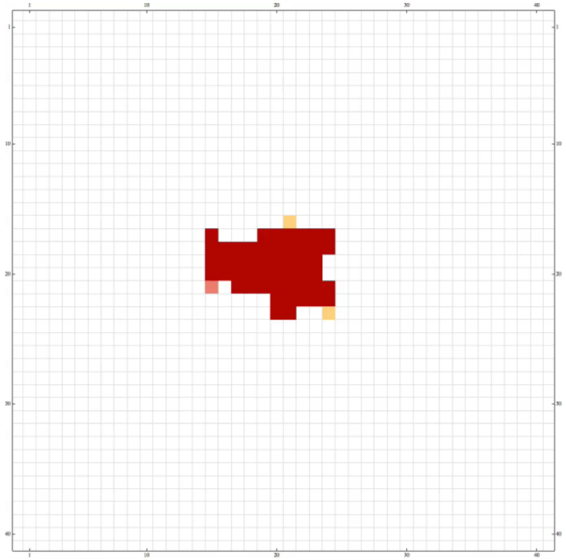

m=25 stored memories

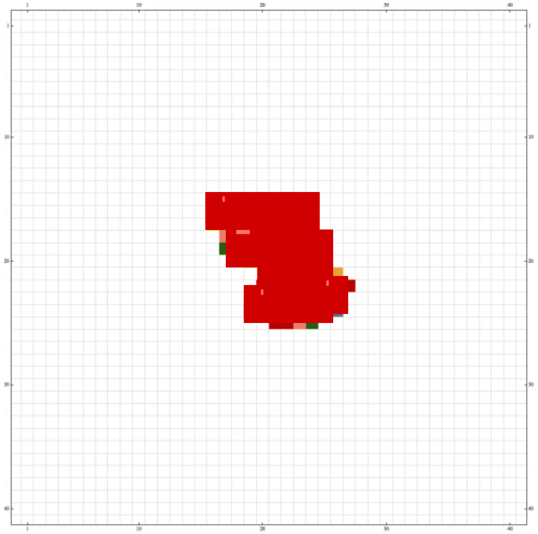

Monomers not shown

Parameters:

N = 400 species (20x20),
Bond energy = E,
Conc. = exp(1.8 E),
T = 0.15 E

Phase diagram

N = 400 components (20x20)

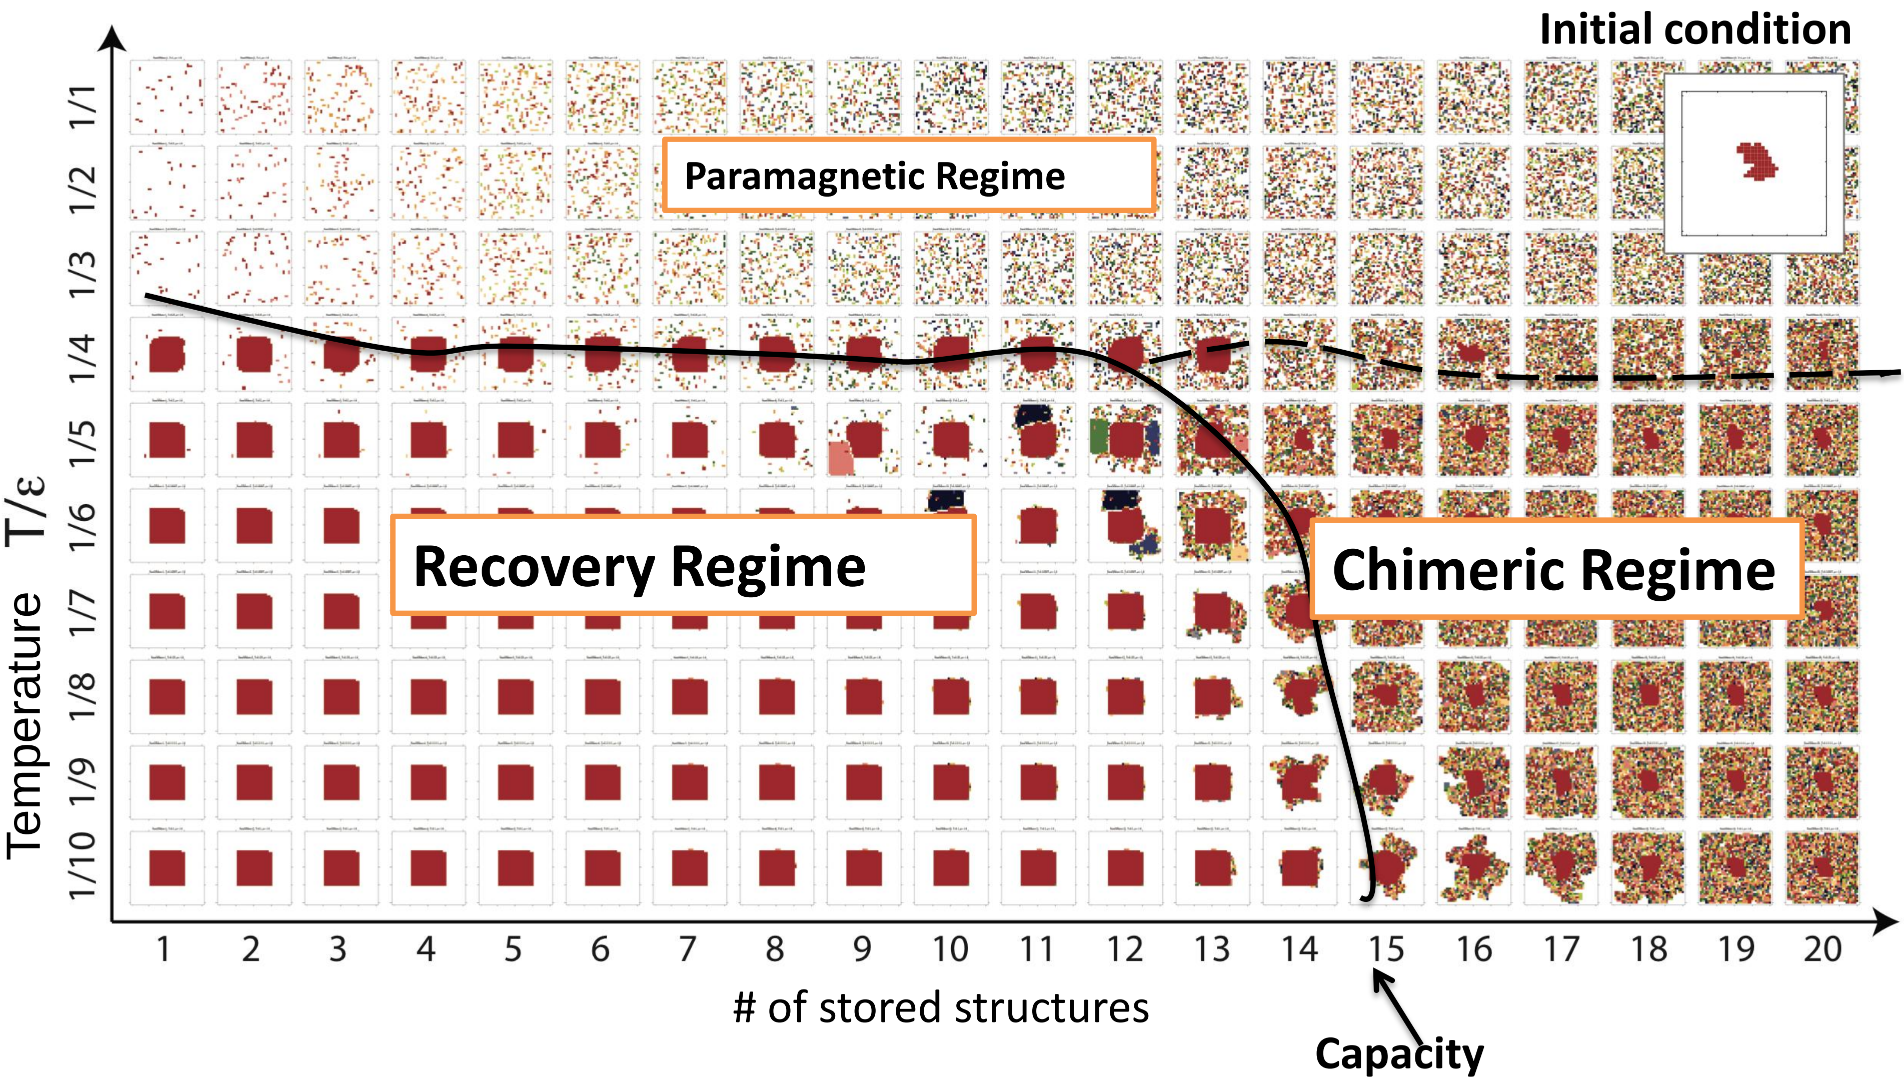

Promiscuity balanced by frustration

Memory 1 Memory 2 Memory 3

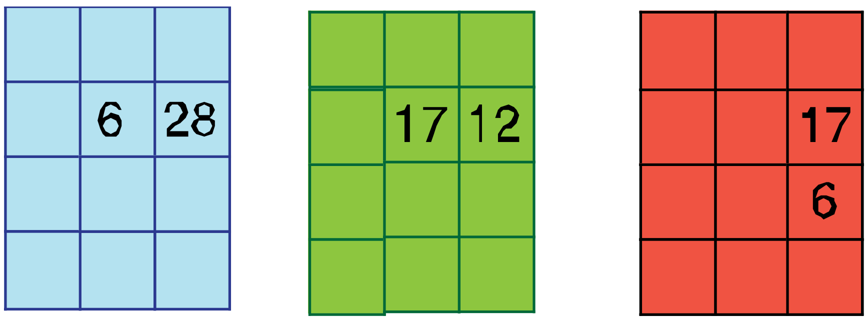

‘m’ stored memories

Total of ‘N’ species

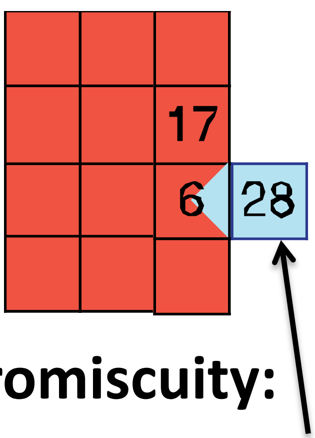

Promiscuity:

‘m’ local choices

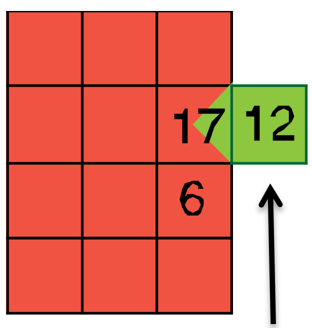

‘m’ local choices

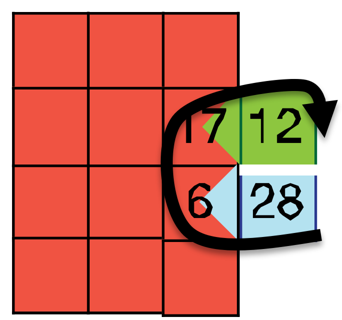

The friend (12) of a friend (17) of a friend (6) .. may **not** be a friend (of 28).

Frustration

Promiscuity balanced by frustration

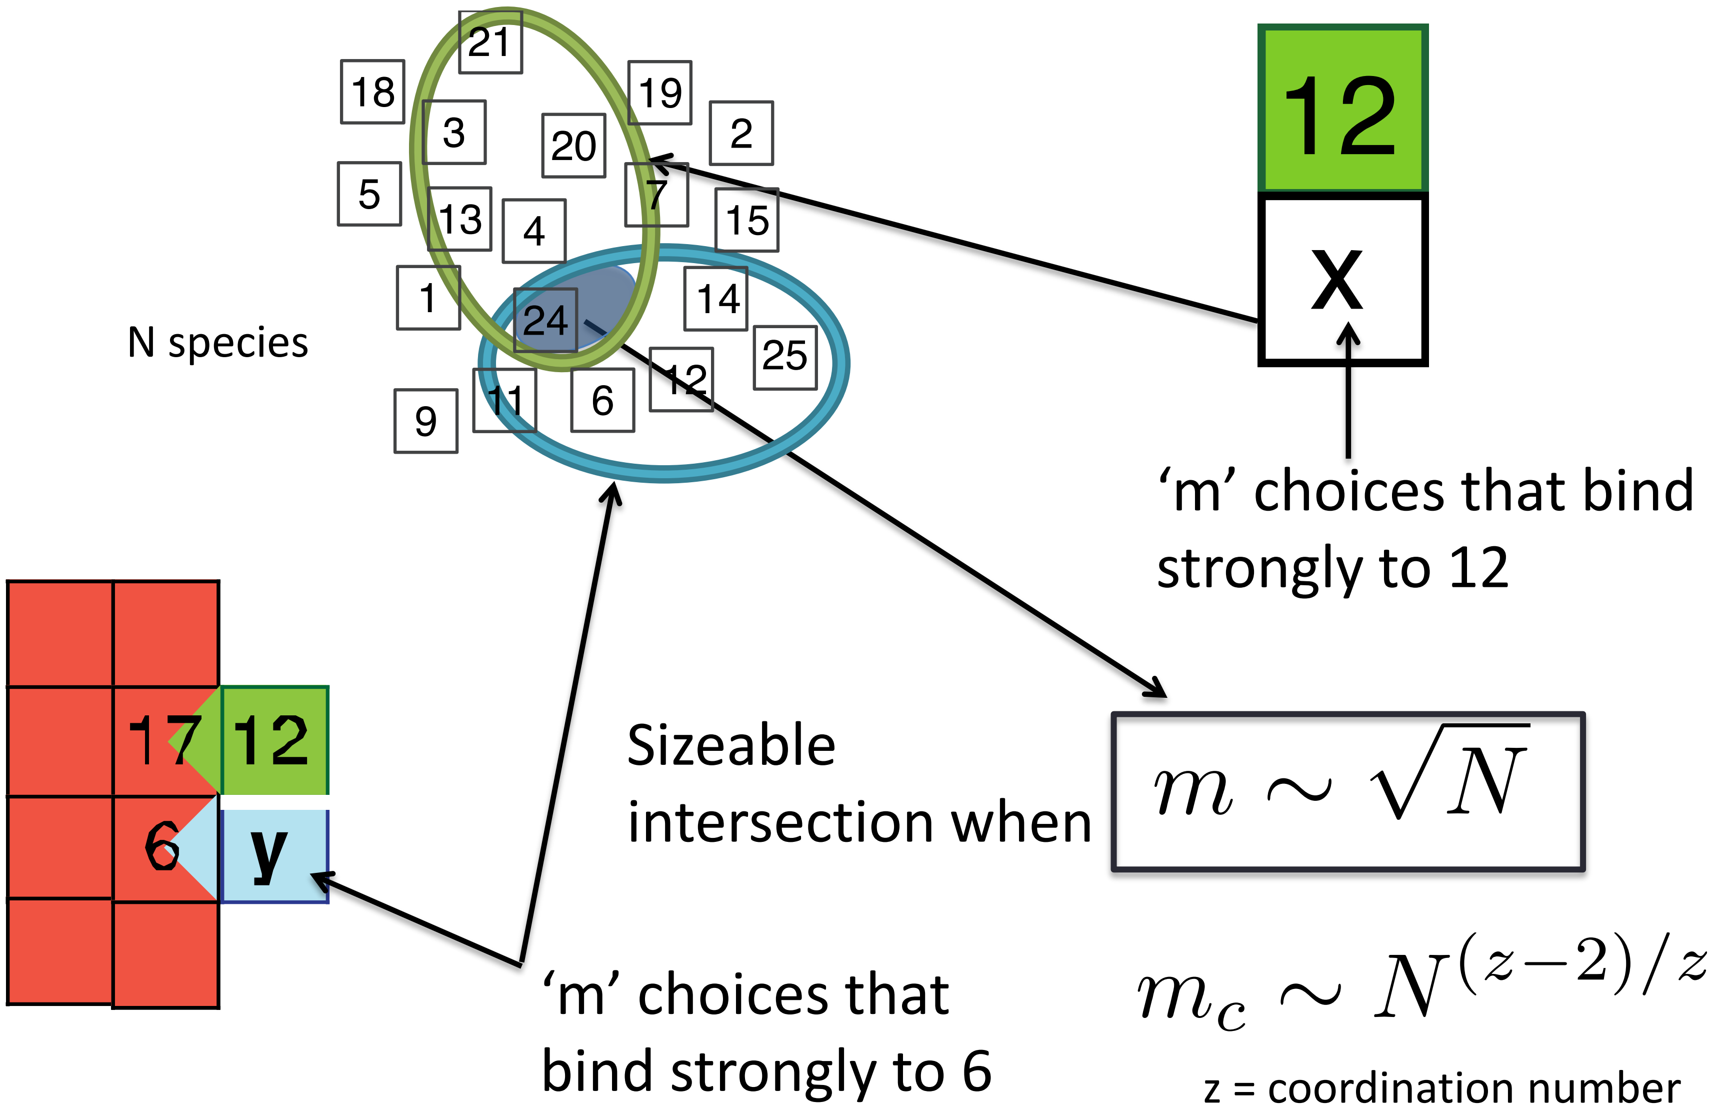

Phases

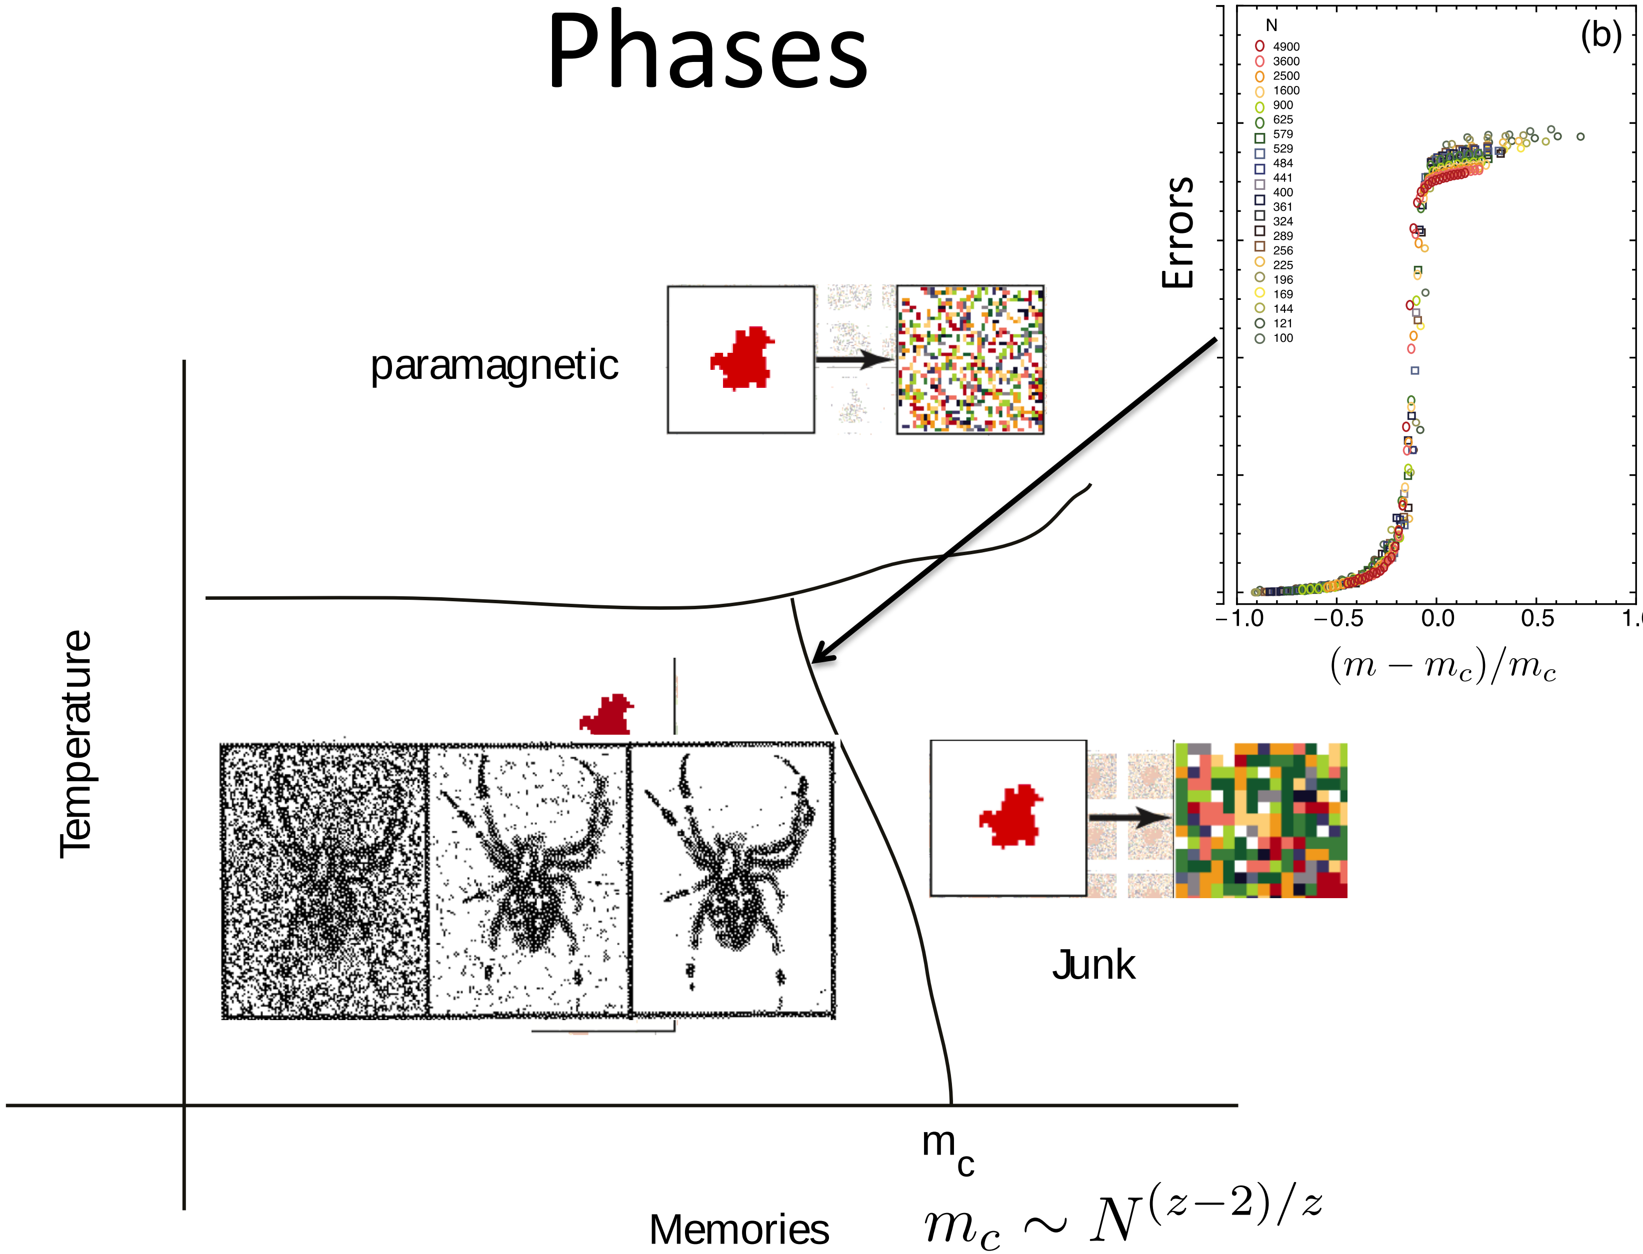

paramagnetic
Temperature
Junk
Memories
m_c
$m_c \sim N^{(z-2)/z}$
Errors
(b)
N
4900
3600
2500
1600
900
625
579
529
484
441
400
361
324
289
256
225
196
169
144
121
100
−1.0
−0.5
0.0
0.5
1.0
$(m - m_c)/m_c$

Pattern recognizer

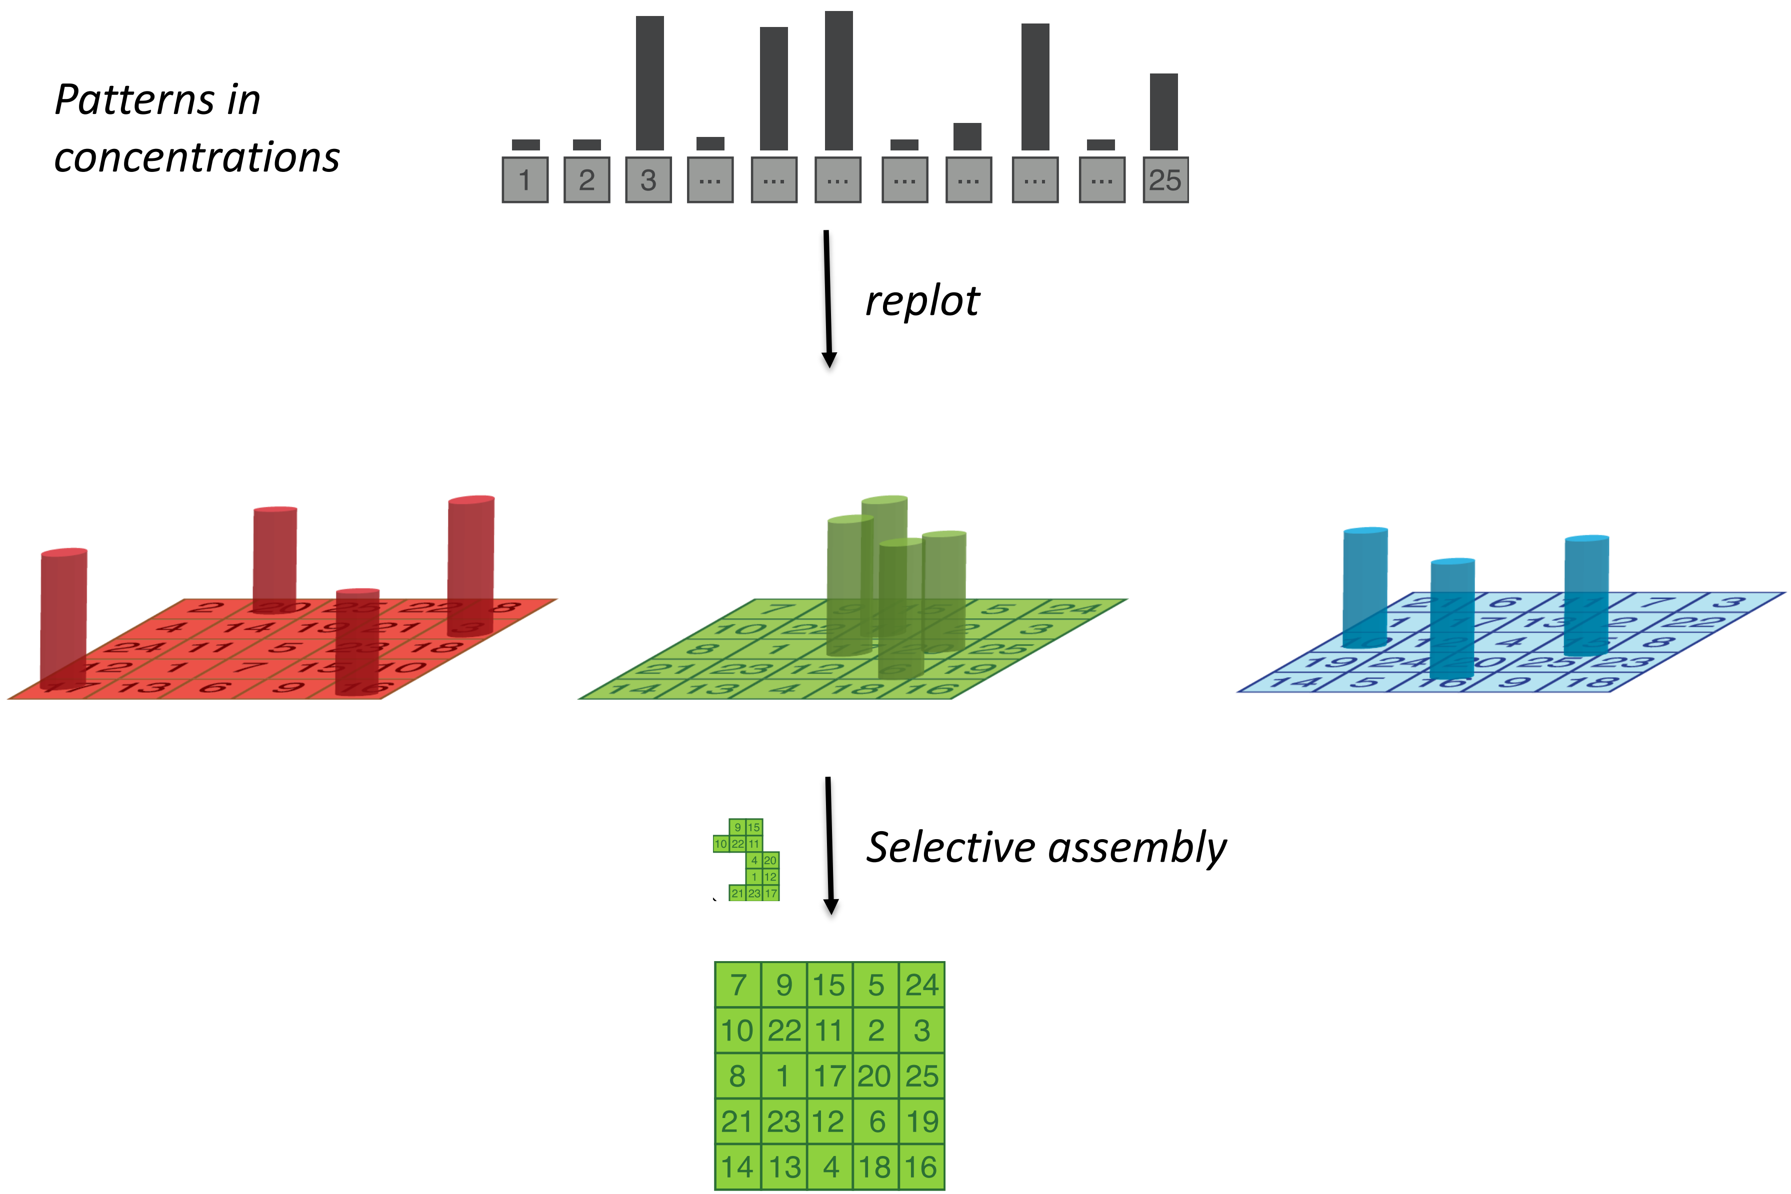

Pattern recognizer

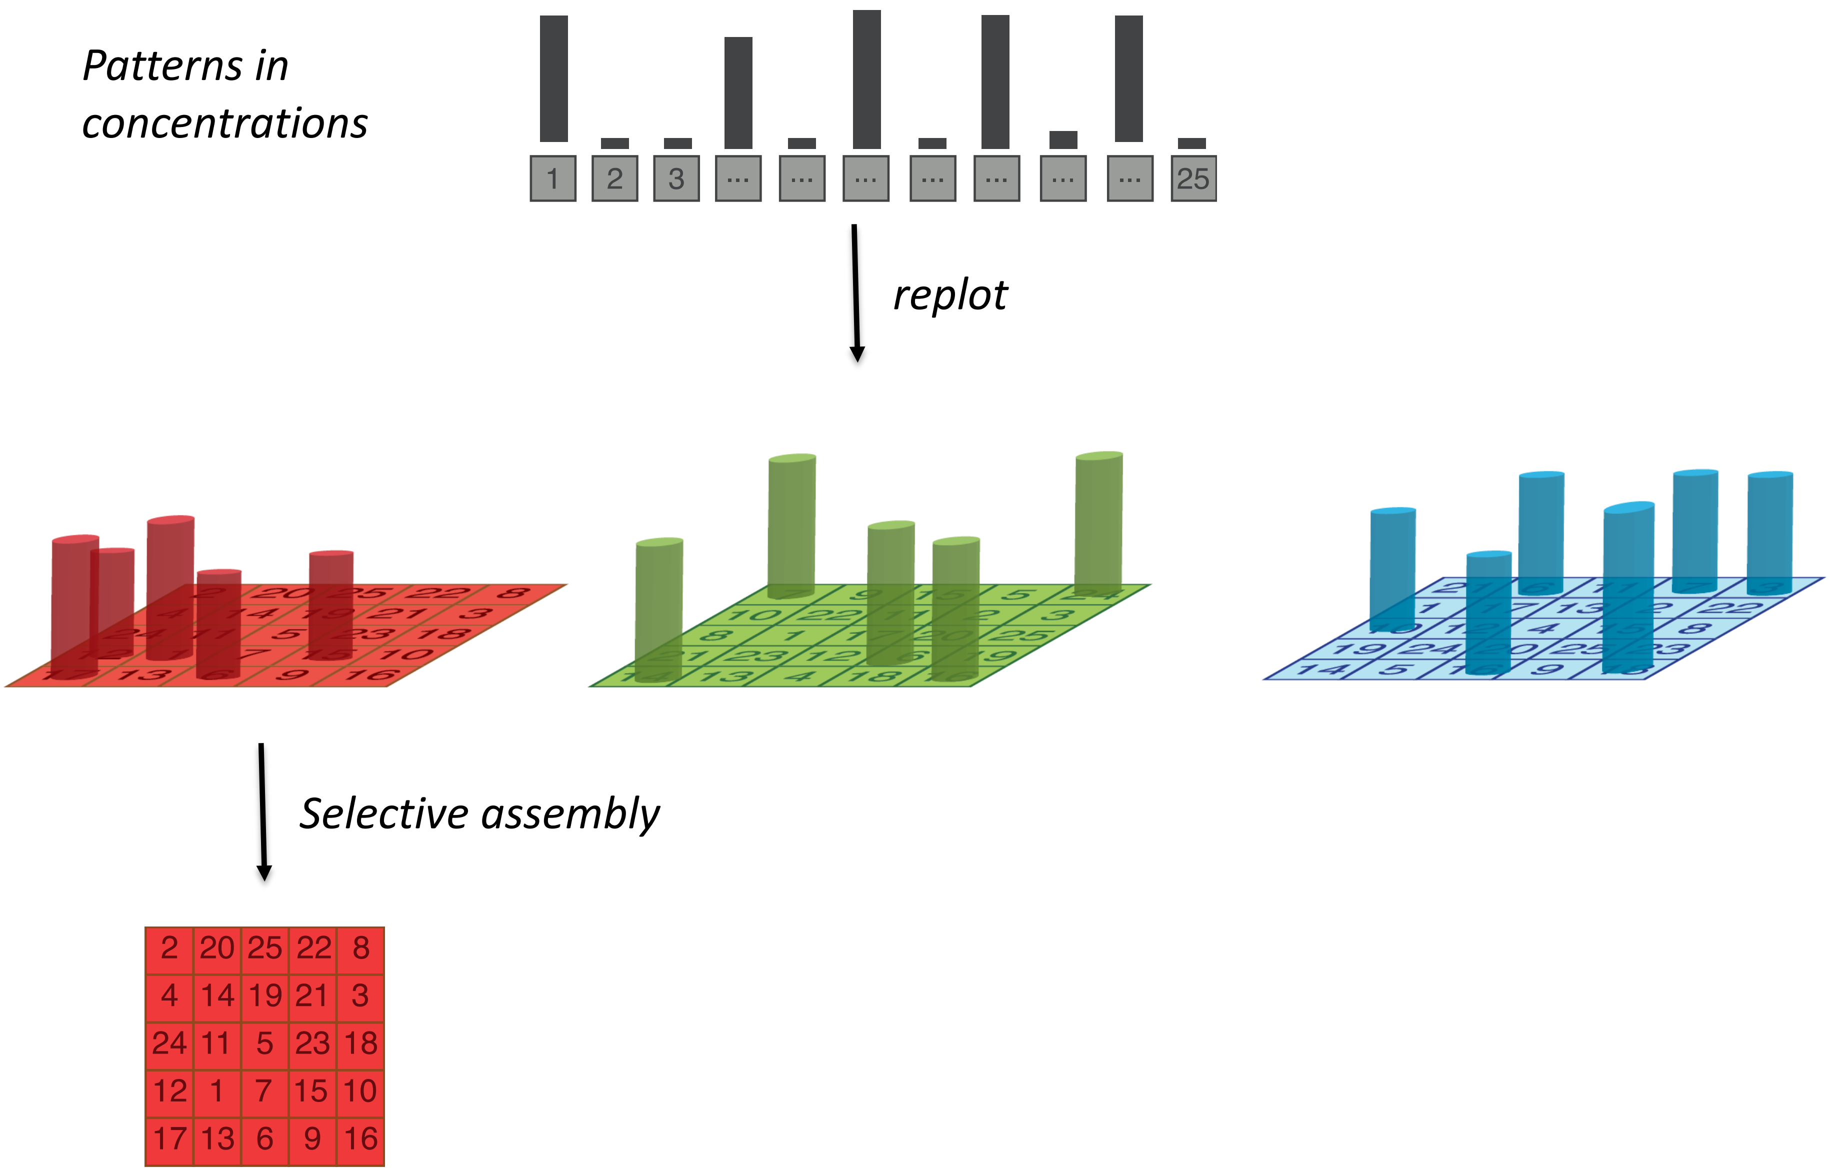

Pattern recognizer

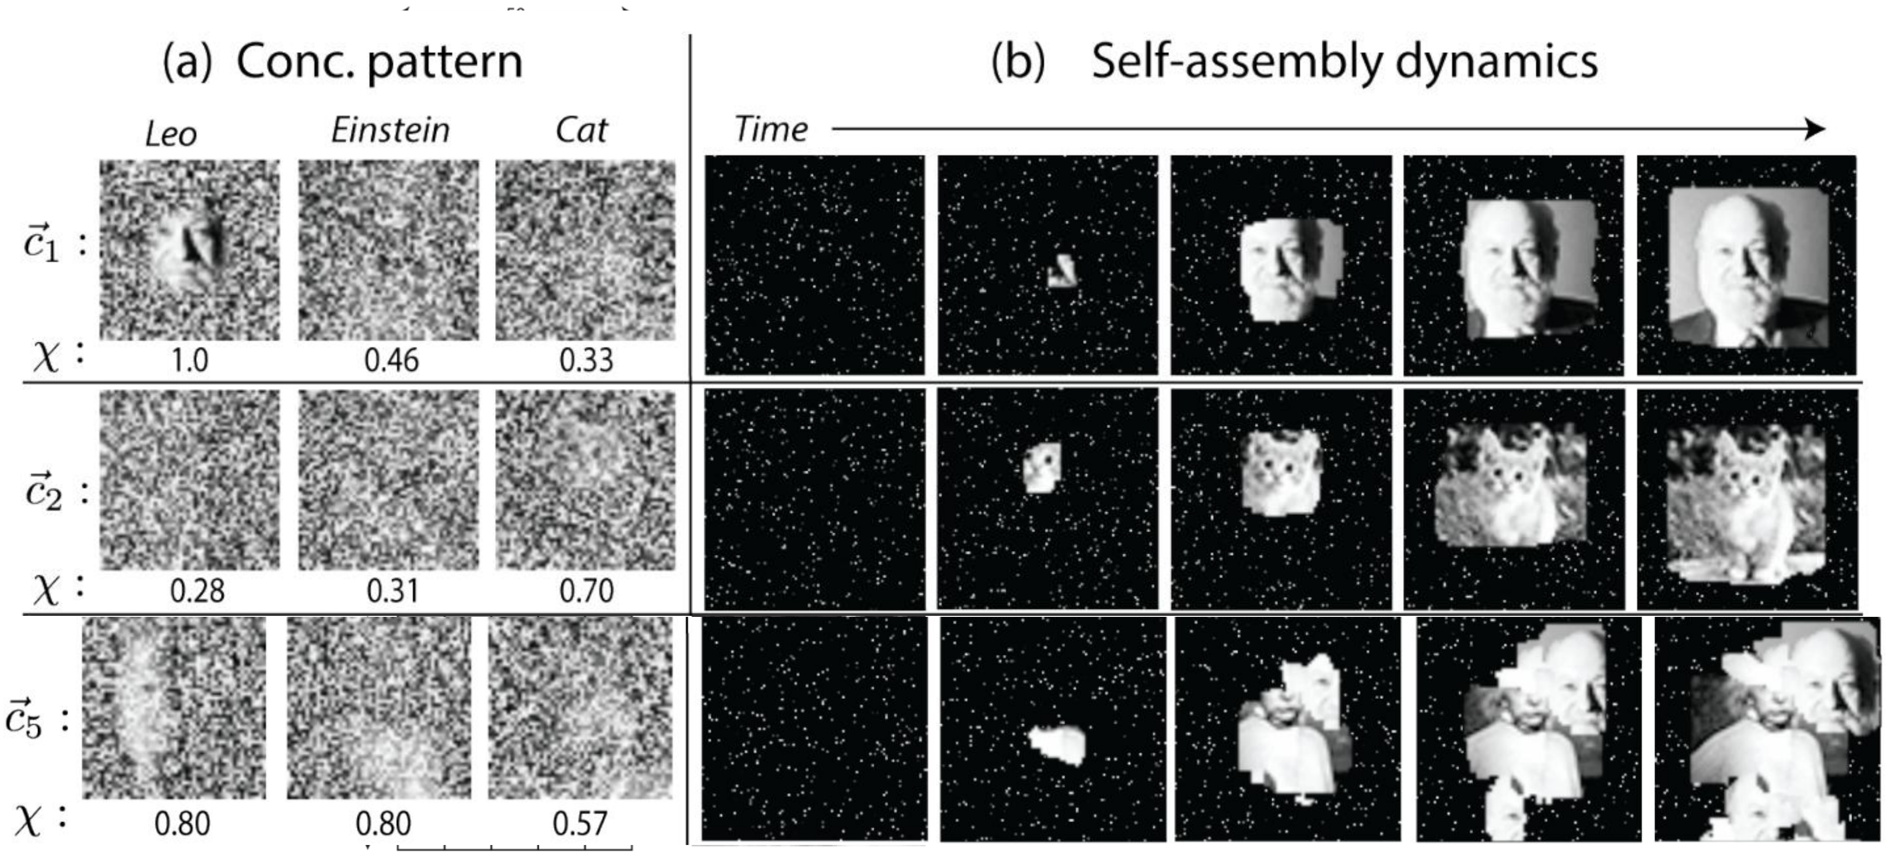

Associative memory

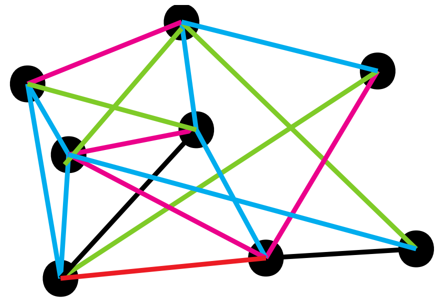

$$J_{ij} = J_{ij} + J_{ij} + J_{ij}$$

Neural networks

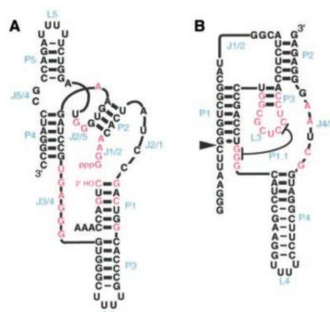

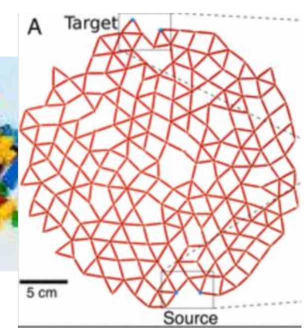

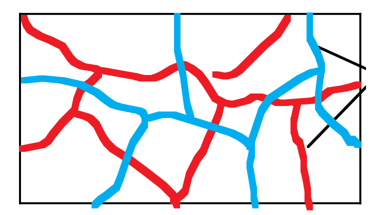

Self-folding polymers
(Ribozymes,
DNA origami)
(Schultes et al 2000)

Self-assembling particles

Mechanical networks
(metamaterials)
(Rocks et al 2017)

Self-folding sheets
(Origami)
(Stern et al 2017)

Self-folding polymers

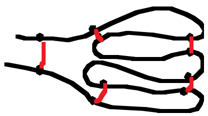

Contacts in desired structure

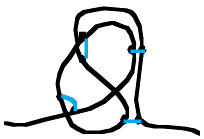

Design
(find sequence)

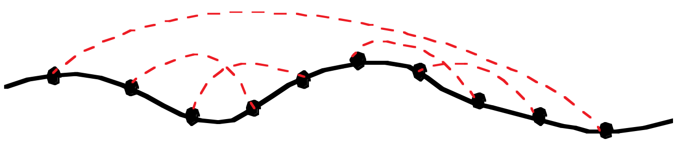

Programmed
interactions

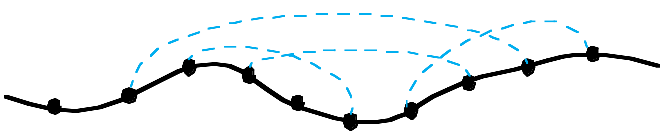

Seq. A

Seq. B

Fold

Fold

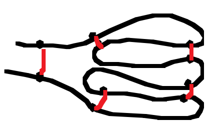

Examples:

- DNA origami (dots = stapling region)
- RNA secondary structure
  (dots = stem regions)

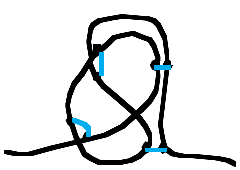

Associative memory in polymer folding

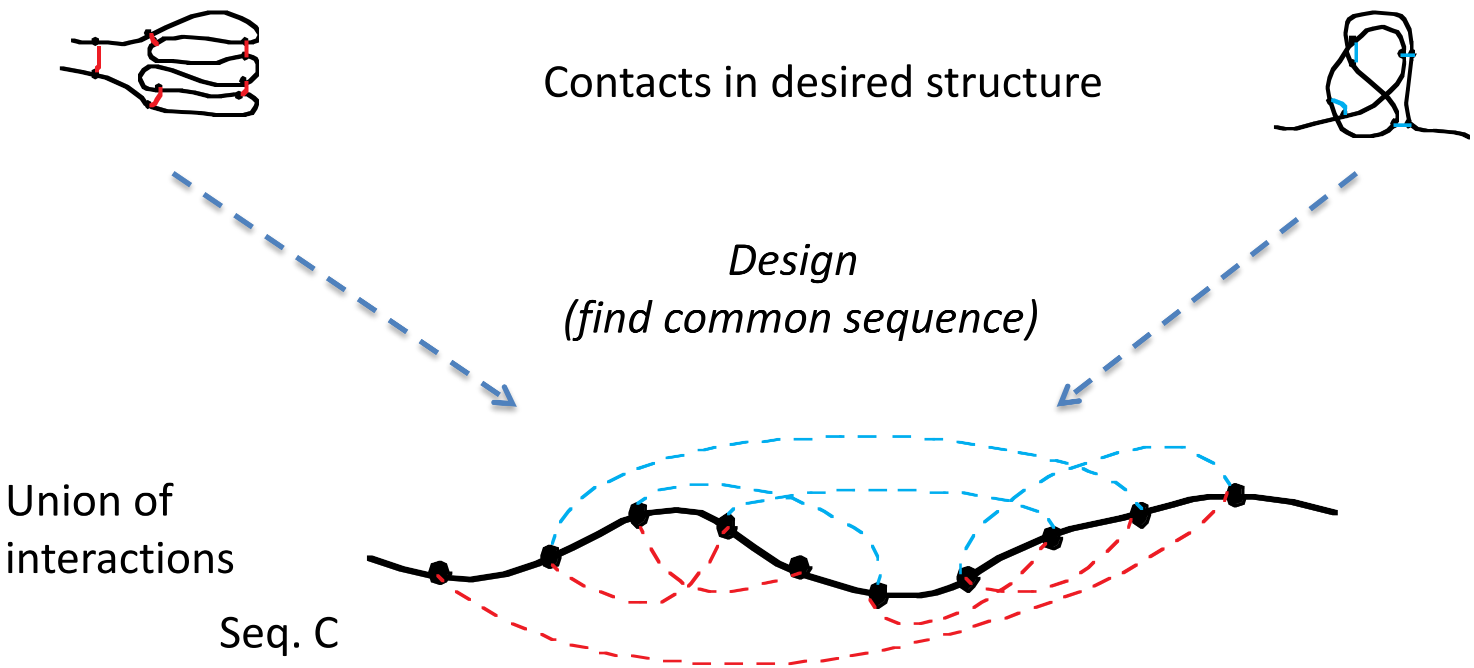

Promiscuous polymers

In how many ways can promiscuous polymers fold?

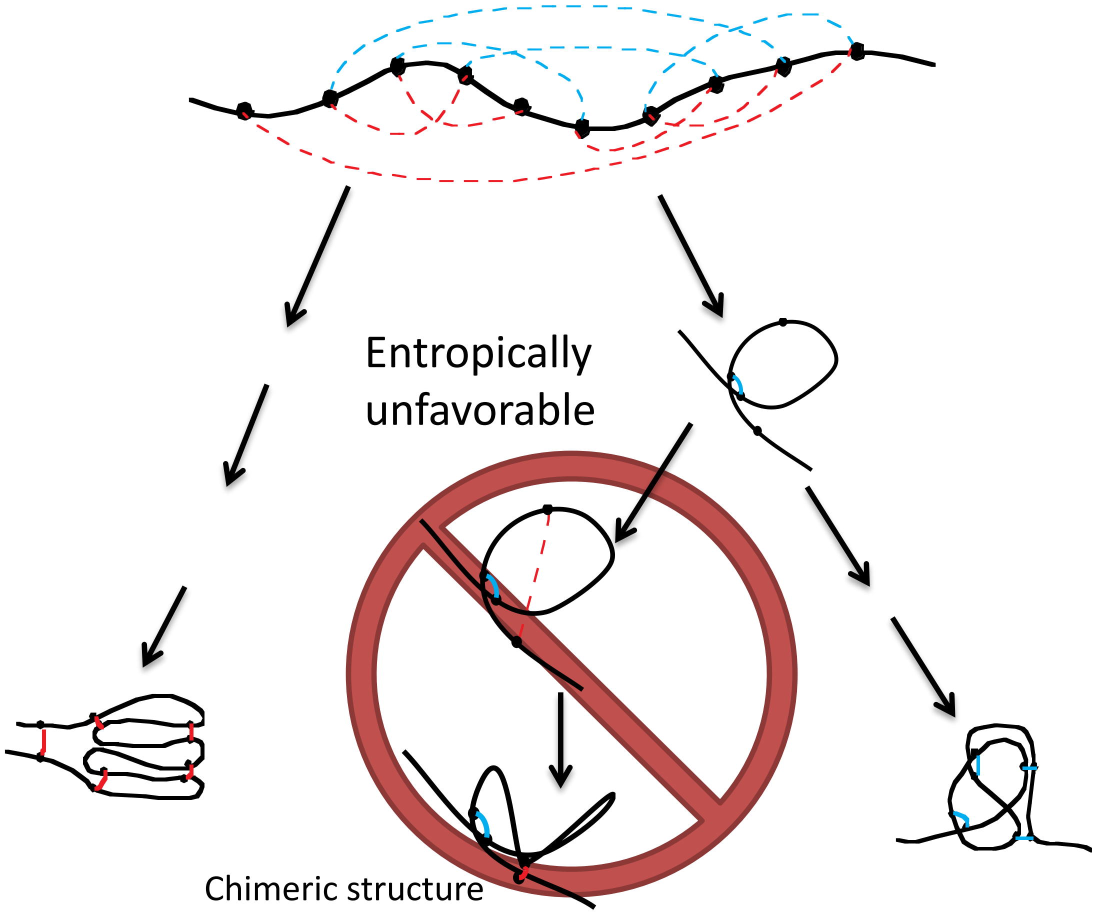

Specific kinetic simulations:
Abkevich et al , JCP 1994
Isambert et al 2000s..

Equilibrium theory:
Ball, Fink PRL 2001

DNA Origami experiments:
Dunn et all, Nature 2015

One Sequence, Two Ribozymes: Implications for the Emergence of New Ribozyme Folds

Erik A. Schultes and David P. Bartel*

Science 2000

Useful evolutionary intermediate

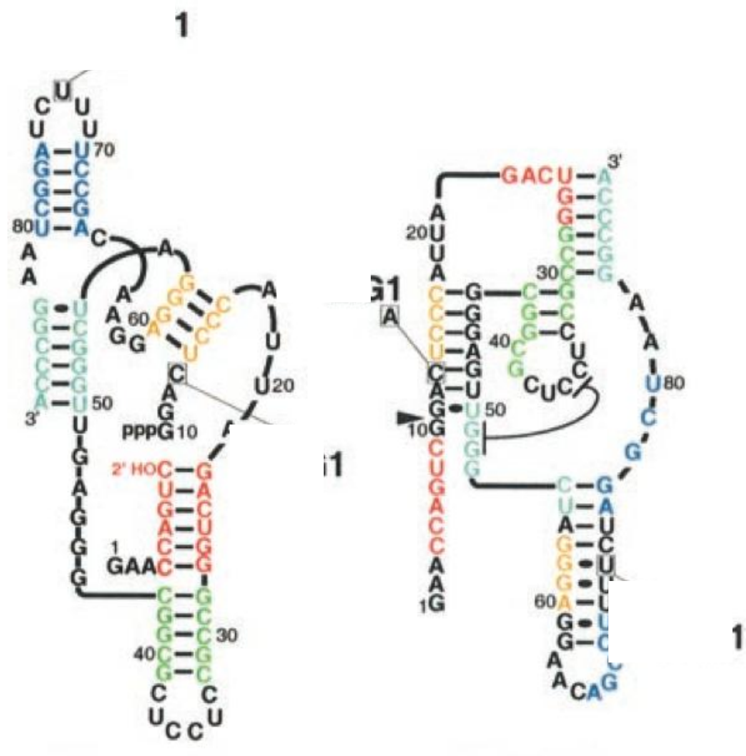

Ligase fold

Cleaving fold *(Hepatitis D Virus ribozyme)*

Self-folding sheets

Tomohiro Tachi

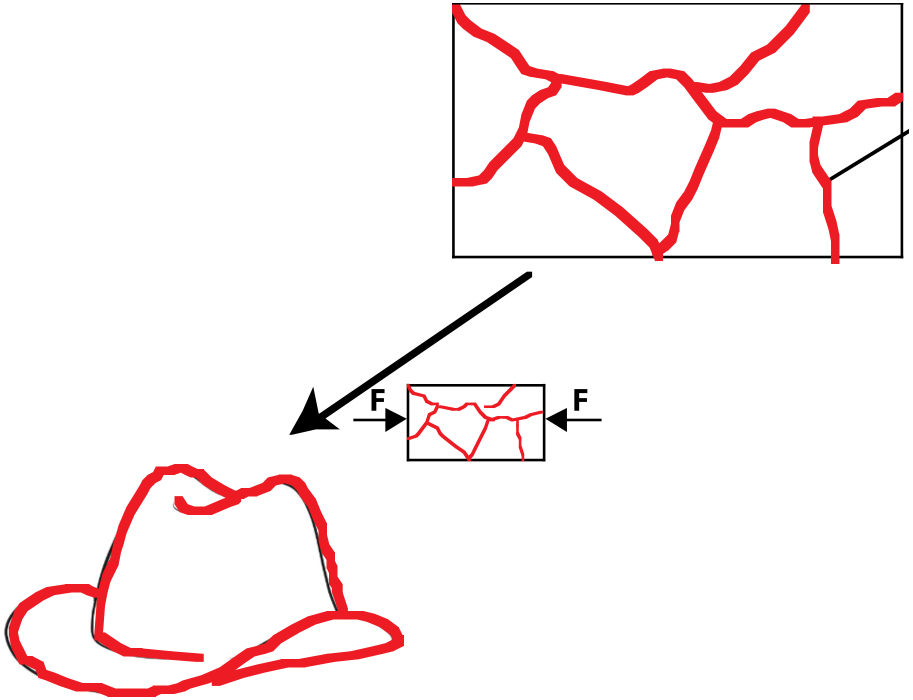

Multiple folding modes

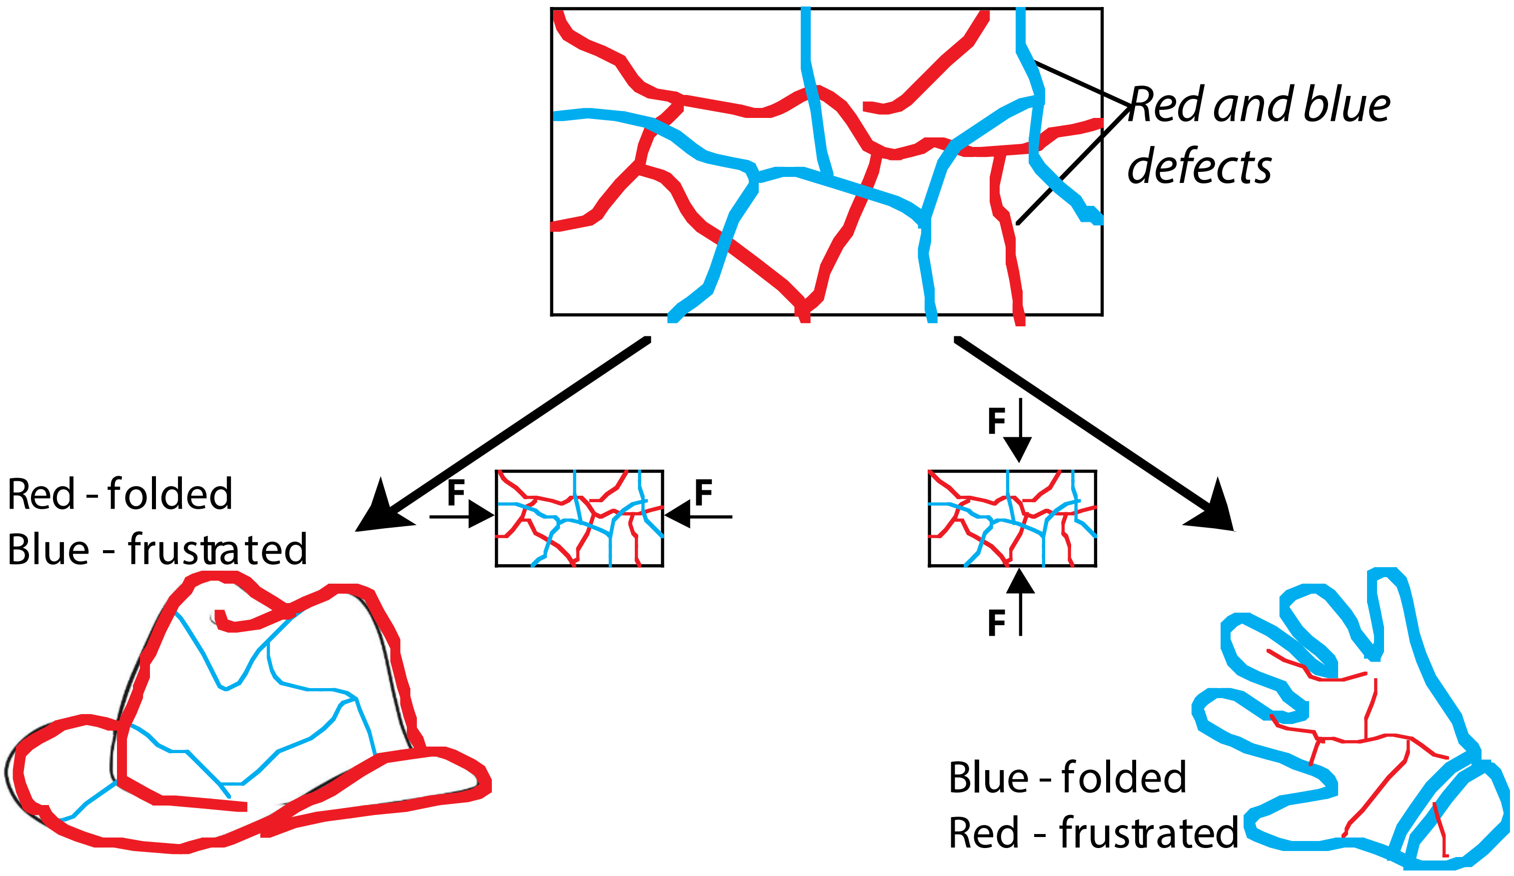

Multiple folding modes

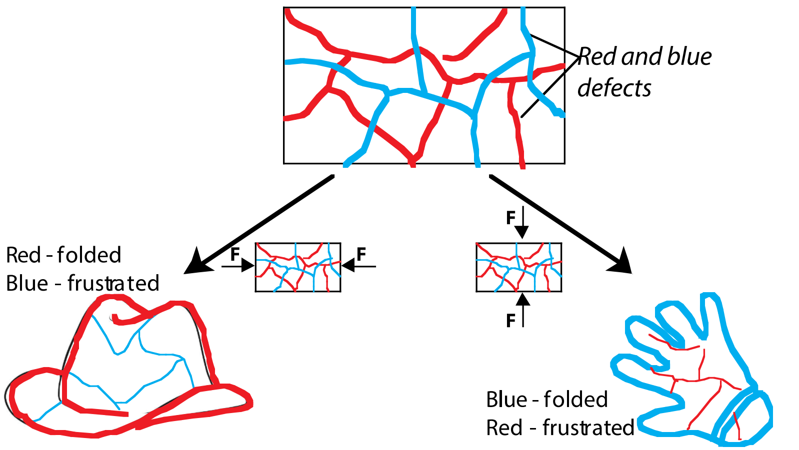

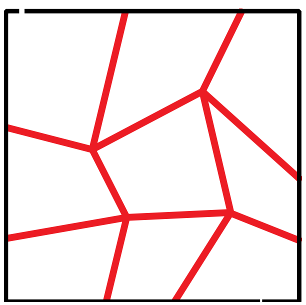

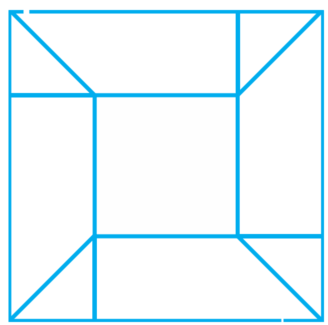

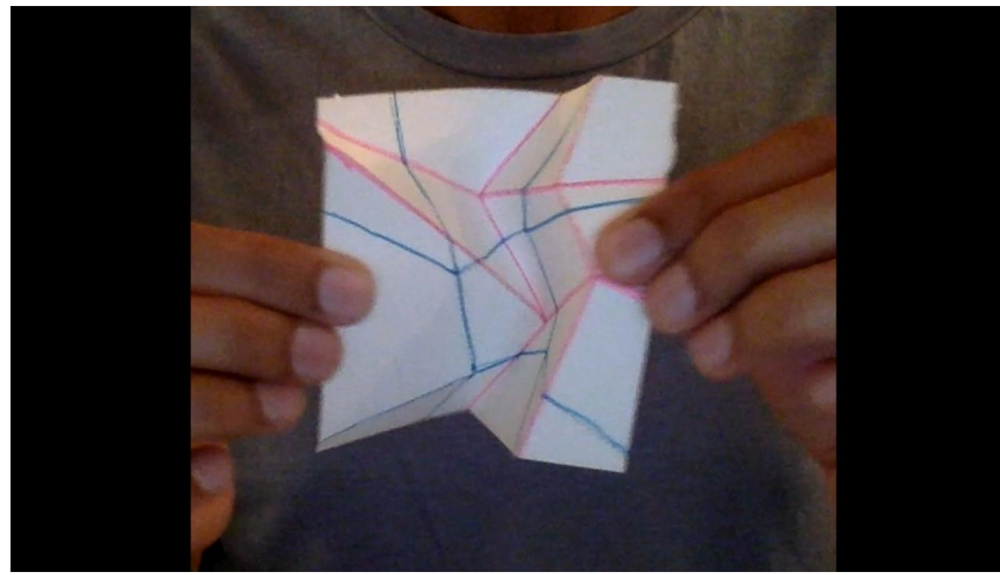

No need to micromanage

Frustrated loops prevent chimeras

State of a crease = Mountain, Valley or Flat

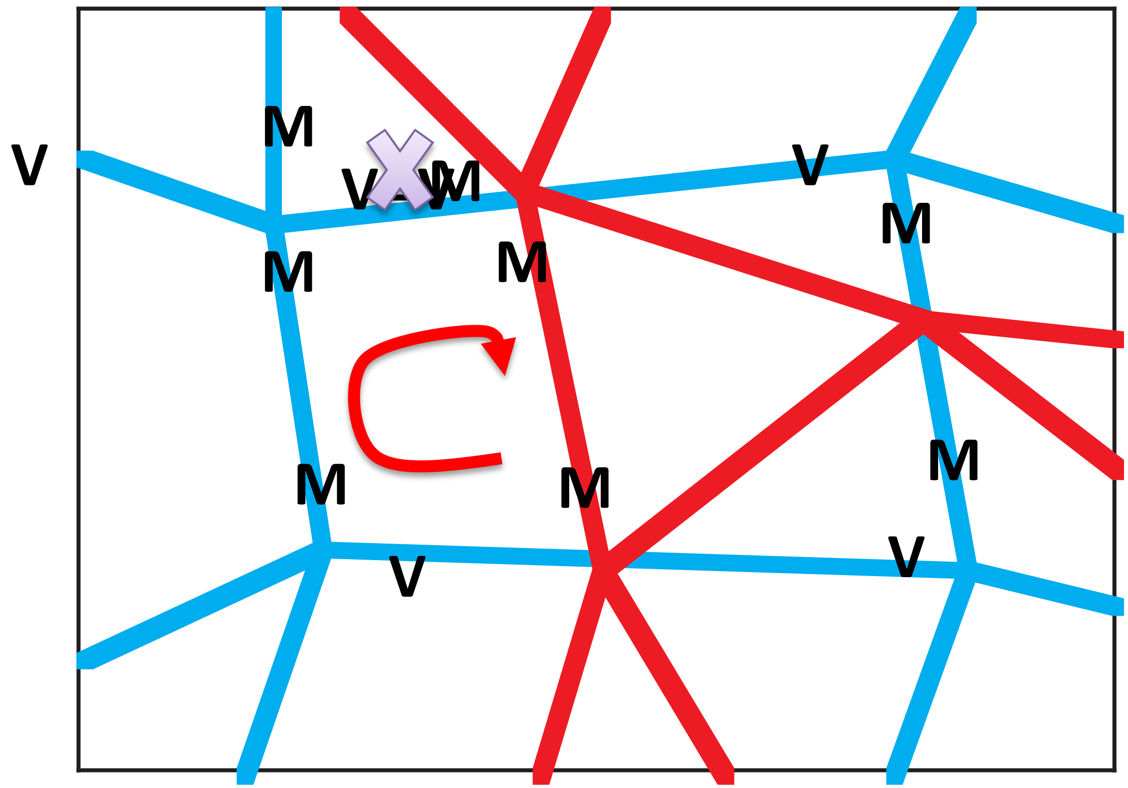

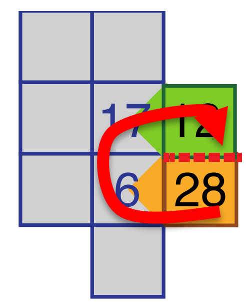

# of folding modes
= # of zero E ground states of disordered frustrated spin-1 system

$$E = \sum_{\text{vertices a}} J^a x_{a1} x_{a2} x_{a3} \ldots$$

Mechanical networks

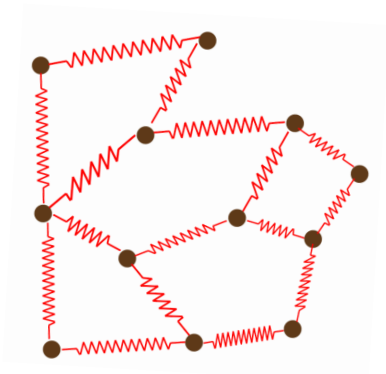

One memory

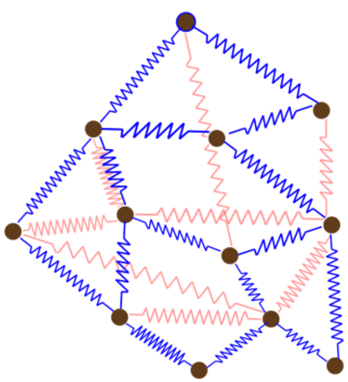

Two memories

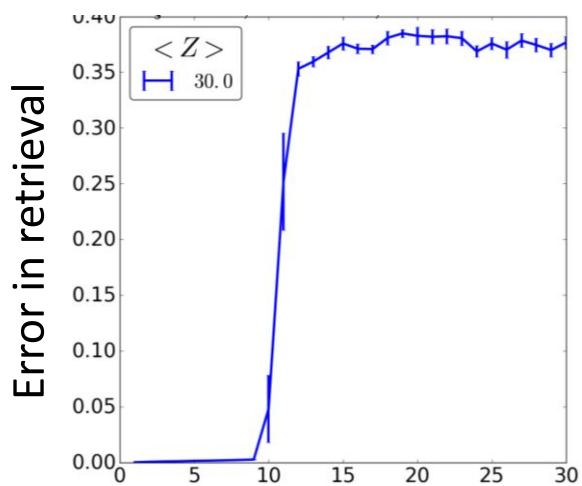

# of memories

Non-linearity of springs

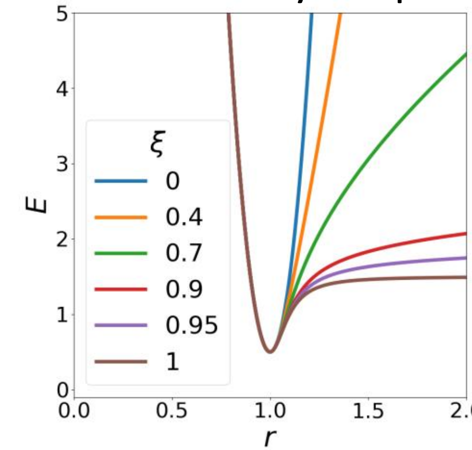

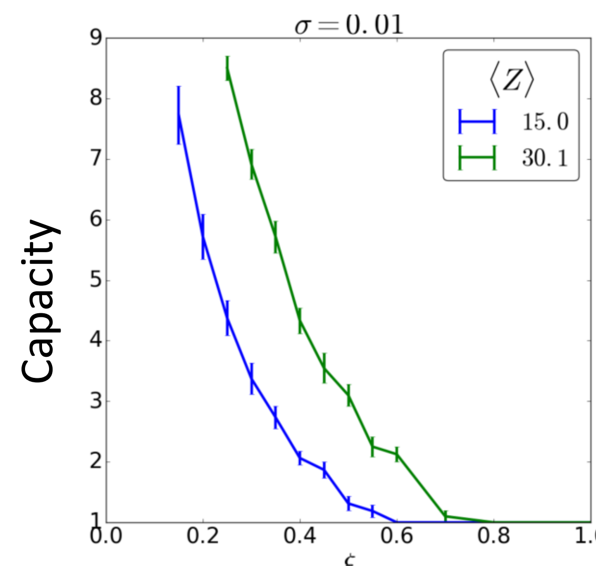

Non-linearity parameter

Sparsity through springs

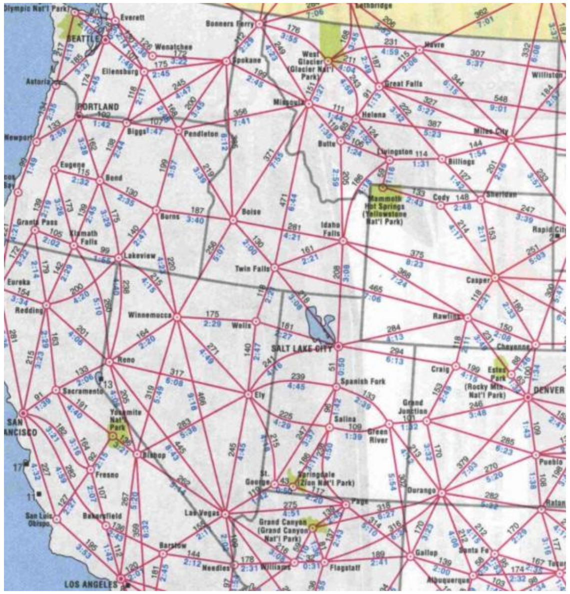

Given: Sufficient pairwise distances between N cities … Reconstruct geography.

Complication: A few distances are *wrong*

L2 minimization: Bad idea $E \sim x^2$

L1 minimization: Best idea $E \sim |x|$

L0 minimization: Better idea $E \sim x^0$

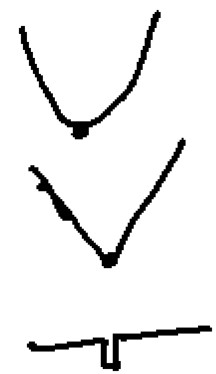

L0 **L1** **L2**

Sparsity emphasized

Non-convex   Convex .

Compressed sensing

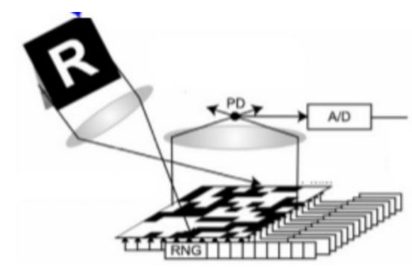

One pixel camera

Sparsity through springs

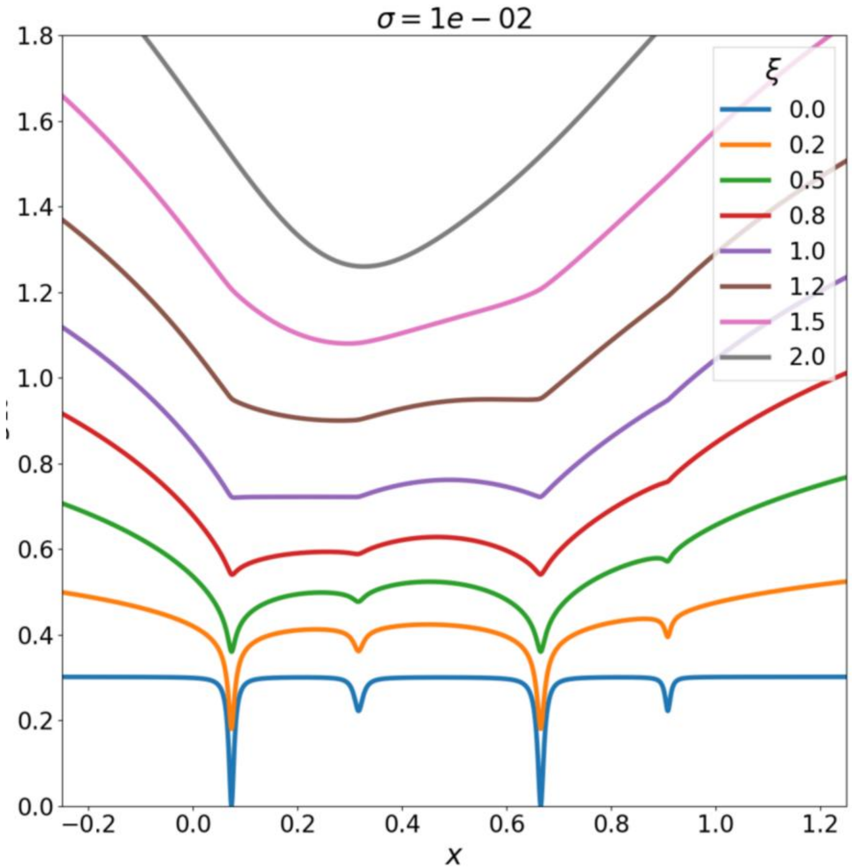

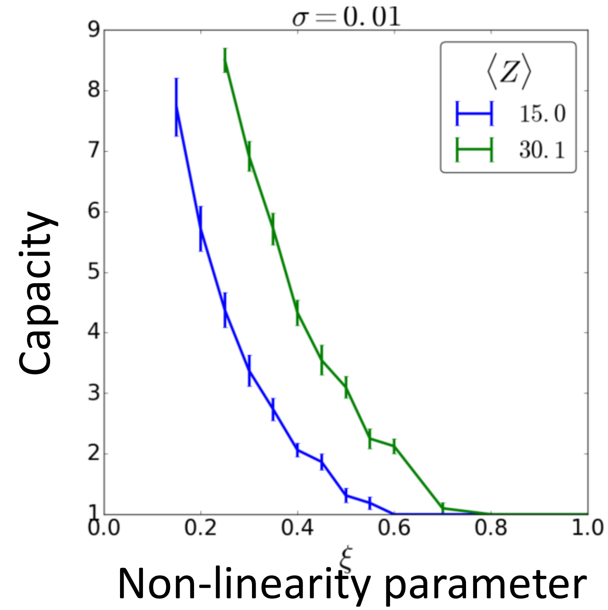

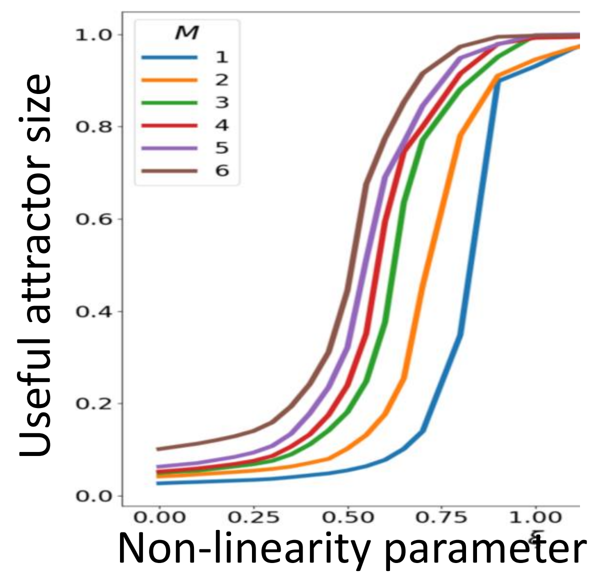

Learning vs memory

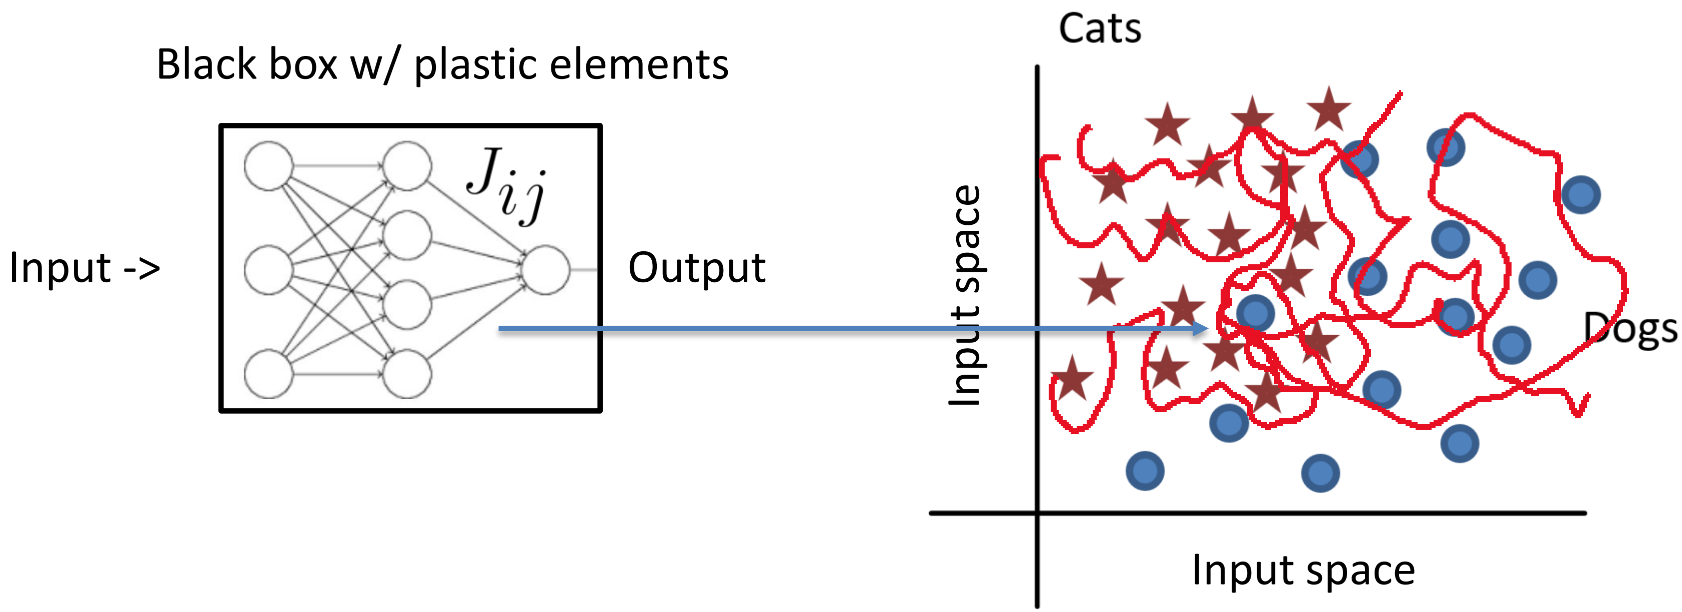

High plasticity

Training phase:
Show examples of inputs that should evoke output
Other inputs should not evoke output

Test phase:
Try other inputs that should evoke output.

Learning vs memory

Black box w/ plastic elements

Input -> 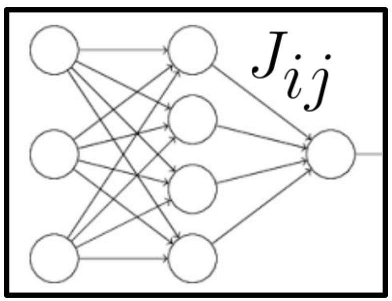 Output

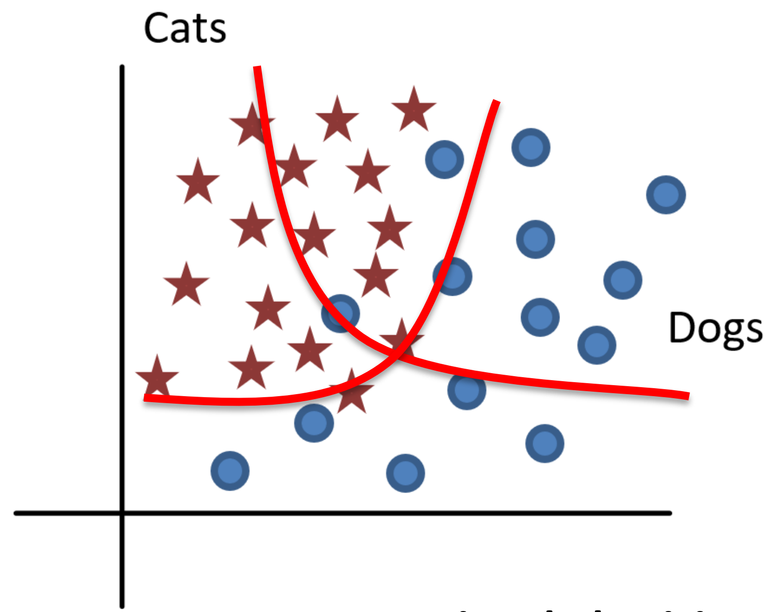

Restricted plasticity

Higher training error
Lower test error

Training phase:
Show examples of inputs that should evoke output
Other inputs should not evoke output

Test phase:
Try other inputs that should evoke output.

Learning vs memory

Vapnik–Chervonenkis (VC) dimension:

Size of largest set of inputs that can always be `shattered'.

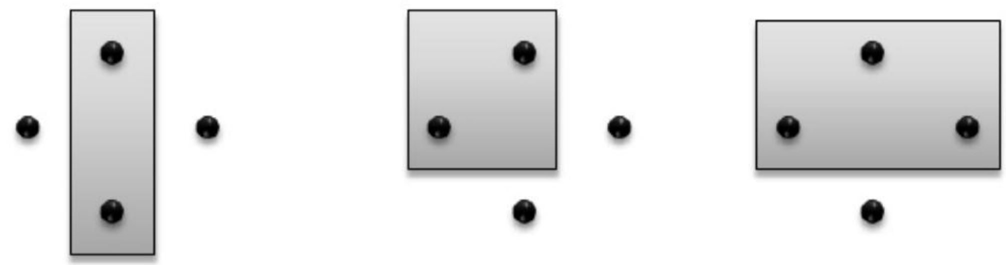

Rectangles can `shatter sets of four points..

Lines can shatter `any' set of three points.. but not sets of four points.

$$\Pr\left(\text{test error} \leqslant \text{training error} + \sqrt{\frac{1}{N}\left[D\left(\log\left(\tfrac{2N}{D}\right)+1\right)-\log\left(\tfrac{\eta}{4}\right)\right]}\right) = 1-\eta,$$

Conclusion:

Higher VC dim => low training error, high test error => more memorization/ less learning

Lower VC dim => high training error, low test error => less memorization / more learning

How to force generalization

Noise ('Dropout')

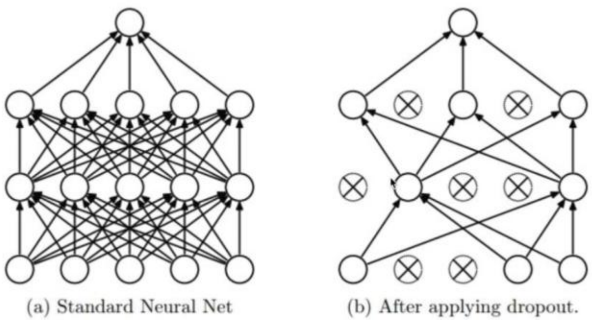

Full network   Random dropout

Randomly turn off (and on) plasticity in different parts during learning.

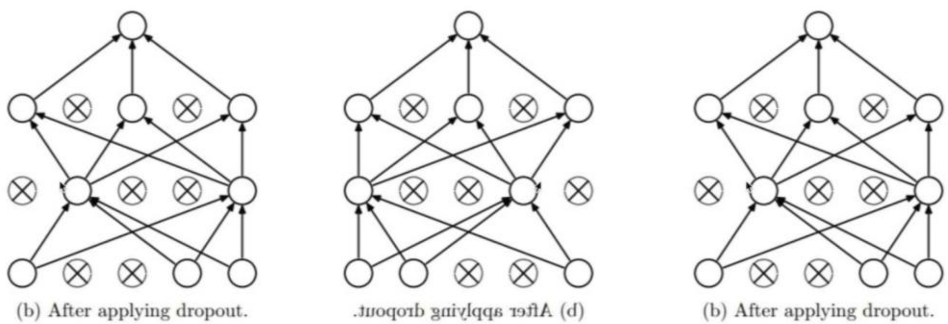

Time during training ->

How to force generalization

Switching environments

Seasonal variation of photo period

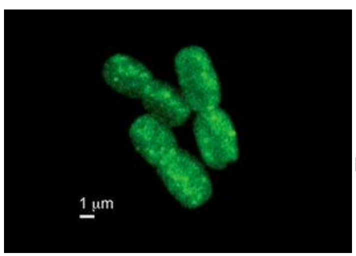

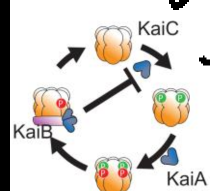

Predict dawn/dusk

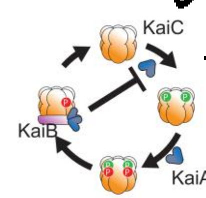

Predict dawn/dusk

S. Elongatus, Rust lab, eLife 2017

Small T
Rapid changes in day length

No fitness pressure
to predict

Intermediate T

Genotypic mem: concept of seasons

Phenotypic mem: day length

Large T
Slow changes in day length

Genotypic mem. of day length (inflexible, memorized)

How to force generalization

Switching environments

`Evolve’ antibody specific to mug

- But ignore handle
- All cups have handles

S. Wang et al, Cell 2015

Answer: Change mugs as a function of time

Time during training ->

VC dim of dynamical systems

Kyle Kawagoe
Ambre Bourdier

Different time series:

Series 1

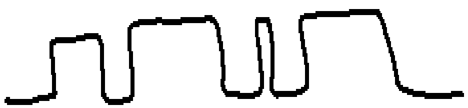

Series 2

Series 3

Can a dynamical system map these to different fixed points?

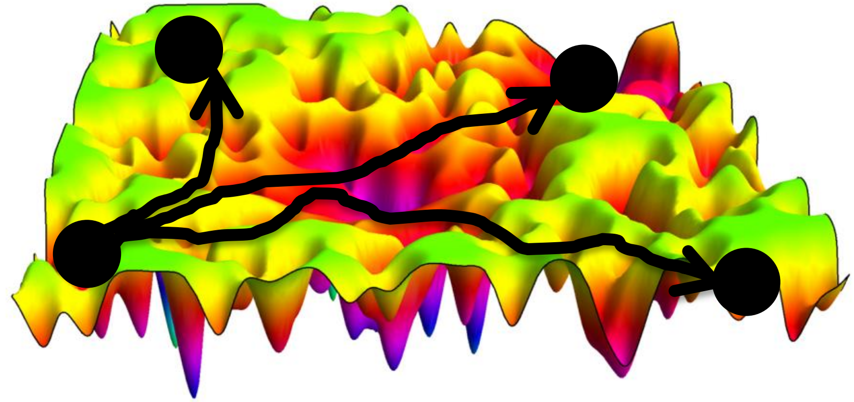

How large a set of time series can be `shattered' by a dynamical system?