ϵM Reconstruction

Machine Reconstruction ...

$\epsilon\mathrm{M}$ Reconstruction:

Any method to go from process $\mathcal{P} \sim \Pr(\overleftrightarrow{\mathrm{S}})$ to its $\epsilon\mathrm{M}$

(1) Analytical: Given model, equations of motion, description, ...

(2) Statistical inference: Given samples of $\mathcal{P}$

(i) Subtree Reconstruction: Time or spacetime data to $\epsilon\mathrm{M}$

(ii) State-splitting (CSSR): Time or spacetime data to $\epsilon\mathrm{M}$

(iii) Spectral (eMSR): Power spectra to $\epsilon\mathrm{M}$

(iv) Optimal Causal Inference: Time or spacetime data to $\epsilon\mathrm{M}$

How to reconstruct an ϵM: Subtree algorithm

Given: Word distributions $\Pr(s^D),\ D = 1, 2, 3, \ldots$

Steps:

(1) Form depth-D parse tree.

(2) Calculate node-to-node transition probabilities.

(3) Causal states: Find morphs $\Pr(\overrightarrow{s}^L | \overleftarrow{s}^K)$ as subtrees.

(4) Label tree nodes with morph (causal state) names.

(5) Extract state-to-state transitions from parse tree.

(6) Assemble into ϵM: $\mathcal{M} = \left\{ \boldsymbol{\mathcal{S}}, \{T^{(s)}, s \in \mathcal{A}\} \right\}$.

Algorithm parameters: D, L, K

Machine Reconstruction ...

How to reconstruct an $\epsilon\mathrm{M}$...

Examples:

(1) Period-1
(2) Fair Coin
(3) Period-2
(4) Golden Mean Process
(5) Even Process

Machine Reconstruction ...

Examples (back to the Prediction Game):

Period-1: $\ldots 11111111111111$

Parse Tree D = 5

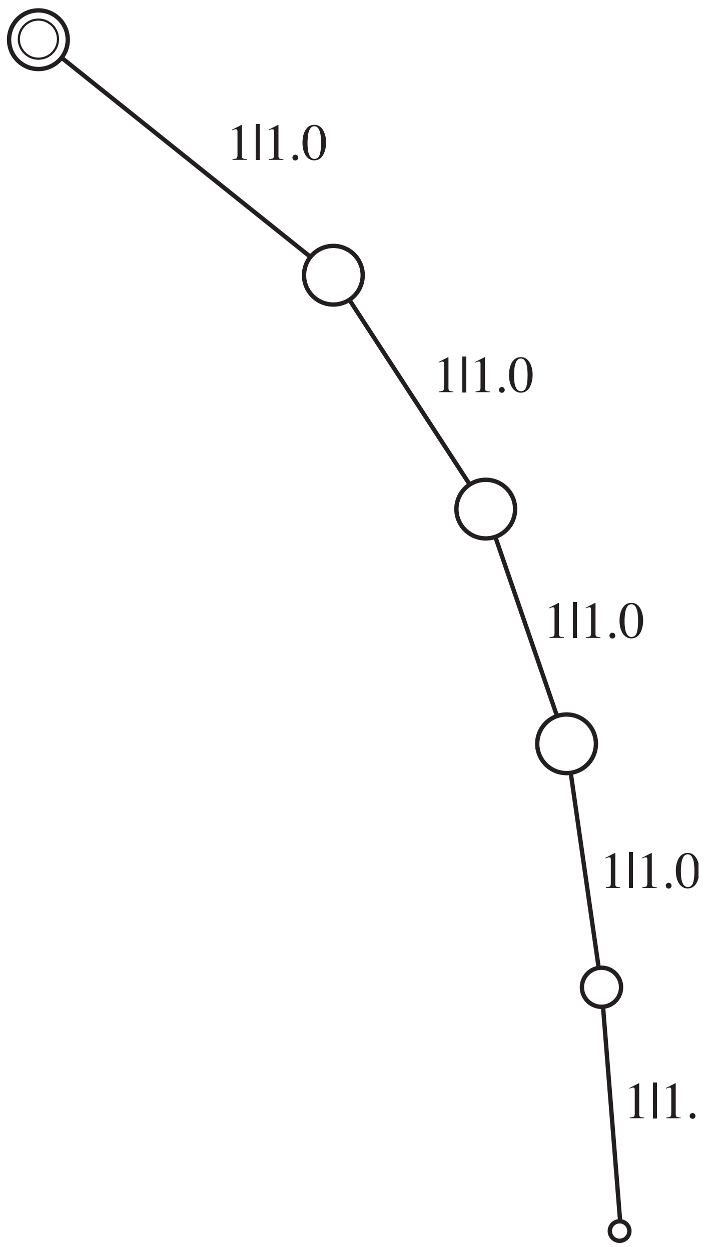

Machine Reconstruction ...

Examples (back to the Prediction Game):

Period-1: $\ldots 11111111111111$

Parse Tree D = 5

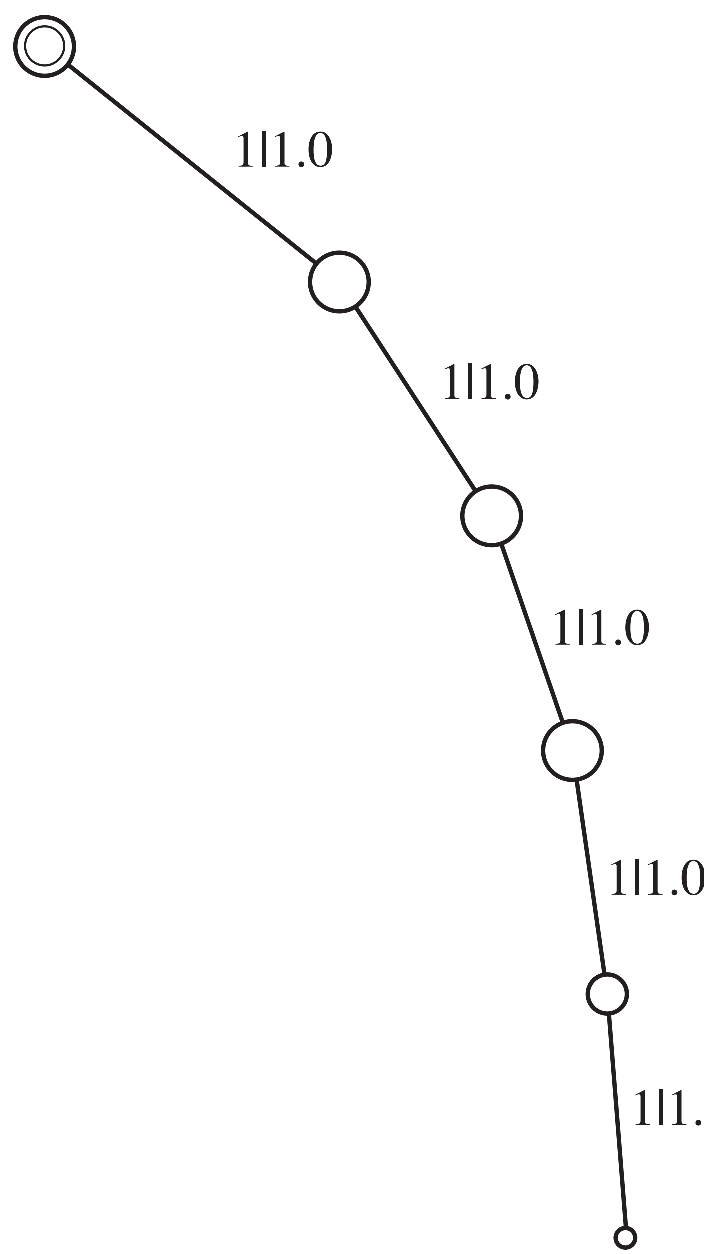

Morph L = 2

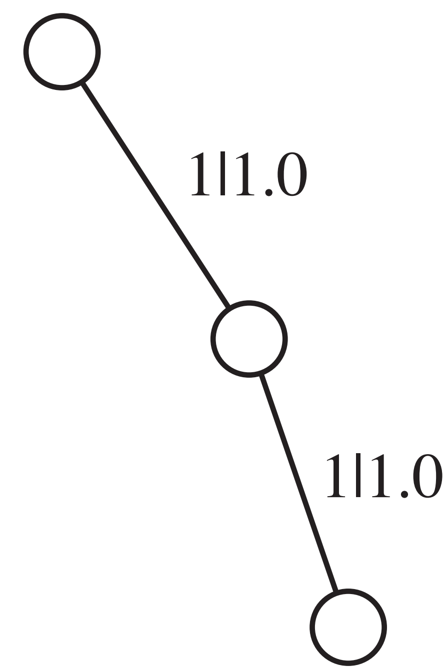

Machine Reconstruction ...

Examples (back to the Prediction Game) ...

Period-1: $\ldots 11111111111111$

Space of histories: A single point.

One future morph: $\{1^+\}$

Morph A distribution: $\Pr(\overrightarrow{S}^L \mid \overleftarrow{s}) = 1$

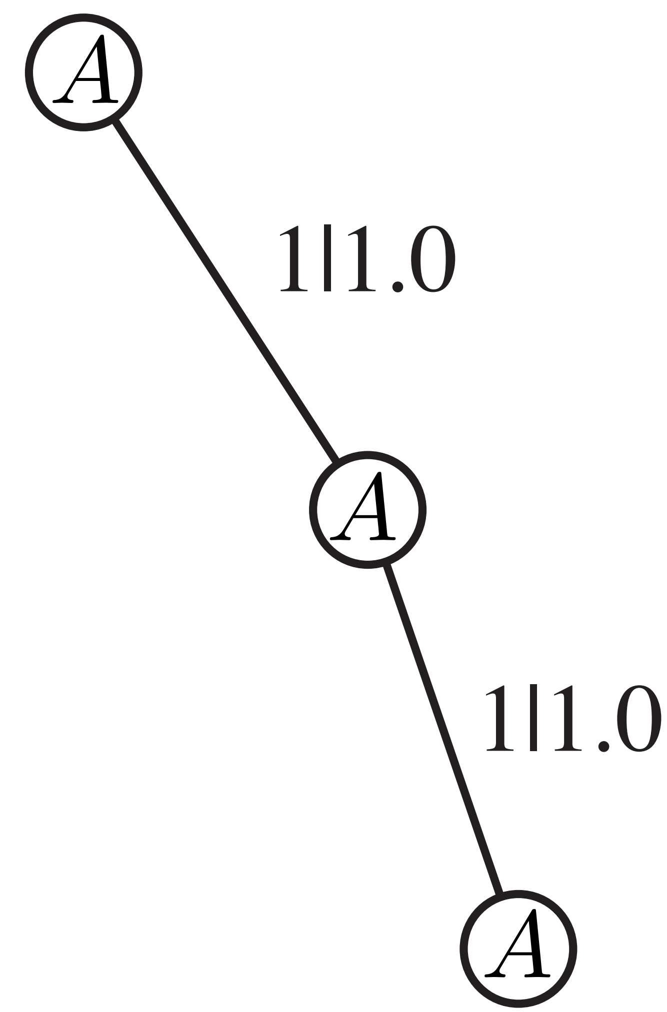

Machine Reconstruction ...

Examples (back to the Prediction Game) ...

Period-1 ...

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

$$\boldsymbol{\mathcal{S}} = \{\mathcal{S}_0 = ...111111\}$$

$$T^{(0)} = (0)$$

$$T^{(1)} = (1)$$

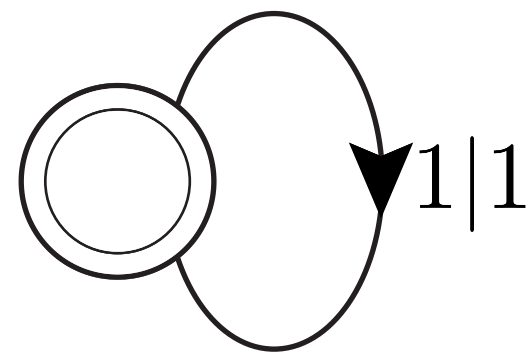

Machine Reconstruction ...

Examples (back to the Prediction Game) ...

Period-1 ...

Causal state distribution: $p_{\mathcal{S}} = (1)$

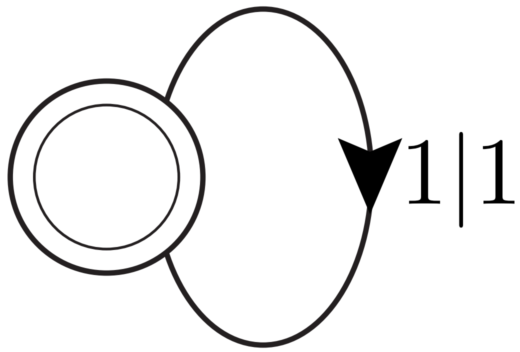

Entropy Rate: $h_\mu = 0 \text{ bits per symbol}$

Statistical Complexity: $C_\mu = 0 \text{ bits}$

Examples (back to the Prediction Game) ...

Fair Coin: $\ldots 0101001110001101$

Parse Tree D = 5

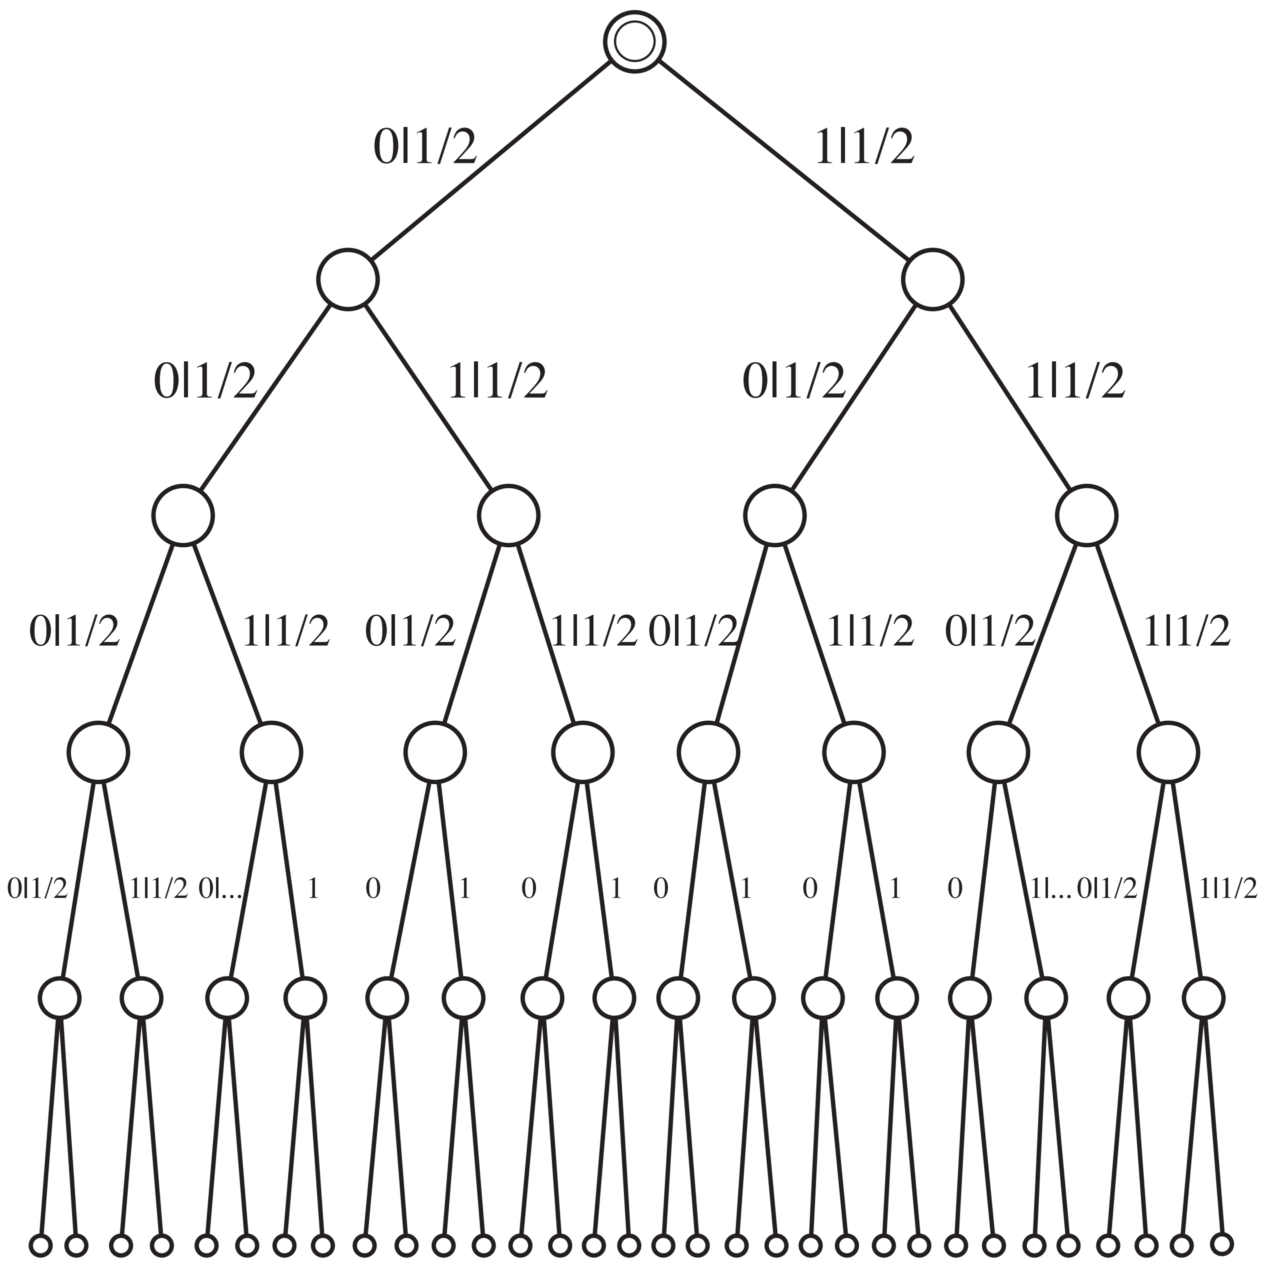

Machine Reconstruction ...

Examples (back to the Prediction Game) ...

Fair Coin: $\ldots 0101001110001101$

Parse Tree D = 5

Morph L = 2

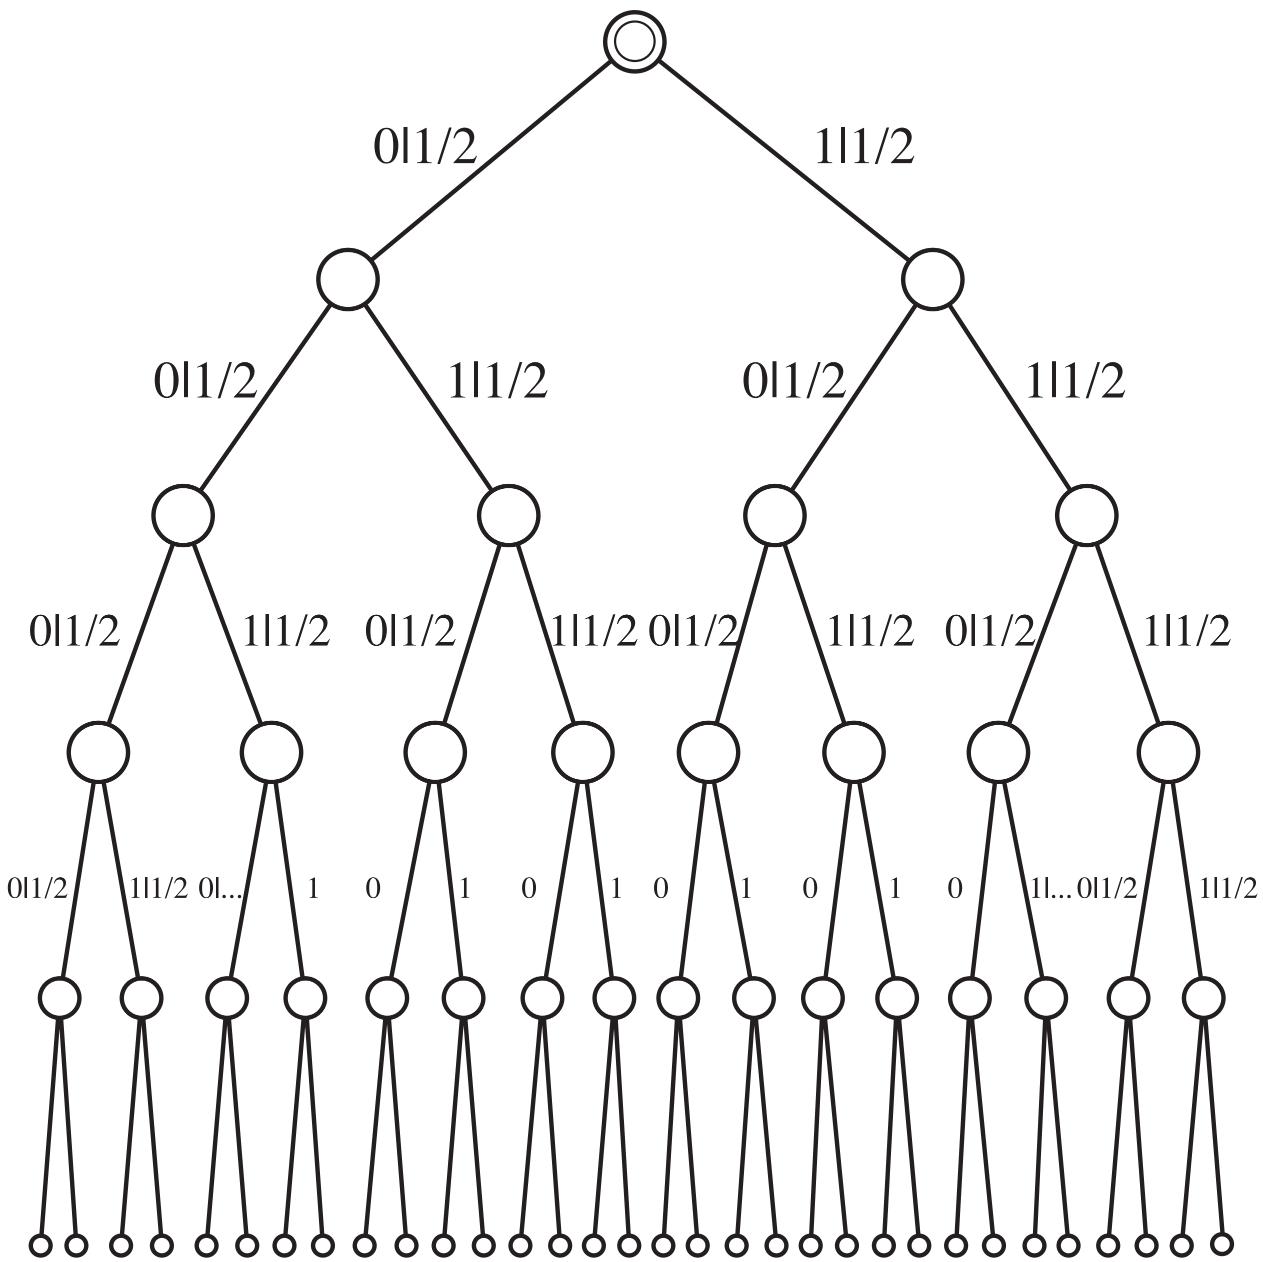

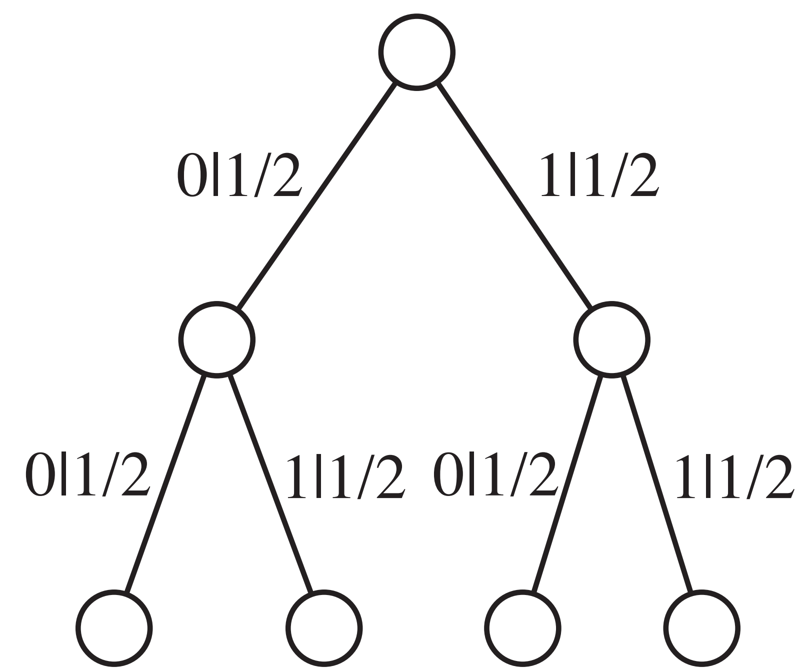

Machine Reconstruction ...

Examples (back to the Prediction Game) ...

Fair Coin ...

Space of histories: $\overleftarrow{\mathbf{S}}^K = \mathcal{A}^K$

One future morph: $\mathcal{A}^L$

Morph A distribution: $\Pr(\overrightarrow{S}^L | \overleftarrow{s}) = 2^{-L}$

Machine Reconstruction ...

Examples (back to the Prediction Game) ...

Fair Coin ...

Label tree nodes
with state names:

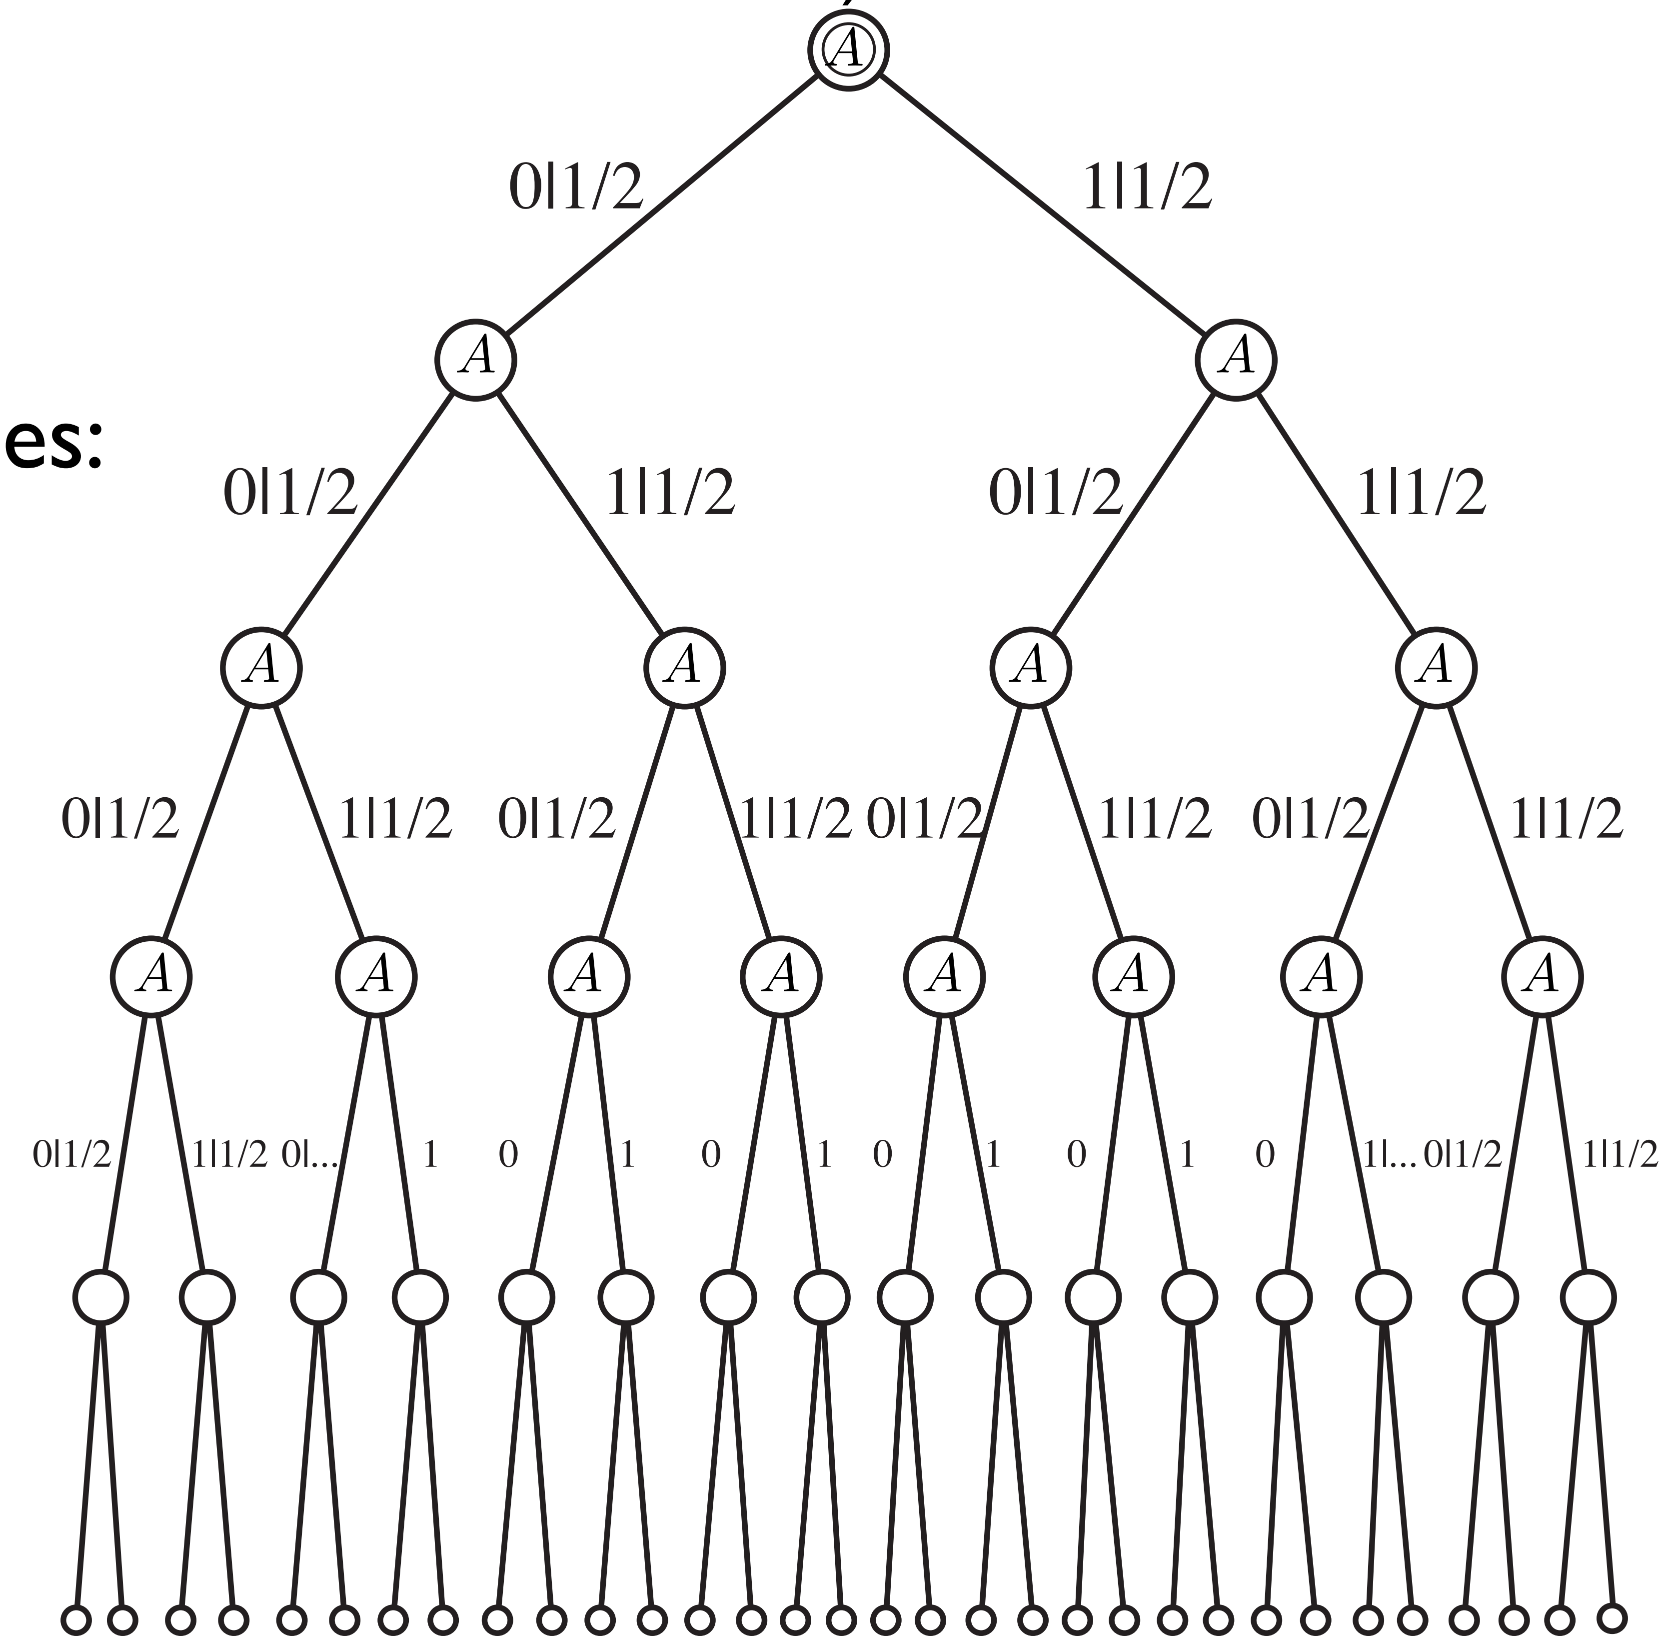

Machine Reconstruction ...

Examples (back to the Prediction Game) ...

Fair Coin ...

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

$$\boldsymbol{\mathcal{S}} = \{\mathcal{S}_0 = \mathcal{A}^L\}$$

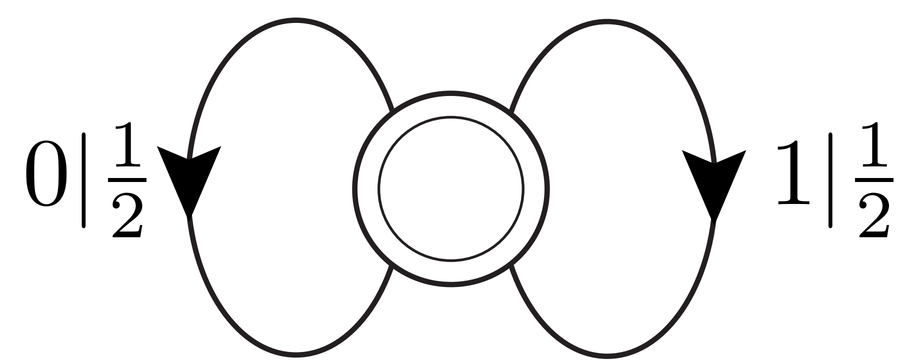

$$T^{(0)} = \left(\tfrac{1}{2}\right)$$

$$T^{(1)} = \left(\tfrac{1}{2}\right)$$

Examples (back to the Prediction Game) ...

Fair Coin ...

Causal state distribution: $p_{\mathcal{S}} = (1)$

Entropy Rate: $h_\mu = 1 \text{ bit per symbol}$

Statistical Complexity: $C_\mu = 0 \text{ bits}$

Machine Reconstruction ...

Examples (back to the Prediction Game) ...

Period-2 Process: $\ldots 01010101010101$

Parse Tree D = 5

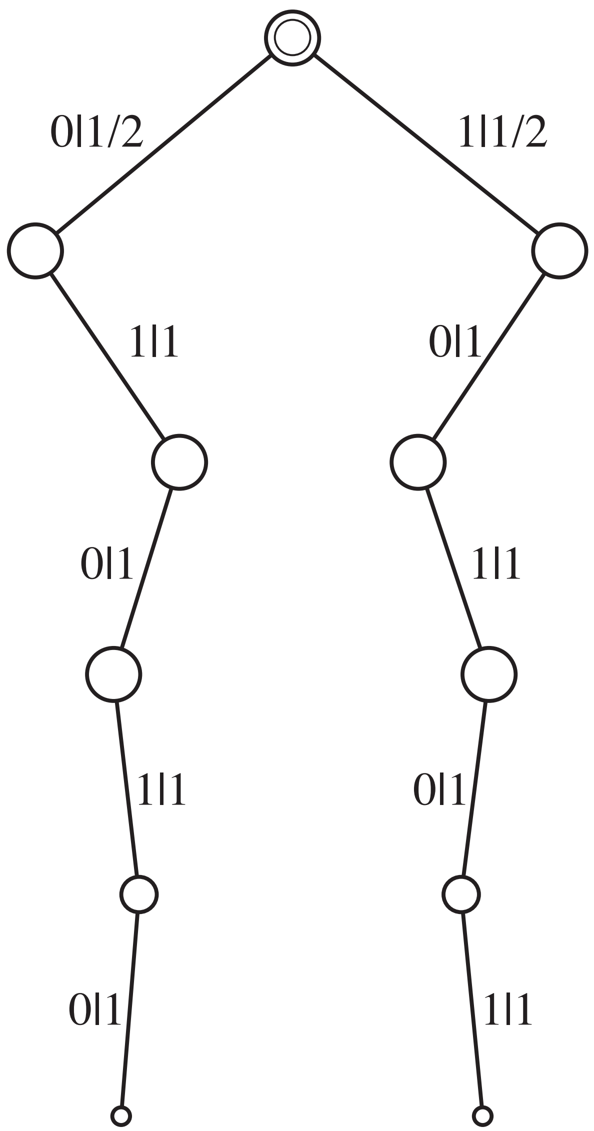

Machine Reconstruction ...

Examples (back to the Prediction Game) ...

Period-2 Process: $\ldots 01010101010101$

Parse Tree D = 5

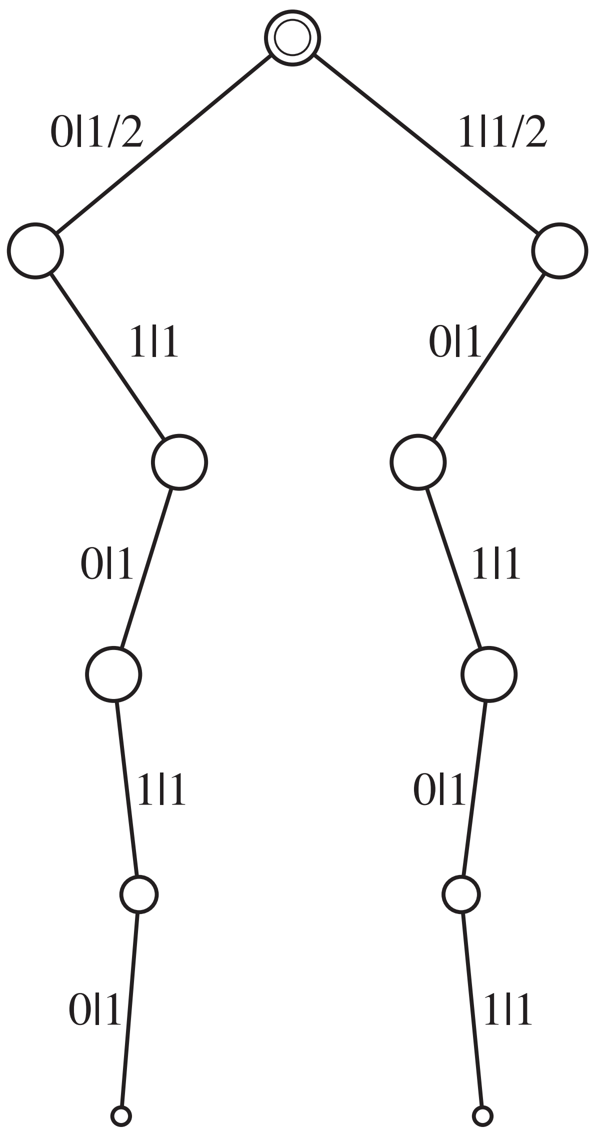

Morphs at L = 2

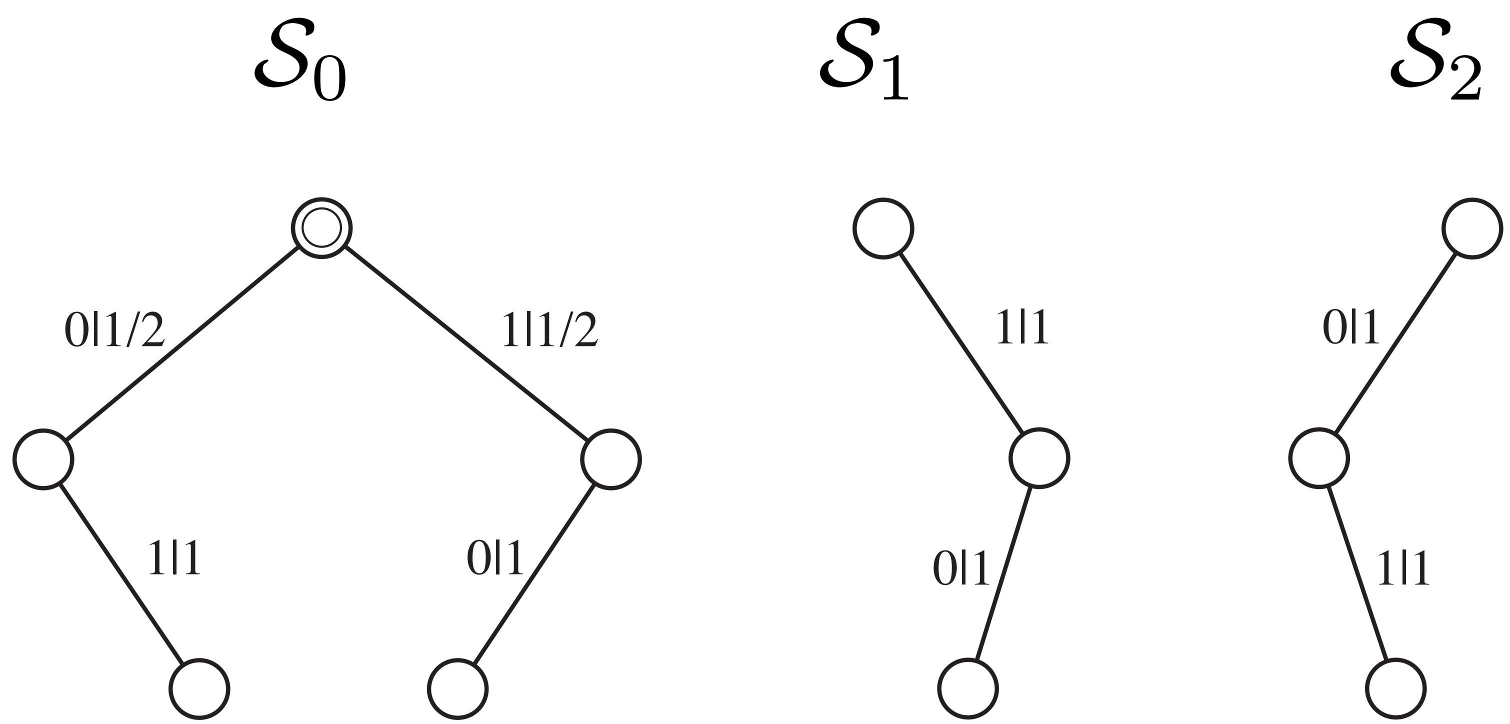

Examples (back to the Prediction Game) ...

Period-2 Process: $\ldots 0101010101010101$

Space of histories: $\overleftarrow{\mathbf{S}} = \{\overleftarrow{s}_0 = \ldots 101010, \overleftarrow{s}_1 = \ldots 010101\}$

Future morphs:

$$\{\overrightarrow{S}^1 | \lambda\} = \{101010\ldots, 010101\ldots\}$$
$$\{\overrightarrow{S}^1 | 0\} = \{101010\ldots\}$$
$$\{\overrightarrow{S}^1 | 1\} = \{010101\ldots\}$$
$$\{\overrightarrow{S}^1 | 10\} = \{101010\ldots\}$$
$$\{\overrightarrow{S}^1 | 01\} = \{010101\ldots\}$$

Morph distributions:

$$\Pr(0|\lambda) = \tfrac{1}{2} \qquad \Pr(1|0) = 1 \quad \Pr(0|0) = 0$$
$$\Pr(1|\lambda) = \tfrac{1}{2} \qquad \Pr(1|1) = 0 \quad \Pr(0|1) = 1$$

Machine Reconstruction ...

Examples (back to the Prediction Game) ...

Period-2 Process ...

Label tree nodes:

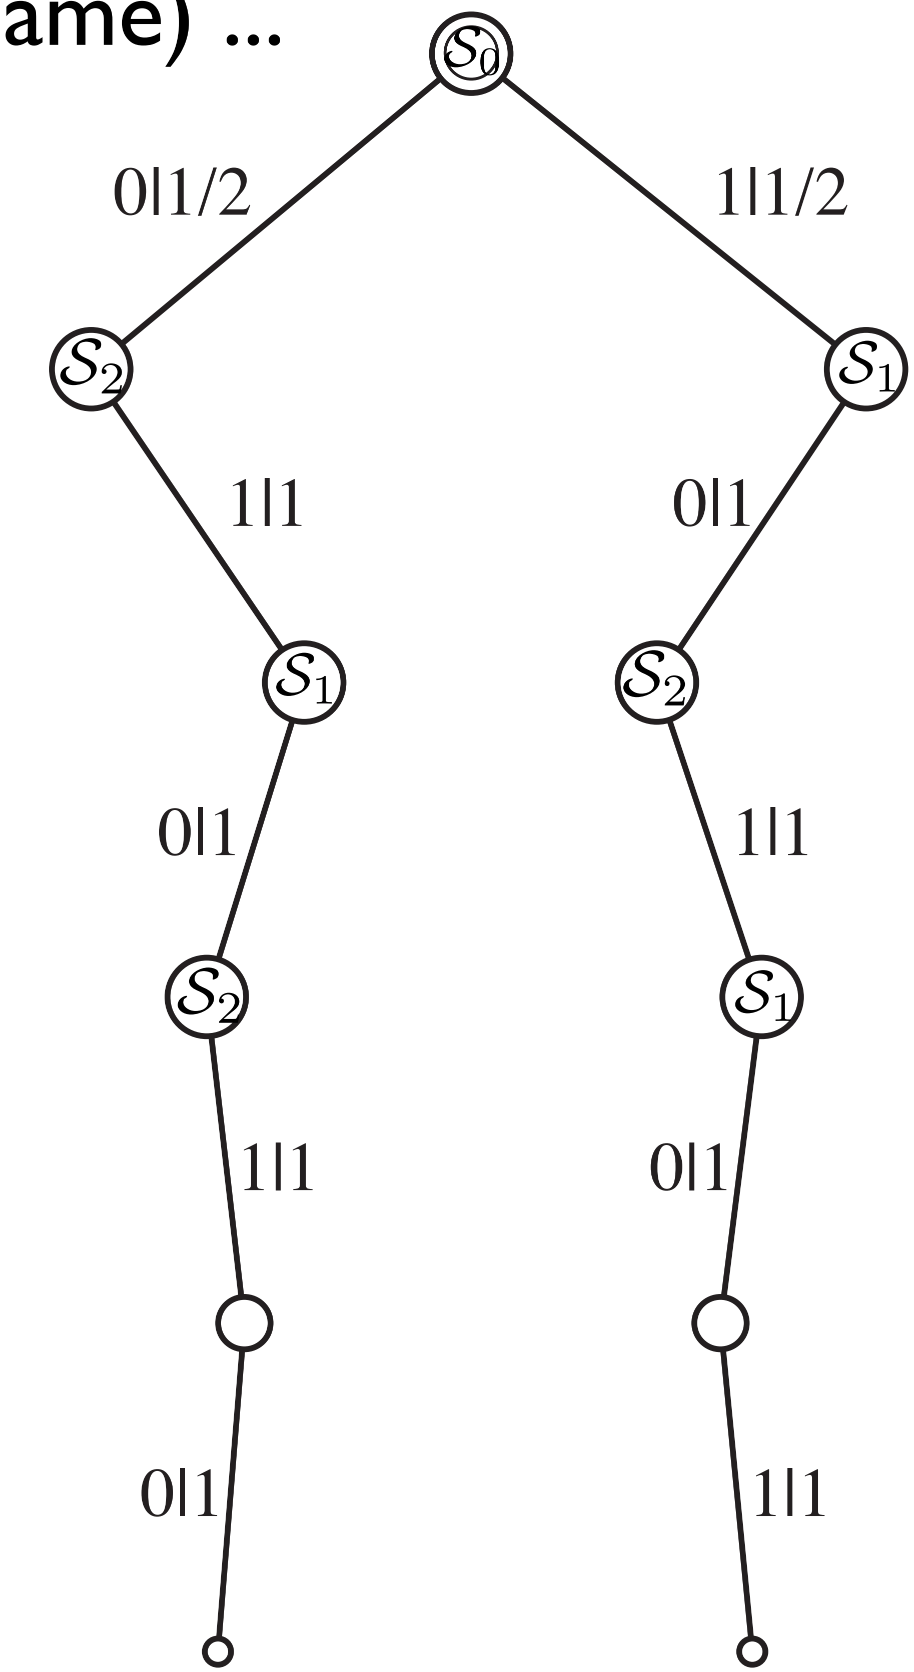

Machine Reconstruction ...

Examples (back to the Prediction Game) ...

Period-2 Process ...

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

$$\boldsymbol{\mathcal{S}} = \{\mathcal{S}_0 = \{\ldots 0101, \ldots 1010\}, \mathcal{S}_1 = \{\ldots 1010\}, \mathcal{S}_2 = \{\ldots 0101\}\}$$

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

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

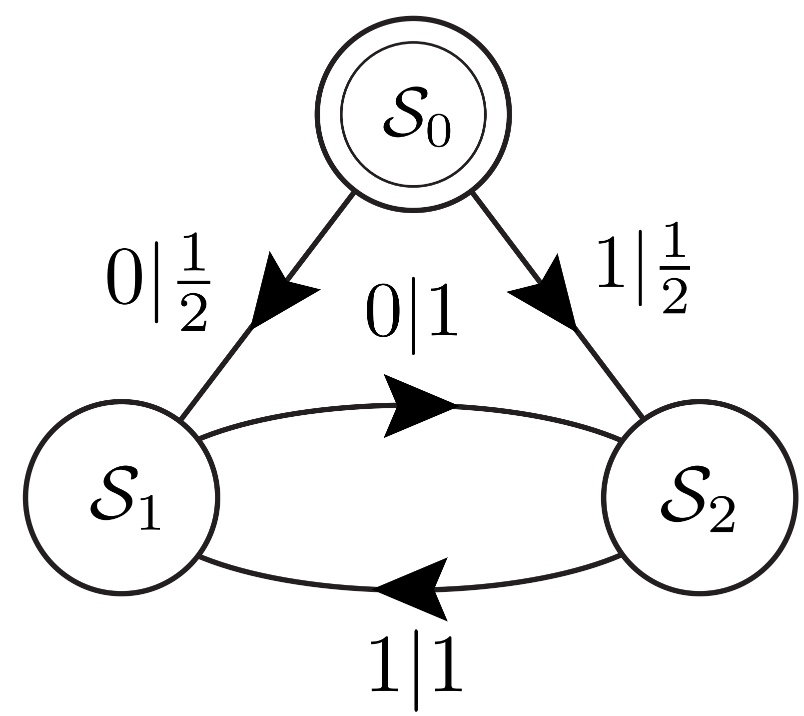

Examples (back to the Prediction Game) ...

Period-2 Process ...

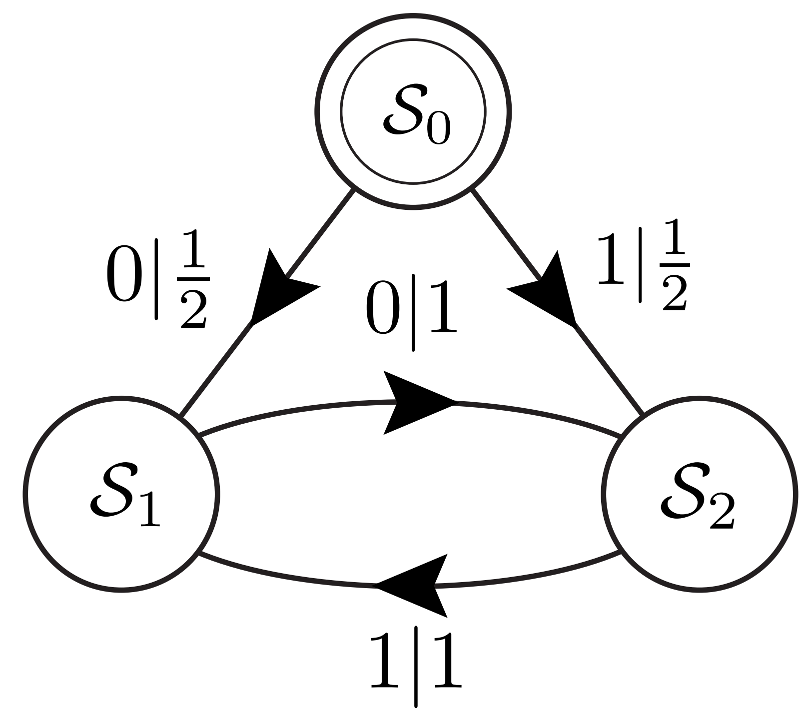

Causal State Distribution:

$$p_{\mathcal{S}} = \left(0, \tfrac{1}{2}, \tfrac{1}{2}\right)$$

Entropy rate: $h_\mu = 0$ bits per symbol

Statistical complexity: $C_\mu = 1$ bit

Machine Reconstruction ...

Examples ...

Golden Mean Process:

Parse Tree D = 5

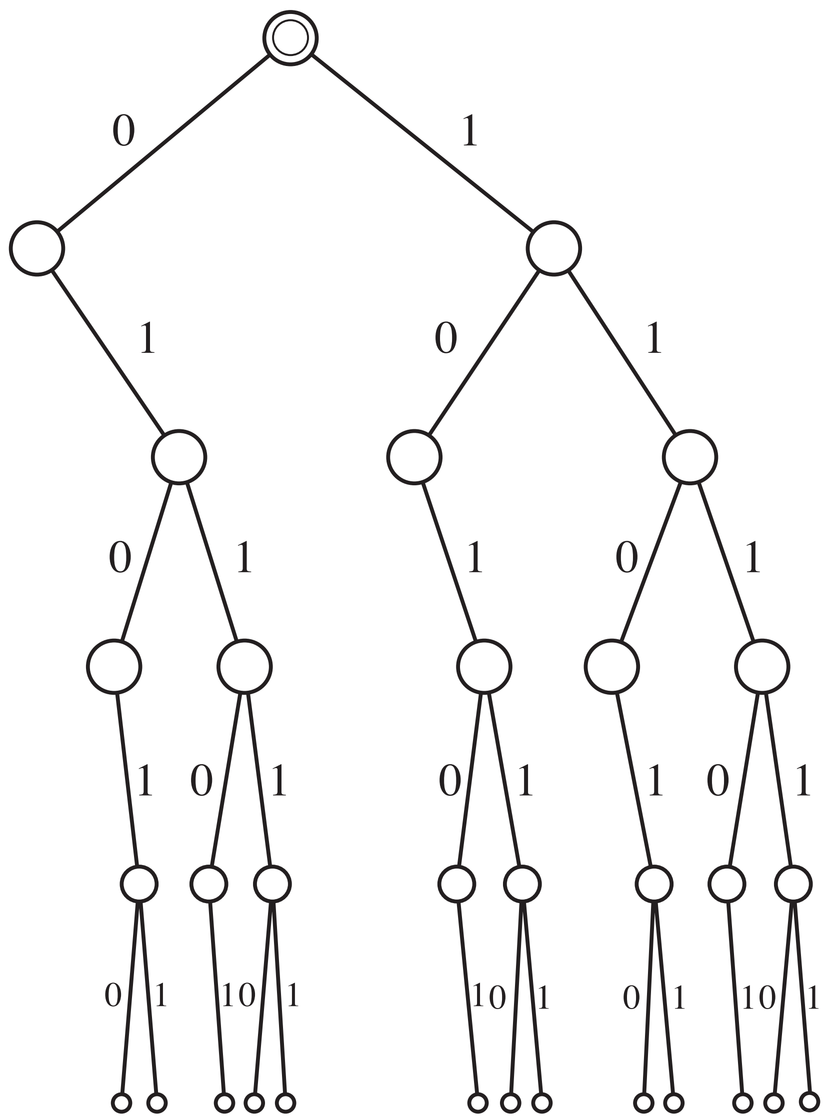

Machine Reconstruction ...

Examples ...

Golden Mean Process: Topological Reconstruction

Parse Tree D = 5

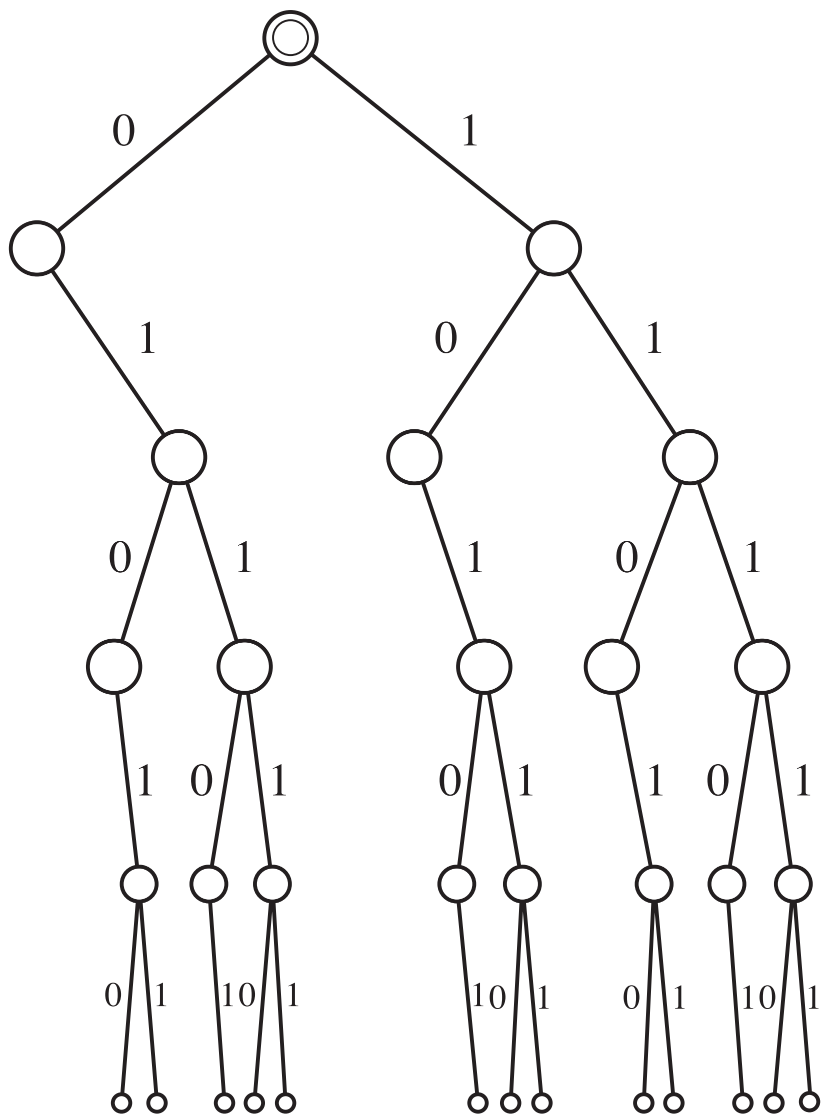

Morphs at L = 2

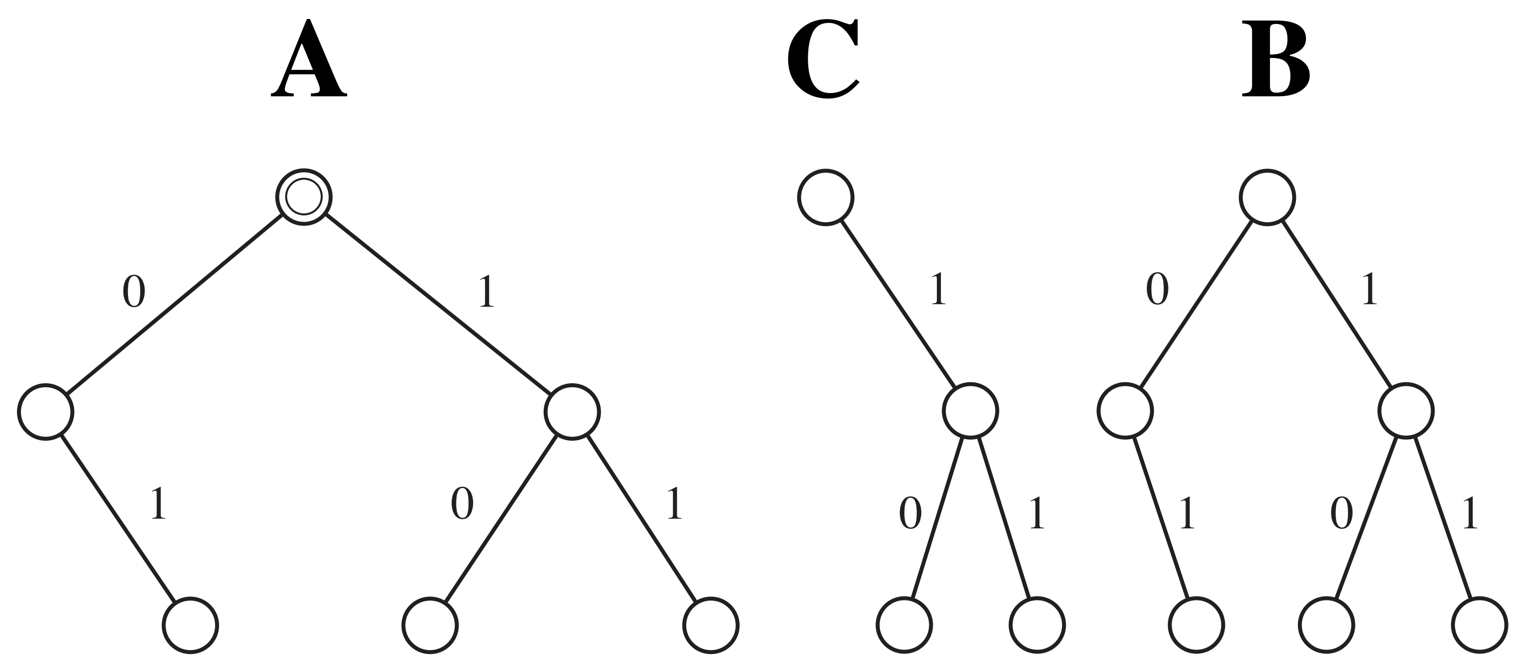

Examples ...

Golden Mean Process ...

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

$$\mathcal{S} = \{A, B, C\}$$

$$T^{(0)} = \begin{pmatrix} 0 & 0 & 1/3 \\ 0 & 0 & 1/2 \\ 0 & 0 & 0 \end{pmatrix}$$

$$T^{(1)} = \begin{pmatrix} 0 & 2/3 & 0 \\ 0 & 1/2 & 0 \\ 0 & 1 & 0 \end{pmatrix}$$

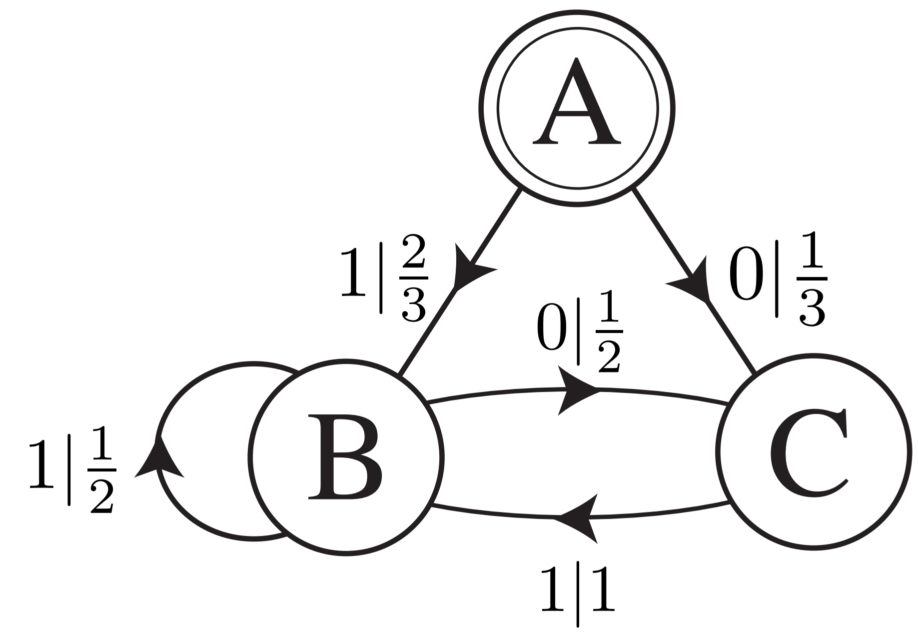

Machine Reconstruction ...

Examples ...

Golden Mean Process ...

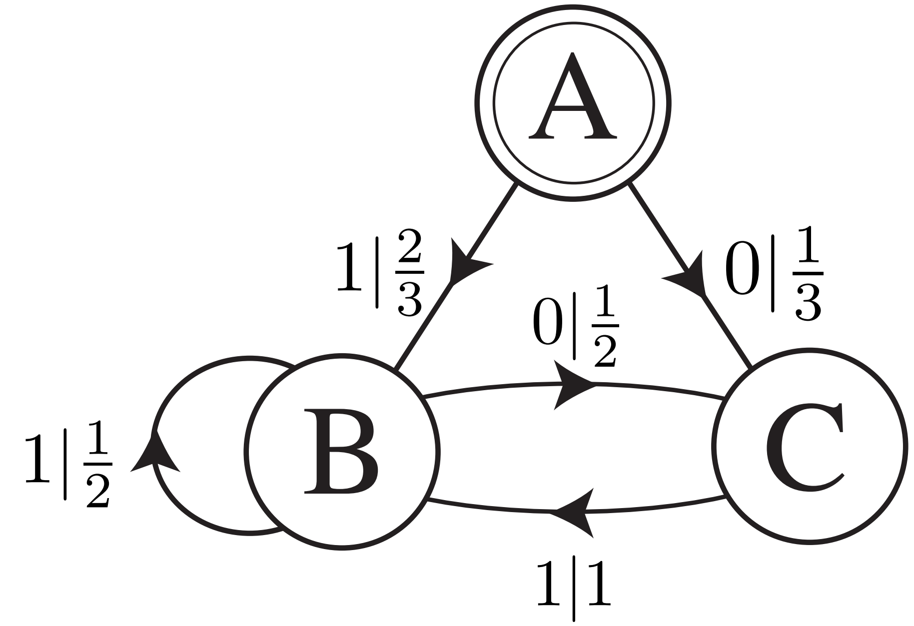

Causal State Distribution:

$$p_{\mathcal{S}} = \left(0, \tfrac{2}{3}, \tfrac{1}{3}\right)$$

Entropy rate: $h_\mu = \frac{2}{3}$ bits per symbol

Statistical complexity: $C_\mu = H(\frac{2}{3})$ bits

Machine Reconstruction ...

Examples ...

Even Process:

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

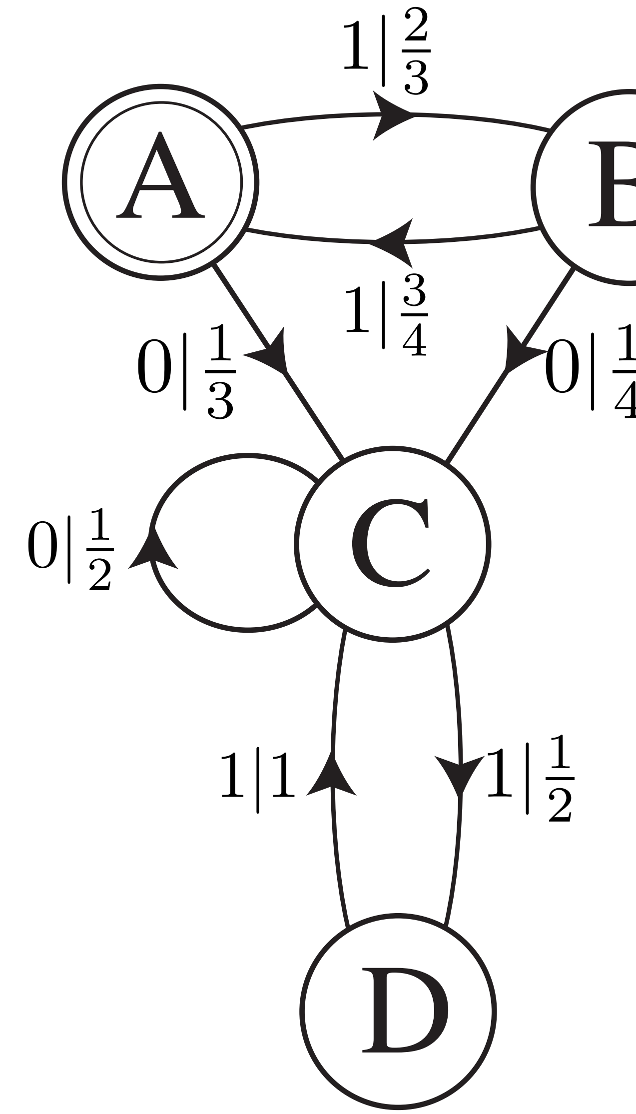

Entropy rate: $h_\mu = \frac{2}{3}$ bits per symbol

Statistical complexity: $C_\mu = H(\frac{2}{3})$ bits

The Learning Channel:

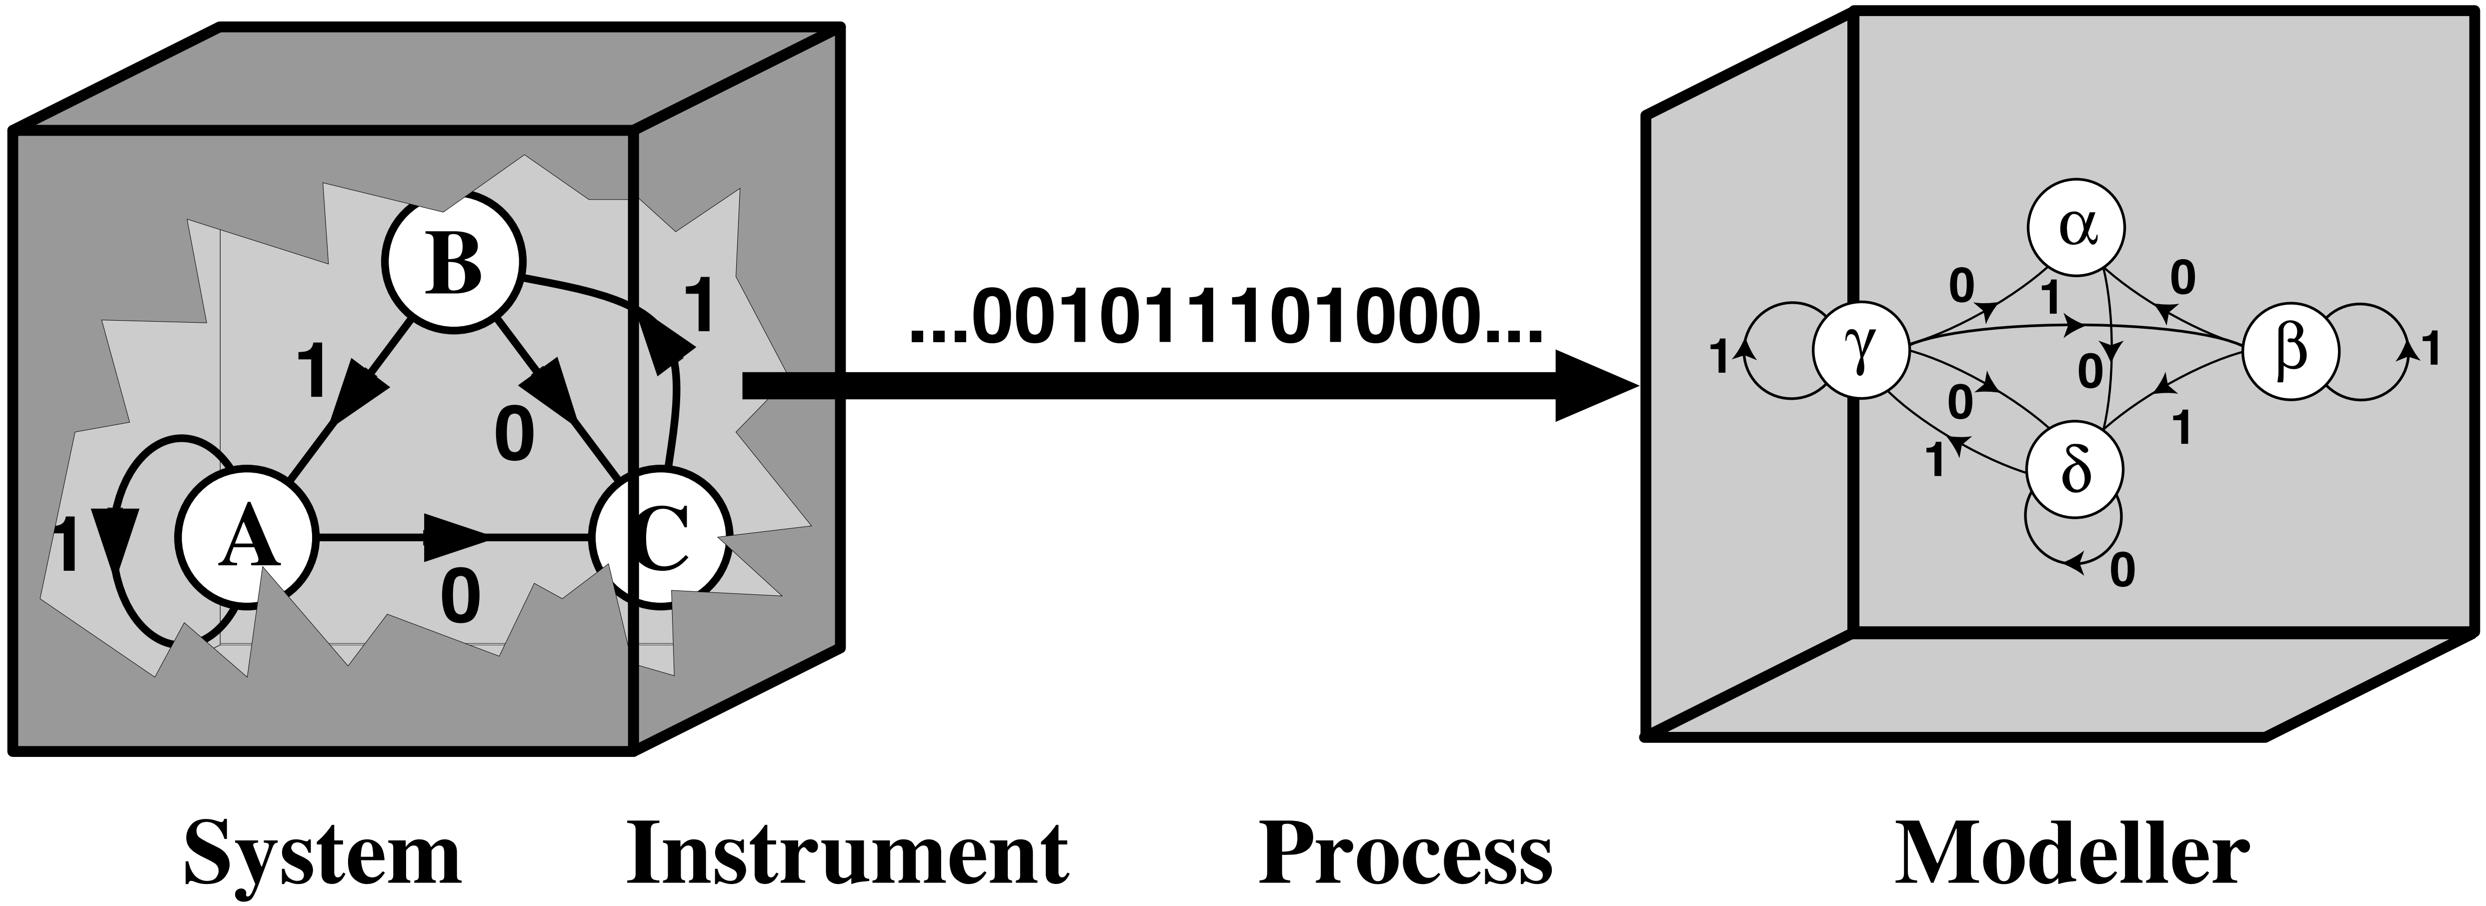

Central questions:

What are the states? Causal States

What is the dynamic? The ϵ-Machine

The ϵ-Machine ...

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

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

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

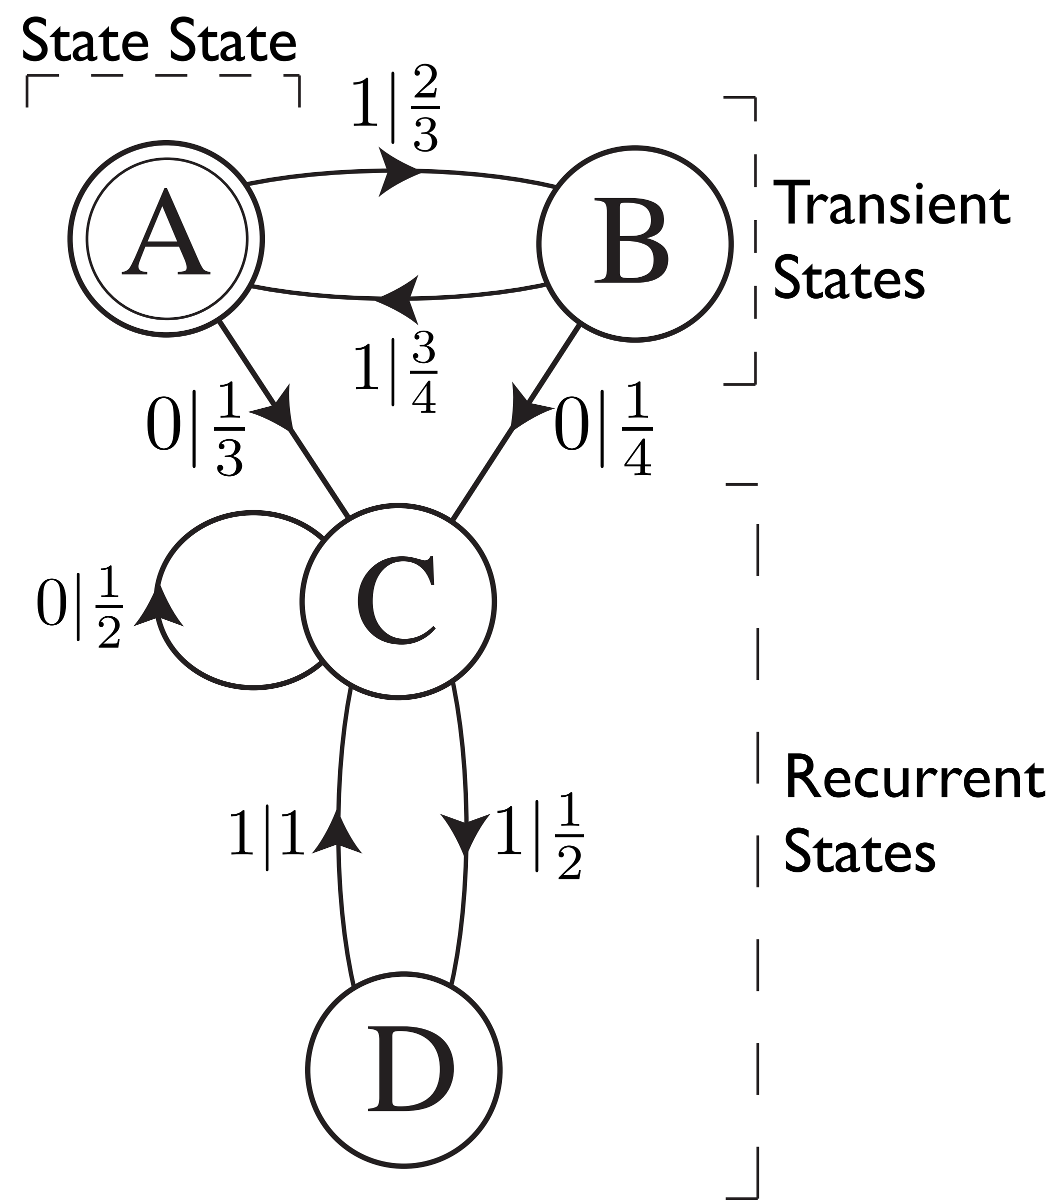

Unique Start State:

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

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

Transient States

Recurrent States

The ϵ-Machine ...

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

ϵ-Machine reproduces the process's word distribution:

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

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

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

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 ...

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

Calculate word distribution from recurrent states: $\mathcal{S}_i \in \mathcal{S}_{\mathrm{recurrent}}$

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

Get $\langle \pi | = (p_{\mathcal{S}_1}, p_{\mathcal{S}_2}, \ldots)$ from $T = \sum_{s \in \mathcal{A}} T^{(s)}$

Then

$$\Pr(s) = \langle \pi | T^{(s)} | 1 \rangle \qquad | 1 \rangle = \begin{pmatrix} 1 \\ 1 \\ \vdots \end{pmatrix}$$

$$\Pr(s_0 s_1) = \langle \pi | T^{(s_0)} T^{(s_1)} | 1 \rangle$$

...

$$\Pr(s^L) = \langle \pi | T^{(s^L)} | 1 \rangle$$

$$T^{(s^L)} = T^{(s_0)} T^{(s_1)} \ldots T^{(s_{L-1})}$$

The ϵ-Machine ...

Properties:

- Conditional independence of future & past
- Unifilar
- Markovian
- Optimal predictor
- Minimal size
- Unique

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 ...

Proof sketch:

$$\begin{aligned}\Pr(\overleftrightarrow{S}\,|\mathcal{S}) &= \Pr(\overleftarrow{S}\overrightarrow{S}\,|\mathcal{S}) \\ &= \Pr(\overrightarrow{S} \mid \overleftarrow{S}, \mathcal{S})\Pr(\overleftarrow{S}\,|\mathcal{S})\end{aligned}$$

Will show: $\Pr(\overrightarrow{S} \mid \overleftarrow{S}, \mathcal{S}) = \Pr(\overrightarrow{S}\,|\mathcal{S})$

$$\mathcal{S} = \epsilon(\overleftarrow{s}\,) \Rightarrow$$

$$\Pr\left(\overrightarrow{S} | \overleftarrow{S} = \overleftarrow{s}\,', \mathcal{S} = \epsilon(\overleftarrow{s}\,)\right) = \Pr\left(\overrightarrow{S} | \overleftarrow{S} = \overleftarrow{s}\right) \quad (\overleftarrow{s}\,' \in [\overleftarrow{s}\,])$$

But, also, $\Pr\left(\overrightarrow{S} \mid \overleftarrow{S} = \overleftarrow{s}\right) = \Pr\left(\overrightarrow{S}\,|\mathcal{S} = \epsilon(\overleftarrow{s}\,)\right)$ (Causal equiv. rel'n)

So, $\Pr\left(\overrightarrow{S} \mid \overleftarrow{S} = \overleftarrow{s}, \mathcal{S} = \sigma\right) = \Pr\left(\overrightarrow{S}\,|\mathcal{S} = \sigma\right)$ $\boxed{\circlearrowleft}$

The ϵ-Machine ...

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

(in automata theory, "deterministic")

That is:

(1) $\mathcal{S}_i \in \boldsymbol{\mathcal{S}}, \ s \in \mathcal{A}$, at most one $\mathcal{S}_j \in \boldsymbol{\mathcal{S}}$:

$$\overleftarrow{s} \in \mathcal{S}_i \Rightarrow \overleftarrow{s} s \in \mathcal{S}_j$$

(2) If there is a next causal state j:

$$\mathcal{S}_{k \neq j} \in \boldsymbol{\mathcal{S}} \Rightarrow T_{ik}^{(s)} = 0$$

(3) If there is not:

$$T_{ij}^{(s)} = 0$$

The ϵ-Machine ...

Unifilarity ...

Proof sketch:

Must show $\overleftarrow{s} \sim \overleftarrow{s}' \Rightarrow \overleftarrow{s}\, s \sim \overleftarrow{s}'\, s$

Futures with symbol prefixed: sF $\quad F \subset \mathcal{A}^\infty$

$$\overleftarrow{s} \sim \overleftarrow{s}' \quad \Rightarrow \quad \Pr\left(\overrightarrow{S} \in sF | \overleftarrow{S} = \overleftarrow{s}\right) = \Pr\left(\overrightarrow{S} \in sF | \overleftarrow{S} = \overleftarrow{s}'\right)$$

$$\Pr\left(\overrightarrow{S}^1 = s, \overrightarrow{S}_1 \in F | \overleftarrow{S} = \overleftarrow{s}\right) = \Pr\left(\overrightarrow{S}^1 = s, \overrightarrow{S}_1 \in F | \overleftarrow{S} = \overleftarrow{s}'\right)$$

$$\Pr\left(\overrightarrow{S}_1 \in F | \overrightarrow{S}^1 = s, \overleftarrow{S} = \overleftarrow{s}\right) \Pr\left(\overrightarrow{S}^1 = s | \overleftarrow{S} = \overleftarrow{s}\right) = \Pr\left(\overrightarrow{S}_1 \in F | \overrightarrow{S}^1 = s, \overleftarrow{S} = \overleftarrow{s}'\right) \Pr\left(\overrightarrow{S}^1 = s | \overleftarrow{S} = \overleftarrow{s}'\right)$$

$$\Pr\left(\overrightarrow{S}_1 \in F | \overleftarrow{S} = \overleftarrow{s}\, s\right) = \Pr\left(\overrightarrow{S}_1 \in F | \overleftarrow{S} = \overleftarrow{s}'\, s\right)$$

(Stationarity and by assumption $\Pr(\overrightarrow{S}^1 = s | \overleftarrow{s}) = 1$ $\Pr(\overrightarrow{S}^1 = s | \overleftarrow{s}') = 1$)

$$\Rightarrow \overleftarrow{s}\, s \sim \overleftarrow{s}'\, s \qquad \boxed{\circlearrowright}$$

Complexity Lecture 2: Computational Mechanics (CSSS 2011); Jim Crutchfield

The ϵ-Machine ...

Unifiliarity ...

Consequence:

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

Entropy rate expression requires this 1-1 mapping.

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

The ϵ-Machine ...

ϵMs are Optimal Predictors:

Compared to any rival effective states R:

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

Proof sketch:

$$H\left[\overrightarrow{S}^{L} | \mathcal{S}\right] = H\left[\overrightarrow{S}^{L} | \overleftarrow{s} \in \mathcal{S}\right]$$

(Causal equiv. rel'n)

$$= H\left[\overrightarrow{S}^{L} | \overleftarrow{s}\right]$$

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

$$R = \eta(\overleftarrow{s})$$

(Data processing inequality)

The ϵ-Machine ...

$\epsilon\mathrm{Ms}$ are Optimal Predictors ...

Lemma:

$$h_\mu(\mathcal{S}) = h_\mu$$

Proof:

$$\begin{aligned}
h_\mu(\mathcal{S}) &= \lim_{L\to\infty} \frac{1}{L} H\left[\overrightarrow{S}^L | \mathcal{S}\right] && \text{(Block entropy)} \\
&= \lim_{L\to\infty} \frac{1}{L} H\left[\overrightarrow{S}^L | \overleftarrow{S}\right] && \text{(Causal equiv. rel'n)} \\
&= \lim_{L\to\infty} \frac{1}{L} L H\left[S | \overleftarrow{S}\right] && \text{(Stationarity)} \\
&= H\left[S | \overleftarrow{S}\right] \\
&= h_\mu && \boxed{\circlearrowleft}
\end{aligned}$$

Corollary: $h_\mu(R) \geq h_\mu$

The ϵ-Machine ...

ϵMs are Optimal Predictors ...

Corollary (Maximal Prescience): $\Pi(R) \leq \Pi(\mathcal{S})$

Rival model: $\Pi(R) = \log_2 |\mathcal{A}| - h_\mu(R)$

But: $\Pi(\mathcal{S}) = \log_2 |\mathcal{A}| - h_\mu = \mathbf{G}$

So: $\Pi(R) \leq \Pi(\mathcal{S})$ $\qquad h_\mu(R) \geq h_\mu$

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. (See below.)

The ϵ-Machine ...

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

Alternative models that are optimal predictors

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

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

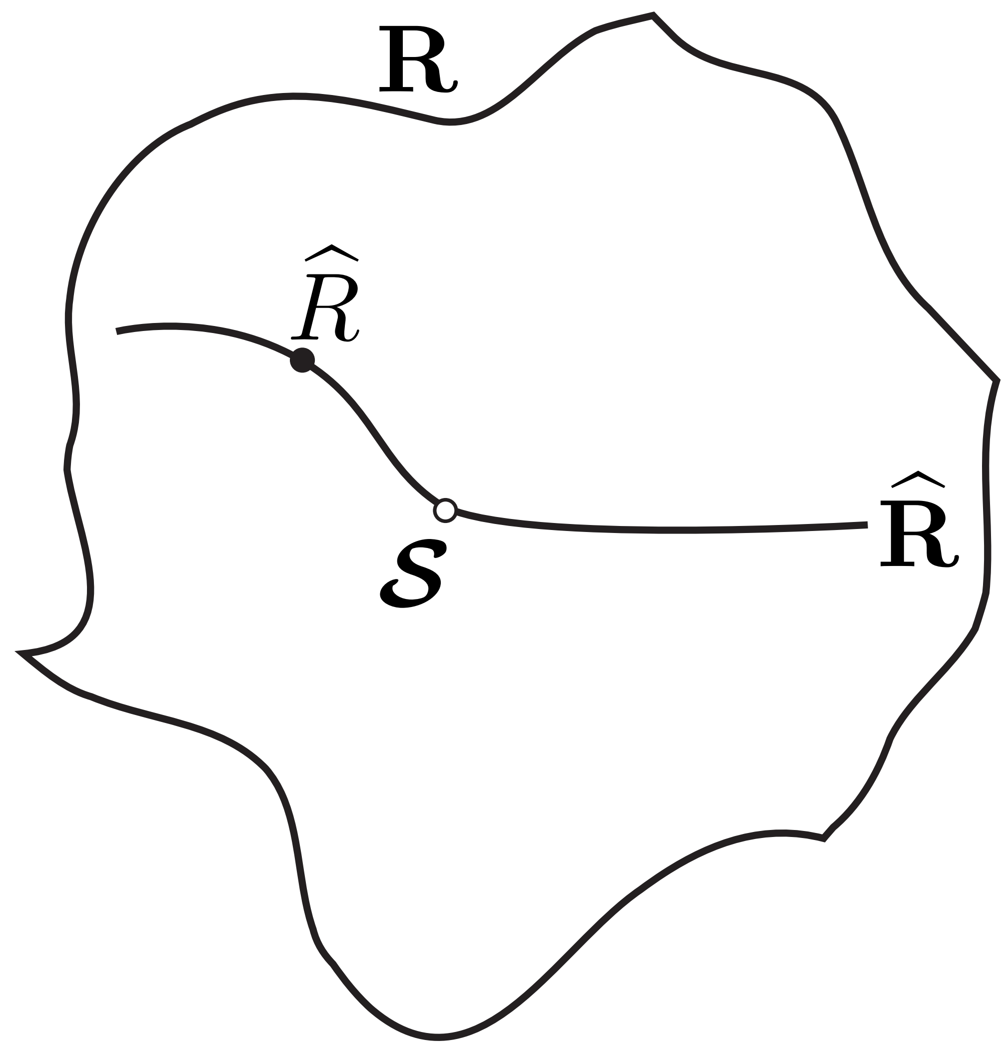

(Prescient rivals are sufficient statistics for process's future.)

The ϵ-Machine ...

Minimal Statistical Complexity:

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

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

Proof sketch:

(1) Prescient rivals are refinements, so

$$\exists g : \mathcal{S} = g(\widehat{R})$$

(2) But

$$H[f(X)] \leq H[X] \Rightarrow H[\mathcal{S}] = H[g(\widehat{R})] \leq H[\widehat{R}]$$

(3) So $C_\mu \leq H[\widehat{R}]$

The ϵ-Machine ...

Minimal Statistical Complexity ...

Consequence:

(1) C_μ measures historical information process stores.

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

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.

The ϵ-Machine ...

Dynamical system's intrinsic computation:

(1) How much of past does process store?

(2) In what architecture is that information stored?

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