# Learning Nonlinear Dynamic Models

**Ruslan Salakhutdinov**

Joint work with **John Langford and Tong Zhang**

# Dynamic Model

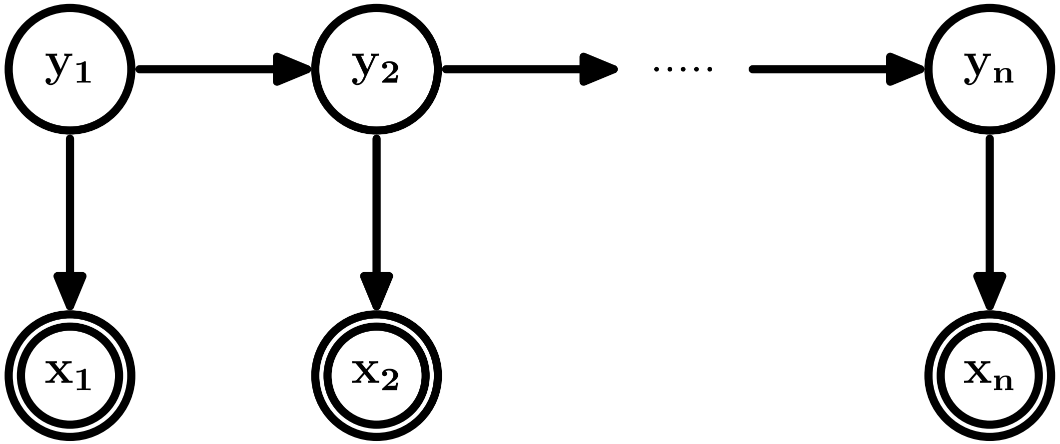


Posterior update:

$$P(\mathbf{y}_{t+1}|X_{1:t}) \propto \sum_{\mathbf{y}_t} P(\mathbf{y}_t|X_{1:t-1})P(\mathbf{x}_t|\mathbf{y}_t)P(\mathbf{y}_{t+1}|\mathbf{y}_t).$$

Prediction of future events:

$$P(\mathbf{x}_{t+1}|X_{1:t}) = \sum_{\mathbf{y}_{t+1}} P(\mathbf{x}_{t+1}|\mathbf{y}_{t+1})P(\mathbf{y}_{t+1}|X_{1:t}).$$

# Nonlinear Dynamic Model

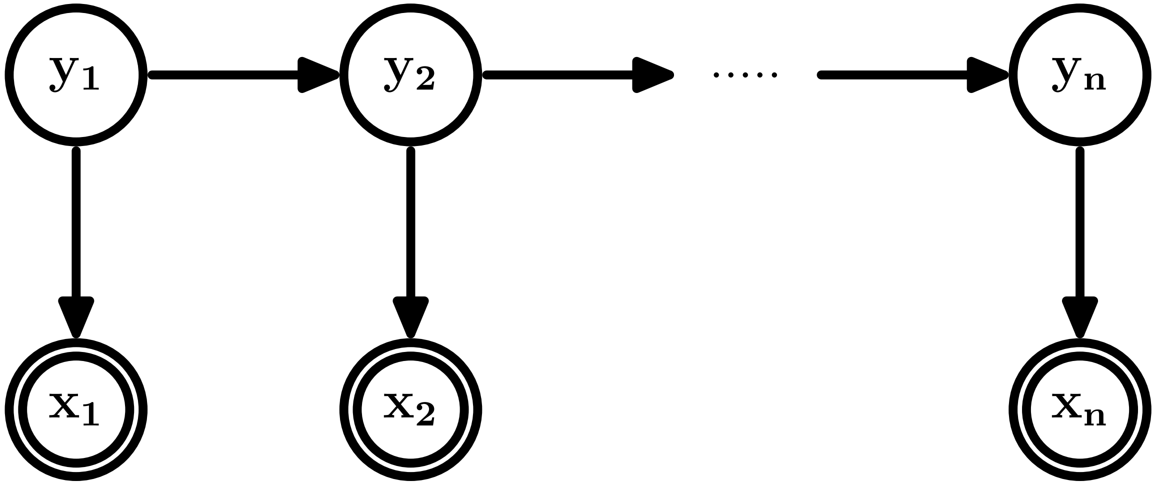


Computing the posterior $P(\mathbf{y}_{t+1}|X_{1:t})$ is difficult.

- Linearize nonlinear function: Extended Kalman Filter.
- Use approximations, e.g. particle filtering.

Inputs are high-dimensional and highly-structured.

# Sufficient Posterior Representation

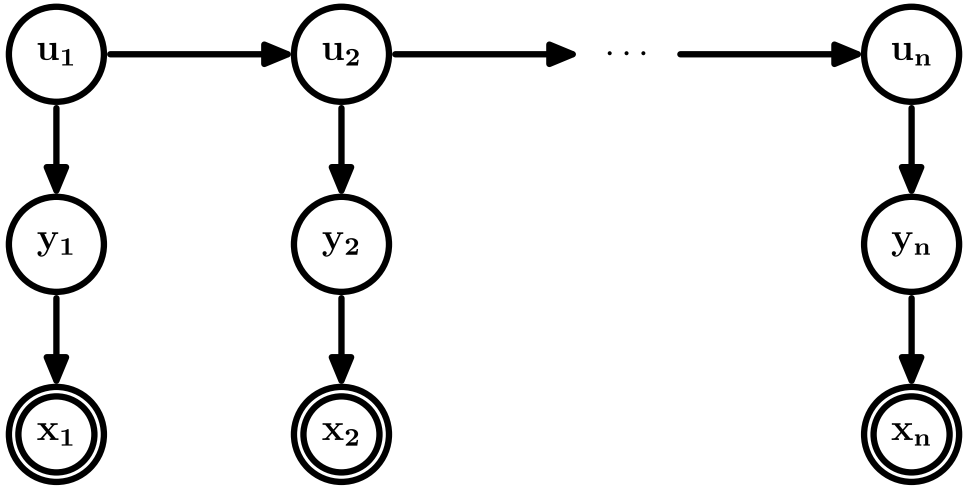


- Posterior $P(\mathbf{y}_{t+1}|X_{1:t})$ is approximated by a family of distributions parameterized by $\mathbf{u}_{t+1} \in \mathcal{U}$:

$$P(\mathbf{y}_{t+1}|X_{1:t}) \approx P(\mathbf{y}_{t+1}|\mathbf{u}_{t+1}).$$

- $\mathbf{u}_{t+1}$ is a sufficient statistic for the posterior $P(\mathbf{y}_{t+1}|X_{1:t})$.
- $\mathbf{u}_{t+1}$ is a deterministic parameter.

# Sufficient Posterior Representation

Sufficient Posterior Representation (SPR):

$$P(\mathbf{x}_{t+1}|X_{1:t}) \approx P(\mathbf{x}_{t+1}|\mathbf{u}_{t+1}).$$

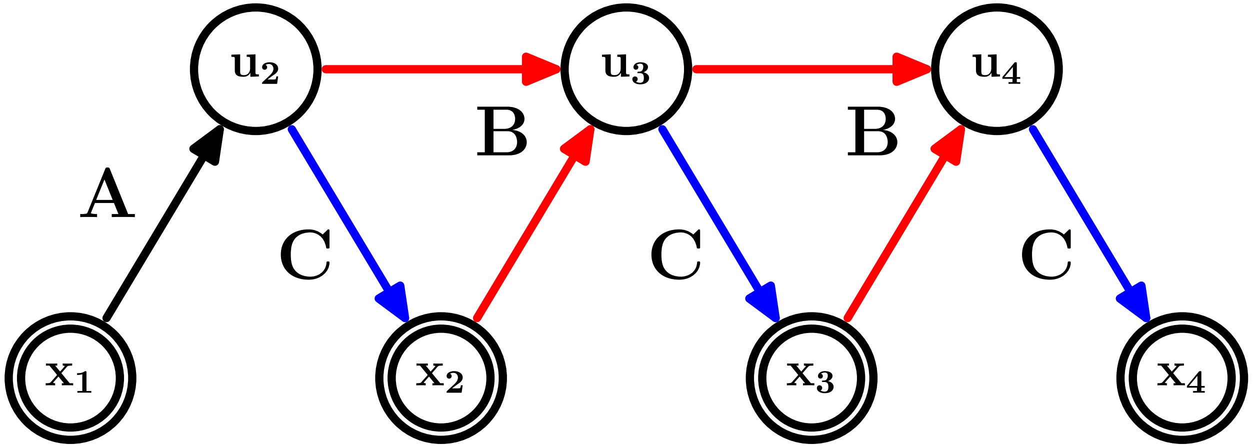


- Posterior update: $\mathbf{u}_{t+1} = B(\mathbf{x}_t, \mathbf{u}_t)$ .

  Give an arbitrary value to the initial state $\mathbf{u}_1$:

$$\mathbf{u}_2 = A(\mathbf{x}_1) = B(\mathbf{x}_1, \mathbf{u}_1) \ .$$

# Sufficient Posterior Representation

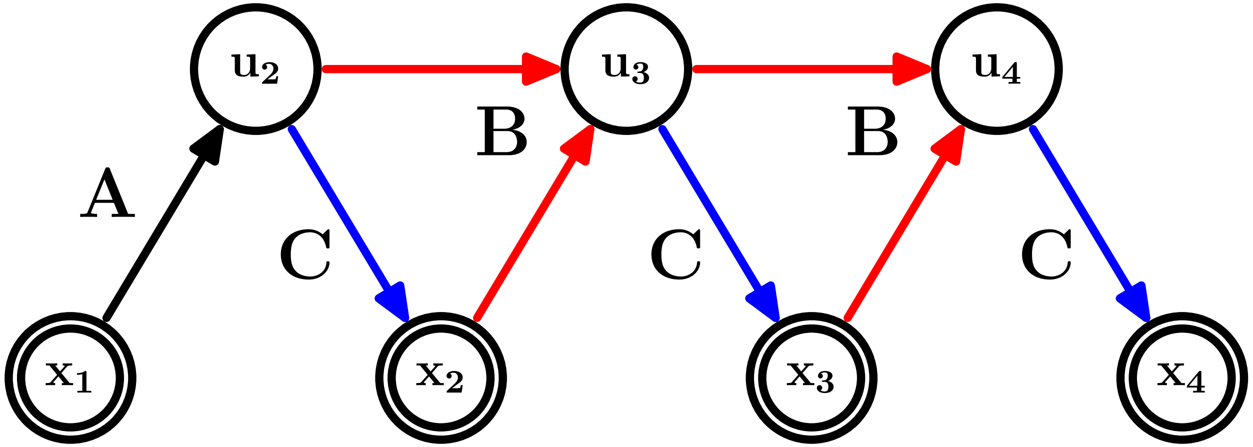


- Prediction:

$$p(\mathbf{x}_{t+1}|X_{1:t}) = C(\mathbf{u}_{t+1}).$$

**Key Observation:** $A$,$B$, and $C$ are deterministic.

# SPR Dynamic Model

A sufficient posterior representation of a dynamic model (SPR-DM) is given by:

- Observed sequence: $\{\mathbf{x}_t\}$
- Unobserved hidden "state": $\{\mathbf{u}_t\}$
- State initialization map: $\mathbf{u}_2 = A(\mathbf{x}_1)$
- State update map: $\mathbf{u}_{t+1} = B(\mathbf{x}_t, \mathbf{u}_t)$
- Prediction map: $p(\mathbf{x}_{t+1}|X_{1:t}) = C(\mathbf{u}_{t+1})$.

# Learning SPR-DM

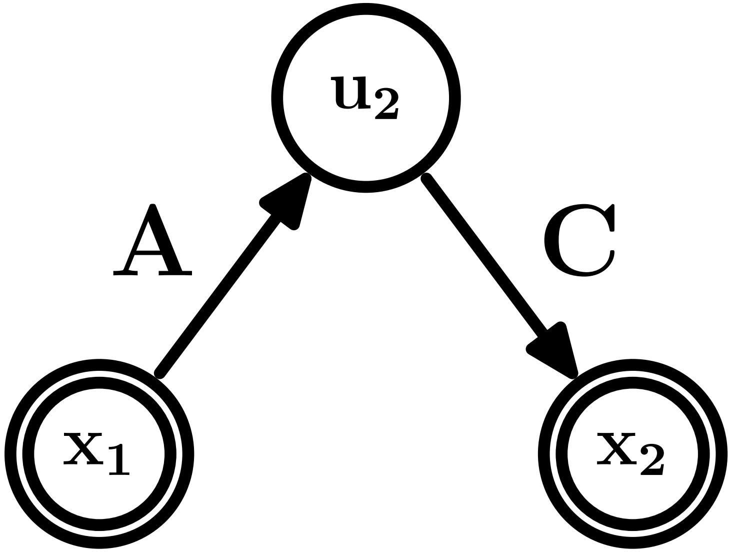


First prediction problem:
$p(\mathbf{x}_2|\mathbf{x}_1) = C(A(\mathbf{x}_1))$.

- State is “that information which summarizes the first observation in predicting the second observation”.
- Internal state can come from any learning algorithm.

# Learning SPR-DM State Evolution

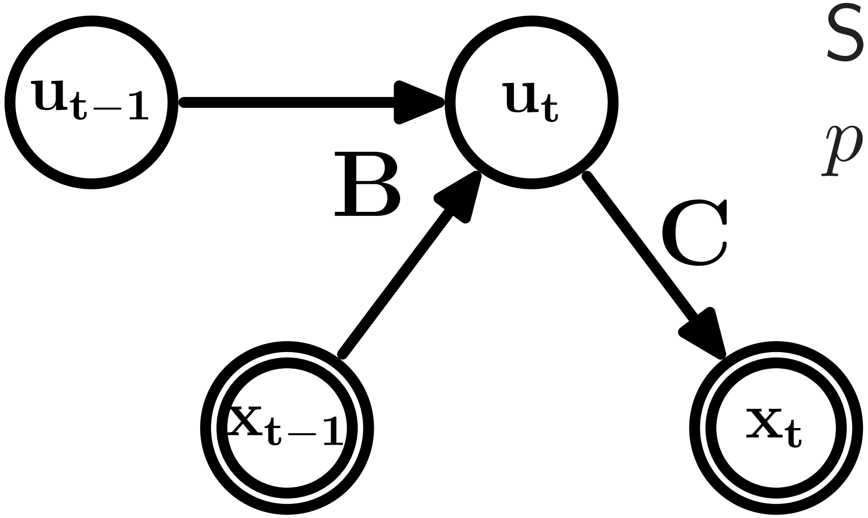


Second prediction problem:
$p(\mathbf{x}_t|X_{1:t-1}) = C(B(\mathbf{x}_{t-1}, \mathbf{u}_{t-1}))$.

- A state and an observation is used to predict the next state, reusing the state prediction from previous step.

# Learning SPR-DM

Pretraining: Local learning.

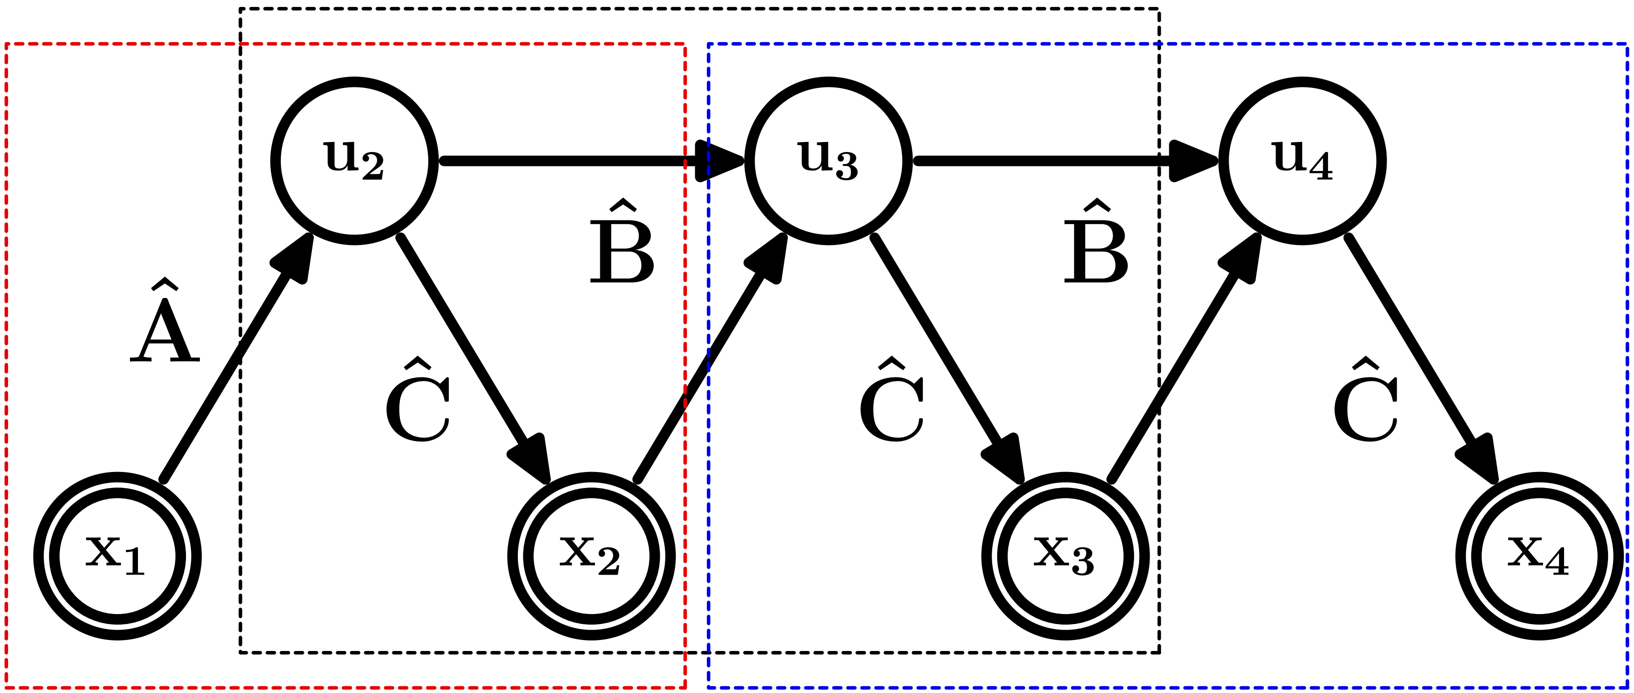


Fine-tuning: Backpropagation through time.

# Invertibility of SPR-DM

- To show consistency, we need the notion of invertibility.

  *The SPR-DM is invertible if there exist a function $R$ such that for all $t$, $R(C(\mathbf{u}_t)) = \mathbf{u}_t$.*

- Let $C = p(X_{t:t+k}|\mathbf{u}_t)$.

  Invertibility: if $\mathbf{u}_t$ and $\mathbf{u}'_t$ induce the same short range behavior $X_{t:t+k}$, then they are identical –

  They induce the same behavior for all $X_{t:\infty}$.

# SPR-DM used in Experiments

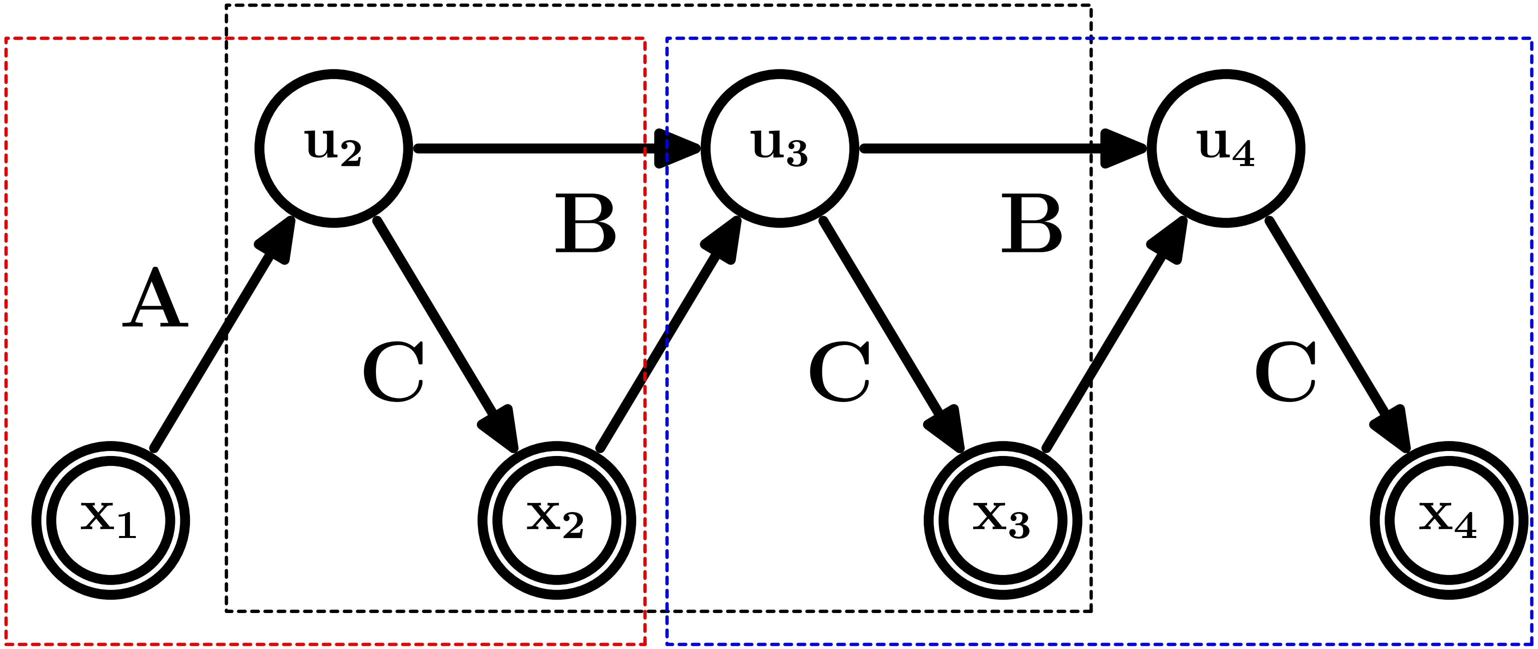


- $\mathbf{u}_2 = \;\; A(\mathbf{x}_1) = \sigma\left(A^\top \mathbf{x}_1 + \mathbf{b}\right)$
- $\mathbf{u}_t = \;\; B(\mathbf{x}_{t-1}, \mathbf{u}_{t-1}) = \sigma\left(B_1^\top \mathbf{x}_{t-1} + B_2^\top \mathbf{u}_{t-1} + \mathbf{b}\right)$
- $\hat{\mathbf{x}}_t = \;\; C(\mathbf{u}_t) = C^\top \mathbf{u}_t + \mathbf{a}$

  where $\sigma(y) = 1/(1+\exp(-y))$.

# Motion Capture Data

- Sequences of 3D joint angles plus body orientation and translation
- Various walking styles: normal, drunk, graceful, gangly, chicken, etc.
- 30 training and 8 test sequences, each of length 50.
- Each time step was represented by a vector of 58 real-valued numbers.

# Motion Capture Data

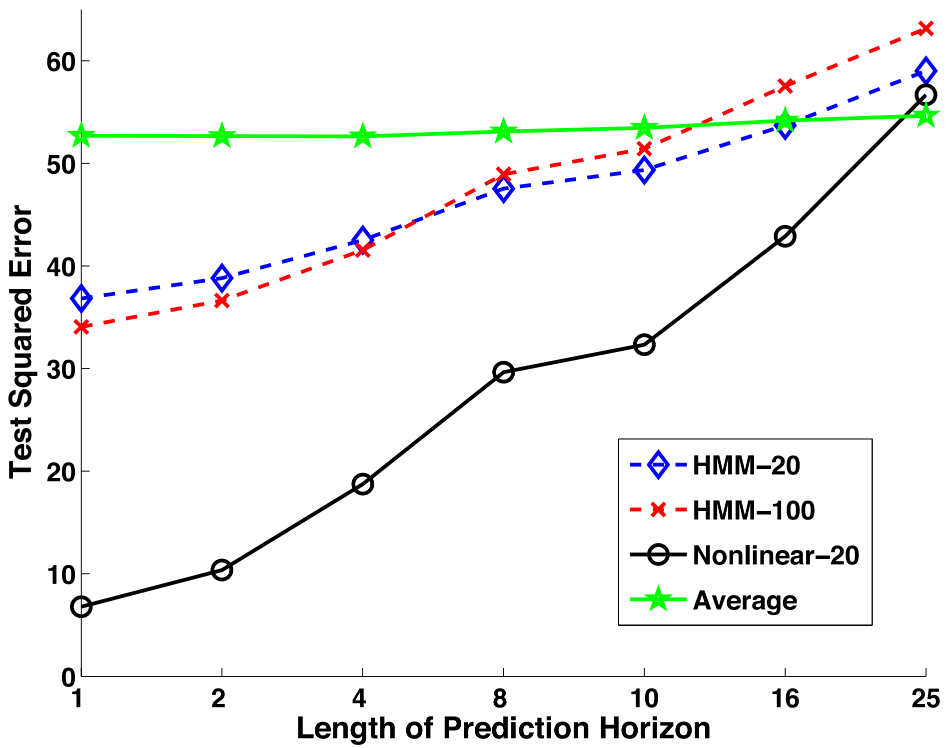


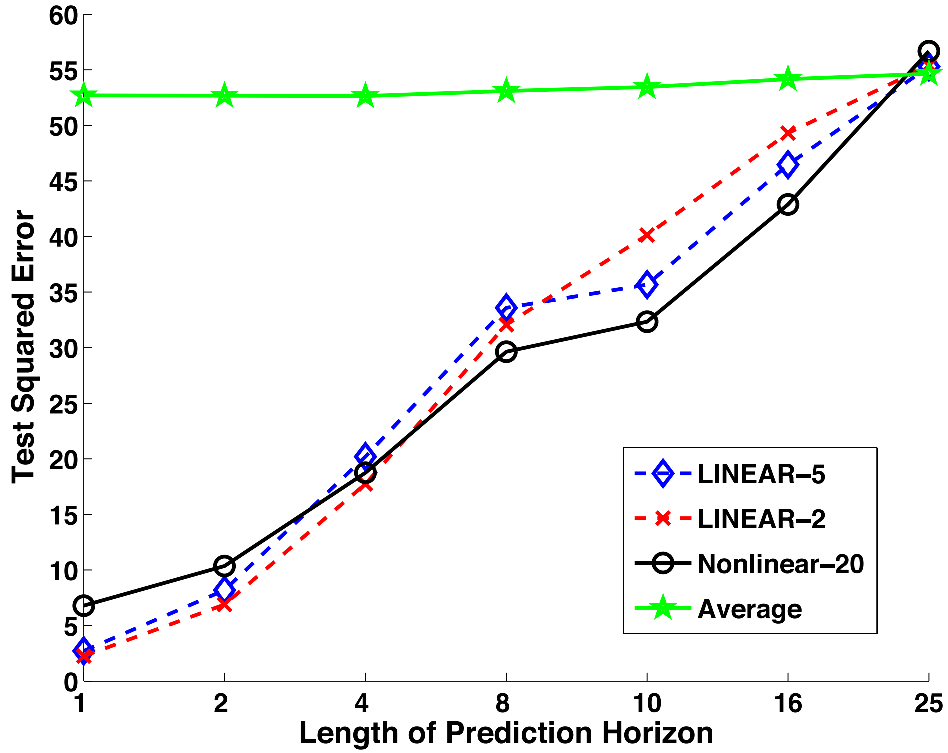


- Comparison: 20-dimensional nonlinear model, 20 and 100-state HMM's, and simple linear models (conditioned on 2 and 5 previous time steps).

# Weizmann Video Data

- Video sequences of nine human subjects.
- Various actions: waving one hand, waving two hands, jumping, and bending.
- 36 training and 10 test sequences, each of length 50.
- Each time step was represented by a vector of 464 real-valued numbers.

# Weizmann Video Data

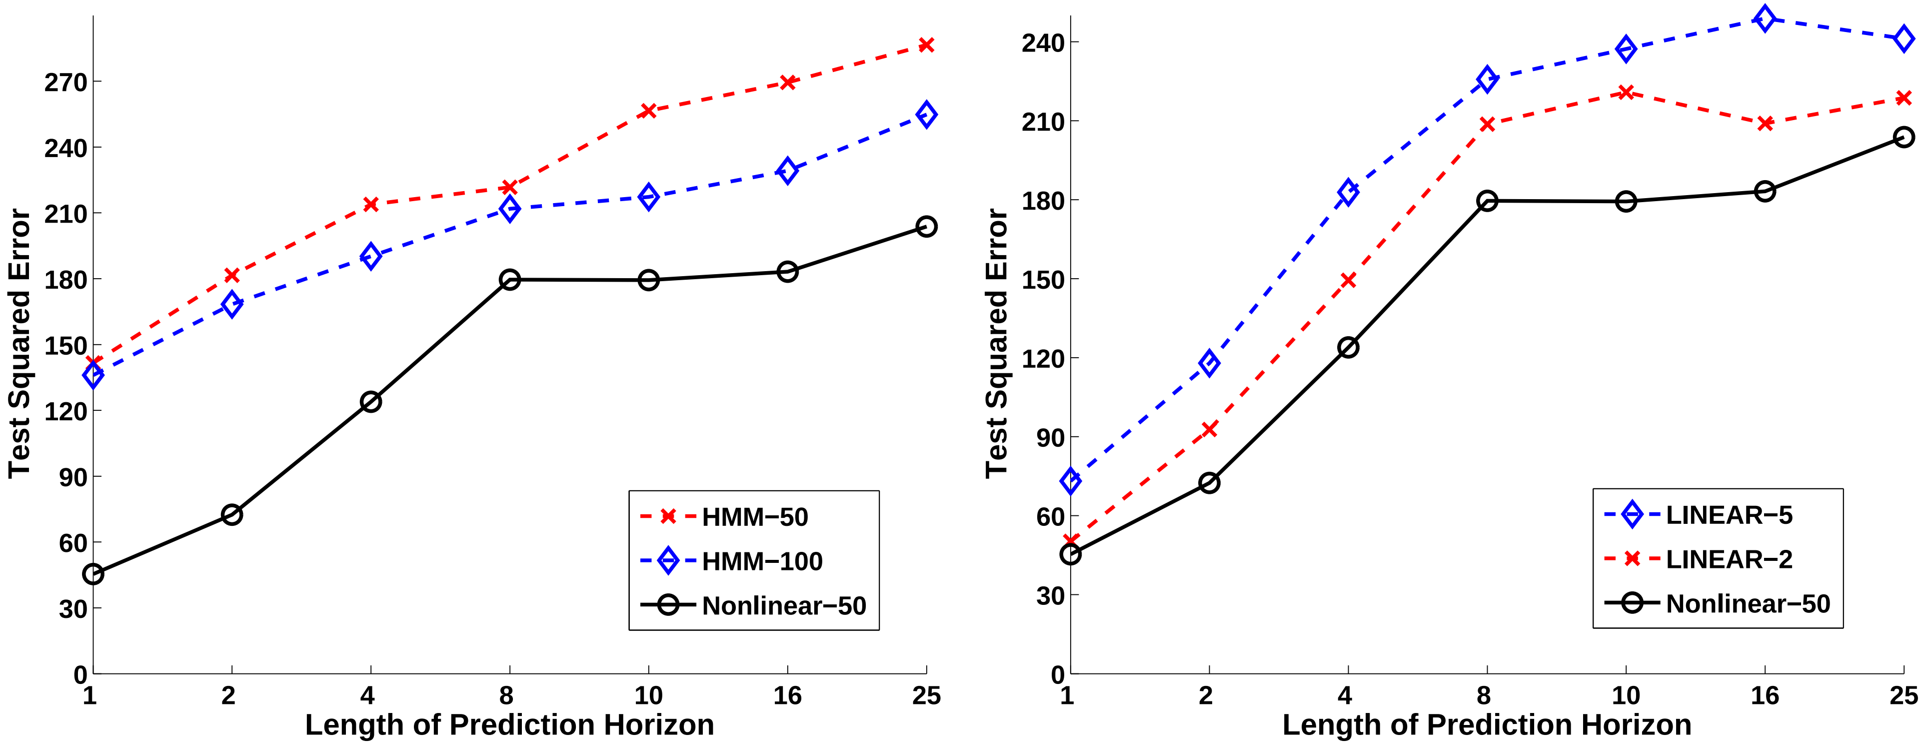


- Comparison: 50-dimensional nonlinear model, 50 and 100-state HMM's, and linear models.

Thank you.