



HAL
open science

Binary classification based Monte Carlo simulation

Elouan Argouarc'H, François Desbouvries

► **To cite this version:**

Elouan Argouarc'H, François Desbouvries. Binary classification based Monte Carlo simulation. *IEEE Signal Processing Letters*, 2024, 31, pp.1449-1453. 10.1109/LSP.2024.3396403 . hal-04591977

HAL Id: hal-04591977

<https://hal.science/hal-04591977v1>

Submitted on 29 May 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Binary classification based Monte Carlo simulation

Elouan Argouarc'h, François Desbouvries, *Senior Member, IEEE*

Abstract—Acceptance-rejection (AR), Independent Metropolis Hastings (IMH) or Importance Sampling (IS) Monte Carlo (MC) algorithms all involve computing *ratios* of two probability density functions (pdf) p_1 and p_0 . On the other hand, classifiers discriminate samples produced by a binary mixture and can be used to approximate the *ratio* of corresponding pdfs. We therefore establish a bridge between simulation and classification, which enables us to propose pdf-free versions of ratio-based simulation algorithms, where the ratio is replaced by a surrogate function computed via a classifier. Our modified samplers are based on very different hypotheses: the knowledge of functions p_1 and p_0 is relaxed (- they may be totally unknown), and is counterbalanced by the availability of a classification function, which can be obtained from a labelled dataset. From a probabilistic modeling perspective, our procedure involves a structured energy based model which can easily be trained and is structurally compatible with the classical samplers.

Index Terms—Neural Classification, Binary Cross Entropy, Stochastic Simulation, Acceptance-Rejection, Importance Sampling, MCMC, Energy Based Models

I. INTRODUCTION

If a and b are two positive numbers,

$$r = \frac{a}{a+b} \in (0, 1) \Leftrightarrow \frac{r}{1-r} = \frac{a}{b} > 0. \quad (1)$$

This equivalence has interesting consequences in Bayesian classification, machine learning and stochastic simulation. Indeed, if a and b are probabilities of two classes in a binary mixture context for a given sample, then ratio $\frac{a}{a+b}$ is the posterior probability which provides with the class probabilities for a given sample, and can be approximated by a parametric classifier r_ϕ trained to distinguish between the two probability distributions. On the other hand, positive ratios $\frac{a}{b}$ play a key role in AR, IMH or IS techniques. Equation (1) relates r to such positive ratios, and tells us that ratio $\frac{a}{b}$ can be computed exactly from r , or, in practice, approximately from r_ϕ , without necessarily knowing a nor b . This observation enables us to propose approximate versions of these algorithms which rely on weaker hypotheses.

Let $\lambda, 1 - \lambda \in (0, 1)$ be the prior probabilities of two categories $k = 1, 0$, distributed resp. $\sim p_1$ and p_0 . Binary classification distinguishes samples from mixture $\lambda p_1 + (1 - \lambda)p_0$ by identifying the pdf which generated them. The appropriate way to classify relies on the posterior probability: x is a sample $\sim p_1$ rather than $\sim p_0$ with probability

$$\Pr(k = 1|x, \lambda, p_0, p_1) = \frac{\lambda p_1(x)}{\lambda p_1(x) + (1 - \lambda)p_0(x)}. \quad (2)$$

Elouan Argouarc'h and François Desbouvries are with SAMOVAR, Télécom SudParis, Institut Polytechnique de Paris, 91120 Palaiseau, France (e-mail: elouan.argouarch, francois.desbouvries@telecom-sudparis.eu). Elouan Argouarc'h is also with Université Paris Saclay, CEA, List F-91120 Palaiseau, France (e-mail: elouan.argouarch@cea.fr)

Indeed, as is well known (see e.g. [1, Chap. 11]), assigning a sample to the label with highest posterior probability is the optimal decision rule in the sense that it minimizes the probability of misclassification.

To compute this posterior probability, one needs to evaluate the pdfs p_1, p_0 and know the prior probability λ but they are often unknown, leaving (2) intractable. If however we dispose of a set $\mathcal{D} = \{(x_i^{(k_i)}, k_i)\}_{i=1}^{N_0+N_1}$ of labelled observations, we can make use of a parametric classifier. So let us assume that we have at our disposal a classifier function r_ϕ , parameterized by ϕ which mimics the unknown posterior pdf:

$$r_\phi(x) \approx \frac{N_1 p_1(x)}{N_1 p_1(x) + N_0 p_0(x)}. \quad (3)$$

Our paper is based on the observation that (3) is equivalent to

$$\frac{N_0}{N_1} \frac{r_\phi(x)}{1 - r_\phi(x)} \approx \frac{p_1(x)}{p_0(x)}, \quad (4)$$

which implies that (typically neural network based) classifiers can also be used for approximating *pdf ratios*.

Equation (4) has already been observed, and exploited in contexts where estimating a ratio of pdfs is relevant. First, classifiers are at the core of adversarial training techniques in which divergence measures involving a ratio are replaced by an approximation based on a classifier [2]. This enables to learn implicit generative models (i.e., with intractable pdfs) [3] [4]. Moreover, classifier based pdf ratio approximation has been applied to estimation of such metrics as Mutual Information [5]. Finally, classifiers based ratios have been applied successfully in statistical hypothesis testing procedures [6], which heavily rely on likelihood-ratio tests.

If p_0 is an instrumental distribution with tractable pdf, then (4) can be turned into an approximation of target pdf p_1 . So classifiers can be used for density estimation, conditional density estimation, or likelihood-to-evidence ratio estimation, making them especially relevant in a likelihood-free inference setting [7][8][9].

However, the question of *sampling* from the corresponding model remains open, and this is precisely the point we discuss in this paper. We realize that pdf ratios also play a key role in such simulation techniques as the AR or Markov Chain Monte Carlo (MCMC) methods, in which samples from instrumental p_0 are transformed into samples from the target p_1 via a sampling mechanism which involves the ratio of the two densities. This establishes a connection between classification and MC sampling, and will enable us to relax the assumption of tractable pdf p_0, p_1 of these sampling algorithms, at the price of approximate sampling. Our approach is therefore completely pdf-free, and as such is especially relevant when the target distribution is unknown or with intractable, noisy, or costly to evaluate pdf (see [10] for a review of MC techniques

in this setting, and [11] for a review of likelihood-free Approximate Bayesian Computation techniques); and/or when the instrumental p_0 is defined by a generative model with implicit pdf [12][3][13]. The rest of this paper is organized as follows. In §II we recall classical ratio-based stochastic simulation algorithms, i.e. the AR, IMH and IS techniques. In §III we show that classifiers computed via the Binary Cross Entropy (BCE) criterion indeed provide with an approximation of the posterior (2). Finally in §IV we propose classification based sampling methods, illustrate our method via simulations¹, and revisit it under the perspective of probabilistic modeling. We end the paper with a conclusion.

II. CLASSICAL RATIO-BASED SAMPLING ALGORITHMS

Stochastic simulation includes a variety of techniques, see e.g. [14]–[20]. In this section we focus on AR, IMH and IS which share in common that they all compute a ratio of pdfs.

A. The AR algorithm

1) *A brief reminder of AR Sampling:* AR Sampling [16, chap. 2] [19, chap. 3] is a simulation algorithm that yields samples distributed according to a target distribution p via samples from a proposal distribution q , which are accepted or rejected as valid samples from p via some acceptance probability. More precisely, let the support of p be inside that of q . This means that there exists a constant $C \geq 1$ such that for all $x \in \mathbb{R}^d$, $p(x) \leq Cq(x)$. Let $X \sim q$, and let k a Bernoulli random variable with parameter $\alpha_{AR}(X) = \frac{p(X)}{Cq(X)}$. AR sampling is based on the fact that $X|k=1$ is distributed according to p . Note that $\Pr(k=1) = \frac{1}{C}$, so the lower the value of C , the higher the acceptance rate.

In order to use the algorithm in practice, we thus need to know pdf p , and build q such that one can sample easily from q and there exists C such that $p(x) \leq Cq(x)$ for all x , we can compute one such value of C , and C is as small as possible. Note finally that the algorithm can easily be adapted to the cases where p and/or q are known up to a (non necessarily common) constant, see e.g. [18, Th. 4.5].

2) *Revisiting AR sampling as optimal binary classification:* As we shall now see, AR sampling is indeed nothing but a binary classification procedure (see also [21, §6] for an application of this principle).

Starting from the target pdf $p(x)$, we find an easy-to-sample distribution Q and constant $C > 1$ such that $Cq(x)$ envelopes $p(x)$. Since $Cq(x) - p(x)$ is non negative, we write $Cq(x)$ as $p(x)$ plus a positive reminder which, up to a constant, is also a pdf; so enveloping $p(x)$ with $Cq(x)$ is nothing but building the implicit binary mixture pdf (see also fig. 1 below)

$$\underbrace{q(x)}_{\text{proposal}} = \frac{1}{C} \underbrace{p(x)}_{\text{target}} + \underbrace{\left(1 - \frac{1}{C}\right) \frac{q(x) - \frac{1}{C}p(x)}{1 - \frac{1}{C}}}_{\text{reminder}} \quad (5)$$

with a priori probabilities $\frac{1}{C}$ and $1 - \frac{1}{C}$. The first component of the mixture is the target pdf p , and the second one is the law of the rejected samples. The score point of AR consists

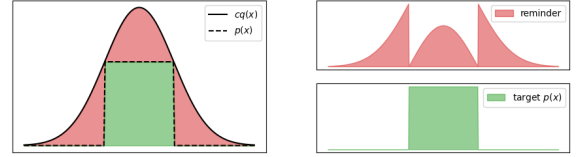


Fig. 1: Enveloping target pdf builds an implicit mixture.

in drawing samples from q *without* needing to sample from its two implicit mixture components (see the r.h.s. of (5)). Accepting (or rejecting) a sample depending on the ratio probability

$$\alpha_{AR}(x) = \frac{p(x)}{Cq(x)} = \frac{\frac{1}{C}p(x)}{\frac{1}{C}p(x) + \left(1 - \frac{1}{C}\right)\frac{q(x) - \frac{1}{C}p(x)}{1 - \frac{1}{C}}} \quad (6)$$

then amounts to classifying the samples with the posterior pdf (compare (6) to (2)).

B. IMH

MCMC algorithms build a Markov chain whose invariant distribution is the target distribution p ; so simulating the chain yields samples asymptotically distributed $\sim p$. The Metropolis-Hastings (MH) algorithm [16] [17] is a particular MCMC method in which the transition is a two-step procedure: given a current state x_t , we propose x^* from $q(\cdot|x_t)$, and then we compute the acceptance probability $\alpha_{MH}(x^*, x_t) = \min\left(1, \frac{p(x^*)q(x_t|x^*)}{p(x_t)q(x^*|x_t)}\right)$. x^* is accepted as the new state x_{t+1} with probability $\alpha_{MH}(x^*, x_t)$; if x^* is rejected then the chain remains in the current state x_t . In practice, $q(\cdot|\cdot)$ plays a crucial role in the performance of the MH algorithm: if not well-tuned, it can lead to a poor exploration of the target distribution.

The IMH algorithm is a simplified version of MH which considers an independent transition. The new point x^* is hence proposed independently of the current state x_t , according to an independent proposal $q(\cdot)$. In this case, the acceptance probability reduces to $\alpha_{IMH}(x^*, x_t) = \min\left(1, \frac{p(x^*)q(x_t)}{p(x_t)q(x^*)}\right)$.

C. IS

Many problems involve computing the expectation of some function f with respect to pdf p : $\mathbb{E}_p[f(x)] = \int f(x)p(x)dx$. In practice the integral can be intractable, so we may need to resort to MC approximations. IS is a technique for reducing the variance of such MC estimates which can be traced back to the 1950's [22] [23] [20, §5.4].

The crude MC estimate reads $\frac{1}{N} \sum_{i=1}^N f(x_i)$, $x_i \stackrel{iid}{\sim} p$. However, on the one hand it is generally difficult to sample directly from p , and on the other hand it can yield a poor estimate when the regions where p is large do not coincide with those where f is large. Rewriting $\mathbb{E}_p[f(x)] = \mathbb{E}_q\left[\frac{p(x)f(x)}{q(x)}\right]$, for some importance distribution q , leads to the IS estimate $\frac{1}{N} \sum_{i=1}^N \frac{p(x_i)}{q(x_i)} f(x_i)$, $x_i \stackrel{iid}{\sim} q$. One can easily show that the pdf which minimizes the variance is, up to a constant, $|f(x)|p(x)$. Even if this optimal importance distribution cannot be used in practice, this tells us that the regions of *importance* are not those where p is large, but rather those where $|f|p$ is large.

¹Code available at github.com/ElouanARGOUARCH/Binary-Classification-Based-Monte-Carlo-Simulation

Note that the IS estimate can be computed only if p and q are known exactly, or known up to a common constant; if this is not the case one can resort to self-normalized IS [24].

Besides being a variance reduction technique, IS can also be seen as a two step sampling procedure for producing samples (approximately) drawn from p , out of samples originally drawn from q . The technique is known as Rubin's SIR mechanism [25] [26] [27] [28, §9.2]: Let $\{x_i\}_{i=1}^N$ be iid samples from $q(x)$ and let $\{\tilde{x}^i\}_{i=1}^M$ be M iid samples from $\sum_{i=1}^N \frac{p(x_i)/q(x_i)}{\sum_{i=1}^N p(x_i)/q(x_i)} \delta_{x_i}(dx)$ (in other words, we draw samples from q , weight each proportionally to $w^u(x_i) = \frac{p(x_i)}{q(x_i)}$, and *resample* iid points from this random discrete distribution). Then $\{\tilde{x}^i\}_{i=1}^M$ become iid. samples from p if $N \rightarrow \infty$.

III. PARAMETRIC CLASSIFIER BY MINIMIZING THE BCE

From now on we consider the setting where λ , p_1 and p_0 are unknown, and we only have the set \mathcal{D} of labeled samples from p_0 and p_1 (resp. with labels $k = 0, 1$), see §I. In this context, for classification purposes we should build a parametric function $r_\phi(x)$ that approximates the posterior pdf. The aim of this section is to show that minimizing a BCE loss indeed yields such a suitable approximation. To see this, first recall the definition of the BCE criterion:

$$\mathcal{L}_{\text{BCE}}(\phi) = - \sum_{i=1}^{N_1} \log(r_\phi(x_i^{(1)})) - \sum_{i=1}^{N_0} \log(1 - r_\phi(x_i^{(0)})), \quad (7)$$

where $r_\phi(x) \triangleq \Pr_\phi(k = 1|x)$ is the probability under model ϕ that the label associated to an observation x is 1. Let $h(x, k)$ be the joint distribution over observations and labels:

$$h(x, k) = \frac{N_k}{N_1 + N_0} \underbrace{p_k(x)}_{h(x|k)}, x \in \mathbb{R}^d, k = 0, 1. \quad (8)$$

Using $r_\phi(x)$, we can build another joint distribution $h_\phi(x, k) = h(x)r_\phi(x)^k(1 - r_\phi(x))^{1-k}$, where $h(x)$ is the x -marginal in (8). The BCE is then, up to additive and multiplicative constants, nothing but an approximate of $D_{\text{KL}}(h(x, k)||h_\phi(x, k)) = \mathbb{E}_{h(x, k)}[\log(h(k|x))] - \mathbb{E}_{h(k, x)}[\log(\Pr_\phi(k|x))]$, where only the last term depends on ϕ . We indeed retrieve the BCE with the MC approximation:

$$\begin{aligned} \mathbb{E}_{h(k, x)}[\log(\Pr_\phi(k|x))] &\stackrel{(8)}{=} \sum_{k=0,1} \frac{N_k}{N_1 + N_0} \mathbb{E}[\log(\Pr_\phi(k|x))] \\ &\approx \frac{1}{N_1 + N_0} \left(\sum_{i=1}^{N_1} \log(r_\phi(x_i^{(1)})) + \sum_{i=1}^{N_0} \log(1 - r_\phi(x_i^{(0)})) \right). \end{aligned}$$

So $\arg \min_\phi D_{\text{KL}}(h(x, k)||h_\phi(x, k)) \approx \arg \min_\phi \mathcal{L}_{\text{BCE}}(\phi)$.

The interest of this interpretation is that, as is well known, a D_{KL} equals zero when the two distributions are equal almost surely. So, if r_ϕ represented any arbitrary function, minimizing $D_{\text{KL}}(h(x, k)||h_\phi(x, k))$ would ensure that $r_\phi(x)^k(1 - r_\phi(x))^{1-k} = h(k|x)$ for all $x \in \mathbb{R}^d$ and $k = 1, 0$, ie that the classifier reaches the target posterior pdf. Of course, in practice, minimizing the BCE does not ensure that this D_{KL} reaches zero. First, we only dispose of a finite number of

labeled observations and minimizing an MC approximation of the D_{KL} does not minimize the D_{KL} itself. Next, the parametric family does not contain $h(k|x)$ in general, in which case we can only ever reach a positive minimum of the D_{KL} . Lastly, standard optimization techniques would only guarantee convergence to a positive local minimum of the D_{KL} . Therefore in practice, minimizing the BCE loss only provides with r_ϕ which approximates the unknown posterior.

IV. USING A BINARY CLASSIFIER FOR (APPROXIMATE) SAMPLING

We now come to the heart of this paper. If p_1 is a pdf of interest in an MC sampling setting, and p_0 a suitable easy-to-sample instrumental distribution - be it the proposal in AR, the independent kernel in IMH, or the importance distribution in IS; then the three sampling algorithms involve the pdf ratio $p_1(x)/p_0(x)$, which is unknown when at least one pdf is intractable. As explained in section III, a parametric binary classifier trained from a set \mathcal{D} of labeled observations computes an approximation of the unknown posterior distribution. However, remember that (3) is equivalent to (4); we thus see that classifiers can also be used for approximating *pdf ratios* of interest, which enables us to propose approximate versions of the sampling algorithms based on this classifier-ratio approximation, *and thus to relax the requirement of tractable pdf, but at the cost of approximate sampling*. Of course, the closer p_0 is to p_1 , the more efficient the sampling algorithms; however, here p_0 is supposed to be given and our problem is not to adjust p_0 from a given p_1 , but rather leverage \mathcal{D} in the case where p_0, p_1 are fixed but unknown pdfs.

Assumptions.

p_1 is the distribution of interest and p_0 a fixed instrumental distribution from which we can propose samples. Ratio $p_1(x)/p_0(x)$ is unknown, but we dispose of the labeled dataset \mathcal{D} , and assume that we can train a binary classification model r_ϕ which minimizes (7).

A. Classifier-based sampling algorithms

A key ingredient for running the algorithms of §II, is the ratio $p_1(x)/p_0(x)$ which appears in $\alpha_{\text{AR}}(x), \alpha_{\text{IMH}}(x, x_t)$ and in $w^u(x)$, so following the idea expressed in (4), we can however make use of a classifier for approximating the unavailable ratio by making the following substitutions:

$$\alpha_{\text{AR}}(x) \leftarrow \frac{1}{\tilde{C}} \frac{r_\phi(x)}{1 - r_\phi(x)} \text{ where } \tilde{C} = \max_{y \in \mathcal{D}} \frac{r_\phi(y)}{1 - r_\phi(y)}; \quad (9)$$

$$\alpha_{\text{IMH}}(x, x_t) \leftarrow \min \left(1, \frac{r_\phi(x)(1 - r_\phi(x_t))}{(1 - r_\phi(x))r_\phi(x_t)} \right); \quad (10)$$

$$w^u(x) \leftarrow \frac{r_\phi(x)}{1 - r_\phi(x)}. \quad (11)$$

Our procedure is summarized by fig. 2: we first train r_ϕ from labeled samples from p_1 and p_0 ; we next use ratio $r_\phi(x)/(1 - r_\phi(x))$ as a surrogate of $p_1(x)/p_0(x)$, which enables us to use the AR, IMH or IS procedure, and thus to turn samples from p_0 into (approximate) samples from p_1 .

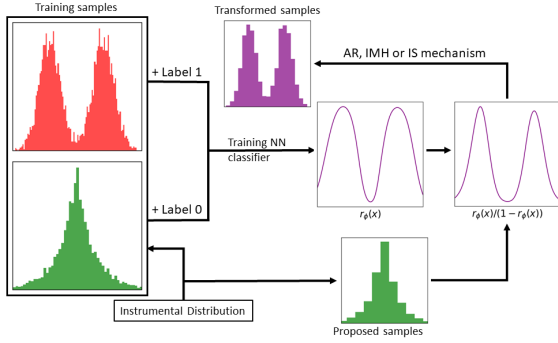


Fig. 2: Summary of the classifier based sampling approach

A main advantage of our approach is that a distribution which is only defined by its sampling procedure and has implicit pdf can be used as instrumental p_0 . Indeed our approach does not require evaluating the pdf p_0 neither during the training of the classifier, nor in the proposed sampling procedures.

B. Illustrating examples

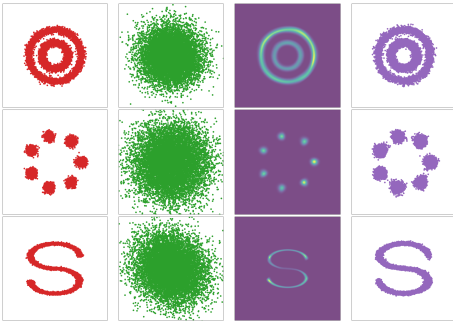


Fig. 3: Density ratio (middle-right) via classification of samples from p_1 (left) and p_0 (middle-left) - approximate samples from p_1 (right) obtained via a ratio based algorithm: AR (top), IMH (middle), IS (bottom)

We illustrate our approach (see fig. 3) on reference 2D examples in order to illustrate the mechanism of (i) obtaining an approximate of the pdf ratio from samples using a feed-forward neural network [29] with 3 hidden layers, 32 hidden units per layers and SiLU activation function that outputs $\text{logit}(r_\phi(x))$ trained according to the BCE criterion; and (ii) sampling from the target distribution via that pdf ratio using the AR, IMH or IS samplers. The instrumental p_0 was set to be a Gaussian with mean and covariance estimated from the samples from p_1 (even though it can be computed, pdf p_0 was not used during the procedure).

C. Probabilistic modelling

So far, we have presented our work as a technique to perform approximate MC sampling; let us now revisit it under the scope of probabilistic modelling. If we rewrite p_1 as $p_1(x) = \frac{p_0(x)(p_1(x)/p_0(x))}{\int p_0(z)(p_1(z)/p_0(z))dz}$, then using (4) amounts to building an approximation p_ϕ of p_1 :

$$p_\phi(x) = \frac{p_0(x)(r_\phi(x)/(1-r_\phi(x)))}{\int p_0(z)(r_\phi(z)/(1-r_\phi(z)))dz}. \quad (12)$$

Our procedure consists in applying the AR, IMH or IS samplers to p_ϕ with proposal p_0 (at least up to the approximation of constant C in the AR case, see (9)). This construction corresponds to a specific energy-based model [30][31][32] with energy function $E_\phi(x) = -\log(p_0(x)) - \text{logit}(r_\phi(x))$. Model p_ϕ inherits the advantages of this energy structure: (i) it can be trained without evaluating the gradient of the numerator of (12) nor of the intractable normalizing constant; (ii) it is structurally compatible with the AR, IMH or IS samplers as sampling from p_ϕ with instrumental p_0 is equivalent to applying the approximate sampling presented in section IV-A. In Fig. 4, we display such an approximation of the distribution of an image. $\text{logit}(r_\phi(x))$ defined via a neural network with 3 hidden layers of size 512 (SiLU activation function), trained according to the BCE criterion, produces a ratio based energy model able to capture details of the target distribution. Samples can effortlessly be obtained via any of the three samplers, with target p_ϕ and instrumental p_0 . Note that computing the unnormalized pdf in (12) (displayed in the middle-right in Fig. 4) indeed requires evaluating pdf p_0 but, again, it is not required for sampling.



Fig. 4: Classifier based energy model: unnormalized pdf (middle-right) and samples (right); obtained from samples (middle-left) from a grayscale image (left) 2D distribution.

V. CONCLUSION

The classical AR, IMH and IS samplers require that both the target p_1 and the easy-to-sample instrumental p_0 are known functions. In practice however, both functions may be either unknown (for p_1) or untractable (for p_0). In this paper we observed that these samplers use p_1 and p_0 only via their ratio $\frac{p_1}{p_0}$ which, in turn, can be approximated by a classifier. We thus showed that one can still approximately sample from p_1 using AR, IMH or IS if functions p_1 and/or p_0 are totally unknown, provided that we dispose of a classifier function, which can be obtained from a set of labeled samples from both distributions. The advantages of our approach are twofold: (i) it is completely pdf-free as compared to standard approaches (neither p_1 nor p_0 needs to be known explicitly); (ii) training reduces to a parametric classification task. From a probabilistic modeling perspective, our approximate samplers coincide with the original ones when applied to some specific energy based approximation of target p_1 which, thanks to its specific structure, can both be trained easily via standard classification, and is structurally compatible with the AR, IMH or IS sampling techniques.

REFERENCES

- [1] C. Giraud, *Introduction to High-Dimensional Statistics*, vol. 139 of *Monographs on Statistics and Applied Probability*. Chapman and Hall/CRC, Boca Raton, FL, 2014.
- [2] X. Nguyen, M. J. Wainwright, and M. I. Jordan, “Estimating divergence functionals and the likelihood ratio by convex risk minimization,” *IEEE Tr. on Inf. Theory*, vol. 56, no. 11, pp. 5847–5861, 2010.
- [3] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, “Generative adversarial networks,” *Communications of the ACM*, vol. 63, no. 11, pp. 139–144, 2020.
- [4] A. Creswell and A. A. Bharath, “Denosing adversarial autoencoders,” *IEEE transactions on neural networks and learning systems*, vol. 30, no. 4, pp. 968–984, 2018.
- [5] M. I. Belghazi, A. Baratin, S. Rajeshwar, S. Ozair, Y. Bengio, A. Courville, and D. Hjelm, “Mutual information neural estimation,” in *International conference on machine learning*, pp. 531–540, PMLR, 2018.
- [6] A. Gretton, K. M. Borgwardt, M. J. Rasch, B. Schölkopf, and A. Smola, “A kernel two-sample test,” *J. of Machine Learning Research*, vol. 13, no. 1, pp. 723–773, 2012.
- [7] C. Durkan, I. Murray, and G. Papamakarios, “On contrastive learning for likelihood-free inference,” in *Int. conf. on machine learning*, pp. 2771–2781, PMLR, 2020.
- [8] J. Hermans, V. Begy, and G. Louppe, “Likelihood-free MCMC with amortized approximate ratio estimators,” in *Int. conf. on Mach. learn.*, pp. 4239–4248, PMLR, 2020.
- [9] O. Thomas, R. Dutta, J. Corander, S. Kaski, and M. U. Gutmann, “Likelihood-free inference by ratio estimation,” *Bayesian Analysis*, vol. 17, no. 1, pp. 1–31, 2022.
- [10] F. Llorente, L. Martino, J. Read, and D. Delgado, “A survey of monte carlo methods for noisy and costly densities with application to reinforcement learning,” *CoRR*, vol. abs/2108.00490, 2021.
- [11] S. A. Sisson, Y. Fan, and M. Beaumont, *Handbook of approximate Bayesian computation*. CRC Press, 2018.
- [12] D. P. Kingma and M. Welling, “Stochastic gradient VB and the variational auto-encoder,” in *Second Int. Conf. on Learning Representations, ICLR*, vol. 19, p. 121, 2014.
- [13] Y. Song, J. Sohl-Dickstein, D. P. Kingma, A. Kumar, S. Ermon, and B. Poole, “Score-based generative modeling through stochastic differential equations,” *arXiv preprint arXiv:2011.13456*, 2020.
- [14] N. Bartoli and P. del Moral, *Simulation et algorithmes stochastiques*. Cépaduès éditions, 2001.
- [15] J. Gentle, *Random Number Generation and Monte Carlo Methods*. Springer, 2004.
- [16] C. Robert and G. Casella, *Monte Carlo Statistical Methods (Springer Texts in Statistics)*. Secaucus, NJ, USA: Springer-Verlag New York, Inc., 2005.
- [17] S. Chib and E. Greenberg, “Understanding the Metropolis-Hastings algorithm,” *The American Statistician*, vol. 49, no. 4, pp. 327–335, 1995.
- [18] A. B. Owen, *Monte Carlo theory, methods and examples*. 2013.
- [19] L. Martino, D. Luengo, and J. Míguez, *Independent random sampling methods*. Springer, 2018.
- [20] J. M. Hammersley and D. C. Handscomb, *Monte Carlo Methods*. London: Methuen and Co., 1964.
- [21] M. Castella, S. Rafi, P. Comon, and W. Pieczynski, “Separation of instantaneous mixtures of dependent sources using classical ICA methods,” *EURASIP Journal on Advances in Signal Processing*, p. n/c, Mar. 2013.
- [22] H. Kahn and A. W. Marshall, “Methods of reducing sample size in Monte Carlo computations,” *J. of the Op. Res. Soc. of Amer.*, vol. 1, no. 5, pp. 263–278, 1953.
- [23] A. W. Marshall, “The use of multi-stage sampling schemes in Monte Carlo computations,” in *Symposium on Monte Carlo Methods* (M. Meyer, ed.), (New York), pp. 123–140, 1956.
- [24] J. Geweke, “Bayesian inference in econometric models using Monte Carlo integration,” *Econometrica*, vol. 57, no. 6, pp. 1317–39, 1989.
- [25] D. B. Rubin, “Using the SIR algorithm to simulate posterior distributions,” in *Bayesian Statistics III* (M. H. Bernardo, K. M. Degroot, D. V. Lindley, and A. F. M. Smith, eds.), Oxford: Oxford University Press, 1988.
- [26] A. E. Gelfand and A. F. M. Smith, “Sampling based approaches to calculating marginal densities,” *J. of the Amer. Stat. Ass.*, vol. 85, no. 410, pp. 398–409, 1990.
- [27] A. F. M. Smith and A. E. Gelfand, “Bayesian statistics without tears : a sampling-resampling perspective,” *The American Statistician*, vol. 46, no. 2, pp. 84–87, 1992.
- [28] O. Cappé, É. Moulines, and T. Rydén, *Inference in Hidden Markov Models*. Springer-Verlag, 2005.
- [29] G. Bebis and M. Georgiopoulos, “Feed-forward neural networks,” *IEEE Potentials*, vol. 13, no. 4, pp. 27–31, 1994.
- [30] S. Zhai, Y. Cheng, W. Lu, and Z. Zhang, “Deep structured energy based models for anomaly detection,” in *Int. Conf. on Machine learning*, pp. 1100–1109, PMLR, 2016.
- [31] M. A. Carreira-Perpinan and G. Hinton, “On contrastive divergence learning,” in *International workshop on artificial intelligence and statistics*, pp. 33–40, PMLR, 2005.
- [32] G. E. Hinton, “A practical guide to training restricted boltzmann machines,” in *Neural Networks: Tricks of the Trade: Second Edition*, pp. 599–619, Springer, 2012.