


CS330 Review Session: Bayesian Meta-Learning



Why Bayesian Meta-Learning?




- Deterministic methods learn a point estimate (e.g. one classifier).
- Bayesian methods show multiple hypotheses.
Useful for:
 - safety-critical settings
 - active learning
 - exploration

What We'll Cover Today

- 1. Amortized Variational Inference**
2. ELBO Derivation for Black-Box Meta-Learning


Out of scope: implementation, MAML-based methods

How to Read Plate Notation




- We see N datapoints generated through the same process
- The parameters outside the plate are shared
- x is observed, z is unobserved

How to Read Plate Notation



- We see N datapoints generated through the same process
- The parameters outside the plate are shared
- x is observed, z is unobserved
- We sample x as $p_{\theta}(z)p_{\theta}(x|z)$ — often just $p(z)$
- Given x , we infer z as $q_{\phi}(z|x)$

Evidence Lower Bound (ELBO)




$$\begin{aligned}\log p(x) &\geq \mathbb{E}_{q(z|x)} [\log p(x, z)] + \mathcal{H}(q(z|x)) \\ &= \mathbb{E}_{q(z|x)} [\log p(x|z)] - D_{KL}(q(z|x) \| p(z))\end{aligned}$$

What We'll Cover Today

1. Amortized Variational Inference
- 2. ELBO Derivation for Black-Box Meta-Learning**

Out of scope: implementation, MAML-based methods

A Black-box Bayesian Model




- Two plates: T tasks, N datapoints per task
- Global parameters θ model task parameters ϕ
 - ϕ = model parameters, model inputs...
- Shorthand: $X = (x_1, \dots, x_N)$, $Y = (y_1, \dots, y_N)$
- For task i , labels predicted as $p(y \mid x, \phi_i)$.

To make predictions on a new task:

(1) infer $q(\phi \mid X, Y)$ (2) predict $p(y \mid x^{new}, \phi)$


Evidence Lower Bound (ELBO)



(whiteboard)

Parameterization

Bayesian black-box meta-learning
with standard, deep variational inference



$$\max_{\theta} \mathbb{E}_{\mathcal{T}_i} \left[\mathbb{E}_{q(\phi_i | \mathcal{D}_i^{\text{tr}}, \theta)} \left[\log p(y_i^{\text{ts}} | x_i^{\text{ts}}, \phi_i) \right] - D_{KL} \left(q(\phi_i | \mathcal{D}_i^{\text{tr}}, \theta) \| p(\phi_i | \theta) \right) \right]$$

Pros:

- + can represent non-Gaussian distributions over y^{ts}
- + produces distribution over functions

Cons:

- Can only represent Gaussian distributions $p(\phi_i | \theta)$
(okay when ϕ_i is latent vector)

- ϕ = network weights
 - “Hypernetwork”
 - Learned prior $p(\phi | \theta)$ is important
- ϕ = inputs to a network
 - Meaning of ϕ is entirely learned
 - Simple prior $p(\phi)$ suffices

What We Covered Today



1. Amortized Variational Inference
2. ELBO Derivation for Black-Box Meta-Learning

Out of scope: implementation, MAML-based methods