Chaos in Neural Networks

What is chaos?

Chaos = sensitive dependence of initial conditions
i.e. instability of trajectory

How did we characterize instability before?

Fixed point: $\dot{x} = F(x)$, find $\dot{x} = 0 = F(x^*)$
then perturb $(x^* + \delta)$ and evolve:

$$\frac{d}{dt}(x^* + \delta) = F(x^* + \delta) \approx \cancel{F(x^*)}^{\,0} + \delta F'$$

$$\dot{\delta} = F'\delta$$

examine eigenvalues of F'

Now: see how fast two trajectories diverge when their initial conditions differ.

e.g. Dynamical system trajectory:

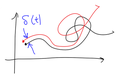

Follow same dynamics, $\dot{x} = F(x)$, but now not at fixed point

Linearize, $\dot{x} \approx Jx$ so $x(t) = x_0 e^{Jt}$

or $y(t) = (x_0 + \delta) e^{Jt}$

then $x - y = \delta(t) = \delta_0 e^{Jt}$

The largest eigenvalue of J is the maximal Lyapunov exponent, and sets the timescale on which trajectories diverge: $\tau \sim 1/\lambda_{max}$

3 dimensions are required for chaos in a continuous dynamical system. Why?

Examples:
- Double pendulum
- 3 orbiting bodies
- Lorenz attractor

Chaos in Random Neural Networks

Sompolinsky, Cristani, Sommers 1988

Tour de force!

They show that large random networks only have chaotic attractors

$\tau \dot{r} = -r + f(Jr)$ rate equation

$Jr = h$ voltage

$\tau \dot{h} = \tau J \dot{r} = J(-r + f(h))$

$\tau \dot{h} = -h + Jf(h)$ voltage equation

$$\tau \dot{h}_i = -h_i + \sum_{j=1}^{N} J_{ij} S_j \qquad S_j = \phi(h_j) \qquad \phi(x) = \tanh(gx)$$

h_j: membrane potential
S_j: firing rate
g: gain in nonlinearity

What is the behavior of this network when $N \to \infty$?

$\langle J_{ij} \rangle = 0$

$\langle J_{ij}^2 \rangle = \frac{J^2}{N}$

more neurons, smaller weights, and for this scaling these effects cancel:

$$\langle (\sum_j J_{ij} S_j)^2 \rangle = \langle \sum_j J_{ij} J_{ij} S_j S_j \rangle \quad \text{Variance of } S$$

$$\sum_j \langle J_{ij}^2 \rangle \langle S_j^2 \rangle \quad \langle S_j^2 \rangle = C(0)$$

$$N \frac{J^2}{N} + 0 = J^2 \text{ indep of } N$$

With many random inputs, this system is "self-averaging". This means that, instead of taking an average over realizations to find its properties, ONE network is representative and has the same properties as the avg.

This is a non-trivial result because fluctuations in a network can prevent self-averaging.

LETS UNPACK THIS:

Consider iid X_i for $i = 1..N$, and take an avg:

$\bar{X} = \frac{1}{N} \sum_i X_i$ $\qquad \sigma^2 = \frac{1}{N} \sum_i (X_i - \bar{X})^2$

Relative variance $\frac{\sigma^2}{\bar{X}^2} \xrightarrow[N \to \infty]{} 0$

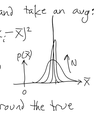

The empirical average concentrates around the true value as $N \to \infty$

This happens even when $\sum_i$ is not a sum over samples but a sum over neurons, if N is large enough.

What is the alternative?

Consider what happens if X_i are correlated.

$X_i = Y_i + Z$ where Z is the same for all i, but Y_i is independent.

Then $\frac{1}{N} \sum_i X_i = \frac{1}{N} \sum_i Y_i + \frac{1}{N} \sum_i Z$

$= \bar{Y} + Z$

Now $\bar{Y}$ will concentrate around the true value, since its an avg but Z will not concentrate: it's just the original number Z

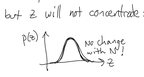

So the avg of X_i will not concentrate either.
This happens when the whole network synchronizes.

Okay, lets assume the network is self-averaging.

Now the input to a neuron is a random pooling of all other neurons, and this input is gaussian:

$$\tau \dot{h}_i = -h_i + \sum_{j=1}^{N} J_{ij} S_j$$

$$= -h_i + \eta_i \quad \text{where } \eta_i \sim N(0, \sigma^2)$$

$$h_i(t) = \int_{-\infty}^{t} e^{t'-t} \eta_i(t') dt' \quad \text{(lowpass filtered)}$$

This noise is independent between neurons, but correlated over time:

$$\langle \eta_i(t) \eta_i(t+\tau) \rangle_t = J^2 C(\tau) \quad \text{autocorrelation}$$

$$C(\tau) = \langle S_i(t) S_i(t+\tau) \rangle$$

$$= \langle \phi(h_i(t)) \phi(h_i(t+\tau)) \rangle$$

This is a self-consistent equation for $C(\tau)$ because $C \to \eta \to h \to S \to C$

They find it is "easier" to solve for autocorrelation of h rather than η:

$$\Delta(\tau) = \langle h(t) h(t+\tau) \rangle_t$$

$$\dot{\Delta} = \langle h(t) \dot{h}(t+\tau) \rangle$$

$$= \langle h(t) (-h(t+\tau) + \eta(t+\tau)) \rangle$$

$$= -\langle h(t) h(t+\tau) \rangle + \langle h(t) \eta(t+\tau) \rangle$$

$$= -\langle h(t) h(t+\tau) \rangle + \langle h(t-\tau) \eta(t) \rangle \quad \text{change variables } t' = t - \tau$$

$$\ddot{\Delta} = -\langle h(t) \dot{h}(t+\tau) \rangle - \langle \dot{h}(t-\tau) \eta(t) \rangle_t$$

$$= -\langle h(t) (-h(t+\tau) + \eta) \rangle - \langle (-h(t-\tau) + \eta(t-\tau)) \eta(t) \rangle_t$$

$$= \langle h(t) h(t+\tau) \rangle - \cancel{\langle h(t) \eta(t+\tau) \rangle} + \cancel{\langle h(t) \eta(t+\tau) \rangle} - \langle \eta(t) \eta(t+\tau) \rangle_t$$

$$\ddot{\Delta} = \Delta - J^2 C(\tau)$$

Cleverness: Use this as equation of motion for $\Delta(t)$ in a potential well $V(\Delta)$!

$$\ddot{\Delta} = -\frac{\partial V}{\partial \Delta} \qquad V = -\frac{1}{2}\Delta^2 + \mathcal{V}(|\Delta|, \Delta(0))$$

Now we can analyze the potential function and learn something without even solving for $\Delta(t)$.

Recall g was the gain parameter

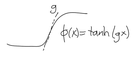

For $gJ < 1$, the potential has the form

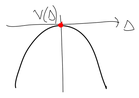

Only one solution, $\Delta(t) = 0 \ \forall t$

FIXED POINT

For $gJ > 1$,

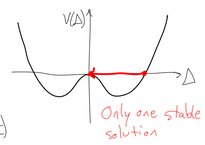

Thus, voltage inputs have a nonzero variance, but at long times are uncorrelated!

Notice the similarity with an Ornstein-Uhlenbeck process — but this is deterministic!

(Contrast $\Delta(t)$ with fixed pts or limit cycles

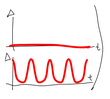

)

This implies that neural networks are chaotic!

Sompolinsky et al also derive Lyapunov exponents, but that is "more complicated".

For finite N, there is a transition to chaos that involves limit cycles of increasing complexity

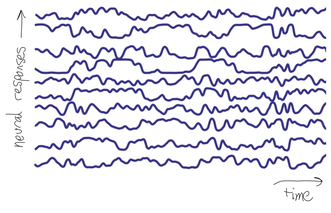

One way to understand the operation of this network is by the eigenvalue spectrum of random matrices

$\tau \dot{r} = -r + Jr$

$J_{ij} \sim N(0, \sigma_J^2 = \frac{1}{N})$

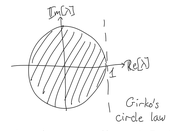

Girko's circle law

So when $gJ > 1$, some eigenvalues cross the $Re[\lambda] = 1$ line and the decay term no longer wins: the corresponding modes (eigenvectors) become unstable

How can a chaotic network COMPUTE anything?

1. Multiple strange attractors can coexist and perhaps encode different things

2. chaos can be suppressed by input

Phase diagram

Chaos can also be suppressed by feedback

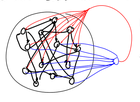

trained feedback acts like external input but is driven by internal activity

"Balanced" regime

This discussion used positive and negative responses for self-averaging. But neurons don't have negative rates!

Average number of inputs $\sim N$

$\pm$ fluctuations $\sim \sqrt{N}$ ← informative part, but swamped by N

Avoid this problem by cancelling the mean:

$$\langle \sum_j w_{ij} r_j \rangle \approx 0$$

$$\sum_j w_{ij} \langle r_j \rangle \approx 0 \xrightarrow[\text{self averaging}]{} \bar{r} \sum_j w_{ij} \approx 0 \Rightarrow \sum_j w_{ij} = 0$$

BALANCED

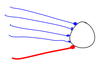

This means that neural firing will be driven by fluctuations about the mean:

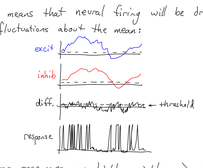

These responses are highly variable, and produce random-seeming fluctuations, even though the system is deterministic.

How is this balance created? If fine-tuning were required, that would be a problem.

Luckily, the balanced state is an emergent property that does not require fine tuning. (in some regimes)

CONCLUSION:

Variable cortical responses arise from fluctuations in synaptic input, and may be caused by chaotic dynamics