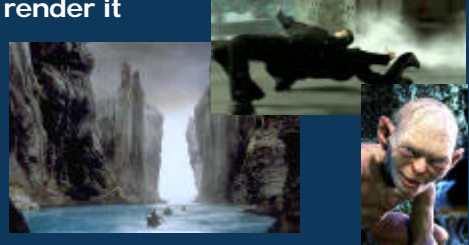


## Bayesian Learning for Computer Graphics


Aaron Hertzmann  
University of Toronto

## Computers are really fast

- If you can create it, you can render it




## How do you create it?




## Two problems

- How do we get the data into the computer?
- How do we manipulate it?



## Data fitting

$$\begin{aligned}
 & \sum_{j,t} \sum_j m_j(q_t) \approx s_j \cdot j_j^2 \\
 & + W_1 \sum_{j,t} \sum_j j_j q_t \approx q_t \approx 1 \cdot j_j^2 \\
 & + W_2 \sum_{j,t} \sum_j j_j \dots j_j^2
 \end{aligned}$$



## Key questions

- How do you fit a model to data?
  - How do you choose weights and thresholds?
  - How do you incorporate prior knowledge?
  - How do you merge multiple sources of information?
  - How do you model uncertainty?

**Bayesian reasoning provides a solution**

## Talk outline

- Bayesian reasoning
- Facial modeling example
- Non-rigid modeling from video

## What is reasoning?

- How do people reason?
- How should computers do it?

## Aristotelian Logic

- If **A** is true, then **B** is true
- **A** is true
- Therefore, **B** is true

**A:** My car was stolen

**B:** My car isn't where I left it

## Real-world is uncertain

Problems with pure logic

- Don't have perfect information
- Don't really know the model
- Model is non-deterministic

## Beliefs

- Let  $B(X)$  = "belief in  $X$ ",
- $B(\neg X)$  = "belief in not  $X$ "

1. An ordering of beliefs exists
2.  $B(X) = f(B(\neg X))$
3.  $B(X) = g(B(X|Y), B(Y))$


## Cox axioms

R.T. Cox, "Probability, frequency, and reasonable expectation," *American J. Physics*, 14(1):1-13, 1946

$$p(\text{TRUE}) = 1$$

$$p(A) = \int_{B=b} p(A; B)$$

$$p(A; B) = p(A|B)p(B)$$



**"Probability theory is nothing more than common sense reduced to calculation."**

- Pierre-Simon Laplace, 1814



## Bayesian vs. Frequentist


- Frequentist ("Orthodox"):
  - Probability = percentage of events in infinite trials
- Medicine, biology: Frequentist
- Astronomy, geology, EE, computer vision: largely Bayesian

## Learning in a nutshell

- Create a mathematical model
- Get data
- Solve for unknowns

## Face modeling

- Blanz, Vetter, "A Morphable Model for the Synthesis of 3D Faces," SIGGRAPH 99



## Generative model

- Faces come from a Gaussian

$$p(\mathbf{S}; \mathbf{p}) = \frac{1}{(2\pi)^{d/2} |\mathbf{p}|} e^{-\frac{1}{2} (\mathbf{S} - \boldsymbol{\mu})^T \mathbf{p}^{-1} (\mathbf{S} - \boldsymbol{\mu})}$$

$$p(\mathbf{f}; \mathbf{g}; \mathbf{p}) = \prod_i p(S_i; \mathbf{p})$$

Learning

$$\arg \max_{\mathbf{g}, \mathbf{p}} p(\mathbf{g}; \mathbf{p} | \mathbf{f}; \mathbf{g})$$



## Bayes Rule

$$P(A; B) = p(A|B)p(B) \\ = p(B|A)p(A)$$

$$P(B|A) = \frac{P(A; B)}{P(A)}$$

$$p(\text{model} | \text{data}) \propto p(\text{data} | \text{model}) p(\text{model})$$

Often:  $p(\text{model} | \text{data}) \propto p(\text{data} | \text{model})$

## Learning a Gaussian

$$\begin{aligned} & \arg \max_{\theta, \mu} p(\theta; \text{data}) \\ &= \arg \max_{\theta, \mu} p(\text{data} | \theta, \mu) \\ &= \arg \max_{\theta, \mu} \prod_i p(S_i | \theta, \mu) \\ &= \arg \max_{\theta, \mu} \prod_i \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{1}{2\sigma^2}(S_i - \mu)^2} \end{aligned}$$

$$\begin{aligned} \theta &= \mu \\ \mu &= \frac{1}{N} \sum_i S_i \end{aligned}$$

## Maximization trick

- Maximize  $p(x)$
- $\leftrightarrow$  minimize  $-\ln p(x)$

## Fitting a face to an image

Generative model

$$p(S | \theta; \mu) = \prod_i \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{1}{2\sigma^2}(S_i - \mu)^2}$$

$$I = \text{Render}(S; \mu) + n$$

$$-\ln p(I | S; \mu; \sigma^2) = \sum_j \frac{1}{2\sigma^2} (I_j - \text{Render}(S; \mu)_j)^2$$



## Fitting a face to an image

Maximize

$$p(S; \mu | I; \theta; \mu; \sigma^2)$$


minimize

$$-\ln p(S; \mu | I; \theta; \mu; \sigma^2)$$

$$\begin{aligned} & \sum_j \frac{1}{2\sigma^2} (I_j - \text{Render}(S; \mu)_j)^2 + \sum_i \frac{1}{2\sigma^2} (S_i - \mu)^2 \\ & + \frac{N}{2} \ln 2\pi\sigma^2 + \frac{1}{2} \ln(2\pi\sigma^2)^N \end{aligned}$$

## Why does it work?

$$\begin{aligned} p(\text{data} | \theta, \mu) &= \prod_i p(x_i | \theta, \mu) \\ \text{s.t. } & \prod_i p(x_i | \theta, \mu) = 1 \end{aligned}$$



## General features

- Models uncertainty
- Applies to any generative model
- Merge multiple sources of information
- Learn *all* the parameters

## Caveats

- Still need to understand the model
- Not necessarily tractable
- Potentially more involved than *ad hoc* methods

## Applications in graphics


- Shape and motion capture
- Learning styles and generating new data

## Learning Non-Rigid 3D Shape from 2D Motion

Joint work with Lorenzo Torresani and Chris Bregler (Stanford, NYU)

## Camera geometry

Orthographic projection



## Non-rigid reconstruction

Input: 2D point tracks  
Output: 3D nonrigid motion

$$p_t = R_t S_t + T_t$$

**Totally ambiguous!**

## Shape reconstruction

Least-squares version

$$p_t = R_t S_t + T_t$$

minimize

$$\sum_t \| p_t - R_t S_t - T_t \|^2 + w \sum_t \| S_t - \hat{S} \|^2$$

## Bayesian formulation

$$p(\mathbf{S}; \mathbf{g}; \mathbf{p}) = \frac{1}{(2\hat{u})^{d/2}} e^{-\frac{1}{2}(\mathbf{S}_t - \hat{\mathbf{g}})^T \mathbf{p}^{-1} (\mathbf{S}_t - \hat{\mathbf{g}}) - 2}$$

$$\mathbf{p}_{:,t} = \mathbf{R}_t \mathbf{S}_{:,t} + \mathbf{T}_t + \mathbf{n} \quad \mathbf{n} \in \mathcal{N}(0; \hat{u}^2)$$

$$\text{maximize } \frac{p(\mathbf{R}; \mathbf{S}; \mathbf{T}; \mathbf{g}; \mathbf{p}; \mathbf{P})}{p(\mathbf{P}; \mathbf{R}; \mathbf{S}; \mathbf{T}) p(\mathbf{S}; \mathbf{g}; \mathbf{p})}$$

## How does it work?


$$\text{Maximize } \frac{p(\mathbf{R}; \mathbf{S}; \mathbf{T}; \mathbf{g}; \mathbf{p}; \mathbf{P})}{p(\mathbf{P}; \mathbf{R}; \mathbf{S}; \mathbf{T}) p(\mathbf{S}; \mathbf{g}; \mathbf{p})}$$

$$\text{Minimize } \sum_t \ln p(\mathbf{P}; \mathbf{R}_t; \mathbf{S}_t; \mathbf{T}_t) - \sum_t \ln p(\mathbf{S}_t; \mathbf{g}; \mathbf{p})$$

$$= \sum_t \ln p_t \left( \sum_j \mathbf{R}_t \mathbf{S}_t - \mathbf{T}_t \right)^2 - 2\hat{u}^2 + \frac{N}{2} \ln \hat{u}^2 + \sum_t \frac{1}{2} (\mathbf{S}_t - \hat{\mathbf{g}})^T \mathbf{p}^{-1} (\mathbf{S}_t - \hat{\mathbf{g}}) + \frac{T}{2} \ln |\mathbf{p}|$$

(actual system is slightly more sophisticated)

## Results



Input  
Least-squares, no reg.  
View 2  
Least-squares, reg.  
Gaussian  
LDS

## Conclusions

- Bayesian methods provide unified framework
- Build a model, and reason about it
- The future of data-driven graphics