Outline

1) Hierarchical Bayes
2) Markov Chain Monte Carlo
3) Gibbs Sampler

Hierarchical Bayes

[Full power of Bayes is realized in large, complex problems with repeat structure, allowing us to pool information across many observations.]

Ex Predict a batter's "true" batting average from n at-bats. $X = \#$ of hits $\sim \text{Binom}(n, \theta)$

Prior info: Most batting avg.s are between 0.1 and 0.3, $\theta = 0.8$ very unlikely. Can represent using a Beta dist., but how to pick α, β?

Solution: Pool info across players $i = 1, \ldots, m$ via hierarchical model

$$\alpha, \beta \sim \lambda_{\alpha,\beta} \qquad (\text{say, indep. } \text{Exp}(1))$$

$$\theta_i \mid \alpha, \beta \overset{\text{iid}}{\sim} \text{Beta}(\alpha, \beta) \qquad i \leq m$$

$$X_i \mid \theta_i \overset{\text{indep}}{\sim} \text{Binom}(n_i, \theta_i) \qquad i \leq m$$

[**Note:** there is always an equivalent model where we marginalize over α, β and just write a more complicated prior on θ. Hierarchical version often gives better intuition or suggests computational strategies]

Graphical form:

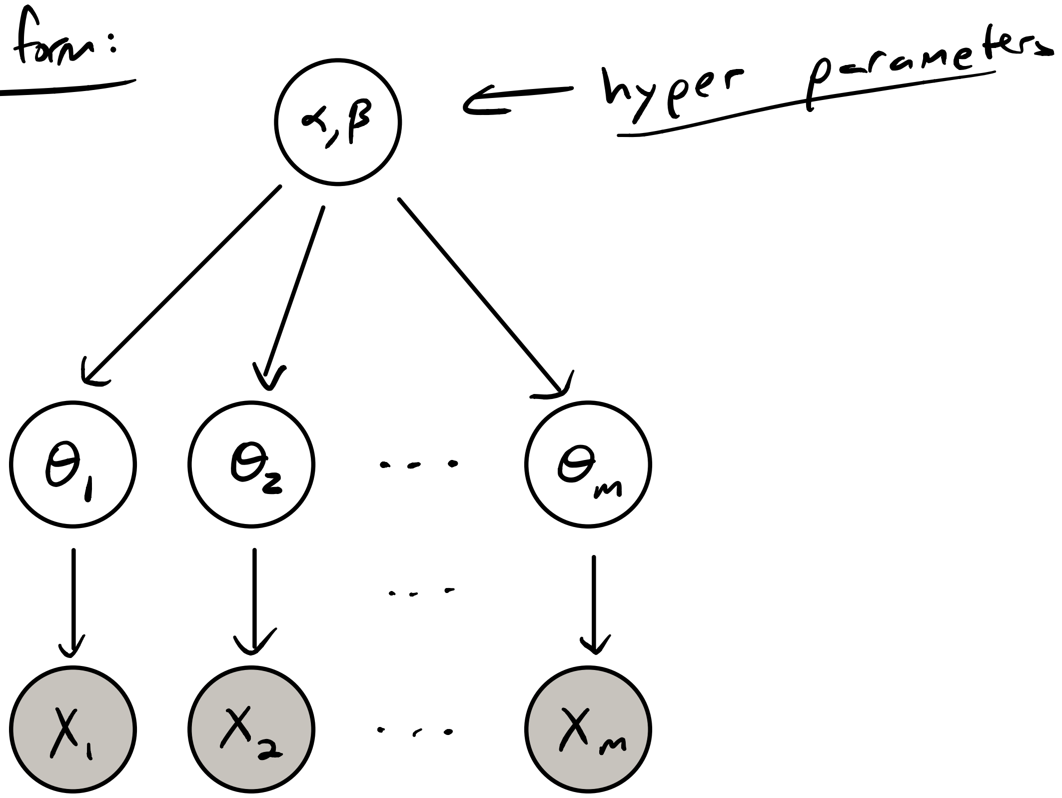

This is a directed graphical model. Implies the distribution may be factorized with one factor for each vertex in a DAG (V, E)

$$p(z_1, \ldots, z_{|V|}) = \prod_{i=1}^{|V|} p_i(z_i \mid \underset{\nwarrow \text{parents}}{Pa(z_i)})$$

For this model,

$$p(\alpha, \beta, \theta_1, \ldots, \theta_m, x_1, \ldots, x_m)$$

$$= p(\alpha, \beta) \cdot \prod_i p(\theta_i \mid \alpha, \beta) \cdot \prod_i p(x_i \mid \theta_i)$$

Practical implication:

$X_2, \ldots, X_m$ indirectly influence the estimate of X_1, by teaching us what values of θ are plausible.

Markov Chain Monte Carlo

Hierarchical models can get very complex very fast, creating big computational headaches

$$\lambda(\theta / x) = \frac{p_\theta(x)\,\lambda(\theta)}{\int_\Omega p_\xi(x)\,\lambda(\xi)\,d\xi} \qquad \begin{array}{l} \longleftarrow \text{usually nice} \\[3ex] \longleftarrow \text{often intractable.} \end{array}$$

Computational strategy: set up a Markov chain with stationary dist $\propto p_\theta(x)\lambda(\theta)$, run it to get approximate samples from $\lambda(\theta|x)$

Definition: A (stationary) Markov chain with trans. kernel $Q(y|x)$ and initial dist. $\pi_0(x)$ is a sequence of r.v.s $X^{(0)}, X^{(1)}, \ldots$ where $X^{(0)} \sim \pi_0$ and $X^{(t+1)} \mid X^{(0)}, \ldots, X^{(t)} \sim Q(\cdot \mid X^{(t)})$

$$Q(y|x) = \mathbb{P}(X^{(t+1)} = y \mid X^{(t)} = x)$$

Marginal dist. of $X^{(1)}$:

$$\pi_1(y) = \mathbb{P}(X^{(1)} = y) = \int_{\mathcal{X}} Q(y|x)\,\pi_0(x)\,d\mu(x)$$

This is a directed graphical model:

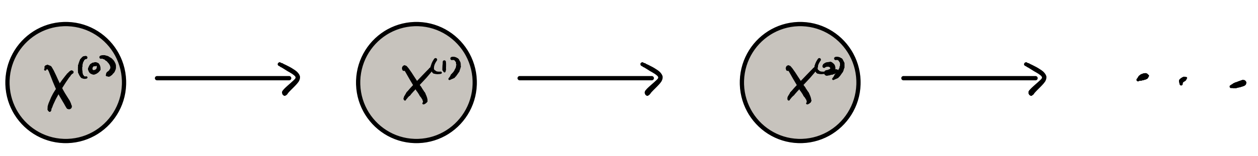

If $\pi(y) = \int_{\mathcal{X}} Q(y|x)\pi(x)\,d\mu(x)$ we say π is a **stationary distribution** for Q

Sufficient condition is **detailed balance:**

$$\pi(x)Q(y|x) = \pi(y)Q(x|y) \quad \forall\, x,y$$

$$\Rightarrow \int_{\mathcal{X}} Q(y|x)\pi(x)\,d\mu(x) = \pi(y)\int_{\mathcal{X}} Q(x|y)\,d\mu(x)$$

$$= \pi(y)$$

A Markov chain with detailed balance is called **reversible** since $(X^{(0)},\dots,X^{(t)}) \overset{D}{=} (X^{(t)},\dots,X^{(0)})$

Theorem: If an MC with stationary dist. π is: for cts $\mathcal{X}$

1) Irreducible: $\forall x,y\ \exists n: p(X^{(n)} \overset{\in A}{=y} \mid X^{(0)}=x) > 0$

2) Aperiodic: $\forall x,\ \gcd\{n>0: p(X^{(n)}=x \mid X^{(0)}=x)>0\}=1$

can be generalized to cts $\mathcal{X}$

Then $\mathcal{L}(X^{(t)}) \xrightarrow{t\to\infty} \pi$ (in TV distance), regardless of π_0 (chain "forgets" π_0)

[Proof beyond scope of our class]

Strategy: Find Q with stationary dist $\lambda(\theta|X)$, start at any X, run chain for a long time

⇝ $X^{(t)} \approx$ sample from posterior, for large t.

Gibbs sampler

Parameter vector $\theta = (\theta_1, \dots, \theta_d)$

Algorithm:

```
Initialize $\theta = \theta^{(0)}$
For $t = 1, \dots, T$:
    For $j = 1, \dots, d$:                                  }
        Sample $\theta_j \sim \lambda(\theta_j \mid \theta_{\sim j}, X)$   } (*)
    Record $\theta^{(t)} = \theta$
```

Variations on (*):

- Update one random coordinate $J^{(t)} \sim \text{Unif}\{0, \dots, d\}$
- Update coordinates in random order

Advantage for hierarchical priors: only need to sample low-dimensional conditional dists:

$$\lambda(\theta_j \mid \theta_{\sim j}, X) \propto p(\theta_j \mid \theta_{Pa(j)}) \cdot \prod_{i : j \in Pa(i)} p(\theta_i \mid \theta_{Pa(i)})$$

Especially easy if using conjugate priors at all levels, often can be parallelized.

MCMC in Practice

In theory: Pick any initialization $\theta^{(0)}$ and valid kernel Q, sample long enough* ⇝ $\theta^{(t)} \approx \lambda(\theta \mid x)$

Do it again N more times ⇝ N samples from $\lambda(\theta \mid x)$

In practice, how do we know we've sampled long enough?

Trace plots: Show how fast the MC mixes

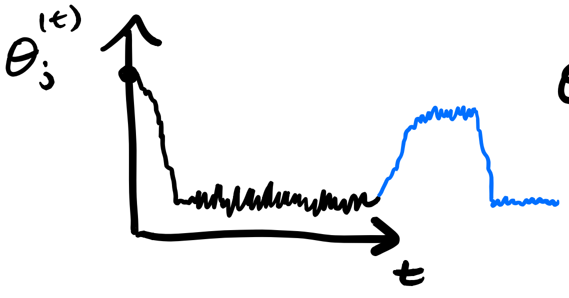

GOOD (?)

Can be deceived!

Esp. for bimodal posterior

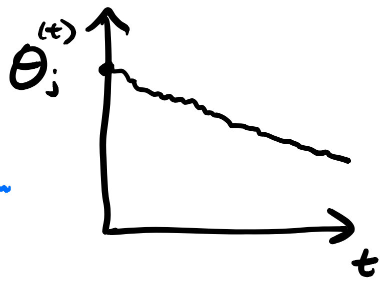

BAD

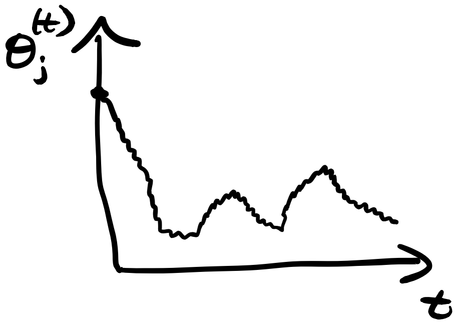

NOT GREAT

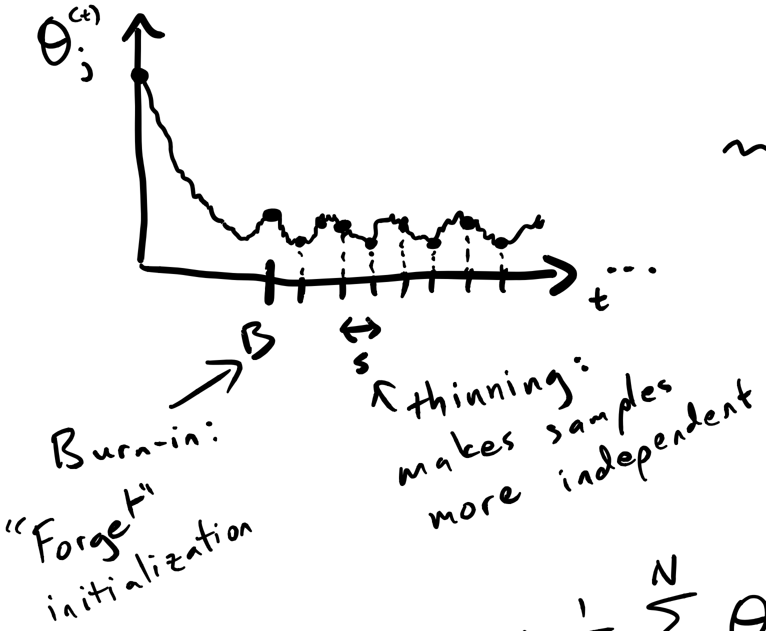

⇝

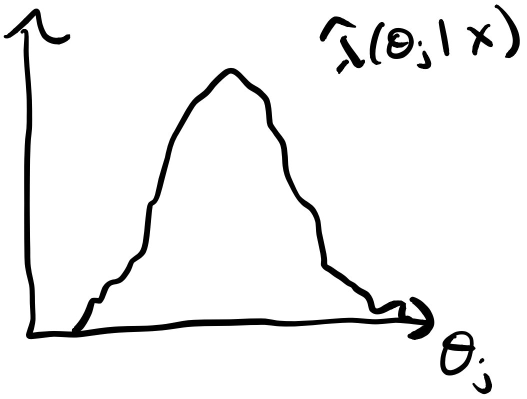

Estimate posterior based on $\{\theta_j^{(B)}, \theta_j^{(B+s)}, \dots, \theta_j^{(B+Ns)}\}$

Posterior mean: $\frac{1}{N+1}\sum_{k=0}^{N} \theta_j^{(B+ks)} \xrightarrow[N\to\infty]{} \mathbb{E}[\theta_j \mid x]$

Implementation details matter!

$$\theta_1, \theta_2 \overset{ind.}{\sim} N(0, 1)$$

$$X_i | \theta \overset{iid}{\sim} N(\theta_1 + \theta_2, 1) \qquad i = 1, \dots, n$$

$$\Rightarrow \begin{pmatrix} \theta_1 \\ \theta_2 \\ \bar{X} \end{pmatrix} \sim N_3\left(0, \begin{pmatrix} 1 & 0 & 1 \\ 0 & 1 & 1 \\ 1 & 1 & 2 + \frac{1}{n} \end{pmatrix}\right)$$

$$\theta | \bar{X} \sim N_2(\mu(\bar{x}), \Sigma(\bar{x}))$$

$$\mu(\bar{x}) = \begin{pmatrix} 1 \\ 1 \end{pmatrix} (2 + \tfrac{1}{n})^{-1} \bar{x} = \frac{n\bar{x}}{2n+1} \cdot \begin{pmatrix} 1 \\ 1 \end{pmatrix}$$

$$\Sigma(\bar{x}) = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} - \begin{pmatrix} 1 \\ 1 \end{pmatrix} (2 + \tfrac{1}{n})^{-1} (1 \ 1)$$

$$= \frac{n+1}{2n+1} \begin{pmatrix} 1 & \frac{-n}{n+1} \\ \frac{-n}{n+1} & 1 \end{pmatrix}$$

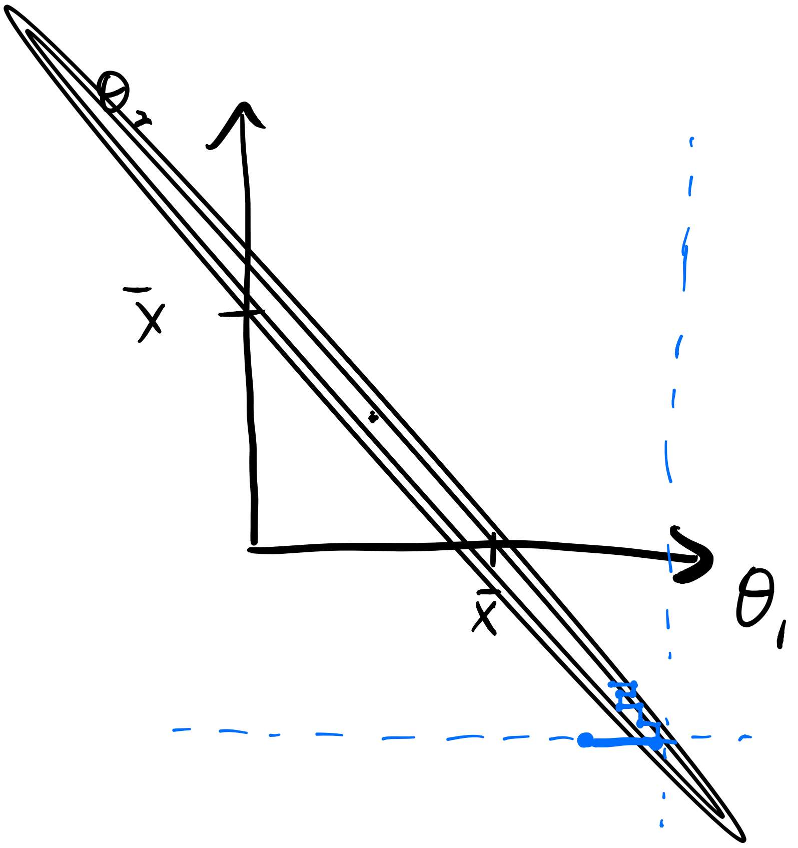

Gibbs takes a long time to mix

Better parameterization:

$$\beta_1 = \theta_1 + \theta_2$$
$$\beta_2 = \theta_1 - \theta_2$$

$$\beta_1 \perp\!\!\!\perp \beta_2 \mid X$$

Gibbs $\Leftrightarrow$ Directly sampling from posterior.

Gaussian Hierarchical Model:

$\tau^2 \sim \lambda(\tau)$ e.g. $\frac{1}{\tau^2} \sim \text{Gamma}(k, s)$

$\theta_i \mid \tau^2 \overset{iid}{\sim} N(0, \tau^2)$ $i \leq n$

$X_i \mid \tau^2, \theta \overset{ind.}{\sim} N(\theta_i, 1)$

Posterior mean:

$$\delta(X_i) = \mathbb{E}[\theta_i \mid X]$$

$$= \mathbb{E}\{\mathbb{E}[\theta_i \mid X, \tau^2\} \mid X]$$

$$= \mathbb{E}\left[\frac{\tau^2}{1+\tau^2} X_i \mid X\right]$$

$$= \underbrace{\left(\mathbb{E}\left[\frac{\tau^2}{1+\tau^2} \mid X\right]\right)}_{\text{Bayes estimate of optimal ``shrinkage'' factor}} \cdot X_i$$

Define $\zeta = \frac{1}{1+\tau^2}$ ($\zeta = 0 \Leftrightarrow$ no shrinkage)

$$X_i \mid \tau^2 \overset{iid}{\sim} N(0, 1+\tau^2)$$

$$\Rightarrow X \mid \zeta \sim N_n(0, \zeta^{-1} I_n)$$

$$= \left(\frac{\zeta}{2\pi}\right)^{n/2} e^{-\zeta/2 \, \|X\|^2}$$

$$\propto_\zeta \text{Gamma}(\underbrace{1+\tfrac{n}{2}}_{\text{shape}}, \underbrace{2/\|X\|^2}_{\text{scale}})$$

← note this is just the likelihood

This likelihood has a sharp peak at $\frac{n+2}{\|X\|^2} \approx \zeta$

Why? $Z \sim \text{Gamma}\left(1+\frac{n}{2},\ {}^{2}/_{\|X\|^2}\right)$

has mean $\frac{n+2}{\|X\|^2}$ $\left(\xrightarrow{n\to\infty} \zeta\right)$

variance $\frac{n+4}{\|X\|^4}$ $\left(\xrightarrow{n\to\infty} 0\right)$

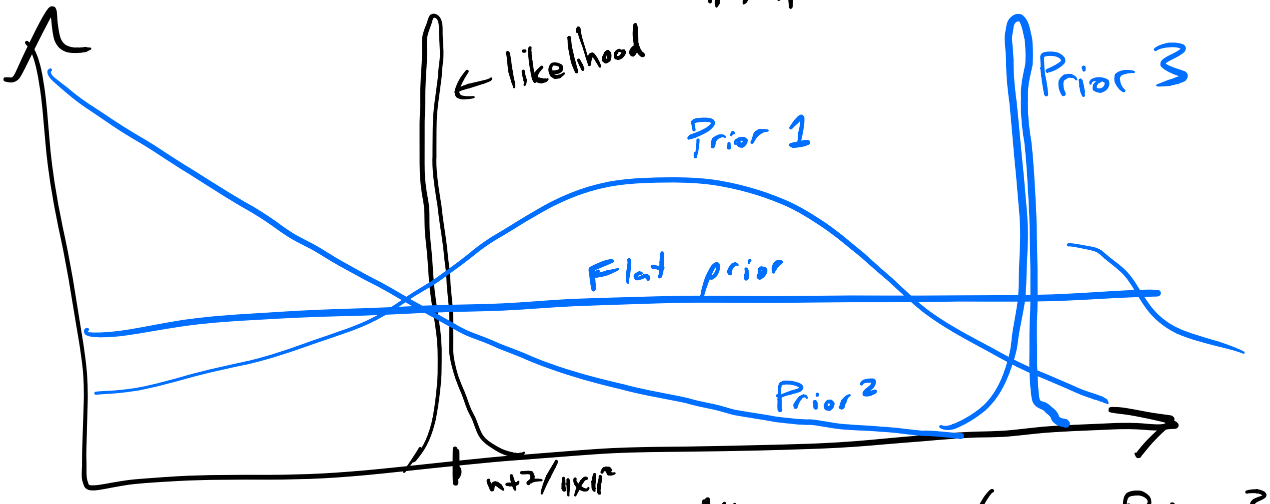

For any reasonably "open-minded" prior (not Prior 3),

$$\mathbb{E}[\zeta|X] \approx \zeta \Rightarrow \hat{\theta}_{\cdot} \approx (1-\zeta)X_i$$

If prior doesn't matter much, why use one?

Could just estimate ζ from data however we want, "plug it in"

Called "Empirical Bayes" a hybrid approach in which hyperparameters treated as fixed, others treated as random.