# Structure or Noise?

Dynamics Days 2008
Knoxville, Tennessee

James P. Crutchfield
Complexity Sciences Center, Director
Physics Department
University of California, Davis
cse.ucdavis.edu/~chaos

5 January 2008

# Why We Must Model I

- Nature spontaneously organizes
- Emergent structures

# Emergent structures

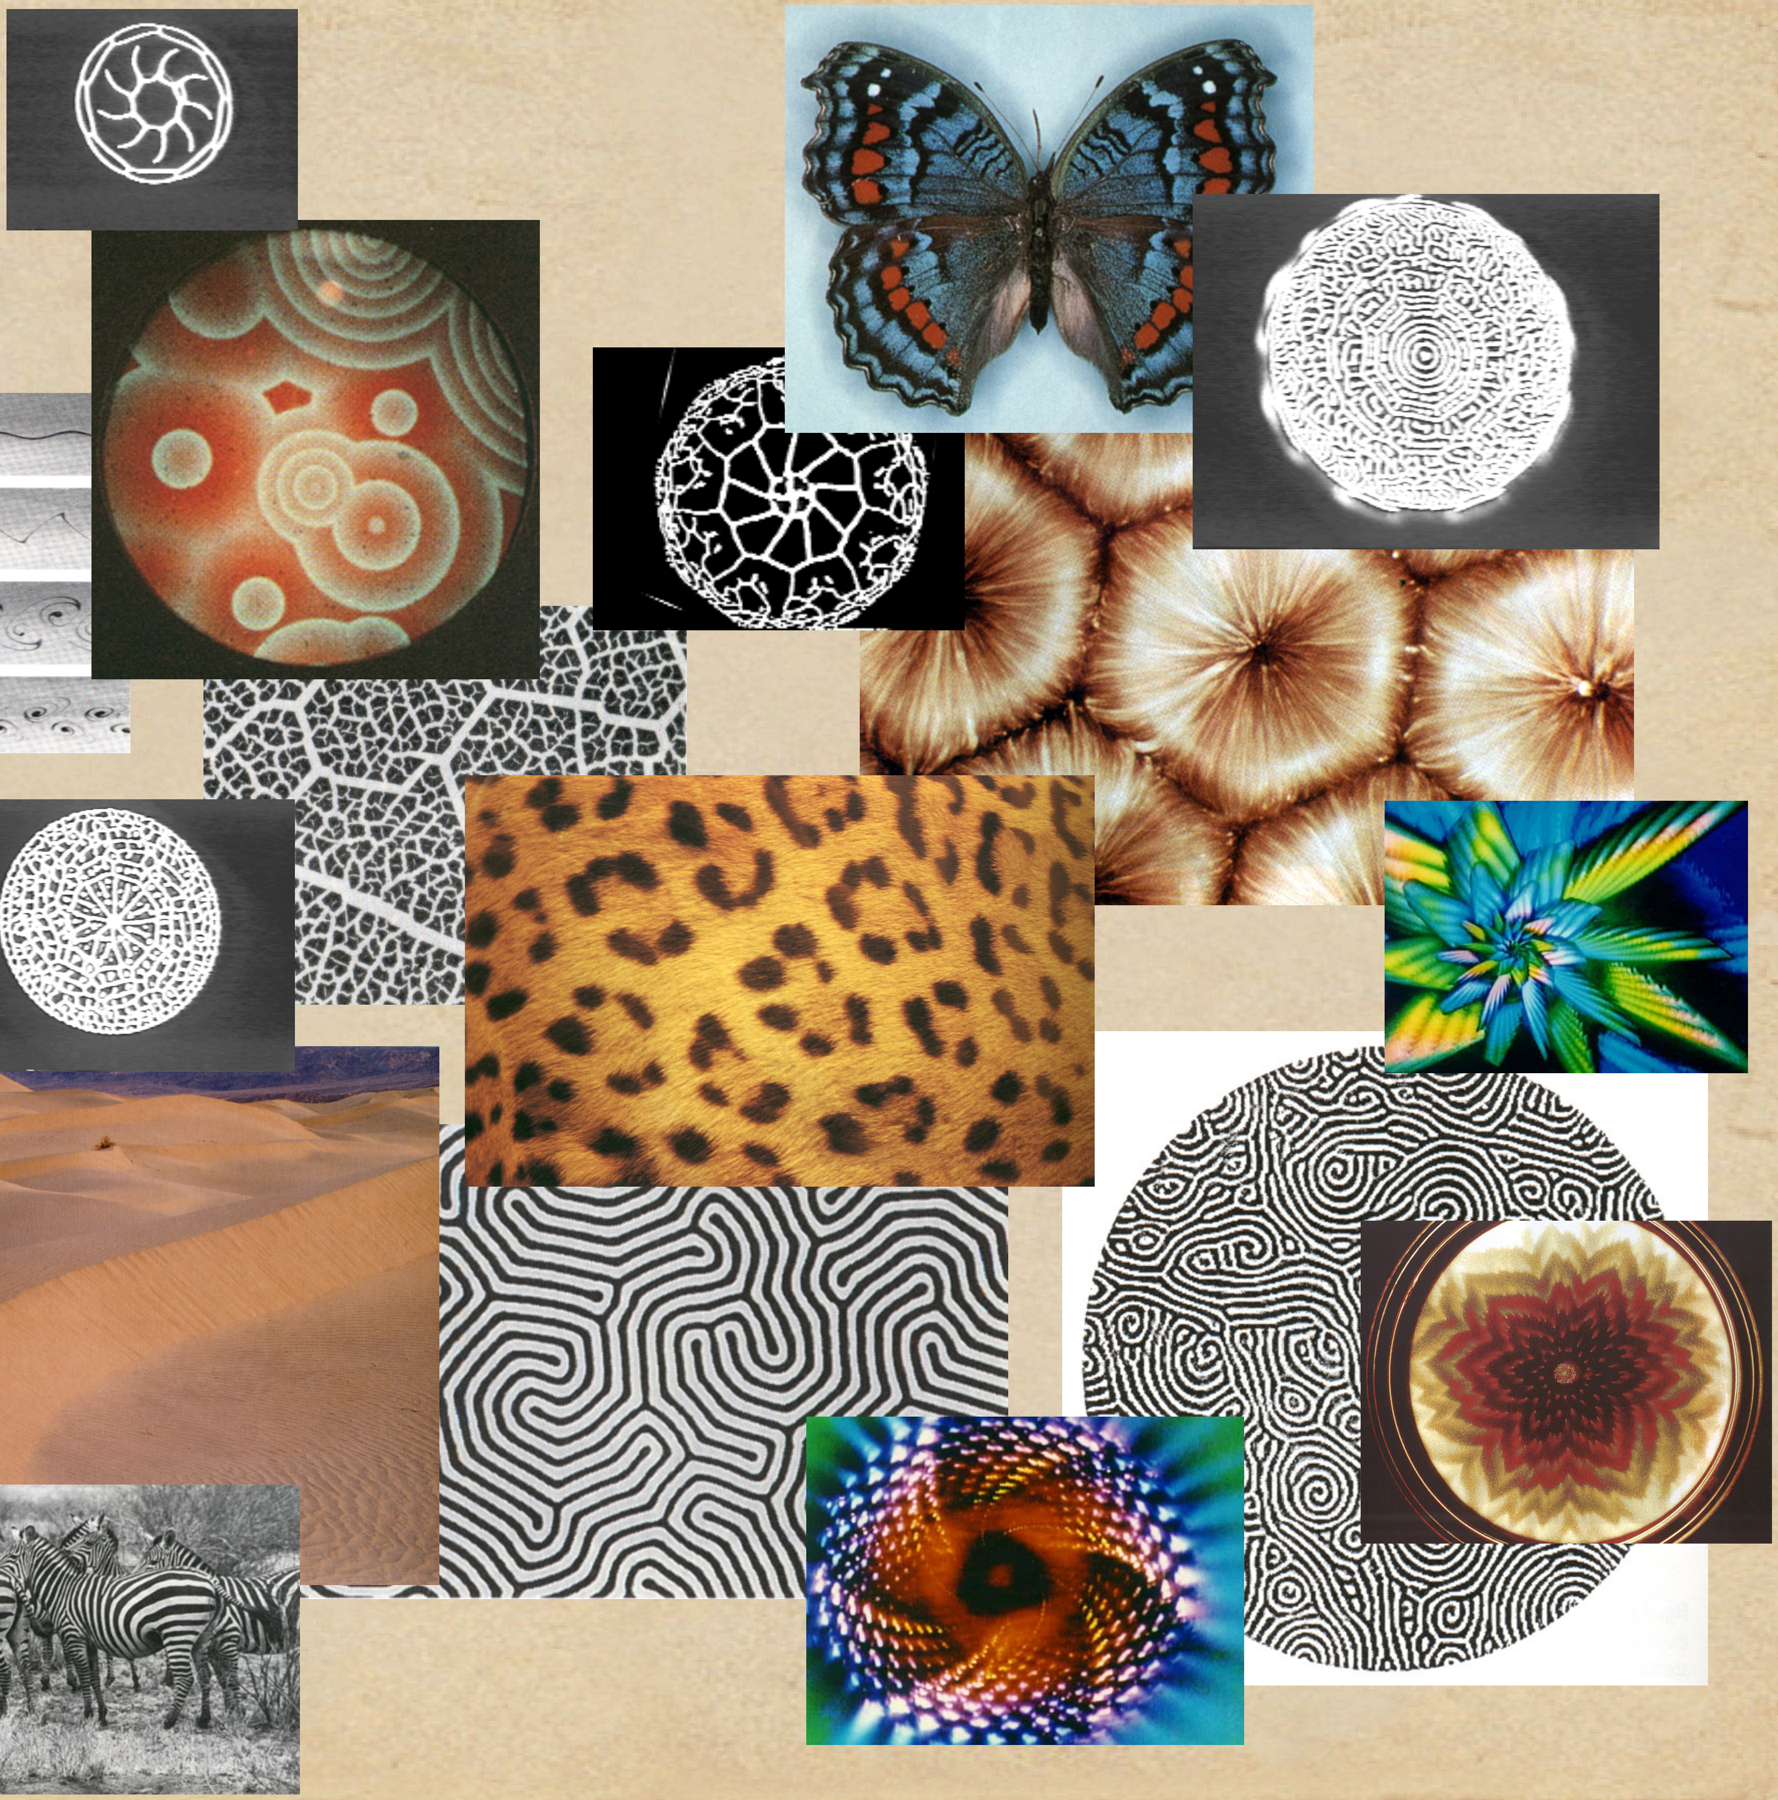

# Why We Must Model 2

- Engineered systems also spontaneously organize
  - Internet route flapping
  - Power-law Internet organization
  - Financial markets crash
  - Power grids fail spectacularly
  - Social pattern formation on the web
  - ...

# Consequence

- Each needs its own explanatory (function) basis
- Problem:

Emergent structures not given directly by the system coordinates or the governing equations of motion

# Why we must Model 3

- Fundamental Mathematics: Intrinsic Randomness
  - Nonlinear dynamical systems [Kolmogorov 1958]:
    - Chaotic systems: Shannon entropy $h_\mu > 0$
  - Kolmogorov-Chaitin [1963] complexity of Data:
    - Size of shortest Turing Machine Program to predict Data
  - KC complexity = Shannon entropy [Brudno 1978] :

$$|\mathrm{Program}| \propto e^{h_\mu |\mathrm{Data}|}$$

# Exponential Increase in Prediction Resources

$$\text{Accuracy} \propto e^{-T}$$

$$|\text{Measurements}| \propto e^{T}$$

$$|\text{Compute time}| \propto e^{T}$$

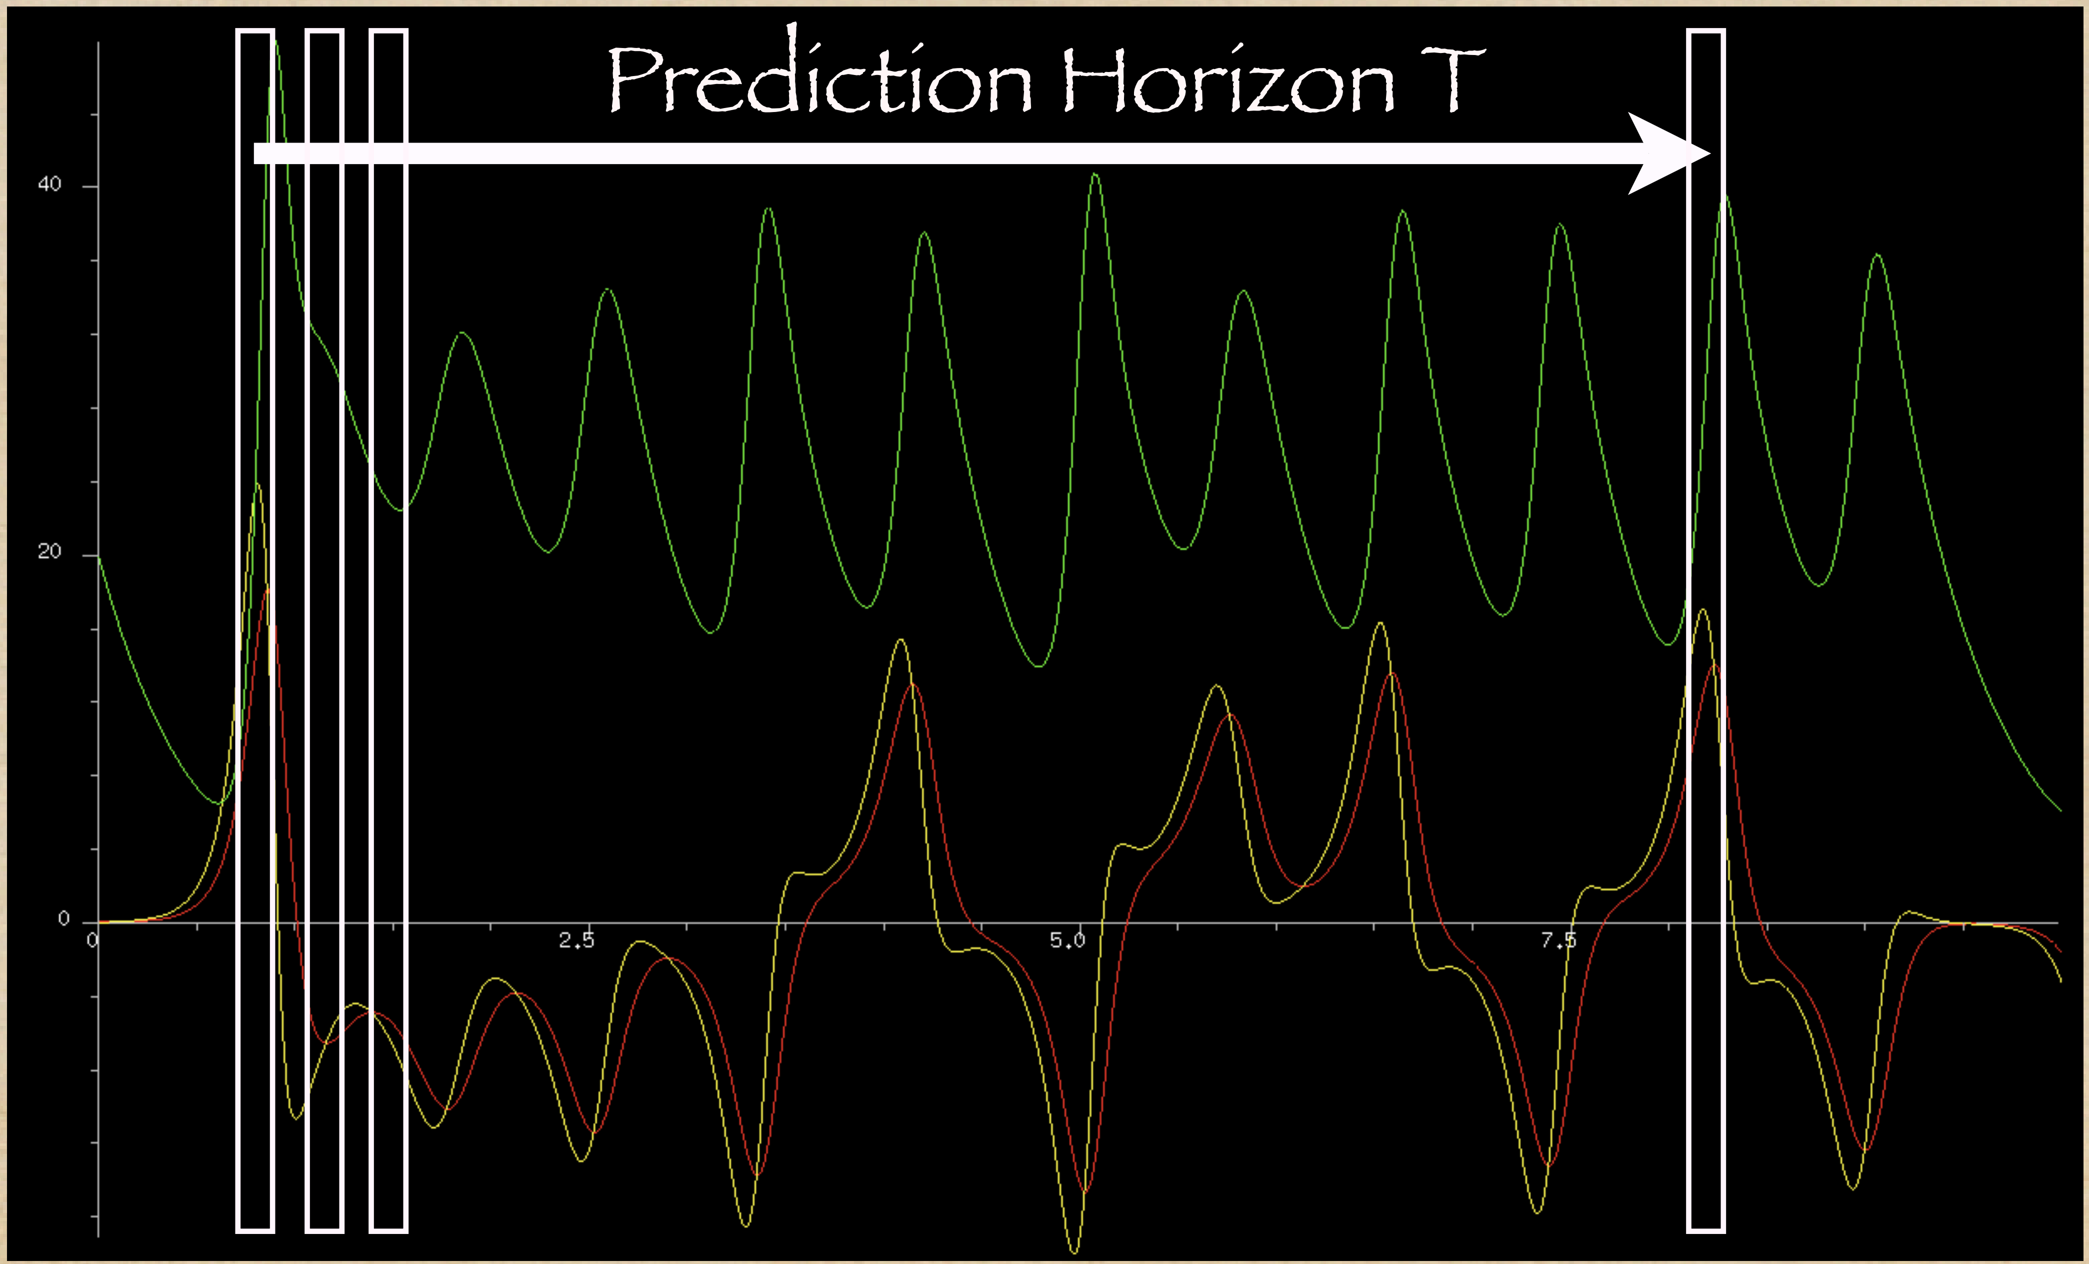

# Consequence

- No short cuts!
  - No closed-form solutions
  - No computational speed-ups
  - Must compute full trajectory
- Right representation is critical for reducing the prediction error as far as possible (but no farther!)

# Past:

# Fundamental in Nonlinear Dynamics!

- Each nonlinear system requires its own representation
- Selecting balance between ascribing structure or noise to a measurement depends on representation
- Fundamental issue: Theory building
- Subsidiary issue:

  Statistical fluctuations due to finite data sample

# Future:
# Fundamental in Designing Multiagent Systems

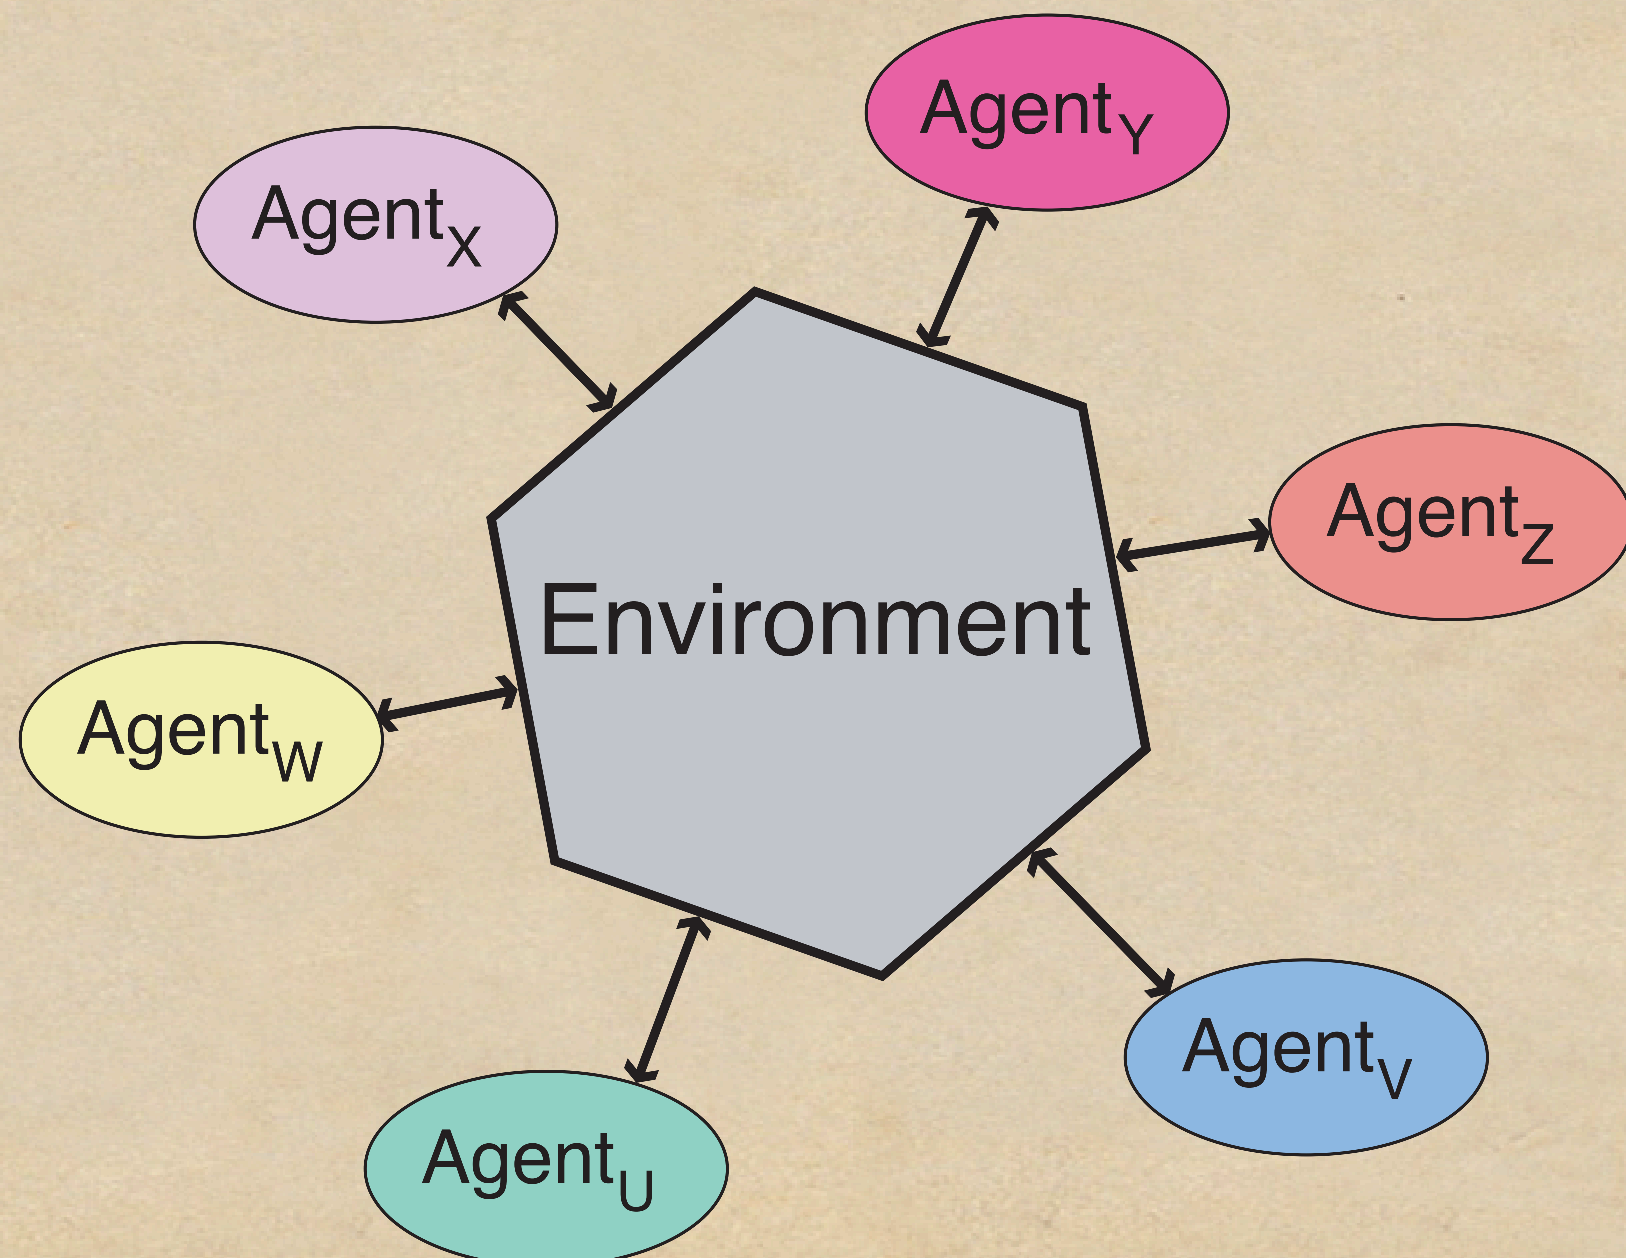

# The Feedback Loop

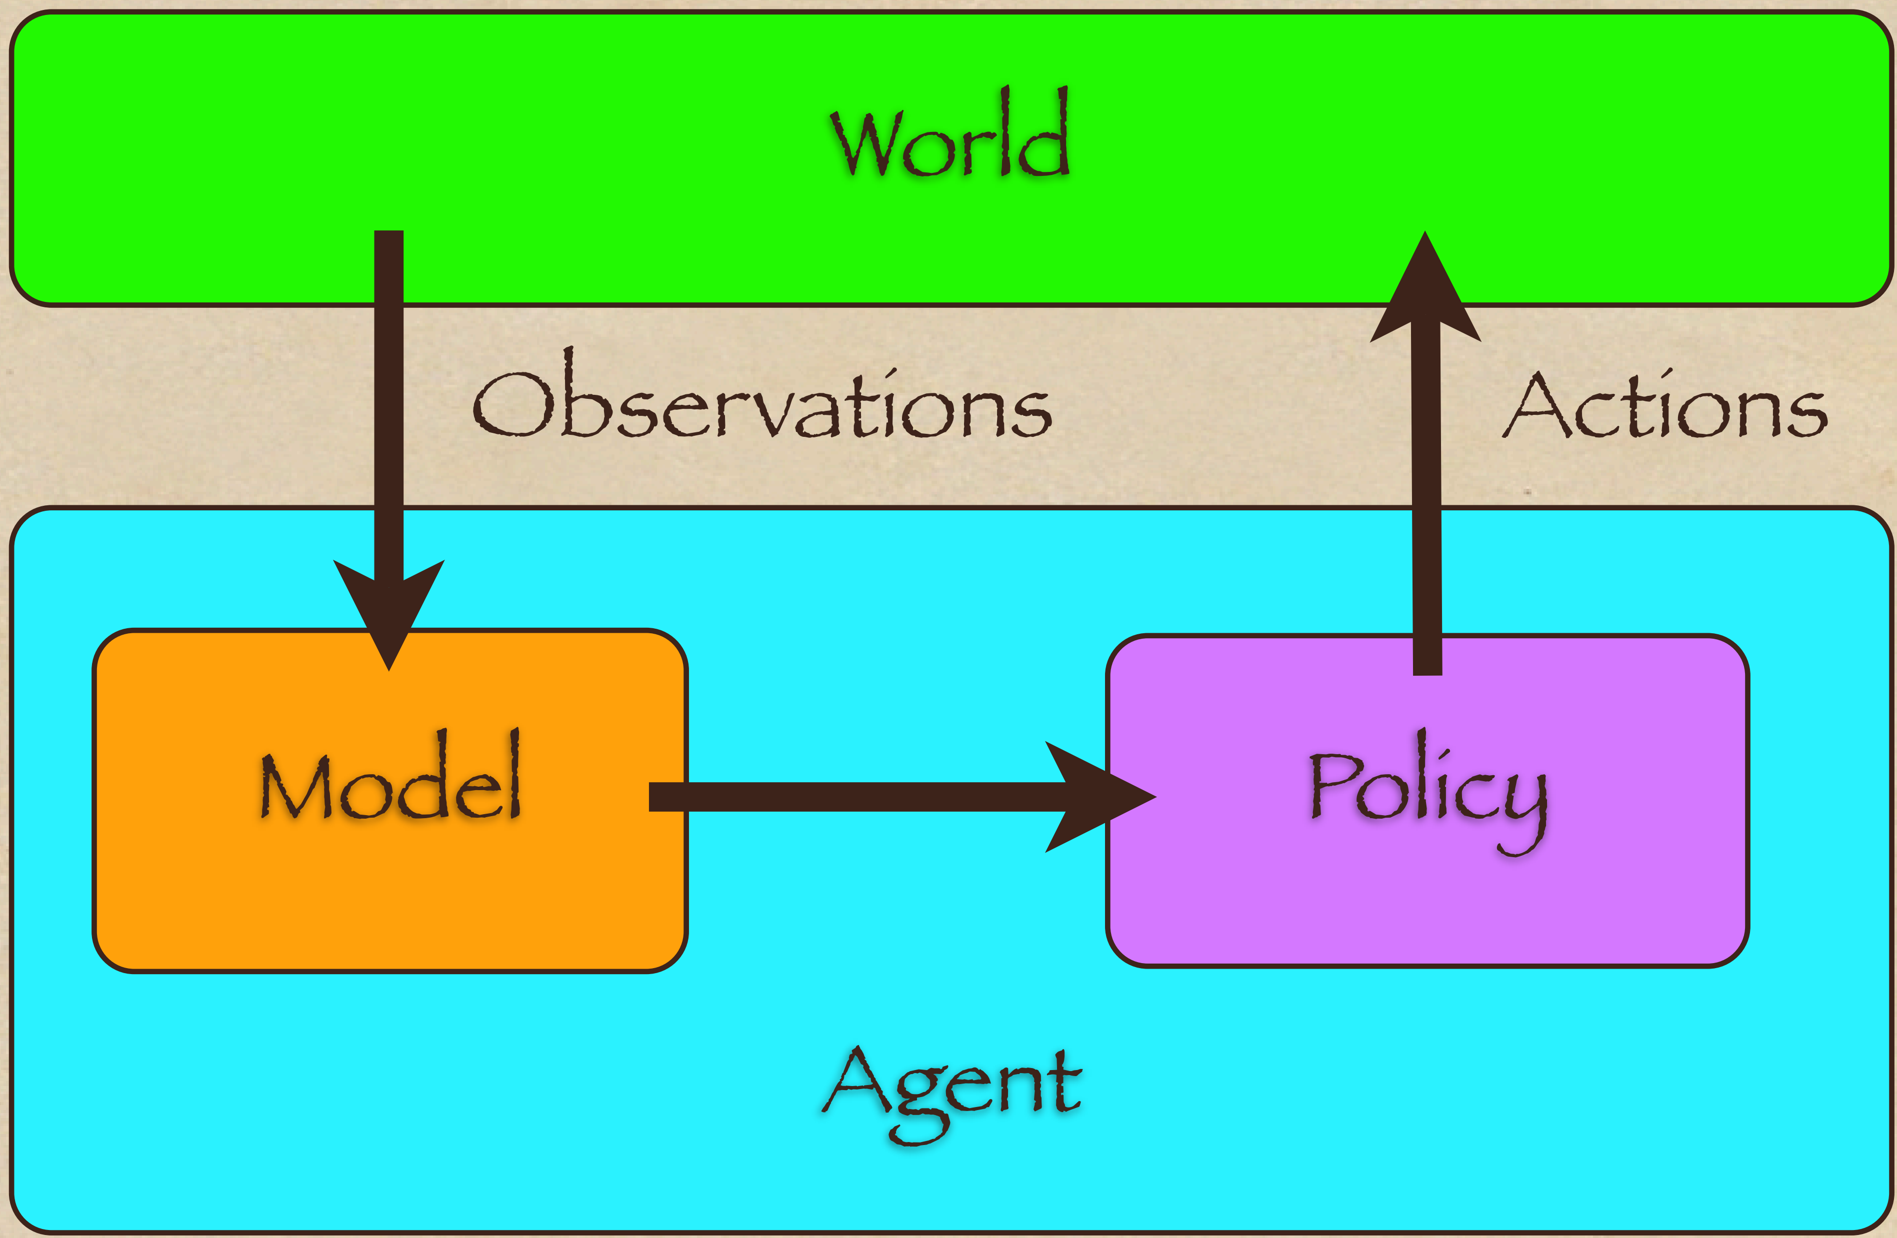

# Knowledge + Action

- A central challenge:

  Actions change the world

  and so

  its statistics,

  and

  what is knowable.

# Approaches

- Modeling:
  - Statistical inference
- Strategizing:
  - Game theory
- Adapting:
  - Reinforcement learning
- Group behavior:
  - Population dynamics (evolution & ecology)
- ...

# Approaches: Sticking points

- Modeling:
  - Statistical inference: static, batch mode
- Strategizing:
  - Game theory: equilibria, no transients
- Adapting:
  - Reinforcement learning: a priori design, brittle
- Group behavior:
  - Population dynamics (evolution & ecology): individuals have no structure (don't learn)
- Where are the basic principles?

# Interactive Learning

(Susanne Still, Chris Ellison, & JPC)

- Problem: Experiment to Learn World Model
  - The world behaves: $\overleftrightarrow{X} = \underset{\text{past}}{\overleftarrow{X}} \; \underset{\text{future}}{\overrightarrow{X}}$
  - Agent learns model of the world: States $\mathcal{R}$
  - Agent take actions $\mathcal{A}$
  - Those actions affect the world
  - Now the world is different!
- How to close the feedback loop?

arxiv.org: [0708.0654](0708.0654) [physics.gen-ph] & [0708.1580](0708.1580)[cs.IT]

# Passively Learning a Model

- Pattern discovery:
  - Learn the world's hidden states $\Pr(\mathcal{R}| \overleftarrow{X})$
- Causal shielding:

$$\Pr(\overleftarrow{X}\overrightarrow{X}) = \Pr(\overleftarrow{X}|\mathcal{R})\Pr(\overrightarrow{X}|\mathcal{R})$$

- Search in the space of models: $\mathcal{R} \in \mathcal{M}$
- Objective function

$$\min_{\Pr(\mathcal{R}|\overleftarrow{X})} \left( I[\overleftarrow{X};\mathcal{R}] + \beta I[\overleftarrow{X};\overrightarrow{X}|\mathcal{R}] \right)$$

$\beta \sim 1/T$

Model: Map from histories to states

Info states contain about histories

Reduce info history has about future

# Passively Learning a Model

- Optimal states $\Pr(\mathcal{R}|\overleftarrow{X})$ are Gibbs distributions:

$$\Pr_{\mathrm{opt}}(\mathcal{R}|\overleftarrow{X}) = \frac{\Pr(\mathcal{R})}{Z(\overleftarrow{X},\beta)} e^{-\beta E(\mathcal{R},\overleftarrow{X})}$$

where

$$E(\mathcal{R},\overleftarrow{X}) = \mathcal{D}\left(\Pr(\overrightarrow{X} \mid \overleftarrow{X}) || \Pr(\overrightarrow{X} | \mathcal{R})\right)$$

$$\Pr(\overrightarrow{X} | \mathcal{R}) = \frac{1}{\Pr(\mathcal{R})} \sum_{\overleftarrow{X}} \Pr(\overrightarrow{X} \mid \overleftarrow{X}) \Pr(\mathcal{R}|\overleftarrow{X}) \Pr(\overleftarrow{X})$$

$$\Pr(\mathcal{R}) = \sum_{\overleftarrow{X}} \Pr(\mathcal{R}|\overleftarrow{X}) \Pr(\overleftarrow{X})$$

# Passively Learning a Model

- Solve these equations self-consistently

- Parametrized family of models $\Pr(\mathcal{R}| \overleftarrow{X})$: $R_\beta$

- Structure or Noise?

  $\beta$ trades-off model size against prediction error

# What Do Solutions Mean?
# Causal Models

- Causal architecture given by $\epsilon$-Machine $M$:
  - Optimal predictor:

    $$h_\mu(M) \leq h_\mu(\mathcal{R})$$

  - Minimal size (within optimal predictors $\widehat{\mathcal{R}}$):

    $$C_\mu(M) \leq C_\mu(\widehat{\mathcal{R}})$$

  - Unique (within min, opt predictors)

JPC & K. Young, Inferring Statistical Complexity, Physical Review Letters 63 (1989) 105-108.
C. R. Shalizi & JPC, Journal Statistical Physics 104 (2001) 817-879.

# Passively Learning a Model

- Theorem: Low-temperature limit

$$\beta \rightarrow \infty$$

Recover $\epsilon$-Machine:

$$R_\beta \rightarrow M$$

- Conclusion: Best causal approximates.

# Passively Learning a Model

## Optimal balance structure & error
## At each level $\beta$ of approximation

Causal Rate Distortion Curve

In theory

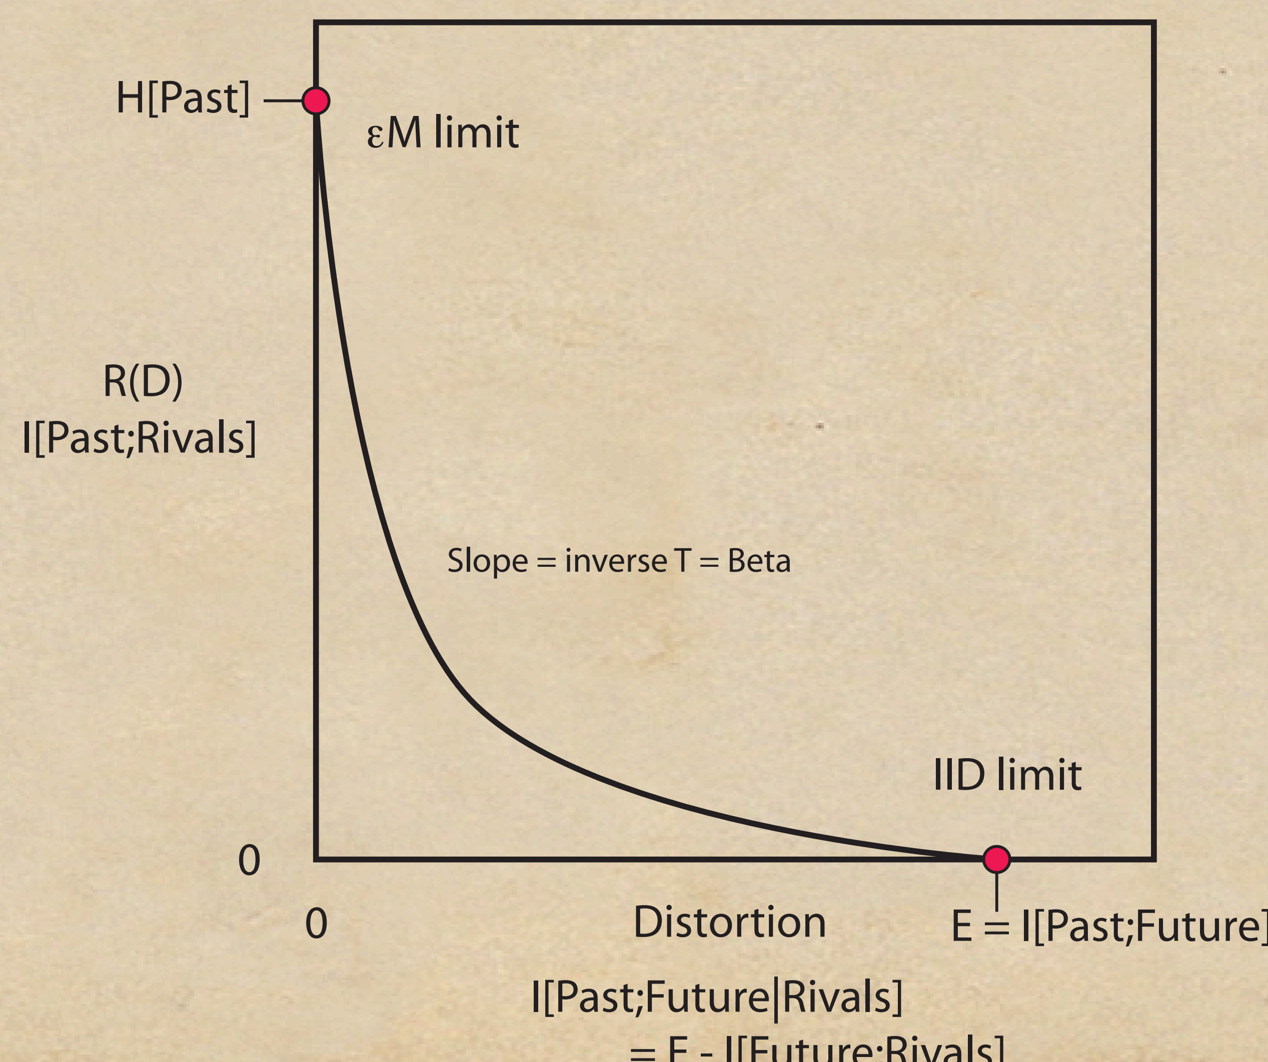

# Passively Learning a Model

## Optimal balance structure & error
## At each level $\beta$ of approximation

In practice: Learn an oo-state world (SNS: simple nondeterminstic source)

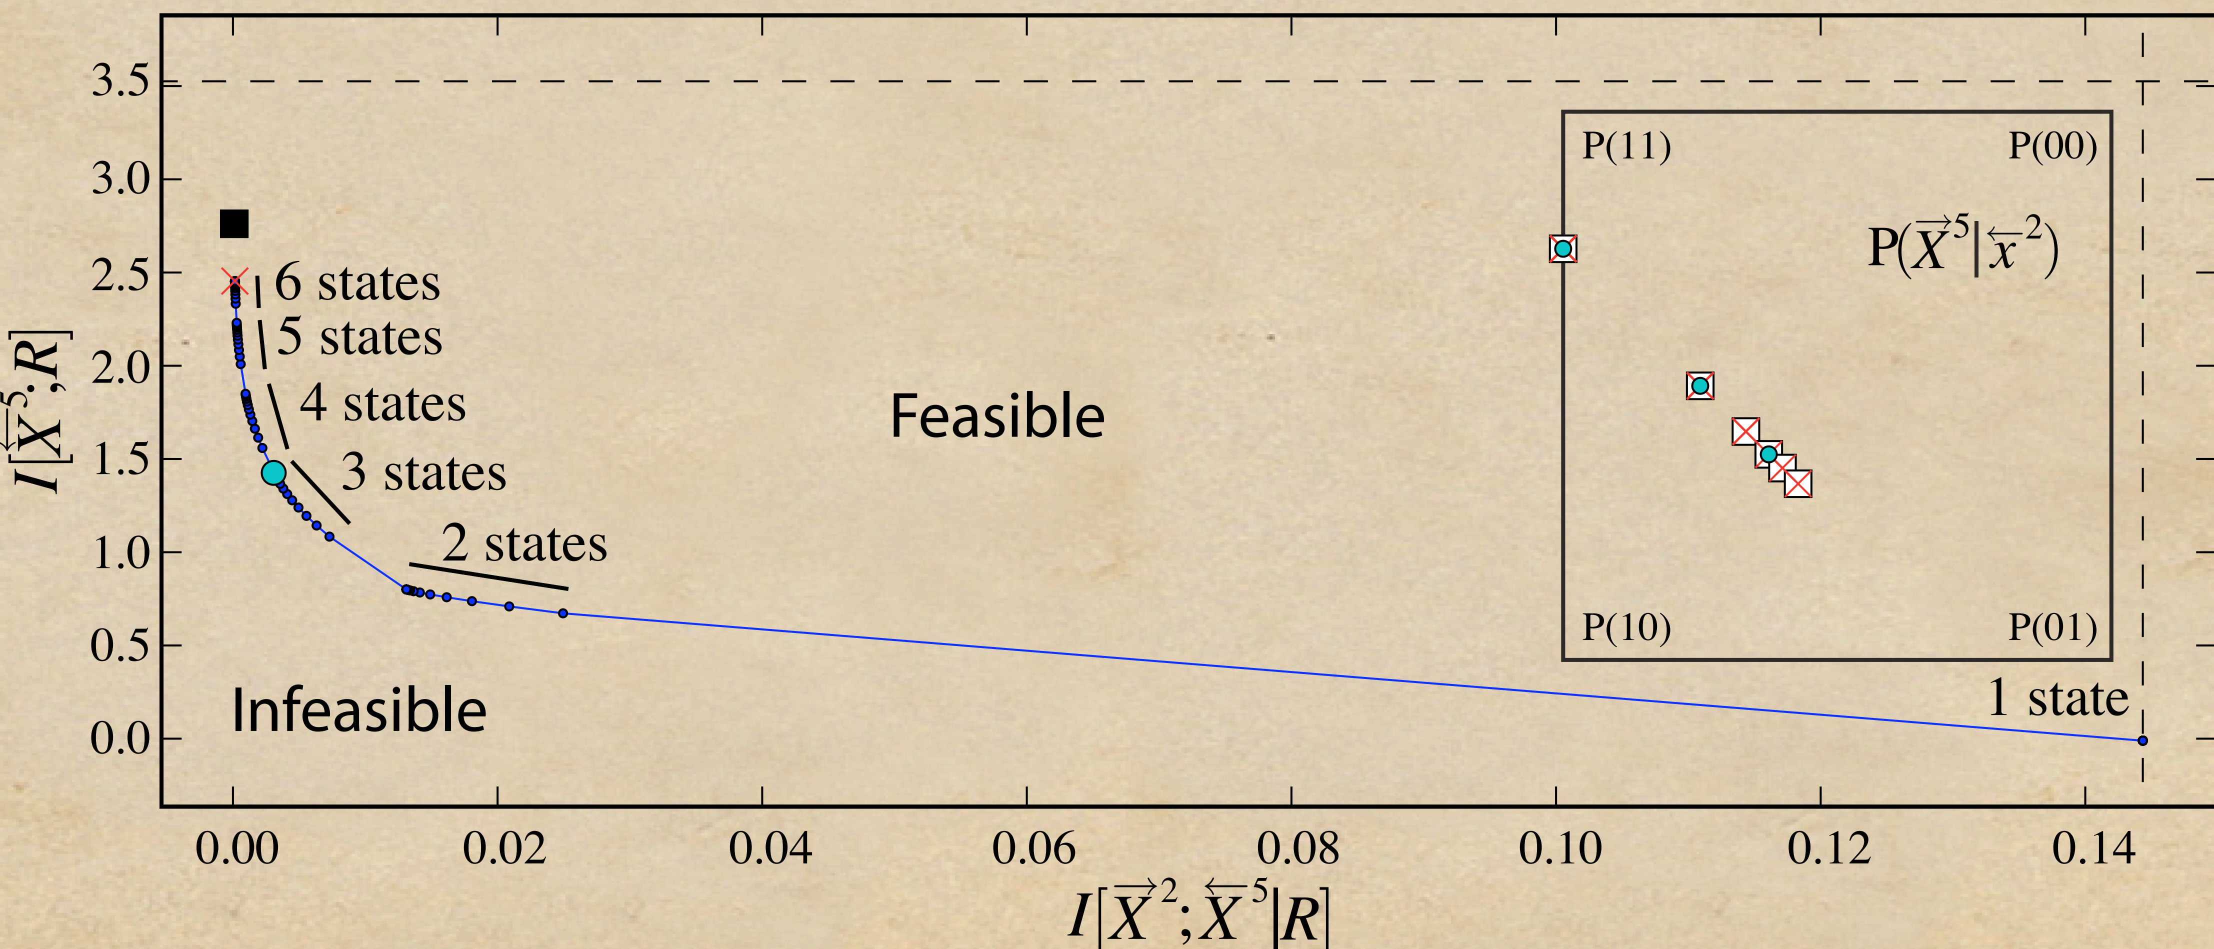

# Interactive Learning

- Decision: Using model, take actions
- Policy: $\Pr(\mathcal{A}|\overleftarrow{X})$ (or from $\mathcal{R}$ )
- Experimentation objective function

$$\max_{\Pr(\mathcal{R}|\overleftarrow{X}),\Pr(\mathcal{A}|\overleftarrow{X})} \Big( I[\{\mathcal{R},\mathcal{A}\};\overrightarrow{X}] - \lambda I[\mathcal{R};\overleftarrow{X}] - \mu I[\mathcal{A};\overleftarrow{X}] \Big)$$

Model: Map from histories to states

Policy: Map from histories to actions

Info states/actions contain about futures

Info states contain about histories

Info actions contain about histories

# Interactive Learning: Results

- Optimal model: Recover causal architecture
- Optimal policies
- Causally equivalent policies
- Curiosity: Take informative actions
- Control: Make world easier to model
- Balance of exploitation and control

arxiv.org: 0708.0654 [physics.gen-ph] & 0708.1580[cs.IT]

# Connections

- Note iLearning subsumes:
  - Causal modeling
  - Game theory
  - Equilibrium economics
  - Reinforcement learning

# Knowledge + Action

- Deviation from knowledge:
  - Model evaluation (e.g., prediction error)
- Valuation: Commitment of resources to action
  - Policy evaluation (e.g., average reward)
- How? Augment objective function
  - Change relative weighting of Lagrange multipliers: $\lambda$ & $\mu$
  - Add new terms: e.g., ...
- Examples:
  - "Science": Need accurate knowledge, at expense of producing it
  - "Politics": Need world to behave, independent of knowledge or cost
  - These are positions on causal rate-distortion curve

# The Feedback Loop

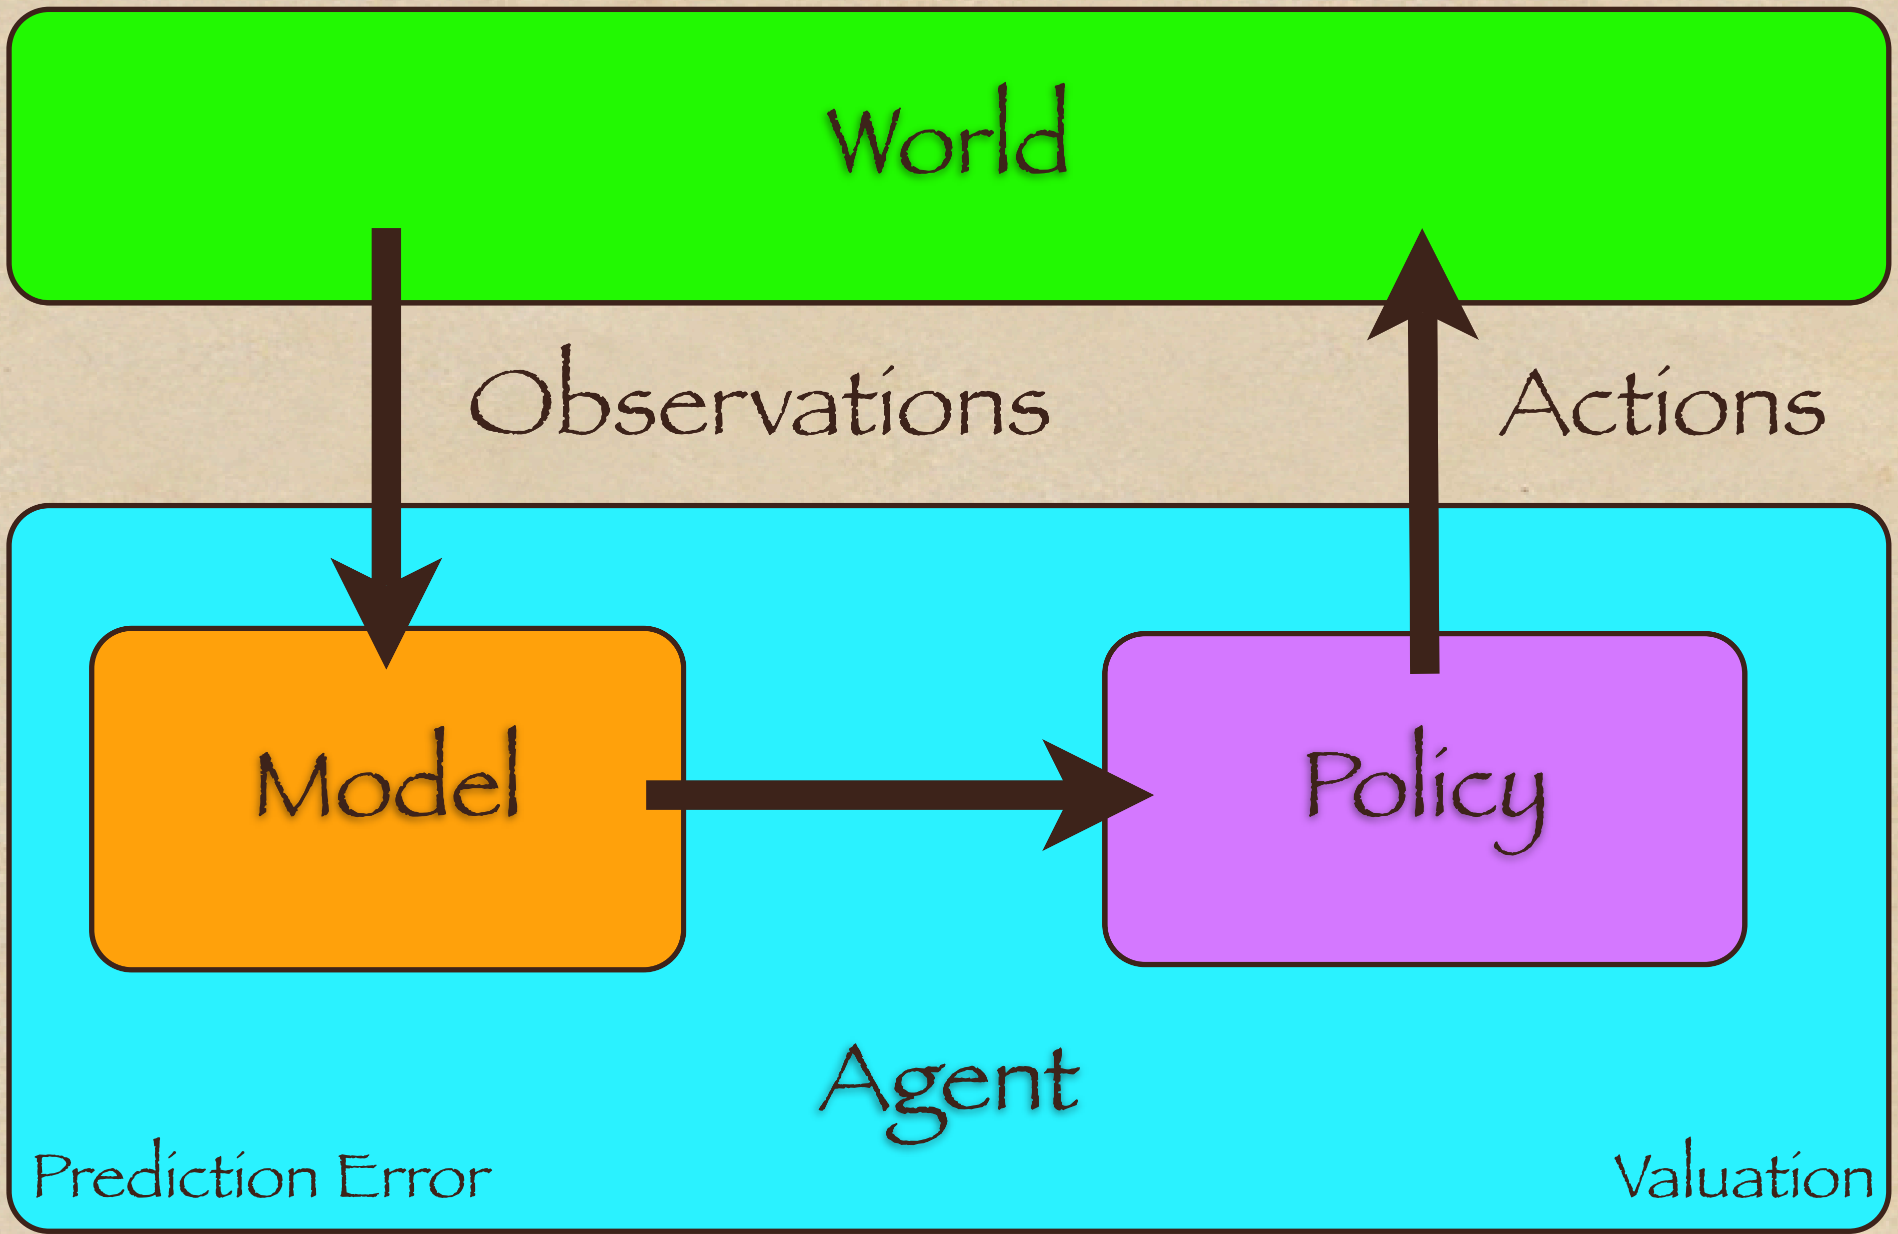


Deviation from knowledge    Commitment of resources to actions

# Main message

- Closing the loop:

  How interaction changes the world &

  how one adapts to those changes
- Theoretical foundations (& algorithms) for closing the feedback loop are now available.

# Conclusion

- Basic principles follow from
  - Information theory (rate distortion)
  - Statistical physics
- Balance exploitation & exploration
- Balance structure & error
- Balance exploitation & control
- Challenge: Fold in risk

# Prospects

- Collective Cognition:
  - Pattern discovery
  - Interactive learning
  - Adaptation dynamics
  - Emergent policy design
  - Multiagent dynamical systems

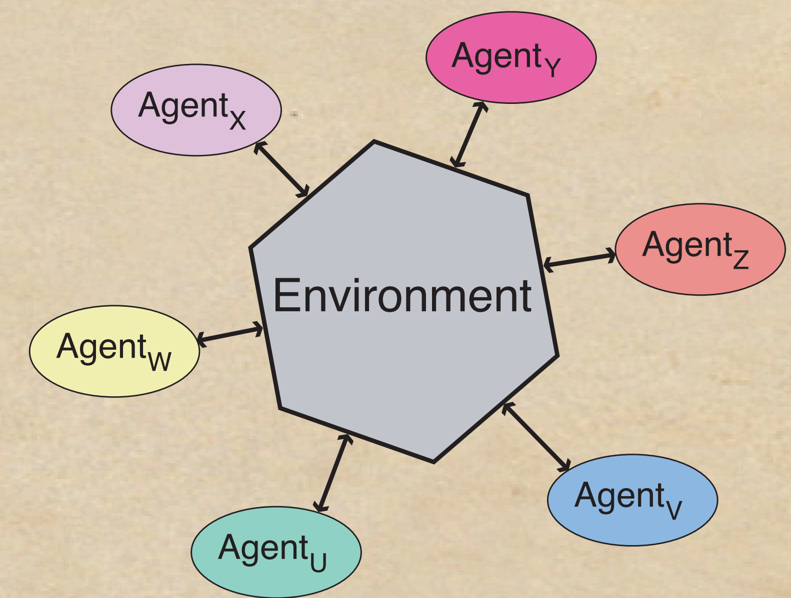

Thanks!