Sampling Methods (contd)

CS772A: Probabilistic Machine Learning

Piyush Rai

Markov Chain Monte Carlo (MCMC)

If the target is a posterior, it will be conditioned on data, i.e., $p(\boldsymbol{z}|\boldsymbol{x})$

- Goal: Generate samples from some target distribution $p(\boldsymbol{z}) = \frac{\tilde{p}(\boldsymbol{z})}{Z_p}$

$\boldsymbol{z}$ usually is high-dim

- Assume we can evaluate $p(\boldsymbol{z})$ at least up to a proportionality constant

Means we can at least evaluate $\tilde{p}(\boldsymbol{z})$

- MCMC uses a Markov Chain which, when converged, starts giving samples from $p(z)$

$$\underbrace{\boldsymbol{z}^{(1)} \to \boldsymbol{z}^{(2)} \to \boldsymbol{z}^{(3)} \to}_{\text{initial samples typically garbage}} \ldots \to \underbrace{\boldsymbol{z}^{(L-2)} \to \boldsymbol{z}^{(L-1)} \to \boldsymbol{z}^{(L)}}_{\text{after convergence, actual samples from } p(\boldsymbol{z})}$$

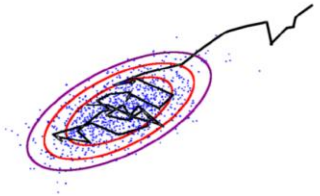

- Given current sample $\boldsymbol{z}^{(\ell)}$ from the chain, MCMC generates the next sample $\boldsymbol{z}^{(\ell+1)}$ as
  - Use a proposal distribution $q(\boldsymbol{z}|\boldsymbol{z}^{(\ell)})$ to generate a candidate sample $\boldsymbol{z}_*$
  - Accept/reject $\boldsymbol{z}_*$ as the next sample based on an acceptance criterion (will see later)
  - If accepted, set $\boldsymbol{z}^{(\ell+1)} = \boldsymbol{z}_*$. If rejected, set $\boldsymbol{z}^{(\ell+1)} = \boldsymbol{z}^{(\ell)}$

Should also have the same support as $p(\boldsymbol{z})$

- Important: The proposal distribution $q(\boldsymbol{z}|\boldsymbol{z}^{(\ell)})$ depends on the previous sample $\boldsymbol{z}^{(\ell)}$

MCMC: The Basic Scheme

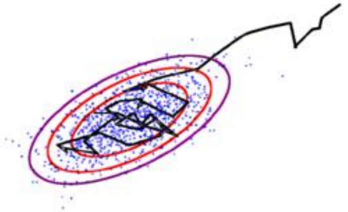

- The chain run infinitely long (i.e., upon convergence) will give ONE sample from $p(\boldsymbol{z})$
- But we usually require several samples to approximate $p(\boldsymbol{z})$

Thus we say that the samples are approximately from the target distribution

MCMC is exact in theory but approximate in practice since we can't run the chain for infinitely long in practice

- This is done as follows
  - Start the chain at an initial $\boldsymbol{z}^{(0)}$
  - Using the proposal $q(\boldsymbol{z}|\boldsymbol{z}^{(\ell)})$, run the chain long enough, say T_1 steps
  - Discard the first $T_1 - 1$ samples (called "burn-in" samples) and take last sample $\boldsymbol{z}^{(T_1)}$
  - Continue from $\boldsymbol{z}^{(T_1)}$ up to T_2 steps, discard intermediate samples, take last sample $\boldsymbol{z}^{(T_2)}$
    - This discarding (called "thinning") helps ensure that $\boldsymbol{z}^{(T_1)}$ and $\boldsymbol{z}^{(T_2)}$ are uncorrelated
  - Repeat the same for a total of S times
  - In the end, we now have S *approximately independent* samples from $p(\boldsymbol{z})$

Will treat it as our first sample from $p(\boldsymbol{z})$

Requirement for Monte Carlo approximation

- Note: Good choices for T_1 and $T_i - T_{i-1}$ (thinning gap) are usually based on heuristics

MCMC: Some Basic Theory

- A first order Markov Chain assumes $p\left(\boldsymbol{z}^{(\ell+1)}|\boldsymbol{z}^{(1)},\ldots,\boldsymbol{z}^{(\ell)}\right)=p(\boldsymbol{z}^{(\ell+1)}|\boldsymbol{z}^{(\ell)})$
- A 1st order Markov Chain $\boldsymbol{z}^{(0)},\boldsymbol{z}^{(1)},\ldots,\boldsymbol{z}^{(L)}$ is a sequence of r.v.'s and is defined by
  - An initial state distribution $p(\boldsymbol{z}^{(0)})$
  - A Transition Function (TF): $T_\ell\left(\boldsymbol{z}^{(\ell)}\rightarrow\boldsymbol{z}^{(\ell+1)}\right)=p(\boldsymbol{z}^{(\ell+1)}|\boldsymbol{z}^{(\ell)})$ – the proposal distribution
- TF is a distribution over the values of next state given the value of the current state
- Assuming $\boldsymbol{z}$ is discrete with K possible values, the TF will be $K\times K$ probability table

Transition probabilities can be defined using a *K*x*K* table if **z** is a discrete r.v. with *K* possible values

	1	2	3	4	5
1	0.4	0.6	0.0	0.0	0.0
2	0.5	0.0	0.5	0.0	0.0
3	0.0	0.3	0.0	0.7	0.0
4	0.0	0.0	0.1	0.3	0.6
5	0.0	0.3	0.0	0.5	0.2

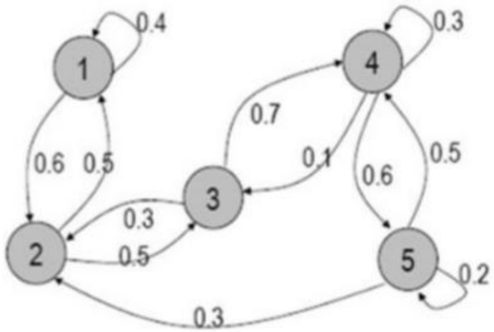

- Homogeneous Markov Chain: The TF is the same for all ℓ , i.e., $T_\ell=T$

MCMC: Some Basic Theory

- Consider the following Markov Chain to sample a discrete r.v. $\boldsymbol{z}$ with 3 possible values

The initial state distribution for $\boldsymbol{z}$

$$p(\boldsymbol{z}^{(0)}) = p\left(z_1^{(0)}, z_2^{(0)}, z_3^{(0)}\right)$$

Probabilities of the initial state taking each of the 3 possible values

$$= [0.5, 0.2, 0.3]$$

$$T = \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0.1 & 0.9 \\ 0.6 & 0.4 & 0 \end{bmatrix}$$

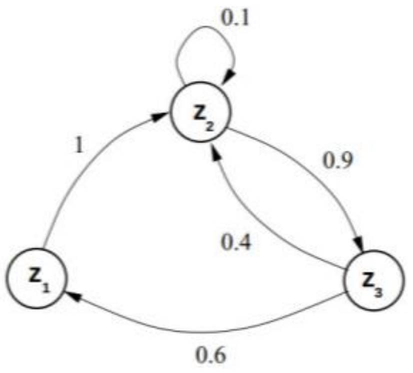

Distribution of $\boldsymbol{z}$ after taking the first step

$$p(\boldsymbol{z}^{(1)}) = p(\boldsymbol{z}^{(0)}) \times T = [0.2, 0.6, 0.2] \quad \text{(rounded to single digit after decimal)}$$

After doing it a few more (say some m) times

Stationary/Invariant Distribution $p(\boldsymbol{z})$ of this Markov Chain

$p(\boldsymbol{z})$ is multinoulli with $\pi = [0.2, 0.4, 0.4]$

$$p(\boldsymbol{z}^{(0)}) \times T^m = [0.2, 0.4, 0.4] \quad \text{(rounded to single digit after decimal)}$$

- $p(\boldsymbol{z})$ being Stationary means no matter what $p(\boldsymbol{z}^{(0)})$ is, we will reach $p(\boldsymbol{z})$
- Such transition functions are desirable in MCMC

MCMC: Some Basic Theory

- A Markov Chain with transition function $\boldsymbol{T}$ has stationary distribution $\boldsymbol{p(z)}$ if $\boldsymbol{T}$ satisfies

Known as the Detailed Balance condition

$$p(\boldsymbol{z})T(\boldsymbol{z}'|\boldsymbol{z}) = p(\boldsymbol{z}')T(\boldsymbol{z}|\boldsymbol{z}')$$

Here $\boldsymbol{T(b|a)}$ denotes the transition probability of going from state $\boldsymbol{a}$ to state $\boldsymbol{b}$

- Detailed Balance ensures "reversibility"
- Integrating out (or summing over) detailed balanced condition on both sides w.r.t. $\boldsymbol{z}'$

Thus $\boldsymbol{p(z)}$ is the stationary distribution of this Markov Chain

$$p(\boldsymbol{z}) = \int p(\boldsymbol{z}')T(\boldsymbol{z}|\boldsymbol{z}')d\boldsymbol{z}'$$

- Thus a Markov Chain with detailed balance always converges to a stationary distribution
- Detailed balance is sufficient but not necessary condition for having a stationary distr.

Some MCMC Algorithms

Metropolis-Hastings (MH) Sampling (1960)

- Suppose we wish to generate samples from a target distribution $p(\boldsymbol{z}) = \frac{\tilde{p}(\boldsymbol{z})}{Z_p}$
- Assume a suitable proposal distribution $q(\boldsymbol{z}|\boldsymbol{z}^{(\tau)})$, e.g., $\mathcal{N}(\boldsymbol{z}|\boldsymbol{z}^{(\tau)}, \sigma^2\boldsymbol{I})$
- In each step, draw $\boldsymbol{z}^*$ from $q(\boldsymbol{z}|\boldsymbol{z}^{(\tau)})$ and accept $\boldsymbol{z}^*$ with probability

Favors acceptance of $\boldsymbol{z}^*$ if it is more probable than $\boldsymbol{z}^{(\tau)}$ (under $p(\boldsymbol{z})$)

$$A(\boldsymbol{z}^*, \boldsymbol{z}^{(\tau)}) = \min\left(1, \frac{\tilde{p}(\boldsymbol{z}^*)\, q(\boldsymbol{z}^{(\tau)}|\boldsymbol{z}^*)}{\tilde{p}(\boldsymbol{z}^{(\tau)})\, q(\boldsymbol{z}^*|\boldsymbol{z}^{(\tau)})}\right)$$

Downweight the probability of acceptance of z^* if the proposal itself favors its generation (i.e., if $q(z^*|z^{(\tau)})$ is high), and upweight if it unfavors the generation

- Transition function of this Markov Chain
  - $T(\boldsymbol{z}^*|\boldsymbol{z}^{(\tau)}) = A(\boldsymbol{z}^*, \boldsymbol{z}^{(\tau)})q(\boldsymbol{z}^*|\boldsymbol{z}^{(\tau)})$ if state changed
  - $T(\boldsymbol{z}^*|\boldsymbol{z}^{(\tau)}) = q(\boldsymbol{z}^{(\tau)}|\boldsymbol{z}^{(\tau)}) + \sum_{\boldsymbol{z}^* \neq \boldsymbol{z}^{(\tau)}}(1 - A(\boldsymbol{z}^*, \boldsymbol{z}^{(\tau)}))q(\boldsymbol{z}^*|\boldsymbol{z}^{(\tau)})$ otherwise

Can Show that this TF satisfied detailed balance

The MH Sampling Algorithm

- Initialize $\boldsymbol{z}^{(1)}$ randomly
- For $\ell = 1,2, \dots, L$
  - Sample $\boldsymbol{z}^* \sim q(\boldsymbol{z}^*|\boldsymbol{z}^{(\ell)})$ and $u \sim \text{Unif}(0,1)$
  - Compute acceptance probability

$$A(\mathbf{z}^*, \mathbf{z}^{(\ell)}) = \min\left(1, \frac{\tilde{p}(\mathbf{z}^*)q(\mathbf{z}^{(\ell)}|\mathbf{z}^*)}{\tilde{p}(\mathbf{z}^{(\ell)})q(\mathbf{z}^*|\mathbf{z}^{(\ell)})}\right)$$

  - If $A\left(\boldsymbol{z}^*, \boldsymbol{z}^{(\ell)}\right) > u$

$$\boldsymbol{z}^{(\ell+1)} = \boldsymbol{z}^*$$

Meaning accepting $\boldsymbol{z}^*$ with probability $A\left(\boldsymbol{z}^*, \boldsymbol{z}^{(\ell)}\right)$

  - Else

$$\boldsymbol{z}^{(\ell+1)} = \boldsymbol{z}^{(\ell)}$$

MH Sampling in Action: A Toy Example..

- Target distribution $p(\boldsymbol{z}) = \mathcal{N}\left(\begin{bmatrix}4\\4\end{bmatrix}, \begin{bmatrix}1 & 0.8\\0.8 & 1\end{bmatrix}\right)$

- Proposal distribution $q(\boldsymbol{z}^{(t)}|\boldsymbol{z}^{(t-1)}) = \mathcal{N}\left(\boldsymbol{z}^{(t-1)}, \begin{bmatrix}0.01 & 0\\0 & 0.01\end{bmatrix}\right)$

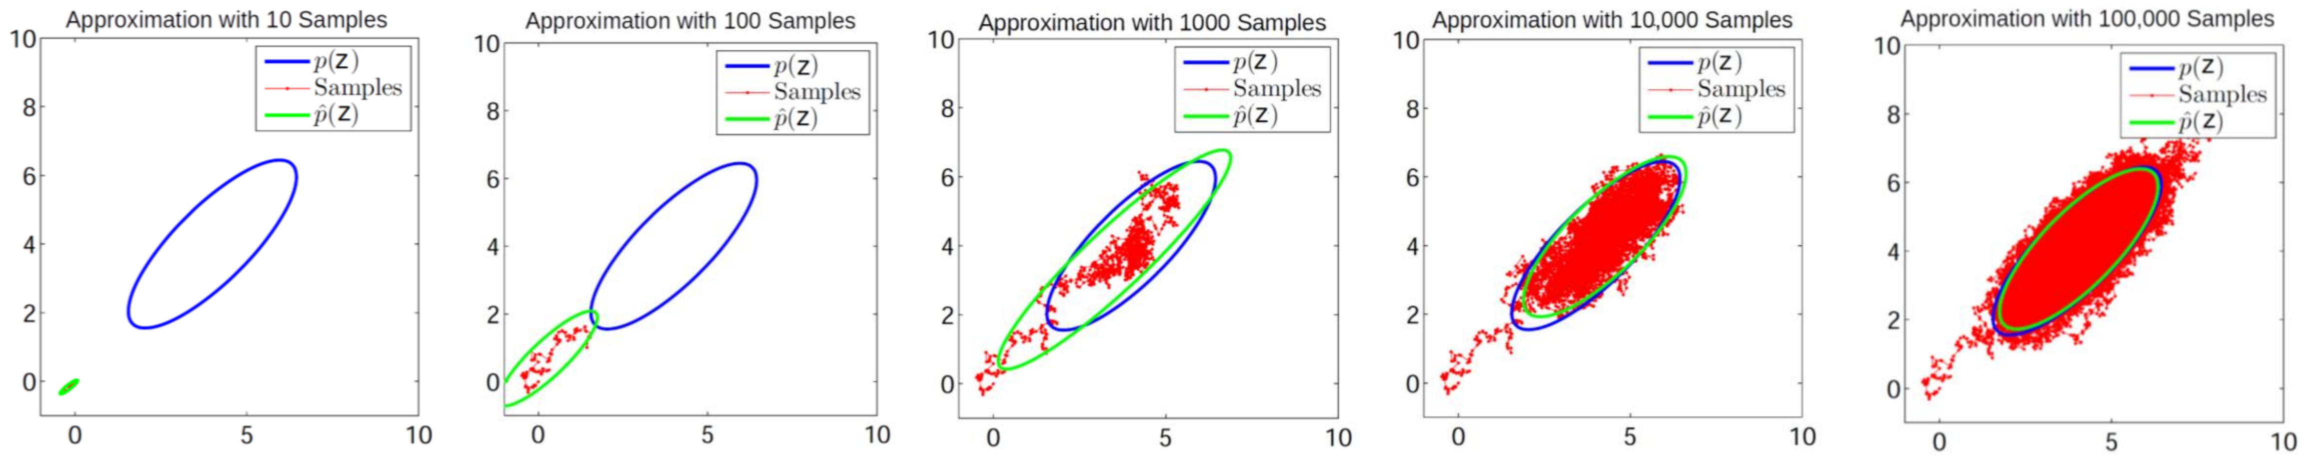

MH Sampling: Some Comments

- If prop. distrib. is symmetric, we get Metropolis Sampling algo (Metropolis, 1953) with

$$A(\mathbf{z}^*, \mathbf{z}^{(\tau)}) = \min\left(1, \frac{\tilde{p}(\mathbf{z}^*)}{\tilde{p}(\mathbf{z}^{(\tau)})}\right)$$

- Some limitations of MH sampling
  - Can sometimes have very slow convergence (also known as slow "mixing")

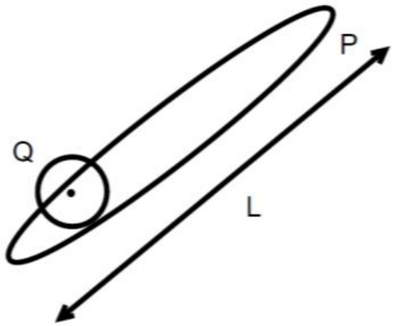

$Q(\mathbf{z}|\mathbf{z}^{(\tau)}) = \mathcal{N}(\mathbf{z}|\mathbf{z}^{(\tau)}, \sigma^2\mathbf{I})$

σ large $\Rightarrow$ many rejections

σ small $\Rightarrow$ slow diffusion

$\sim \left(\frac{L}{\sigma}\right)^2$ iterations required for convergence

  - Computing acceptance probability can be expensive*, e.g., if $p(\mathbf{z}) = \frac{\tilde{p}(\mathbf{z})}{Z_p}$ is some target posterior then $\tilde{p}(\mathbf{z})$ would require computing likelihood on all the data points (expensive)

*Austerity in MCMC Land: Cutting the Metropolis-Hastings Budget (Korattikara et al, 2014), Firefly Monte Carlo: Exact MCMC with Subsets of Data {(Maclaurin and Adams, 2015)

Gibbs Sampling (Geman & Geman, 1984)

- Goal: Sample from a joint distribution $\boldsymbol{p(z)}$ where $\boldsymbol{z} = [z_1, z_2, \dots, z_M]$
- Suppose we can't sample from $\boldsymbol{p(z)}$ but can sample from each conditional $p(z_i|\boldsymbol{z}_{-i})$
  - In Bayesian models, can be done easily if we have a locally conjugate model
- For Gibbs sampling, the proposal is the conditional distribution $p(z_i|\boldsymbol{z}_{-i})$
- Gibbs sampling samples from these conditionals in a cyclic order
- Gibbs sampling is equivalent to MH sampling with acceptance prob. = 1

Hence no need to compute it

$$A(\boldsymbol{z}^*, \boldsymbol{z}) = \frac{p(\boldsymbol{z}^*)q(\boldsymbol{z}|\boldsymbol{z}^*)}{p(\boldsymbol{z})q(\boldsymbol{z}^*|\boldsymbol{z})} = \frac{p(z_i^*|\boldsymbol{z}^*_{-i})p(\boldsymbol{z}^*_{-i})p(z_i|\boldsymbol{z}^*_{-i})}{p(z_i|\boldsymbol{z}_{-i})p(\boldsymbol{z}_{-i})p(z_i^*|\boldsymbol{z}_{-i})} = 1$$

where we use the fact that $\boldsymbol{z}^*_{-i} = \boldsymbol{z}_{-i}$

Since only one component is changed at a time

Gibbs Sampling: Sketch of the Algorithm

- $\boldsymbol{M}$: Total number of variables, $\boldsymbol{T}$: number of Gibbs sampling iterations

1. Initialize $\{z_i : i = 1, \ldots, M\}$
2. For $\tau = 1, \ldots, T$:
   - Sample $z_1^{(\tau+1)} \sim p(z_1 | z_2^{(\tau)}, z_3^{(\tau)}, \ldots, z_M^{(\tau)})$.
   - Sample $z_2^{(\tau+1)} \sim p(z_2 | z_1^{(\tau+1)}, z_3^{(\tau)}, \ldots, z_M^{(\tau)})$.

     $\vdots$
   - Sample $z_j^{(\tau+1)} \sim p(z_j | z_1^{(\tau+1)}, \ldots, z_{j-1}^{(\tau+1)}, z_{j+1}^{(\tau)}, \ldots, z_M^{(\tau)})$.

     $\vdots$
   - Sample $z_M^{(\tau+1)} \sim p(z_M | z_1^{(\tau+1)}, z_2^{(\tau+1)}, \ldots, z_{M-1}^{(\tau+1)})$.

Assuming $\boldsymbol{z} = [z_1, z_2, \ldots, z_M]$

CP of each component of $\boldsymbol{z}$ uses the most recent values (from this or the previous iteration) of all the other components

Each iteration will give us one sample $\boldsymbol{z}^{(\tau)}$ of $\boldsymbol{z} = [z_1, z_2, \ldots, z_M]$

- Note: Order of updating the variables usually doesn't matter (but see "Scan Order in Gibbs Sampling: Models in Which it Matters and Bounds on How Much" from NIPS 2016)

Gibbs Sampling: A Simple Example

- Can sample from a 2-D Gaussian using 1-D Gaussians

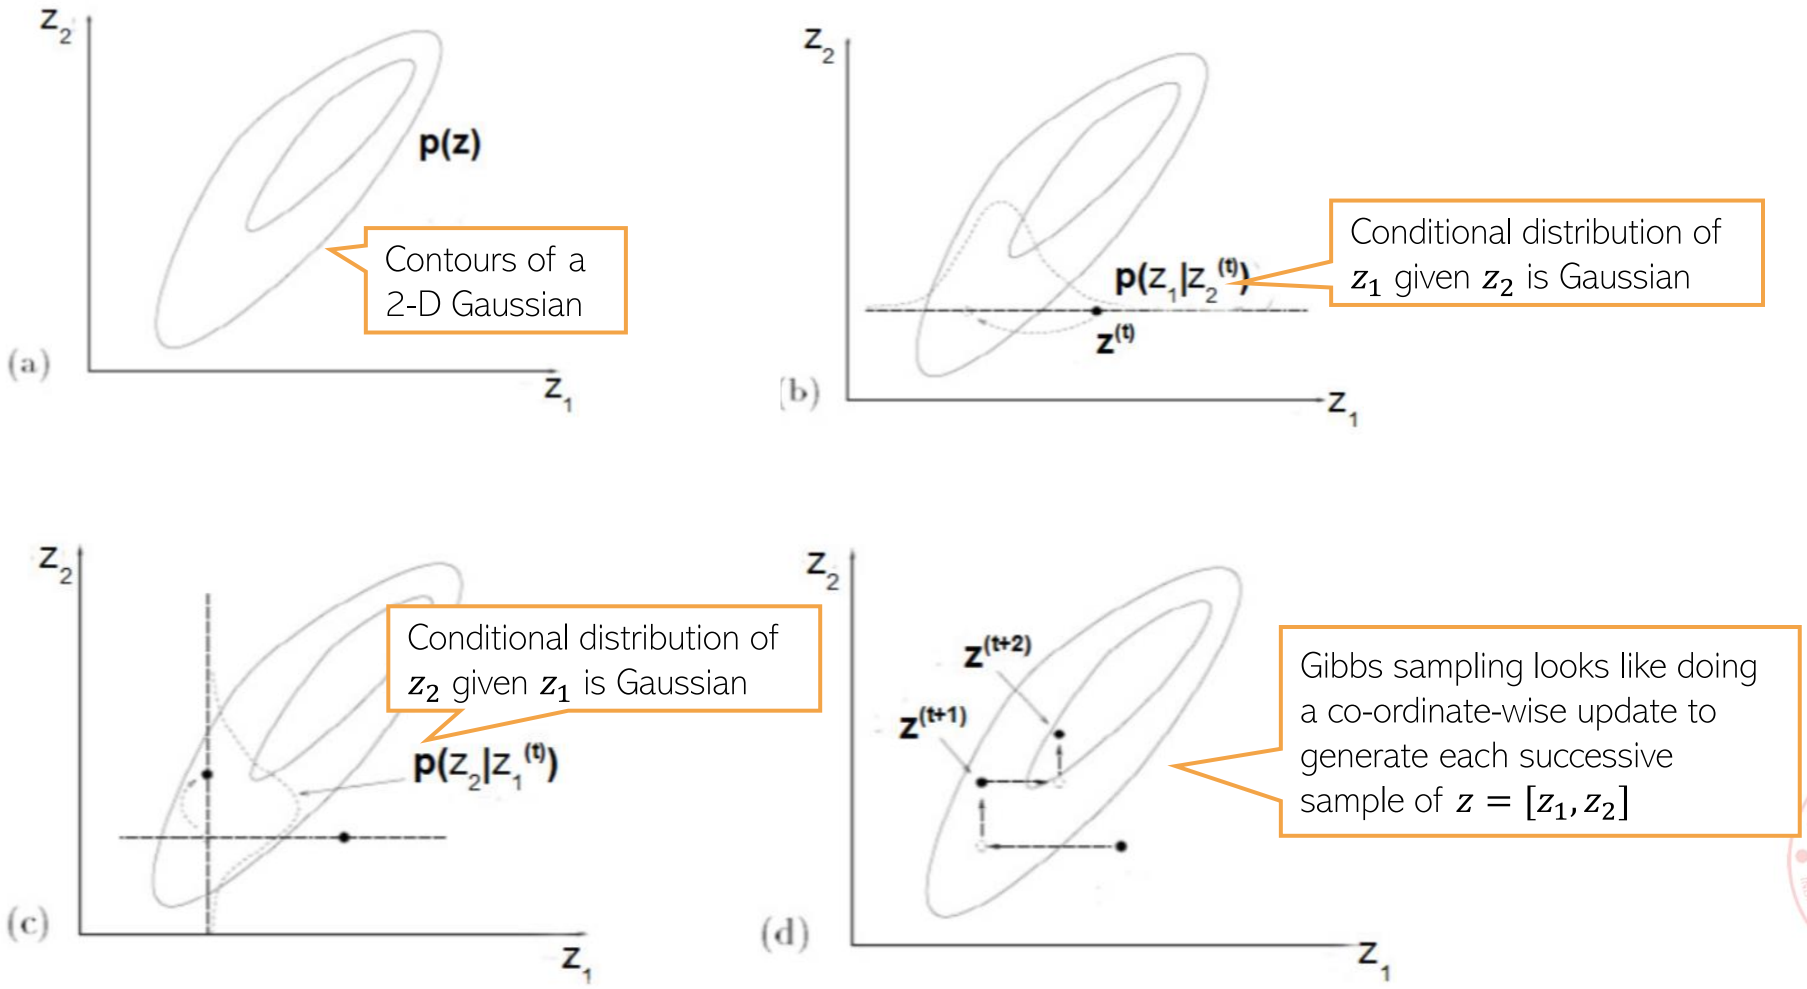

Gibbs Sampling: Another Simple Example

- Bayesian linear regression: $p(y_n|\boldsymbol{x}_n, \boldsymbol{w}, \beta) = \mathcal{N}(y_n|\boldsymbol{w}^\top \boldsymbol{x}_n, \beta^{-1})$, $p(\boldsymbol{w}) = \mathcal{N}(\boldsymbol{w}|0, \lambda^{-1}I)$, $p(\lambda) = \text{Gamma}(\lambda|a, b)$, $p(\beta) = \text{Gamma}(\beta|c, d)$. Gibbs sampler for $p(w, \lambda, \beta|\boldsymbol{X}, \boldsymbol{y})$ will be
- Initialize λ, β as $\lambda^{(0)}, \beta^{(0)}$. For iteration $t = 1,2, \dots, T$
  - Generate a random sample of $\boldsymbol{w}$ by sampling from its CP as

$$\boldsymbol{w}^{(t)} \sim \mathcal{N}(\boldsymbol{w}|\boldsymbol{\mu}^{(t-1)}, \boldsymbol{\Sigma}^{(t-1)}) \quad \text{where} \quad \boldsymbol{\Sigma}^{(t-1)} = \left(\beta^{(t-1)}\boldsymbol{X}^\top\boldsymbol{X} + \lambda^{(t-1)}\right)^{-1}, \quad \boldsymbol{\mu}^{(t-1)} = \left(\boldsymbol{X}^\top\boldsymbol{X} + \frac{\lambda^{(t-1)}}{\beta^{(t-1)}}\right)^{-1}\boldsymbol{X}^\top\boldsymbol{y}$$

  - Generate a random sample of λ by sampling from its CP as

$$\lambda^{(t)} \sim \text{Gamma}\left(\lambda|a + \frac{D}{2}, b + \frac{\boldsymbol{w}^{(t)^\top}\boldsymbol{w}^{(t)}}{2}\right)$$

  - Generate a random sample of β by sampling from its CP as

$$\beta^{(t)} \sim \text{Gamma}\left(\beta|c + \frac{N}{2}, d + \frac{\|\boldsymbol{y} - \boldsymbol{X}\boldsymbol{w}^{(t)}\|^2}{2}\right)$$

> Note: Assuming these are post-burnin samples and thinning (if any) is also considered

- The posterior's approximation is the set of collected samples $\{\boldsymbol{w}^{(t)}, \lambda^{(t)}, \beta^{(t)}\}_{t=1}^{T}$