

An Introduction to Computational Mechanics


1. Computational Mechanics provides another way of measuring an object's complexity or regularities.
2. Unlike the excess entropy, computational mechanics makes use of the models of formal computation to provide a direct, structural accounting of a system's intrinsic information processing.
3. Computational Mechanics lets us see how a system stores, transmits, and manipulates information.

Context:

- As before, we have a long sequence of symbols, s_1, s_2, s_3, \dots , from a binary alphabet. Assume a stationary probability distribution over the sequence.

Measurement Channel

Consider again a long sequence of measurements:



- On the left is “nature”—some system's state space.
- The act of measurement projects the states down to a lower dimension and discretizes them.
- They then reach the observer on the right.
- Figure source: Crutchfield, In Modeling Complex Systems. L. Lam and H.C. Morris, eds. Springer-Verlag, 1992: 66-10.
- Task: What can the observer infer about the intrinsic computation, the pattern or complexity, of the observed process?

An initial example: The Prediction Game


- Your task is to observe a sequence, and then come up with a way of predicting, as best you can, subsequent values of the sequence.
- The sequence might have non-zero entropy rate, so perfect prediction might be impossible.
- We will begin by focusing at some length on the following example:

... 10111110101110111010111 ...

- After some squinting, you will probably notice that Consider a every other symbol is 1. The other symbols are 0 or 1 with equal probability.

Initial example, continued

- The machine that can reproduce this sequence is:



- This machine has two states, A and B.
- From A, one sees a 1 with probability 1.
- From B, one sees a 1 with probability $1/2$, and a 0 with probability $1/2$.
- Only two states are required. Why?

Initial Example, why two states?

- Why are only two states necessary? And what exactly do we mean by “state”?
- There are many particular observed sequences which give one equivalent information about the future sequences
- For example, if you see 1010, or 1110 or simply 0, in all cases you know with certainty that a 1 is next.
- The idea is that it only makes sense to distinguish between historical sequences that give rise to different predictive information.
- There will usually be many sequences that give the same predictive information. Group these sequences together into a **state**.
- We will formalize this notion of state on the following slide.

Causal States

- How much of the left half \overleftarrow{S} is needed to predict the right half \overrightarrow{S} ?
- Only need to distinguish between \overleftarrow{S} 's that give rise to different states of knowledge about \overrightarrow{S} .
- Two \overleftarrow{S} 's that give rise to the same state of knowledge are equivalent:

$$\overleftarrow{S}_i \sim \overleftarrow{S}_j \text{ iff } \Pr(\overrightarrow{S} | \overleftarrow{S}_i) = \Pr(\overrightarrow{S} | \overleftarrow{S}_j).$$
- Equivalence classes induced by \sim are **Causal States**, minimal sets of aggregate variables necessary for optimal prediction of \overrightarrow{S} .
- In our initial example, 0, 10, 1010, 1011, and many others, are all equivalent under \sim , and hence belong to the same state, called **A**.

ϵ -Machines


- The causal states together with the probability of transitions between causal states are an ϵ -**machine**, a minimal model capable of statistically reproducing the original configuration.
- The ϵ -machine tells us **how the system computes**.
- The “ ϵ ” reminds us that the measurement symbols upon which the machine is formed may be distorted via noise or the discretization process.
- Technical Note: Unlike the examples from yesterday, the ϵ -machines are models of **stochastic** computation.

Statistical Complexity

1. Let $\Pr(\mathcal{S})$ denote the asymptotic distribution of causal states.
2. Then the *statistical complexity* $C_\mu \equiv H[\mathcal{S}]$.
3. C_μ is the minimum amount of memory needed to statistically reproduce the process.

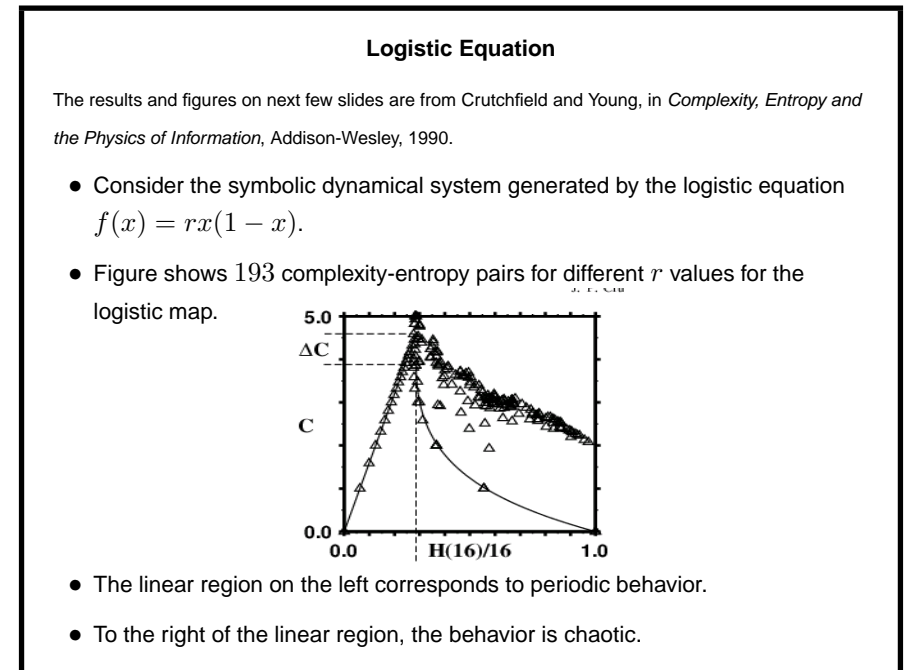
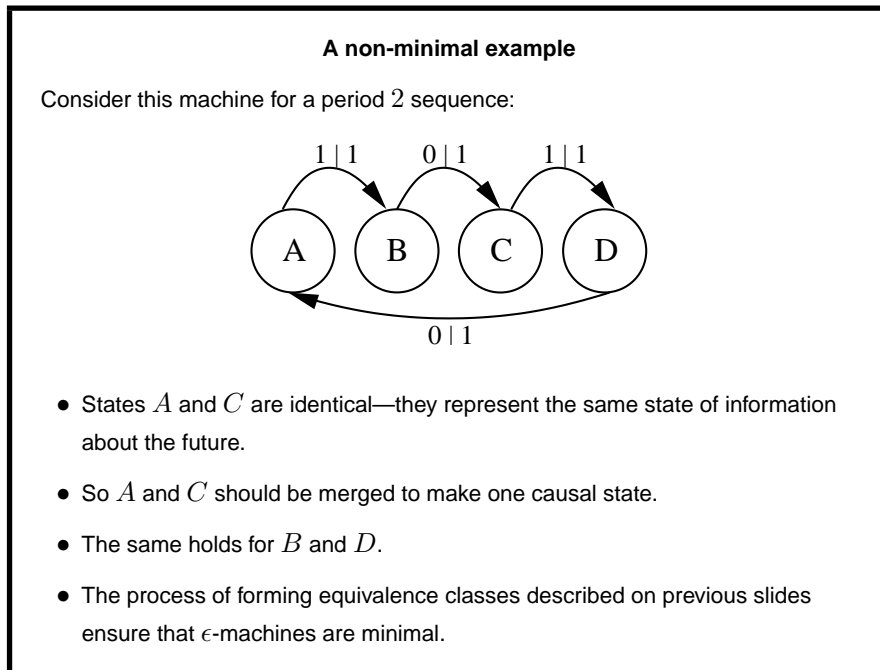
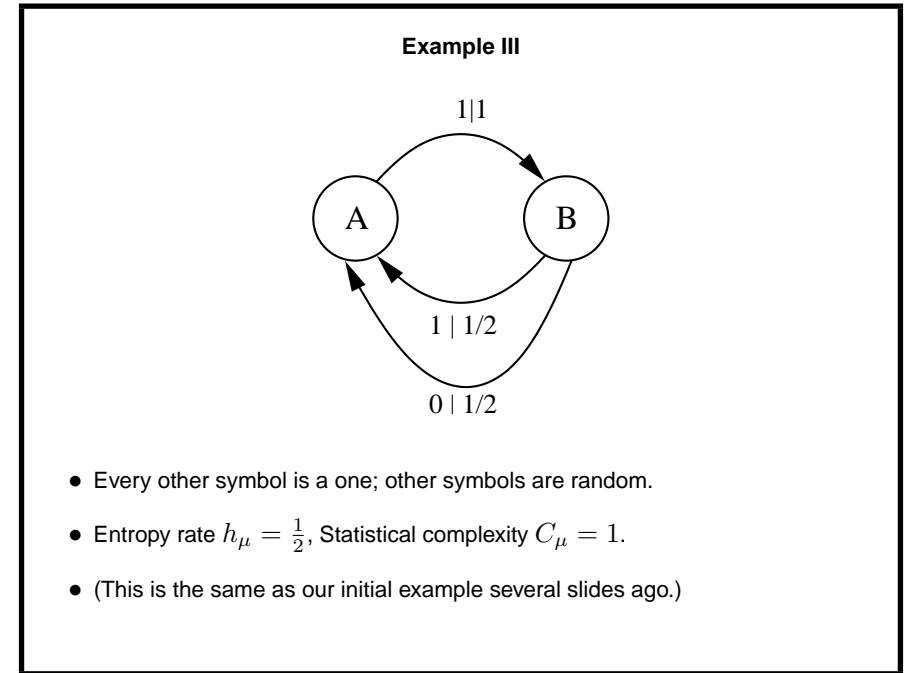
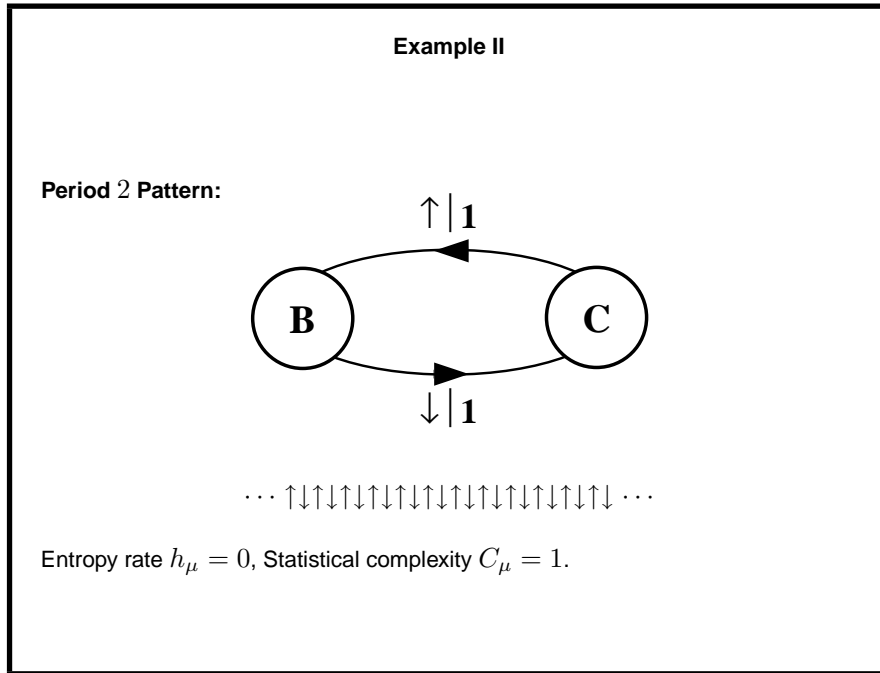
Example I

Fair Coin:




... HHTHTHTTTHTHTHTTHTHH ...

Entropy rate $h_\mu = 1$, Statistical Complexity $C_\mu = 0$.



Logistic Equation: Critical Machine


- What happens as the periods get larger?
- At the period-doubling accumulation point the number of state V in the ϵ -machine diverges.



- This suggests that there is no longer a finite representation at the lowest level of the Chomsky Hierarchy.
- The nature of the divergence leads one to a higher-level computational model for the system.


Critical Machine, continued

- ϵ -machine estimated with a window size of $L = 16$:



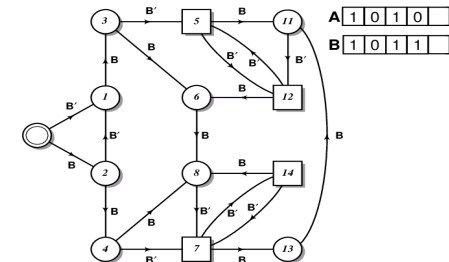
Critical Machine, continued

- Deterministic chains replaced by equivalent strings:



- Looking at the regularities in the machine leads one to a finite representation at a higher computational level.

Critical Machine: Finite Representation



- Two string registers, A and B . Start with A containing 0 and B containing 0.
- Squares are new type of state. When making a transition from a square, update string registers: $A \rightarrow BB$, and $B \rightarrow BA$.
- B' is B with the last bit flipped. (E.g., if $B = 1010$, $B' = 1011$.)
- **Main Point:** Diverging model size at lower level of Chomsky hierarchy necessitated a new model at a higher level.
- Innovation in response to emergence.

Computational Mechanics Conclusions:

Questions:

- What are patterns and how can we discover them?
- What does it mean to say a system is organized?

Summary:

- Computation theory classifies sets of sequences by considering how difficult it is to recognize them.
- Causal states and ϵ -machines adapt computation theory for use in a probabilistic setting.
- The ϵ -machine provides an answer to the question: What patterns are present in a system?

Conclusions, continued

- The ϵ -machine can be inferred directly from observed data.
- The ϵ -machine reconstruction pattern can discover patterns—even patterns that we haven't seen before.
- In his next lecture, Cosma will discuss an algorithm to do infer ϵ -machines from time series.
- Jim Crutchfield, in his lectures, will discuss computational mechanics and its applications in much more detail.

For many references, including applications, algorithms, and proofs, see references in bibliography, which is posted on the wiki.