# CS329 Probabilistic Robotics

Today's Topic: 3D Mapping with EM

Sebastian Thrun
Carnegie Mellon University
www.cs.cmu.edu/~thrun

## The Problem

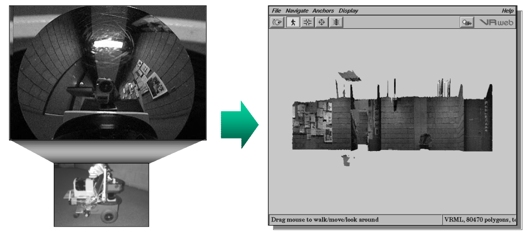

- Forward-pointed laser (2D information)
- Upward-pointed laser (3D information)
- Panoramic camera

## The (Almost) Raw Range Data

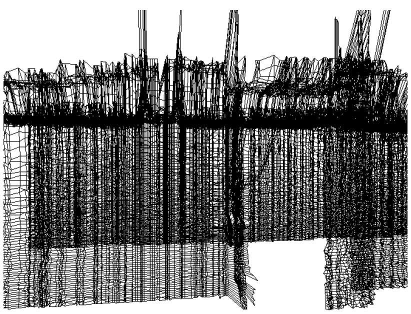

## 3D Mapping

- Observation: World mostly composed of planar surfaces

- Problem: How to generate multi-plane model?

## The 3D Multi-Plane Mapping Problem

- Entails a range of problems
  - Parameter estimation: Each surface presently possesses 6 parameters $(x, y, z, \theta, \phi, \psi)$
  - Outlier identification: Shall we explain every single measurement by a planar surface?
  - Correspondence: What measurement corresponds to what surface?
  - Model selection: How many surfaces should there be?
  - Robot localization: Noise in robot motion has a systematic effect on multiple measurements (this spoils any independence assumption)

## Concurrent Mapping and Localization: To be Discussed in Later Cllases

occupancy grid map

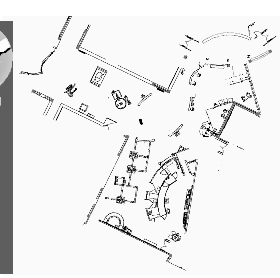

CAD map

## The 3D Multi-Plane Mapping Problem

- Entails a range of problems
  - Parameter estimation: Each surface presently possesses 6 parameters ($x, y, z, \theta, \phi, \psi$)
  - Outlier identification: Shall we explain every single measurement by a surface?
  - Correspondence: What measurement corresponds to what surface?
  - Model selection: How many surfaces should there be?
  - Robot localization: Noise in robot motion has a systematic effect on multiple measurements (this spoils any independence assumption)

## 3D Structure Mapping (Raw Data)

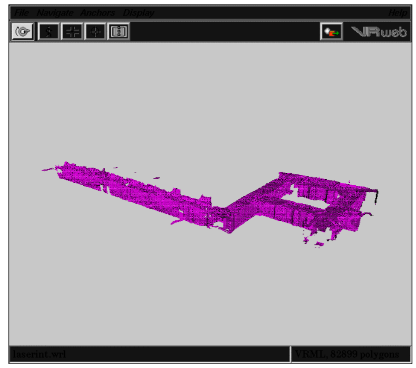

## 3D Texture Mapping (Raw Data)


## Fine-Grained Structure

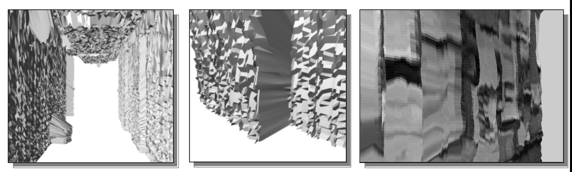

- Can we do better?

## Fine-Grained Structure (Wall data): Scatter plot of surface normals

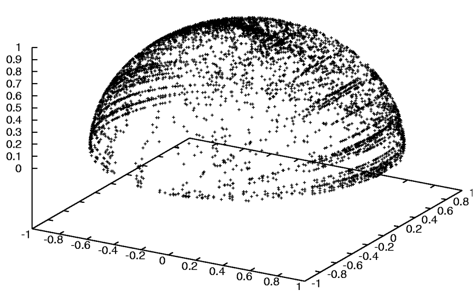


By W. Burgard and D. Hähnel

## The Remaining Problem

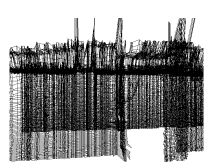

Data

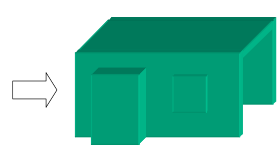

Mostly Multi-Planar model

## Clustering Data with K-Means

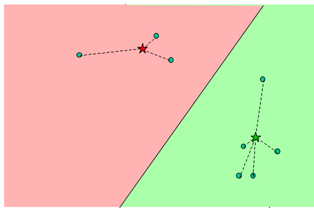

## EM: Generalized K-Means

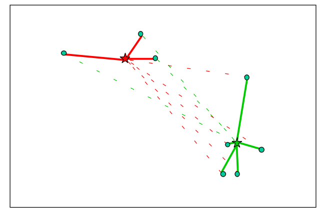

## EM for Fitting Multi-Planar Models

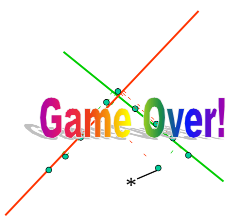


## Probabilistic Model

- 3D Model: $\theta = \{\theta_1, \theta_2, \ldots, \theta_J\}$

Planar surface in 3D

$$\theta_j = \langle \alpha_j, \beta_j \rangle \in \Re^3 \times \Re$$

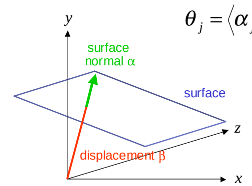


Distance point-surface

$$\mathrm{dist}(\theta_j, z_i) = |\alpha_j \cdot z_i - \beta|$$

## Measurement Model

- Case 1: Measurement $z_i$ caused by plane $\theta_j$

$$p(z_i \mid \theta_j) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{1}{2}\frac{(\alpha_j \cdot z_i - \beta_j)^2}{\sigma^2}}$$

- Case 2: $z_i$ caused by non-planar object

$$p(z_i \mid \theta_*) = \frac{1}{z_{\max}} = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{1}{2}\ln\frac{z_{\max}^2}{2\pi\sigma^2}}$$

## Measurement Model

$$p(z_i \mid \theta, c_1, \ldots, c_J, c_*) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{1}{2}\left(c_* \ln\frac{z_{\max}^2}{2\pi\sigma^2} + \sum_{j=1}^{J} c_j \frac{(\alpha_j \cdot z_i - \beta_j)^2}{\sigma^2}\right)}$$

correspondence variables $C$:

$$c_*, c_j \in \{0,1\}$$

$$c_* + \sum_{j=1}^{J} c_j = 1$$

$$\Rightarrow p(Z \mid \theta, C) = \prod \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{1}{2}\left(c_{i*} \ln\frac{z_{\max}^2}{2\pi\sigma^2} + \sum_{j=1}^{J} c_{ij} \frac{(\alpha_j \cdot z_i - \beta_j)^2}{\sigma^2}\right)}$$

## Expected Log-Likelihood Function

$$p(Z\mid\theta,C) = \prod_{i=1} \frac{1}{\sqrt{2\pi\sigma^2}}\, e^{-\frac{1}{2}\left(c_{i*}\ln\frac{z_{\max}^2}{2\pi\sigma^2}+\sum_{j=1}^{J} c_{ij}\frac{(\alpha_j\cdot z_i-\beta_j)^2}{\sigma^2}\right)}$$

…after some simple math

$$E_c\left[\ln p(Z,C\mid\theta)\mid Z,\theta\right] = \sum_i \begin{pmatrix} \ln\frac{1}{(J+1)\sqrt{2\pi\sigma^2}} \\ -\frac{1}{2}E[c_{i*}\mid Z,\theta]\ln\frac{z_{\max}^2}{2\pi\sigma^2} \\ -\frac{1}{2}E[c_{ij}\mid Z,\theta]\frac{(\alpha_j\cdot z_i-\beta_j)^2}{\sigma^2} \end{pmatrix}$$

## The EM Algorithm

- Define: $Q(\theta'\mid\theta) = E_c[\ln p(Z,C\mid\theta')\mid Z,\theta] = \sum_i \begin{pmatrix} \ln\frac{1}{(J+1)\sqrt{2\pi\sigma^2}} \\ -\frac{1}{2}E[c_{i*}\mid Z,\theta]\ln\frac{z_{\max}^2}{2\pi\sigma^2} \\ -\frac{1}{2}E[c_{ij}\mid Z,\theta]\frac{(\alpha_j'\cdot z_i-\beta_j')^2}{\sigma^2} \end{pmatrix}$
- Initialization: Pick random $\theta$
- E-Step: Calculate expectations in $Q(\theta^{[n+1]}\mid\theta^{[n]})$
- M-Step: Generate New Map $\theta^{[n+1]} = \underset{\theta}{\operatorname{argmax}}\ Q(\theta\mid\theta^{[n]})$

## The E (Expectation) Step

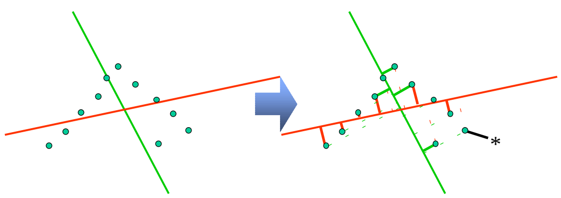


## E-Step

- Calculate expectations with fixed model $\theta$:

$$E[c_{i*}\mid Z,\theta^{[n]}] \propto e^{-\frac{1}{2}\ln\frac{z_{\max}^2}{2\pi\sigma^2}}$$

$$E[c_{ij}\mid Z,\theta^{[n]}] \propto e^{-\frac{1}{2}\frac{(\alpha_j\cdot z_i-\beta_j)^2}{\sigma^2}}$$

(normalize so that $E[c_{i*}\mid Z,\theta^{[n]}]+\sum_{j=1}^{J}E[c_{ij}\mid Z,\theta^{[n]}]=1$)

## The M (Maximization) Step

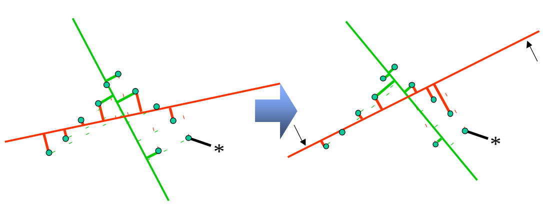


## M-Step

- Maximize

$$Q(\theta'\mid\theta) = E_c[\ln p(Z,C\mid\theta')\mid Z,\theta] = \sum \begin{pmatrix} \ln\frac{1}{(J+1)\sqrt{2\pi\sigma^2}} \\ -\frac{1}{2}E[c_{i*}\mid Z,\theta]\ln\frac{z_{\max}^2}{2\pi\sigma^2} \\ -\frac{1}{2}E[c_{ij}\mid Z,\theta]\frac{(\alpha_j'\cdot z_i-\beta_j')^2}{\sigma^2} \end{pmatrix}$$

subject to $\alpha_j'\cdot\alpha_j'=1$

- Is equivalent to minimizing

$\sum\sum_{j=1}^{J} E[c_{ij}\mid Z,\theta^{[n]}](\alpha_j'\cdot z_i-\beta_j')^2$ subject to $\alpha_j'\cdot\alpha_j'=1$

## M-Step: Solve via Lagrange Multipliers
Notice: the notation has been simplified

$$\sum_i \sum_{j=1}^{J} E[c_{ij}](\alpha_j \cdot z_i - \beta_j)^2 \to \min \qquad \forall j : \alpha_j \cdot \alpha_j = 1$$

- Define

$$L = \sum_i \sum_{j=1}^{J} E[c_{ij}](\alpha_j \cdot z_i - \beta_j)^2 + \sum_{j=1}^{J} \lambda_j \alpha_j \cdot \alpha_j$$

- And observe that

$$\frac{\partial L}{\partial \alpha_j} = E[c_{ij}](\alpha_j \cdot z_i - \beta_j) z_i - \lambda_j \alpha_j \overset{!}{=} 0$$

$$\frac{\partial L}{\partial \beta_j} = E[c_{ij}](\alpha_j \cdot z_i - \beta_j) \overset{!}{=} 0$$

$$\alpha_j \cdot \alpha_j = 1$$

## M-Step: Solve via Lagrange Multipliers

- Solve for $\beta$:

$$\frac{\partial L}{\partial \beta_j} = E[c_{ij}](\alpha_j \cdot z_i - \beta_j) \overset{!}{=} 0 \qquad \Rightarrow \beta_j = \frac{\sum_k E[c_{kj}]\alpha_j \cdot z_k}{\sum_k E[c_{kj}]\alpha_j}$$

- Substitute back:

$$\frac{\partial L}{\partial \alpha_j} = E[c_{ij}](\alpha_j \cdot z_i - \beta_j) z_i - \lambda_j \alpha_j \overset{!}{=} 0$$

$$\Rightarrow \frac{\partial L}{\partial \alpha_j} = E[c_{ij}](\alpha_j \cdot z_i - \frac{\sum_k E[c_{kj}]\alpha_j \cdot z_k}{\sum_k E[c_{kj}]\alpha_j}) z_i - \lambda_j \alpha_j \overset{!}{=} 0$$

- Is of the linear form:

$$A_j \alpha = \lambda_j \alpha_j$$

- Solution: two Eigenvectors of $A$ with smallest Eigenvalues.

## The 3D Multi-Plane Mapping Problem

- Entails a range of problems
  - Parameter estimation: Each surface presently possesses parameters $(x, y, z, \theta, \phi, \psi)$
  - Outlier identification: Shall we explain every single measurement by a surface?
  - Correspondence: What measurement corresponds to what surface?
  - Model selection: How many surfaces should there be?
  - Robot localization: Noise in robot motion has a systematic effect on multiple measurements (this spoils any independence assumption)

## Choosing the "Right" Number of Planes: AIC

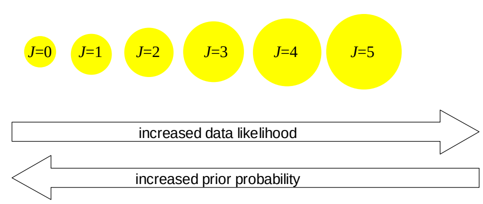


$$\log p(J \mid d) = const + \log p(d \mid J) + \log p(J)$$

## Model Selection

Approximately every 20 iterations of EM:

- Start new surfaces
  - Near any set of collinear measurements
- Terminate unsupported surfaces
  - If not supported by enough measurements
  - If density of measurements too low
  - If two planes are too close to each other

## Quality of 3D model

Without EM With EM

## Quality of 3D model

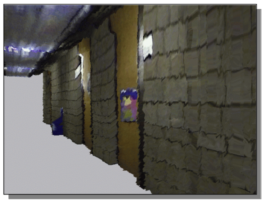

Without EM

With EM

## Numerical Results

$J$ (number of surfaces in $\theta$)

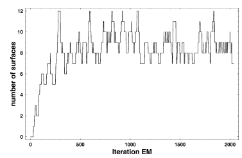


fraction measurements explained by $\theta$

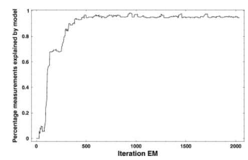


## Advantages and Limitations

- Increased compactness, accuracy ☺
- Not real-time ☹
- Only flat surfaces ☹

## Towards Online EM

- Each Measurement is only considered in $\leq N$ E-steps
- When doing so, each measurement only compared to $\leq K$ surfaces
- Each surface calculated from $\leq M$ measurements
- New surfaces only started from recent measurements, discarded if not growing fast
- …with small $N$, $K$, and $M$

## Online EM and Model Selection

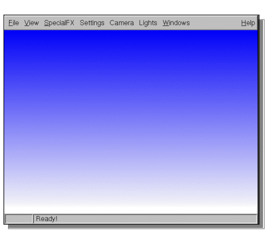

raw data

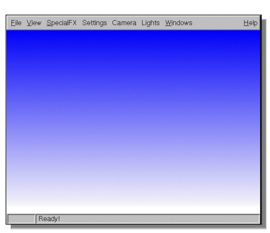

mostly planar map with Online EM

## Relation to Bayes Filters??

- Maximizes posterior likelihood

$$\operatorname*{argmax}_{\theta} E_{c_{1,t}}[\log p(z_{1,t}, c_{1,t} \mid \theta) \mid z_{1,t}, \theta] + \log p(\theta)$$

$$\approx \operatorname*{argmax}_{\theta} p(\theta \mid z_{1,t})$$

instead of computing posterior incrementally

$$p(\theta \mid z_{1,t}) = f(p(\theta \mid z_{1,t-1}), z_t)$$

- What does this imply?