Select Lab



A Constraint Generation Approach to Learning Stable Linear Dynamical Systems

Sajid M. Siddiqi Byron Boots Geoffrey J. Gordon

Carnegie Mellon University

NIPS 2007 poster W22



steam

Application: Dynamic Textures¹



- Videos of moving scenes that exhibit **stationarity** properties
- Dynamics can be captured by a low-dimensional model
- Learned models can efficiently simulate realistic sequences
- Applications: compression, recognition, synthesis of videos

¹S. Soatto, D. Doretto and Y. Wu. *Dynamic Textures*. *Proceedings of the ICCV*, 2001

• Input data sequence $Y = \begin{bmatrix} y_1 \ y_2 \dots y_T \end{bmatrix} \qquad y_t \in \mathcal{R}^m$

Linear Dynamical Systems



Input data sequence Y = [y₁ y₂...y_T] y_t ∈ R^m Assume a latent variable that explains the data X = [x₁ x₂...x_T] x_t ∈ Rⁿ



Input data sequence Y = [y₁ y₂...y_T] y_t ∈ R^m
Assume a latent variable that explains the data X = [x₁ x₂...x_T] x_t ∈ Rⁿ



- Input data sequence $Y = \begin{bmatrix} y_1 \ y_2 \dots y_T \end{bmatrix} \qquad y_t \in \mathcal{R}^m$
- Assume a **latent variable** that explains the data $X = \begin{bmatrix} x_1 \ x_2 \dots x_T \end{bmatrix} \qquad x_t \in \mathcal{R}^n$
- Assume a Markov model on the latent variable



- Input data sequence $Y = \begin{bmatrix} y_1 \ y_2 \dots y_T \end{bmatrix} \qquad y_t \in \mathcal{R}^m$
- Assume a **latent variable** that explains the data $X = \begin{bmatrix} x_1 \ x_2 \dots x_T \end{bmatrix} \qquad x_t \in \mathcal{R}^n$
- Assume a Markov model on the latent variable

Linear Dynamical Systems²



• State and observation models:

$$\begin{aligned} x_{t+1} &= Ax_t + w_t & w_t \sim N(0, Q) \\ y_t &= Cx_t + v_t & v_t \sim N(0, R) \end{aligned}$$

- Dynamics matrix: $A \in \mathcal{R}^{n \times n}$
- Observation matrix: $C \in \mathcal{R}^{m \times n}$

² Kalman, R. E. (1960). A new approach to linear filtering and prediction problems. Trans.ASME-JBE

Linear Dynamical Systems²

<u>This talk:</u>

- Learning LDS parameters from data while ensuring a **stable** dynamics matrix *A* more efficiently and accurately than previous methods

• State and observation models:

$$\begin{aligned} x_{t+1} &= A x_t + w_t & w_t \sim N(0, Q) \\ y_t &= C x_t + v_t & v_t \sim N(0, R) \\ \end{aligned}$$
Dynamics matrix: $A \in \mathcal{R}^{n \times n}$
Observation matrix: $C \in \mathcal{R}^{m \times n}$

² Kalman, R. E. (1960). A new approach to linear filtering and prediction problems. Trans.ASME-JBE

Learning Linear Dynamical Systems

• Suppose we have an estimated state sequence

$$\widehat{X}_{1:\tau} = [\widehat{x}_1 \ \widehat{x}_2 \dots \widehat{x}_\tau] \in \mathcal{R}^{n \times \tau}$$

Learning Linear Dynamical Systems

• Suppose we have an estimated state sequence

$$\widehat{X}_{1:\tau} = [\widehat{x}_1 \ \widehat{x}_2 \dots \widehat{x}_\tau] \in \mathcal{R}^{n \times \tau}$$

• Define *state reconstruction error* as the objective:

$$J(A) = ||AX_{1:\tau-1} - X_{2:\tau}||_F^2$$

Learning Linear Dynamical Systems

• Suppose we have an estimated state sequence

$$\widehat{X}_{1:\tau} = [\widehat{x}_1 \ \widehat{x}_2 \dots \widehat{x}_\tau] \in \mathcal{R}^{n \times \tau}$$

• Define *state reconstruction error* as the objective: $J(A) = ||AX_{1:\tau-1} - X_{2:\tau}||_F^2$

• We would like to learn \hat{A} such that $x_{t+1} \approx \hat{A} x_t$

i.e.
$$\hat{A} = min_A J(A)$$

• Subspace ID uses matrix decomposition to estimate the state sequence

- Subspace ID uses matrix decomposition to estimate the state sequence
- Build a *Hankel matrix* **D** of stacked observations

$$\mathcal{D} = \begin{bmatrix} y_1 & y_2 & y_3 & \cdots \\ y_2 & y_3 & y_4 & \cdots \\ y_3 & y_4 & y_5 & \cdots \\ \vdots & \vdots & \vdots & \ddots \end{bmatrix}$$

³ P. Van Overschee and B. De Moor Subspace Identification for Linear Systems. Kluwer, 1996

- Subspace ID uses matrix decomposition to estimate the state sequence
- Build a *Hankel matrix* **D** of stacked observations



³ P. Van Overschee and B. De Moor Subspace Identification for Linear Systems. Kluwer, 1996

• In expectation, the Hankel matrix is inherently **low-rank!**

• In expectation, the Hankel matrix is inherently **low-rank!**

$$E(\mathcal{D}) = \begin{bmatrix} Cx_1 & Cx_2 & Cx_3 & \cdots \\ Cx_2 & Cx_3 & Cx_4 & \cdots \\ Cx_3 & Cx_4 & Cx_5 & \cdots \\ \vdots & \vdots & \vdots & \ddots \end{bmatrix} = \begin{bmatrix} Cx_1 & Cx_2 & Cx_3 & \cdots \\ CAx_1 & CAx_2 & CAx_3 & \cdots \\ CA^2x_1 & CA^2x_2 & CA^2x_3 & \cdots \\ \vdots & \vdots & \vdots & \ddots \end{bmatrix}$$

• In expectation, the Hankel matrix is inherently low-rank!

$$E(\mathcal{D}) = \begin{bmatrix} Cx_1 & Cx_2 & Cx_3 & \cdots \\ Cx_2 & Cx_3 & Cx_4 & \cdots \\ Cx_3 & Cx_4 & Cx_5 & \cdots \\ \vdots & \vdots & \vdots & \ddots \end{bmatrix} = \begin{bmatrix} Cx_1 & Cx_2 & Cx_3 & \cdots \\ CAx_1 & CAx_2 & CAx_3 & \cdots \\ CA^2x_1 & CA^2x_2 & CA^2x_3 & \cdots \\ \vdots & \vdots & \vdots & \ddots \end{bmatrix} \\ = \begin{bmatrix} C \\ CA \\ CA^2 \\ CA^3 \\ \vdots \end{bmatrix} \begin{bmatrix} x_1 & x_2 & x_3 & x_4 & \cdots \end{bmatrix}$$

• In expectation, the Hankel matrix is inherently low-rank!

$$E(\mathcal{D}) = \begin{bmatrix} Cx_1 & Cx_2 & Cx_3 & \cdots \\ Cx_2 & Cx_3 & Cx_4 & \cdots \\ Cx_3 & Cx_4 & Cx_5 & \cdots \\ \vdots & \vdots & \vdots & \ddots \end{bmatrix} = \begin{bmatrix} Cx_1 & Cx_2 & Cx_3 & \cdots \\ CAx_1 & CAx_2 & CAx_3 & \cdots \\ CA^2x_1 & CA^2x_2 & CA^2x_3 & \cdots \\ \vdots & \vdots & \vdots & \ddots \end{bmatrix}$$

Can use SVD to obtain the low-dimensional state sequence
$$= \begin{bmatrix} C \\ CA \\ CA^2 \\ CA^3 \\ \vdots \end{bmatrix} \begin{bmatrix} x_1 & x_2 & x_3 & x_4 & \cdots \end{bmatrix}$$

• In expectation, the Hankel matrix is inherently low-rank!

 $E(\mathcal{D}) = \begin{bmatrix} Cx_1 & Cx_2 & Cx_3 & \cdots \\ Cx_2 & Cx_3 & Cx_4 & \cdots \\ Cx_3 & Cx_4 & Cx_5 & \cdots \\ \vdots & \vdots & \vdots & \ddots \end{bmatrix} = \begin{bmatrix} Cx_1 & Cx_2 & Cx_3 & \cdots \\ CAx_1 & CAx_2 & CAx_3 & \cdots \\ CA^2x_1 & CA^2x_2 & CA^2x_3 & \cdots \\ \vdots & \vdots & \vdots & \ddots \end{bmatrix}$ $= \begin{vmatrix} C \\ CA \\ CA^2 \\ CA^3 \\ \vdots \end{vmatrix} \begin{bmatrix} x_1 & x_2 & x_3 & x_4 & \cdots \end{bmatrix}$ • Can use SVD to obtain the low-dimensional state sequence For *D* with *k* observations per column, **I**]

³ P. Van Overschee and B. De Moor Subspace Identification for Linear Systems. Kluwer, 1996

• In expectation, the Hankel matrix is inherently low-rank!



³ P. Van Overschee and B. De Moor Subspace Identification for Linear Systems. Kluwer, 1996

• Lets train an LDS for **steam** textures using this algorithm, and simulate a video from it!

$Y = [y_1 \ y_2 \dots y_T]$



 $x_t \, 2 \, R^{40}$

Simulating from a learned model





The model is unstable

Notation

- $\lambda_1, \ldots, \lambda_n$: eigenvalues of A ($|\lambda_1| > \ldots > |\lambda_n|$)
- v_1, \ldots, v_n : unit-length eigenvectors of A
- $\sigma_1, \dots, \sigma_n$: singular values of A ($\sigma_1 > \sigma_2 > \dots \sigma_n$)
- S_λ: matrices with $|λ_1| \cdot 1$ S_σ: matrices with $σ_1 \cdot 1$

• a matrix *A* is **stable if** $|\lambda_1| \cdot 1$, i.e. if $A \in S_{\lambda}$

• a matrix *A* is **stable if** $|\lambda_1| \cdot 1$, i.e. if $A \in S_{\lambda}$

$$A_1 = \left[\begin{array}{rrr} 0.3 & 10\\ 0 & 0.3 \end{array} \right]$$

$$A_2 = \left[\begin{array}{cc} 0.3 & 9.9\\ 0.1 & 0.3 \end{array} \right]$$

• a matrix *A* is **stable if** $|\lambda_1| \cdot 1$, i.e. if $A \in S_{\lambda}$

$$A_1 = \begin{bmatrix} 0.3 & 10 \\ 0 & 0.3 \end{bmatrix}$$
$$|\lambda_1| = 0.3$$

$$A_2 = \begin{bmatrix} 0.3 & 9.9 \\ 0.1 & 0.3 \end{bmatrix}$$
$$|\lambda_1| = 1.295$$

• a matrix A is stable if $|\lambda_1| \cdot 1$, i.e. if $A \in S_{\lambda}$



• a matrix *A* is **stable if** $|\lambda_1| \cdot 1$, i.e. if $A \in S_{\lambda}$



• But S_{λ} is non-convex!





α





⁴S. L. Lacy and D. S. Bernstein. Subspace identification with guaranteed stability using constrained optimization. In *Proc. of the ACC* (2002), *IEEE Trans. Automatic Control* (2003)

• Lets train an LDS for steam again, this time constraining *A* to be in S_{σ}

Simulating from a Lacy-Bernstein stable texture model





Model is over-constrained. Can we do better?

Our Approach

- Formulate the S_σ approximation of the problem as a semi-definite program (SDP)
- Start with a **quadratic program (QP) relaxation** of this SDP, and incrementally **add constraints**
- Because the SDP is an inner approximation of the problem, we **reach stability early**, before reaching the feasible set of the SDP
- We **interpolate** the solution to return the best stable matrix possible



objective function contours



objective function contours



• *A*₁: unconstrained QP solution (least squares estimate)

objecțive function contours



• *A*₁: unconstrained QP solution (least squares estimate)

objecțive function contours



- *A*₁: unconstrained QP solution (least squares estimate)
- *A*₂: QP solution after 1 constraint (happens to be stable)

objecțive function contours



- *A*₁: unconstrained QP solution (least squares estimate)
- *A*₂: QP solution after 1 constraint (happens to be stable)
- A_{final} : Interpolation of stable solution with the last one

objecțive function contours



- *A*₁: unconstrained QP solution (least squares estimate)
- *A*₂: QP solution after 1 constraint (happens to be stable)
- *A_{final}*: Interpolation of stable solution with the last one *A_{previous method}*: Lacy Bernstein (2002)

• Lets train an LDS for **steam** using constraint generation, and simulate ...

Simulating from a Constraint Generation stable texture model





Model captures more dynamics and is still stable • Least Squares

Constraint Generation





Empirical Evaluation

- Algorithms:
 - Constraint Generation CG (our method)
 - Lacy and Bernstein (2002) -LB-1
 - finds a $\sigma_1 \cdot 1$ solution
 - Lacy and Bernstein (2003)–LB-2
 - solves a similar problem in a transformed space
- Data sets
 - Dynamic textures
 - Biosurveillance baseline models (see paper)

Reconstruction error

• *Steam* video



number of latent dimensions

Running time

• *Steam* video



number of latent dimensions

Conclusion

- A novel **constraint generation algorithm** for learning **stable linear dynamical systems**
- SDP relaxation enables us to optimize over a **larger set** of matrices while being more efficient

Conclusion

- A novel **constraint generation algorithm** for learning **stable linear dynamical systems**
- SDP relaxation enables us to optimize over a **larger set of matrices** while being **more efficient**
- Future work:
 - Adding stability constraints to EM
 - Stable models for more structured dynamic textures



• Thank you!

Subspace ID with Hankel matrices

Stacking multiple observations in *D* forces latent states to model the future _____



e.g. annual sunspots data with 12-year cycles



• The state space of a stable LDS lies inside some ellipse



- The state space of a stable LDS lies inside some ellipse
- The set of matrices that map a particular ellipse into itself (and hence are stable) is convex



- The state space of a stable LDS lies inside some ellipse
- The set of matrices that map a particular ellipse into itself (and hence are stable) is convex
- If we knew in advance which ellipse contains our state space, finding A would be a convex problem. But we don't



- The state space of a stable LDS lies inside some ellipse
- The set of matrices that map a particular ellipse into itself (and hence are stable) is convex
- If we knew in advance which ellipse contains our state space, finding A would be a convex problem.
 But we don't
- ...and the set of all stable matrices is non-convex







