Intelligence as an Emergent Phenomenon

Lecture 12 I400/I590

Artificial Life as an approach to Artificial Intelligence

Larry Yaeger

Professor of Informatics, Indiana University

Emergence

- Some complex phenomena associated with water:
  - Crystalline ice structures, snowflakes
  - Phase transition to steam, boiling
  - Whirlpools, vortices, turbulence
- All of these complex phenomena are consequences of the same, simple set of local rules that govern the interaction of molecules of water
- Yet the relation between the molecular level rules and the complex macroscopic phenomena is far from obvious
- This is the essence, and a perfect example, of *emergent phenomena*

Intelligence

- "It would be convenient if intelligence were an emergent behavior of randomly connected neurons in the same sense that snowflakes and whirlpools are emergent behaviors of water molecules."—Hillis
- It might then be possible to build a thinking machine just by hooking up a sufficient number of neurons in a suitably architected network

Need to Know

- If true, it allows for the possibility of constructing intelligence without first understanding it
  - Understanding intelligence is difficult
  - And probably a long way off
- So this possibility is very attractive to the would-be builder of thinking machines

But How?

- Unfortunately, the idea does not suggest any specific approach to the problem of designing a mind
- Emergence, in itself, offers neither
  - Guidance on how to construct such a system
  - An explanation of how or why it works

Embrace the Confusion

- For some, emergence offers a way to believe in deterministic, physical causality, while simultaneously maintaining the impossibility of a purely reductionist explanation of thought
- But, as Hillis puts it, anyone who looks to emergent systems as a way of defending human thought from the scrutiny of science is likely to be disappointed

Models of Emergence

- There is an ongoing, modern interest in emergence within computer models of
  - Simulated neural networks
    - Modeling a collection of biological neurons, such as the brain
  - Spin glasses
    - Modeling molecular crystals
  - Cellular automata
    - Modeling ontogenesis, bacterial growth, etc.
  - Evolutionary algorithms
    - Modeling phylogenesis, speciation, engineering problems
- In all cases, both the model and the system being modeled produce dramatic examples of emergent behavior

Parallel Processing

- In addition to parallel distributed processing models, we are beginning to think in terms of parallel computers
- Though, perhaps, this is not a deep philosophical shift, it is of great practical significance, as it enables the study of large emergent systems experimentally
- Even serial processing has sped up to such an extent that personal computers are capable of significant pseudo-parallel modeling and experimentation

Philosophize or Practice?

- Hillis says he has often been asked which “side” he is on, in certain ideological schisms
- “Not being a philosopher, my inclination is to focus on the practical aspects of this question:
  - How would we go about constructing an emergent intelligence?
  - What information would we need to know in order to succeed?
  - How can this information be determined by experiment?”

Symbols

- Hillis comments that "the emergent system that I can most easily imagine would be an implementation of symbolic thought rather than a refutation of it"
- Symbolic thought would be an emergent property of the system

A Parable of Songs and Apes

- [Read from the text]

Symbiosis

- It is not uncommon for two biological species to live together so interdependently that they appear to be a single organism
  - Lichens are symbionts of a fungus and an alga so closely intertwined they can only be separated under a microscope
  - Bean plants need bacteria living in their roots to fix nitrogen they extract from the soil, while the bacteria require nutrients from the bean plants
  - Single-celled *Paramecium bursarra* uses green algae living inside itself to synthesize food
  - The Portuguese Man-'O-War is really a colony of organisms involved in a complex set of mutually beneficial living arrangements

Symbiosis

- Slime molds such as *Dictyostelium* live as individual, single-celled organisms when food is plentiful, but self-organize into slugs and stalks with fruiting bodies when resources are scarce

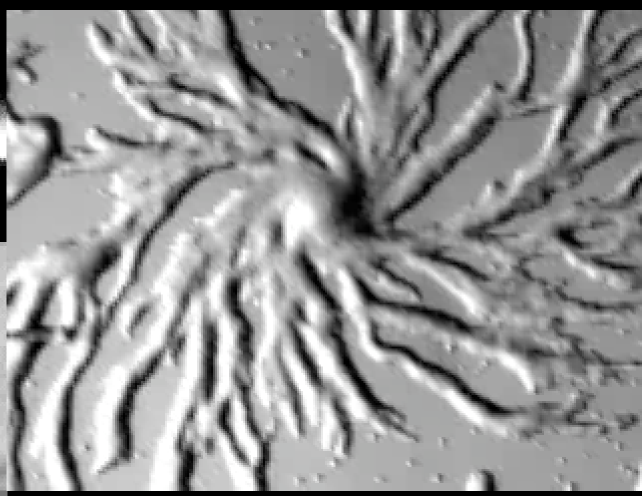

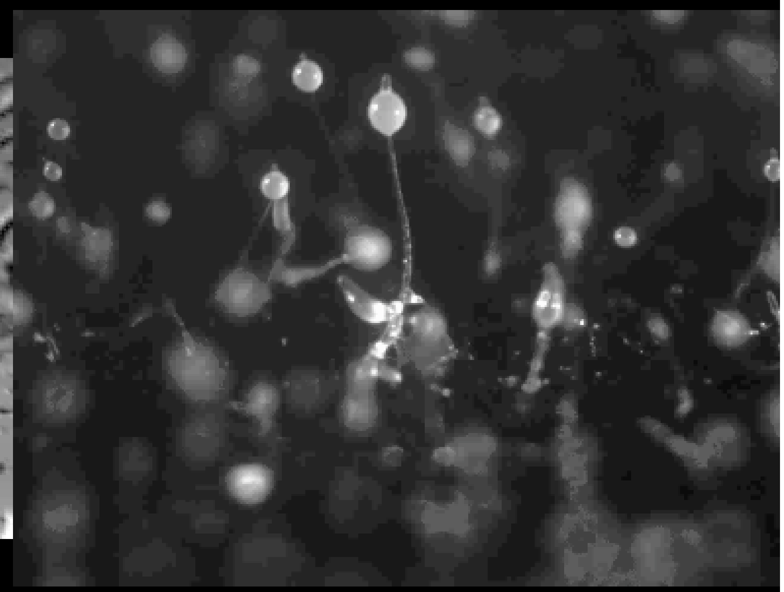

Symbiosis

- There is extensive evidence that mitochondria in all eukaryotes (organisms with nucleated cells) derive from symbiotic bacteria (Ivan Wallin 1920, Lynn Margulis 1981)
  - Mitochondria have their own DNA
  - They have their own, unique protein-building ribosomes
  - Like bacteria:
    - Their DNA has a ring structure
    - They reproduce (inside our cells) by binary fission
    - They produce ATP by oxidation
  - Sequences of mitochondrial DNA closely resemble those in Eubacteria lines

Deeper Symbiosis?

- Freeman Dyson (*The Origins of Life*, 1985) suggests that all modern biological life is a symbiotic combination of two very different self-reproducing systems—drops of oil and free-floating DNA
- Oil drops absorbed raw chemicals from their surroundings, and split when it reached a certain size, with roughly half its constituents going to each part
- Self-catalyzing DNA (or perhaps RNA) parasitically invaded the oil drops, using the concentrated chemicals to reproduce and piggybacking on the "cell division" of the drops
- Perhaps the DNA also contributed something, but, in any event, the two coevolved into the mutually beneficial symbiosis we know today as life

Parallels

- Biological symbiosis is conceptually similar to Hillis's two-part story of intelligence
- Both posit a preexisting homeostatic mechanism that is infected by an opportunistic parasite
- Both consist of two parts that reproduce according to very different rules, but which coevolve so successfully that the resulting symbiont appears to be a single entity

Parallels

- Hillis suggests that choosing between ***emergence*** and ***symbolic computation*** in the study of intelligence is like choosing between ***metabolism*** and ***genetic replication*** in the study of life
- Just as…
    - metabolic systems provide a substrate in which the genetic system can work
- so…
    - an emergent system may provide a substrate in which the symbolic system can operate

Symbolic Approach

- Currently the metabolic system of life is far too complex for us to fully understand or model
- By comparison, the Mendelian rules of genetic replication are almost trivial
- It is possible to model and study them without worrying about the metabolism that supports them
- Similarly, symbolic thought may be fruitfully studied without worrying about the details of the emergent system that supports it
  - This is the traditional, symbolic approach to AI, and has yielded "the most" progress (according to Hillis)

Emergent Approach

- The other approach is to build a model of the emergent substrate of intelligence
- It would not need to mimic in detail all relevant biological mechanisms
- But it would need to provide those emergent properties that are necessary to support the operations of thought

What Properties?

- What is the minimum we would need to know in order to construct an artificial brain?
- Hillis's first fundamental criterion is the ***size*** of the system
  - How big should it be?

Size from Storage Mechanism

- One way to estimate the necessary size is by understanding the storage mechanism used by the brain
- If information is stored primarily by modifications to synapses, then we can start by
  - Counting synapses — 10^{15}
  - Estimating capacity of individual synapse — few bits
- However, much capacity may be unused or used inefficiently
- So even as an upper bound, this estimate, say 10^{16} bits, may be less than informative

Size from Statistical Sampling

- A second method is to measure the information in symbolic knowledge by statistical sampling
- E.g., estimate the size of a person's vocabulary by testing on words randomly sampled from a dictionary
  - The fraction of test words known estimates the fraction of total words known
  - So estimate the total vocabulary by multiplying the test fraction times the total dictionary size

Size from Statistical Sampling

- Could similarly estimate how many facts in the *Encyclopaedia Britannica* are known by an individual
- However, such estimates only address predetermined bodies of knowledge
  - Our encyclopedic estimate gives no measure of facts known to the individual but not contained in the encyclopedia
  - These methods offer no measure of non-symbolic, sensory knowledge (visual, auditory, olfactory, tactile/kinesthetic)
- But perhaps such a method could provide an informative lower bound

Size from Statistical Sampling

- A related measure might come from the game of 20 questions
- Each answer provides one bit of information
- Skillful players require almost all 20 questions
- So knowledge *in common* between the two players is on the order of 2^{20} bits, or about one million (10^6) bits
- But the questions are not perfect, the objects are not chosen at random, it addresses only the shared knowledge, and ignores other sensory modalities
- Perhaps a refined version could be developed and used to provide another lower bound

Size from Acquisition Rate

- A third approach is to estimate the average rate of information acquisition and multiply by lifespan
- Experiments on memorizing random sequences of symbols suggest a maximum rate of memorization (of this type) is about one “chunk” per second
  - A chunk, in this context, is safely < 100 bits
- If true, a 20 year old human, learning at the maximum rate for 16 hours/day (and never forgetting) would know less than 50 billion bits of information, or 5×10^{10} bits

Size from Acquisition Rate

- A problem with this estimate is that the experiment only measures information coming through one sensory channel under one set of circumstances
- If visual imagery were stored directly and completely, it would be necessary to dramatically increase this 100 bit/second limit
- However, in experiments measuring the ability of exceptional individuals to store eidetic (extraordinarily accurate and vivid) images of random-dot stereograms, subjects are given 5 minutes to memorize 100x100 dots
  - This is about 33 bits/second
  - Hillis suggests that memorizing only a few hundred bits per image would probably suffice

Size from Acquisition Rate

- Even if we accept reports of extraordinary feats of memory, such as those of Luria's showman in *Mind of the Mnemonist*, at face value…
- The maximum rate of commitment to memory never seems to exceed a few bits/second
- Experiments should be able to refine this number
- Even if we knew the maximum rate of memorization exactly, the rate averaged over a lifetime would probably be much less
- But knowing the maximum rate could establish an upper bound on storage requirements

Estimate of Storage Requirements

- Based on these sketchy data, Hillis estimates that an intelligent machine would require about 10^9 bits of storage, plus or minus two orders of magnitude
- As a “would-be builder of thinking machines”, Hillis finds this number encouragingly small, since it is well within the capacity of modern computers
- As a “human with an ego”, he finds it distressing that his entire lifetime of memories could be placed on a reel of magnetic tape
- He hopes that experimental evidence will further illuminate the issue

What Properties?

- Hillis's second fundamental criterion for the emergence of intelligence is the required rate of computation
  - How fast does it need to be?

Speed from Bandwidth

- Unlike information theory's bits, for processing speed there is no agreed-upon metric
- It is difficult, if not impossible, to define a unit of measure that does not depend on the manner in which data is represented
- So, for now, try to answer, "Given an efficiently stored representation of human knowledge, what rate of access to that storage (in bits per second) is required to achieve humanlike performance?"
  - *Efficiently stored representation* means any representation requiring only a multiplicative constant of storage over the number of bits of information

Bandwidth from Knowledge Access

- Hillis speculates that human-level performance will require large fractions of an individual's knowledge to be accessed several times per second
- Given a representation of acquired intelligence with a "realistic" representation efficiency of 10%, and a need to access all of that information 10 times/sec…
- Hillis estimates that the 10^9 bits of memory we calculated earlier would require a memory bandwidth of about 10^{11} bits/sec
- This corresponds to about one bit/sec per neuron in the cerebral cortex

Computer Memory Bandwidth

- Memory bandwidth of a conventional sequential computer at the time Hillis was writing (1988) was in the range of 10^6 to 10^8 bits/sec
- The original 65,536-processor Connection Machine could access its memory at about 10^{11} bits/sec
  - "It is not entirely coincidence that this fits well with the estimate above."
- 1.25 GHz frontside bus speeds with standard DDR SDRAM in modern computers has increased this to about 0.8×10^{11} bits/sec
  - 1.6×10^{11} in common dual-processor, dual-bus configurations
  - Special purpose buses and RAM in GPUs exceed this

What Properties?

- Hillis's third fundamental question regarding the emergence of intelligence is
  - What sensory-motor functions are necessary to sustain symbolic intelligence?

Sensory-Motor Needs

- Hillis points out that "an ape is a complex sensory-motor machine, and it is possible that much of this complexity is necessary to sustain intelligence"
- Large portions of the brain seem to be devoted to visual, auditory, and motor processing, and it is unknown how much of this machinery is needed for thought

Sensory-Motor Needs

- A person who is blind and deaf or totally paralyzed can undoubtedly be intelligent
  - But this does not prove that the portion of the brain devoted to these functions is unnecessary for thought
  - It may be, for example, that a blind person takes advantage of the visual processing apparatus of the brain for spatial reasoning
- fMRI data show some brain areas involved in both imagining and actually moving body parts

Sensory-Motor Needs

- It should be possible to identify certain functions as unnecessary for thought by studying patients who have localized brain damage, but unimpaired cognitive abilities
- For example, damages to certain parts of visual cortex impair binocular fusion, but demonstrate no obvious impairment of cognitive function
- Hillis says "one can imagine metaphorically whittling away at the brain until it is reduced to its essential core"
- But, accidental damage to the brain rarely limits itself to a single area of the brain completely and exclusively

Sensory-Motor Needs

- Hillis suggests it may be more productive to assume that all sensory-motor apparatus is unnecessary to thought until proven otherwise
- This does *not* mean sensory-motor apparatus and function are unimportant for the evolution and application of intelligence
  - "One can believe in the necessity of the opposable thumb for the development of intelligence without doubting a human capacity for thumbless thought."

Levels of Understanding

- Hillis notes that although questions of capacity and scope are necessary to define the magnitude of the task of constructing an emergent intelligence, ultimately the key question is one of understanding
- We may not need *all* the details, but it seems likely that we will need to understand some of the basic principles upon which the emergent substrate of intelligence is based

Gaining Understanding

- One way to develop an understanding of such a substrate is to study the properties of specific emergent systems
- This kind of experimental study is currently being conducted on several classes of man-made systems
  - Neural networks
  - Spin glasses
  - Cellular automata
  - Evolutionary systems
  - Adaptive automata

Gaining Understanding

- Another way to develop this understanding is to study biological systems, which are our only real examples of intelligence and our only real examples of an emergent system that has produced intelligence
- Disciplines providing useful information of this type include
  - Neurophysiology
  - Cognitive psychology
  - Evolutionary biology

Gaining Understanding

- A third approach would be to develop a theoretical understanding of the requirements of intelligence and the phenomena of emergence
- Theoretical work proceeds in domains including
  - Logic and computability
  - Linguistics
  - Dynamical systems theory

Limited Understanding

- One should not conclude, however, that a complete reductionist understanding is necessary for the creation of intelligence
- "Even a little understanding could go a long way toward the construction of an emergent system."

Limited Understanding

- An excellent example of the practical application of limited understanding is the success at modeling the emergent behavior of fluids with cellular automata
- Vortical flows are difficult to analyze
  - Equations (Navier-Stokes) have been known for a century, but except for a few simple cases they cannot be solved analytically
  - The common method of solution is a numerical integration of a discretized form of a subset of the equations tailored to the flow regime of interest

Cellular Automata Fluids

- With sufficient computational power, it is possible to simulate fluids, despite very limited understanding, by simulating billions of very simple colliding particles that represent molecules in the fluid
  - These are not models of detailed molecular interactions (which would be horrifically difficult and computationally expensive)
  - Rather, a few simple constraints, such as conservation of energy and of particle number, are sufficient to reproduce large-scale behavior
  - A system of simplified particles that obey these two laws, but are otherwise unrealistic, can reproduce the same emergent phenomena as real physics

Cellular Automata Fluids

- Particles of unit mass, moving only at unit speed, on a hexagonal lattice, and colliding according to the rules of billiard balls produce
  - Laminar flow
  - Vortex streams
  - Turbulence

  that are “indistinguishable from the behavior of real fluids”

Cellular Automata Fluids

- The detailed rules of interaction are very different from the interactions of real molecules, yet the emergent phenomena are the same
- The emergent phenomena of fluid flow can be created without understanding the details of the forces between molecules or the equations that describe fluid flow
- Note, however, that the model was created by physicists who knew a lot about molecules and fluids
  - That knowledge informed the selection of features that were important to implement (and those which were not)

Even More Limited Understanding

- Physics is an unusually exact science. Perhaps a better example of an emergent system that can be simulated with only a limited understanding is evolutionary biology
- Despite significant limits on our current understanding of real biological evolution, we are able to use evolutionary algorithms that produce interesting emergent behaviors
- Hillis evolved programs to sort numbers, quite successfully, yet cannot understand the details of the resulting algorithm's evolutionary history or even how it works

Putting Limited Understanding to Work

- The fluid flow and simulated evolution examples suggest that it is possible to make a great deal of use of a small amount of understanding
- If a similar process produces a system of emergent intelligence, we may very well have a similar lack of understanding about how it works
    - However, it may then provide us with the tools to analyze an intelligent system in detail and develop such an understanding
    - Like Randy Beer's "frictionless mind" metaphor—while frictionless physics is not at all realistic, it allows us to abstract real physics in such a way as to yield powerful insights into the way the world works

Emergent Intelligence

- Hillis speculates that such an emergent system would not be intelligent in and of itself, but rather a “metabolic substrate” in which intelligence would grow
- In terms of the apes and the songs, he suggests that the emergent portion of the system would play the role of the ape, or at least that part of the ape that hosts the songs
- This artificial mind might need to be inoculated with human knowledge, not unlike teaching a child
  - Like a child, the emergent mind would presumably be susceptible to bad ideas as well as good
  - Hillis suggests the result would be not so much an artificial intelligence, but rather a human intelligence sustained within an artificial mind

An Emergent Dream

- "Of course, I understand that this is just a dream, and I will admit that I am propelled more by hope than by the probability of success. But if this artificial mind can sustain itself and grow of its own accord, then for the first time human thought will live free of bones and flesh, giving this child of mind an earthly immortality denied to us."
- And the songs will finally have evolved a more perfect host

Credits

- Lecture based largely on the reading assignment:
  Hillis, D. W., Intelligence as an Emergent Behavior, p. 175-189, *Daedalus*, Journal of the American Academy of Arts and Sciences, special issue on Artificial Intelligence, Winter 1988
- *Dictyostelium* movies from http://www.dictybase.org/

References

- Symbiotic origin of Mitochondria in Eukaryotic cells http://cas.bellarmine.edu/tietjen/images/origin_of_mitochondria_in_eukary.htm
- *The Emergent Chain*, Gary Lucas (unpublished)