Biology as Information Dynamics

John Baez

Biological Complexity: Can It Be Quantified?
Beyond Center
February 2, 2017

IT'S ALL RELATIVE — EVEN INFORMATION!

When you learn something, how much information do you gain?

It depends on what you believed before!

We can model hypotheses as probability distributions. When you update your prior hypothesis *p* to a new one *q*, how much information did you gain?

This much:

$$I(q, p) = \sum_{i=1}^{n} q_i \log \left(\frac{q_i}{p_i} \right)$$

This is the **information of *q* relative to *p***, also called the 'information gain' or 'Kullback–Leibler divergence'.

$$I(q, p) \geq 0 \qquad \text{and} \qquad I(q, p) = 0 \iff q = p$$

For example, suppose we flip a coin you think is fair. Your prior hypothesis is this:

$$p_H = \frac{1}{2} \qquad p_T = \frac{1}{2}$$

Then you learn it landed heads up:

$$q_H = 1 \qquad q_T = 0$$

The relative information is 1 bit:

$$I(q, p) = 1 \log \left(\frac{1}{1/2} \right) + 0 \log \left(\frac{0}{1/2} \right) = \log 2$$

where we define $0 \log 0 = 0$. You have gained 1 bit of information.

But suppose you think there's only a 25% chance of heads:

$$p_H = \frac{1}{4} \qquad p_T = \frac{3}{4}$$

Then you learn the coin landed heads up:

$$q_H = 1 \qquad q_T = 0$$

Now the relative information is higher:

$$I(q, p) = 1 \log \left(\frac{1}{1/4} \right) + 0 \log \left(\frac{0}{3/4} \right) = \log 4 = 2 \log 2$$

You have gained 2 bits of information!

FREE ENERGY AS RELATIVE INFORMATION

The free energy of a system at temperature T is

$$F = E - TS$$

where E is the expected value of its energy and S is its entropy. In equilibrium this is minimized.

But in fact, the free energy of a system is proportional to *the information it contains, relative to equilibrium*.

'Equilibrium' is a hypothesis about the system's state, in which we guess that the probability of it being in its ith state is

$$p_i \propto \exp(-E_i/kT)$$

Here E_i is the energy of the ith state and k is Boltzmann's constant.

But suppose we choose a different hypothesis, and say probability of finding the system in its ith state is q_i.

A fun calculation shows that

$$F(q) - F(p) = kT\, I(q, p)$$

where $F(q)$ is the free energy for q and $F(p)$ is the free energy in equilibrium.

So: the extra free energy of a system out of equilibrium is proportional to *the information it has, relative to equilibrium*.

As a system approaches equilibrium, this information goes away.

NATURAL SELECTION AND RELATIVE INFORMATION

Suppose we have self-replicating entities of different kinds:

- ▶ molecules of different chemicals
- ▶ organisms belonging to different species
- ▶ genes of different alleles
- ▶ restaurants belonging to different chains
- ▶ people with different beliefs
- ▶ game-players with different strategies
- ▶ etc.

I'll call them **organisms** of different **species**.

Let P_i be the population of the ith species, as a function of time.

Suppose the **replicator equation** holds:

$$\frac{dP_i}{dt} = F_i\, P_i$$

where $F_i = F_i(P_1, \ldots, P_n)$, the **fitness** of the ith species, can depend on the populations of all the species.

The probability that a randomly chosen organism belongs to the *i*th species is

$$p_i = \frac{P_i}{\sum_j P_j}$$

If we think of each species as a *strategy*, the probability distribution p_i can be seen as a *hypothesis about the best strategy*. The quality of a strategy is measured by its fitness F_i. We can show

$$\frac{dp_i}{dt} = (F_i - \overline{F})p_i$$

where $\overline{F} = \sum_j F_j p_j$ is the **mean fitness**.

This is a continuous-time version of *Bayesian hypothesis updating*:

- Marc Harper, The replicator equation as an inference dynamic, arXiv:0911.1763.

Now suppose that q is some *fixed* probability distribution of species, while $p_i(t)$ changes as above.

A fun calculation shows that

$$\frac{d}{dt} I(q, p(t)) = -\sum_i \left(F_i - \overline{F} \right) q_i$$

But what does this equation *mean?*

$$\frac{d}{dt}I(q, p(t)) = -\sum_i \left(F_i - \overline{F}\right) q_i$$

$I(q, p(t))$ is the 'information left to learn' in going from $p(t)$ to q.

$\sum_i \left(F_i - \overline{F}\right) q_i$ is the average 'excess fitness' of a small 'invader' population with distribution q_i. It says how much fitter the invaders are, on average, than the general population $P_i(t)$.

If this average excess fitness is ≥ 0 for all choices of $P_i(t)$, we call q a **dominant** distribution.

So: if there is a dominant distribution q, the actual probability distribution of species $p(t)$ tends to approach this — and the information 'left to learn', $I(q, p(t))$, tends to decrease!
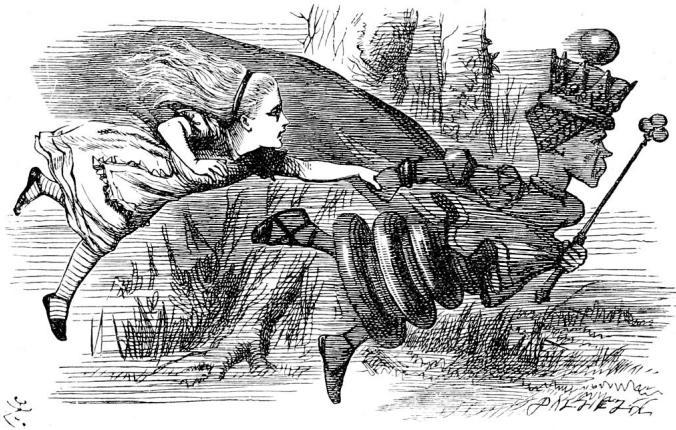
But real biology is much more interesting! At all levels, biological systems only approach steady states in limited regions for limited times.

Even using just the replicator equation

$$\frac{dP_i}{dt} = F_i P_i$$

we can already see more complex phenomena.

THE RED QUEEN HYPOTHESIS

"Now, here, you see, it takes all the running you can do, to keep in the same place."

The **Red Queen Hypothesis**: replicators must keep changing simply to survive amid other changing replicators.

For example, in males of the common side-blotched lizard, orange beats blue, blue beats yellow, and yellow beats orange:
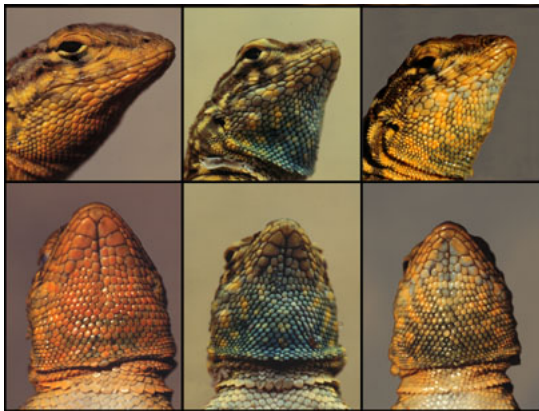
We can model this using the replicator equation by assuming lizards play randomly chosen opponents in the rock-paper-scissors game, with fitness determined by the game's outcome:
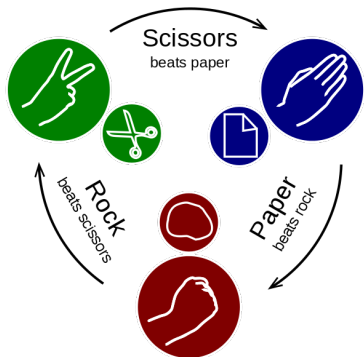
▶ Sinervo and Lively, The rock-paper-scissors game and the evolution of alternative male strategies, *Nature* **380** (1996).

The replicator equation gives this dynamics for the probability distribution of strategies 'rock', 'paper' and 'scissors':
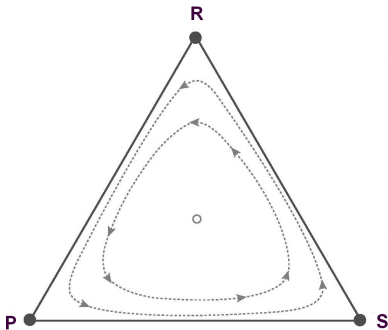
There is a steady state, but it is not an attractor. In general, *the population never stops learning new information!*

INFORMATION GEOMETRY

How can we quantify the *rate of learning?*

Here we face a 'paradox':

For any probability distribution $p(t)$ that changes smoothly with time, we have

$$\frac{d}{dt} I(p(t), p(t_0)) \bigg|_{t=t_0} = 0$$

for all times t_0.

"To first order, you're never learning anything."

However, as long as the velocity $\dot{p}(t_0)$ is nonzero, we have

$$\left.\frac{d^2}{dt^2} I(p(t), p(t_0))\right|_{t=t_0} > 0$$

"To second order, you're always learning something... unless your opinions are fixed."

This lets us define a 'rate of learning' — that is, the 'speed' of the changing probability distribution $p(t)$.

Namely, define the length of the vector $\dot{p}(t_0)$ by

$$\|\dot{p}(t_0)\|^2 = \frac{d^2}{dt^2} I(p(t), p(t_0)) \bigg|_{t=t_0}$$

This notion of length defines a 'Riemannian metric' on the space of probability distributions: the **Fisher information metric**. This makes the space of probability distributions round:
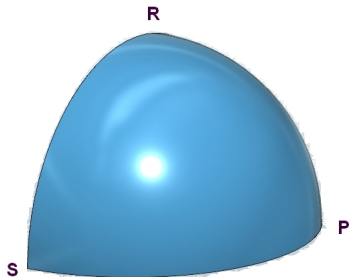
Now suppose we have populations obeying the replicator equation:

$$\frac{dP_i}{dt} = F_i P_i$$

so the probability distribution of species evolves via

$$\frac{dp_i}{dt} = (F_i - \overline{F})p_i$$

Then a fun calculation shows

$$\left\| \frac{dp}{dt} \right\|^2 = \sum_i (F_i - \overline{F})^2 p_i$$

The square of the rate of learning is the variance of the fitness!

This is a clean, general statement of **Fisher's fundamental theorem of natural selection**.

These are some tiny steps toward a theory of *biology as information dynamics*.