Variational Discriminator Bottleneck: Improving Imitation Learning, Inverse RL, and GANs by Constraining Information Flow

Xue Bin Peng & Angjoo Kanazawa & Sam Toyer & Pieter Abbeel & Sergey Levine

ICLR 2019

Motivation

The issues of GAN:

- Discriminator is prone to overpowering the generator, which will provide the generator with uninformative gradients for improvement

- Hard to train (unstable or poor performance)

Motivation

How to balance the generator and discriminator?

- To regularize the internal representation by using the information bottleneck

- Modulate the accuracy of discriminator and maintain useful and informative gradient

Preliminaries

Deep Variational information bottleneck

- Given a dataset {x,y} with features x and labels y, the standard maximum likelihood estimate q(y|x) can be determined

$$\min_{q} \quad \mathbb{E}_{\mathbf{x},\mathbf{y}\sim p(\mathbf{x},\mathbf{y})}\left[-\log q(\mathbf{y}|\mathbf{x})\right]$$

- Prone to overfitting

Preliminaries

- The bottleneck can be incorporated by first introducing an encoder that maps the features x to a latent distribution over Z. E(z|x) here is the distribution
- And then enforcing an upper bound **I_c** on the mutual information between the encoding and the original features I(X,Z). Then the objective:

$$\begin{aligned} J(q, E) = \min_{q, E} \quad & \mathbb{E}_{\mathbf{x},\mathbf{y} \sim p(\mathbf{x},\mathbf{y})} \left[\mathbb{E}_{\mathbf{z} \sim E(\mathbf{z}|\mathbf{x})} \left[-\log q(\mathbf{y}|\mathbf{z}) \right] \right] \\ \text{s.t.} \quad & I(X, Z) \leq I_c . \end{aligned}$$

Preliminaries

- Note that the model q(y|z) now maps samples from the latent distribution z to the label y. The mutual information is defined according to

$$I(X,Z) = \int p(\mathbf{x},\mathbf{z}) \log \frac{p(\mathbf{x},\mathbf{z})}{p(\mathbf{x})p(\mathbf{z})} \, d\mathbf{x} \, d\mathbf{z} \; = \int p(\mathbf{x})E(\mathbf{z}|\mathbf{x}) \log \frac{E(\mathbf{z}|\mathbf{x})}{p(\mathbf{z})} \, d\mathbf{x} \, d\mathbf{z}$$

- A variational lower bound can be obtained by using an approximation r(z) of the marginal p(z)

$$I(X,Z) \leq \int p(\mathbf{x})E(\mathbf{z}|\mathbf{x}) \log \frac{E(\mathbf{z}|\mathbf{x})}{r(\mathbf{z})} \, d\mathbf{x} \, d\mathbf{z} = \mathbb{E}_{\mathbf{x} \sim p(\mathbf{x})} \left[\mathrm{KL} \left[E(\mathbf{z}|\mathbf{x}) || r(\mathbf{z}) \right] \right].$$

Preliminaries

- A variational lower bound = An upper bound on I(X|Z)

= An upper bound on the regularized objective:

$$\tilde{J}(q,E) = \min_{q,E} \quad \mathbb{E}_{\mathbf{x},\mathbf{y}\sim p(\mathbf{x},\mathbf{y})}\left[\mathbb{E}_{\mathbf{z}\sim E(\mathbf{z}|\mathbf{x})}\left[-\log q(\mathbf{y}|\mathbf{z})\right]\right]$$
$$\text{s.t.} \quad \mathbb{E}_{\mathbf{x}\sim p(\mathbf{x})}\left[\mathrm{KL}\left[E(\mathbf{z}|\mathbf{x})||r(\mathbf{z})\right]\right] \leq I_c.$$

- This can be solved by introducing the Lagrange multiplier

$$\min_{q,E} \quad \mathbb{E}_{\mathbf{x},\mathbf{y}\sim p(\mathbf{x},\mathbf{y})}\left[\mathbb{E}_{\mathbf{z}\sim E(\mathbf{z}|\mathbf{x})}\left[-\log q(\mathbf{y}|\mathbf{z})\right]\right] + \beta\left(\mathbb{E}_{\mathbf{x}\sim p(\mathbf{x})}\left[\mathrm{KL}\left[E(\mathbf{z}|\mathbf{x})||r(\mathbf{z})\right]\right] - I_c\right).$$

VDB

- Standard GAN:

$$\max_G \min_D \quad \mathbb{E}_{\mathbf{x}\sim p^*(\mathbf{x})}\left[-\log\left(D(\mathbf{x})\right)\right] + \mathbb{E}_{\mathbf{x}\sim G(\mathbf{x})}\left[-\log\left(1 - D(\mathbf{x})\right)\right]$$

- Introduce an encoder E into discriminator that maps a sample x into a stochastic encoding z, and then apply the information bottleneck $\mathbf{I_c}$ on the mutual information I(X,Z)

$$J(D,E) = \min_{D,E} \quad \mathbb{E}_{x\sim p^*(\mathbf{x})}\left[\mathbb{E}_{\mathbf{z}\sim E(\mathbf{z}|\mathbf{x})}\left[-\log\left(D(\mathbf{z})\right)\right]\right] + \mathbb{E}_{\mathbf{x}\sim G(\mathbf{x})}\left[\mathbb{E}_{\mathbf{z}\sim E(\mathbf{z}|\mathbf{x})}\left[-\log\left(1 - D(\mathbf{z})\right)\right]\right]$$

$$\text{s.t.} \quad \mathbb{E}_{\mathbf{x}\sim \tilde{p}(\mathbf{x})}\left[\mathrm{KL}\left[E(\mathbf{z}|\mathbf{x}) \| r(\mathbf{z})\right]\right] \leq I_c,$$

$$\tilde{p} = \frac{1}{2}p^* + \frac{1}{2}G$$

VDB

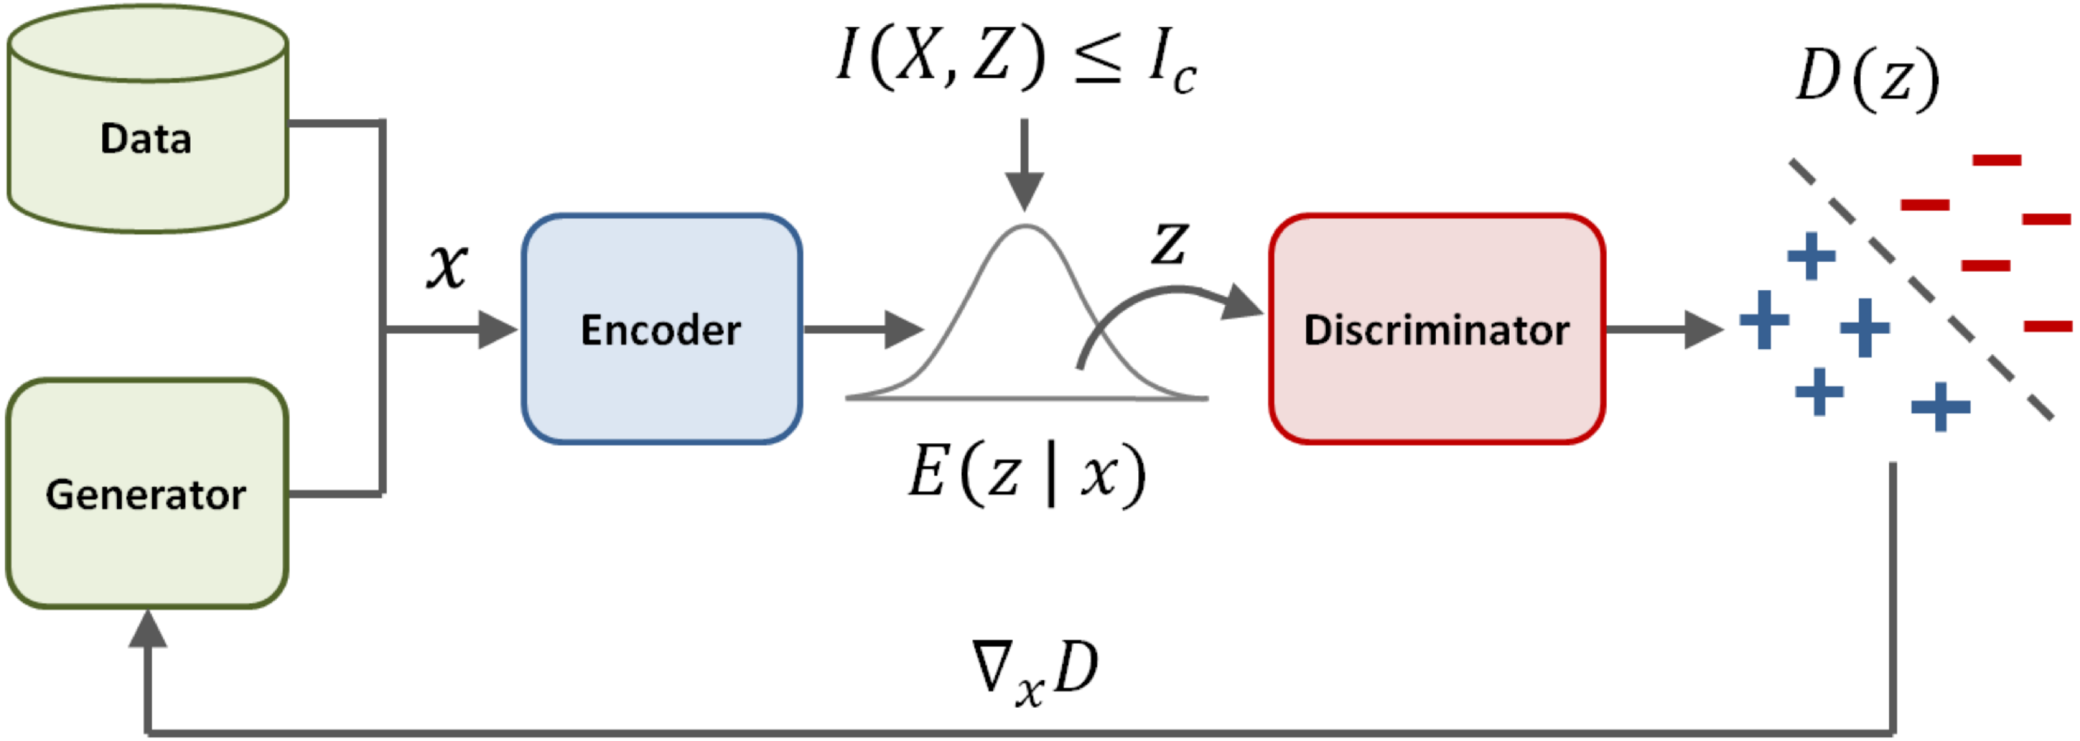

VDB

- Objective (Discriminator) with dual gradient descent:

$$J(D, E) = \min_{D,E} \max_{\beta \geq 0} \quad \mathbb{E}_{\mathbf{x} \sim p^*(\mathbf{x})} \left[\mathbb{E}_{\mathbf{z} \sim E(\mathbf{z}|\mathbf{x})} \left[-\log \left(D(\mathbf{z}) \right) \right] \right] + \mathbb{E}_{\mathbf{x} \sim G(\mathbf{x})} \left[\mathbb{E}_{\mathbf{z} \sim E(\mathbf{z}|\mathbf{x})} \left[-\log \left(1 - D(\mathbf{z}) \right) \right] \right]$$
$$+ \beta \left(\mathbb{E}_{\mathbf{x} \sim \tilde{p}(\mathbf{x})} \left[\mathrm{KL} \left[E(\mathbf{z}|\mathbf{x}) || r(\mathbf{z}) \right] \right] - I_c \right).$$

$$D, E \leftarrow \arg\min_{D,E} \mathcal{L}(D, E, \beta)$$
$$\beta \leftarrow \max \left(0, \ \beta + \alpha_\beta \left(\mathbb{E}_{\mathbf{x} \sim \tilde{p}(\mathbf{x})} \left[\mathrm{KL} \left[E(\mathbf{z}|\mathbf{x}) || r(\mathbf{z}) \right] \right] - I_c \right) \right)$$

VDB

- Objective (Generator) with approximation, directly use the mean:

$$\max_{G} \quad \mathbb{E}_{\mathbf{x} \sim G(\mathbf{x})} \left[-\log\left(1 - D(\mu_E(\mathbf{x}))\right) \right]$$

Experiments

- video

Conclusion

- Prevent overfitting in GAN/VAE is crucial
- The VDB framework provides a way to achieve better representation
- How to select the I_c ?
- Any potential way to combine pre-trained representation with VDB?