Probabilistic Fitting

- Generative probabilistic model
  - Tells a story about how stochastic data comes to be
  - Darts fall around the center of the board, but where exactly?
  - Consider a model with parameters, $\boldsymbol{\theta}$
  - Consider an observation, x_i
  - We denote the probability of seeing x_i under the model by:

$$p(x_i \mid \Theta)$$

Read "given" or "conditioned on"
Restricts to the case of $\boldsymbol{\theta}$

Defined by $P(A \mid B) = \frac{P(A,B)}{P(B)}$

Probabilistic Fitting

- Multiple observations
  - Suppose we have multiple observations, in a vector $\mathbf{x}$
  - What is the probability of $\mathbf{x}$?
- If observations are independent then probability is the product of the individual observations
  - Essentially a definition, but is consistent with intuition
  - The observations are conditionally independent **given** the model
- So, the probability of x is then:

$$p(\mathbf{x} \mid \Theta) = \prod p(x_i \mid \Theta)$$

Probabilistic Fitting

- So, given the model, we have the probability of observing the data

$$p(\mathbf{x} \mid \Theta) = \prod p(x_i \mid \Theta)$$

- But what we really want is the probability of the model (parameters) given the data!
- Bayes rule comes to the rescue!

Bayes Rule

- Bayes rule: $P(A \mid B) = \frac{P(B \mid A)P(A)}{P(B)}$
- Proof $P(A,B) = P(B \mid A)P(A) = P(A \mid B)P(B)$
- With our notation: $P(\Theta \mid \mathbf{x}) = \frac{P(\mathbf{x} \mid \Theta)P(\Theta)}{P(\mathbf{x})}$

likelihood function for the parameters

prior probability (often taken to be uniform)

$$P(\Theta \mid \mathbf{x}) = \frac{P(\mathbf{x} \mid \Theta) P(\Theta)}{P(\mathbf{x})}$$

posterior probability

normalizer, often is not of interest

Common special case
$P(\Theta \mid \mathbf{x}) \propto P(\mathbf{x} \mid \Theta)$

Know the words in **red**

Probabilistic Fitting

- If we assume **uniform** prior, then we can find the posterior density for the parameters by:

$$p(\Theta \mid \mathbf{x}) \propto p(\mathbf{x} \mid \Theta)$$

- Now the objective is to find the parameters Θ such that this *likelihood* is maximum
- Note--this is the same as finding the parameters which minimize the **negative log likelihood**

Probabilistic fitting with independence and uniform prior

Finding the "best" model under simple circumstances

$\underset{\Phi}{\text{maximize}}\ \mathrm{p}(\Theta \mid \mathbf{x})$ (one definition of best Θ)

$\underset{\Phi}{\text{maximize}}\ \mathrm{p}(\mathbf{x} \mid \Theta)$ (by Bayes rule, uniform prior)

$\underset{\Phi}{\text{minimize}}\ -\log(\mathrm{p}(\mathbf{x} \mid \Theta))$ (log is monotonic increasing)

$\underset{\Phi}{\text{minimize}}\ -\log\left(\prod p(x_i \mid \Theta)\right)$ (by independence)

$\underset{\Phi}{\text{minimize}}\ -\sum \log(p(x_i \mid \Theta))$ (high school math)

- Back to lines: ax+by+c=0 where a^2+b^2=1
- Algebraic fact: Distance squared from (x,y) to this line is (ax+by+c)^2
- **Generative model** for lines: Choose point on line, and then, with probability proportional to p(d), **normally distributed** (Gaussian), go a distance d from the line.
- Now the probability of an observed (x,y) is given by

$$p((x,y) \mid \Theta) \propto \exp\left(-\frac{(ax+by+c)^2}{2\sigma^2}\right)$$

Lines

Convenient formula for line
$ax+by+c=0$
where a^2+b^2=1

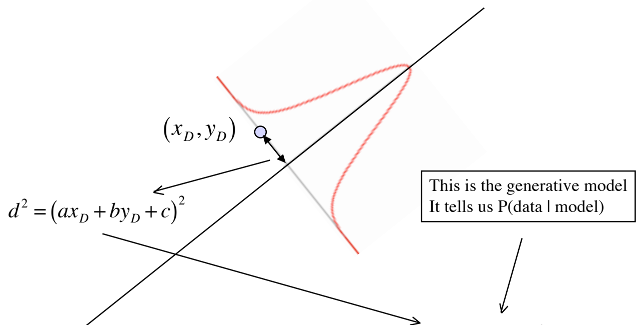

$$p((x_D, y_D) \mid \Theta) \propto \exp\left(-\frac{(ax_D + by_D + c)^2}{2\sigma^2}\right)$$

We have the probability density of the observed (x,y) given by

$$p((x,y) \mid \Theta) \propto \exp\left(-\frac{(ax+by+c)^2}{2\sigma^2}\right)$$

The negative log is

$$\frac{(ax+by+c)^2}{2\sigma^2}$$

And the negative log likelihood of multiple observations is

$$\frac{1}{2\sigma^2}\sum_i (ax_i + by_i + c)^2$$

From the previous slide, we had that the negative log likelihood of multiple observations is given by

$$\frac{1}{2\sigma^2}\sum_i (ax_i + by_i + c)^2 \qquad (\text{where } a^2 + b^2 = 1)$$

This should be recognizable as homogeneous least squares

Thus we have shown that least squares is maximum likelihood estimation under normality (Gaussian) error statistics!

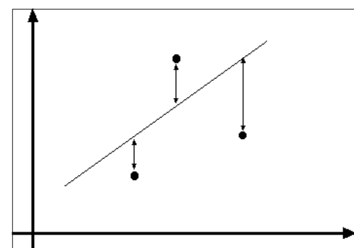

Line fitting by minimizing error == maximum likelihood estimation

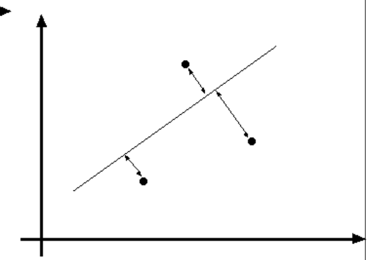

Fitting curves other than lines

- In principle, an easy generalization
  - Assuming Gaussian error statistics, Euclidean distance is a good measure
  - The probability of obtaining a point, given a curve, is given by a negative exponential of distance squared
- In practice, this can be hard
  - It can be difficult to compute the distance between a point and a curve
  - Circles, ellipses, and a few others are not too hard
  - Otherwise, craft an approximation
  - §15.3 has more

More on the Bayesian Method

- Recall that a generative probabilistic model
  - Tells a story about how stochastic data comes to be
  - Provides likelihood given data given model

$$p(\{\mathbf{x}_i\} \mid \Theta)$$

- Bayes rule
  - Tells us how to go *from* data given model *to* model given data
  - Tell us how to combine prior knowledge and evidence from data
  - Gives a probability distribution for an answer → $p(\Theta \mid \mathbf{x})$
    - Ideal for further reasoning
    - Supports various estimates (see cartoon on next slide)
    - Supports "risk" functions

Bayesian Estimators

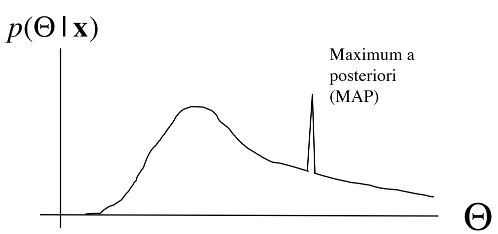

Information from Priors and Data

- Recall that vision problems do not have unique solutions!
  - We have to choose solutions suggested both by data and by what we believe (world knowledge)
  - What we believe about the world is the the prior