Variational Inference and Generative Models

CS 285

Instructor: Sergey Levine UC Berkeley



Today's Lecture

- 1. Probabilistic latent variable models
- 2. Variational inference
- 3. Amortized variational inference
- 4. Generative models: variational autoencoders
- Goals
 - Understand latent variable models in deep learning
 - Understand how to use (amortized) variational inference

Probabilistic models

p(x)









Latent variable models in general



Latent variable models in RL

conditional latent variable models for multi-modal policies



latent variable models for model-based RL



 $p(o_t|x_t)$ actually models $p(x_{t+1}|x_t)$ and $p(x_1)$

 $p(x_t)$ latent space has *structure*

Other places we'll see latent variable models

Using RL/control + variational inference to model human behavior



Using generative models and variational inference for exploration



How do we train latent variable models?

the model: $p_{\theta}(x)$

the data:
$$\mathcal{D} = \{x_1, x_2, x_3, \dots, x_N\}$$

maximum likelihood fit:

$$\theta \leftarrow \arg\max_{\theta} \frac{1}{N} \sum_{i} \log p_{\theta}(x_i)$$

$$p(x) = \int p(x|z)p(z)dz$$

$$\theta \leftarrow \arg \max_{\theta} \frac{1}{N} \sum_{i} \log \left(\int p_{\theta}(x_i|z) p(z) dz \right)$$

completely intractable

Estimating the log-likelihood

alternative: *expected* log-likelihood:

 $\theta \leftarrow \arg\max_{\theta} \frac{1}{N} \sum_{i} E_{z \sim p(z|x_i)}[\log p_{\theta}(x_i, z)]$

but... how do we calculate $p(z|x_i)$?

intuition: "guess" most likely z given x_i , and pretend it's the right one

...but there are many possible values of z so use the distribution $p(z|x_i)$



Variational Inference

The variational approximation

but... how do we calculate $p(z|x_i)$?

what if we approximate with $q_i(z) = \mathcal{N}(\mu_i, \sigma_i)$

can bound $\log p(x_i)!$

$$\log p(x_i) = \log \int_z p(x_i|z)p(z)$$
$$= \log \int_z p(x_i|z)p(z)\frac{q_i(z)}{q_i(z)}$$
$$= \log E_{z \sim q_i(z)} \left[\frac{p(x_i|z)p(z)}{q_i(z)}\right]$$



The variational approximation

but... how do we calculate $p(z|x_i)$?

can bound $\log p(x_i)!$

 $\log p(x_i) = \log \int_z p(x_i|z)p(z)$ $= \log \int_z p(x_i|z)p(z)\frac{q_i(z)}{q_i(z)}$ $= \log E_{z \sim q_i(z)} \begin{bmatrix} \frac{p(x_i|z)p(z)}{q_i(z)} \end{bmatrix}$ $\geq E_{z \sim q_i(z)} \begin{bmatrix} \log \frac{p(x_i|z)p(z)}{q_i(z)} \end{bmatrix} = E_{z \sim q_i(z)}[\log p(x_i|z) + \log p(z)] + \mathcal{H}_x(q_{ij})(z)[\log q_i(z)]$

Jensen's inequality

 $\log E[y] \ge E[\log y]$

A brief aside...

Entropy:

$$\mathcal{H}(p) = -E_{x \sim p(x)}[\log p(x)] = -\int_{x} p(x)\log p(x)dx$$

Intuition 1: how random is the random variable?

Intuition 2: how large is the log probability in expectation *under itself*

what do we expect this to do? $E_{z \sim q_i(z)}[\log p(x_i|z) + \log p(z)] + \mathcal{H}(q_i)$







A brief aside...

KL-Divergence:

$$D_{\mathrm{KL}}(q||p) = E_{x \sim q(x)} \left[\log \frac{q(x)}{p(x)} \right] = E_{x \sim q(x)} [\log q(x)] - E_{x \sim q(x)} [\log p(x)] = -E_{x \sim q(x)} [\log p(x)] - \mathcal{H}(q)$$

Intuition 1: how *different* are two distributions?

Intuition 2: how small is the expected log probability of one distribution under another, minus entropy?

why entropy?



The variational approximation

 $\mathcal{L}_i(p,q_i)$

 $\log p(x_i) \ge E_{z \sim q_i(z)} [\log p(x_i|z) + \log p(z)] + \mathcal{H}(q_i)$

what makes a good $q_i(z)$? approximate in what sense? why?

intuition: $q_i(z)$ should approximate $p(z|x_i)$ compare in terms of KL-divergence: $D_{\text{KL}}(q_i(z)||p(z|x))$

 $D_{\mathrm{KL}}(q_{i}(x_{i})||p(z|x_{i})) = E_{z \sim q_{i}(z)} \left[\log \frac{q_{i}(z)}{p(z|x_{i})} \right] = E_{z \sim q_{i}(z)} \left[\log \frac{q_{i}(z)p(x_{i})}{p(x_{i},z)} \right]$ $= -E_{z \sim q_{i}(z)} [\log p(x_{i}|z) + \log p(z)] + E_{z \sim q_{i}(z)} [\log q_{i}(z)] + E_{z \sim q_{i}(z)} [\log p(x_{i})]$ $= -E_{z \sim q_{i}(z)} [\log p(x_{i}|z) + \log p(z)] - \mathcal{H}(q_{i}) + \log p(x_{i})$ $= -\mathcal{L}_{i}(p,q_{i}) + \log p(x_{i})$ $\log p(x_{i}) = D_{\mathrm{KL}}(q_{i}(z)||p(z|x_{i})) + \mathcal{L}_{i}(p,q_{i})$ $\log p(x_{i}) \geq \mathcal{L}_{i}(p,q_{i})$

The variational approximation

 $\mathcal{L}_i(p,q_i)$

 $\log p(x_i) \ge E_{z \sim q_i(z)}[\log p(x_i|z) + \log p(z)] + \mathcal{H}(q_i)$

 $\log p(x_i) = D_{\mathrm{KL}}(q_i(z) || p(z|x_i)) + \mathcal{L}_i(p, q_i)$ $\log p(x_i) \ge \mathcal{L}_i(p, q_i)$ $D_{\mathrm{KL}}(q_i(z) || p(z|x_i)) = E_{z \sim q_i(z)} \left[\log \frac{q_i(z)}{p(z|x_i)} \right] = E_{z \sim q_i(z)} \left[\log \frac{q_i(z)p(x_i)}{p(x_i, z)} \right]$ $= -E_{z \sim q_i(z)} [\log p(x_i|z) + \log p(z)] - \mathcal{H}(q_i) + \log p(x_i)$ $-\mathcal{L}_i(p, q_i) \qquad \text{independent of } q_i!$

 \Rightarrow maximizing $\mathcal{L}_i(p, q_i)$ w.r.t. q_i minimizes KL-divergence!

How do we use this?

$$\mathcal{L}_i(p, q_i)$$

$$\log p(x_i) \ge E_{z \sim q_i(z)}[\log p_\theta(x_i|z) + \log p(z)] + \mathcal{H}(q_i)$$

$$\theta \leftarrow \arg \max_{\theta} \frac{1}{N} \sum_{i} \log p_{\theta}(x_i)$$

$$\theta \leftarrow \arg\max_{\theta} \frac{1}{N} \sum_{i} \mathcal{L}_i(p, q_i)$$

for each x_i (or mini-batch): calculate $\nabla_{\theta} \mathcal{L}_i(p, q_i)$: sample $z \sim q_i(z)$ $\nabla_{\theta} \mathcal{L}_i(p, q_i) \approx \nabla_{\theta} \log p_{\theta}(x_i | z)$ $\theta \leftarrow \theta + \alpha \nabla_{\theta} \mathcal{L}_i(p, q_i)$ how? update q_i to maximize $\mathcal{L}_i(p, q_i)$ let's say $q_i(z) = \mathcal{N}(\mu_i, \sigma_i)$ use gradient $\nabla_{\mu_i} \mathcal{L}_i(p, q_i)$ and $\nabla_{\sigma_i} \mathcal{L}_i(p, q_i)$ gradient ascent on μ_i, σ_i

What's the problem?

for each x_i (or mini-batch):

calculate $\nabla_{\theta} \mathcal{L}_i(p, q_i)$: sample $z \sim q_i(z)$ $\nabla_{\theta} \mathcal{L}_i(p, q_i) \approx \nabla_{\theta} \log p_{\theta}(x_i | z)$

 $\theta \leftarrow \theta + \alpha \nabla_{\theta} \mathcal{L}_i(p, q_i)$

update q_i to maximize $\mathcal{L}_i(p, q_i)$

let's say $q_i(z) = \mathcal{N}(\mu_i, \sigma_i)$ use gradient $\nabla_{\mu_i} \mathcal{L}_i(p, q_i)$ and $\nabla_{\sigma_i} \mathcal{L}_i(p, q_i)$ gradient ascent on μ_i, σ_i

How many parameters are there? $|\theta| + (|\mu_i| + |\sigma_i|) \times N$ intuition: $q_i(z)$ should approximate $p(z|x_i)$ what if we learn a network $q_i(z) = q(z|x_i) \approx p(z|x_i)$?



Amortized Variational Inference

What's the problem?

for each x_i (or mini-batch):

calculate $\nabla_{\theta} \mathcal{L}_i(p, q_i)$: sample $z \sim q_i(z)$ $\nabla_{\theta} \mathcal{L}_i(p, q_i) \approx \nabla_{\theta} \log p_{\theta}(x_i | z)$

 $\theta \leftarrow \theta + \alpha \nabla_{\theta} \mathcal{L}_i(p, q_i)$

update q_i to maximize $\mathcal{L}_i(p, q_i)$

let's say $q_i(z) = \mathcal{N}(\mu_i, \sigma_i)$ use gradient $\nabla_{\mu_i} \mathcal{L}_i(p, q_i)$ and $\nabla_{\sigma_i} \mathcal{L}_i(p, q_i)$ gradient ascent on μ_i, σ_i

How many parameters are there? $|\theta| + (|\mu_i| + |\sigma_i|) \times N$ intuition: $q_i(z)$ should approximate $p(z|x_i)$ what if we learn a network $q_i(z) = q(z|x_i) \approx p(z|x_i)$?



Amortized variational inference



for each x_i (or mini-batch): calculate $\nabla_{\theta} \mathcal{L}(p_{\theta}(x_i|z), q_{\phi}(z|x_i))$: sample $z \sim q_{\phi}(z|x_i)$ $\nabla_{\theta} \mathcal{L} \approx \nabla_{\theta} \log p_{\theta}(x_i|z)$ $\theta \leftarrow \theta + \alpha \nabla_{\theta} \mathcal{L}$ $\phi \leftarrow \phi + \alpha \nabla_{\phi} \mathcal{L}$ how do we calculate this?

Amortized variational inference

for each x_i (or mini-batch):

calculate $\nabla_{\theta} \mathcal{L}(p_{\theta}(x_{i}|z), q_{\phi}(z|x_{i}))$: sample $z \sim q_{\phi}(z|x_{i})$ $\nabla_{\theta} \mathcal{L} \approx \nabla_{\theta} \log p_{\theta}(x_{i}|z)$ $\theta \leftarrow \theta + \alpha \nabla_{\theta} \mathcal{L}$ $\phi \leftarrow \phi + \alpha \nabla_{\phi} \mathcal{L}$ $\int \mathcal{L}_{i} = E_{z \sim q_{\phi}(z|x_{i})}[\log p_{\theta}(x_{i}|z) + \log p(z)] + \mathcal{H}(q_{\phi}(z|x_{i}))$ $J(\phi) = E_{z \sim q_{\phi}(z|x_{i})}[r(x_{i}, z)]$

can just use policy gradient!

What's wrong with this gradient?

$$\nabla J(\phi) \approx \frac{1}{M} \sum_{j} \nabla_{\phi} \log q_{\phi}(z_j | x_i) r(x_i, z_j)$$

The reparameterization trick

Is there a better way?

$$J(\phi) = E_{z \sim q_{\phi}(z|x_{i})}[r(x_{i}, z)] \qquad q_{\phi}(z|x) = \mathcal{N}(\mu_{\phi}(x), \sigma_{\phi}(x))$$

$$= E_{\epsilon \sim \mathcal{N}(0,1)}[r(x_{i}, \mu_{\phi}(x_{i}) + \epsilon \sigma_{\phi}(x_{i}))] \qquad z = \mu_{\phi}(x) + \epsilon \sigma_{\phi}(x)$$
estimating $\nabla_{\phi} J(\phi)$:
sample $\epsilon_{1}, \dots, \epsilon_{M}$ from $\mathcal{N}(0, 1)$ (a single sample works well!) $\epsilon \sim \mathcal{N}(0, 1)$
 $\nabla_{\phi} J(\phi) \approx \frac{1}{M} \sum_{j} \nabla_{\phi} r(x_{i}, \mu_{\phi}(x_{i}) + \epsilon_{j} \sigma_{\phi}(x_{i}))$ independent of ϕ !

most autodiff software (e.g., TensorFlow) will compute this for you!

Another way to look at it...

$$\begin{aligned} \mathcal{L}_{i} &= E_{z \sim q_{\phi}(z|x_{i})}[\log p_{\theta}(x_{i}|z) + \log p(z)] + \mathcal{H}(q_{\phi}(z|x_{i})) \\ &= E_{z \sim q_{\phi}(z|x_{i})}[\log p_{\theta}(x_{i}|z)] + \underbrace{E_{z \sim q_{\phi}(z|x_{i})}[\log p(z)] + \mathcal{H}(q_{\phi}(z|x_{i}))]}_{-D_{\mathrm{KL}}(q_{\phi}(z|x_{i})||p(z))} & \longleftarrow \text{ this often has a convenient analytical form (e.g., KL-divergence for Gaussians)} \\ &= E_{z \sim q_{\phi}(z|x_{i})}[\log p_{\theta}(x_{i}|z)] - D_{\mathrm{KL}}(q_{\phi}(z|x_{i})||p(z)) \\ &= E_{\epsilon \sim \mathcal{N}(0,1)}[\log p_{\theta}(x_{i}|\mu_{\phi}(x_{i}) + \epsilon \sigma_{\phi}(x_{i}))] - D_{\mathrm{KL}}(q_{\phi}(z|x_{i})||p(z)) \\ &\approx \log p_{\theta}(x_{i}|\mu_{\phi}(x_{i}) + \epsilon \sigma_{\phi}(x_{i})) - D_{\mathrm{KL}}(q_{\phi}(z|x_{i})||p(z)) \\ x_{i} & \longleftarrow \begin{array}{c} \mu_{\phi}(x_{i}) & \longleftarrow \\ \phi & \phi(x_{i}) & \oplus \\$$

Reparameterization trick vs. policy gradient

Policy gradient

- Can handle both discrete and continuous latent variables
- High variance, requires multiple samples & small learning rates
- Reparameterization trick
 - Only continuous latent variables
 - Very simple to implement
 - Low variance

$$J(\phi) \approx \frac{1}{M} \sum_{j} \nabla_{\phi} \log q_{\phi}(z_j | x_i) r(x_i, z_j)$$

$$\nabla_{\phi} J(\phi) \approx \frac{1}{M} \sum_{j} \nabla_{\phi} r(x_i, \mu_{\phi}(x_i) + \epsilon_j \sigma_{\phi}(x_i))$$

Variational Inference in Deep RL



Using the variational autoencoder

$$q_{\phi}(z|x) = \mathcal{N}(\mu_{\phi}(x), \sigma_{\phi}(x)) \quad \phi \stackrel{(\mathbf{z})}{\longleftarrow} \quad p_{\theta}(x|z) = \mathcal{N}(\mu_{\theta}(z), \sigma_{\theta}(z))$$

$$p(x) = \int p(x|z)p(z)dz$$

why does this work?

sampling: $z \sim p(z)$

 $x \sim p(x|z)$

$$\mathcal{L}_i = E_{z \sim q_\phi(z|x_i)}[\log p_\theta(x_i|z)] - D_{\mathrm{KL}}(q_\phi(z|x_i)||p(z))$$



 $p_{\theta}(x|z)$

 \boldsymbol{z}

Example applications

Representation learning

 \boldsymbol{z} is a representation of \boldsymbol{s}





1. Train VAE on states in replay buffer \mathcal{R} 2. Run RL, using z as the state instead of s

Why is this a good idea?



Sample algorithm

- ▶ 1. Collect transition (s, a, s', r), add to \mathcal{R}
 - 2. Update $p_{\theta}(s|z)$ and $q_{\phi}(z|s) \le m/\beta$ batch from \mathcal{R}
 - **3**. Update Q(z, a) w/ batch from \mathcal{R}

This also provides a great way to use prior data!

Conditional models

$$\mathcal{L}_{i} = E_{z \sim q_{\phi}(z|x_{i}, y_{i})} [\log p_{\theta}(y_{i}|x_{i}, z) + \log p(z|x_{i})] + \mathcal{H}(q_{\phi}(z|x_{i}, y_{i}))$$

just like before, only now generating y_i and everything is conditioned on x_i

at test time:



can optionally depend on x

p(z)





Example applications

Multimodal imitation learning

p(a|s,z)(TACOCODDD 000000 XXXXXXXXXX (00000)XXXXXXXXXX 00000000 $z \sim \mathcal{N}(0, \mathbf{I})$ p(z)

Image: state stat

Learning Latent Plans from Play

COREY LYNCH	MOHI KHANSARI	TED XIAO	VIKASH KUMAR	JONATHAN TOMPSON	SERGEY LEVINE	PIERRE SERMANET
Google Brain	Google X	Google Brain	Google Brain	Google Brain	Google Brain	Google Brain



Example applications

Multimodal imitation learning



Learning Fine-Grained Bimanual Manipulation with Low-Cost Hardware

Tony Z. Zhao¹ Vikash Kumar³ Sergey Levine² Chelsea Finn¹ ¹ Stanford University ² UC Berkeley ³ Meta







State space models





Example applications

Representation learning and model-based RL

1. Learn state space model and plan in the latent space

Embed to Control: A Locally Linear Latent Dynamics Model for Control from Raw Images

Manuel Watter*Jost Tobias Springenberg*Martin RiedmillerJoschka BoedeckerGoogle DeepMindUniversity of Freiburg, GermanyLondon, UK{watterm, springj, jboedeck}@cs.uni-freiburg.deriedmiller@google.com

AE VE VE VE STANDOS

Swing-up with the E2C algorithm

SOLAR: Deep Structured Latent Representations for Model-Based Reinforcement Learning



Learning Latent Dynamics for Planning from Pixels

Danijar Hafner¹² Timothy Lillicrap³ Ian Fischer⁴ Ruben Villegas¹⁵ David Ha¹ Honglak Lee¹ James Davidson¹









(c) Cheetah

(d) Finger

(e) Cup

(f) Walker

Example applications

Representation learning and model-based RL

1. Learn state space model and run RL in the state space

Stochastic Latent Actor-Critic: Deep Reinforcement Learning with a Latent Variable Model

Alex X. Lee^{1,2} Anusha Nagabandi¹ Pieter Abbeel¹ Sergey Levine¹ ¹University of California, Berkeley ²DeepMind {alexlee_gk,nagaban2,pabbeel,svlevine}@cs.berkeley.edu

true rollouts

samples



DREAM TO CONTROL: LEARNING BEHAVIORS BY LATENT IMAGINATION





