Intrinsic Computation

Jim Crutchfield
Complexity Sciences Center
Physics Department
University of California at Davis

Complex Systems Summer School
Santa Fe Institute
Santa Fe, NM
21 June 2017

Information Roadmap for a Complex Process

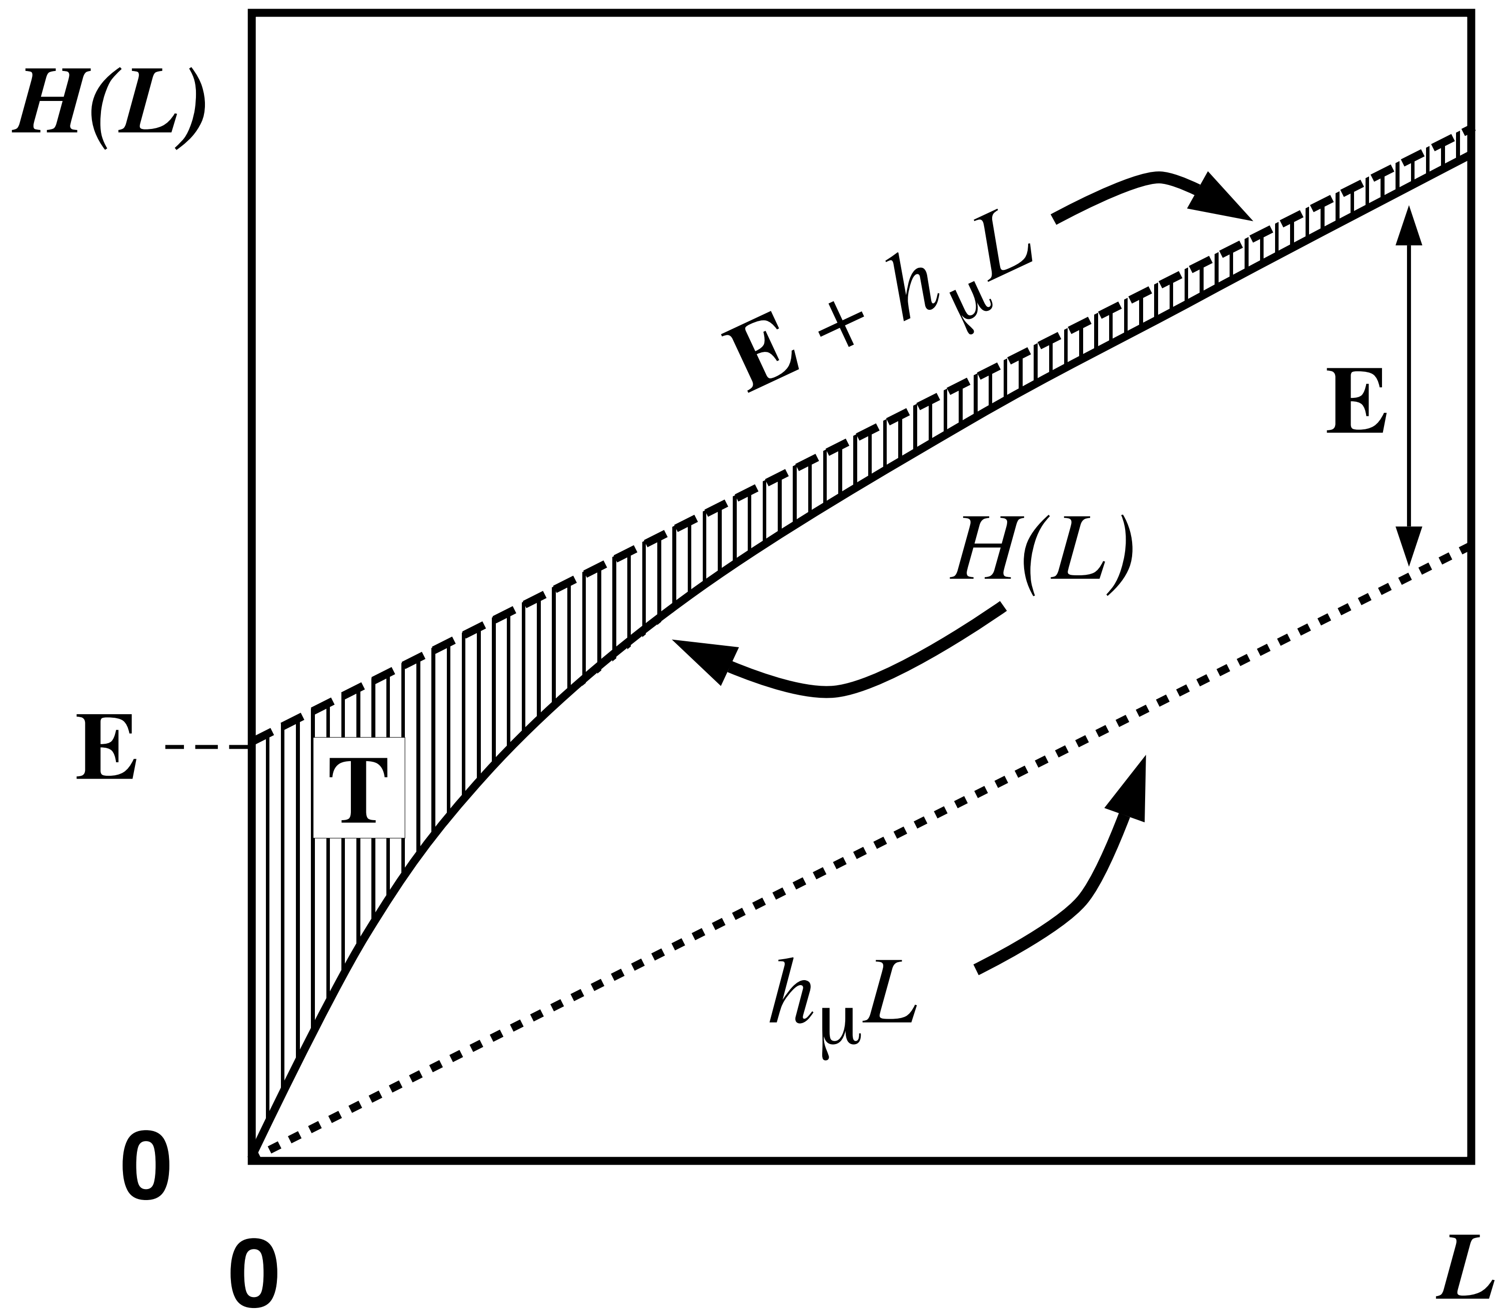

What's wrong with information theory?

The Learning Channel:

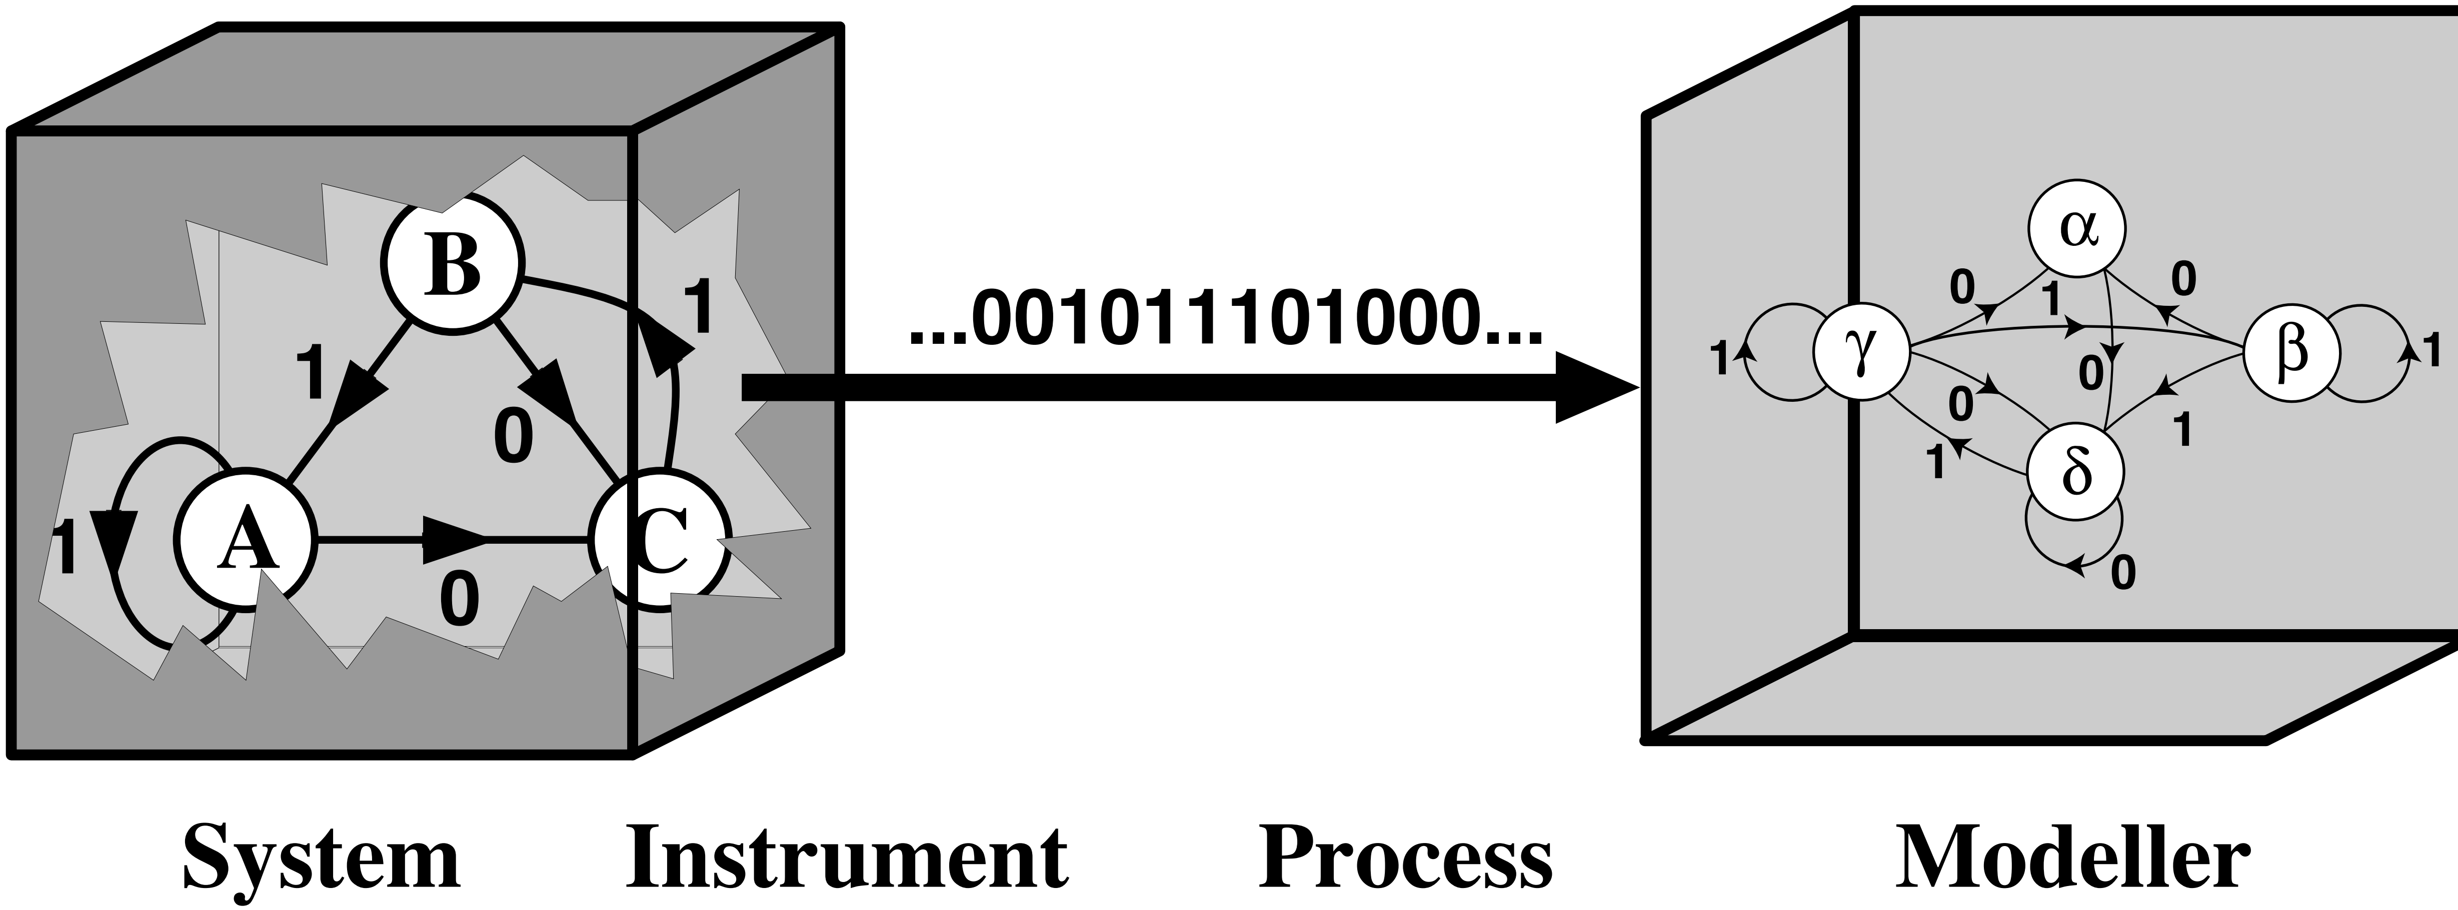

Central questions:

What are the states?

What is the dynamic?

The Learning Channel ...
The Prediction Game

Rules:

1. I give you a data stream (an observed past sequence).
2. You predict its future.
3. You give a model (states & transitions) describing the process.

The Learning Channel ...

The Prediction Game ...

Process I:

Past: $\ldots 111111111111$

Your prediction is?

Future: $111111111111\ldots$

Your model (states & dynamic) is?

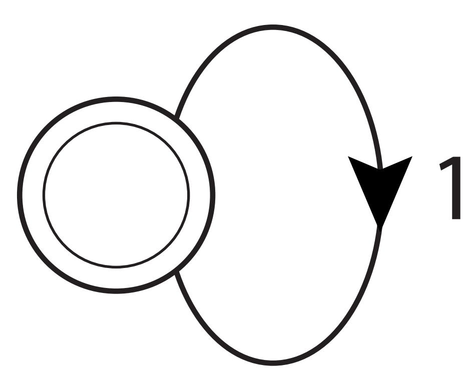

The Learning Channel ...

The Prediction Game ...

Process II:

Past: $\ldots 10110010001101110$

Your prediction is?

Analysis: All words of length L occur & equally often

Future: Well, anything can happen, how about?

$$0101011101001101\ldots$$

Your model is?

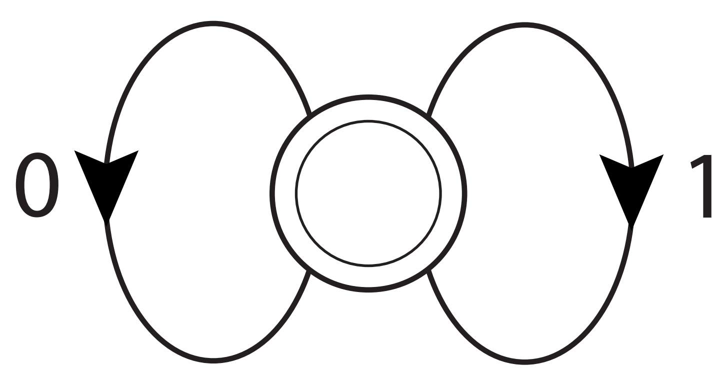

The Learning Channel ...

The Prediction Game ...

Process III:

Past: $\ldots 1010101010101010$

Your prediction is?

Future: $1010101010101010 \ldots$

Your model is?

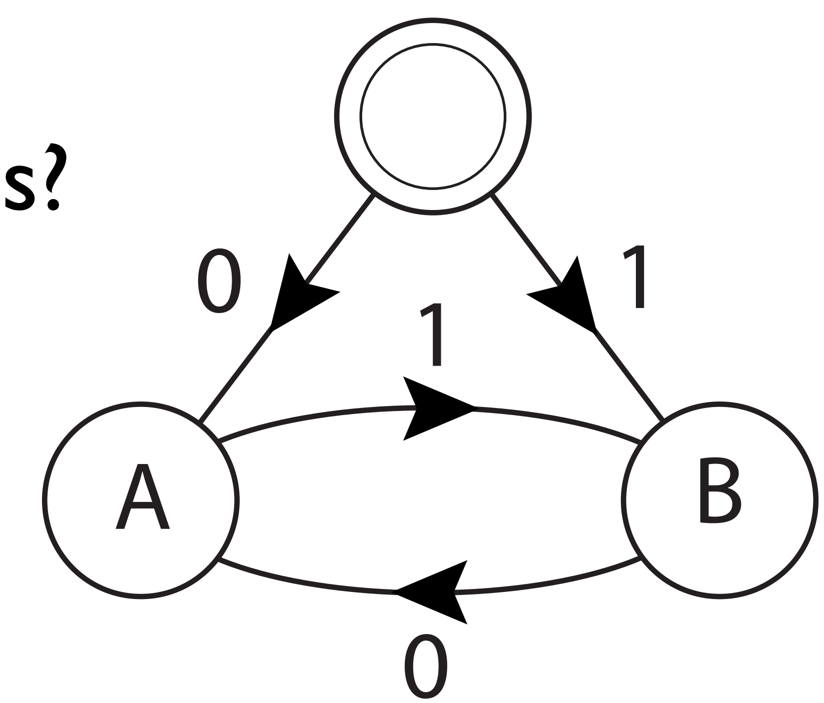

The Learning Channel ...

Theory? Algorithms?

Computational Mechanics

The Learning Channel ...

Goal:

Predict the future $\overrightarrow{S}$

using information from the past $\overleftarrow{S}$

But what "information" to use?

We want to find the effective "states"

and the dynamic (state-to-state mapping)

How to define "states", if they are hidden?

All we have are sequences of observations

Over some measurement alphabet $\mathcal{A}$

These symbols only indirectly reflect the hidden states

The Learning Channel ...

Effective States:

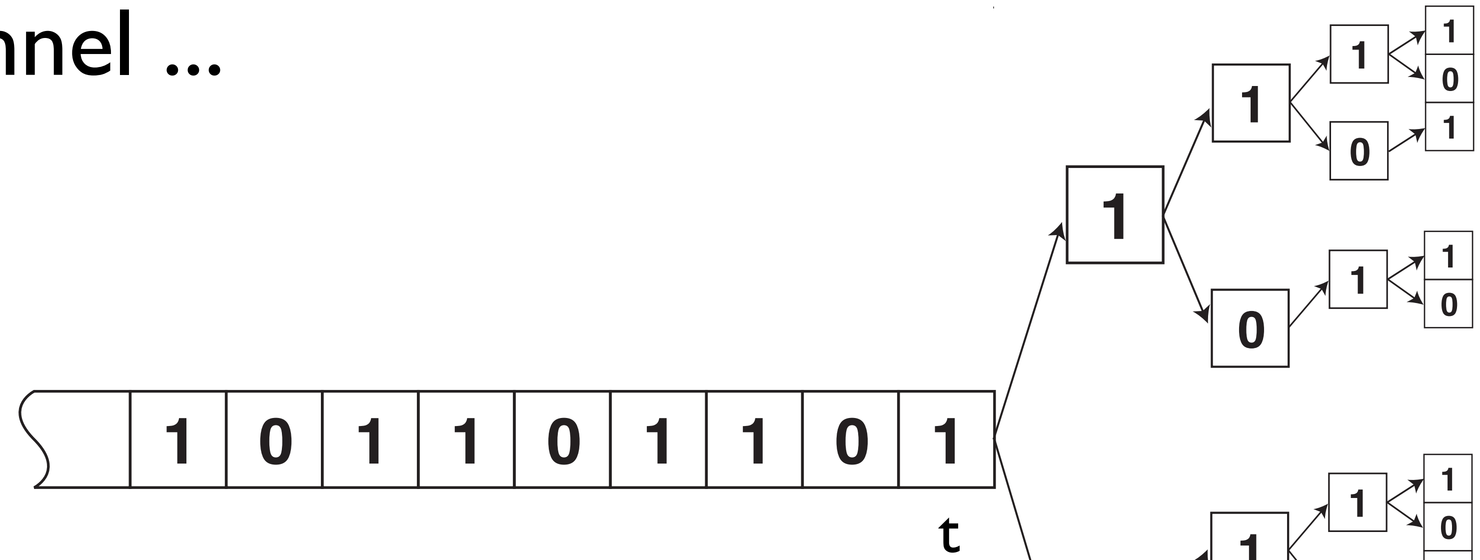

Process is in different "states"
when futures look different
State(t) $\nsim$ State(t+1)

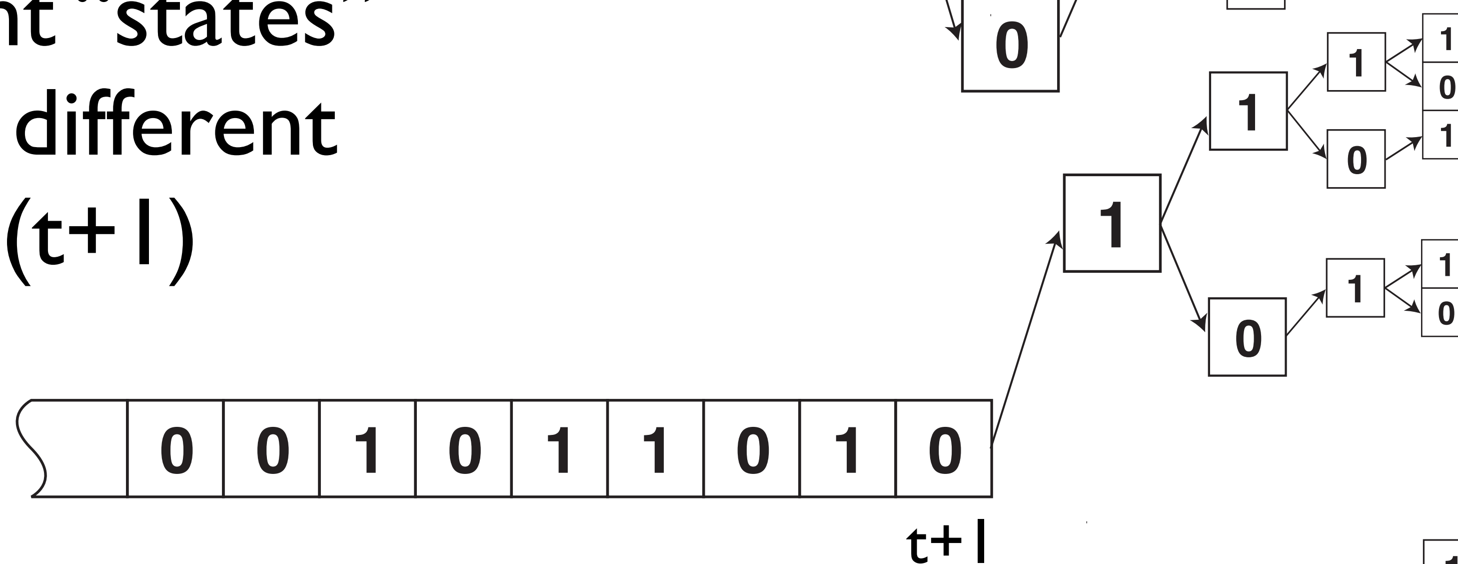

Process is in the same "state"
when the future looks the same:
State(t) ~ State(t+2)

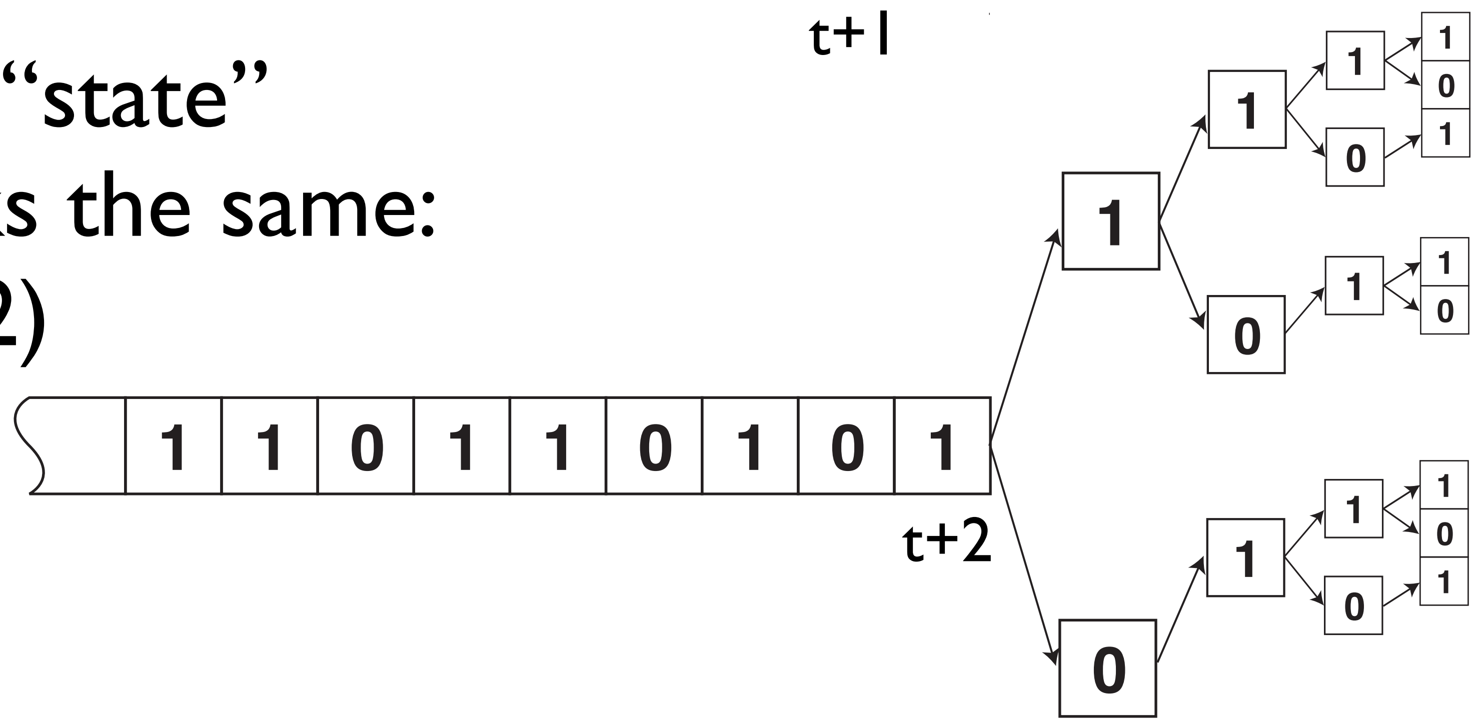

The Learning Channel ...

Effective for what?

What's a prediction?

A mapping from the past to the future.

Process $\Pr(\overleftrightarrow{S})$: $\overleftrightarrow{S} = \overleftarrow{S}\,\overrightarrow{S}$

Future: $\overrightarrow{S}^L$

Particular past: $\overleftarrow{s}$

Future Morph: $\Pr(\overrightarrow{S}^L \mid \overleftarrow{s})$ (the most general mapping)

Refined goal:

Predict as much about the future $\overrightarrow{S}$,

using as little of the past $\overleftarrow{S}$ as possible.

The Learning Channel ...

How Effective are the Effective States?

Candidate "rival" model R (Think … some HMM)

A given mapping from pasts to future morphs

How to measure goodness?

Effective Prediction Error:

$$H[\overrightarrow{S}^L | R]$$

Uncertainty about future given effective states

Effective Prediction Error Rate:

$$h_\mu(R) = \lim_{L \to \infty} \frac{H[\overrightarrow{S}^L | R]}{L}$$

Entropy rate given effective states

The Learning Channel ...
How Effective are the Effective States?

Statistical Complexity of the Effective States:

$$C_\mu(R) = H[R] = H(\Pr(R))$$

Interpretations:

Uncertainty in state.

Shannon information one gains when told effective state.

Model "size" $\propto \log_2(\text{number of states})$

Historical memory used by R.

The Learning Channel ...

Goals Restated:

Question 1:

Can we find effective states that give good predictions?

$$H[\overrightarrow{S}^L | R] = H[\overrightarrow{S}^L \mid \overleftarrow{S}]$$

or

$$h_\mu(R) = h_\mu$$

Question 2:

Can we find the smallest such set?

$$\min C_\mu(R)$$

The Learning Channel ...

Effective States:

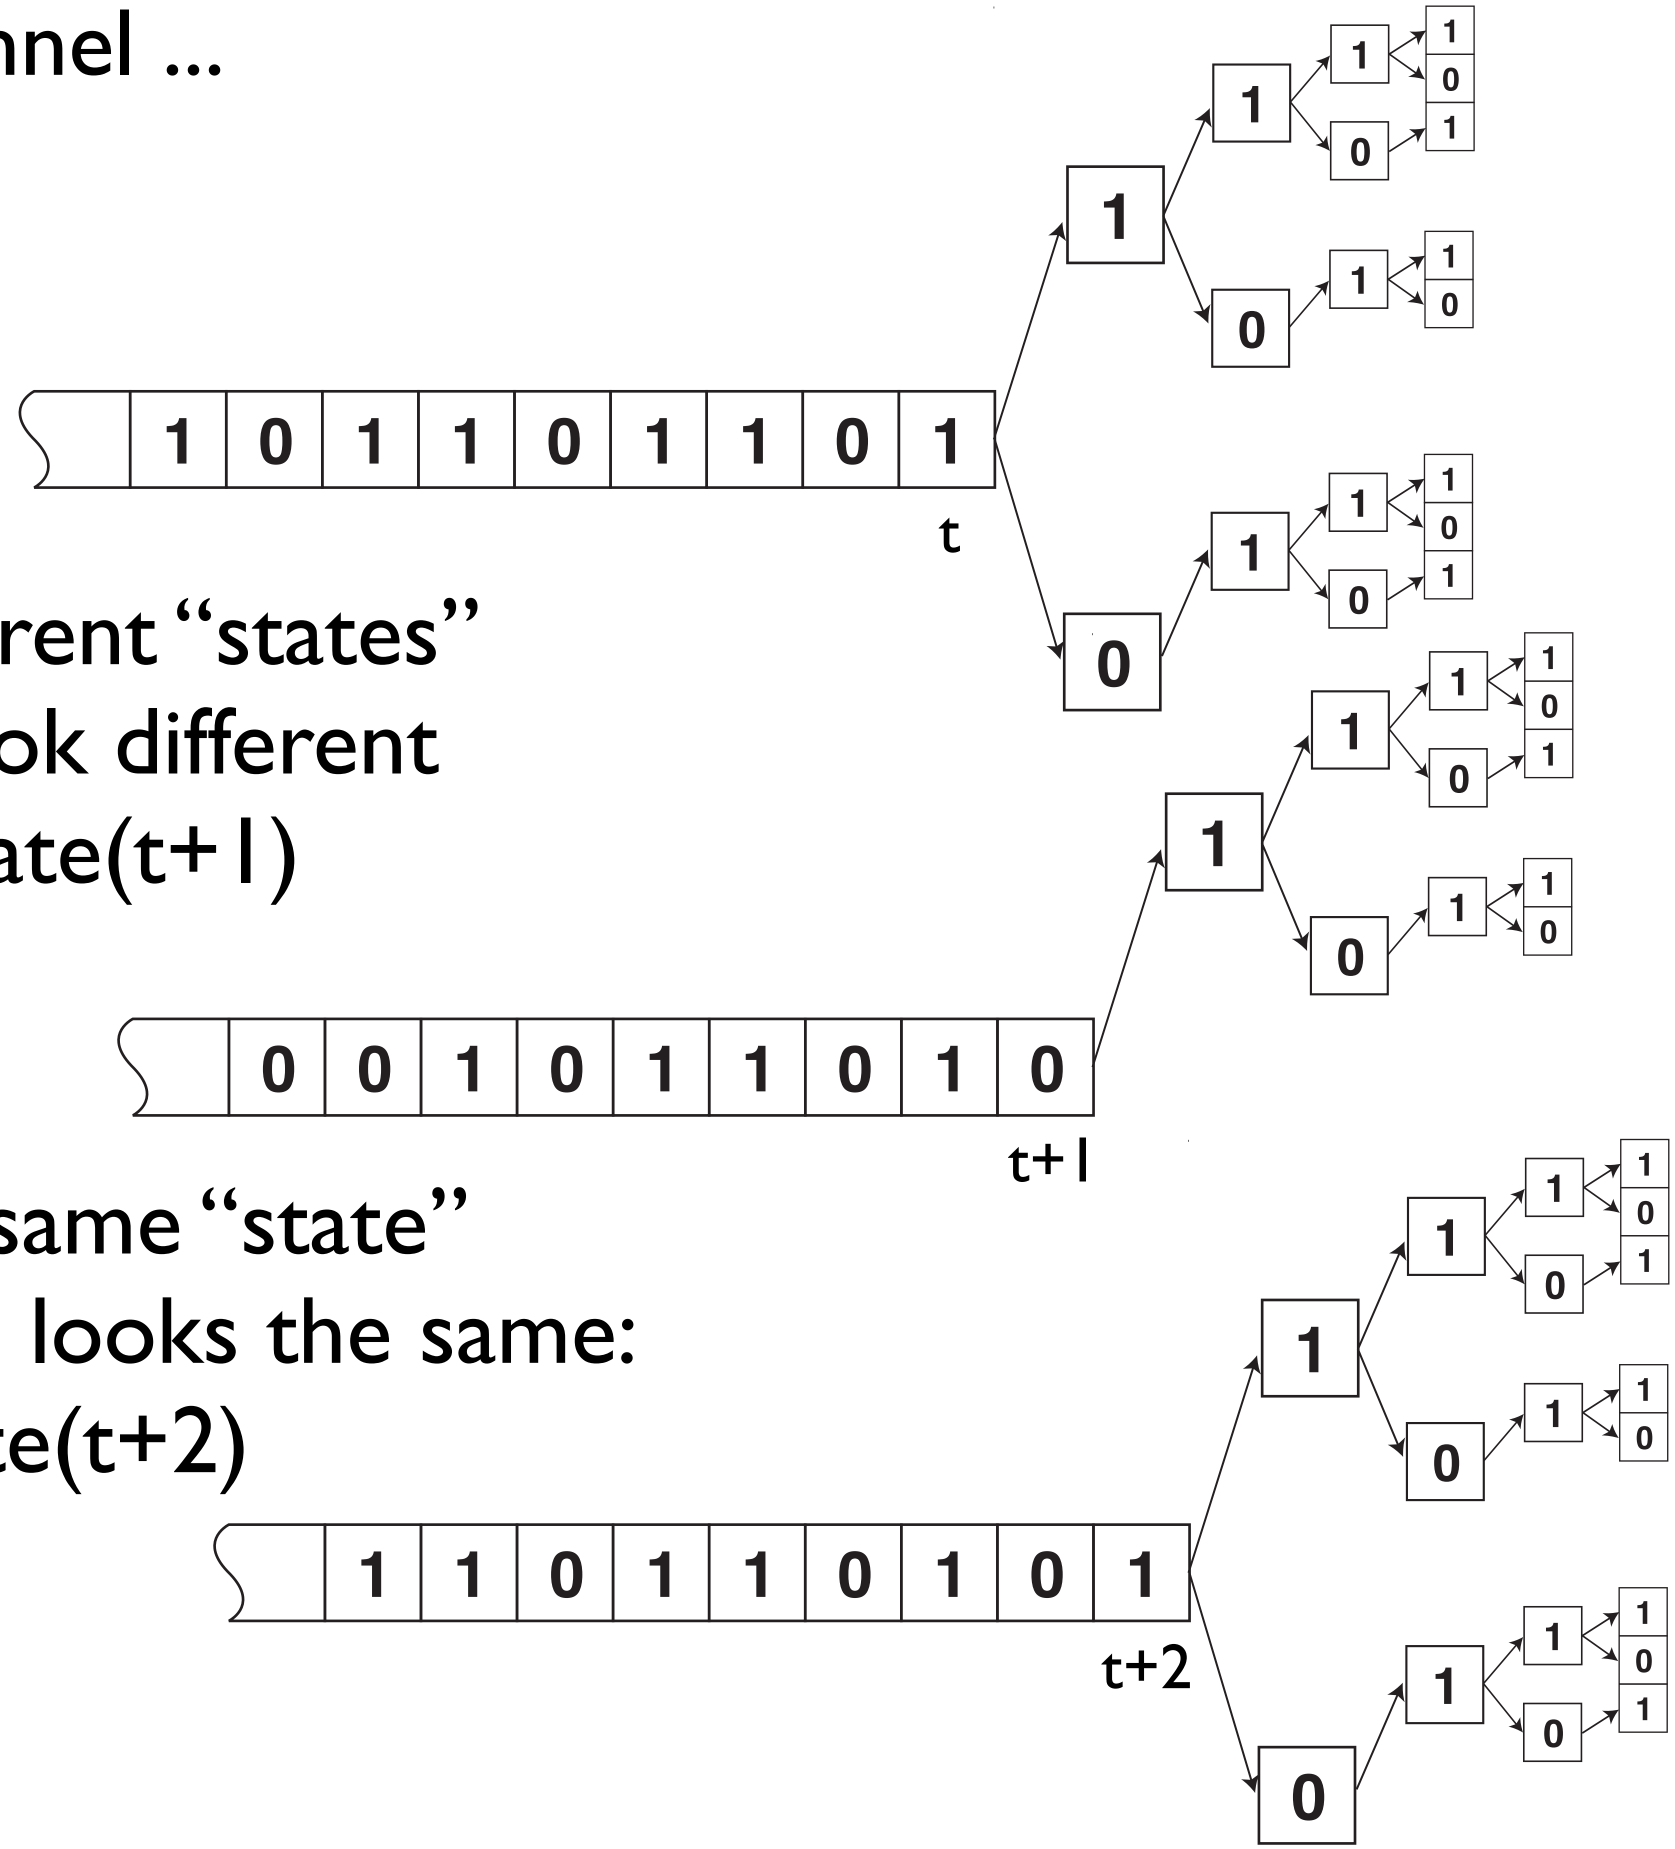

Process is in different “states”
when futures look different
State(t) $\nsim$ State(t+1)

Process is in the same “state”
when the future looks the same:
State(t) ~ State(t+2)

The Learning Channel ...

Causal States:

Causal State:

Set of pasts with same morph $\Pr(\overrightarrow{S} \mid \overleftarrow{s})$.

Set of histories that lead to same predictions.

Predictive equivalence relation:

$$\overleftarrow{s}' \sim \overleftarrow{s}'' \iff \Pr(\overrightarrow{S} \mid \overleftarrow{S} = \overleftarrow{s}') = \Pr(\overrightarrow{S} \mid \overleftarrow{S} = \overleftarrow{s}'')$$

$$\overleftarrow{s}', \overleftarrow{s}'' \in \overleftarrow{\mathbf{S}}$$

The Learning Channel ...

Causal State Components

Causal State = Pasts with same morph: $\Pr(\overrightarrow{S} \mid \overleftarrow{s})$

$$\mathcal{S} = \{ \overleftarrow{s}' : \overleftarrow{s}' \sim \overleftarrow{s} \}$$

Set of causal states:

$$\boldsymbol{\mathcal{S}} = \overleftarrow{\mathbf{S}} / \sim \;\; = \{ \mathcal{S}_0, \mathcal{S}_1, \mathcal{S}_2, \ldots \}$$

Causal state map:

$$\epsilon : \overleftarrow{\mathbf{S}} \rightarrow \boldsymbol{\mathcal{S}}$$

$$\epsilon(\overleftarrow{s}) = \{ \overleftarrow{s}' : \overleftarrow{s}' \sim \overleftarrow{s} \}$$

We've answered the first part of the modeling goal:

We have the effective states!

Now,

What is the dynamic?

The Learning Channel ...

Causal State Dynamic ...

Causal-state Filtering:

$$\overleftrightarrow{s} = \ldots \; s_{-3} \quad s_{-2} \quad s_{-1} \quad s_0 \quad s_1 \quad s_2 \quad s_3 \quad \ldots$$

$$\epsilon(\overleftrightarrow{s}) = \ldots \epsilon(\overleftarrow{s}_{-3})\epsilon(\overleftarrow{s}_{-2})\epsilon(\overleftarrow{s}_{-1})\epsilon(\overleftarrow{s}_{0})\epsilon(\overleftarrow{s}_{1})\epsilon(\overleftarrow{s}_{2})\epsilon(\overleftarrow{s}_{3}) \ldots$$

$$\overleftrightarrow{\mathcal{S}} = \ldots \quad \mathcal{S}_{t=-3} \; \mathcal{S}_{t=-2} \; \mathcal{S}_{t=-1} \; \mathcal{S}_{t=0} \; \mathcal{S}_{t=1} \; \mathcal{S}_{t=2} \; \mathcal{S}_{t=3} \quad \ldots$$

Causal-state process:

$$\Pr(\overleftrightarrow{\mathcal{S}})$$

The Learning Channel ...

Causal State Dynamic ...

Conditional transition probability:

$$
\begin{aligned}
T_{ij}^{(s)} &= \Pr(\mathcal{S}_j, s | \mathcal{S}_i) \\
&= \Pr\left(\mathcal{S} = \epsilon(\overleftarrow{s}\, s) | \mathcal{S} = \epsilon(\overleftarrow{s})\right)
\end{aligned}
$$

State-to-State Transitions:

$$\{T_{ij}^{(s)} : s \in \mathcal{A},\ i, j = 0, 1, \ldots, |\boldsymbol{\mathcal{S}}|\}$$

The ϵ-Machine ...

$$\text{Process} \;\Rightarrow\; \text{Predictive equivalence} \;\Rightarrow\; \epsilon - \text{Machine}$$

$$\Pr(\overleftrightarrow{S}) \;\Rightarrow\; \overleftarrow{\mathbf{S}} \,/ \sim \;\Rightarrow\; \epsilon - \text{Machine}$$

$$\mathcal{M} = \left\{ \boldsymbol{\mathcal{S}}, \{T^{(s)}, s \in \mathcal{A}\} \right\}$$

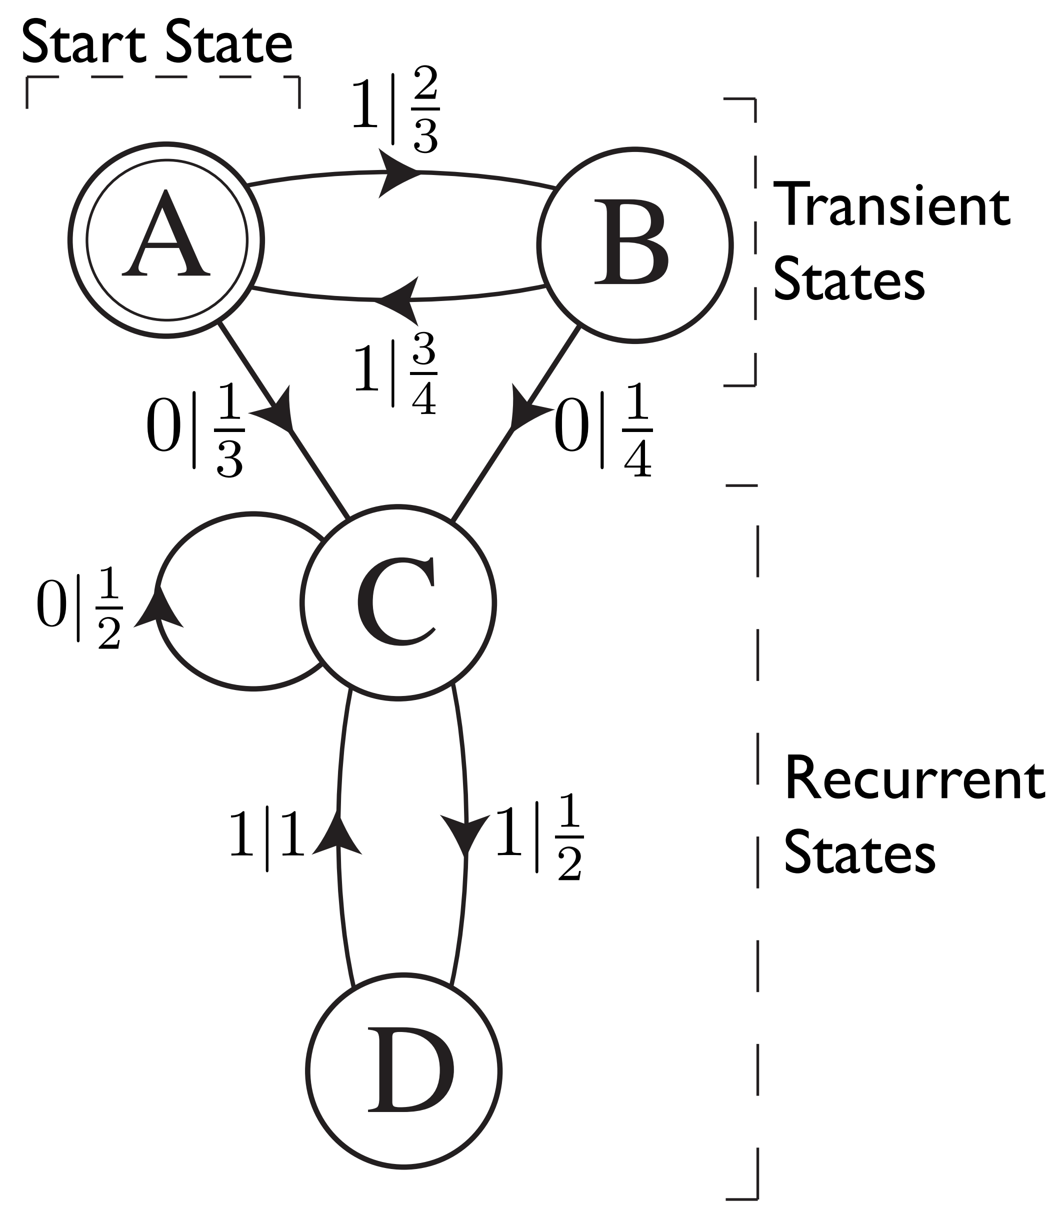

Unique Start State:

$$\mathcal{S}_0 = [\lambda]$$

$$\Pr(\mathcal{S}_0, \mathcal{S}_1, \mathcal{S}_2, \ldots) = (1, 0, 0, \ldots)$$

Transient States

Recurrent States

The Learning Channel ...

The ϵ-Machine of a Process ...

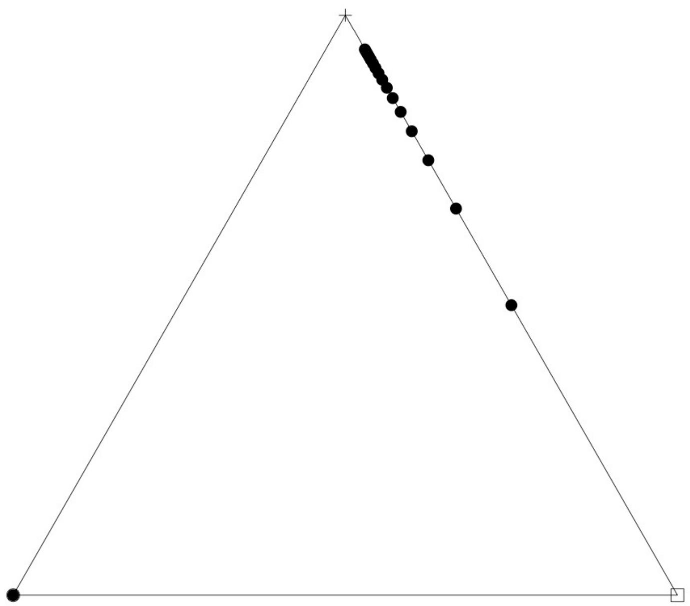

Denumerable Causal States

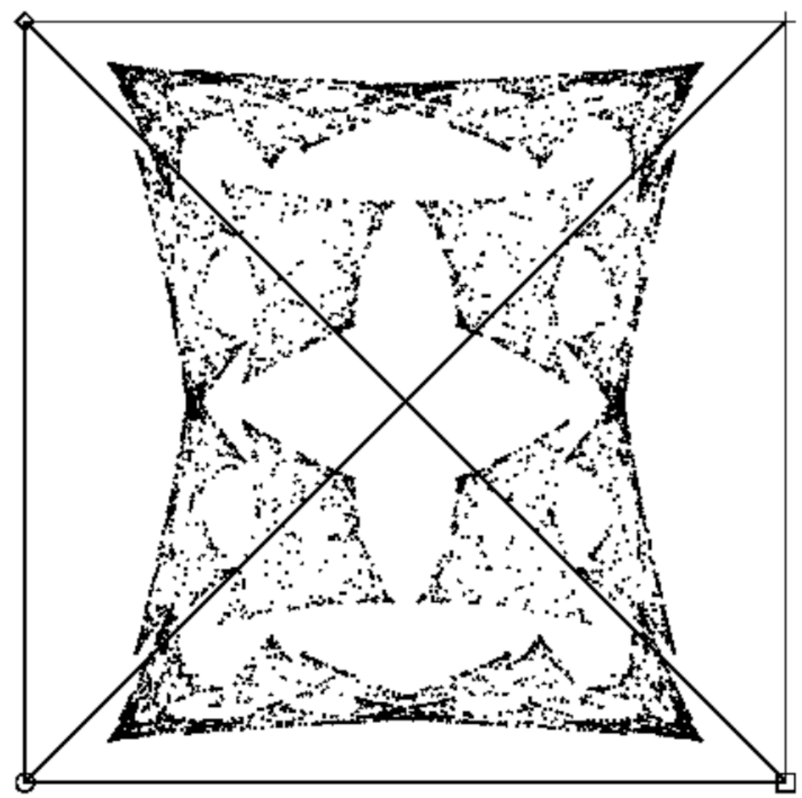

Fractal

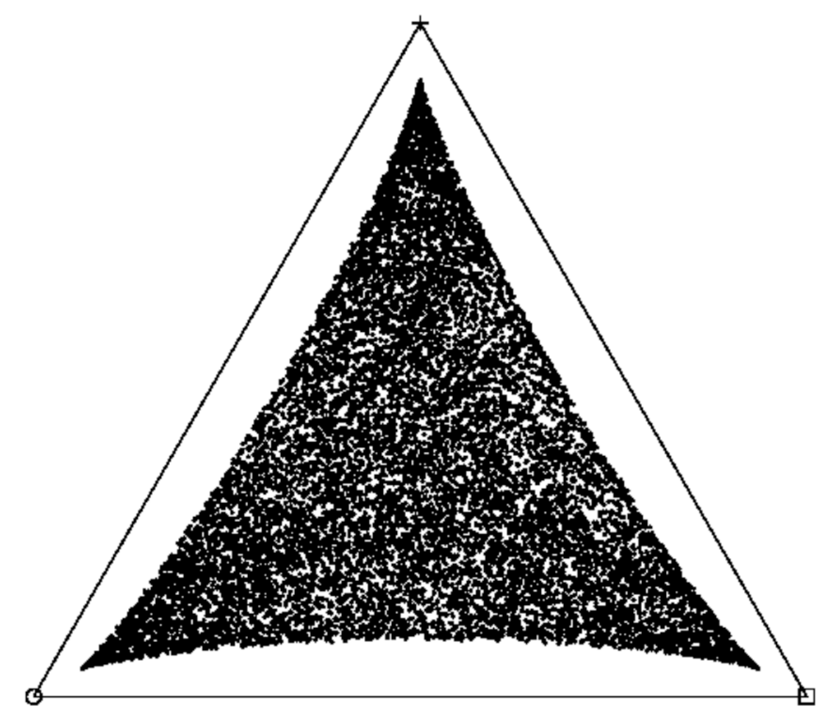

Continuous

The Learning Channel:

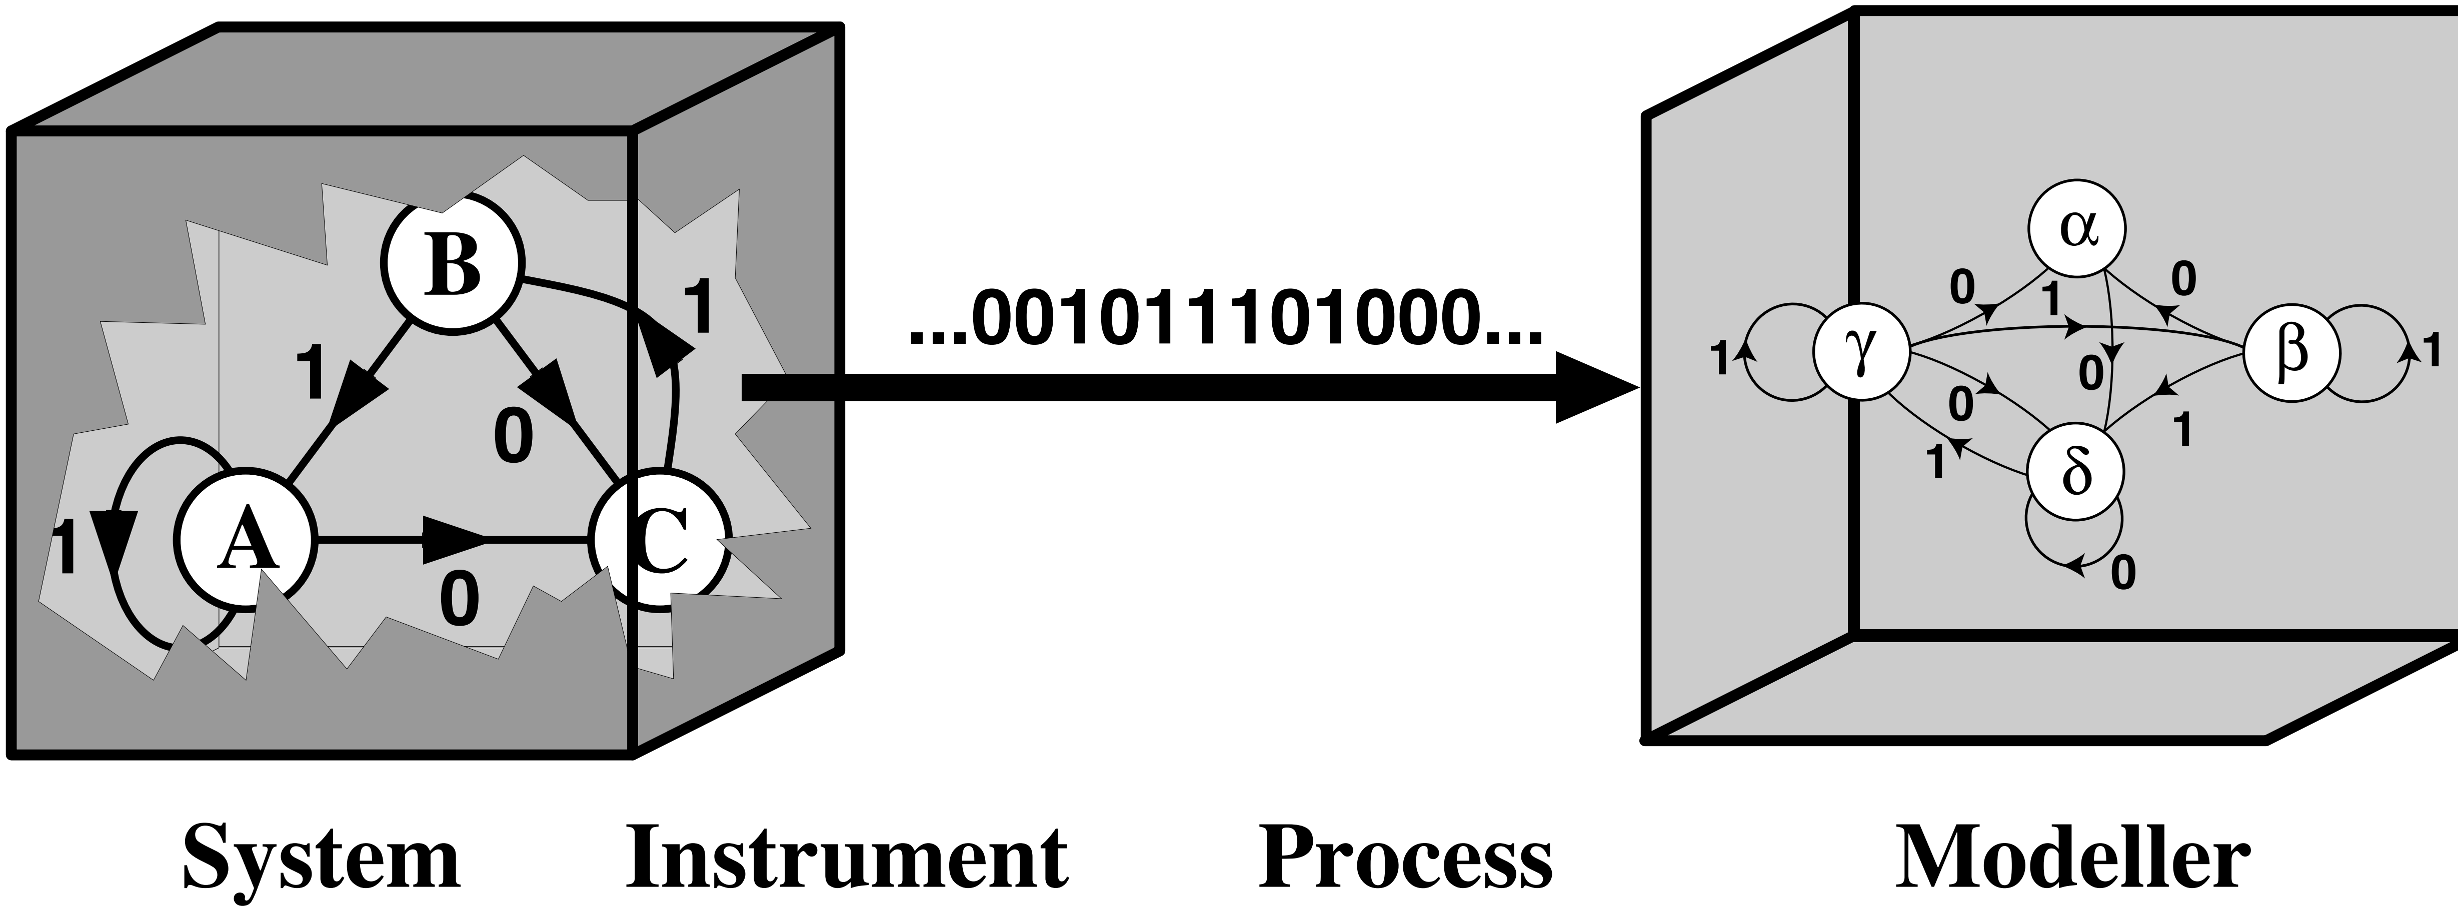

Central questions:

What are the states? Causal States

What is the dynamic? The ϵ-Machine

The ϵ-Machine

The ϵ-Machine ...

A Model of a Process $\Pr(\overleftrightarrow{S})$:

ϵ-Machine reproduces the process's word distribution:

$$\Pr(s^1), \Pr(s^2), \Pr(s^3), \ldots$$

$$s^L = s_1 s_2 \ldots s_L \qquad \mathcal{S}(t=0) = \mathcal{S}_0$$

$$\Pr(s^L) = \Pr(\mathcal{S}_0)\Pr(\mathcal{S}_0 \rightarrow_{s=s_1} \mathcal{S}(1))\Pr(\mathcal{S}(1) \rightarrow_{s=s_2} \mathcal{S}(2)) \cdots \Pr(\mathcal{S}(L-1) \rightarrow_{s=s_L} \mathcal{S}(L))$$

Initially, $\Pr(\mathcal{S}_0) = 1.$

$$\Pr(s^L) = \prod_{l=1}^{L} T^{(s_l)}_{i=\epsilon(s^{l-1}),j=\epsilon(s^l)}$$

The ϵ-Machine ...

Causal shielding:

Past and future are independent given causal state

Process: $\Pr(\overleftrightarrow{S}) = \Pr(\overleftarrow{S}\overrightarrow{S})$

$$\Pr(\overleftarrow{S}\overrightarrow{S} \,|\mathcal{S}) = \Pr(\overleftarrow{S} \,|\mathcal{S}) \Pr(\overrightarrow{S} \,|\mathcal{S})$$

Causal states shield past & future from each other.

Similar to states of a Markov chain, but for hidden processes.

The ϵ-Machine ...

$\epsilon\mathrm{Ms}$ are Unifilar: $(\mathcal{S}_t, s) \rightarrow$ unique $\mathcal{S}_{t+1}$

(in automata theory, "deterministic")

Consequence:

Unifilarity: 1-1 map between state-sequences & symbol-sequences.

Entropy rate expression requires this 1-1 mapping.

Can use $\epsilon\mathrm{M}$ to calculate entropy rate h_μ.

(Any unifilar presentation will do.)

The ϵ-Machine ...

ϵMs are Optimal Predictors:

Compared to any rival effective states R:

$$H\left[\overrightarrow{S}^{L} | R\right] \geq H\left[\overrightarrow{S}^{L} | \mathcal{S}\right]$$

The ϵ-Machine ...

Prescient Rivals $\widehat{\mathbf{R}}$:

Alternative models that are optimal predictors

$$\widehat{R} \in \widehat{\mathbf{R}}$$

$$H[\overrightarrow{S}^L | \widehat{R}] = H[\overrightarrow{S}^L | \mathcal{S}]$$

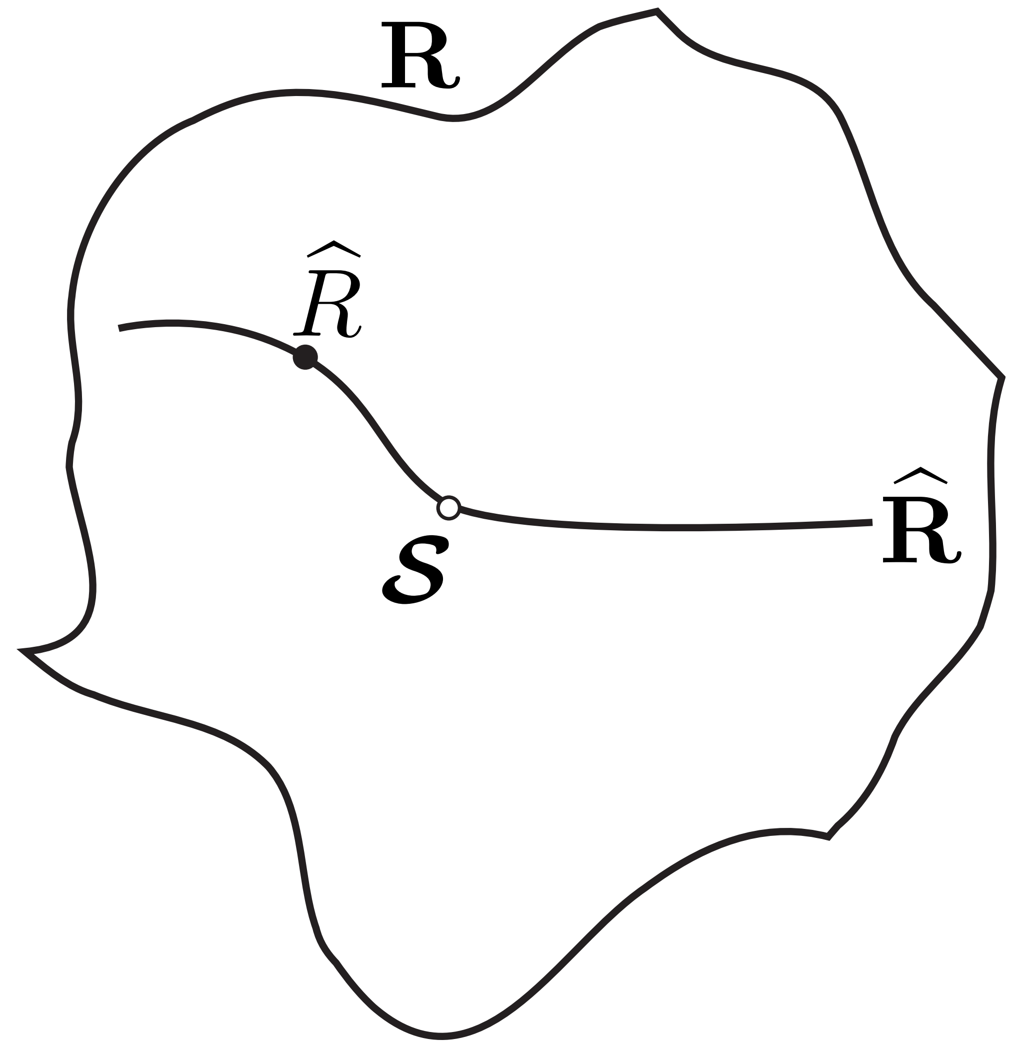

(Prescient rivals are sufficient statistics for process's future.)

The ϵ-Machine ...

Minimal Statistical Complexity:

For all prescient rivals, $\epsilon\mathrm{M}$ is the smallest:

$$C_\mu(\widehat{R}) \geq C_\mu(\mathcal{S})$$

Consequence:

(1) C_μ measures historical information process stores.

(2) This would not be true, if not minimal representation.

Remarks:

(1) Causal states contain every difference (in past) that makes a difference (to future) (Bateson "information")

(2) Causal states are sufficient statistics for the future.

The ϵ-Machine ...

Summary:

$\epsilon\mathrm{M}$:

(1) Optimal predictor: Lower prediction error than any rival.

(2) Minimal size: Smallest of the prescient rivals.

(3) Unique: Smallest, optimal, unifilar predictor is equivalent.

(4) Model of the process: Reproduces all of process's statistics.

(5) Causal shielding: Renders process's future independent of past.

Measures of Intrinsic Computation

Measures of Intrinsic Computation ...

A complex process's intrinsic computation:

(1) How much of past does process store?

$$C_\mu$$

(2) In what architecture is that information stored?

$$\left\{ \boldsymbol{\mathcal{S}}, \{T^{(s)}, s \in \mathcal{A}\} \right\}$$

(3) How is stored information used to produce future behavior?

$$h_\mu$$

Measures of Intrinsic Computation ...

Measures of Structural Complexity:

Information Measures		Interpretation
Entropy Rate	h_μ	Intrinsic Randomness
Excess Entropy	$\mathbf{E}$	Info: Past to Future
Total Predictability	$\mathbf{G}$	Redundancy
Transient Information	$\mathbf{T}$	Synchronization

How related to statistical complexity C_μ?

How to get from $\epsilon\mathrm{M}$?

Measures of Intrinsic Computation ...

Measures from the $\epsilon\mathrm{M}$:

Entropy Rate of a Process:

$$h_\mu(\Pr(\overleftrightarrow{S})) = \lim_{L \to \infty} \frac{H(L)}{L}$$

Entropy Rate given $\epsilon\mathrm{M}$:

$$h_\mu(\boldsymbol{\mathcal{S}}) = - \sum_{\mathcal{S} \in \boldsymbol{\mathcal{S}}} \Pr(\mathcal{S}) \sum_{s \in \mathcal{A}, \mathcal{S}' \in \boldsymbol{\mathcal{S}}} T^{(s)}_{\mathcal{S}\mathcal{S}'} \log_2 T^{(s)}_{\mathcal{S}\mathcal{S}'}$$

where $\Pr(\mathcal{S})$ is casual-state asymptotic probability.

Possible only due to $\epsilon\mathrm{M}$ unifilarity!

1-1 mapping between measurement sequences & internal paths.

Measures of Intrinsic Computation ...

Measures from the $\epsilon\mathrm{M}$...

Statistical Complexity of a Process:

$$C_\mu(\boldsymbol{\mathcal{S}}) = -\sum_{\mathcal{S} \in \boldsymbol{\mathcal{S}}} \Pr(\mathcal{S}) \log_2 \Pr(\mathcal{S})$$

where $\Pr(\mathcal{S})$ is casual-state asymptotic probability.

Meaning:

Shannon information in the causal states.

The amount of historical information a process stores.

The amount of structure in a process.

Measures of Intrinsic Computation ...

Measures from the $\epsilon\mathrm{M}$...

Excess Entropy: Three versions, all equivalent for 1D processes

$$\mathbf{E} = \lim_{L \to \infty} [H(L) - h_\mu L]$$

$$\mathbf{E} = \sum_{L=1}^{\infty} [h_\mu(L) - h_\mu]$$

$$\mathbf{E} = I[\overleftarrow{S}; \overrightarrow{S}]$$

How to get, given $\epsilon\mathrm{M}$?

Special cases: When $\epsilon\mathrm{M}$ is IID, periodic, or spin chain.

General case: Need a new framework.

Measures of Intrinsic Computation ...

Measures from the $\epsilon\mathrm{M}$ ...

Bound on Excess Entropy:

$$\mathbf{E} \leq C_\mu$$

Proof sketch:

(1) $\mathbf{E} = I[\overrightarrow{S}; \overleftarrow{S}] = H[\overrightarrow{S}] - H[\overrightarrow{S} \mid \overleftarrow{S}]$

(2) Causal States: $H[\overrightarrow{S} \mid \overleftarrow{S}] = H[\overrightarrow{S} | \mathcal{S}]$

(3) $\mathbf{E} = H[\overrightarrow{S}] - H[\overrightarrow{S} | \mathcal{S}]$

$$= I[\overrightarrow{S}; \mathcal{S}]$$

$$= H[\mathcal{S}] - H[\mathcal{S} | \overrightarrow{S}]$$

$$\leq H[\mathcal{S}] = C_\mu$$

Measures of Intrinsic Computation ...

Measures from the ϵM ...

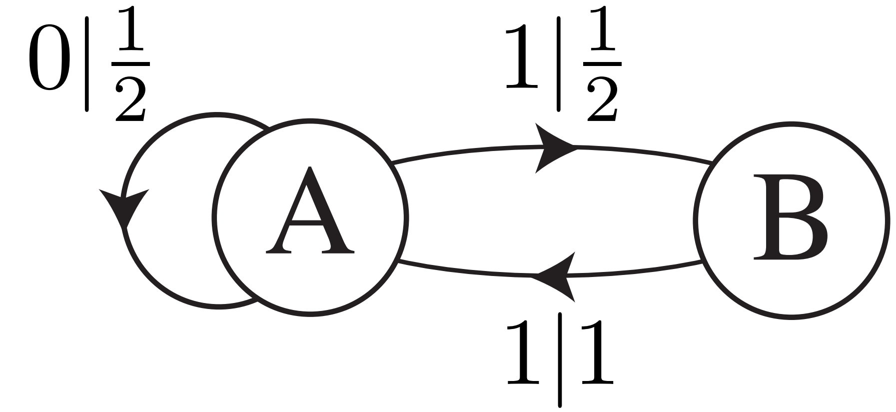

Bound on Excess Entropy ...

But, the bound is saturated!

$$T^{(0)} = \begin{pmatrix} \frac{1}{2} & 0 \\ 0 & 0 \end{pmatrix}$$

Even Process:

$$T^{(1)} = \begin{pmatrix} 0 & \frac{1}{2} \\ 1 & 0 \end{pmatrix}$$

$$C_\mu = H(2/3) \approx 0.9182$$

$$\mathbf{E} \approx 0.9182$$

$$\pi_V = (2/3, 1/3)$$

When does this occur?

In general, need a new framework for answering this question: Directional Computational Mechanics.

Measures of Intrinsic Computation ...

Measures from the $\epsilon\mathrm{M}$ ...

Bound on Excess Entropy ...

Consequence: The Cryptographic Limit

Can have $\mathbf{E} \to 0$ when $C_\mu \gg 1$.

Excess entropy is *not* the process's stored information.

$\mathbf{E}$ is the *apparent* information,
as revealed in *measurement sequences*.

Statistical complexity *is* stored information.

Measures of Intrinsic Computation ...

Measures from the ϵM ...

Bound on Excess Entropy ...

Executive Summary:

C_μ is the amount of information the process uses

to *communicate*

$\mathbf{E}$ bits of information from the past to the future.

Measures of Intrinsic Computation ...

Measures from the $\epsilon\mathrm{M}$ ...

Bound on Excess Entropy: $\mathbf{E} \leq C_\mu$

Consequence:

The inequality is Why We Must Model.

Cannot simply use sequences as states.

There is internal structure not expressed by this.

Information Diagrams for Processes

Information Diagrams for Processes

Process I-diagram:

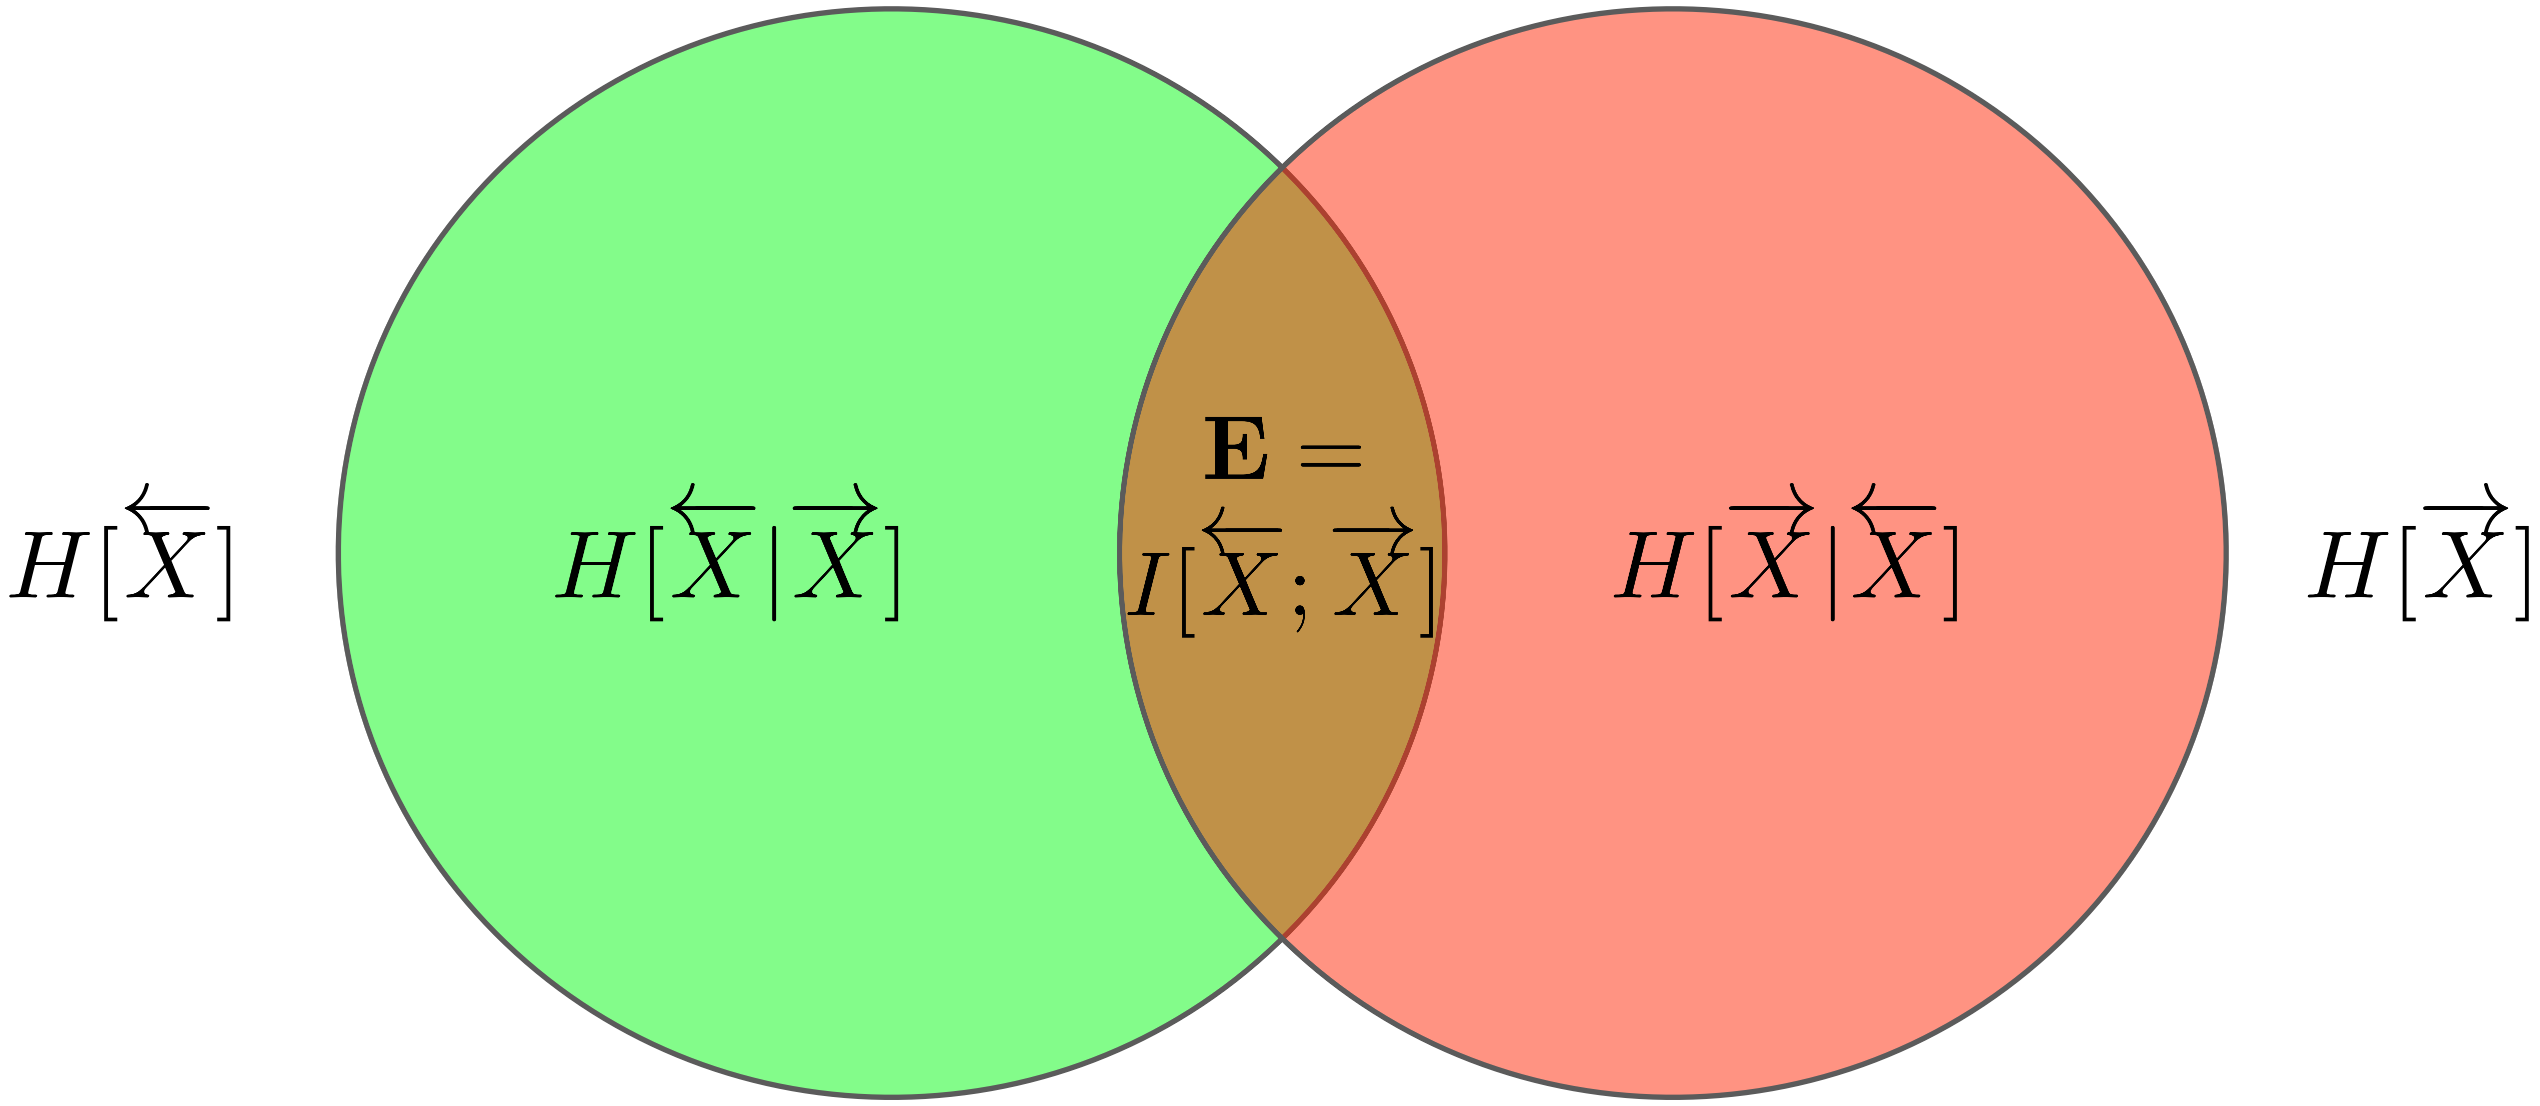

Information Diagrams for Processes

Process I-diagram using ε-machine:

Start with 3-variable I-diagram and whittle down:

Past as composite random variable: $\overleftarrow{X}$

Future as composite random variable: $\overrightarrow{X}$

Causal states: $\mathcal{S} \in \boldsymbol{\mathcal{S}}$

Information measures:

$$H[\overleftarrow{X}] \quad H[\overrightarrow{X}] \quad H[\mathcal{S}] \quad \cdots \quad I[\overrightarrow{X};\overleftarrow{X};\mathcal{S}] \quad \cdots \quad H[\overrightarrow{X},\overleftarrow{X},\mathcal{S}]$$

There are $8 = 2^3$ atomic information measures.

Information Diagrams for Processes

Process I-diagram using ε-machine ...

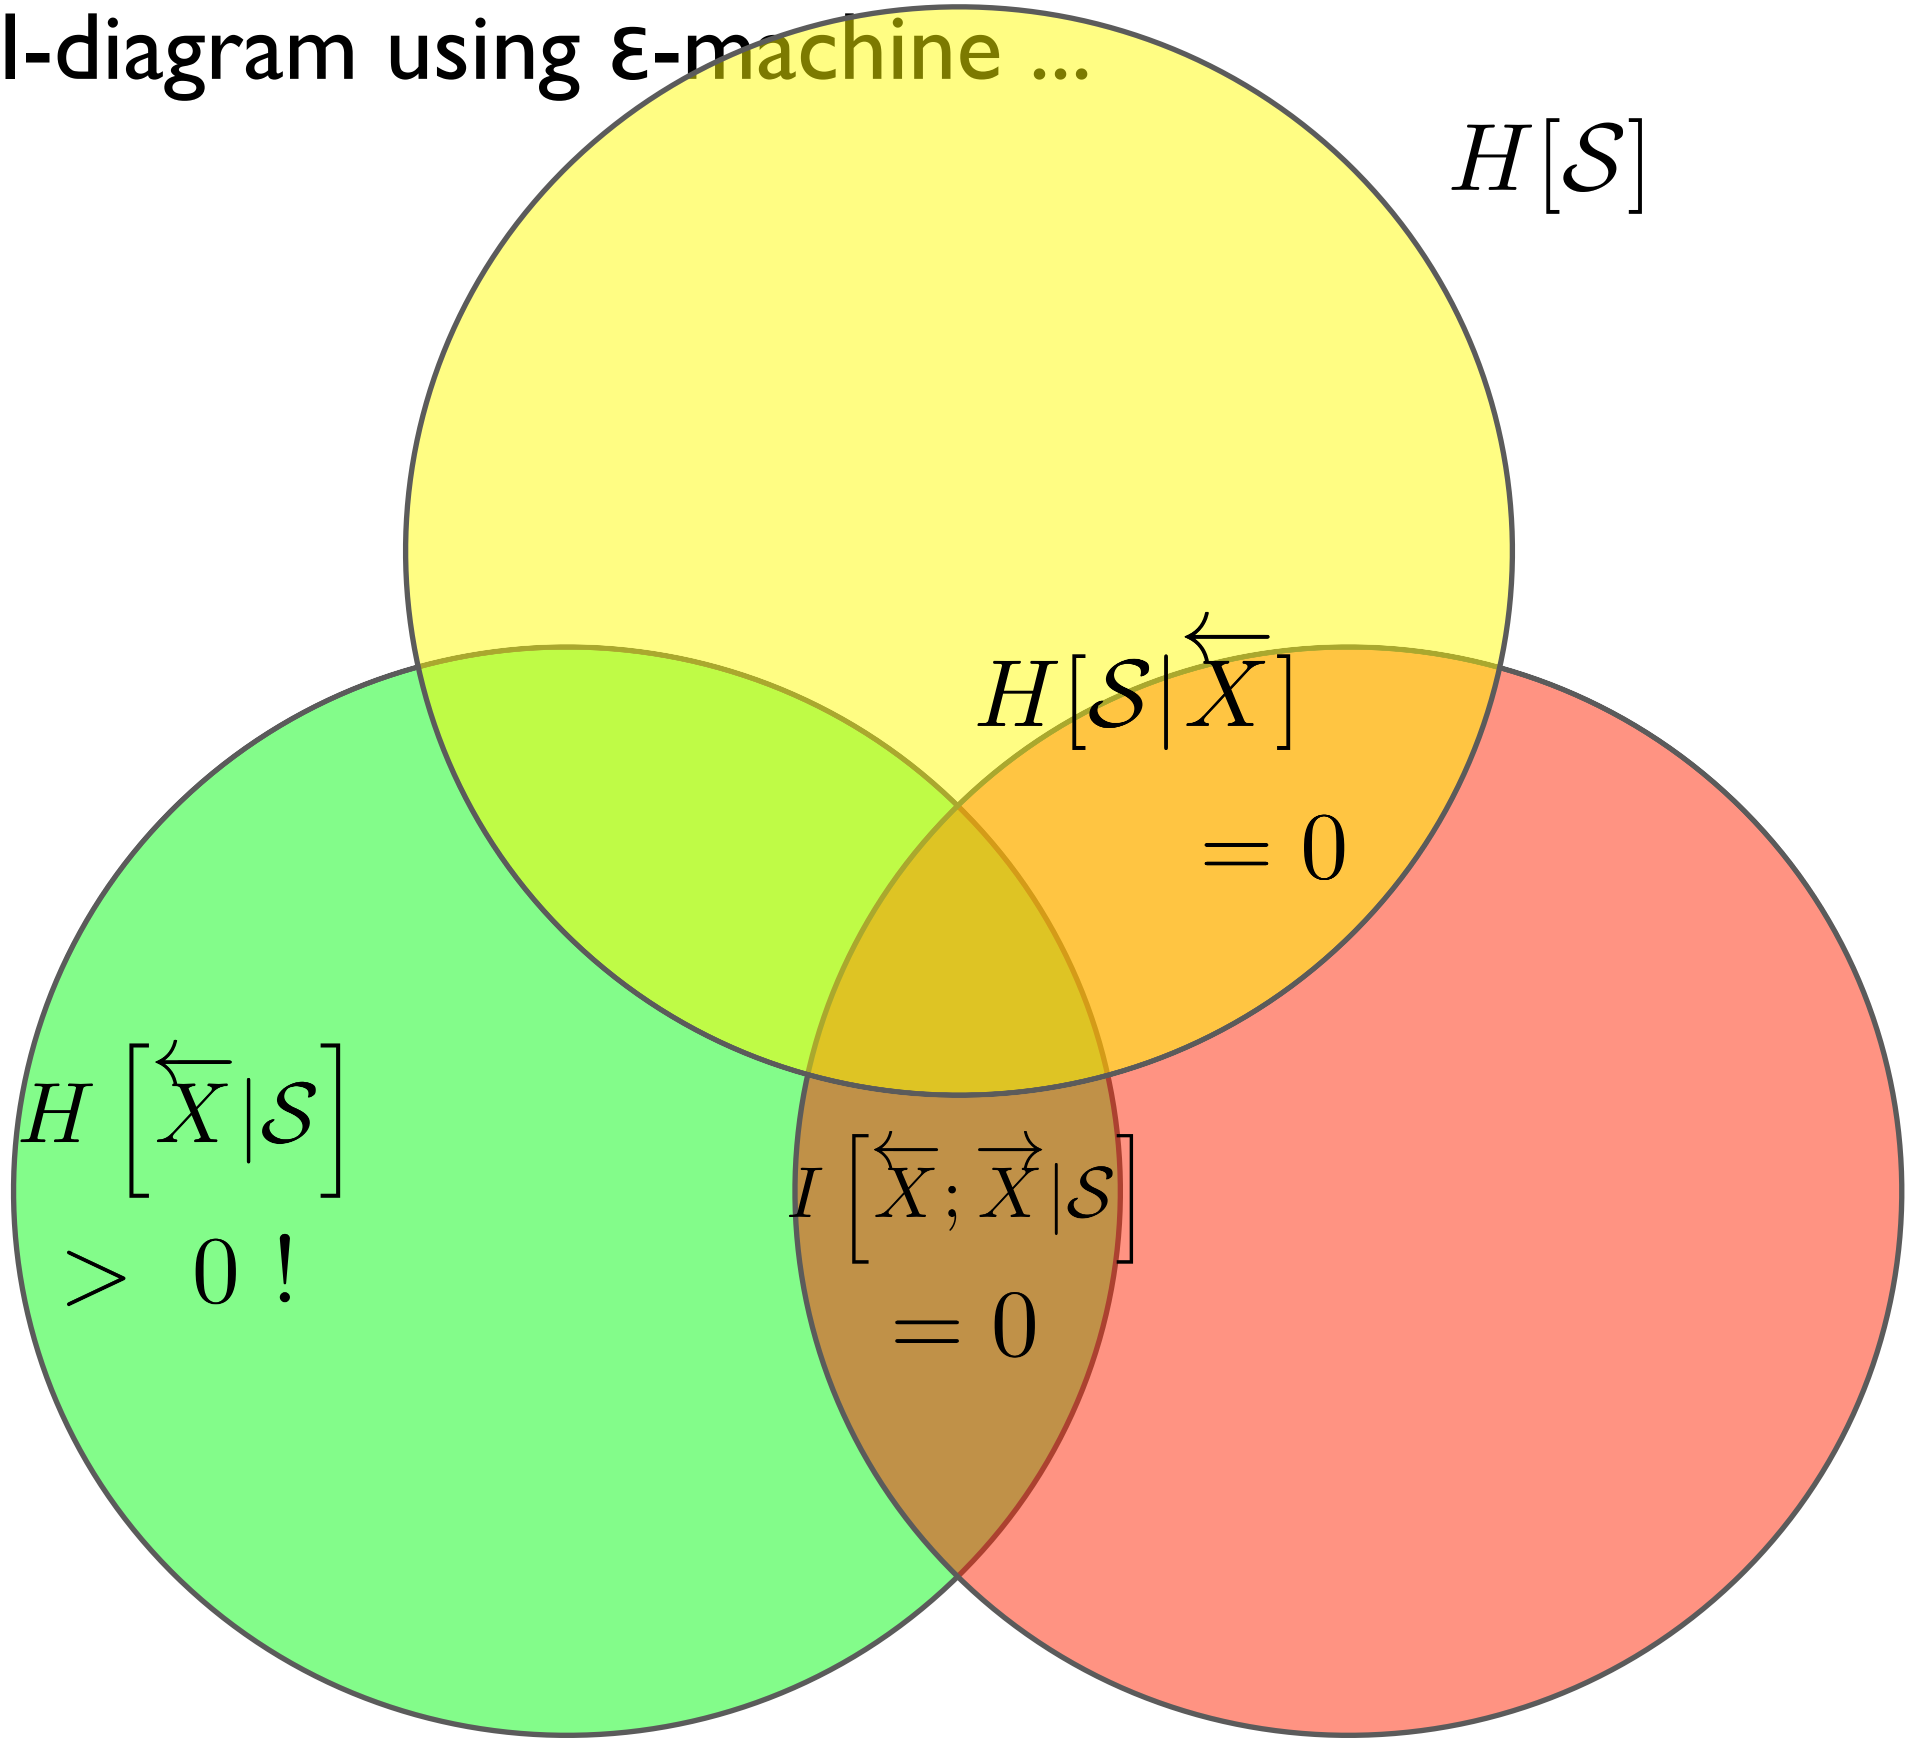

Information Diagrams for Processes

ε-Machine I-diagram:

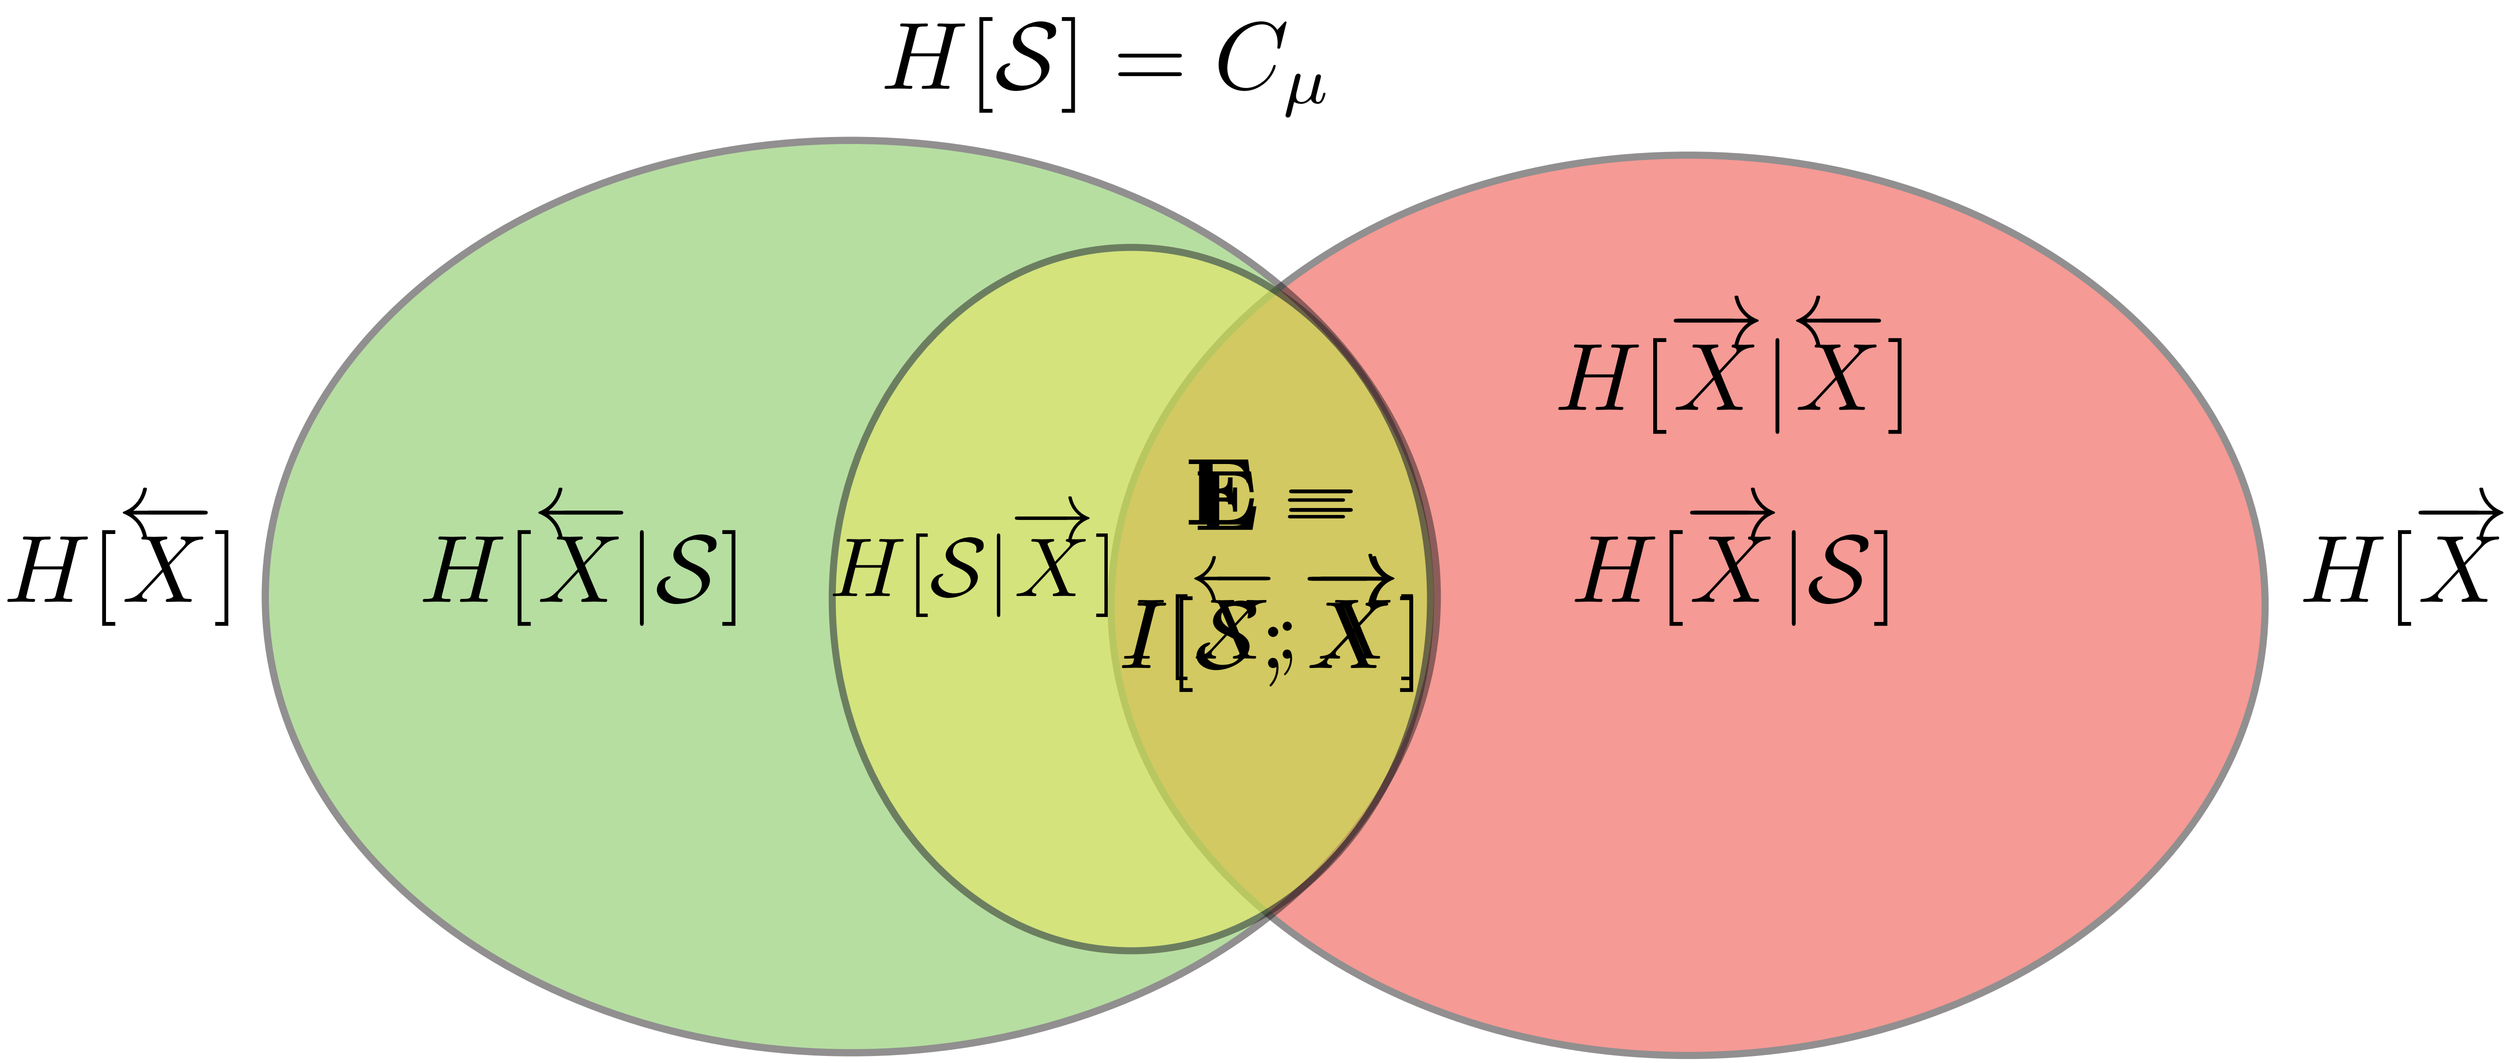

What is this Mystery Wedge?

What do we know about this?

Information Diagrams for Processes

What is $H[\overrightarrow{X}|\mathcal{S}]$?

Unpredictability: $H[\overrightarrow{X}^L|\mathcal{S}] = L h_\mu$

Proof Sketch:

$$\begin{aligned}
H[\overrightarrow{X}^L|\mathcal{S}] &= H[\overrightarrow{X}^L|\overleftarrow{X}] \\
&= H[X_0 X_1 \ldots X_{L-1}|\overleftarrow{X}] \\
&= H[X_1 \ldots X_{L-1}|\overleftarrow{X} X_0] + H[X_0|\overleftarrow{X}] \\
&= H[X_1 \ldots X_{L-1}|\overleftarrow{X}] + H[X_0|\overleftarrow{X}] \\
&\vdots \\
&= H[X_{L-1}|\overleftarrow{X}] + \cdots + H[X_1|\overleftarrow{X}] + H[X_0|\overleftarrow{X}] \\
&= L H[X_0|\overleftarrow{X}] \\
&= L h_\mu
\end{aligned}$$

Information Diagrams for Processes

What is Mystery Wedge? $H[\mathcal{S}|\overrightarrow{X}]$

Uncertainty of causal state given future. Implications?

Recall Bound on Excess Entropy: $\mathbf{E} \leq C_\mu$

Proof sketch:
$$
\begin{aligned}
\mathbf{E} &= I[\overleftarrow{X};\overrightarrow{X}] \\
&= H[\overrightarrow{X}] - H[\overrightarrow{X}|\overleftarrow{X}] \\
&= H[\overrightarrow{X}] - H[\overrightarrow{X}|\mathcal{S}] \\
&= I[\overrightarrow{X};\mathcal{S}] \\
&= H[\mathcal{S}] - H[\mathcal{S}|\overrightarrow{X}] \\
&\leq H[\mathcal{S}] \\
&= C_\mu
\end{aligned}
$$

I am the Mystery Wedge!

Information Diagrams for Processes

What is Mystery Wedge? $H[\mathcal{S}|\overrightarrow{X}]$

Wedge is the inaccessibility of hidden state information!

$$H[\mathcal{S}|\overrightarrow{X}] = C_\mu - \mathbf{E}$$

Wedge controls Internal - Observed

The Process Crypticity:

$$\chi = C_\mu - \mathbf{E}$$

Controls how much internal state information is observable.

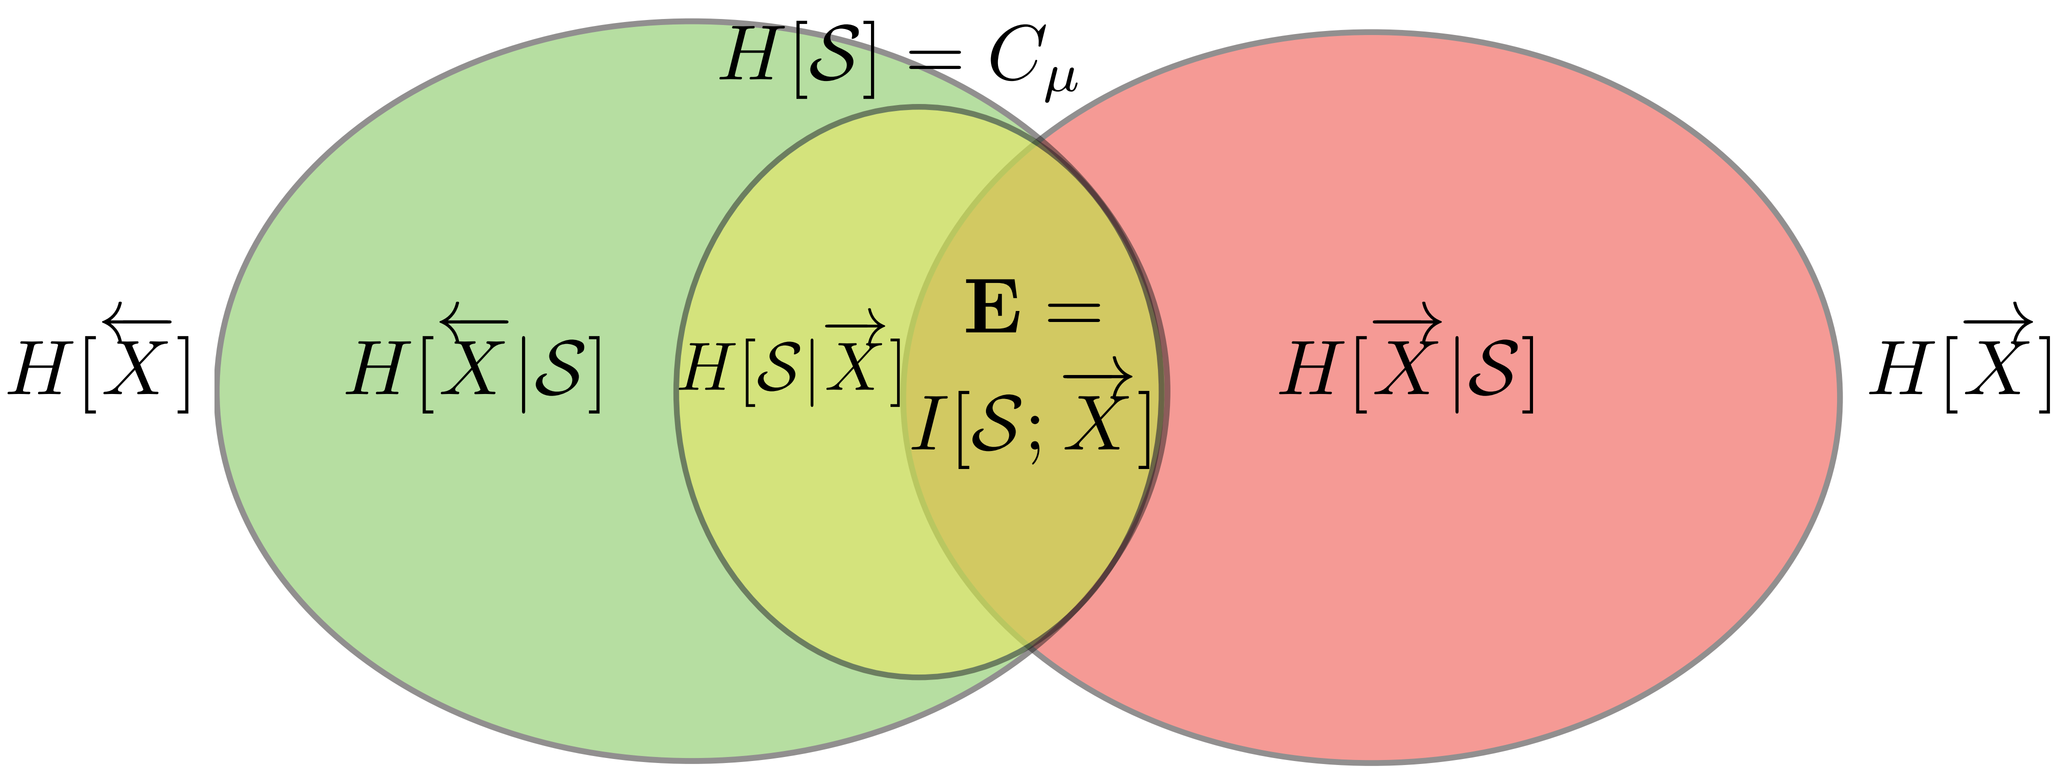

How to get $\mathbf{E}$ from ϵM?

Directional Computational Mechanics

- Theorem:

$$\mathbf{E} = I[\mathcal{S}^+ ; \mathcal{S}^-]$$

- Effective transmission capacity of channel between forward and reverse processes.
- Time agnostic representation: The **BiMachine**.

J. P. Crutchfield, C. J. Ellison, and J. R. Mahoney, "Time's Barbed Arrow: Irreversibility, Crypticity, and Stored Information", Physical Review Letters **103**:9 (2009) 094101.

Forward-Reverse ε-Machine Information Diagram

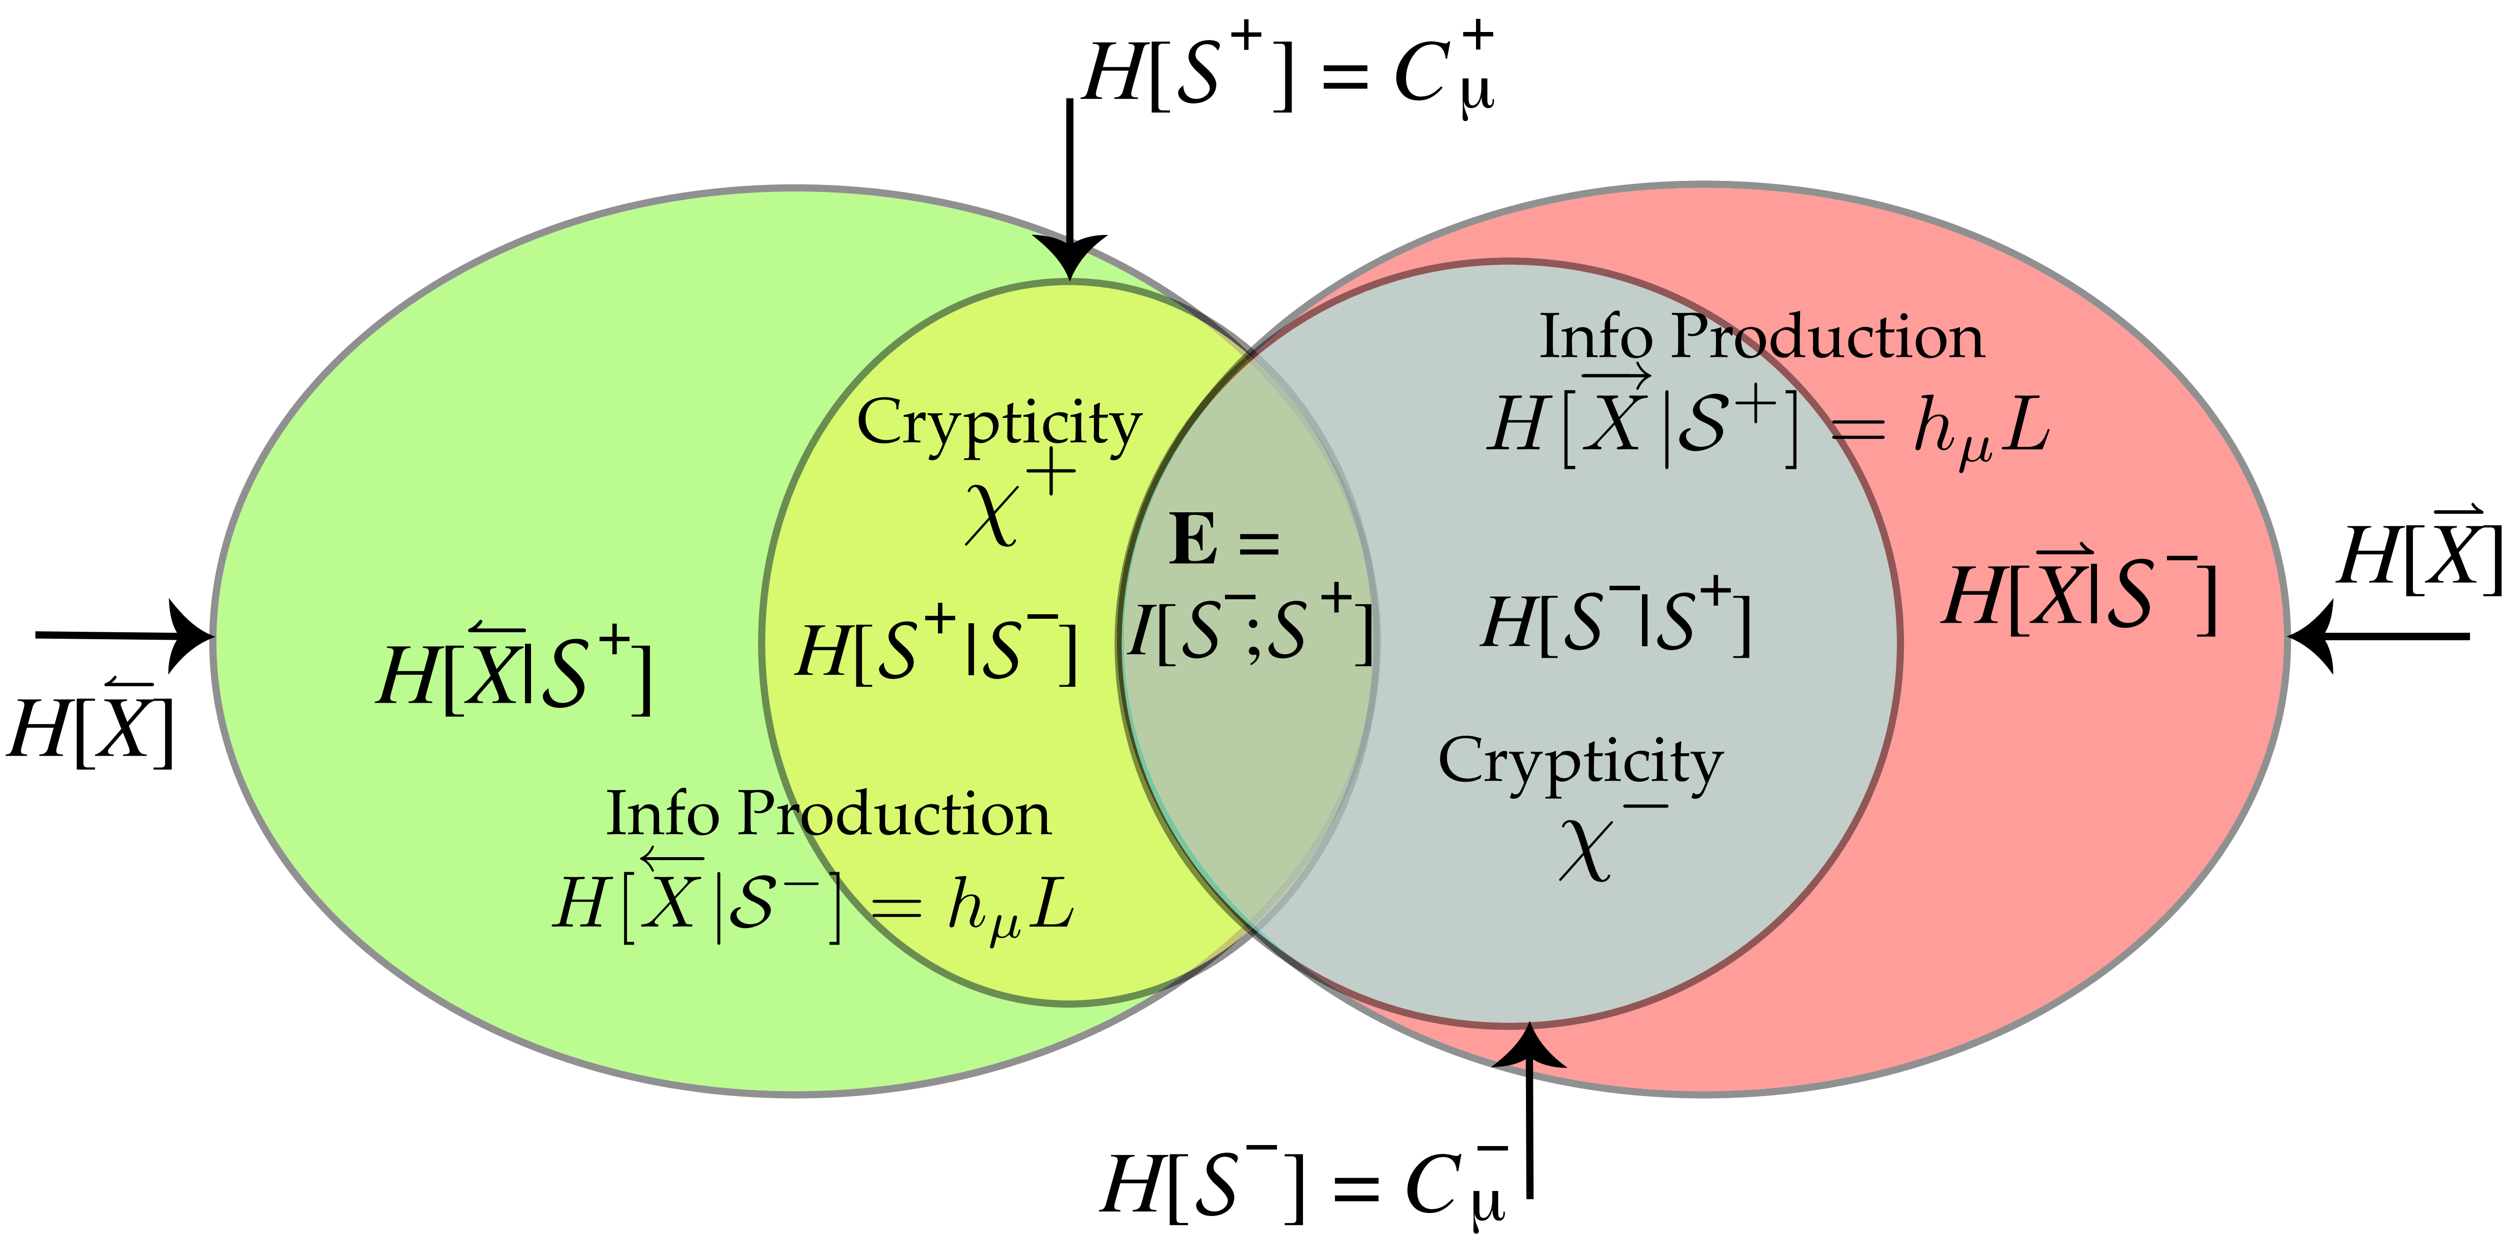

Intrinsic Computation ...

Analysis narrative:

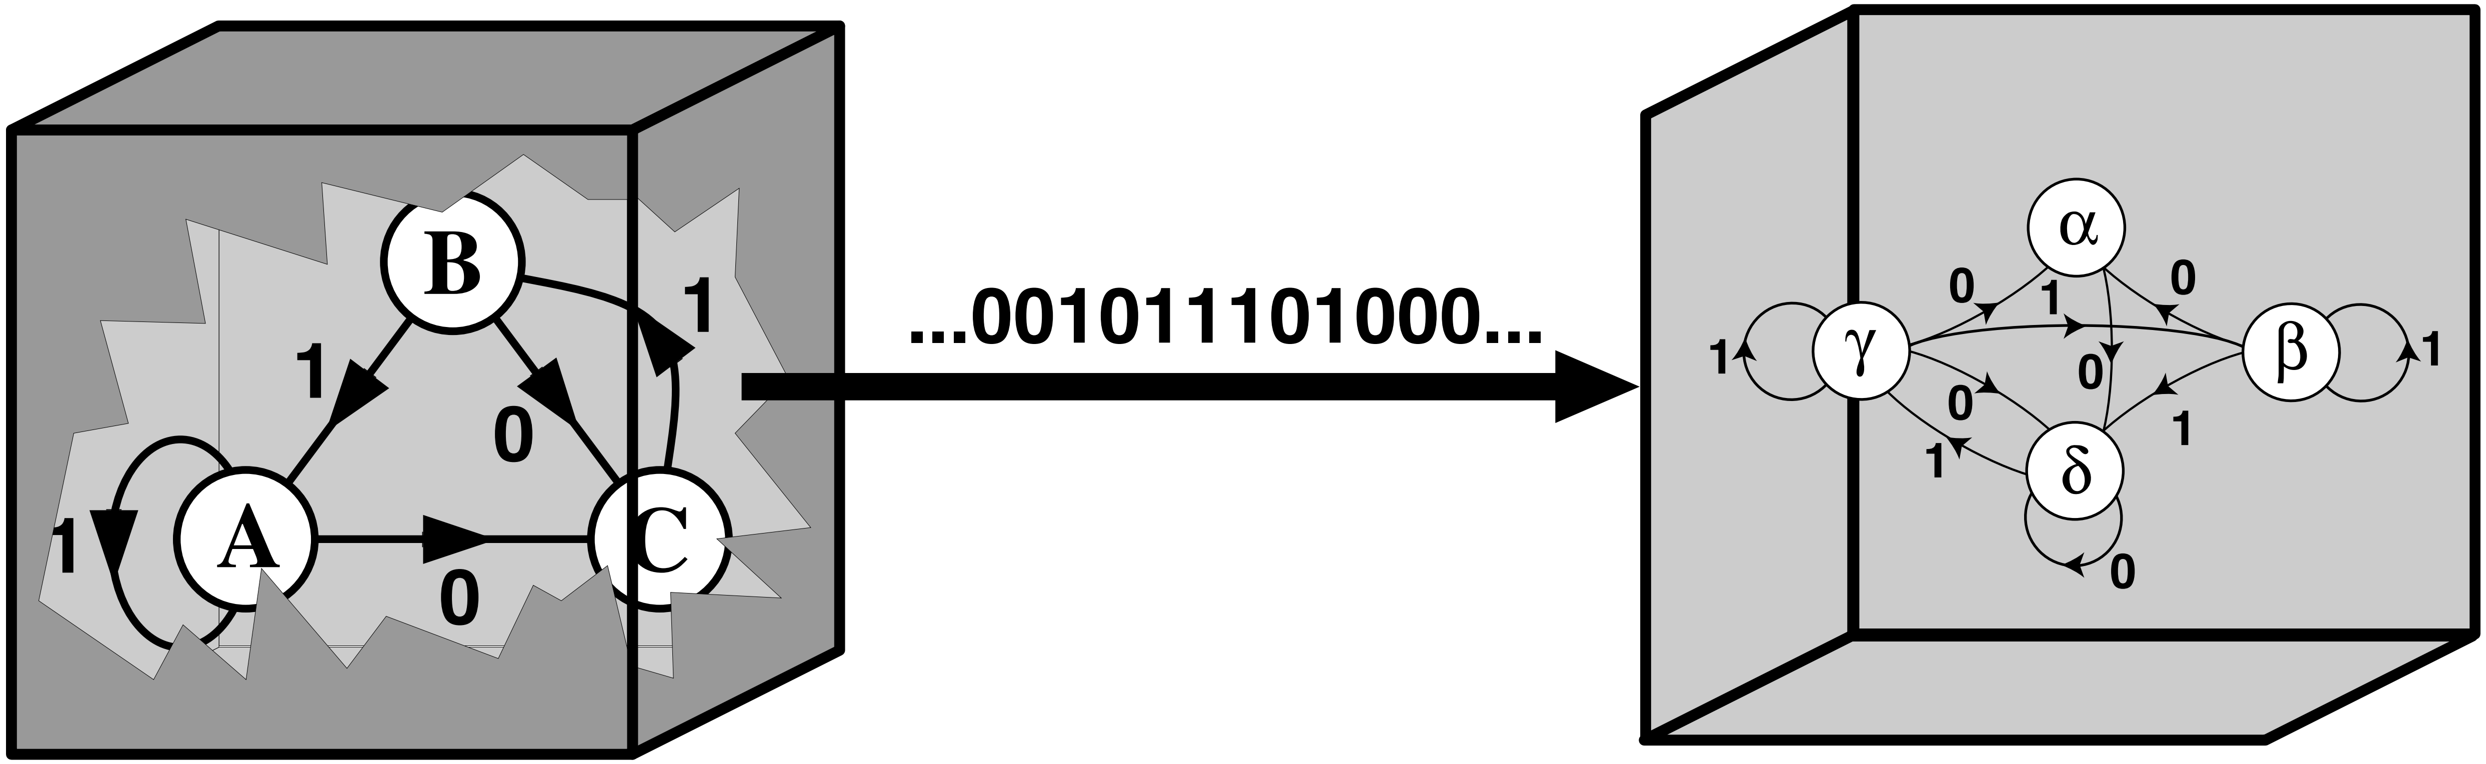

System **Instrument** **Process** **Modeller**

Forms of Chaos:
- Deterministic sources of novelty
- Mechanisms that produce unpredictability
- Sensitive dependence on initial condition
- Sensitive dependence on parameter

Measurement Theory:
- Partitions
- Optimal Instrument:

$$\max_{\{\mathcal{P}\}} h_\mu$$

$$\min_{\{\mathcal{P}\}} C_\mu$$

How random?

$$\lambda, H(L), h_\mu$$

How structured?

$$C_\mu, \mathbf{E}, \mathbf{T}, \mathbf{G}, \mathcal{R}$$

Universal model:
- $\epsilon - \mathrm{Machine}$
- Pattern defined
- Causal Architecture
- Intrinsic Computation

Intrinsic Computation ...

A system is unpredictable
if it has positive entropy rate: $h_\mu > 0$

A system is complex
if it has positive structural complexity measures: $C_\mu > 0$

A system is emergent
if its structural complexity measures increase over time:

$$C_\mu(t') > C_\mu(t), \text{ if } t' > t$$

A system is hidden
if its crypticity is positive: $\chi > 0$

Algorithmic Basis of Information ...

Kolmogorov-Chaitin Complexity versus Statistical Complexity

KC Complexity versus Statistical Complexity

We saw that:
KC complexity of typical realizations from an information source grows proportional to the Shannon entropy rate:

$$K(x) \propto h_\mu |x|$$

Thus, KC complexity is a measure of randomness.

KC Complexity versus Statistical Complexity

What's the relationship to Statistical Complexity?

Since randomness drives Kolmogorov-Chaitin complexity, let's discount for generating randomness:

Programs consist of model m and data d (random part unexplained by m).

Sophistication of object:

$$S_k(x) = \min\{|m| : p = m + d \text{ and } |p| - K(x) \leq k\}$$

Also, uncomputable.

KC Complexity versus Statistical Complexity

Consider the average sophistication:

$$S(\ell) = \langle S_0(x_{0:\ell}) \rangle$$

It is statistical complexity:

$$C_\mu \propto_{\ell \gg 1} S(\ell)$$

Since program = model + data:

$$K(\ell) = S(\ell) + \langle |d| \rangle_{x_{0:\ell}}$$

We have:

$$K(\ell) \approx C_\mu + h_\mu \ell$$

Since a process has a structure, as ℓ gets large,
with probability 1 each possible $x_{0:\ell}$ has the same model.

KC Complexity versus Statistical Complexity

Recall the Block Entropy

$$H(\ell) \approx C_\mu + h_\mu \ell$$

Similar scaling.

$K(\ell)$ versus $H(\ell)$:

$\mathbf{E}$ quantifies the amount of information observed as ℓ gets large, whereas C_μ quantifies how much information it takes to predict as ℓ gets large.

KC Complexity versus Statistical Complexity

Kolmogorov-Chaitin Theory versus Computational Mechanics

First, ε−machine describes distribution over a system's behaviors, including individual realizations.

Second, one can exactly calculate the Shannon entropy rate for a system's behaviors.

Third, the computational model is a probabilistic UTM: a Bernoulli-Turing Machine.

KC Complexity versus Statistical Complexity

Computational Mechanics was introduced to be a calculable, quantitative version of KC Complexity Theory.

Constructive! For finite eMs, all complexity/information measures
- can be calculated in closed form.
- $O(1)$ computational complexity.

So, much computational complexity in KC Theory and in Information Theory obviated.

Thermodynamics of Adaptive Complex Systems: An Application

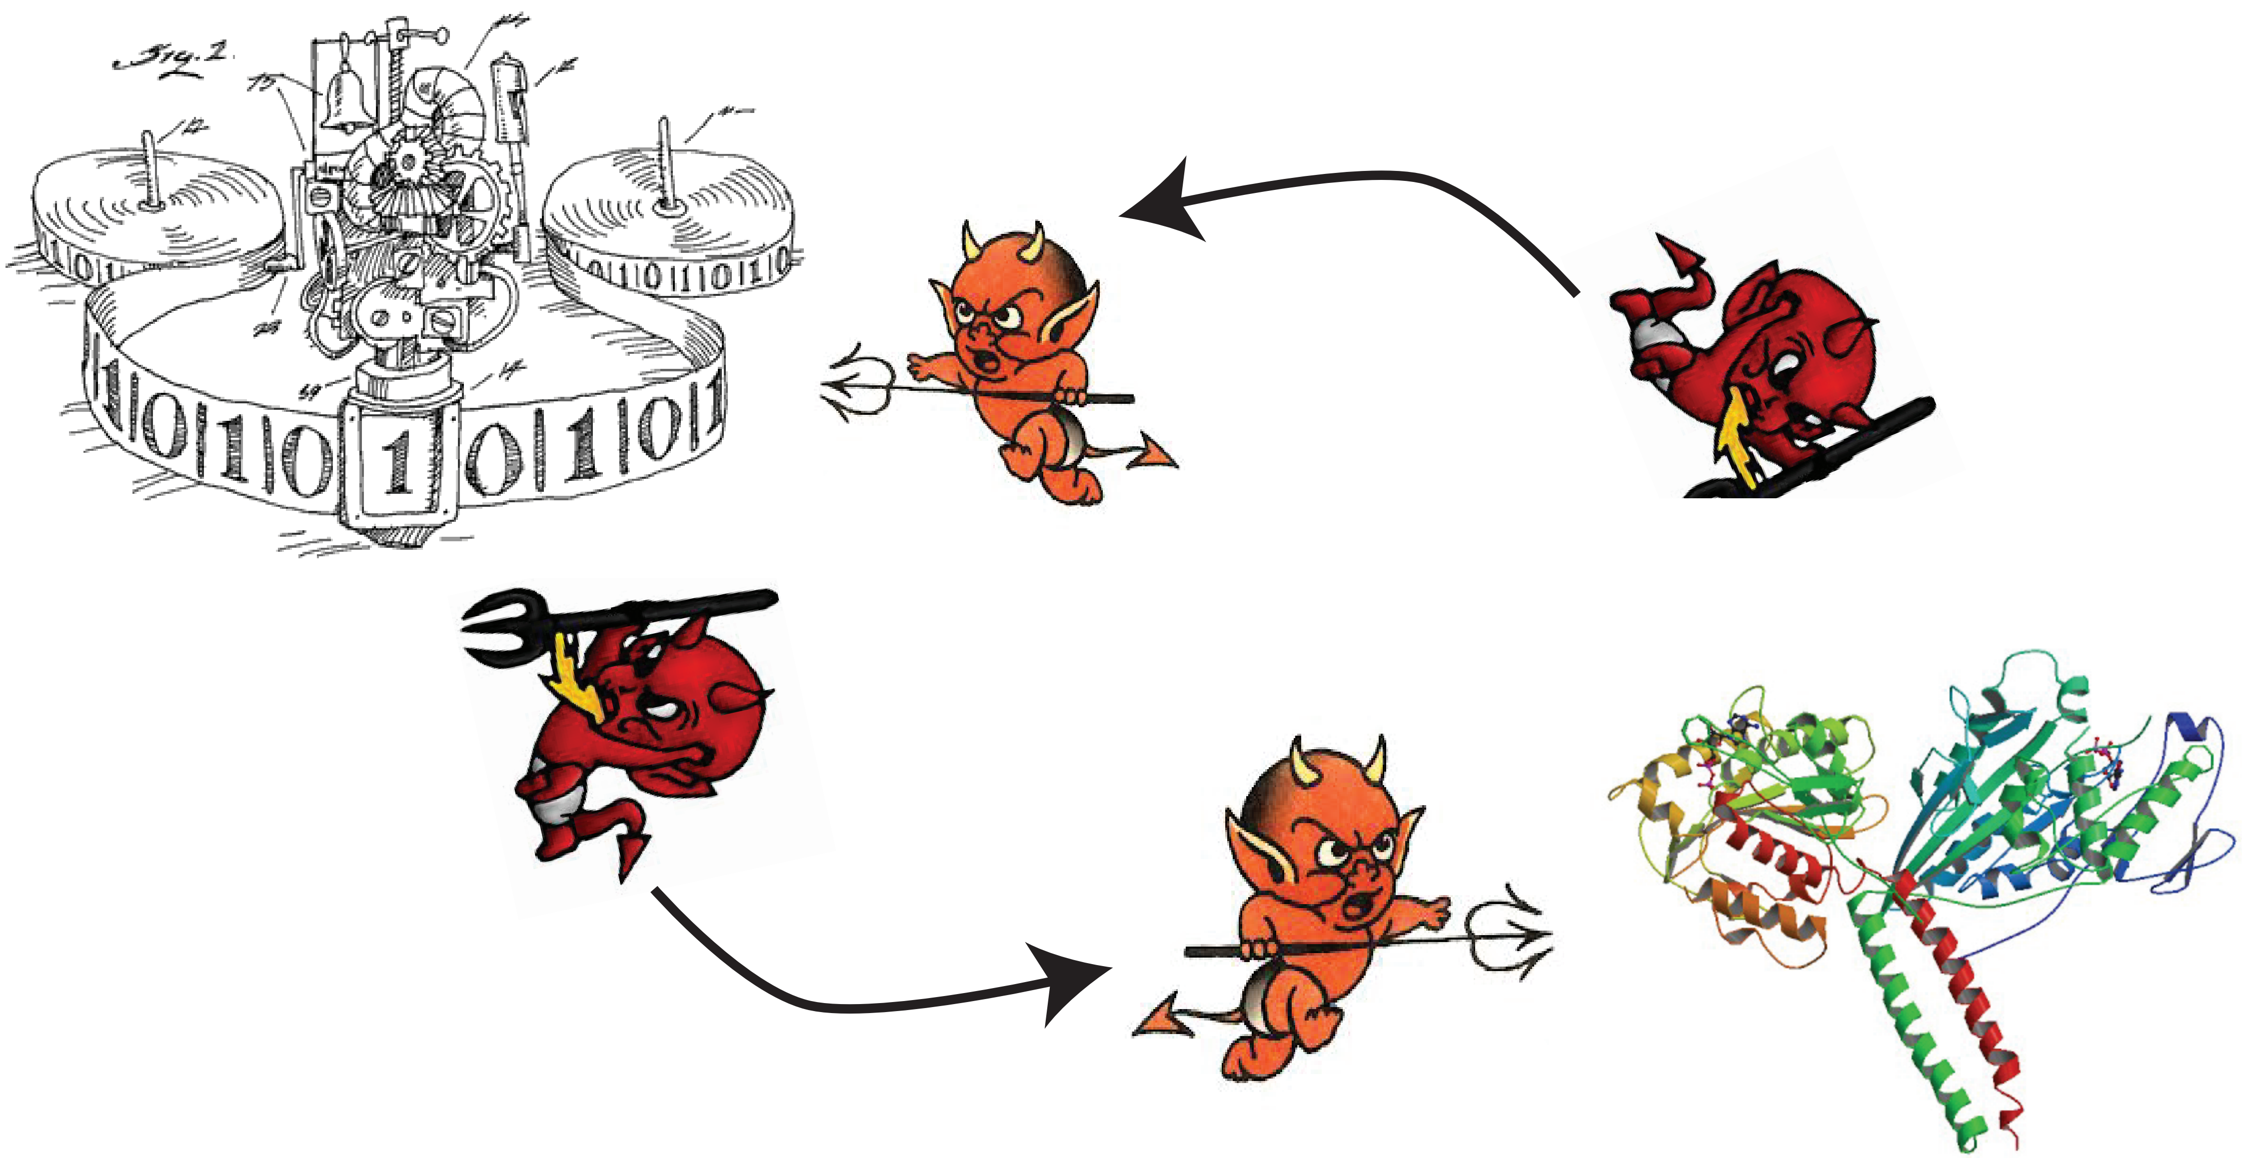

Maxwell's Demon

James Clerk &
Katherine Maxwell
(1865)

... if we conceive of a being whose faculties are so sharpened that he can follow every molecule in its course, such a being, whose attributes are as essentially finite as our own, would be able to do what is impossible to us. ... Now let us suppose that ... a vessel is divided into two portions, A and B, by a division in which there is a small hole, and that a being, who can see the individual molecules, opens and closes this hole, so as to allow only the swifter molecules to pass from A to B, and only the slower molecules to pass from B to A. He will thus, without expenditure of work, raise the temperature of B and lower that of A, in contradiction to the second law of thermodynamics. ...

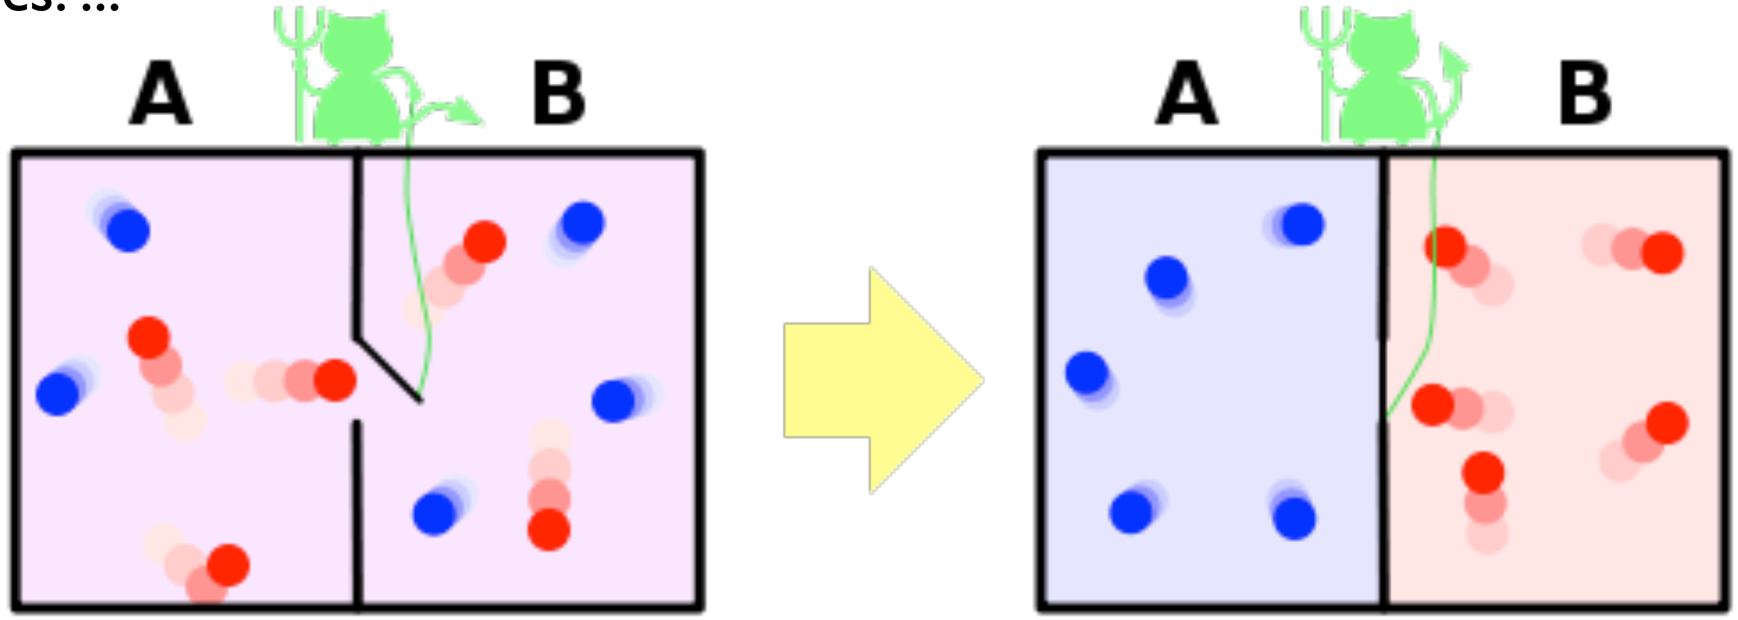

Maxwell's Demon

- Demon creates order out of chaos.

- Uses molecular information to convert heat to temperature difference & so to useful work.

Szilard's Engine:

"ON THE DECREASE OF ENTROPY IN A THERMODYNAMIC SYSTEM BY THE INTERVENTION OF INTELLIGENT BEINGS",
Leo Szilard, Zeitschrift fur Physik **65** (1929) 840-866.

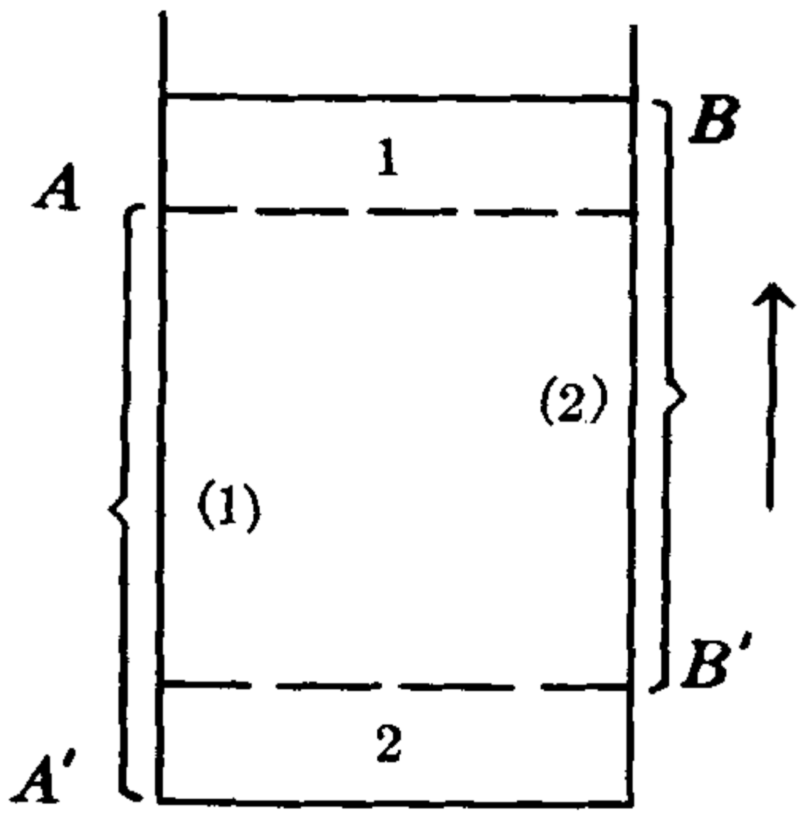

FIG. 1

... we must conclude *that the intervention which establishes the coupling between y and x, the measurement of x by y, must be accompanied by a production of entropy.*

... a simple inanimate device can achieve the same essential result as would be achieved by the intervention of intelligent beings. We have examined the "biological phenomena" of a nonliving device and have seen that it generates exactly that quantity of entropy which is required by thermodynamics.

Information Engines

A Technological Interlude or Why your Lap(top) is Hot!

Landauer Principle: $Q_{\text{erase}} \geq k_{\text{Boltzmann}} T \ln 2$

Erasing a bit dissipates energy to the heat bath.

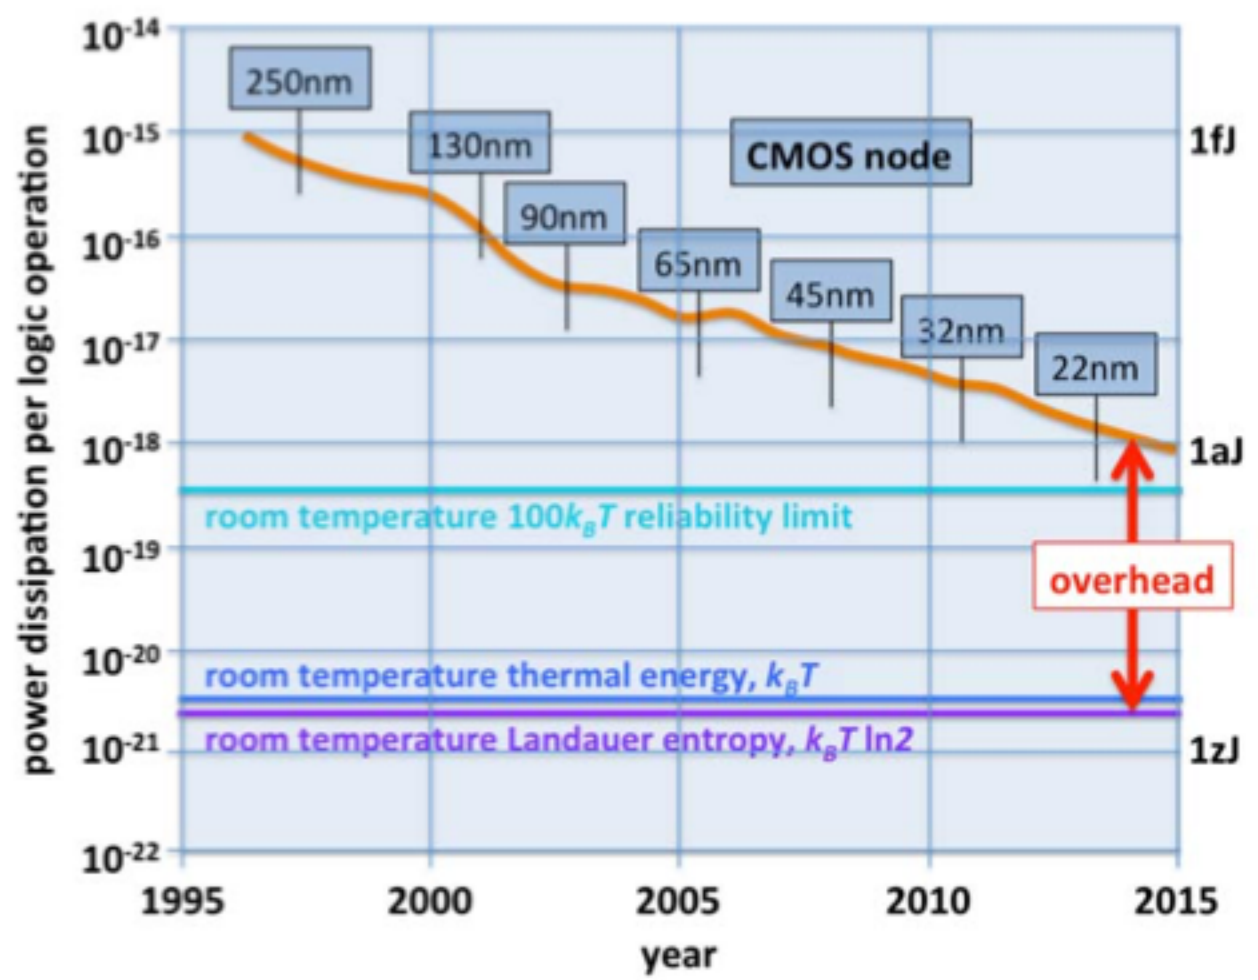

Figure 1. Evolution over time of the energy per CMOS logic operation. Energy overhead (see text) for the latest node is shown by the red arrow.

Szilard's Engine is a Chaotic Dynamical System

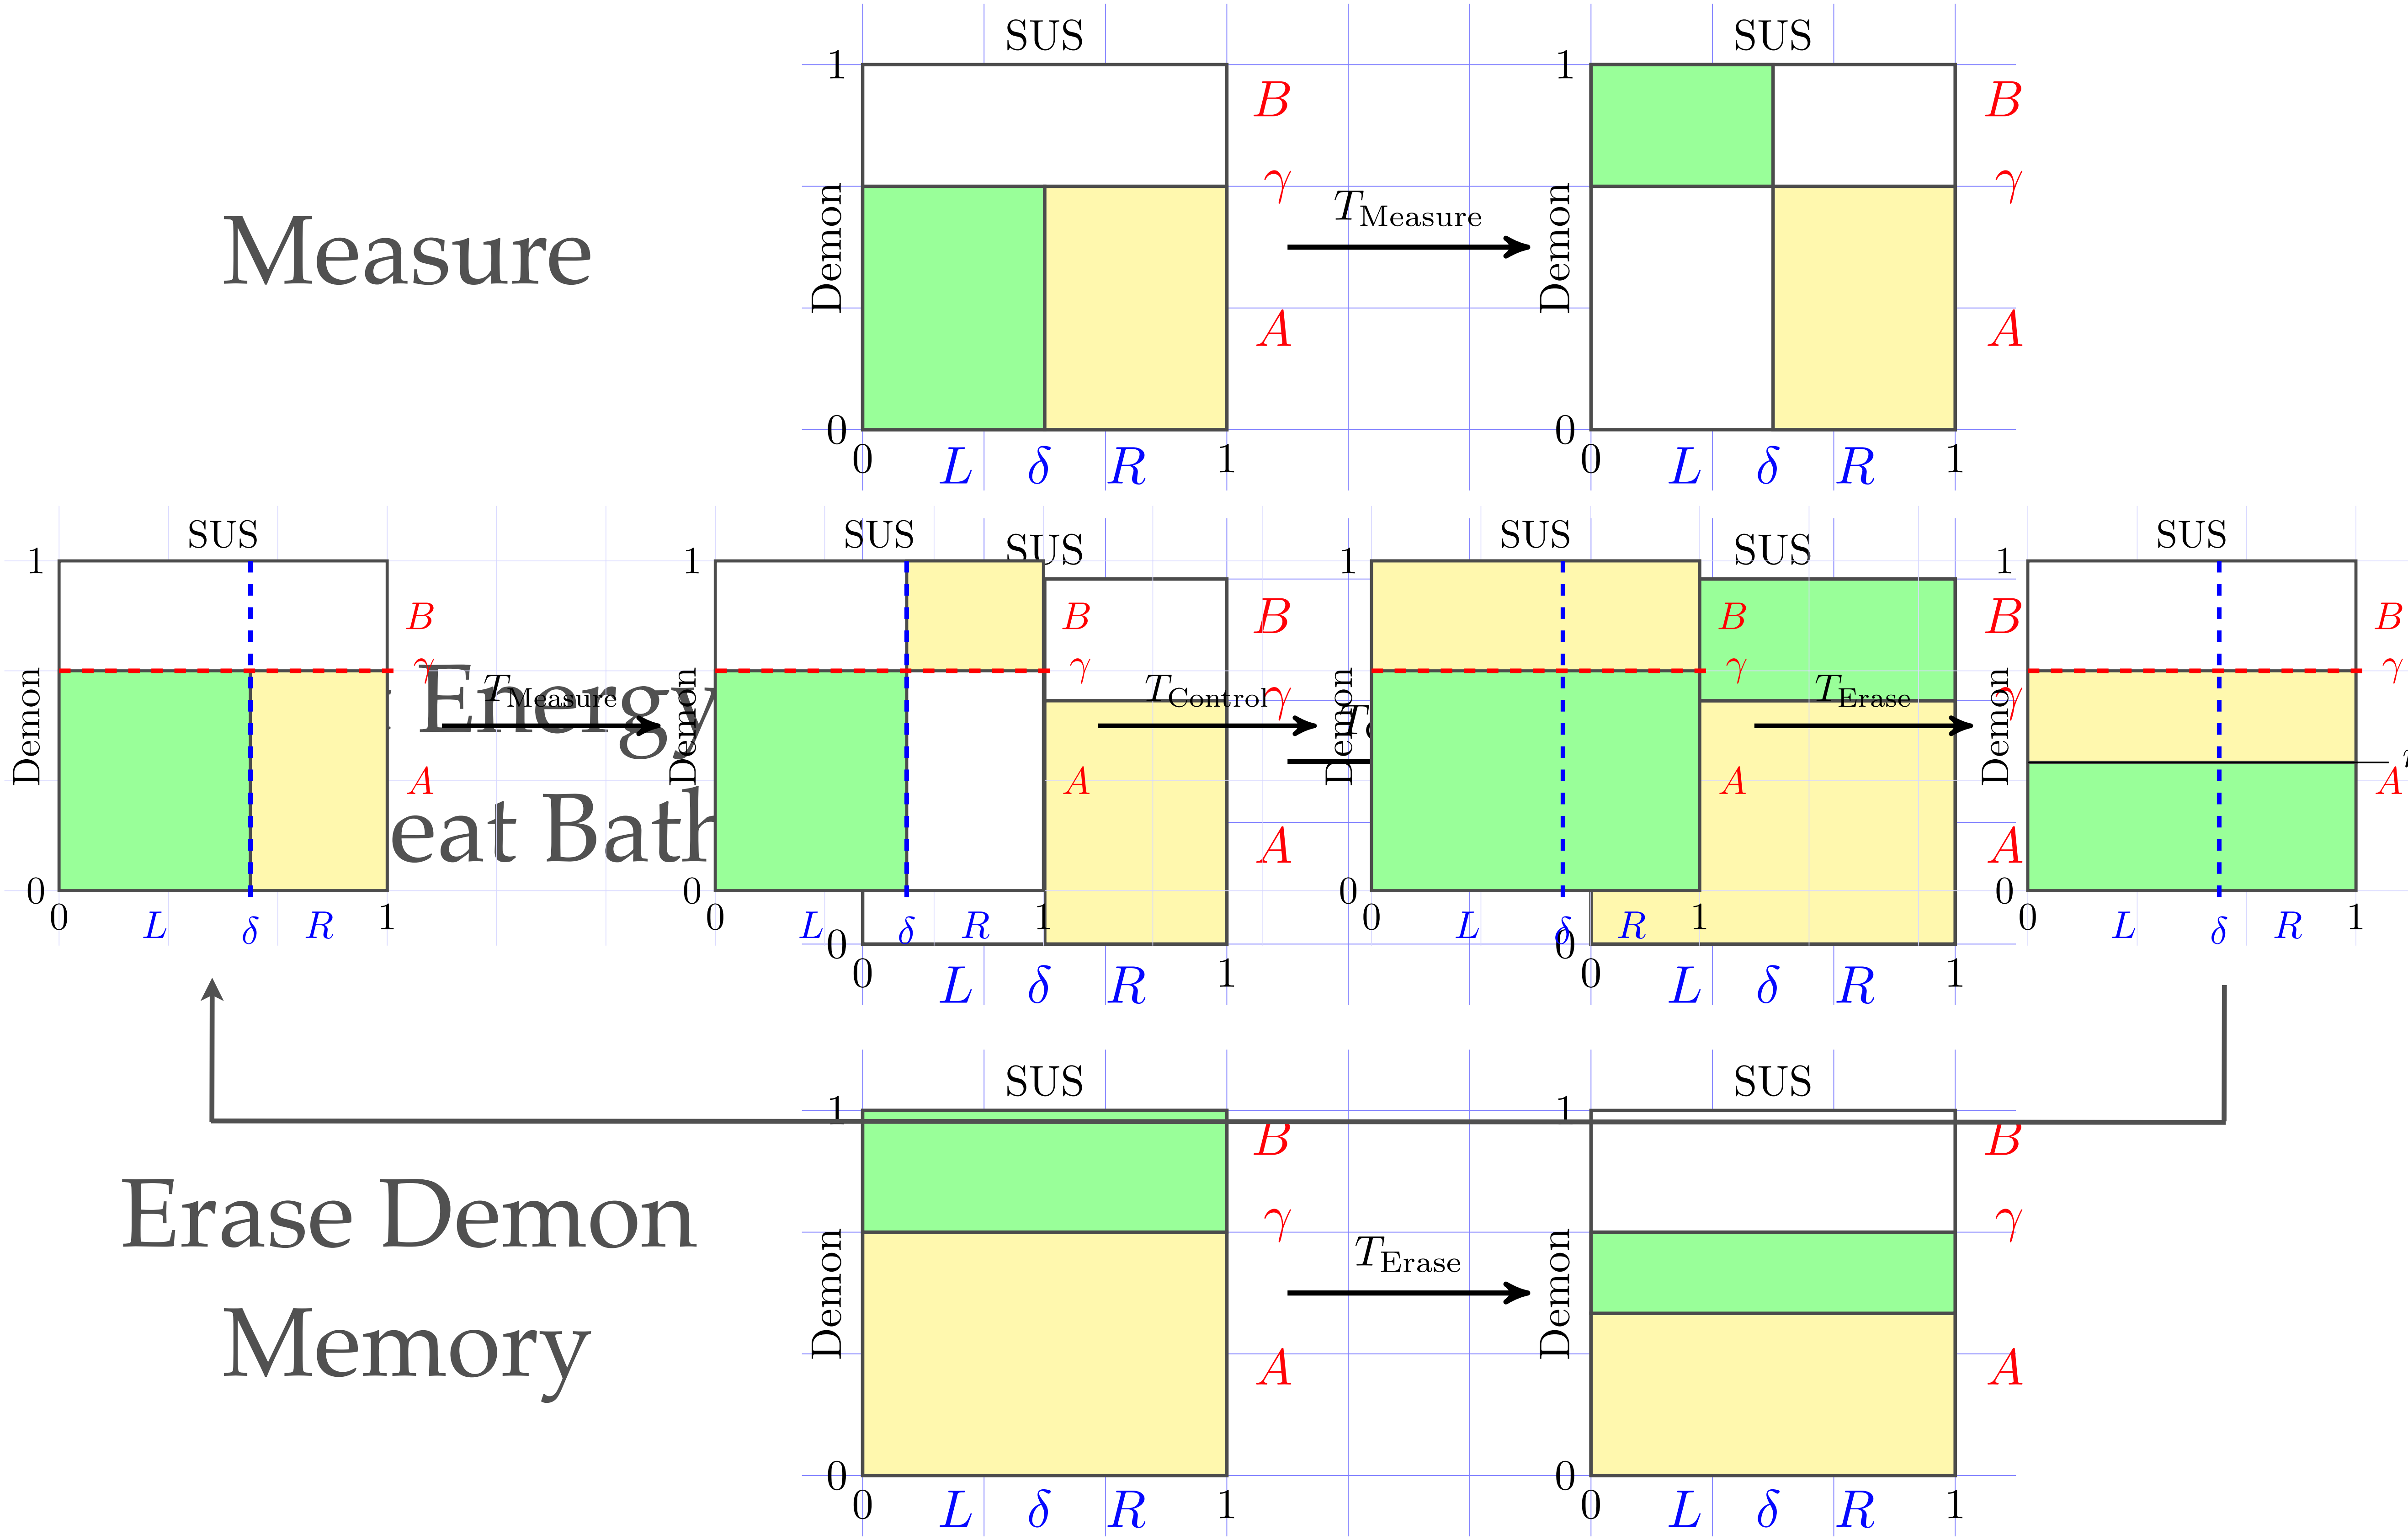

Boyd & Crutchfield, arxiv:1506.04327

Szilard Engine is a Chaotic Map

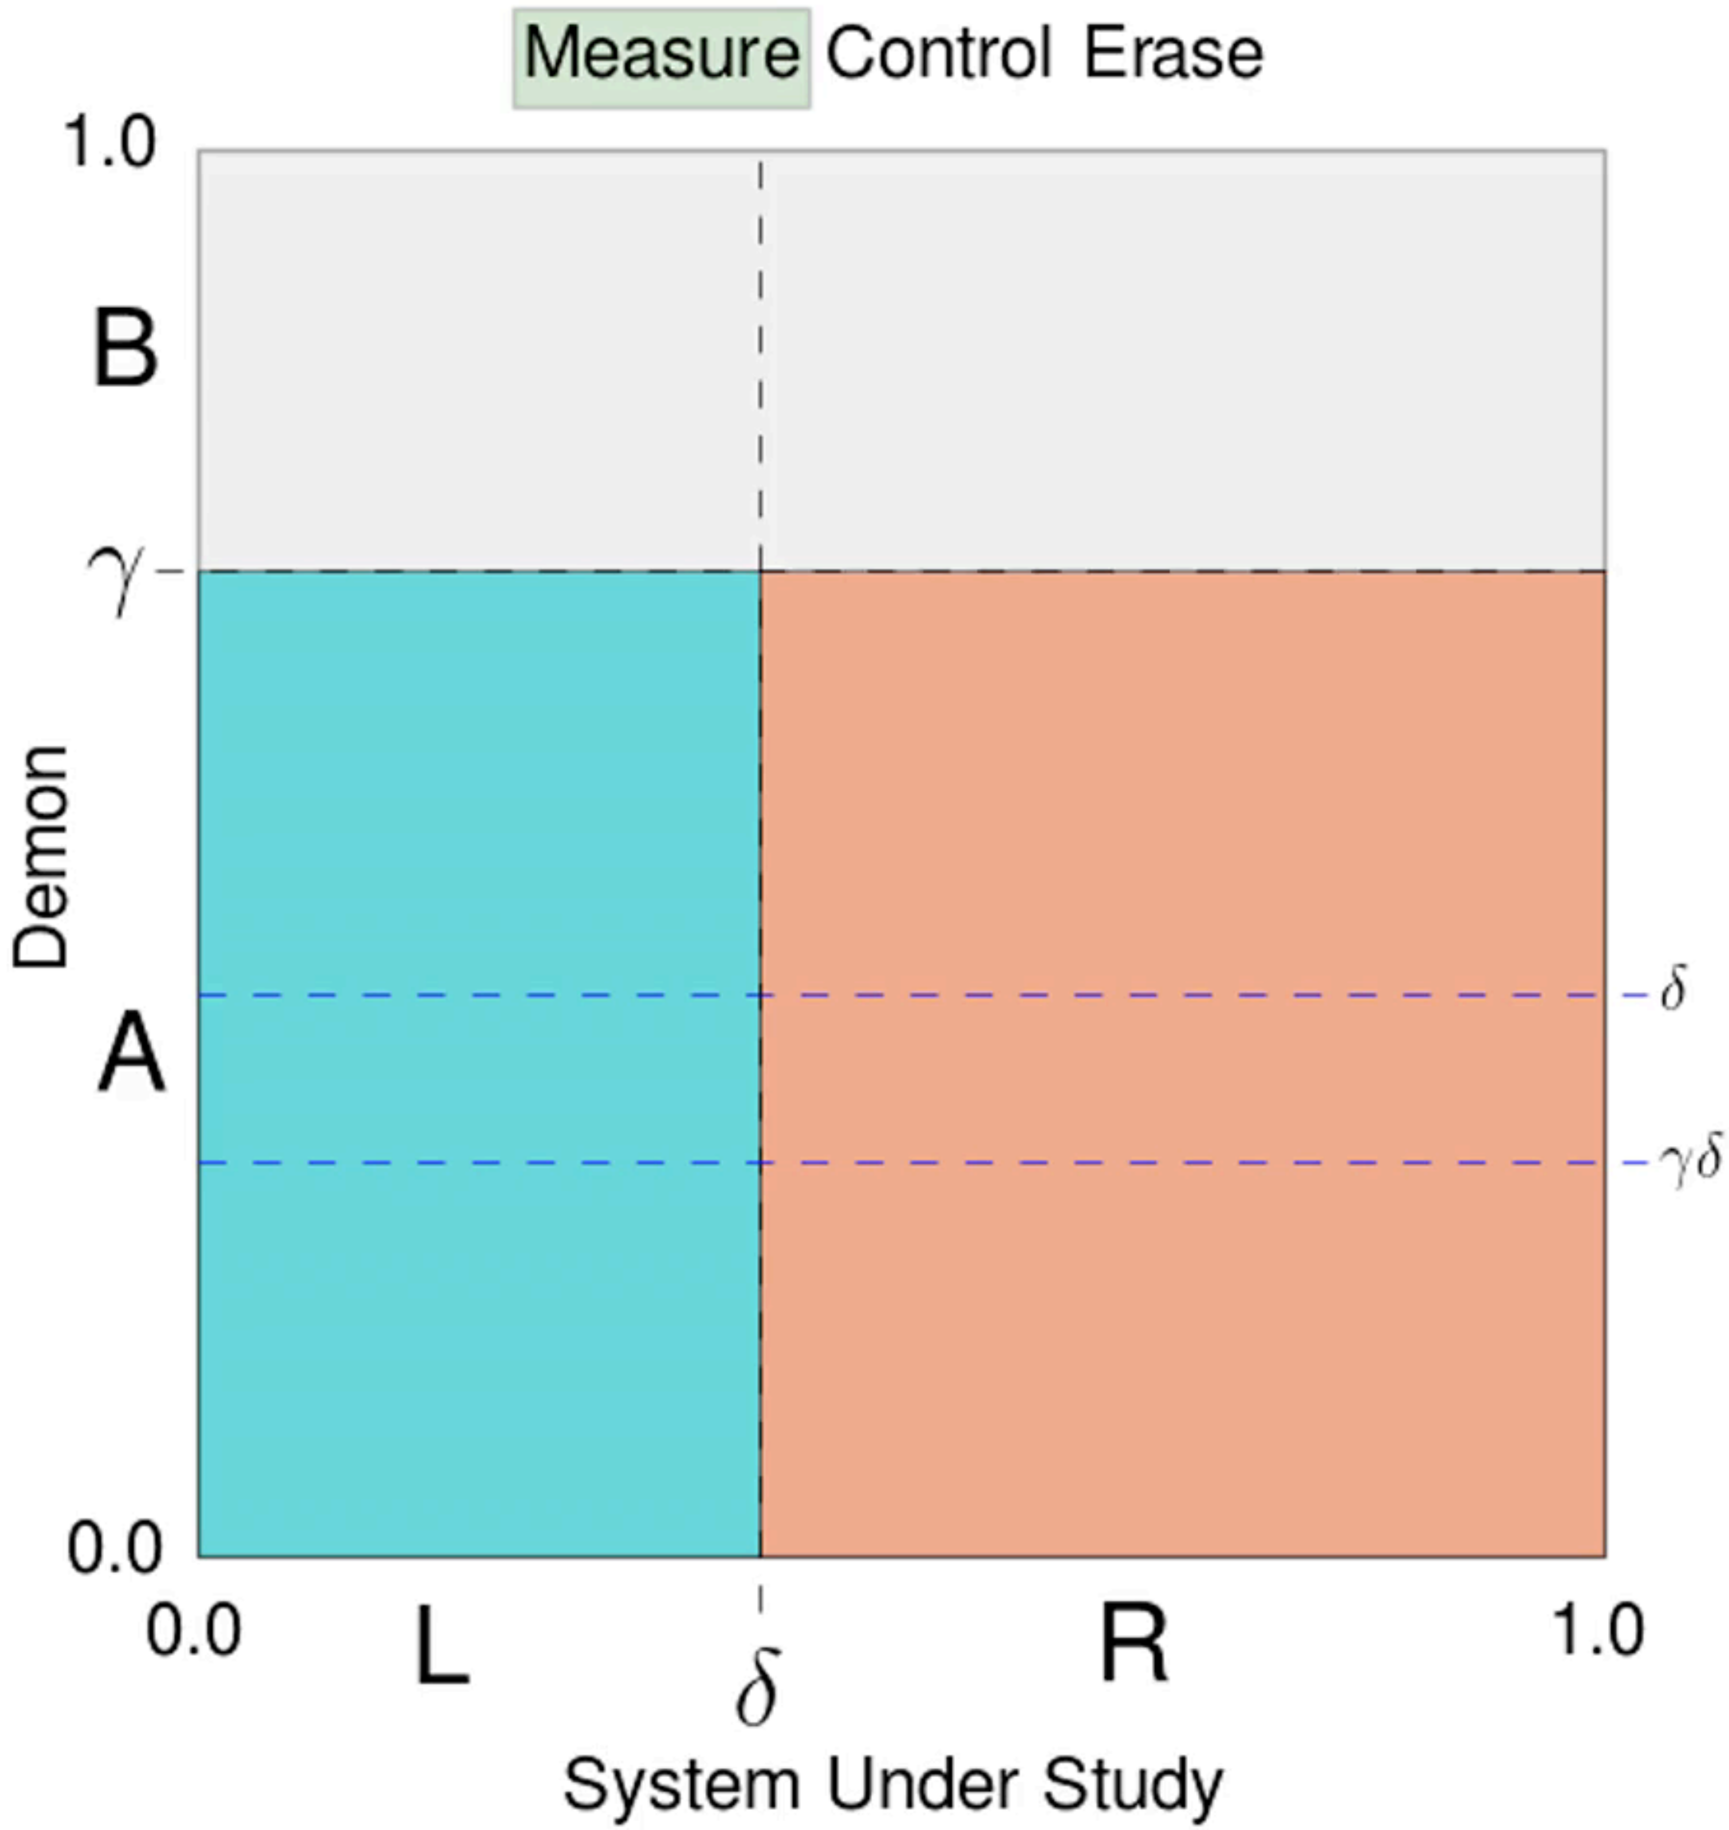

Beyond Landauer:
Energy Dissipation during Erasure and Measurement!

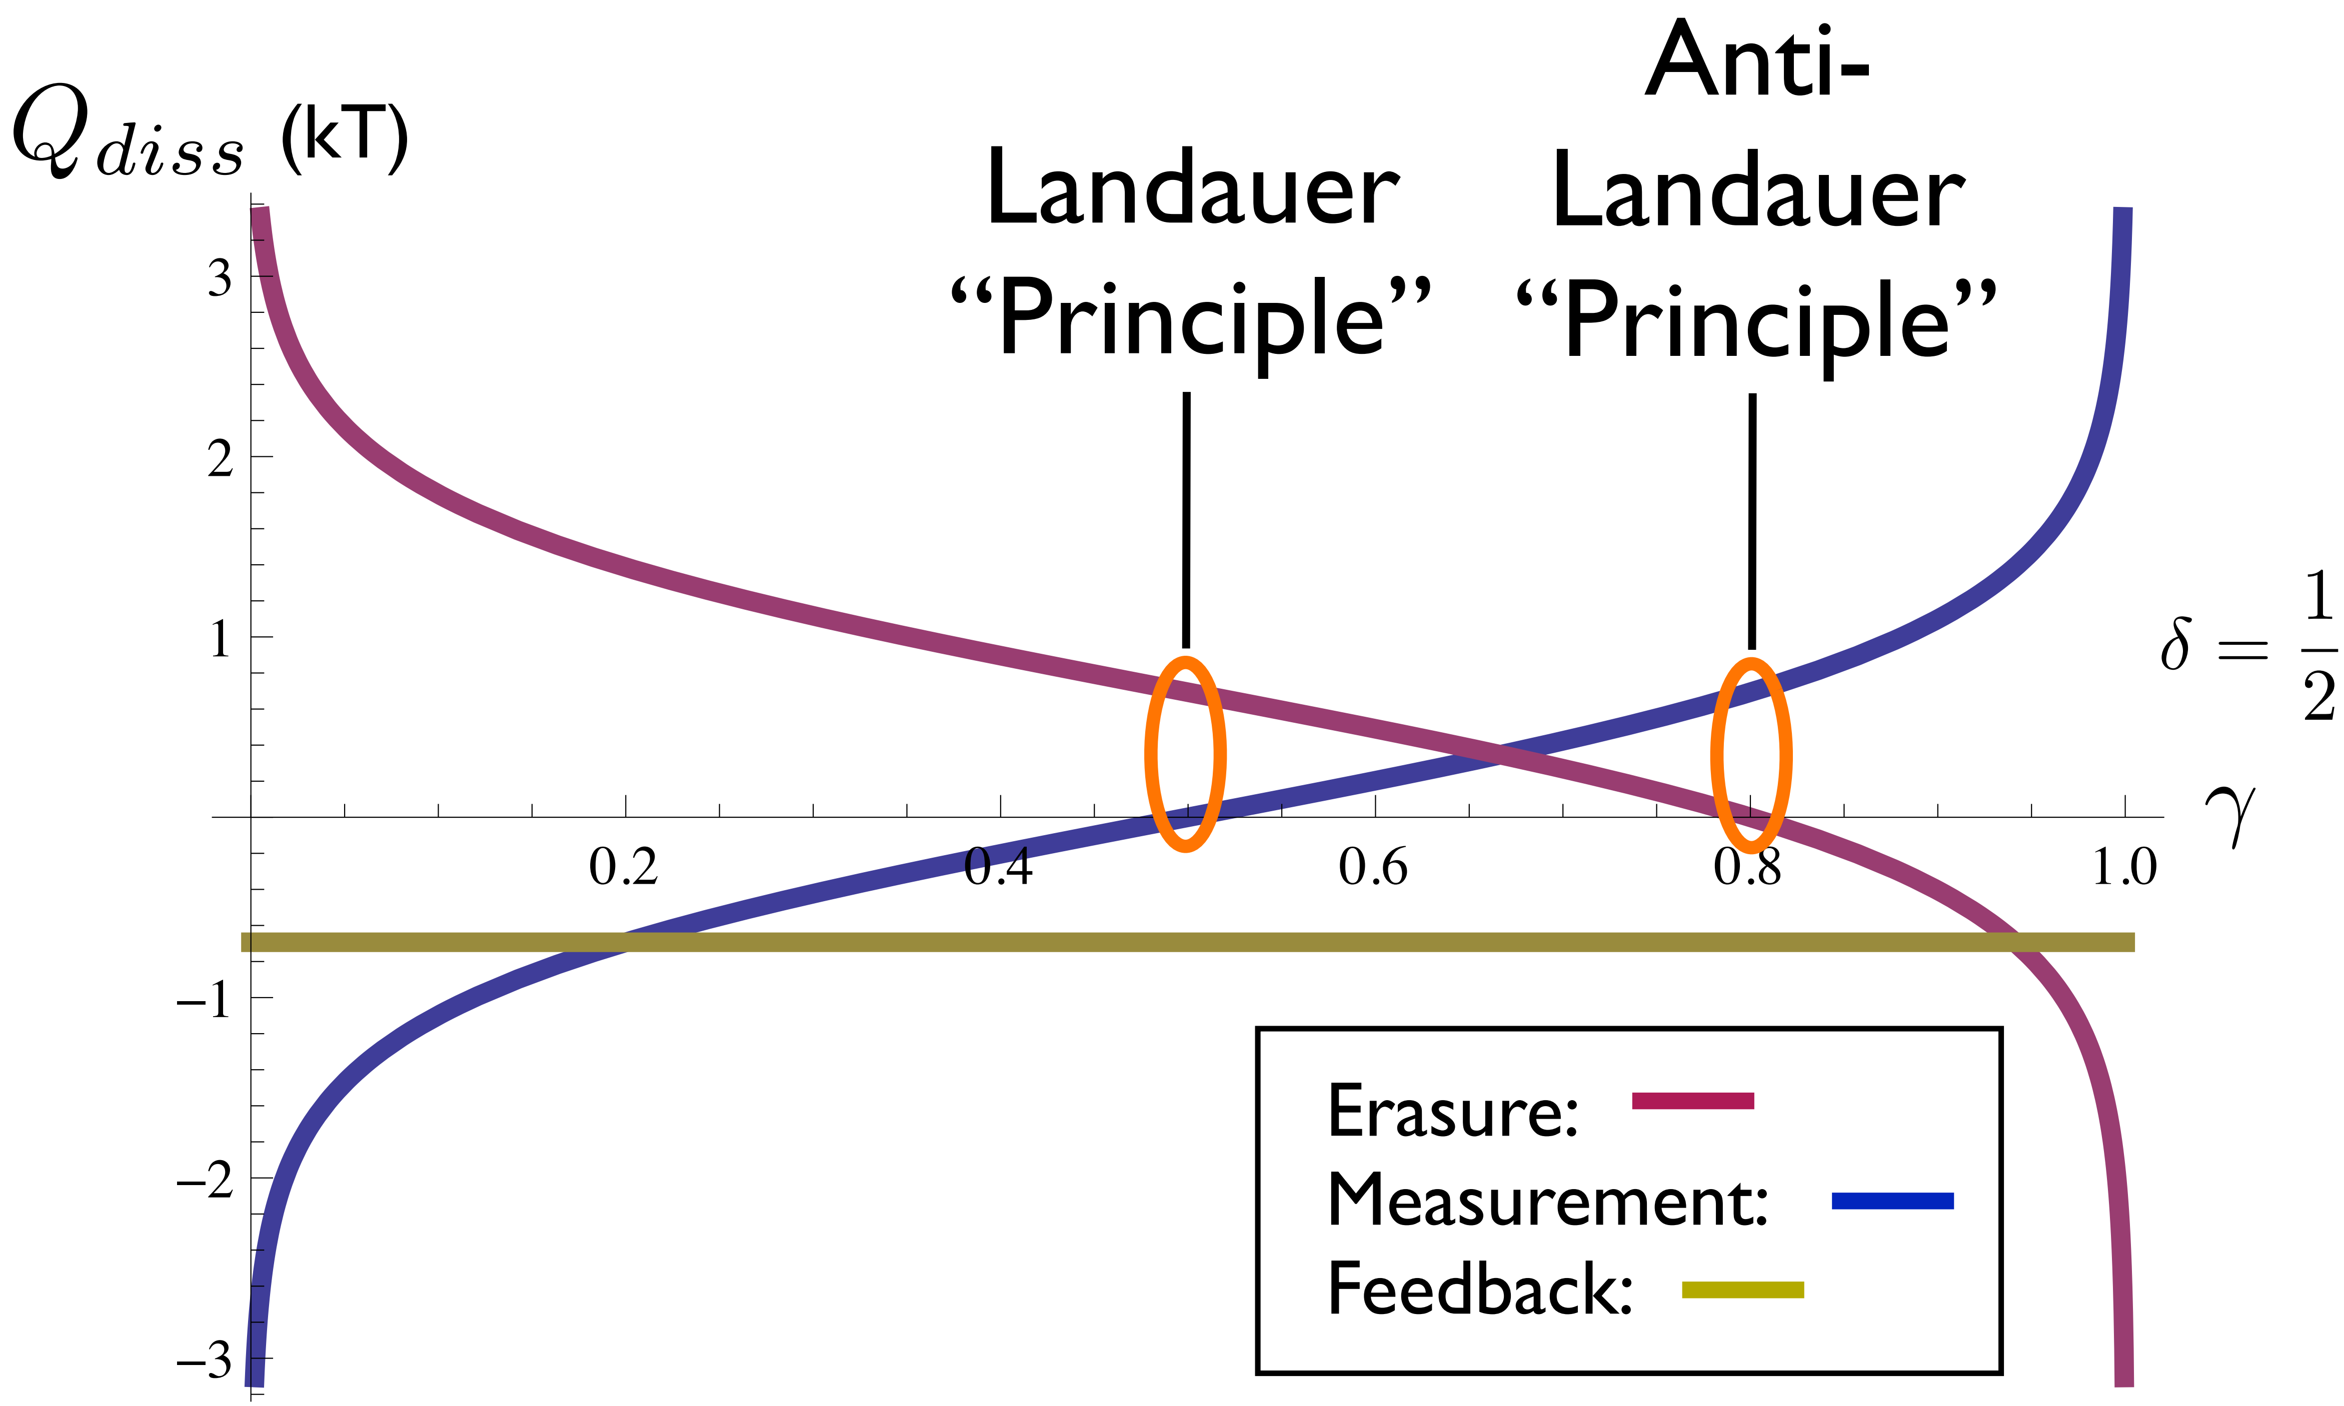

Complexity!

Information Theory for Complex Systems

Yesterday:

- Complex Processes
- Information in Processes

Today:

- Memory in Processes

Intrinsic Computation

- Measuring Structure
- Intrinsic Computation
- Optimal Models
- Physics of Information

See online course:

http://csc.ucdavis.edu/~chaos/courses/ncaso/