EVALUATING AND TUNING PRECISION AND RECALL FOR GENERATIVE MODELS DATA SCIENCE LAB

Alexandre Vérine

Junior AI Fellow of University PSL, Centre des Données, ENS-PSL

October 29, 2025



CONTEXT

In the past few years, generative models have made significant progress in various domains of application.




eg: Code Completion

1 / 81


CONTEXT

In the past few years, generative models have made significant progress in various domains of application.



CONTEXT


In the past few years, generative models have made significant progress in various domains of application.



PR FOR GENERATIVE MODELS

CONTEXT

In the past few years, generative models have made significant progress in various domains of application.



THE VARIOUS PERFORMANCES OF GENERATIVE MODELS **MOTIVATION**

As the generation becomes better, the evaluation becomes more challenging.

DALL-E 2 (2023)





Midjourney v5 (2023)

Prompt: *A dog playing with a child.*

PR FOR GENERATIVE MODELS 2 / 81

THE VARIOUS PERFORMANCES OF GENERATIVE MODELS MOTIVATION

As the generation becomes better, the evaluation becomes more challenging.

DALL-E 2 (2023)



Midjourney v5 (2023)



Prompt: A dog playing with a child.

PR FOR GENERATIVE MODELS 2 / 81

THE VARIOUS PERFORMANCES OF GENERATIVE MODELS

TRADITIONAL METRICS

Two metrics are often used to evaluate generative models:

- ► Inception Score (IS) [24]
- ► Fréchet Inception Distance (FID) [15]

Question: Can they separately evaluate quality and diversity of generative models?

INCEPTION SCORE

Assume that we have a pre-trained classifier Inception-v3:

Definition 1.1 (Inception Score)

Let denote $\mathbb{P}(Y|x)$ be the conditional class distribution of an image x given by the Inception-v3 model and $\mathbb{P}(Y)$ the class distribution in dataset sampled from P. The Inception Score is defined as:

$$\mathrm{IS}(\widehat{P}) = \exp\left(\mathbb{E}_{\boldsymbol{x} \sim \widehat{P}}\left[\mathcal{D}_{\mathrm{KL}}(\mathbb{P}(Y|\boldsymbol{x})||\mathbb{P}(Y))\right]\right)$$

where \mathcal{D}_{KL} is the Kullback-Leibler divergence.

With *H* being the entropy function, IS can be reformulated as:

$$\log\left(\mathrm{IS}(\widehat{P})\right) = H\left(\mathbb{E}_{\boldsymbol{x} \sim \widehat{P}}\left[\mathbb{P}(Y|\boldsymbol{x})\right]\right) - \mathbb{E}_{\boldsymbol{x} \sim \widehat{P}}\left[H\left(\mathbb{P}(Y|\boldsymbol{x})\right)\right].$$

PR FOR GENERATIVE MODELS

4/81

INCEPTION SCORE

$$\log\left(\mathrm{IS}(\widehat{P})\right) = H\left(\mathbb{E}_{\boldsymbol{x} \sim \widehat{P}}\left[\mathbb{P}(Y|\boldsymbol{x})\right]\right) - \mathbb{E}_{\boldsymbol{x} \sim \widehat{P}}\left[H\left(\mathbb{P}(Y|\boldsymbol{x})\right)\right].$$

Inception Score is maximized:

- ▶ if for every $x \sim \widehat{P}$, $H_Y(\mathbb{P}(Y|x))$ is minimized. This suggests that the classification model is highly confident in predicting a singular label per image, which implies that the images are clearly recognizable and of high quality.
- if the entropy of $\mathbb{E}_{x \sim \widehat{p}}[\mathbb{P}(Y|x)]$ is maximized. The label predictions must be uniformly distributed over all possible labels. This indicates that the generative model generates heterogeneous labels, thereby ensuring diversity.




Figure. Source: medium.com

PR FOR GENERATIVE MODELS 5 / 81

IS FALLS SHORT

However, in practical applications, the Inception Score has revealed several shortcomings [11, 6, 3, 23]:

- ► The IS is not sensitive to measuring intraclass diversity. In particular, if the model generates only one high-quality image per class, the IS will be high.
- ▶ The IS does not directly measure the realism of individual images. If the images are peripherally saturating, noisy, or distorted, and if the classification confidence is high, they are evaluated as realistic.
- ► The IS is biased toward the classes represented in ImageNet, since the Inception-v3 model is trained on this dataset. For instance, if the goal is to evaluate models that generate faces such as the CelebA dataset, IS will favor models generating faces with glasses, sunglasses, or cowboy hat, since these attributes are ImageNet classes.
- ▶ IS is not necessarily optimal when the generated images are identical to the dataset. Since it does not directly compare the generated distribution with the true data distribution, Barratt and Sharma [3] shows that IS(*P*) is not always optimal.

FRÉCHET INCEPTION DISTANCE

The Fréchet Inception Distance (FID) compares the statistics of generated samples to real samples based on the features extracted using the output of the last pooling layer of Inception-v3, which is in dimension 2048, we denote it $\phi: \mathbb{R}^d \to \mathbb{R}^m$. We take two sets of samples $\{x_1^{\mathrm{real}}, \dots, x_N^{\mathrm{real}}\}$ drawn from both P and $\{x_1^{\mathrm{fake}}, \dots, x_N^{\mathrm{fake}}\}$ drawn from \widehat{P} . We consider the empirical mean and covariance of the latent representation of the samples:

$$\boldsymbol{\mu} = \frac{1}{N} \sum_{i=1}^{N} \phi(\mathbf{x}_i^{\text{real}}) \quad \text{and} \quad \boldsymbol{\Sigma} = \frac{1}{N-1} \sum_{i=1}^{N} \left(\phi(\mathbf{x}_i^{\text{real}}) - \boldsymbol{\mu} \right) \left(\phi(\mathbf{x}_i^{\text{real}}) - \boldsymbol{\mu} \right)^{\top},$$

$$\hat{\boldsymbol{\mu}} = \frac{1}{N} \sum_{i=1}^{N} \phi(\mathbf{x}_i^{\text{fake}}) \quad \text{and} \quad \hat{\boldsymbol{\Sigma}} = \frac{1}{N-1} \sum_{i=1}^{N} \left(\phi(\mathbf{x}_i^{\text{fake}}) - \hat{\boldsymbol{\mu}} \right) \left(\phi(\mathbf{x}_i^{\text{fake}}) - \hat{\boldsymbol{\mu}} \right)^{\top}$$

Definition 1.2 (Fréchet Inception Distance)

The Fréchet Inception Distance is defined as the Wasserstein-2 distance between two multivariate Gaussians $\mathcal{N}(\mu, \Sigma)$ and $\mathcal{N}(\hat{\mu}, \hat{\Sigma})$:

$$\text{FID} = \|\boldsymbol{\mu} - \hat{\boldsymbol{\mu}}\|^2 + \text{Tr}\left(\boldsymbol{\Sigma} + \widehat{\boldsymbol{\Sigma}} - 2(\boldsymbol{\Sigma}\widehat{\boldsymbol{\Sigma}})^{1/2}\right)$$

PR FOR GENERATIVE MODELS


FID FALLS SHORT

FID is still the most widely used metric to evaluate generative models, but it has several limitations [15, 6, 10]:

- ▶ The FID compares statistical summaries (mean and covariance) of the latent distributions of Inception, a discriminative model. Therefore, it may not capture all aspects of image quality, such as texture and local structure, that are perceptible to humans.
- ► FID assumes that the latent representations follow a Gaussian distribution, which may not hold true in practice. This assumption can lead to inaccurate assessments of the similarity between real and generated data distributions.
- ► FID does not distinguish between different types of error in image generation. For example, it treats a noisy object the same as a completely wrong object being generated, which may not align with human judgment.
- ▶ Similarly to IS, FID accounts for both quality and diversity, but without a clear trade-off.

FID FALLS SHORT

For example, Kynkäänniemi et al. [19] highlights the limitations of FID with samples drawn from StyleGAN:



OUTLINE

In this presentation, we discuss on evaluating, optimizing and improving quality and diversity of generative models:

- 1. **Evaluating:** How can we assess quality and diversity independently in Generative Models?
 - 2. **Tuning:** Can we optimize a specific trade-off between quality and diversity?
- 3. **Improving:** How can we improve the quality and diversity of a pre-trained generative models?


PR FOR GENERATIVE MODELS

CONTEXT AND MOTIVATION CONTEXT


Evaluating:

How can we assess quality and diversity independtly in Generative Models?


FRAMEWORK




▶ Assumption: There is an unknown *target distribution P* in $\mathcal{X} \subset \mathbb{R}^d$.




▶ Assumption: There is an unknown *target distribution P* in $\mathcal{X} \subset \mathbb{R}^d$.




- ▶ Assumption: There is an unknown *target distribution P* in $\mathcal{X} \subset \mathbb{R}^d$.
- ▶ Goal: Learn a *parameterized distribution* \widehat{P} that approximate P:





- ▶ Assumption: There is an unknown *target distribution P* in $\mathcal{X} \subset \mathbb{R}^d$.
- ▶ Goal: Learn a *parameterized distribution* \widehat{P} that approximate P:
 - 1. Consider a distribution Q in a latent space $\mathcal{X} \subset \mathbb{R}^m$, usually $\mathcal{N}(0, I_m)$.




- ▶ Assumption: There is an unknown *target distribution P* in $\mathcal{X} \subset \mathbb{R}^d$.
- ▶ Goal: Learn a *parameterized distribution* \widehat{P}_G that approximate P:
 - 1. Consider a distribution Q in a latent space $\mathcal{X} \subset \mathbb{R}^m$, usually $\mathcal{N}(0, I_m)$.
 - 2. Take a generator model G represented by a neural network. Take $\widehat{P}_G = G \# Q$.

- ▶ Assumption: There is an unknown *target distribution* P in $\mathcal{X} \subset \mathbb{R}^d$.
- ▶ Goal: Learn a parameterized distribution \widehat{P}_G that approximate P:
 - 1. Consider a distribution Q in a latent space $\mathcal{X} \subset \mathbb{R}^m$, usually $\mathcal{N}(0, I_m)$.
 - 2. Take a generator model G represented by a neural network. Take $\widehat{P}_G = G \# Q$.
 - 3. Compute G^{opt} that minimize a dissimilarity measure D between P and \widehat{P}_G :

$$G^{\text{opt}} = \operatorname*{argmin}_{G} D(P, \widehat{P}_{G})$$

FRAMEWORK




- ▶ Assumption: There is an unknown *target distribution P* in $\mathcal{X} \subset \mathbb{R}^d$.
- ▶ Goal: Learn a *parameterized distribution* \widehat{P} that approximate P:
 - 1. Consider a distribution Q in a latent space $\mathcal{X} \subset \mathbb{R}^m$, usually $\mathcal{N}(0, I_m)$.
 - 2. Take a generator model G represented by a neural network. Take $\widehat{P} = G \# Q$.
 - 3. Compute G^{opt} that minimize a dissimilarity measure D between P and \widehat{P} :


$$G^{\text{opt}} = \operatorname*{argmin}_{G} D(P, \widehat{P})$$

PR FOR GENERATIVE MODELS


IN PRACTICE

$$\widehat{P} \neq P$$






IN PRACTICE



Low Diversity

IN PRACTICE



Low Diversity

Low Quality

METRICS TO EVALUATE QUALITY AND DIVERSITY

To assess models, we use the notion of Precision and Recall, inspired from Information Retrieval:

Quality

Diversity

METRICS TO EVALUATE QUALITY AND DIVERSITY

To assess models, we use the notion of Precision and Recall, inspired from Information Retrieval:

What proportion of generated samples are realistic?

PRECISION AND RECALL FOR GENERATIVE MODELS


METRICS TO EVALUATE QUALITY AND DIVERSITY


To assess models, we use the notion of Precision and Recall, inspired from Information Retrieval:



What proportion of generated samples are realistic? What proportion of real samples can be generated?

FOR FINITE SUPPORT





FOR FINITE SUPPORT

Definition 1.3 (Support-Based Precision and Recall - [19].)


For any distributions $P \in \mathcal{P}(\mathcal{X})$ and $\widehat{P} \in \mathcal{P}(\mathcal{X})$, we say that the distribution P has precision $\bar{\alpha}$ at recall $\bar{\beta}$ with respect to \widehat{P} if


$$\bar{\alpha} := \widehat{P}(\operatorname{Supp}(P))$$
 and $\bar{\beta} := P(\operatorname{Supp}(\widehat{P})).$

FOR FINITE SUPPORT

Precision for finite support is the proportion of generated data that lies on the support of the real data:

$$\bar{\alpha} = \widehat{P}(\operatorname{Supp}(P)).$$







FOR FINITE SUPPORT

Precision for finite support is the proportion of generated data that lies on the support of the real data:

$$\bar{\alpha} = \widehat{P}(\operatorname{Supp}(P)).$$







FOR FINITE SUPPORT

Recall for finite support is the proportion of the support of the real data that is covered by the generated data:

$$\bar{\beta} = P(\operatorname{Supp}(\widehat{P})).$$







FOR FINITE SUPPORT


Recall for finite support is the proportion of the support of the real data that is covered by the generated data:

$$\bar{\beta} = P(\operatorname{Supp}(\widehat{P})).$$







ESTIMATING THE PRECISION AND RECALL




ESTIMATING THE PRECISION AND RECALL



ESTIMATING THE PRECISION AND RECALL







19 / 81

PR FOR GENERATIVE MODELS

ESTIMATING THE PRECISION AND RECALL




ESTIMATING THE PRECISION AND RECALL



PR FOR GENERATIVE MODELS

FOR LLMS



IN PRACTICE

MNIST Dataset [32]

```
881033484

364255647

9886556

98865

98845

95445

95445

95445

9631

9631

9631
```

PR FOR GENERATIVE MODELS 21 / 81

IN PRACTICE

Low Precision High Recall

\$145007428 1120915955 0681164441 7199429173 7309238459 6406519169 9657328511 8836908516 8886128010

Precision: 0.54 Recall: 0.91


MNIST Dataset [32]

881033799 3647551711 9884755111 988455411 954455799 96756799 967567 High Precision Low Recall

Precision: 0.80 Recall: 0.70

FOR LLMS


On open-ended generation, the quality and diversity of LLMs can also be evaluated using Precision and Recall: Bronnec et al. [8]




PR FOR GENERATIVE MODELS 22 / 81


FOR LLMS


We can also evaluate the quality and diversity of LLMs on Chatbot open-ended generation. We can for instance check the impact of In-Context examples on the quality and diversity of the generated text. For instance on Wikipedia Biographies generation:







FOR INIFINITE SUPPORT





FOR INIFINITE SUPPORT





Both distributions have **perfect** Precision *and* Recall.

DEFINITION

Definition 1.4 (PR-Curve for Generative Models - Sajjadi et al. [23], Simon et al. [25])

Let $P, \widehat{P} \in \mathcal{P}(\mathcal{X})$ be two distributions such that $P, \widehat{P} \ll \mu$. The PR-Curve is the set $PRD(P, \widehat{P})$ defined as:

$$PRD(P, \widehat{P}) = \{(\alpha_{\lambda}, \beta_{\lambda}) \mid \lambda \in [0, \infty]\}$$

with:

$$\alpha_{\lambda} = \int_{\mathcal{X}} \min \left(\lambda p(\mathbf{x}), \widehat{p}(\mathbf{x}) \right) d\mu(\mathbf{x}) \quad and \quad \beta_{\lambda} = \int_{\mathcal{X}} \min \left(p(\mathbf{x}), \widehat{p}(\mathbf{x}) / \lambda \right) d\mu(\mathbf{x}).$$

DEFINITION

For the Precision, λp is compared to \hat{p} for different threshold $\lambda \in [0, +\infty]$:

$$\alpha_{\lambda} = \int_{\mathcal{X}} \min\left(\lambda p(\mathbf{x}), \widehat{p}(\mathbf{x})\right) d\mu(\mathbf{x}) \tag{1}$$

DEFINITION

For the Precision, λp is compared to \hat{p} for different threshold $\lambda \in [0, +\infty]$:

$$\alpha_{\lambda} = \int_{\mathcal{X}} \min\left(\lambda p(\mathbf{x}), \widehat{p}(\mathbf{x})\right) d\mu(\mathbf{x}) \tag{1}$$

DEFINITION

For the Recall, p is compared to \hat{p}/λ for different threshold $\lambda \in [0, +\infty]$:

$$\beta_{\lambda} = \int_{\mathcal{X}} \min \left(p(\mathbf{x}), \widehat{p}(\mathbf{x}) / \lambda \right) d\mu(\mathbf{x}) \tag{2}$$

DEFINITION

For the Recall, p is compared to \hat{p}/λ for different threshold $\lambda \in [0, +\infty]$:

$$\beta_{\lambda} = \int_{\mathcal{X}} \min \left(p(\mathbf{x}), \widehat{p}(\mathbf{x}) / \lambda \right) d\mu(\mathbf{x}) \tag{2}$$

EXAMPLES




Figure. Learning distribution with low recall and high precision.

EXAMPLES




Figure. Learning distribution with high recall and low precision.

EXAMPLES




Figure. Learning distribution with low recall and low precision.

EXAMPLES




Figure. Learning distribution with high recall and high precision.

28 / 81

PR-CURVE AND SUPPORT-BASED PRECISION AND RECALL

RELATION

The PR-Curve is a generalization of the Precision and Recall for finite support:


Theorem 1.5 (Support-based and PR-Curves - Siry et al. [26])

Let $P, \widehat{P} \in \mathcal{P}(\mathcal{X})$ be two distributions. Then, the support-based Precision and Recall $(\bar{\alpha}, \bar{\beta})$ are related to the PR-Curve values $PRD(P, \widehat{P})$ for $\lambda = 0$ and $\lambda = \infty$:

$$\bar{\alpha} = \max_{\lambda} \alpha_{\lambda} = \alpha_{\infty} \quad and \quad \bar{\beta} = \max_{\lambda} \beta_{\lambda} = \beta_{0}.$$


PR-CURVE AND SUPPORT-BASED PRECISION AND RECALL

RELATION



PR-CURVE AND SUPPORT-BASED PRECISION AND RECALL

RELATION




IN PRACTICE

Low Precision High Recall



Precision: 0.54 Recall: 0.91

PR-Curves



High Precision Low Recall



Precision: 0.80 Recall: 0.70

IN NLP




Figure. PR-Curve for distributions journal articles: AG News.

ON THE PLATFORM

Metrics used to evaluate your models are:

- ► FID
- ► Precision (for finite support)
- ► Recall (for finite support)
- ▶ (Obviously) the visual inspection of the generated samples.

PR FOR GENERATIVE MODELS

Tuning:

How can we tune a model to a specific trade-off between Precision and Recall?

TRUNCATION

Hard Trunctation Karras et al. [17]

Soft Trunctation Kingma and Dhariwal [18]

HARD TRUNCATION



Figure. From left to right: $\psi = 0.0$, $\psi = 0.3$ $\psi = 0.7$ $\psi = 1.0$.




Figure. Source: [19]

SOFT TRUNCATION




Figure. Soft-Truncation on BigGAN. Source:[7].

36 / 81

TEMPERATURE SCALING

In LLMs, we can tune the temperature. Let's assume that the model is outputting a categorical distribution with probability $\widehat{P}(x_l|x_{< i})$ for a token x_l with a given context $x_{< i}$. We can tune the temperature t > 0 as follows:

$$\widehat{P}^t(x_l|\mathbf{x}_{< i}) = \frac{\widehat{P}(x_l|\mathbf{x}_{< i})^{1/t}}{\sum_{x_i \in \mathcal{V}^L} \widehat{P}(x_i|\mathbf{x}_{< i})^{1/t}}.$$

TEMPERATURE SCALING

In LLMs, we can tune the temperature. Let's assume that the model is outputting a categorical distribution with probability $\widehat{P}(x_l|x_{< i})$ for a token x_l with a given context $x_{< i}$. We can tune the temperature t > 0 as follows:

$$\widehat{P}^t(x_l|\mathbf{x}_{< i}) = \frac{\widehat{P}(x_l|\mathbf{x}_{< i})^{1/t}}{\sum_{x_i \in \mathcal{V}^L} \widehat{P}(x_i|\mathbf{x}_{< i})^{1/t}}.$$





Figure. Effect on the PR-Curve for different temperatures for Coda Llama 2. Source: Verine et al. [30].

TRAINING A GENERATIVE MODEL

IN GENERAL

Traditionally, the goal is to minimize a dissimilarity mea*sure* between the target distribution *P* and the learned distribution \widehat{P} :

$$\min_{G} D(P, \widehat{P}_{G}) \tag{3}$$




Distribution Space

TRAINING A GENERATIVE MODEL

WITH *f*-DIVERGENCES

Traditionally, the goal is to minimize an f-divergence between the target distribution P and the learned distribution \widehat{P} :

$$\min_{G} \mathcal{D}_{f}(P \| \widehat{P}_{G}) \tag{3}$$



Distribution Space

f-DIVERGENCES

DEFINITION

Definition 2.1 (*f*-divergences)

For any two probability distributions P and \widehat{P} in $\mathcal{P}(\mathcal{X})$ such that $P, \widehat{P} \ll \mu$. Let p and \widehat{p} be the Radon-Nikodym densities of P and \widehat{P} with respect to μ , respectively. Let f be any convex lower semi-continuous function $f:[0,\infty]\to]-\infty,+\infty]$ such that f(1)=0, the f-divergence between P and \widehat{P} is

$$\mathcal{D}_{f}(P||\widehat{P}) = \int_{\mathcal{X}} \widehat{p}(\mathbf{x}) f\left(\frac{p(\mathbf{x})}{\widehat{p}(\mathbf{x})}\right) d\mu(\mathbf{x}). \tag{4}$$

f-DIVERGENCES

DEFINITION

Definition 2.1 (*f*-divergences)

For any two probability distributions P and \widehat{P} in $\mathcal{P}(\mathcal{X})$ such that $P, \widehat{P} \ll \mu$. Let p and \widehat{p} be the Radon-Nikodym densities of P and \widehat{P} with respect to μ , respectively. Let f be any convex lower semi-continuous function $f:[0,\infty]\to]-\infty,+\infty]$ such that f(1)=0, the f-divergence between P and \widehat{P} is

$$\mathcal{D}_f(P||\widehat{P}) = \int_{\mathcal{X}} \widehat{p}(\mathbf{x}) f\left(\frac{p(\mathbf{x})}{\widehat{p}(\mathbf{x})}\right) d\mu(\mathbf{x}). \tag{4}$$

Usual divergences are *f*-divergences:

- ► Kullback-Leibler (KL),
- ► Reverse Kullback-Leibler (rKL),
- ► Jensen-Shannon (JS),
- ► Total Variation (TV),
- ightharpoonup α -divergences.

ESTIMATING *f*-DIVERGENCES

DUAL VARIATIONAL FORM

f-divergences are *hardly tractable* in practice. However, they can be approximated by a dual approximation.

ESTIMATING *f*-DIVERGENCES

DUAL VARIATIONAL FORM

f-divergences are *hardly tractable* in practice. However, they can be approximated by a dual approximation.

- ▶ $f^*(t) = \sup_{u \in \mathbb{R}} \{tu f(u)\}$ be the Fenchel conjugate of f.
- $ightharpoonup \mathcal{T}$ be the set of all measurable functions $\mathcal{X} \to \mathbb{R}$.

ESTIMATING *f*-DIVERGENCES

DUAL VARIATIONAL FORM

f-divergences are *hardly tractable* in practice. However, they can be approximated by a dual approximation.

- $f^*(t) = \sup_{u \in \mathbb{R}} \{tu f(u)\}$ be the Fenchel conjugate of f.
- $ightharpoonup \mathcal{T}$ be the set of all measurable functions $\mathcal{X} \to \mathbb{R}$.

Theorem 2.2 (Dual variational form of an f-divergence- Nguyen et al. [20])

Let $P, \widehat{P} \in \mathcal{P}(\mathcal{X})$ two distributions such that P is absolutely continuous with respect to \widehat{P} and f a suitable generator function. The f-divergence between P and \widehat{P} admits a dual variational form:

$$\mathcal{D}_{f}(P||\widehat{P}) = \sup_{T \in \mathcal{T}} \left(\mathbb{E}_{\mathbf{x} \sim P} \left[T(\mathbf{x}) \right] - \mathbb{E}_{\mathbf{x} \sim \widehat{P}} \left[f^{*}(T(\mathbf{x})) \right] \right). \tag{5}$$

We use $T^{\text{opt}} \in \mathcal{T}$ to denote the function that achieves the supremum.

TRAINING A GENERATIVE MODEL WITH *f*-DIVERGENCES

USING THE DUAL VARIATIONAL FORM

$$\min_{G} \max_{T} \underbrace{\mathbb{E}_{x \sim P} \left[T(x) \right] - \mathbb{E}_{x \sim \widehat{P}_{G}} \left[f^{*}(T(x)) \right]}_{\mathcal{D}_{f,T}^{\text{dual}}} \tag{6}$$

Training a Generative Model with f-divergences

USING THE DUAL VARIATIONAL FORM

$$\min_{G} \max_{T} \underbrace{\mathbb{E}_{x \sim P} \left[T(x) \right] - \mathbb{E}_{x \sim \widehat{P}_{G}} \left[f^{*}(T(x)) \right]}_{\mathcal{D}_{f,T}^{\text{dual}}} \tag{6}$$

- ▶ The discriminator *T* is trained *to estimate* the divergence.
- ▶ The generator *G* is trained *to minimize* the divergence.

Training a Generative Model with f-divergences

USING THE DUAL VARIATIONAL FORM

$$\min_{G} \max_{T} \underbrace{\mathbb{E}_{x \sim P} \left[T(x) \right] - \mathbb{E}_{x \sim \widehat{P}_{G}} \left[f^{*}(T(x)) \right]}_{\mathcal{D}_{f,T}^{\text{dual}}} \tag{6}$$

- ▶ The discriminator *T* is trained *to estimate* the divergence.
- ▶ The generator *G* is trained *to minimize* the divergence.

TRAINING A GENERATIVE MODEL WITH *f*-DIVERGENCES

USING THE DUAL VARIATIONAL FORM

$$\min_{G} \max_{T} \mathbb{E}_{\boldsymbol{x} \sim P} \left[\log \left(D(\boldsymbol{x}) \right) \right] - \mathbb{E}_{\boldsymbol{x} \sim \widehat{P}_{G}} \left[f^{*} \left(\log(D(\boldsymbol{x})) \right) \right]$$
 (6)

- ▶ The discriminator *T* is trained *to estimate* the divergence.
- ▶ The generator *G* is trained *to minimize* the divergence.
- ▶ With $T(x) = \log(D(x))$ with $D(x) \in [0, 1]$.

TRAINING A GENERATIVE MODEL WITH *f*-DIVERGENCES

USING THE DUAL VARIATIONAL FORM

By doing so, we can rewrite the optimization problem as:

$$\min_{G} \max_{T} \mathbb{E}_{\boldsymbol{x} \sim P} \left[\log \left(D(\boldsymbol{x}) \right) \right] + \mathbb{E}_{\boldsymbol{x} \sim \widehat{P}_{G}} \left[\log \left(1 - D(\boldsymbol{x}) \right) \right) \right] \tag{6}$$

- ▶ The discriminator *T* is trained *to estimate* the divergence.
- ▶ The generator *G* is trained *to minimize* the divergence.
- ▶ With $T(x) = \log(D(x))$ with $D(x) \in [0, 1]$.
- $f^*(t) = f_{\text{IS}}^*(t) = -\log(1 \exp(t))$ for the Jensen-Shannon divergence.

We recover the original GAN framework.

Training a Generative Model with f-divergences

USING THE DUAL VARIATIONAL FORM


$$\min_{G} \max_{T} \underbrace{\mathbb{E}_{\mathbf{x} \sim P} \left[T(\mathbf{x}) \right] - \mathbb{E}_{\mathbf{x} \sim \widehat{P}_{G}} \left[f^{*}(T(\mathbf{x})) \right]}_{\mathcal{D}_{f,T}^{\text{dual}}} \tag{6}$$

- ▶ The discriminator *T* is trained *to estimate* the divergence.
- ▶ The generator *G* is trained *to minimize* the divergence.
- ► Generative Adversarial Networks [12] for the Jensen-Shannon divergence.
- Extended to other *f*-divergences by Nowozin et al. [21].
- ▶ Extend to other generative models such as Normalizing Flows by Grover et al. [13].

Effect of the f-divergence on the learned distribution

All *f*-divergences are not equal:

$$\mathcal{D}_{f}(P||\widehat{P}) = \mathbb{E}_{\mathbf{x} \sim \widehat{P}} \left[f\left(\frac{p(\mathbf{x})}{\widehat{p}(\mathbf{x})}\right) \right]$$




PR FOR GENERATIVE MODELS

42 / 81

Effect of the f-divergence on the learned distribution


All *f*-divergences are not equal:



Penalizing low values of $\frac{p(x)}{\widehat{p}(x)}$

Effect of the f-divergence on the learned distribution


All *f*-divergences are not equal:



Penalizing low values of $\frac{p(x)}{\widehat{p}(x)}$ Favors high precision

Examples of f-divergence minimization

Examples of f-divergence minimization



TUNING PRECISION AND RECALL IN GENERATIVE MODELS CONTRIBUTIONS

Can we optimize a specific trade-off between Precision and Recall?

TUNING PRECISION AND RECALL IN GENERATIVE MODELS

CONTRIBUTIONS

Can we optimize a specific trade-off between Precision and Recall?


▶ What is the relation between the Precision-Recall curve and *f*-divergences?

DEFINITION

Definition 2.3 (PR-Divergence generator function f_{λ})

Given a trade-off parameter $\lambda \in [0, +\infty]$, we define the generator function $f_{\lambda} : [0, +\infty] \to]-\infty, +\infty]$ given by

$$f_{\lambda}(u) = \begin{cases} \max(\lambda u, 1) - \max(\lambda, 1) & \text{for } \lambda \in [0, +\infty[, \\ \mathbb{1}_{\{u=0\}} & \text{for } \lambda = +\infty. \end{cases}$$
 (7)



PR FOR GENERATIVE MODELS

PROPERTIES

Proposition 2.4 (PR-Divergence)

For any distributions $P, \widehat{P} \in \mathcal{P}(\mathcal{X})$ such that $P, \widehat{P} \ll \mu$, then for any $\lambda \in [0, +\infty]$ the PR-Divergence defined as

$$\mathcal{D}_{\lambda\text{-PR}}(P\|\widehat{P}) = \int_{\mathcal{X}} \widehat{p}(\mathbf{x}) f_{\lambda}\left(\frac{p(\mathbf{x})}{\widehat{p}(\mathbf{x})}\right) d\mu(\mathbf{x})$$
(8)

belongs to the class of f-divergences.

LINKING THE PR-DIVERGENCE TO THE PR-CURVE

Theorem 2.5 (PR-Curves as a function of $\mathcal{D}_{\lambda\text{-PR}}$)

Given $P, \widehat{P} \in \mathcal{P}(\mathcal{X})$ such that $P, \widehat{P} \ll \mu$ and $\lambda \in [0, +\infty]$, the PR-Curve ∂PRD is related to the PR-Divergence $\mathcal{D}_{\lambda-PR}(P\|\widehat{P})$ as follows.

$$\alpha_{\lambda}(P\|\widehat{P}) = \min(1,\lambda) - \mathcal{D}_{\lambda\text{-PR}}(P\|\widehat{P}).$$

$$eta_{\lambda}(P\|\widehat{P}) = \min(1,\lambda) - \mathcal{D}_{\lambda-PR}(\widehat{P}\|P).$$

LINKING THE PR-DIVERGENCE TO THE PR-CURVE

Theorem 2.5 (PR-Curves as a function of $\mathcal{D}_{\lambda\text{-PR}}$)

Given $P, \widehat{P} \in \mathcal{P}(\mathcal{X})$ such that $P, \widehat{P} \ll \mu$ and $\lambda \in [0, +\infty]$, the PR-Curve ∂PRD is related to the PR-Divergence $\mathcal{D}_{\lambda-PR}(P\|\widehat{P})$ as follows.

$$\alpha_{\lambda}(P\|\widehat{P}) = \min(1,\lambda) - \mathcal{D}_{\lambda\text{-PR}}(P\|\widehat{P}).$$

$$eta_{\lambda}(P\|\widehat{P}) = \min(1,\lambda) - \mathcal{D}_{\lambda-PR}(\widehat{P}\|P).$$

A direct consequence of Theorem 2.5:

$$\underset{\widehat{P} \in \mathcal{P}(\mathcal{X})}{\operatorname{argmin}} \, \mathcal{D}_{\lambda\text{-PR}}(P \| \widehat{P}) = \underset{\widehat{P} \in \mathcal{P}(\mathcal{X})}{\operatorname{argmax}} \, \alpha_{\lambda}(P \| \widehat{P}).$$


EXPLAINING QUALITY/DIVERSITY

CONNECTION BETWEEN PR-DIVERGENCE AND f-DIVERGENCES

Theorem 2.6 (*f*-divergences as a weighted average of PR-Divergences)


For any $P, \widehat{P} \in \mathcal{P}(\mathcal{X})$ supported on all \mathcal{X} and any $\lambda \in [0, +\infty]$, then:

$$\mathcal{D}_{\!f}(P\|\widehat{P}) = \int_0^\infty \frac{1}{\lambda^3} f''\left(\frac{1}{\lambda}\right) \mathcal{D}_{\lambda\text{-PR}}(P\|\widehat{P}) \mathrm{d}\lambda,$$



OPTIMIZING THE PR-DIVERGENCE

EXAMPLES



PR FOR GENERATIVE MODELS

OPTIMIZING THE PR-DIVERGENCE


EXAMPLES

PR for Generative Models

49 / 81

OPTIMIZING THE PR-DIVERGENCE WITH OUR APPROACH

IN PRACTICE



OPTIMIZING THE PR-DIVERGENCE WITH OUR APPROACH

TRAINING GANS

Model		CIFAR-10 32 × 32			CelebA 64 × 64				
		FID	P	R	FID	P	R		
Baseline	Big-	13.37	86.51	65.66	9.16	78.41	51.42		
GAN									
$\lambda = 0.05$		13.29	81.10	70.63	-	-	-		
$\lambda = 0.1$		11.62	81.78	74.58	-	-	-		
$\lambda = 0.2$		13.36	84.85	65.13	8.79	83.37	44.07		
$\lambda = 0.5$		14.50	83.27	68.23	6.03	77.60	55.98		
$\lambda = 1.0$		14.03	83.04	69.35	13.07	81.70	36.85		
$\lambda = 2.0$		16.94	84.93	59.79	14.23	82.98	32.87		
$\lambda = 5.0$		32.54	83.39	56.94	22.45	83.96	25.81		
$\lambda = 10.0$		39.69	84.11	39.29	-	-	-		
$\lambda = 20.0$		67.03	90.03	21.81	-	-			







$$\lambda = 0.1$$

$$\lambda = 10$$

OPTIMIZING THE PR-DIVERGENCE WITH OUR APPROACH

FINE-TUNING GANS

Model	ImageNet 128 × 128			FFHQ 256 × 256		
	FID	P	R	FID	P	R
Baseline BigGAN	9.83	28.04	41.21	41.41	65.57	10.17
Soft $\psi = 0.7$	11.39	23.04	31.13	56.43	76.59	4.87
Soft $\psi = 0.5$	15.49	20.20	19.83	82.05	84.48	1.58
Hard $\psi = 2.0$	9.69	25.83	39.89	43.32	68.84	8.66
Hard $\psi = 1.0$	12.12	21.86	35.42	56.19	76.44	4.76
Hard $\psi=0.5$	15.21	21.13	29.55	71.32	80.99	4.84
$\lambda = 0.2$	9.92	26.69	42.04	35.66	78.70	9.45
$\lambda = 0.5$	10.82	26.83	42.38	35.24	78.41	9.66
$\lambda = 1.0$	20.42	29.72	28.21	35.91	78.95	8.32
$\lambda = 2.0$	20.21	30.27	30.49	36.33	81.10	8.69
$\lambda = 5.0$	20.76	30.87	28.38	38.16	84.31	8.52

When λ increases, $\begin{cases} \text{Precision} \uparrow \\ \text{Recall} \downarrow \end{cases}$ with better performances than truncation.

IMPROVING PRECISION AND RECALL IN GENERATIVE MODELS


Improving:

How can we improve the quality and diversity of a pre-trained generative models?

GENERAL SETTING

To sample a point from the learned distribution \widehat{P} :

- ▶ Sample $z \sim Q$.
- ► Compute x = G(z).




GENERAL SETTING

To sample a point from the learned distribution \widehat{P} :

- ▶ Sample $z \sim Q$.
- ightharpoonup Compute x = G(z).

$$P \neq \widehat{P}$$




GENERAL SETTING

To sample a point from the learned distribution \widehat{P} :

- ▶ Sample $z \sim Q$.
- ightharpoonup Compute x = G(z).

We have an estimation of $\frac{p(x)}{\widehat{p}(x)}$ using $\nabla f^*(T(x))$.




REJECTION SAMPLING

To sample a point from the refined distribution \widetilde{P} :

- ▶ Sample $z \sim Q$.
- ightharpoonup Compute x = G(z).
- ightharpoonup Accept x with probability a(x).

Using $\frac{p(x)}{\widehat{p}(x)}$ in a(x) allows sampling from P.




REJECTION SAMPLING

To sample a point from the refined distribution \widetilde{P} :

- ▶ Sample $z \sim Q$.
- ightharpoonup Compute x = G(z).
- ightharpoonup Accept x with probability a(x).

The acceptance rate is:

$$\mathbb{E}_{\widehat{p}}\left[a(\boldsymbol{x})\right]$$
.




REJECTION SAMPLING

To sample a point from the refined distribution \widetilde{P} :

- ▶ Sample $z \sim Q$.
- ightharpoonup Compute x = G(z).
- ightharpoonup Accept x with probability a(x).

It defines a new distribution \widetilde{P} .



REJECTION SAMPLING IN HIGH DIMENSION

BUDGETED REJECTION SAMPLING

TUNING THE ACCEPTANCE RATE

Definition 3.1 (Discriminator Rejection Sampling (DRS) - Azadi et al. [2])

Let $\gamma \in \mathbb{R}$ *, the acceptance probability is:*

$$a_{\mathrm{DRS}}(\mathbf{x}) = \frac{r(\mathbf{x})}{r(\mathbf{x})(1 - e^{\gamma}) + Me^{\gamma}}.$$

If γ < 0, then the acceptance rate increases.

BUDGETED REJECTION SAMPLING

TUNING THE ACCEPTANCE RATE

Definition 3.1 (Discriminator Rejection Sampling (DRS) - Azadi et al. [2])

Let $\gamma \in \mathbb{R}$ *, the acceptance probability is:*


$$a_{\mathrm{DRS}}(\mathbf{x}) = \frac{r(\mathbf{x})}{r(\mathbf{x})(1 - e^{\gamma}) + Me^{\gamma}}.$$

If γ < 0, then the acceptance rate increases.

DEFINITION

Traditionally, the goal is:

$$\min_{G} \quad \mathcal{D}_f(P \| \widehat{P}_G)$$




Distribution Space

DEFINITION

Traditionally, the goal is:

$$\min_{G} \quad \mathcal{D}_f(P \| \widehat{P}_G)$$

An acceptance function a(x) such that the acceptance rate is greater than 1/K defines a refined distribution \widetilde{P}_a in a convex set that contains \widehat{P}_G .




Distribution Space

DEFINITION

With a given \widehat{P}_G , our goal is:

$$\min_{a} \quad \mathcal{D}_{f}(P \| \widetilde{P}_{a})
\text{s.t.} \begin{cases} \mathbb{E}_{\widehat{P}} \left[a(\mathbf{x}) \right] \ge 1/K, \\ \forall \mathbf{x} \in \mathcal{X}, \ 0 \le a(\mathbf{x}) \le 1. \end{cases}$$
(9)



Distribution Space

DEFINITION

Theorem 3.2 (Optimal Acceptance Function)

For a sampling budget $K \ge 1$ and finite \mathcal{X} , the solution is,

$$a_{\text{OBRS}}(\mathbf{x}) = \min\left(\frac{p(\mathbf{x})}{\widehat{p}(\mathbf{x})}\frac{c_K}{M}, 1\right),$$
 (10)

where $c_K \geq 1$ is such that $\mathbb{E}_{x \sim \widehat{p}}[a_{OBRS}(x)] = 1/K$.

DEFINITION

Theorem 3.2 (Optimal Acceptance Function)

For a sampling budget $K \ge 1$ and finite \mathcal{X} , the solution is,

$$a_{\text{OBRS}}(\mathbf{x}) = \min\left(\frac{p(\mathbf{x})}{\widehat{p}(\mathbf{x})}\frac{c_K}{M}, 1\right),$$
 (10)

where $c_K \geq 1$ is such that $\mathbb{E}_{x \sim \widehat{p}}[a_{OBRS}(x)] = 1/K$.

EFFECT OF THE OBRS

Proposition 3.3 (Precision and Recall Improvement)

Let $K \leq M$ be the budget for the OBRS. For any $(\alpha, \beta) \in PRD(P, \widehat{P})$ we have $(\alpha', \beta) \in PRD(P, \widetilde{P}_{a_{OBRS}})$ with $\alpha' = \min\{1, K\alpha\}$.

PR FOR GENERATIVE MODELS

EFFECT OF THE OBRS

Proposition 3.3 (Precision and Recall Improvement)

Let $K \leq M$ be the budget for the OBRS. For any $(\alpha, \beta) \in PRD(P, \widehat{P})$ we have $(\alpha', \beta) \in PRD(P, \widetilde{P}_{a_{OBRS}})$ with $\alpha' = \min\{1, K\alpha\}$.

EFFECT OF THE OBRS

Proposition 3.3 (Precision and Recall Improvement)


Let $K \leq M$ be the budget for the OBRS. For any $(\alpha, \beta) \in PRD(P, \widehat{P})$ we have $(\alpha', \beta) \in PRD(P, \widetilde{P}_{a_{OBRS}})$ with $\alpha' = \min\{1, K\alpha\}$.


EFFECT OF THE OBRS

Proposition 3.3 (Precision and Recall Improvement)

Let $K \leq M$ be the budget for the OBRS. For any $(\alpha, \beta) \in PRD(P, \widehat{P})$ we have $(\alpha', \beta) \in PRD(P, \widetilde{P}_{a_{OBRS}})$ with $\alpha' = \min\{1, K\alpha\}$.

IN PRACTICE





1/K	FID	P	R
0.25	1.57	78.48	86.73
0.50	1.58	78.23 77.94	86.05
0.75	1.77	77.94	86.54
1	1.97	77.91	86.62

GAN on CelebA

Diffusion Model on CIFAR-10

OTHER METHODS TO IMPROVE PRECISION AND RECALL

BOOSTING

Boosting Generative models:



Figure. Left: Samples from the dataset given high weights by the discriminator. Right: Samples from the dataset given low weights by the discriminator. The next model will focus on the sample on the right. Source: Tolstikhin et al. [28]

- Tolstikhin et al. [28]
- ► Grover and Ermon [14]

OTHER METHODS TO IMPROVE PRECISION AND RECALL

GRADIENT ASCENT

Using the discriminator as a classifier and perform a gradient descent:





Figure. Source: Ansari et al. [1]

- Ansari et al. [1]
- ► Tanaka [27]
- ▶ Che et al. [9]

OTHER METHODS TO IMPROVE PRECISION AND RECALL

GAUSSIAN MIXTURES

Training a Gaussian Mixture $\mathcal{N}(\mu_k, \sigma I)$ in the latent space:



0.95

Figure. Source: Ben-Yosef and Weinshall [4]

- ▶ Ben-Yosef and Weinshall [4]
- ▶ Pandeva and Schubert [22]
- ► Alternative idea: Use Expectation-Maximization Bishop [5]

RECAP

References to evaluate generative models:

- ► FID: Heusel et al. [15]
- ▶ PR-Curves: Sajjadi et al. [23]
- ▶ Support based metrics: Kynkäänniemi et al. [19]

Methods to tune precision and recall:

- ► Truncation: Karras et al. [17], Kingma and Dhariwal [18]
- ► *f*-GAN: Nowozin et al. [21]
- ▶ PR-GAN: Verine et al. [31]

PR FOR GENERATIVE MODELS

RECAP

Methods to improve precision and recall:

- ▶ Rejecting samples: Azadi et al. [2], Verine et al. [31], Turner et al. [29], Tanaka [27]
- ▶ Boosting: Tolstikhin et al. [28], Grover and Ermon [14]
- ► Gradient Ascent: Ansari et al. [1], Tanaka [27], Che et al. [9]
- Latent Space Reshaping: Ben-Yosef and Weinshall [4], Pandeva and Schubert [22], Issenhuth et al. [16]
- ► EM in the latent space: Bishop [5]

PR FOR GENERATIVE MODELS

CONCLUSION

Thanks!

PR for Generative Models

How to SLURM?

DATA SCIENCE LAB

Constant Bourdrez

PhD Student, Centre des Données, ENS-PSL

October 29, 2025



PR FOR GENERATIVE MODELS 68 / 81

WHAT IS SLURM?



- Open-source workload manager for high-performance clusters.
- ► Schedules and queues jobs; matches jobs to available nodes.
- ► Allocates requested resources (CPU, GPU, memory) reliably.
- Provides monitoring and accounting tools (e.g., squeue, sacct).

PR FOR GENERATIVE MODELS 69 / 81

TRADITIONAL SLURM WORKFLOW





Figure. Typical SLURM job submission and execution workflow.

PR for Generative Models 70 / 81

SLURM BACKFILL

- ▶ Jobs at the top of the queue have highest priority.
- ► Slurm can run lower-priority jobs without delaying higher-priority jobs.
- ▶ Helps smaller jobs run and prevents large jobs from blocking the system.
- Especially useful for multi-node job scheduling.



PR FOR GENERATIVE MODELS 71 / 81

SUBMITTING A SLURM JOB

```
name="GAN TRAINING"
outdir="outputs"
echo "Launching job $name"
sbatch <<EOT</pre>
#!/bin/bash
#SBATCH -p mesonet
#SBATCH -N 1
#SBATCH -c 28
#SBATCH --gres=gpu:${n_gpu}
#SBATCH --time=00:20:00
#SBATCH --mem=256G
#SBATCH --account=m25146
#SBATCH -- job-name=$ {name}
#SBATCH --cpus-per-task=4
#SBATCH --output=$outdir/path.txt
source activate venv/bin/activate
python my_program.py
EOT
```

PR FOR GENERATIVE MODELS 72 / 81

BASIC COMMANDS

Submit a job:

sbatch myscript.sh

chmod +x myscript.sh
./myscript.sh

Check your jobs:

squeue -u \$USER

Cancel a job:

scancel JOB_ID

PR FOR GENERATIVE MODELS 73 / 81

GENERATING AN SSH KEY FOR MESONET

Step 1: Generate a new SSH key

```
ssh-keygen -t ed25519 -C "your_email@example.com"
```

- ▶ Press Enter to accept the default file location. (.ssh/id_ed25519)
- ▶ Enter a passphrase if you want extra security (optional).

Step 2: Copy the public key

```
cat .ssh/id_ed25519.pub
```

This prints your public key in the terminal.

Step 3: Add your key to MesoNET

- 1. Log in to your MesoNET account.
- 2. Go to the **SSH Keys** section and paste the public key you copied from the terminal.
- 3. Save your changes.

CONNECTING TO JULIET (MESONET)

Prerequisites:

- ► A valid **MesoNET account**
- ► An **SSH key** associated with your account

Connect via SSH:

```
ssh username@juliet.mesonet.fr
```

Replace username with your MesoNET ID. Optional SSH configuration:

```
Host juliet
Hostname juliet.mesonet.fr
User username
IdentityFile [PathToYourSSHKey]
IdentitiesOnly yes
```

Replace username and [PathToYourSSHKey] with your information.

CONCLUSION

Thanks!

PR FOR GENERATIVE MODELS

REFERENCES I

- [1] Ansari, A. F., Ang, M. L., and Soh, H. (2021). Refining Deep Generative Models via Discriminator Gradient Flow. arXiv:2012.00780 [cs, stat].
- [2] Azadi, S., Olsson, C., Darrell, T., Goodfellow, I., and Odena, A. (2019). Discriminator Rejection Sampling. arXiv:1810.06758 [cs, stat].
- [3] Barratt, S. and Sharma, R. (2018). A Note on the Inception Score. arXiv:1801.01973 [cs, stat].
- [4] Ben-Yosef, M. and Weinshall, D. (2018). Gaussian Mixture Generative Adversarial Networks for Diverse Datasets, and the Unsupervised Clustering of Images. *arXiv:1808.10356* [cs, stat]. arXiv: 1808.10356.
- [5] Bishop, C. M. (2006). Pattern Recognition and Machine Learning (Information Science and Statistics). Springer-Verlag.
- [6] Borji, A. (2022). Pros and cons of GAN evaluation measures: New developments. *Computer Vision and Image Understanding*, 215:103329.
- [7] Brock, A., Donahue, J., and Simonyan, K. (2019). Large Scale GAN Training for High Fidelity Natural Image Synthesis. arXiv:1809.11096 [cs, stat].
- [8] Bronnec, F. L., Verine, A., Negrevergne, B., Chevaleyre, Y., and Allauzen, A. (2024). Exploring Precision and Recall to assess the quality and diversity of LLMs. *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics*. arXiv:2402.10693 [cs].

PR FOR GENERATIVE MODELS

REFERENCES II

- [9] Che, T., Zhang, R., Sohl-Dickstein, J., Larochelle, H., Paull, L., Cao, Y., and Bengio, Y. (2020). Your GAN is Secretly an Energy-based Model and You Should Use Discriminator Driven Latent Sampling. In *Advances in Neural Information Processing Systems*, volume 33, pages 12275–12287. Curran Associates, Inc.
- [10] Chong, M. J. and Forsyth, D. (2020). Effectively Unbiased FID and Inception Score and where to find them. arXiv:1911.07023 [cs].
- [11] Frolov, S., Hinz, T., Raue, F., Hees, J., and Dengel, A. (2021). Adversarial text-to-image synthesis: A review. *Neural Networks*, 144:187–209.
- [12] Goodfellow, I. J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. (2014). Generative Adversarial Networks. In 27th Conference on Neural Information Processing Systems (NeurIPS 2014). arXiv: 1406.2661.
- [13] Grover, A., Dhar, M., and Ermon, S. (2018). Flow-GAN: Combining Maximum Likelihood and Adversarial Learning in Generative Models. arXiv:1705.08868 [cs, stat].
- [14] Grover, A. and Ermon, S. (2017). Boosted Generative Models. arXiv:1702.08484 [cs, stat].
- [15] Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., and Hochreiter, S. (2017). GANs Trained by a Two Time-Scale Update Rule Converge to a Local Nash Equilibrium. In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.

REFERENCES III

- [16] Issenhuth, T., Tanielian, U., Picard, D., and Mary, J. (2022). Latent reweighting, an almost free improvement for GANs. In 2022 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV), pages 3574–3583, Waikoloa, HI, USA. IEEE.
- [17] Karras, T., Laine, S., Aittala, M., Hellsten, J., Lehtinen, J., and Aila, T. (2020). Analyzing and Improving the Image Quality of StyleGAN. arXiv:1912.04958 [cs, eess, stat].
- [18] Kingma, D. P. and Dhariwal, P. (2018). Glow: Generative Flow with Invertible 1x1 Convolutions. In 32nd Conference on Neural Information Processing Systems (NeurIPS 2018), Montréal, Canada., volume 31.
- [19] Kynkäänniemi, T., Karras, T., Laine, S., Lehtinen, J., and Aila, T. (2019). Improved Precision and Recall Metric for Assessing Generative Models. In *33rd Conference on Neural Information Processing Systems (NeurIPS 2019), Vancouver, Canada.* arXiv: 1904.06991.
- [20] Nguyen, X., Wainwright, M. J., and Jordan, M. I. (2009). On surrogate loss functions and \$f\$-divergences. *The Annals of Statistics*, 37(2). arXiv:math/0510521.
- [21] Nowozin, S., Cseke, B., and Tomioka, R. (2016). f-GAN: Training Generative Neural Samplers using Variational Divergence Minimization. arXiv:1606.00709 [cs, stat].
- [22] Pandeva, T. and Schubert, M. (2019). MMGAN: Generative Adversarial Networks for Multi-Modal Distributions. *arXiv:1911.06663 [cs, stat]*. arXiv: 1911.06663.

PR FOR GENERATIVE MODELS

REFERENCES IV

- [23] Sajjadi, M. S. M., Bachem, O., Lucic, M., Bousquet, O., and Gelly, S. (2018). Assessing Generative Models via Precision and Recall. In 32nd Conference on Neural Information Processing Systems (NeurIPS 2018), Montréal, Canada. arXiv: 1806.00035.
- [24] Salimans, T., Goodfellow, I., Zaremba, W., Cheung, V., Radford, A., and Chen, X. (2016). Improved Techniques for Training GANs. arXiv:1606.03498 [cs].
- [25] Simon, L., Webster, R., and Rabin, J. (2019). Revisiting precision recall definition for generative modeling. In *Proceedings of the 36th International Conference on Machine Learning*, pages 5799–5808. PMLR. ISSN: 2640-3498.
- [26] Siry, R., Webster, R., Simon, L., and Rabin, J. (2023). On the Theoretical Equivalence of Several Trade-Off Curves Assessing Statistical Proximity. *Journal of Machine Learning Research*, 24.
- [27] Tanaka, A. (2019). Discriminator optimal transport. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- [28] Tolstikhin, I., Gelly, S., Bousquet, O., Simon-Gabriel, C.-J., and Schölkopf, B. (2017). AdaGAN: Boosting Generative Models. arXiv:1701.02386 [cs, stat].
- [29] Turner, R., Hung, J., Frank, E., Saatchi, Y., and Yosinski, J. (2019). Metropolis-Hastings Generative Adversarial Networks. In *Proceedings of the 36th International Conference on Machine Learning*, pages 6345–6353. PMLR. ISSN: 2640-3498.

REFERENCES V

- [30] Verine, A., Bronnec, F. L., Zheng, K., Allauzen, A., Chevaleyre, Y., and Negrevergne, B. (2025). Improving Diversity in Language Models: When Temperature Fails, Change the Loss. ICML 2025. arXiv. arXiv:2508.09654 [cs].
- [31] Verine, A., Negrevergne, B., Pydi, M. S., and Chevaleyre, Y. (2023). Precision-Recall Divergence Optimization for Generative Modeling with GANs and Normalizing Flows. *Advances in Neural Information Processing Systems*, 36:32539–32573.
- [32] Yann LeCun, Corinna Cortes, and Burges, C. (2010). MNIST handwritten digit database. ATT Labs, 2.