Creating **artificial systems** to study **natural systems**

CS369 // Computational Science

Matthew Egbert
m.egbert@auckland.ac.nz

Computational Models

Computational models are often designed to be comparable to a *target* real-world phenomenon.

Through an iterative process the model can be improved upon until it can predict what the real world system will do in various circumstances.

A trade off is common between being simple, realistic and predictive. *What is the minimal model that is maximally predictive?*

TARGET
Belousov-Zhabotinsky Chemical Oscillator

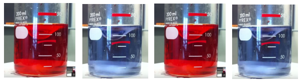

consists of ~18 chemical reactions

3-variable **MODEL**

I $A + Y \longrightarrow X + P$
II $X + Y \longrightarrow 2P$
III $A + X \longrightarrow 2X + 2Z$
IV $2X \longrightarrow A + P$
V $B + Z \longrightarrow \frac{1}{2} f Y$

$$\frac{d[X]}{dt} = k_I[A][Y] - k_{II}[X][Y] + k_{III}[A][X] - 2k_{IV}[X]^2$$
$$\frac{d[Y]}{dt} = -k_I[A][Y] - k_{II}[X][Y] + \frac{1}{2} f k_V[B][Z]$$
$$\frac{d[Z]}{dt} = 2k_{III}[A][X] - k_V[B][Z]$$

Numerical Integration

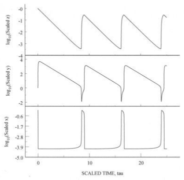

https://en.wikipedia.org/wiki/Oregonator
http://www.scholarpedia.org/article/Oregonator

Complex Systems

Many systems of interest are ***extremely* complex**, involving many components, involved in complex forms of feedback, etc.

And this leads to some serious challenges!

1. Prohibitive computational complexity
2. Free parameters -- *How do we assign values to unknowns?*
3. When models become too big, they become almost as hard to understand as their target!

Lots of work to be done here, and progress continues to be made.

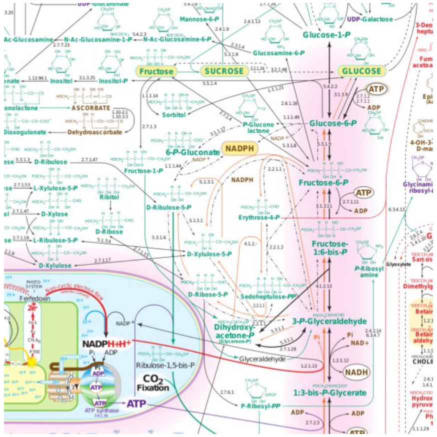

The targets of models

Sometimes the goal of a model is the **quantitative** comparison to or prediction of a *target* real-world system.

How will this bridge stand up to wind, while carrying various loads? etc.

But this is not the only way that models are used.

Qualitative recapitulation

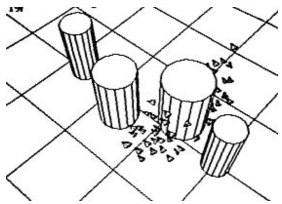

An alternative approach is to create an artificial system that captures (only) **qualitatively** some interesting property, and to study that.

Something that we can learn from, that might later help us better understand the real-world phenomenon.

We will see an example of a computational model that serves this role in a moment.

Modelling an idea

Sometimes, the target is an *idea* rather than a real-world phenomenon.

What are the implications of a theory or set of assumptions?

What would happen if you had a system with some configuration X?

We will also take a look at an example of this kind of computational model.

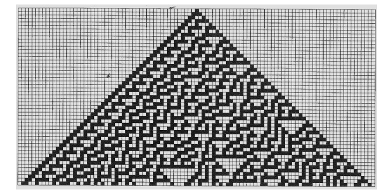

Artificial Life

I will now present a some classic examples of *Artificial Life* investigations that centre on computational models.

Part of my goal is to communicate the variety of ways that developing and studying artificial systems can be helpful in trying to understand the natural world...

thought experiment: what are the implications of a set of rules or assumptions?

sufficiency proof: X suffices to produce Y

tool for communication of ideas: e.g. by providing a concrete example of a concept

intuition pump: the model is used to facilitates an intuitive understanding of a system

creativity pump: the model is used to produce new theories or hypotheses

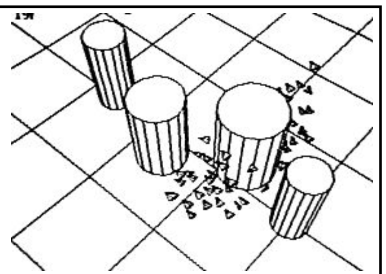

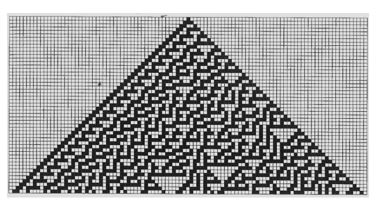

Cellular Automata

Wolfram's Minimal
1D Cellular Automata

A systematic investigation
In 1990s & 2000s, Stephen Wolfram (creator of Mathematica) carried out a systematic investigation of a mathematical abstraction called **cellular automata**.

The simplest systems he investigated had **one infinite, discrete spatial dimension**, where each cell is in one of two states

- "on" = black
- "off" = white

Time (discrete updates) can be shown on a vertical axis.

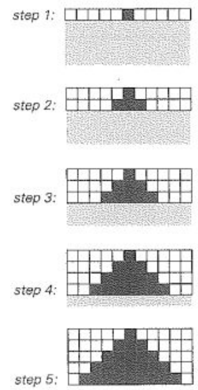

*Stephen Wolfram **A New Kind of Science** 2002. Publisher: Wolfram Media Place: Champaign, IL ISBN: 1-57955-008-8*

Rule 254

Neighbourhood is

- current cell
- neighbour on left
- neighbour on right

There are $2^3 = 8$ possible previous states ("inputs"). And for a given rule set the output for each of those possibilities will be either on or off.

A representation of the rule for the cellular automaton shown above. The top row in each box gives one of the possible combinations of colors for a cell and its immediate neighbors. The bottom row then specifies what color the center cell should be on the next step in each of these cases. In the numbering scheme described in Chapter 3, this is cellular automaton rule 254.

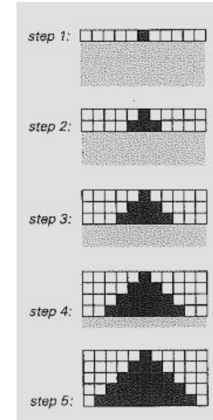

Rule Numbers

There are 256 possible rule sets (for a 2-state CA with a 3-cell neighbourhood.) To enumerate them, Wolfram constructed a binary number that corresponds to which initial-conditions produce which states.

But of the $2^8 = 256$ elementary cellular automata, there are only 88 fundamentally inequivalent rules. The others are mirror-images or inversions of other rules (e.g. where all on states are off states and vice-versa).

e.g. rule 0 (every input produces a white cell) is structurally equivalent to rule 255 (every input produces a black cell)

128 64 32 16 8 4 2 1

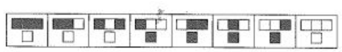

0 0 0 1 1 1 1 0 = 30

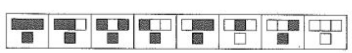

1 1 1 1 1 0 1 0 = 250

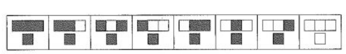

1 1 1 1 1 1 1 0 = 254

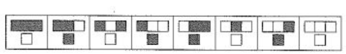

0 1 0 1 1 0 1 0 = 90

Rule 254

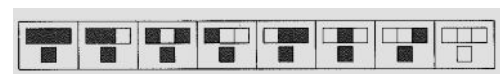

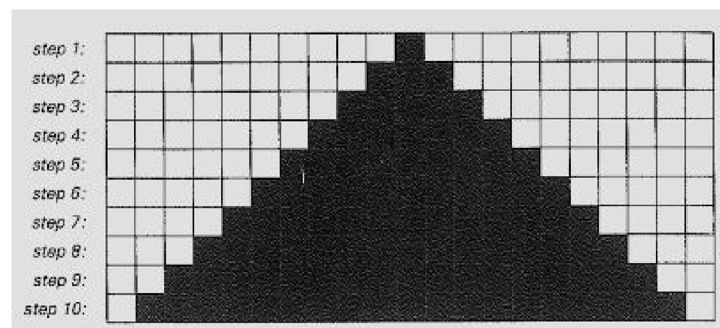

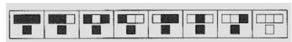

Rule 254

Rule 250

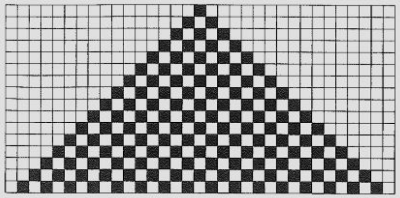

A cellular automaton with a slightly different rule. The rule makes a particular cell black if either of its neighbors was black on the step before, and makes the cell white if both its neighbors were white. Starting from a single black cell, this rule leads to a checkerboard pattern. In the numbering scheme of Chapter 3, this is cellular automaton rule 250.

Rule 90

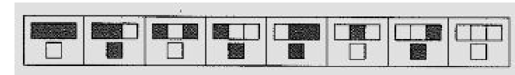

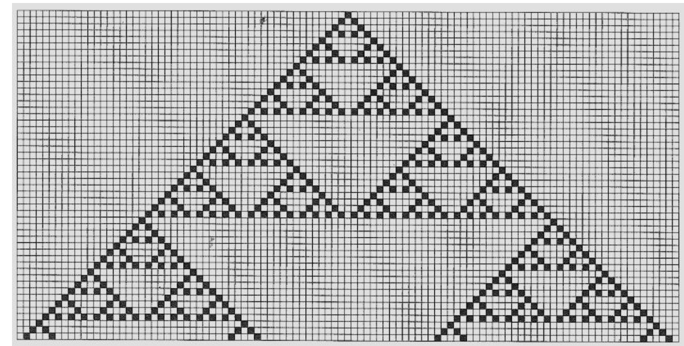

Rule 30

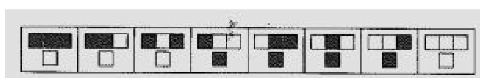

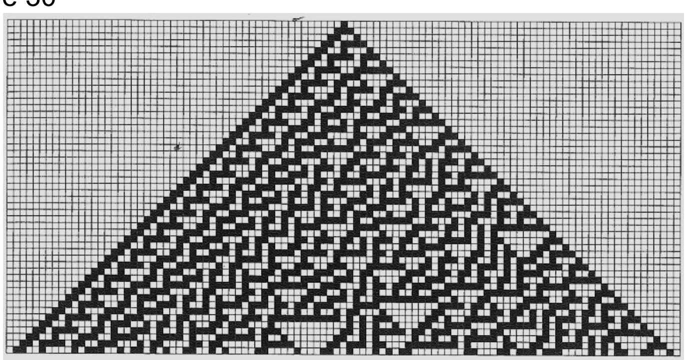

Rule 30

First 500 steps

Some signs of temporary order on the left-hand side, but chaotic on the right-hand side.

"even continuing for a million steps many aspects of the pattern obtained seem perfectly random according to standard mathematical and statistical tests." (p. 30 Wolfram 2002)

This rule is used as a random-number generator in *Mathematica*.

...out of simple rules, Chaos!
Non-predictability!

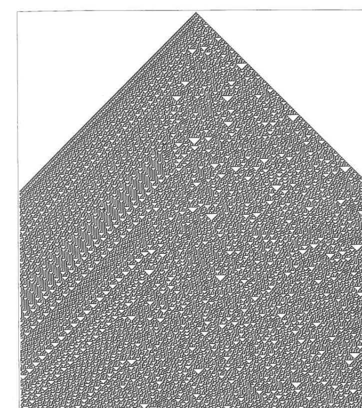

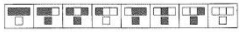

Rule 110

This CA only grows down and to the left.

On the right, we see 150 updates of this rule from a single on cell.

But does it create **organized** or **chaotic** dynamics?

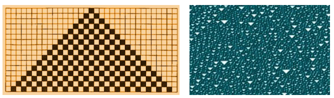

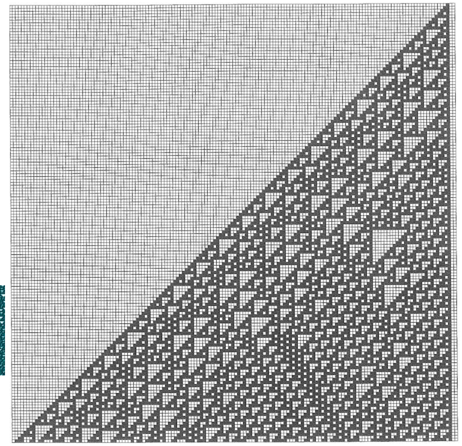

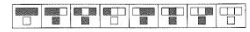

Rule 110

Now we see 700 steps.

Is this **organized** or **chaotic**?

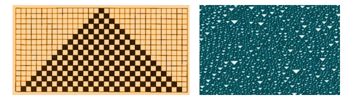

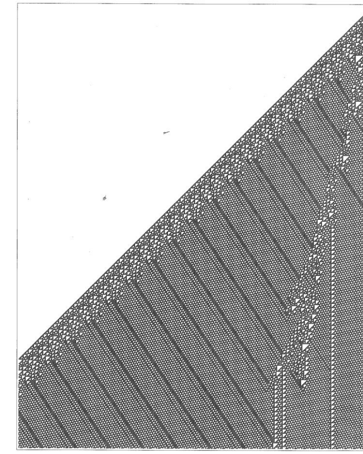

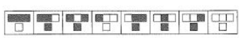

Rule 110

Now we see steps 700-1400

Is this **organized** or **chaotic**?

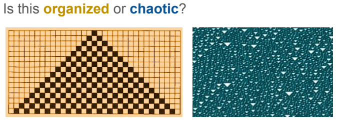

Is it ever going to stop?

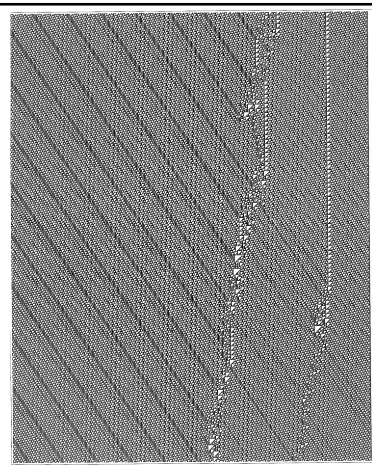

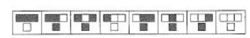

Rule 110

Now we see steps 1400-2100

Is this organized or chaotic?

Is it ever going to stop?

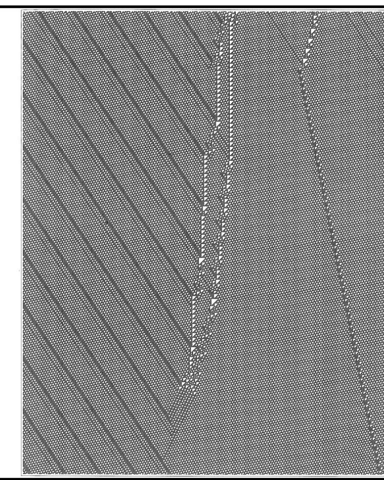

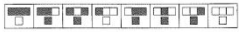

Rule 110

Now we see steps 2100-2800

Is this organized or chaotic?

Is it ever going to stop?

Remember you are effectively seeing the entire system in these images. The stuff to the left is periodic and regular. There is nothing to the right (it only grows to the left).

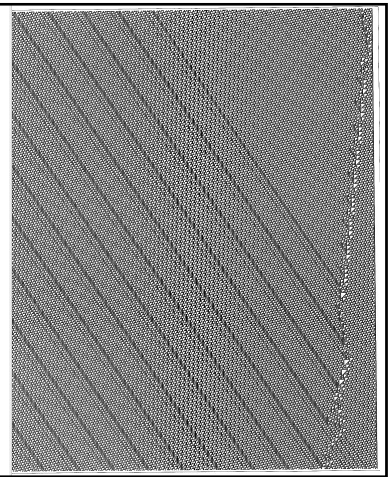

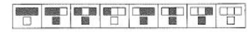

Rule 110

Now we see steps 2800-3400.

What has happened?

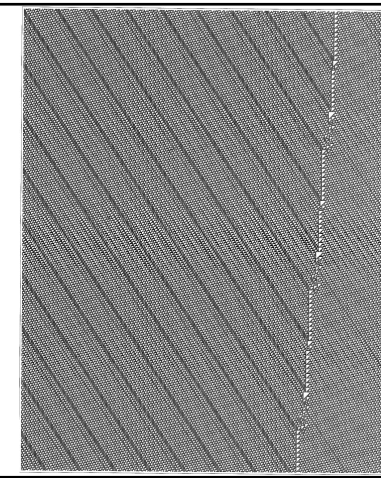

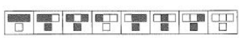

Rule 110

Now we see steps 2800-3400.

What has happened?
The single initially on cell produces an extended transient, lasting 2780 steps. At this point, a regular recurring structure emerges. (The part to the left expands forever, but the pattern is now predictable.)

Is this **organized** or **chaotic**?

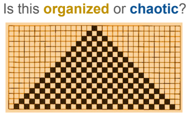

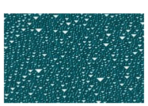

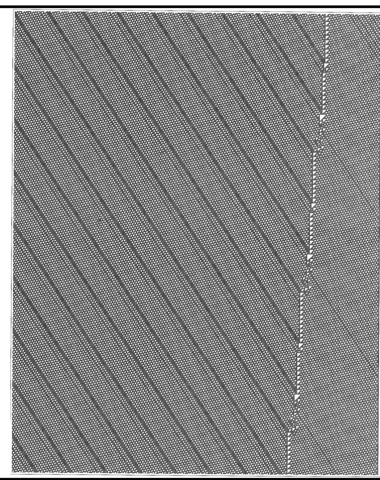

Rule 110

Wolfram calls this class neither chaotic nor ordered, but **COMPLEX** (more on this in a second).

Sidenote: this particular rule has been proven to be **universal**, "effectively capable of emulating any other system"
(http://mathworld.wolfram.com/Universality.html)

c.f. a Universal Turing Machine is capable of simulating the operation of any Turing machine on any input.

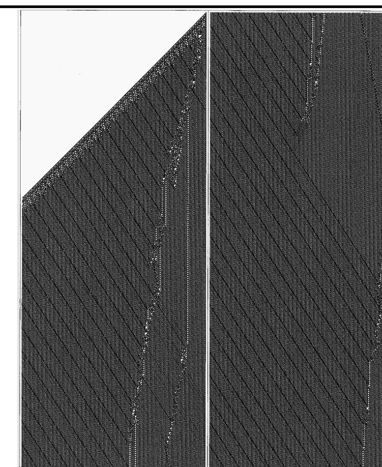

Wolfram's Classification Scheme

Q: What does a given CA do when started from **random** initial conditions?

Each initial condition produces a different result, but a given CA ruleset tends to produce dynamics that fall into one of four classes.

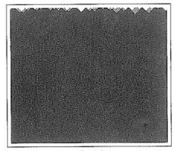
class 1

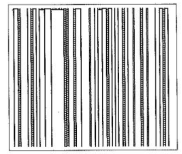
class 2

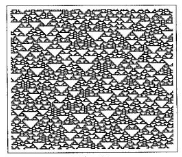
class 3

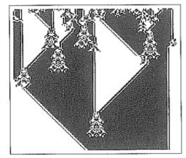
class 4

Wolfram's Classification Scheme

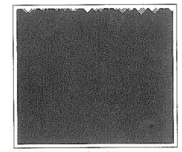
class 1

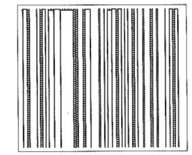
class 2

Class 1: *"Homogenous"*
Behavior is very simple and almost all initial conditions lead to exactly the same uniform final state.

Class 2: *"Regular"*
There are many different possible final states, but all of them consist just of a certain set of simple structures that either remain the same forever or repeat every few steps.

paraphrased following Wolfram **A New Kind of Science** (2002), p 231

Wolfram's Classification Scheme

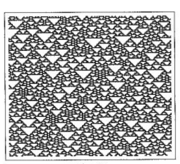
class 3

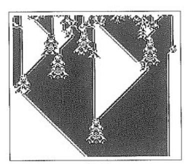
class 4

Class 3: *"Chaotic"*
The behaviour is more complicated and seems in many respects random, although triangles and other small-scale structures are essentially always at some level seen.

Class 4: *"Complex"*
A mixture of order and randomness; localized structures are produced which on their own are fairly simple, but these structures move around and interact with each other in very complicated ways.

Think about what this investigation into cellular automata does for us as scientists.

It provides a simple and "complete" (we know all of the parts) model that we can use to describe different kinds of dynamics (chaos, complex, periodic) etc.

A point of comparison for natural systems.

It provides an example of a very simple system that is essentially not predictable.

It also (unexpectedly) connects to real-world phenomena...

Self-similar "fractal"

Without any complex rules, and without any long-distance communication a fractal is generated.

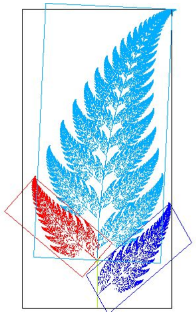

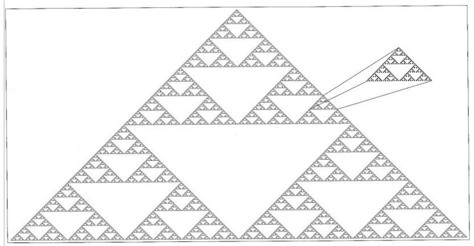

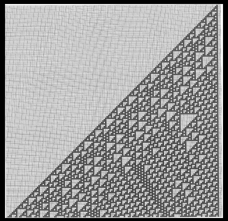

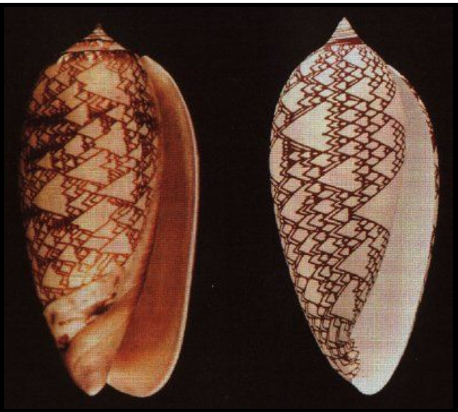

Which of these shells is real?

Meinhardt, H. (1995). The Algorithmic Beauty of Sea Shells. Springer Verlag. (p. 179, 180)

Boids

Swarms

What produces these amazing patterns of behaviour?

Are there leaders and followers?

Are they responding to their environment?

What purpose does swarming serve?

Is the same mechanism(s) at play in schools of fish as in flocks of birds?

http://gencept.com/the-coordinated-beauty-of-thousands-of-birds-moving-in-sync

Boids

Created by Craig Reynolds in 1986, and presented at one of the very first *Artificial Life* Workshops shortly thereafter.

Regarded as a classic example of Artificial Life, demonstrating how complex life-like behaviour can come from simple rules.

Instead of describing things at the high-level (the flock) the system is defined at the low-level (the birds and the simple rules that determine their behaviour). The simple rules result in complex **collective behaviour.**

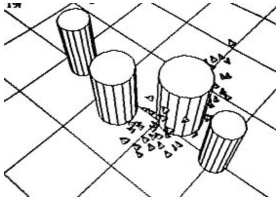

Images taken from Craig Reynold's website. http://www.red3d.com/cwr/boids/

Neighbourhood

Boids change their behaviour in response to the position and velocity of other, **nearby** boids.

The neighborhood is parameterized by:

- a **distance** (measured from the center of the boid)
- an **angle**, measured from the boid's direction of flight (its **heading**)

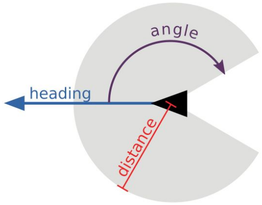

Image is a recreation of one from Craig Reynold's website. http://www.red3d.com/cwr/boids/

Rule 1: Separation

Steer to avoid being too close to flock-mates -- *collision avoidance.*

Accelerate away from each flock-mate at rate inversely proportional to proximity. i.e. the closer you are to a flock-mate, the more you accelerate away from that flock-mate.

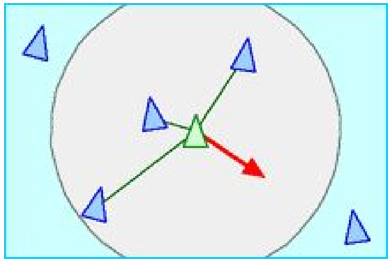

Images taken from Craig Reynold's website. http://www.red3d.com/cwr/boids/

Rule 2: Alignment

Steer toward the average heading of local flock-mates.

Calculate the mean velocity of all flock-mates. Update velocity to be a little bit more similar to that value.

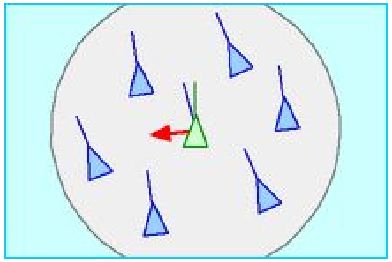

Images taken from Craig Reynold's website. http://www.red3d.com/cwr/boids/

Rule 3: Cohesion

Move toward the average position of local flock-mates.

Calculate the mean position of flock-mates and accelerate toward that position.

DEMO

http://www.emergentmind.com/

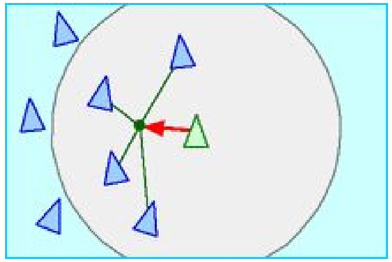

Images taken from Craig Reynold's website. http://www.red3d.com/cwr/boids/

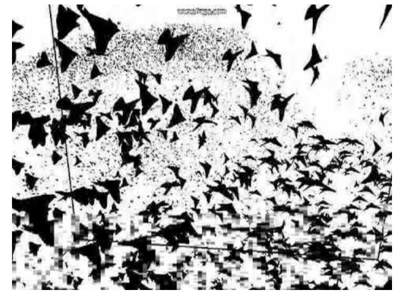

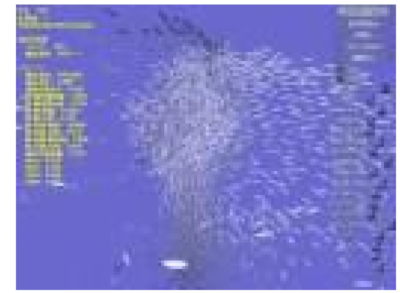

What is learned?

Prediction
We cannot use the model to precisely predict the real-world system. (Although there may be some qualitative predictions that could be made.) But…

Sufficiency Proof
The three simple rules are **sufficient** for this kind of system to display the complex kind of swarming that we witness in Nature. "Life-like" complex (but not random) behaviour. This is an example of...

Emergence
Complex, yet organized global behavior can arise from the interaction of simple local rules.

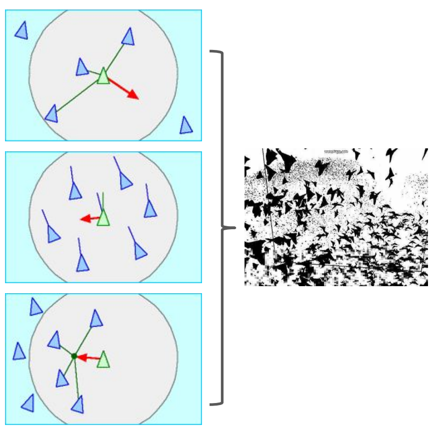

Swarms

Entertainment
Computer animation and games

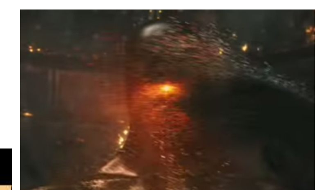

The Day the Earth Stood Still

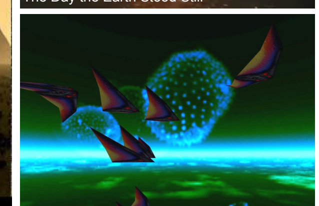

Half-Life

Emergent engineering

How do you design a simple set of rules to produce a desired collective behaviour?

Biology
Insight into group and social dynamics (e.g. ants and termites).

Engineering
How can you get a swarm of robots to perform a particular collective behaviour?

Termites rush to a damaged area of the nest.
U.S. Department of Agriculture
(taken by Scott Bauer)
Damage to a nest of Formosan subterranean termites brings hoards of workers and soldiers with dark, oval shaped heads scrambling to repair the hole.

Artificial Life

"Artificial Life [AL] is the study of man-made systems that exhibit behaviors characteristic of natural living systems. It complements the traditional biological sciences concerned with the analysis of living organisms by attempting to synthesize ***life-like behaviors*** *within computers and other artificial media. By extending the empirical foundation upon which biology is based beyond the carbon-chain life that has evolved on Earth, Artificial Life can contribute to theoretical biology by* ***locating life-as-we-know-it within the larger picture of life-as-it-could-be****."*

Christopher Langton