


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?





Digital Michaelangelo Project



Steven Schkolne


Two problems

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



Data fitting

$$\begin{aligned} & \sum_{j;t} m_j(q_t) \approx s_{j;t} j^2 \\ & + w_1 \sum_{j;t} j q_t \approx q_{t-1} j^2 \\ & + w_2 \sum_{j;t} j :: j^2 \end{aligned}$$



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

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

$$p(A) = \sum_{B=b_i} p(A; B)$$

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



Photo by Jack Engeman

“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

$$p(\mathbf{S}_i | \mathbf{g}; \mathbf{p}) = \frac{1}{(2\pi)^{d_j} |\mathbf{p}|} e^{-\frac{1}{2} (\mathbf{s}_i - \hat{\mathbf{s}})^T \mathbf{p}^{-1} (\mathbf{s}_i - \hat{\mathbf{s}})}$$

$$p(\mathbf{f} | \mathbf{S}_i, \mathbf{g}; \mathbf{p}) = \prod_i p(\mathbf{S}_i | \mathbf{g}; \mathbf{p})$$

Learning

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



\mathbf{S}_i

Bayes Rule

$$\begin{aligned}P(A; B) &= p(A|B)p(B) \\ &= p(B|A)p(A)\end{aligned}$$

$$P(B|A) = p(A|B)p(B)/p(A)$$

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

Oftentimes: $p(\text{model}|\text{data}) / p(\text{data}|\text{model})$

Learning a Gaussian

$$\begin{aligned}
 & \arg \max_{\theta, \mu} p(\theta; \mu | \{S_i\}) \\
 &= \arg \max_{\theta, \mu} p(\{S_i\} | \theta; \mu) \\
 &= \arg \max_{\theta, \mu} \prod_i p(S_i | \theta; \mu) \\
 &= \arg \max_{\theta, \mu} \prod_i \frac{1}{(2\pi)^{d/2} |\Sigma|} e^{-\frac{1}{2} (S_i - \theta)^T \Sigma^{-1} (S_i - \theta)}
 \end{aligned}$$

$$\theta = \frac{1}{N} \sum_i S_i$$

$$\Sigma = \frac{1}{N} \sum_i (S_i - \theta)(S_i - \theta)^T$$

Maximization trick

- Maximize $p(x)$
<-> minimize $\ln p(x)$

Fitting a face to an image

Generative model

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

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

$$p(I; S; \mu; \hat{\Sigma}) = \frac{1}{(2\pi)^{d/2} |\Sigma|} e^{-\frac{1}{2} \|I - \text{Render}(S; \mu)\|^2_{\Sigma^{-1}}}$$



Fitting a face to an image

Maximize

$$p(S; \mathcal{I}; \mathcal{S}; \mathbf{p}; \hat{u}^2)$$

minimize

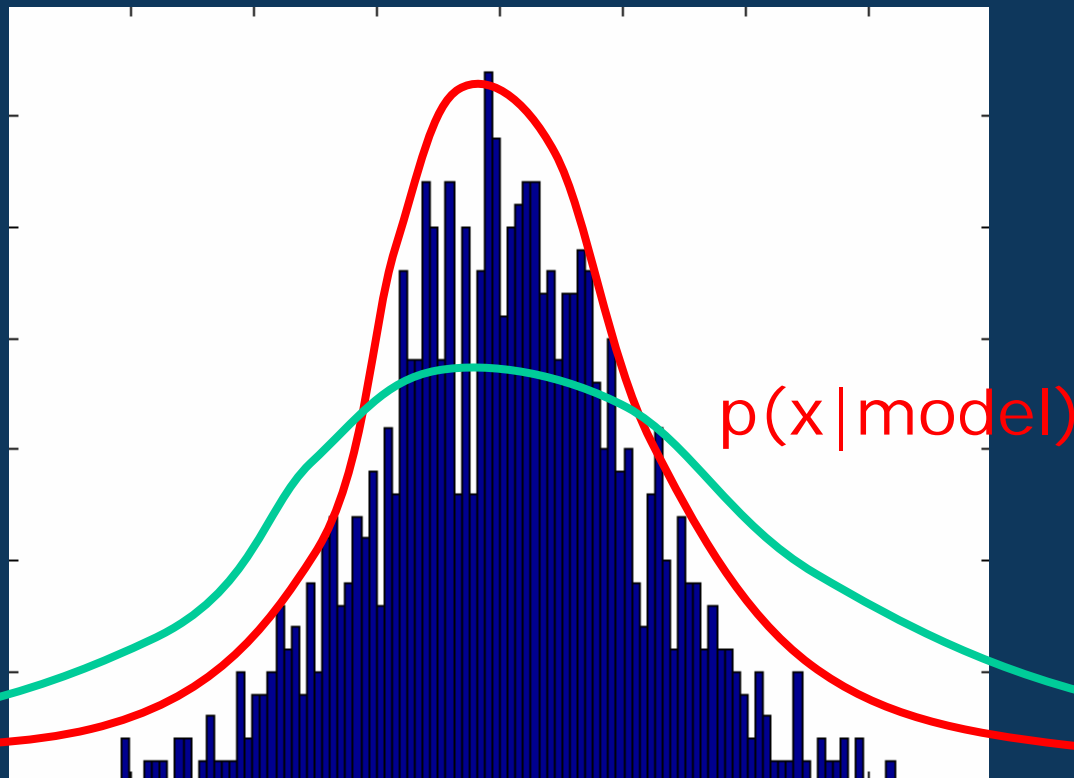
$$\rightarrow \ln p(S; \mathcal{I}; \mathcal{S}; \mathbf{p}; \hat{u}^2) =$$

$$\|\mathcal{I} - \text{Render}(S; \mathcal{I})\|^2 = 2\hat{u}^2 + (S - \mathcal{S})^T \mathbf{p}^{-1} (S - \mathcal{S}) = 2 \\ + \frac{N}{2} \ln 2\hat{u}^2 + \frac{1}{2} \ln (2\hat{u})^d \|\mathbf{p}\|$$

Why does it work?

$$p(f_{x_i | g_j \text{ model}}) = \int p(x_j | \text{model})$$

$$s:t: \int p(x_j | \text{model}) = 1$$



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

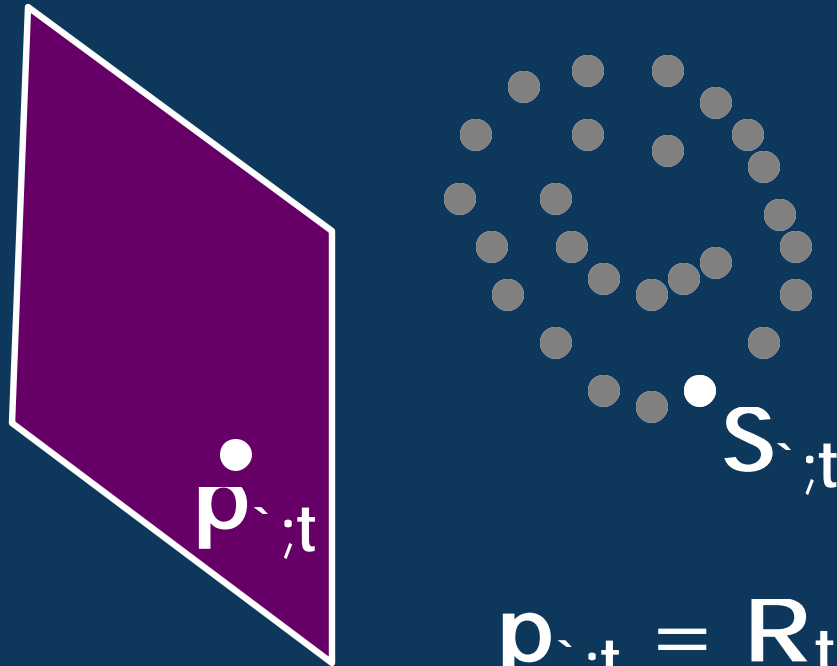
- Shape and motion capture
- Learning styles and generating new data

Learning Non-Rigid 3D Shape from 2D Motion

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

Camera geometry

Orthographic projection



$$p_{;t} = R_t S_{;t} + T_t$$

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

$$\mathbf{p}_t = \mathbf{R}_t \mathbf{S}_t + \mathbf{T}_t$$

minimize

$$\sum_t \|\mathbf{p}_t - \mathbf{R}_t \mathbf{S}_t - \mathbf{T}_t\|^2 + w \sum_t \|\mathbf{S}_t - \mathbf{s}_t\|^2$$

Bayesian formulation

$$p(\mathbf{S}_t | \mathcal{S}; \mathbf{p}) = \frac{1}{(2\hat{u})^{d_j} p_j} e^{-\frac{1}{2} (\mathbf{s}_t - \hat{\mathbf{g}})^T \mathbf{p}^{-1} (\mathbf{s}_t - \hat{\mathbf{g}})}$$

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

$$\begin{aligned} &\text{maximize} && p(\mathbf{R}; \mathbf{S}; \mathbf{T}; \mathcal{S}; \mathbf{p} | \mathbf{P}) \\ &&& / p(\mathbf{P} | \mathbf{R}; \mathbf{S}; \mathbf{T}) p(\mathcal{S}; \mathbf{p}) \end{aligned}$$

How does it work?

Maximize $p(R; S; T; \mathcal{S} | P)$
 $\quad \quad \quad / \quad p(P | R; S; T) p(\mathcal{S} | p)$


Minimize

$$\sum_t \ln p(P_t | R_t; S_t; T_t) - \sum_t \ln p(S_t | \mathcal{S}; p)$$

$$= \sum_t \left[\ln p_t \left(\frac{R_t S_t}{T_t} \right)^2 + \frac{N}{2} \ln \hat{u}^2 \right. \\ \left. + \frac{1}{2} (S_t - \mathcal{S})^T p^{-1} (S_t - \mathcal{S}) + \frac{T}{2} \ln |p| \right]$$

(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**