

---

# Adversarial Learning of a Sampler Based on an Unnormalized Distribution

---

Chunyu Li<sup>1</sup> Ke Bai<sup>2</sup> Jianqiao Li<sup>2</sup> Guoyin Wang<sup>2</sup> Changyou Chen<sup>3</sup> Lawrence Carin<sup>2</sup>  
<sup>1</sup>Microsoft Research, Redmond      <sup>2</sup>Duke University      <sup>3</sup>University at Buffalo

## Abstract

We investigate adversarial learning in the case when only an unnormalized form of the density can be accessed, rather than samples. With insights so garnered, adversarial learning is extended to the case for which one has access to an unnormalized form  $u(x)$  of the target density function, but no samples. Further, new concepts in GAN regularization are developed, based on learning from samples or from  $u(x)$ . The proposed method is compared to alternative approaches, with encouraging results demonstrated across a range of applications, including deep soft Q-learning.

## 1 Introduction

Significant progress has been made recently on generative models capable of synthesizing highly realistic data samples [Goodfellow et al., 2014, Oord et al., 2016, Kingma and Welling, 2014]. If  $p(x)$  represents the true underlying probability distribution of data  $x \in \mathcal{X}$ , most of these models seek to represent draws  $x \sim p(x)$  as  $x = h_\theta(\epsilon)$  and  $\epsilon \sim q_0$ , with  $q_0$  a specified distribution that may be sampled easily [Goodfellow et al., 2014, Radford et al., 2016]. The objective is to learn  $h_\theta(\epsilon)$ , modeled typically via a deep neural network. Note that the model doesn't impose a form on (or attempt to explicitly model) the density function  $q_\theta(x)$  used to implicitly model  $p(x)$ .

When learning  $h_\theta(\epsilon)$  it is typically assumed that one has access to a set of samples  $\{x_i\}_{i=1,N}$ , with each  $x_i$  drawn i.i.d. from  $p(x)$ . While such samples are often available, there are other important settings for which one may wish to learn a generative model for  $p(x)$ , without access to associated

samples. An important example occurs when one has access to an *unnormalized* distribution  $u(x)$ , with  $p(x) = u(x)/C$  and normalizing constant  $C$  unknown. The goal of sampling from  $p(x)$  based on  $u(x)$  is a classic problem in physics, statistics and machine learning [Hastings, 1970, Gelman et al., 1995]. This objective has motivated theoretically exact (but expensive) methods like Markov chain Monte Carlo (MCMC) [Brooks et al., 2011, Welling and Teh, 2011], and approximate methods like variational Bayes [Hoffman et al., 2013, Kingma and Welling, 2014, Rezende et al., 2014] and expectation propagation [Minka, 2001, Li et al., 2015]. A challenge with methods of these types (in addition to computational cost/approximations) is that they are means of drawing samples or approximating density forms based on  $u(x)$ , but they do not directly yield a model like  $x = h_\theta(\epsilon)$  and  $\epsilon \sim q_0$ , with the latter important for many fast machine learning implementations.

A recently developed, and elegant, means of modeling samples based on  $u(x)$  is Stein variational gradient descent (SVGD) [Liu and Wang, 2016]. SVGD also learns to draw a set of samples, and an amortization step is used to learn  $x = h_\theta(\epsilon)$  and  $\epsilon \sim q_0$  based on the SVGD-learned samples [Wang and Liu, 2016, Feng et al., 2017, Y. Pu and Carin, 2017]. Such amortization may also be used to build  $h_\theta(\epsilon)$  based on MCMC-generated samples [Li et al., 2017b]. While effective, SVGD-based learning of this form may be limited computationally by the number of samples that may be practically modeled, limiting accuracy. Further, the two-step nature by which  $x = h_\theta(\epsilon)$  is manifested may be viewed as less appealing.

In this paper we develop a new extension of generative adversarial networks (GANs) [Goodfellow et al., 2014] for settings in which we have access to  $u(x)$ , rather than samples drawn from  $p(x)$ . The formulation, while new, is simple, based on a recognition that many existing GAN methods constitute different means of estimating a function of a likelihood ratio [Kanamori et al., 2010, Mohamed and L., 2016,

Uehara et al., 2016]. The likelihood ratio is associated with the true density function  $p(x)$  and the model  $q_\theta(x)$ . Since we do not have access to  $p(x)$  or  $q_\theta(x)$ , we show, by a detailed investigation of  $f$ -GAN [Nowozin et al., 2016], that many GAN models reduce to learning  $g_0(p(x)/q_\theta(x))$ , where  $g_0(\cdot)$  is a general monotonically increasing function.  $f$ -GAN is an attractive model for uncovering underlying principles associated with GANs, due to its generality, and that many existing GAN approaches may be viewed as special cases of  $f$ -GAN. With the understanding provided by an analysis of  $f$ -GAN, we demonstrate how  $g_0(p(x)/q_\theta(x))$  may be estimated via  $u(x)$ , and an introduced reference distribution  $p_r(x)$ . As discussed below, the assumptions on  $p_r(x)$  are that it is easily sampled, it has a known functional form, and it represents a good approximation to  $q_\theta(x)$ .

For the special case of variational inference for latent models, the proposed formulation recovers the adversarial variational Bayes (AVB) [Mescheder et al., 2017] setup. However, we demonstrate that the proposed approach has more applicability than inference. Specifically, we demonstrate its application to soft Q-learning [Haarnoja et al., 2017], and it leads to the first general purpose adversarial policy algorithm in reinforcement learning. We make a favorable comparison in this context to the aforementioned SVGD formulation.

An additional contribution of this paper concerns regularization of adversarial learning, of interest when learning based on samples or on an unnormalized distribution  $u(x)$ . Specifically, we develop an entropy-based regularizer. When learning based on  $u(x)$ , we make connections to simulated annealing regularization methods used in prior sampling-based models. We also introduce a bound on the entropy, applicable to learning based on samples or  $u(x)$ , and make connections to prior work on cycle consistency used in GAN regularization.

## 2 Traditional GAN Learning

We begin by discussing GAN from the perspective of the  $f$ -divergence [Nguyen et al., 2010a], which has resulted in  $f$ -GAN [Nowozin et al., 2016].  $f$ -GAN is considered because many popular GAN methods result as special cases, thereby affording the opportunity to identify generalizable components that may extended to new settings. Considering continuous probability density functions  $p(x)$  and  $q(x)$  for  $x \in \mathcal{X}$ , the  $f$ -divergence is defined as  $D_f(p||q) = \int_{\mathcal{X}} q(x) f\left[\frac{p(x)}{q(x)}\right] dx$ , where  $f: \mathbb{R}_+ \rightarrow \mathbb{R}$  is a convex, lower-semicontinuous function satisfying  $f(1) = 0$ . Different choices of

$f[r(x)]$ , with  $r(x) = p(x)/q(x)$ , yield many common divergences; see [Nowozin et al., 2016] and Table 1.

An important connection has been made between the  $f$ -divergence and generative adversarial learning, based on the inequality [Nguyen et al., 2010a]

$$D_f(p||q) \geq \sup_{T \in \mathcal{T}} [\mathbb{E}_{x \sim p}[T(x)] - \mathbb{E}_{x \sim q}[f^*(T(x))]] \quad (1)$$

where  $f^*(t)$  is the convex conjugate function, defined as  $f^*(t) = \sup_{u \in \text{dom}_f} \{ut - f(u)\}$ , which has an analytic form for many choices of  $f$  [Nowozin et al., 2016]. Further, under mild conditions, the bound is tight when  $T(x) = f'\left[\frac{p(x)}{q(x)}\right]$  where  $f'(r)$  is the derivative of  $f(r)$ . Even if we know  $f'(r)$  we cannot evaluate  $T(x) = f'\left[\frac{p(x)}{q(x)}\right]$  explicitly, because  $q(x)$  and/or  $p(x)$  are unknown.

Note that to compute the bound in (1), we require expectations wrt  $p$  and  $q$ , which we effect via sampling (this implies we only need samples from  $p$  and  $q$ , and do not require the explicit form of the underlying distributions). Specifically, assume  $p$  corresponds to the true distribution we wish to model, and  $q_\theta$  is a model distribution with parameters  $\theta$ . We seek to learn  $\theta$  by minimizing the bound of  $D_f(p||q_\theta)$  in (1), with draws from  $q_\theta$  implemented as  $x = h_\theta(\epsilon)$  with  $\epsilon \sim q_0$ , where  $q_0$  is a probability distribution that may be sampled easily (*e.g.*, uniform, or isotropic Gaussian [Goodfellow et al., 2014]). The learning problem consists of solving

$$(\hat{\theta}, \hat{\phi}) = \underset{\theta}{\text{argmin}} \underset{\phi}{\text{argmax}} \left[ \mathbb{E}_{x \sim p}[T_\phi(x)] - \mathbb{E}_{\epsilon \sim q_0}[f^*(T_\phi(h_\theta(\epsilon)))] \right] \quad (2)$$

where  $T_\phi(x)$  is typically a (deep) neural network with parameters  $\phi$ , with  $h_\theta(\epsilon)$  defined similarly. Attempting to solve (2) produces  $f$ -GAN [Nowozin et al., 2016].

One typically solves this minimax problem by alternating between update of  $\theta$  and  $\phi$  [Nowozin et al., 2016, Goodfellow et al., 2014]. Note that the update of  $\theta$  only involves the second term in (2), corresponding to  $\underset{\theta}{\text{argmax}} \mathbb{E}_{\epsilon \sim q_0}[f^*(T_\phi(h_\theta(\epsilon)))]$ . Recall that the bound in (1) is tight when  $T_\phi(x) = f'[p(x)/q_{\theta_{n-1}}(x)]$  [Nguyen et al., 2010b], where  $\theta_{n-1}$  represent parameters  $\theta$  from the previous iteration. Hence, assuming  $T_{\phi_n}(x) = f'[p(x)/q_{\theta_{n-1}}(x)]$ , we update  $\theta$  as

$$\theta_n = \underset{\theta}{\text{argmax}} \mathbb{E}_{\epsilon \sim q_0} g[p(h_\theta(\epsilon))/q_{\theta_{n-1}}(h_\theta(\epsilon))] \quad (3)$$

where  $g(r) = f^*(f'(r))$ .

Different choices of  $f$  yield a different optimal function  $g(r)$  (see Table 1). However, in each case  $\theta$  is updated such that samples from  $q_\theta$  yield an increase




Figure 1: Illustration of learning  $q_\theta$  in the two different settings of the target  $p(x)$ . (a) Learning from an unnormalized distribution, as in RAS; (b) Learning from samples, as in the traditional GANs.

Table 1: Functions  $g(r)$  and  $f(r)$  corresponding to particular  $f$ -GAN setups.

$f$ -Divergence	$f(r)$	$g(r)$ in $\theta$ update
Kullback-Leibler (KL)	$r \log r$	$r$
Reverse KL	$-\log r$	$\log r$
Squared Hellinger	$(\sqrt{r} - 1)^2$	$\sqrt{r}$
Total variation	$ r - 1 /2$	$\frac{1}{2} \text{sign}(r - 1)$
Pearson $\chi^2$	$(r - 1)^2$	$(r - 1)^2 + 2r$
Neyman $\chi^2$	$(r - 1)^2/r$	$-1/r$
GAN	$r \log r - (r + 1) \log(r + 1)$	$-\log \left[ \frac{1}{1+r} \right]$

in the likelihood ratio  $r_{\theta_{n-1}}(x) = p(x)/q_{\theta_{n-1}}(x)$ , implying samples from  $q_\theta$  better match  $p(x)$  than they do  $q_{\theta_{n-1}}(x)$ . Recall that the likelihood ratio  $r_{\theta_{n-1}}(x)$  is the optimal means of distinguishing between samples from  $p(x)$  and  $q_{\theta_{n-1}}(x)$  [Van Trees, 2001, Neyman and Pearson, 1933]. Hence,  $r_{\theta_{n-1}}(x)$  is a critic, approximated through  $T_{\phi_n}(x)$ , that the actor  $q_\theta$  seeks to maximize when estimating  $\theta_n$ .

We may alternatively consider

$$\begin{aligned} \phi_n &= \operatorname{argmax}_\phi \{ \mathbb{E}_{x \sim p(x)} \log[\sigma(w_\phi(x))] \\ &\quad + \mathbb{E}_{\epsilon \sim q_0} \log[1 - \sigma(w_\phi(h_{\theta_{n-1}}(\epsilon)))] \} \quad (4) \\ \theta_n &= \operatorname{argmax}_\theta \mathbb{E}_{\epsilon \sim q_0} g_0[w_{\phi_n}(h_\theta(\epsilon))] \quad (5) \end{aligned}$$

where now  $g_0(r)$  is an *arbitrary* monotonically increasing function of  $r$ ,  $\sigma(\cdot)$  is the sigmoid function. From [Kanamori et al., 2010, Mescheder et al., 2017, Gutmann and Hyvärinen, 2010], the solution to (4) is

$$w_\phi(x) = \log[p(x)/q_{\theta_{n-1}}(x)], \quad (6)$$

where model  $w_\phi(x)$  is assumed to have sufficient capacity to represent the likelihood ratio for all  $x \in \mathcal{X}$ . Hence, here  $w_\phi(x)$  replaces  $T_\phi(x)$  from  $f$ -GAN, and the solution to  $w_\phi(x)$  is a particular function of the likelihood ratio. If  $g_0(w_\phi(x)) = w_\phi(x)$  this corresponds to learning based on minimizing the reverse KL divergence  $\text{KL}(q_\theta \| p)$ . When  $g_0(\cdot) = \log[\sigma(\cdot)]$ , one recovers the original GAN [Goodfellow et al., 2014], for which learning corresponds to  $(\hat{\theta}, \hat{\phi}) = \operatorname{argmin}_\theta \operatorname{argmax}_\phi \{ \mathbb{E}_{x \sim p(x)} \log[\sigma(w_\phi(x))] + \mathbb{E}_{\epsilon \sim q_0} \log[1 - \sigma(w_\phi(h_\theta(\epsilon)))] \}$ .

In (4)-(5) and in  $f$ -GAN, respective estimation of  $w_\phi(x)$  and  $T_\phi(x)$  yields approximation of a function

of a likelihood ratio; such an estimation appears to be at the heart of many GAN models. This understanding is our launching point for extending the range of applications of adversarial learning.

### 3 Unnormalized-Distribution GAN

In the above discussion, and in virtually all prior GAN research, access is assumed to samples from target distribution  $p(x)$ . In many applications samples from  $p(x)$  are unavailable, but the *unnormalized*  $u(x)$  is known, with  $p(x) = u(x)/C$  but with constant  $C = \int u(x)dx$  intractable. A contribution of this paper is a recasting of GAN to cases for which we have  $u(x)$  but no samples from  $p(x)$ , recognizing that most GAN models require an accurate estimate of the underlying likelihood ratio.

We consider the formulation in (4)-(5) and for simplicity set  $g_0(w_\phi(x)) = w_\phi(x)$ , although any choice of  $g_0(\cdot)$  may be considered as long as its monotonically increasing. The update of  $\theta$  remains as in (5), and we seek to estimate  $\log[p(x)/q_{\theta_{n-1}}(x)]$  based on knowledge of  $u(x)$ . Since  $\log[p(x)/q_{\theta_{n-1}}(x)] = \log[u(x)/q_{\theta_{n-1}}(x)] - \log C$ , for the critic it is sufficient to estimate  $\log[u(x)/q_{\theta_{n-1}}(x)]$ . Toward that end, we introduce a *reference* distribution  $p_r(x)$ , that (i) may be sampled easily, and (ii) has an explicit functional form that may be evaluated. The reference distribution can be connected to both importance sampling and the reference ratio method developed in bioinformatics [Hamelryck et al., 2010]. We have

$$\log\left[\frac{u(x)}{q_{\theta_{n-1}}(x)}\right] = \underbrace{\log\left[\frac{p_r(x)}{q_{\theta_{n-1}}(x)}\right]}_{\mathcal{F}_1} + \underbrace{\log\left[\frac{u(x)}{p_r(x)}\right]}_{\mathcal{F}_2} \quad (7)$$

where  $\mathcal{F}_2$  may be evaluated explicitly. We learn  $\mathcal{F}_1$  via (4), with  $\mathbb{E}_{x \sim p(x)}$  changed to  $\mathbb{E}_{x \sim p_r(x)}$ . Therefore, learning becomes alternating between the following two updates:

$$\begin{aligned} \phi_n &= \operatorname{argmax}_\phi \{ \mathbb{E}_{x \sim p_r(x)} \log[\sigma(w_\phi(x))] \\ &\quad + \mathbb{E}_{\epsilon \sim q_0} \log[1 - \sigma(w_\phi(h_{\theta_{n-1}}(\epsilon)))] \} \quad (8) \end{aligned}$$

$$\theta_n = \operatorname{argmax}_\theta \mathbb{E}_{\epsilon \sim q_0} \left[ w_{\phi_n}(h_\theta(\epsilon)) + \log\left[\frac{u(h_\theta(\epsilon))}{p_r(h_\theta(\epsilon))}\right] \right] \quad (9)$$

We call this procedure *reference-based adversarial sampling* (RAS) for unnormalized distributions. One should carefully note its distinction from the traditional GANs<sup>1</sup>, which usually learn to draw samples to mimic the given samples of a target distribution. To illustrate the difference, we visualize the learning schemes for the two settings in Figure 1.

The parameters of reference distribution  $p_r(x)$  are estimated using samples from  $q_\theta$ . We consider different forms of  $p_r$  depending on the application.

- **Unconstrained domains** For the case when the support  $\mathcal{X}$  of the target distribution is unconstrained, we model  $p_r(x)$  as a Gaussian distribution with diagonal covariance matrix, with mean and variance components estimated via samples from  $q_{\theta_{n-1}}$ , drawn as  $x = h_{\theta_{n-1}}(\epsilon)$  with  $\epsilon \sim q_0$ .
- **Constrained domains** In some real-world applications the support  $\mathcal{X}$  is bounded. For example, in reinforcement learning, the action often resides within a finite interval  $[c_1, c_2]$ . In this case, we propose to represent each dimension of  $p_r$  as a generalized Beta distribution  $\text{Beta}(\hat{\alpha}_0, \hat{\beta}_0, c_1, c_2)$ . The shape parameters are estimated using method of moments:  $\hat{\alpha}_0 = \bar{a} \left( \frac{\bar{a}(1-\bar{a})}{\bar{v}} - 1 \right)$  and  $\hat{\beta}_0 = (1 - \bar{a}) \left( \frac{\bar{a}(1-\bar{a})}{\bar{v}} - 1 \right)$ , where  $\bar{a} = \frac{\bar{a}' - c_1}{c_2 - c_1}$  and  $\bar{v} = \frac{\bar{v}'}{(c_2 - c_1)^2}$ , and  $\bar{a}'$  and  $\bar{v}'$  are sample mean and variance, respectively.

## 4 Entropy Regularization

Whether we perform adversarial learning based on samples from  $p(x)$ , as in Sec. 2, or based upon an unnormalized distribution  $u(x)$ , as in Sec. 3, the update of parameters  $\theta$  is of the form  $\theta_n = \text{argmax}_\theta \mathbb{E}_{\epsilon \sim q_0} g_0[\log(p(h_\theta(\epsilon))/q_{\theta_{n-1}}(h_\theta(\epsilon)))]$ , where  $\log(p(x)/q_{\theta_{n-1}}(x))$  is approximated as in (4) or its modified form (for learning from an unnormalized distribution).

A well-known failure mode of GAN is the tendency of the generative model,  $x = h_\theta(\epsilon)$  with  $\epsilon \sim q_0$ , to under-represent the full diversity of data that may be drawn  $x \sim p(x)$ . Considering  $\theta_n = \text{argmax}_\theta \mathbb{E}_{\epsilon \sim q_0} g_0[\log(p(h_\theta(\epsilon))/q_{\theta_{n-1}}(h_\theta(\epsilon)))]$ ,  $\theta_n$  will seek to favor synthesis of data  $x$  for which  $q_{\theta_{n-1}}(x)$  is small and  $p(x)$  large. When learning  $q_\theta(x)$  in this manner, at iteration  $n$  the model  $q_{\theta_n}$  tends to favor synthesis of a subset of data  $x$  that are probable from  $p(x)$  and less probable from  $q_{\theta_{n-1}}(x)$ . This subset of data that  $q_{\theta_n}$  models well can change with  $n$ , with the iterative learning continuously moving to model a subset of the data  $x$  that are probable via  $p(x)$ . This subset can be very small, in the worst case yielding

a model that always generates the *same* single data sample that looks like a real draw from  $p(x)$ ; in this case  $h_\theta(\epsilon)$  yields the same or near-same output for all  $\epsilon \sim q_0$ , albeit a realistic-looking sample  $x$ .

To mitigate this failure mode, it is desirable to add a regularization term to the update of  $\theta$ , encouraging that the entropy of  $q_{\theta_n}$  be large at each iteration  $n$ , discouraging the model from representing (while iteratively training) a varying small subset of the data supported by  $p(x)$ . Specifically, consider the regularized update of (5) as:

$$\theta_n = \text{argmax}_\theta \mathbb{E}_{\epsilon \sim q_0} g_0[w_{\phi_n}(h_\theta(\epsilon))] + \beta H(q_\theta) \quad (10)$$

where  $H(q_\theta)$  represents the entropy of the distribution  $q_\theta$ , for  $\beta > 0$ . The significant challenge is that  $H(q_\theta) = -\mathbb{E}_{x \sim q_\theta} \log q_\theta(x)$ , but by construction we lack an explicit form for  $q_\theta(x)$ , and hence the entropy may not be computed directly. Below we consider two means by which we may approximate  $H$ , one of which is explicitly appropriate for the case in which we learn based upon the unnormalized  $u(x)$ , and the other of which is applicable to whether we learn via samples from  $p(x)$  or based on  $u(x)$ .

In the case for which  $p(x) = u(x)/C$  and  $u(x)$  is known, we may consider approximating or replacing  $H(q_\theta)$  with  $-\mathbb{E}_{x \sim q_\theta} \log u(x) + \log C$ , and the term  $\log C$  may be ignored, because it doesn't impact the regularization in (10); we therefore replace the entropy  $H(q_\theta)$  with the cross entropy  $-\mathbb{E}_{x \sim q_\theta} \log p(x)$ . The first term in (10) tends to encourage the model to learn to draw samples where  $p(x)$ , or  $u(x)$ , is large, while the second term discourages over-concentration on such high-probability regions, as  $-\mathbb{E}_{x \sim q_\theta} \log p(x)$  becomes large when  $q_\theta$  encourages samples near lower probability regions of  $p(x)$ . This will ideally yield a spreading-out of the samples encouraged by  $q_\theta$ , with high-probability regions of  $p(x)$  modeled well, but also regions spreading out from these high-probability regions.

To gain further insight into (10), we again consider the useful case of  $g_0[w_{\phi_n}(h_\theta(\epsilon))] = w_{\phi_n}(h_\theta(\epsilon))$  and assume the ideal solution  $w_{\phi_n}(x) = \log[p(x)/q_{\theta_{n-1}}(x)]$ . In this case cross-entropy-based regularization may be seen as seeking to maximize wrt  $\theta$  the function

$$\begin{aligned} \mathbb{E}_{x \sim q_\theta} \log[p(x)/q_\theta(x)] & - \beta \mathbb{E}_{x \sim q_\theta} \log p(x) \\ & = \mathbb{E}_{x \sim q_\theta} \log[p(x)^{1-\beta}/q_\theta(x)] \end{aligned}$$

For the special case of  $p(x) = \exp[-E(x)]/C$ , with  $E(x) > 0$  an ‘‘energy’’ function, we have  $p(x)^{1-\beta} = \exp[-\frac{1}{T_\beta} E(x)]/C$  with  $T_\beta = 1/(1 - \beta)$ . Hence, the cross-entropy regularization is analogous to annealing, with  $\beta \in [0, 1)$ ;  $\beta \rightarrow 1_-$  corresponds to high ‘‘temperature’’  $T_\beta$ , which as  $\beta \rightarrow 0_+$  is lowered and with  $p(x)^{1-\beta} \rightarrow p(x)$ . When  $\beta > 0$  the peaks in  $p(x)$  are

<sup>1</sup>We refer to generative models learned via samples as GAN, and generative models learned via an unnormalized distribution as RAS.

“flattened out,” allowing the model to yield samples that “spread out” and explore the diversity of  $p(x)$ . This interpretation suggests learning via (10), with the cross-entropy replacement for  $H(q_\theta)$ , with  $\beta$  near 1 one at the start, and progressively reducing  $\beta$  toward 0 (corresponding to lowering temperature  $T_\beta$ ).

The above setup assumes we have access to  $u(x)$ , which is not the case when we seek to learn  $q_\theta$  based on samples of  $p$ . Further, rather than replacing  $H(q_\theta)$  by the cross-entropy, we may wish to approximate  $H(q_\theta)$  based on samples of  $q_\theta$ , which we have via  $x = h_\theta(\epsilon)$  with  $\epsilon \sim q_0$  (with this estimated via samples from  $p(x)$  or based on  $u(x)$ ). Toward that end, consider the following lemma.

**Lemma 1** *Let  $t_\xi(\epsilon|x)$  be a probabilistic inverse mapping associated with the generator  $q_\theta(x)$ , with parameters  $\xi$ . The mutual information between  $x$  and  $\epsilon$  satisfies*

$$I(x; \epsilon) = H(q_\theta) \geq H(q_0) + \mathbb{E}_{\epsilon \sim q_0} \log t_\xi(\epsilon|h_\theta(\epsilon)). \quad (11)$$

The proof is provided in the Supplement Material (SM). Since  $H(q_0)$  is a constant wrt  $(\theta, \xi)$ , one may seek to maximize  $\mathbb{E}_{\epsilon \sim q_0} \log t_\xi(\epsilon|h_\theta(\epsilon))$  to increase the entropy  $H(q_\theta)$ . Hence, in (10) we replace the entropy term with  $\mathbb{E}_{\epsilon \sim q_0} \log t_\xi(\epsilon|h_\theta(\epsilon))$ .

In practice we consider  $t_\xi(\epsilon|x) = \mathcal{N}(\epsilon; \mu_\xi(x), I)$ , where here  $I$  is the identity matrix, and  $\mu_\xi(x)$  is a vector mean. Hence,  $H(q_\theta)$  in (10) is replaced by  $-\mathbb{E}_{\epsilon \sim q_0} \|\epsilon - \mu_\xi(h_\theta(\epsilon))\|_2^2$ . Note that a failure mode of GAN, as discussed above, corresponds to many or all  $\epsilon \sim q_0$  being mapped via  $x = h_\theta(\epsilon)$  to the same output. This is discouraged via this regularization, as such behavior makes it difficult to simultaneously minimize  $\mathbb{E}_{\epsilon \sim q_0} \|\epsilon - \mu_\xi(h_\theta(\epsilon))\|_2^2$ . This regularization is related to cycle-consistency [Li et al., 2017a]. However, the justification of the negative cycle-consistency as a lower bound of  $H(q_\theta)$  is deemed a contribution of this paper (not addressed in [Li et al., 2017a]).

## 5 Related Work

**Use of a reference distribution** We have utilized a readily-sampled reference distribution, with known density function  $p_r(x)$ , when learning to sample from an unnormalized distribution  $u(x)$ . The authors of [Gutmann and Hyvärinen, 2010] also use such a reference distribution to estimate the probability distribution associated with observed data samples. However, [Gutmann and Hyvärinen, 2010] considered a distinct problem, for which one wished to fit observed samples to a *specified* unnormalized distribution. Here we employ the reference distribution in the context of learning to sample from a known  $u(x)$ , with no empirical samples from  $p(x)$  provided.

**Adversarial variational Bayes** In the context of variational Bayes analysis, the adversarial variational Bayes (AVB) [Mescheder et al., 2017] was proposed for posterior inference of variational autoencoders (VAEs) [Kingma and Welling, 2014]. Assume we are given a *parametric* generative model  $p_\theta(x|z)$  with prior  $p(z)$  on latent code  $z$ , designed to model observed data samples  $\{x_i\}_{i=1,N}$ . There is interest in designing an inference arm, capable of efficiently inferring a distribution on the latent code  $z$  given observed data  $x$ . Given observed  $x$ , the posterior distribution on the code is  $p_\theta(z|x) = p_\theta(x|z)p(z)/p_\theta(x) \propto p_\theta(x|z)p(z)$ , where  $p_\theta(x) = \int p_\theta(x|z)p(z)dz$ , and  $u_\theta(z;x) = p_\theta(x|z)p(z)$  represents an *unnormalized* distribution of the latent variable  $z$ , which also depends on the data  $x$ .

One may show that if the procedure in Sec. 3 is employed to draw samples from  $p_\theta(z|x)$ , based on the unnormalized  $u_\theta(z;x)$ , one exactly recovers AVB [Mescheder et al., 2017]. The AVB considered  $g_0(w_\phi) = w_\phi$  within our framework. We do not consider the application to inference with VAEs, as the experiments in [Mescheder et al., 2017] are applicable to the framework we have developed. The generality of the RAS is made more clear in our paper. We show its applicability to reinforcement learning in Sec. 6, and broaden the discussion on the type of adaptive reference distributions in Sec. 3, with extensions to constrained domain sampling.

**Regularization** The term  $\mathbb{E}_{\epsilon \sim q_0} \log t_\xi(\epsilon|h_\theta(\epsilon))$  employed here was considered in [Li et al., 2017a, Chen et al., 2018b, Zhu et al., 2017], but the use of it as a bound on the entropy of  $q_\theta$  is new. From Lemma 1 we see that  $\mathbb{E}_{\epsilon \sim q_0} \log t_\xi(\epsilon|h_\theta(\epsilon))$  is also a bound on the mutual information between  $x$  and  $\epsilon$ , maximization of which is the same goal as InfoGAN [Chen et al., 2016]. However, unlike in [Chen et al., 2016], here the mapping  $\epsilon \rightarrow x$  is deterministic, where in InfoGAN it is stochastic. Additionally, the goal here is to encourage diversity in generated  $x$ , which helps mitigate mode collapse, where in InfoGAN the goal was to discover latent semantic concepts.

**Stein variational gradient descent (SVGD)** In the formulation of (4)-(5), if one sets  $g_0(w_\phi) = w_\phi$ , then the learning objective corresponds to minimizing the reverse KL divergence  $\text{KL}(q_\theta(x)||p(x))$ . SVGD [Liu and Wang, 2016] also addresses this goal given unnormalized distribution  $u(x)$ , with  $p(x) = u(x)/C$ . Like for the proposed approach, the goal is not to explicitly learn a functional form for  $q_\theta(x)$ , rather the goal of SVGD is to learn to draw samples from it. We *directly* learn a sampler model via  $x = h_\theta(\epsilon)$  and  $\epsilon \sim q_0$ , where in [Wang and Liu, 2016] a specified set of samples is adjusted sequentially to correspond to draws from the unnormalized distribution  $u(x)$ . In




Figure 2: Comparison of different GAN variants. The GAN models and corresponding entropy-regularized variants are visualized in the same color; in each case, the left result is unregularized, and the right employs entropy regularization. The black dots indicate the means of the distributions.

this setting, one assumes access to a set of samples  $\{x_i\}$  drawn from some distribution, and these samples are updated deterministically as  $x'_i = x_i + \mu\gamma(x_i)$  where  $\mu > 0$  is a small step size, and  $\gamma(x)$  is a nonlinear function, assumed described by a reproducing kernel Hilbert space (RKHS) with given kernel  $k(x, x')$ . In this setting, the samples are updated  $\{x_i\} \rightarrow \{x'_i\}$ , with a deterministic function  $\gamma(x)$  that is evaluated in terms of  $u(x)$  and  $\nabla_x u(x)$ . While this process is capable of transforming a specific set of samples such that they ultimately approximate samples drawn from  $p(x)$ , we do not have access to a model  $x = h_\theta(\epsilon)$  that allows one to draw new samples quickly, on demand. Consequently, within the SVGD framework, a model  $x = h_\theta(\epsilon)$  is learned separately as a second ‘‘amortization’’ step. The two-step character of SVGD should be contrasted with the direct approach of the proposed model to learn  $x = h_\theta(\epsilon)$ . SVGD has been demonstrated to work well, and therefore it is a natural model against which to compare, as considered below.

## 6 Experimental Results

The Tensorflow code to reproduce the experimental results is at [github](https://github.com/ChunyuanLI/RAS)<sup>2</sup>.

### 6.1 Effectiveness of Entropy Regularization

#### 6.1.1 Learning based on samples

We first demonstrate that the proposed entropy regularization improves mode coverage when learning based on samples. Following the design in [Metz et al., 2017], we consider a synthetic dataset of samples drawn from a 2D mixture of 8 Gaussians. The results on real datasets are reported in SM.

We consider the original GAN and three state-of-the-art GAN variants: Unrolled-GAN [Metz et al., 2017], D2GAN [Nguyen et al., 2017] and Spectral Normaliza-

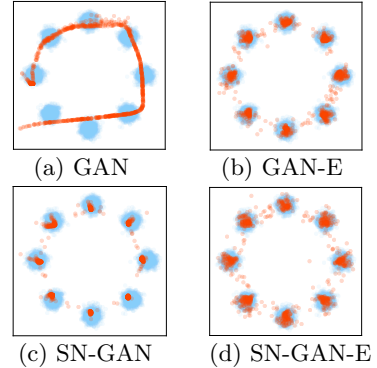


Figure 3: Generated samples.

tion (SN)-GAN [Miyato et al., 2018]. For simplicity, we consider the case when  $g_0(\cdot)$  is an identity function, and this form of GAN is denoted as *adversarially learned likelihood-ratio* (ALL) in Fig. 2.

For all variants, we study their entropy-regularized versions, by adding the entropy bound in (11), when training the generator. If not specifically mentioned, we use a *fix-and-decay* scheme for  $\beta$  for all experiments: In total  $T$  training iterations, we first fix  $\beta = 1$  in the first  $T_0$  iteration, then linearly decay it to 0 in the rest  $T - T_0$  iterations. On this 8-Gaussian dataset,  $T = 50k$  and  $T_0 = 10k$ .


Twenty runs were conducted for each algorithm. Since we know the true distribution in this case, we employ the symmetric KL divergence as a metric to quantitatively compare the quality of generated data. In Fig. 2 we report the distribution of divergence values for all runs. We add the entropy bound to each variant, and visualize their results as violin plots with gray edges (the color for each variant remains for comparison). The largely decreased mean and reduced variance of the divergence show that the entropy annealing yields significantly more consistent and reliable solutions, across all methods. We plot the generated samples in Fig. 3. We visualize the generated samples of the original GAN in Fig. 3(a). The samples ‘‘struggle’’ between covering all modes and separating modes. This issue is significantly reduced by ALL with entropy regularization, as shown in Fig. 3(b). SN-GAN (Fig. 3(c)) generates samples that concentrate only around the centroid of the mode. However, after adding our entropy regularizer (Fig. 3(d)), the issue is alleviated and the samples spread out.

#### 6.1.2 Learning based on an unnormalized distribution

When the unnormalized form of a target distribution is given, we consider two types of entropy regularization to improve our RAS algorithm: (i)  $E_{cc}$ : the cycle-

<sup>2</sup><https://github.com/ChunyuanLI/RAS>





(a) Target (b) RAS (c) RAS+E<sub>cc</sub> (d) RAS+E<sub>cc</sub>  
 Figure 4: Entropy regularization for unnormalized distributions.

consistency-based regularization; (ii)  $E_{cc}$ : the cross-entropy-based regularization. To clearly see the advantage of the regularizers, we fix  $\beta = 0.5$  in this experiment. Figure 4 shows the results, with each case shown in one row. The target distributions are shown in column (a), the sampling results of RAS are shown in column (b). RAS can reflect the general shape of the underlying distribution, but tends to concentrate on the high density regions. The two entropy regularizers are shown in (c) and (d). The entropy encourages the samples to spread out, leading to better approximation, and  $E_{cc}$  appears to yield best performance.

### 6.1.3 Comparison of two learning settings

In traditional GAN learning, we have a finite set of  $N$  samples with the empirical distribution  $p'(x)$  to learn from, each sample drawn from the true distribution  $p(x)$ . It is known that the optimum of GANs yields the marginal distribution matching  $q_\theta(x) = p'(x)$  [Goodfellow et al., 2014]; it also implies that the performance of  $q_\theta(x)$  in is limited by  $p'(x)$ . In contrast, when we learn from an unnormalized form as in RAS, the likelihood ratio is estimated using samples drawn from  $p_r(x)$  and from  $q_\theta$ . Hence, we can draw as many samples as desired to get an accurate likelihood-ratio estimation, which further enables  $q_\theta$  to approach  $p(x)$ . This means RAS can potentially provide better approximation, when  $u(x)$  is available.

We demonstrate this advantage on the above 8-Gaussian distribution. We train GAN on  $p'(x)$  with  $N = 100, 1000, 10000, 100000$  samples, and train RAS on  $u(x)$ . Note that the samples from  $p_r$  and  $q_\theta$  are drawn in an online fashion to train RAS. With an appropriate number of iterations ( $T = 50k$ ) to assure convergence, in total  $T \cdot B \approx 50M$  samples were used to estimate the likelihood ratio in (8), where  $B = 1024$  is the minibatch size.

In the evaluation stage, we draw 20k samples from  $q_\theta$  for each model, and compute the symmetric KL di-




Figure 5: Comparison of learning via GAN and RAS.

vergence against the true distribution. The results are shown in Figure 5. As an illustration for the ideal performance, we draw 20k samples from the target distribution and show its divergence as the black line. The GAN gradually performs better, as more target samples are available in training. However, they are still worse than RAS by a fairly large margin.

### 6.2 Sampling from Constrained Domains

To show that RAS can draw samples when  $\mathcal{X}$  is bounded, we apply it to sample from the distributions with the support  $[-1, 1]$ . The details for the functions and decay of  $\beta$  are in SM. We adopt the Beta distribution as our reference, whose parameters are estimated using the method of moments (see Sec. 3). The activation function in the last layer of the generator is chosen as  $\tanh$ . As a baseline, we naively use an empirical Gaussian as the reference. We also compare with the standard SVGD [Liu and Wang, 2016] and the amortized SVGD methods [Wang and Liu, 2016], in which 512 particles are used.

Figure 6 shows the comparison. Note that since the support of the Beta distribution is defined in an interval, our RAS can easily match this reference distribution, leading the adversary to accurately estimate the likelihood ratio. Therefore, it closely approximates the target, as shown in Figure 6(a). Alternatively, when a Gaussian reference is considered, the adversarial ratio estimation can be inaccurate in the low density regions, resulting in degraded sampling performance shown in Figure 6(b). Since SVGD is designed for sampling in unconstrained domains, a principled mechanism to extend it for a constrained domain is less clear. Figure 6(c) shows SVGD results, and a substantial percentage of particles fall out of the desired domain. The amortized SVGD method adopts an  $\ell_2$  metric to match the generator’s samples to the SVGD targets, it collapses to the distribution mode, as in Figure 6(d). We observed that the amortized MCMC results [Li et al., 2017b, Chen et al., 2018a] are similar to the amortized SVGD [Li et al., 2018].

### 6.3 Soft Q-learning

Soft Q-learning (SQL) has been proposed recently [Haarnoja et al., 2017], with reinforcement learning (RL) policy based on a general class of distributions, with the goal of representing complex, multimodal behavior. An agent can take an action  $a \in \mathcal{A}$  based on a policy  $\pi(a|s)$ , defined as the probability of taking action  $a$  when in state  $s$ . It is shown in [Haarnoja et al., 2017] that the target policy has a known unnormalized density  $u(a; s)$ .

**SQL-SVGD** To take actions from the optimal policy (*i.e.*, sampling), learning  $\theta$  in

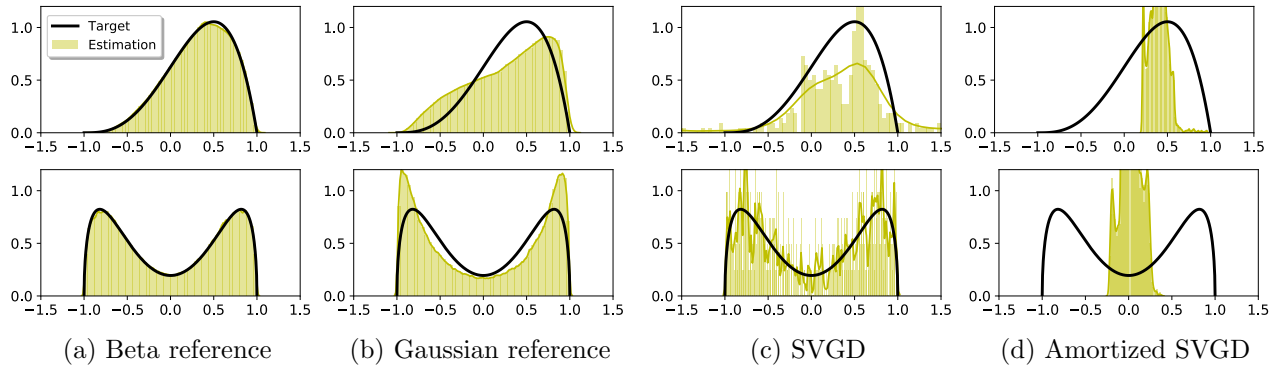


Figure 6: Sampling from constrained domains

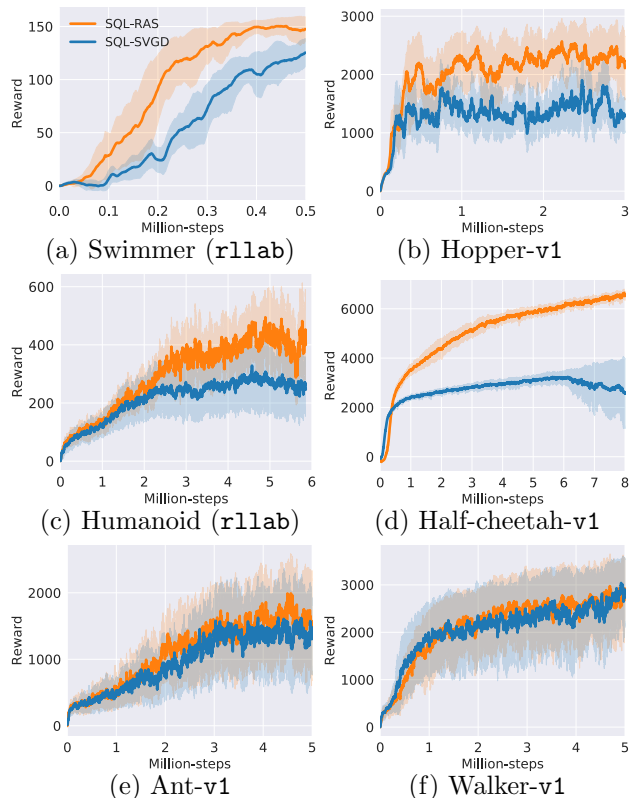


Figure 7: Soft Q-learning on MuJoCo environments.

[Haarnoja et al., 2017] is performed via amortized SVGD in two separated steps: (i) the samples of  $u(a|s)$  are first drawn using SVGD by minimizing  $KL(q_\theta(a|s)||u(a;s))$ ; (ii) these samples are then used as the target to update  $\theta$  under an  $\ell_2$  amortization metric. We call this procedure as SQL-SVGd.

**SQL-RAS** Alternatively, we apply our RAS algorithm to replace the amortized SVGD. When the action space is unconstrained, we may use the Gaussian reference  $p_r$ . However, the action space is often constrained in continuous control, with each dimension in an interval  $[c_1, c_2]$ . Hence, we adopt the Beta-distribution reference for RAS.

Following [Haarnoja et al., 2018], we compare RAS with amortized SVGD on six continuous control tasks: Hopper, Half-cheetah, Ant and Walker from the OpenAI gym benchmark suite [Brockman et al., 2016], as well as the Swimmer and Humanoid tasks in the rllab implementation [Duan et al., 2016]. Note that the action space is constrained in  $[-1, 1]$  for all the tasks. The dimension of the action space ranges from 2 to 21 on the different tasks. The higher-dimension environments are usually harder to solve. All hyperparameters used in this experiment are listed in SM.

Figure 7 shows the total average return of evaluation rollouts during training. We train 3 different instances of each algorithm, with each performing one evaluation rollout every 1k environment steps. The solid curves corresponds to the mean and the shaded regions to the standard deviation. Overall, it shows that RAS significantly outperforms amortized SVGD on four tasks both in terms of learning speed and the final performance. This includes the most complex benchmark, the 21-dimensional Humanoid (rllab). On other two tasks, the two methods perform comparably. In the SQL setting, learning a good stochastic policy with entropy maximization can help training. It means that RAS can better estimate the target policy.

## 7 Conclusions

We introduce a reference-based adversarial sampling method as a general approach to draw from unnormalized distributions. It allows us to extend GANs from traditional sample-based learning setting to this new setting, and provide novel methods for important downstream applications, *e.g.*, Soft Q-learning. RAS can also be easily used for constrained domain sampling. Further, an entropy regularization is proposed to improve the sample quality, applicable to learning from samples or an unnormalized distribution. Extensive experimental results show the effectiveness of the entropy regularization. In Soft Q-learning, RAS provides performance comparable to, if not better than, its alternative method amortized SVGD.



**Acknowledgements** We thank Rohith Kudithipudi, Ruiyi Zhang, Yulai Cong and Ricardo Henao for helpful feedback/editing. We acknowledge anonymous reviewers for proofreading and improving the manuscript.

## References

- [Brockman et al., 2016] Brockman, G., Cheung, V., Petterson, L., Schneider, J., Schulman, J., Tang, J., and Zaremba, W. (2016). Openai gym. *arXiv preprint arXiv:1606.01540*.
- [Brooks et al., 2011] Brooks, S., Gelman, A., Jones, G., and Meng, X.-L. (2011). *Handbook of Markov Chain Monte Carlo*.
- [Chen et al., 2018a] Chen, C., Li, C., Chen, L., Wang, W., Pu, Y., and Duke, L. C. (2018a). Continuous-time flows for efficient inference and density estimation. In *International Conference on Machine Learning*, pages 823–832.
- [Chen et al., 2018b] Chen, L., Dai, S., Pu, Y., Li, C., Su, Q., and Carin, L. (2018b). Symmetric variational autoencoder and connections to adversarial learning. *AISTATS*.
- [Chen et al., 2016] Chen, X., Duan, Y., Houthoofd, R., Schulman, J., Sutskever, I., and Abbeel, P. (2016). InfoGAN: Interpretable representation learning by information maximizing generative adversarial nets. In *NIPS*.
- [Duan et al., 2016] Duan, Y., Chen, X., Houthoofd, R., Schulman, J., and Abbeel, P. (2016). Benchmarking deep reinforcement learning for continuous control. In *ICML*.
- [Feng et al., 2017] Feng, Y., Wang, D., and Liu, Q. (2017). Learning to draw samples with amortized stein variational gradient descent. *UAI*.
- [Gelman et al., 1995] Gelman, A., Carlin, J. B., S., S. H., and Rubin, D. B. (1995). Bayesian data analysis. *London: Chapman and Hall*.
- [Goodfellow et al., 2014] Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. (2014). Generative adversarial nets. In *NIPS*.
- [Gutmann and Hyvärinen, 2010] Gutmann, M. and Hyvärinen, A. (2010). Noise-contrastive estimation: A new estimation principle for unnormalized statistical models. In *AISTATS*.
- [Haarnoja et al., 2017] Haarnoja, T., Tang, H., Abbeel, P., and Levine, S. (2017). Reinforcement learning with deep energy-based policies. *ICML*.
- [Haarnoja et al., 2018] Haarnoja, T., Zhou, A., Abbeel, P., and Levine, S. (2018). Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. *ICML*.
- [Hamelryck et al., 2010] Hamelryck, T., Borg, M., Paluszewski, M., Paulsen, J., Frellsen, J., Andreetta, C., Boomsma, W., Bottaro, S., and Ferkinghoff-Borg, J. (2010). Potentials of mean force for protein structure prediction vindicated, formalized and generalized. *PLoS one*.
- [Hastings, 1970] Hastings, W. (1970). Monte Carlo sampling methods using Markov Chains and their applications. *Biometrika*.
- [Heusel et al., 2017] Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Klambauer, G., and Hochreiter, S. (2017). GANs trained by a two time-scale update rule converge to a Nash equilibrium. *NIPS*.
- [Hoffman et al., 2013] Hoffman, M. D., Blei, D. M., Wang, C., and Paisley, J. (2013). Stochastic variational inference. *The Journal of Machine Learning Research*.
- [Kanamori et al., 2010] Kanamori, T., Suzuki, T., and Sugiyama, M. (2010). Theoretical analysis of density ratio estimation. *IEICE Trans. Fund. Electronics, Comm., CS*.
- [Kingma and Welling, 2014] Kingma, D. P. and Welling, M. (2014). Auto-encoding variational Bayes. *ICLR*.
- [Krizhevsky et al., 2012] Krizhevsky, A., Sutskever, I., and Hinton, G. E. (2012). Imagenet classification with deep convolutional neural networks. In *NIPS*.
- [Li et al., 2018] Li, C., Li, J., Wang, G., and Carin, L. (2018). Learning to sample with adversarially learned likelihood-ratio.
- [Li et al., 2017a] Li, C., Liu, H., Chen, C., Pu, Y., Chen, L., Henao, R., and Carin, L. (2017a). ALICE: Towards understanding adversarial learning for joint distribution matching. *NIPS*.
- [Li et al., 2015] Li, Y., Hernández-Lobato, J. M., and Turner, R. E. (2015). Stochastic expectation propagation. In *NIPS*.
- [Li et al., 2017b] Li, Y., Turner, R. E., and Liu, Q. (2017b). Approximate inference with amortised MCMC. *arXiv preprint arXiv:1702.08343*.
- [Liu and Wang, 2016] Liu, Q. and Wang, D. (2016). Stein variational gradient descent: A general purpose Bayesian inference algorithm. In *NIPS*.

- [Liu et al., 2015] Liu, Z., Luo, P., Wang, X., and Tang, X. (2015). Deep learning face attributes in the wild. In *ICCV*.
- [Mescheder et al., 2017] Mescheder, L., Nowozin, S., and Geiger, A. (2017). Adversarial variational Bayes: Unifying variational autoencoders and generative adversarial networks. In *ICML*.
- [Metz et al., 2017] Metz, L., Poole, B., Pfau, D., and Sohl-Dickstein, J. (2017). Unrolled generative adversarial networks. *ICLR*.
- [Minka, 2001] Minka, T. P. (2001). Expectation propagation for approximate Bayesian inference. In *UAI*.
- [Miyato et al., 2018] Miyato, T., Kataoka, T., Koyama, M., and Yoshida, Y. (2018). Spectral normalization for generative adversarial networks. In *ICLR*.
- [Mohamed and L., 2016] Mohamed, S. and L., B. (2016). Learning in implicit generative models. *NIPS workshop on adversarial training*.
- [Neyman and Pearson, 1933] Neyman, J. and Pearson, E. S. (1933). On the problem of the most efficient tests of statistical hypotheses. *Phil. Trans. R. Soc. Lond. A*, 231(694-706):289–337.
- [Nguyen et al., 2017] Nguyen, T., Le, T., Vu, H., and Phung, D. (2017). Dual discriminator generative adversarial nets. *NIPS*.
- [Nguyen et al., 2010a] Nguyen, X., Wainwright, M., and Jordan, M. (2010a). Estimating divergence functionals and the likelihood ratio by convex risk minimization. *IEEE Trans. Info. Theory*.
- [Nguyen et al., 2010b] Nguyen, X., Wainwright, M. J., and Jordan, M. I. (2010b). Estimating divergence functionals and the likelihood ratio by convex risk minimization. *IEEE Transactions on Information Theory*.
- [Nowozin et al., 2016] Nowozin, S., Cseke, B., and Tomioka, R. (2016). f-GAN: Training generative neural samplers using variational divergence minimization. *NIPS*.
- [Oord et al., 2016] Oord, A., Kalchbrenner, N., and Kavukcuoglu, K. (2016). Pixel recurrent neural network. In *ICML*.
- [Radford et al., 2016] Radford, A., Metz, L., and Chintala, S. (2016). Unsupervised representation learning with deep convolutional generative adversarial networks. In *ICLR*.
- [Rezende et al., 2014] Rezende, D. J., Mohamed, S., and Wierstra, D. (2014). Stochastic backpropagation and approximate inference in deep generative models. In *ICML*.
- [Uehara et al., 2016] Uehara, M., Sato, I., Suzuki, M., Nakayama, K., and Matsuo, Y. (2016). Generative adversarial nets from a density ratio estimation perspective. *arXiv preprint arXiv:1610.02920*.
- [Van Trees, 2001] Van Trees, H. L. (2001). *Detection, estimation, and modulation theory*. John Wiley & Sons.
- [Wang and Liu, 2016] Wang, D. and Liu, Q. (2016). Learning to draw samples: With application to amortized MLE for generative adversarial learning. In *arXiv:1611.01722v2*.
- [Welling and Teh, 2011] Welling, M. and Teh, Y. W. (2011). Bayesian learning via stochastic gradient Langevin dynamics. In *ICML*.
- [Y. Pu and Carin, 2017] Y. Pu, Z. Gan, R. H. C. L. S. H. and Carin, L. (2017). VAE learning via Stein variational gradient descent. *NIPS*.
- [Zhu et al., 2017] Zhu, J.-Y., Park, T., Isola, P., and Efros, A. (2017). Unpaired image-to-image translation using cycle-consistent adversarial networks. *ICCV*.