

Today on MA40189:

- **Gibbs sampler**: sample from a high-dimensional distribution by sampling from low-dimensional conditionals
- distribution of interest is $\pi(\theta)$ where $\theta = (\theta_1, \dots, \theta_d)$
- for each $p = 1, \dots, d$ obtain $\theta_p^{(t)}$ from **conditional distribution**

$$\pi(\theta_p \mid \theta_1^{(t)}, \dots, \theta_{p-1}^{(t)}, \theta_{p+1}^{(t-1)}, \dots, \theta_d^{(t-1)})$$

- Gibbs sampler can be viewed as a **special case** of the **Metropolis-Hastings algorithm**:
 - each iteration t consists of d Metropolis-Hastings steps
 - each with an **acceptance probability** of **one**
- sampling from the joint distribution of $\theta = (\theta_1, \theta_2)$ where

$$f(\theta_1, \theta_2) \propto \binom{n}{\theta_1} \theta_2^{\theta_1 + \alpha - 1} (1 - \theta_2)^{n - \theta_1 + \beta - 1}$$

for $\theta_1 \in \{0, 1, \dots, n\}$ and $0 \leq \theta_2 \leq 1$

- $f(\theta_1 \mid \theta_2) \propto f(\theta_1, \theta_2)$ as a function of θ_1
- $f(\theta_2 \mid \theta_1) \propto f(\theta_1, \theta_2)$ as a function of θ_2