

# Understanding Deep Generative Models with Generalized Empirical Likelihoods

Suman Ravuri, Mélanie Rey, Shakir Mohamed  
DeepMind  
London, UK

{ravuris,melanierey,shakir}@deepmind.com

Marc Peter Deisenroth  
University College London  
London, UK

m.deisenroth@ucl.ac.uk

Understanding how well a deep generative model captures a distribution of high-dimensional data remains an important open challenge. It is especially difficult for certain model classes, such as Generative Adversarial Networks and Diffusion Models, whose models do not admit exact likelihoods. In this work, we demonstrate that generalized empirical likelihood (GEL) methods offer a family of diagnostic tools that can identify many deficiencies of deep generative models (DGMs). We show, with appropriate specification of moment conditions, that the proposed method can identify *which* modes have been dropped, the degree to which DGMs are mode imbalanced, and whether DGMs sufficiently capture intra-class diversity. We show how to combine techniques from Maximum Mean Discrepancy and Generalized Empirical Likelihood to create not only distribution tests that retain per-sample interpretability, but also metrics that include label information. We find that such tests predict the degree of mode dropping and mode imbalance up to 60% better than metrics such as improved precision/recall.

## 1. Introduction

In an era that has witnessed deep generative models (DGMs) produce photorealistic images from text descriptions [55, 62, 63], speech rivaling that of professional voice actors [8, 20, 49], and text seemingly indistinguishable from writing on the internet [5, 50], it perhaps seems quaint to focus on better evaluation of such models. One might reasonably claim that such evaluation may have been useful prior to the development of these models, but such evaluation is certainly less important now. They might further bolster their claim by noting that researchers have already proposed reasonable metrics for these modalities — Fréchet Inception Distance for images [26], Mean Opinion Score (MOS) for speech, and perplexity for text —, and the results have been so compelling that researchers and practitioners have begun to deploy such models in downstream tasks [1, 3, 48, 80]. So why should we focus on better evaluation now?

We are advocating for better and more nuanced evaluation now precisely *because* DGMs have reached sufficient matu-

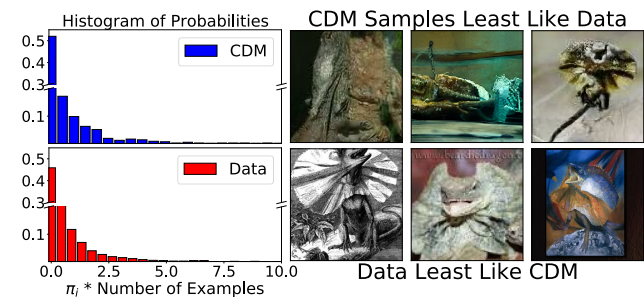


Figure 1. Empirical Likelihood methods can identify model samples outside of the data distribution, and data samples outside of the model distribution. In this example, we use a two-sample generalized empirical likelihood test to evaluate how well a Cascaded Diffusion Model trained on ImageNet captures the “Frilled Lizard” class. The three samples on the bottom right show samples from the data distribution the model is not likely to represent, and those on the top right show samples from the model not likely to be in the data distribution. These examples have 0 probability in the empirical likelihood test.

ity to be used in downstream tasks. When Deep Generative Models struggled to produce realistic  $32 \times 32$  images, a researcher working on the model may not have found nuanced evaluation useful. Now that DGMs can credibly produce realistic megapixel images from text input, that same researcher may find more nuanced evaluation helpful to understand what deficiencies still exist. This view reflects recent trends in generative model evaluation, as researchers have now started to adopt metrics such as precision/recall [42] to better understand *how* the model is misspecified.

Our approach comprises three parts: a moment condition on the DGM and data distributions that tells us if the model captured a salient property of the data distribution; a score that indicates how well a moment condition is satisfied; and a decomposition of that score that shows how each test point contributed. From these three elements, we can diagnose how a DGM is misspecified.

We observe that many evaluation metrics can be expressed as a set of moment conditions — for example, Fréchet Inception Distance can be recast as the condition

that the first and second moments of Inception v3 Pool3 (further denoted as “Pool3”) features match. Then, with a set of moment conditions as our specification, we employ the machinery of Empirical Likelihood methods [51] to provide us both the aggregate and per-sample scores.

The Empirical Likelihood is a moment condition test that approaches the problem by answering the following question: how much must the data distribution change in order to satisfy the moment condition? Typically, researchers use these tests to determine whether a moment condition is satisfied. We find that beyond just identifying *whether* a moment condition is satisfied, under certain conditions on the moments, the changed distribution can tell us *how* a DGM is misspecified. Our main contribution is elucidating the conditions under which we can use Empirical Likelihood methods as a diagnostic tool. In the process, we introduce four new tests and show the following:

**Mode dropping.** We show how to create moment conditions that identify *which* modes have dropped, without access to label information from the model.

**Mode imbalance.** We experimentally show that when a DGM samples one mode more frequently than others, GEL methods can predict the degree of this mode imbalance.

**Interpretable tests.** We detail how to use moment conditions from Maximum Mean Discrepancy [22] approaches to create interpretable tests.

**Improper label conditioning.** By including label information, we create tests that identify when a label-conditioned generative model ignores its conditioning label.

## 2. Background

In this section, we provide a brief background on the Empirical Likelihood and moment conditions useful for generative model evaluation. Both of these areas have a rich history of their own. We refer the reader to the excellent monograph by [52] for a broader discussion on Empirical Likelihood and [25] for moment restrictions.

**Generalized Empirical Likelihood** The Empirical Likelihood (EL) [51] is a classical nonparametric method of statistical inference. Originally, EL was proposed as a method of inference on the mean: given  $n$  independent samples  $x_1, \dots, x_n \in \mathbb{R}^d$  from unknown distribution  $p$ , EL determines whether the mean of  $p$  is equal to a known constant  $\mathbf{c} \in \mathbb{R}^d$ . EL is unique in its approach, as it models the samples with a weighted empirical distribution  $P_\pi(x) = \sum_{i=1}^n \pi_i \mathbb{I}_{[x_i=x]}$ , with weights  $\pi_i$  that satisfy the rules of probability (namely that  $\sum_i \pi_i = 1, \pi_i \geq 0$ ). It finds, among all distributions  $P_\pi$  that match the mean condition  $\mathbb{E}_{P_\pi}[x] \equiv \sum_{i=1}^n \pi_i x_i = \mathbf{c}$ , the one that maximizes the likelihood  $\prod_i \pi_i$ . This “empirical likelihood” can be

expressed as the solution to the convex problem

$$\{\pi \mid \sum_i \pi_i = 1, \pi_i \geq 0\} \max \sum_{i=1}^n \log \pi_i \quad \text{s.t.} \quad \mathbb{E}_{x \sim P_\pi}[x] = \mathbf{c} \quad (1)$$

We denote by  $\pi^*$  and  $P_{\pi^*}$  the weights and implied distribution, respectively, that solve Eq. (1), and abbreviate “*subject to*” to “*s.t.*”. Intuitively, if  $\mathbb{E}_{x \sim p}[x] = \mathbf{c}$ , then we expect  $P_{\pi^*}$  to be “close” to the empirical distribution  $\hat{P}_n(x) \equiv \sum_{i=1}^n n^{-1} \mathbb{I}_{[x_i=x]}$ .<sup>1</sup> Eq. (1) makes this notion precise by measuring the closeness of the two distributions with the KL divergence  $D_{KL}(\hat{P}_n \parallel P_\pi)$ . If no distribution satisfies the mean constraint (which is equivalent to the mean not lying in the interior of the convex hull of  $\{x_i\}$ )<sup>2</sup>, then by convention the empirical likelihood and KL divergence are  $-\infty$  and  $\infty$ , respectively.

We use two extensions for generative model evaluation. The first extension is replacing the mean condition  $\mathbb{E}_{P_\pi}[x] = \mathbf{c}$  with a moment condition  $\mathbb{E}_{x \sim P_\pi}[\mathbf{m}(x; \mathbf{c})] = \mathbf{0}$  [54], making the method much more general. The second is replacing the KL divergence with one from the Cressie-Read family<sup>3</sup> [11]. Using other members relaxes the condition that  $\pi_i$  be strictly positive, and is useful for identifying mode dropping. The resulting objective is called the *Generalized Empirical Likelihood* (GEL):

$$\{\pi \mid \sum_i \pi_i = 1, \pi_i \geq 0\} \min D(\hat{P}_n \parallel P_\pi) \quad \text{s.t.} \quad \mathbb{E}_{x \sim P_\pi}[\mathbf{m}(x; \mathbf{c})] = \mathbf{0} \quad (2)$$

In this work, we use two members of the family, shown in the top panel of Tab. 1: the “exponential tilting” (ET) objective, and the “Euclidean likelihood”.<sup>4</sup> We also use a third extension, a two-sample version of the above test, but we defer discussion of this extension to Sec. 3.2.

**Moment Conditions** We note that many evaluation metrics are statistics of moments. They may be as simple as checking if low-order moments match. Indeed, many popular metrics for generative model evaluation are statistics of a low-order of moments: for example, Fréchet Inception Distance and Kernel Inception Distance are statistics of the first two and three moments, respectively. They may also be as complex as a full distribution test: for a likelihood-based model, the condition that  $\mathbb{E}_x[\nabla_\theta \log p_\theta(x)] = \mathbf{0}$  implies the likelihood is maximized. By decoupling the moment condition from the statistic, we can construct alternative tests based on GEL objectives. The

<sup>1</sup>In absence of the mean constraint  $\sum_{i=1}^n \pi_i x_i = \mathbf{c}$ ,  $\pi^* = n^{-1} \mathbf{1}$ , the implied distribution  $P_{\pi^*}(x)$  is the empirical distribution, also known as the nonparametric maximum likelihood estimate of the sample [37].

<sup>2</sup>Appendix A.2 includes further discussion of the convex hull condition.

<sup>3</sup>This family of divergences are of the form  $CR(\lambda) = \frac{2}{\lambda(\lambda+1)} \sum_i [(n\pi_i)^{-\lambda} - 1]$ .

<sup>4</sup>The latter objective is mainly of intellectual interest: as the solution is proportional to the Hotelling T-square statistic [29], we can recast any T-square statistic as a GEL.

Name	Objective	Divergence	$\pi_i > 0$	$\pi_i = 0$	$\pi_i < 0$
Empirical Likelihood	$\prod_i \pi_i$	$D_{KL}(\hat{P}_n \parallel P_\pi)$	Y	N	N
Exponential Tilting	$-\sum_i \pi_i \log \pi_i$	$D_{KL}(P_\pi \parallel \hat{P}_n)$	Y	Y	N
Euclidean Likelihood	$\frac{1}{2} \sum_i (\pi_i - n^{-1})^2$	$\frac{1}{2} \sum_i (\pi_i - n^{-1})^2$	Y	Y	Y

Name	Population Statistic $S$	Condition at $S = 0$	Equivalent $\mathbf{m}(x; \mathbf{c})$	$\mathbf{c}$
Mean	$(\mu_x - \mathbf{c})^\top (\mu_x - \mathbf{c})$	$\mu_x = \mathbf{c}$	$x - \mathbf{c}$	$\mathbf{c}$
Score Function	$\ \mathbb{E}[\nabla_\theta \log p_\theta(x)]\ ^2$	Maximum Likelihood	$\nabla_\theta \log p_\theta(x)$	$\mathbf{0}$
Fréchet Inception Distance	$\ \mu_x - \mu_y\ ^2 + \text{tr}(\Sigma_x + \Sigma_y - 2(\Sigma_x \Sigma_y)^{1/2})$	$\mu_x = \mu_y, \Sigma_x = \Sigma_y$	$[\phi(x), \phi(x)\phi(x)^\top] - \mathbf{c}$	$\mathbb{E}_y[[\phi(y), \phi(y)\phi(y)^\top]]$
Mean Embedding	$\sum_i (\mathbb{E}_x[k(x, t_i)] - \mathbb{E}_y[k(y, t_i)])^2$	$P = Q \implies S = 0$ a.s., $P \neq Q \implies S \neq 0$ a.s.	$k(x, t_i) - k(y, t_i) \quad i = 1, \dots, W$	$\mathbf{0}$

Table 1. Top: Generalized Empirical Likelihood objectives and valid values of  $\pi_i$ . Bottom: Common statistics as moment conditions.

bottom of Tab. 1 shows some common examples and the implied moment conditions.

There exist two issues with applying the Empirical Likelihood to the moment conditions in the table: with the exception of the score function, the moment conditions are not sufficient to distinguish between two distributions; and the dimensionality of the moment conditions may be extremely high, leading to statistical and computational challenges. While the first issue might not be particularly problematic, as we may be more interested in usability than in theoretical robustness, the second issue is more pressing, as we aim to keep the dimensionality relatively low.

A particularly appealing set of moment restrictions that addresses these issues are those given by the Maximum Mean Discrepancy (MMD) [22]

$$D^2(p, q) = \mathbb{E}_{x_1, x_2} [k(x_1, x_2)] + \mathbb{E}_{y_1, y_2} [k(y_1, y_2)] - 2 \mathbb{E}_{x, y} [k(x, y)] \quad (3)$$

where  $x, x_1, x_2 \sim p$  and  $y, y_1, y_2 \sim q$ . With the appropriate choice of kernel — such as when the kernel is characteristic [19] or universal —  $D^2(p, q) = 0$  iff  $p = q$ , and, unlike competing approaches, unbiased estimators exist. Despite these appealing properties, metrics based on MMD, such as Kernel Inception Distance (KID) [2], perhaps due to the  $O(n^2)$  computational cost, have not enjoyed widespread adoption. More importantly, a straightforward inclusion of the MMD constraint into Eq. (2) is computationally difficult, as the constraint is a nonlinear function of  $\pi_i$ .

Instead, we focus on an alternative characterization based on Mean Embeddings (ME) [10, 30]. In this approach, one tests equality of distributions  $p$  and  $q$  by comparing *mean embeddings*  $\mu_p(t) \equiv \mathbb{E}_x[k(x, t)]$  and  $\mu_q(t) \equiv \mathbb{E}_y[k(y, t)]$ , where  $t$  is a *witness point* sampled from a third distribution  $r$ . A somewhat remarkable property is that, if  $k$  is characteristic, analytic, and integrable, and  $r$  is absolutely continuous with respect to the Lebesgue measure,  $\sum_w (\mu_p(t_w) - \mu_q(t_w))^2 > 0$  a.s. if  $p \neq q$  and  $\sum_w (\mu_p(t_w) - \mu_q(t_w))^2 = 0$  a.s. if  $p = q$ .<sup>5</sup>

Previous works [10, 30] define a semimetric  $n\bar{z}^\top \Sigma_z^{-1} \bar{z}$ ,

<sup>5</sup>One can interpret this condition colloquially as satisfying the axiom of coincidence with probability 1. The aforementioned work establishes that the statistic also satisfies the other axioms of a distance “with probability 1” (this notion of distance is called a *random metric*).

where  $\bar{z} \equiv \frac{1}{n} \sum_i z_i$ ,  $\Sigma_z \equiv \frac{1}{n-1} \sum_i (z_i - \bar{z})(z_i - \bar{z})^\top$  and  $z_i = [k(x_i, t_1) - k(y_i, t_1), \dots, k(x_i, t_W) - k(y_i, t_W)]^\top$ . This T-square statistic can be written as the solution (up to a constant) to the Euclidean likelihood with moment condition  $\mathbb{E}_z[z] = \mathbf{0}$ .

$$\min_{\{\pi\}_{\sum_i \pi_i = 1}} \frac{1}{2} \sum_{i=1}^n (\pi_i - n^{-1})^2 \quad \text{subject to} \quad (4)$$

$$\sum_{i=1}^n \pi_i \begin{bmatrix} k(x_i, t_1) - k(y_i, t_1) \\ \vdots \\ k(x_i, t_W) - k(y_i, t_W) \end{bmatrix} = \mathbf{0}$$

Note that the moment condition is linear in  $\pi$  and is within the GEL framework. In the following section, we modify Eq. (4) to better diagnose DGM deficiencies.

### 3. GEL for Evaluating Generative Models

**Setup** We assume that we have samples  $x_i$  from the data distribution  $p$ , and access to a DGM with distribution  $q$ , from which we can generate samples  $y_j$ . In the GEL approach to generative model evaluation, we make three choices: the constant  $\mathbf{c}$ , which is a function of the generative model distribution  $q$  and embeds relevant information about the DGM; the associated moment function  $\mathbf{m}(x; \mathbf{c})$ , whose expectation equals  $\mathbf{0}$  should relevant information about  $p$  and  $q$  match; and the divergence  $D(\hat{P}_n \parallel P_\pi)$ , which forms the objective. Before describing the proposed metrics, we provide a motivating example.

#### 3.1. A Motivating Example

To illustrate the utility of GEL method for DGM evaluation, we create a simple test that identifies *which* modes the generative model has dropped. Ours is a GEL test of the mean:  $\mathbf{c} = \mathbb{E}_{y \sim q}[\phi(y)]$  is the mean of a feature function  $\phi$ ,  $\mathbf{m}(x; \mathbf{c}) = \phi(x) - \mathbf{c}$  is the moment function, and the reverse KL  $D_{KL}(P_\pi \parallel \hat{P}_n)$  is the divergence. The resulting problem is an exponentially tilting GEL.

The test requires a fourth choice, an appropriate feature space  $\phi$ . The key idea for choosing  $\phi$  is that if a test set contains samples  $x$  from a missing mode,  $\phi$  must be designed such that the weight  $\pi_i$  of feature  $\phi(x_i)$  goes to zero. The following lemma offers some insight into one approach.

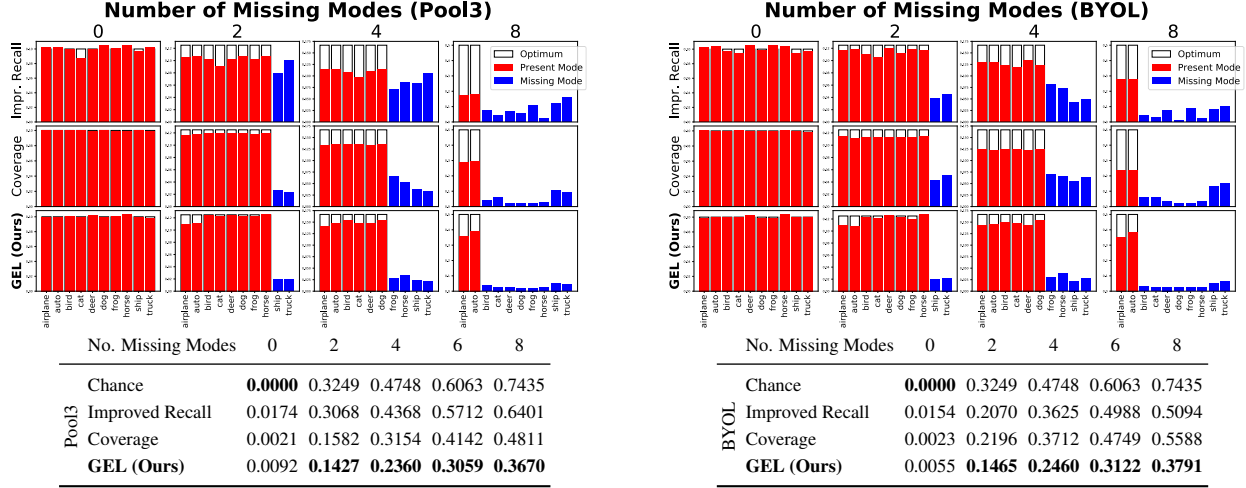


Figure 2. Mode dropping detection on CIFAR-10 data using Pool3 features (Left panel) or BYOL features (Right panel). We vary the number of missing modes in the training set and compare three different metrics computed on the test set: aggregate improved recall “weights” (Top row), aggregate Coverage “weights” (Second) and GEL weights (Third). Weights in the figures rescaled such that the optimal weight for a mode present in the training set is 1. Tables at the bottom show the Hellinger distance between ground truth probabilities and weights obtained using the three metrics.

**Lemma 1.** Assume that the true data distribution  $p$  is a mixture of the model distribution  $q$  and another distribution  $v$ . We consider an estimation  $\hat{\mathbb{E}}_q[\phi(y)]$  of the model mean obtained using samples  $y_i, i \in \mathcal{I} \subset \mathbb{N}$  from  $q$ ; and a test set composed of samples from  $p$ . The test set can be split into samples  $\{s_1, \dots, s_m\} \sim v$  and samples  $\{y_{m+1}, \dots, y_n\} \sim q$ . Then, the mean equality condition is

$$\hat{\mathbb{E}}_q[\phi(y)] = \sum_{i=1}^m \pi_i \phi(s_i) + \sum_{i=m+1}^n \pi_i \phi(y_i) \quad (5)$$

If the convex hull  $\text{Conv}\{\phi(s_i), i = 1, \dots, m\}$  does not intersect with  $\text{Span}\{\phi(y_i), i \in \mathcal{I} \cup \{m+1, \dots, n\}\}$  then  $\pi_i = 0, i = 1, \dots, m$ .

The proof can be found in Appendix B. Note that deep features, such as those from the Pool3 layer of the Inception v3 network [74], the FC2 layer of the VGG16 network [68], or BYOL features [23], approximately satisfy these properties. As features lie in the non-negative orthant (as it is an output of ReLU),  $\mathbf{0}$  does not lie within the convex hull of any set of points. Since the features of different classes are approximately linearly separable, it is unlikely that linear combination of features of a dropped mode would be part of the solution that satisfies the mean constraint.

In addition to the features, we chose the reverse KL divergence as our objective, since by construction we require  $\pi_i$  to be 0 for certain test set samples. Note that we need not worry about the normalization of  $\phi(x)$ , since we obtain the same  $\pi_i$  for any full-rank linear transformation of  $\phi$ .

*Experiment:* To illustrate the performance of the method when modes are missing, we conduct a controlled study. We

construct synthetic generative models from the CIFAR-10 training set [41]. Each of these “generative models” contains 5,000 total examples per class, with a varying number of missing classes (ranging from 0 to 8). We use the CIFAR-10 test set as the data distribution.

As shown in Fig. 2, GEL weights are highly sensitive to which modes were dropped, and can be used as a diagnostic tool for evaluating models.<sup>6</sup> We note that our approach is sensitive to *which* modes were dropped even though the features used were not trained on CIFAR-10 data. We also compare the method to the improved recall metric [42], and per-class coverage probabilities [47] (the latter with four nearest neighbors, as parameter gave us the recommended .95 coverage for 0 dropped modes). The former substantially overestimates the probability of the missing mode. The latter, while performing better than recall, overestimates probabilities for multiple missing modes. GEL outperforms both and moreover, does not require tuning of the highly sensitive nearest neighbor parameter  $k$ . Finally, as shown on the right-hand side of Fig. 2, GEL is not hugely reliant on the Pool3 space: it performs similarly on Bootstrap Your Own Latent (BYOL) features.

### 3.2. Proposed Tests

We extend the test in Sec. 3.1 in various ways to improve upon the diagnostic accuracy of the metric, and to make it more broadly applicable for generative model evaluation. We focus on three specific areas. First, we use ideas from MMD

<sup>6</sup>Empirical likelihoods are also highly sensitive, but they also end up at the boundary of the convex hull after a few modes are dropped.



and ME to create GEL test of distributions. Second, we further use ideas from kernelization to include label information in the constraint, creating a test for conditional generative models. Finally, to make the metric more robust to model misspecification, improve interpretability of the metric, and ensure model comparison, we introduce a two-sample test and a modification of the one-sample test.

**A Distribution Test by combining GEL and ME** If the test in Sec. 3.1 is the test of the mean, then by analogy, one only needs to replace the mean condition to create GEL tests of distributions. We use moment conditions based on Mean Embeddings (ME) referenced in Sec. 2 for the appropriate moment restrictions, which importantly are linear in  $\pi$ .

To render the test useful for DGM evaluation, we make three design choices. First, similar to the mean test in Sec. 3.1, we use a feature space. Second, to improve interpretability, we use multiple witness points (typically 50–1,024), which are sampled from a validation set of the data distribution.<sup>7</sup> Finally, we use the exponential kernel  $k(x, y) = \exp(x^\top y/d)$  (where  $d$  is the dimension of the vectors), which is analytic and characteristic on compact sets of  $\mathbb{R}^d$  [46].<sup>8</sup> This gives us the **Kernel GEL (KGEL)** test

$$\min_{\{\pi \mid \sum_i \pi_i = 1, \pi_i \geq 0\}} D_{KL}(P_\pi \parallel \hat{P}_n) \quad \text{subject to} \quad \sum_{i=1}^n \pi_i \begin{bmatrix} k(\phi(x_i), \phi(t_1)) \\ \vdots \\ k(\phi(x_i), \phi(t_W)) \end{bmatrix} = \mathbb{E}_{y \sim q} \begin{bmatrix} k(\phi(y), \phi(t_1)) \\ \vdots \\ k(\phi(y), \phi(t_W)) \end{bmatrix} \quad (6)$$

**Including Label Information via Kernelization** A second benefit of kernelization is that it allows us to easily include label information in the metric, enabling us to create tests for conditional generative models. The extension is straightforward: if we consider  $x = (x^{(i)}, x^{(l)})$  to include both the image  $x^{(i)}$  and its associated label  $x^{(l)}$ , respectively, then we can construct kernels  $k(x, y)$  that include both image and label information. We use the product kernel  $k(x, y) = k_i(x^{(i)}, y^{(i)})k_l(x^{(l)}, y^{(l)})$ , where  $k_i$  and  $k_l$  are image and label kernels, respectively. If  $k_i$  and  $k_l$  are universal, then the product kernel is also universal and therefore characteristic [73]. A test on the joint distribution is simply Eq. (6) with the appropriate product kernel.<sup>9</sup>

The choice of label kernel depends on the type of label information. If the dataset contains a small number of labels, the delta kernel,  $k_l(x^{(l)}, t^{(l)}) = \mathbb{I}_{[x^{(l)}=t^{(l)}]}$  is an appropriate

<sup>7</sup>Technically, this violates the condition that the witness distribution be absolutely continuous with respect to the Lebesgue measure. We could easily fix the violation by adding a small amount of Gaussian noise to the witness points. This, however, makes no practical difference to the resulting GEL solution.

<sup>8</sup>Also, the kernel is trivially integrable on compact sets of  $\mathbb{R}^d$ .

<sup>9</sup>As we use exponential kernel, which is characteristic but not universal, the product kernel may not be characteristic. We find this distinction to be more of a theoretical than practical concern.

choice (as it is universal [46]).  $k_l$  assigns 0 similarity to two points that do not have the same label. We use this test in Sec. 4.1 to identify samples where DGMs ignore label information.

When a label hierarchy is available, such as for ImageNet, we can construct a kernel that measures similarity between labels in a more fine-grained way. First, we associate a label with the path from the root node to the label leaf node, encoding it as a string. We then compute similarity between labels using any string kernel, with one simple option being the special case of the Smith-Waterman score using no gap penalty and an identity substitution matrix [78]. We provide further details in Appendix D.

**Accounting for Model Misspecification with Two-sample Tests** DGMs may not only fail to represent part of the data distribution, but also may generate samples outside of the support of the data distribution. One can perform the previously described tests by letting  $c$  be a function of data, and  $x$  DGM samples. There exists, however, a larger issue: if the DGM produces many out of distribution examples (or a single example sufficiently out of distribution), then there may not exist a  $\pi$  that satisfies the moment restriction.<sup>10</sup>

We address these issues by using a two-sample extension to GEL [53]. In this version, we assign a first set of probabilities  $\pi_i$  to test points  $x_i$ , and a second set of probabilities  $\psi_j$  to model points  $y_j$ . The two-sample variant of the mean test, which we denote **Two-sample Generalized Empirical Likelihood (GEL2)** is

$$\min_{\pi, \psi \geq 0} D(\hat{P}_n \parallel P_\pi) + D(\hat{P}_m \parallel P_\psi) \quad \text{subject to} \quad \sum_{i=1}^n \pi_i \phi(x_i) = \sum_{j=1}^m \psi_j \phi(y_j), \quad \sum_{i=1}^n \pi_i = \sum_{j=1}^m \psi_j = 1. \quad (7)$$

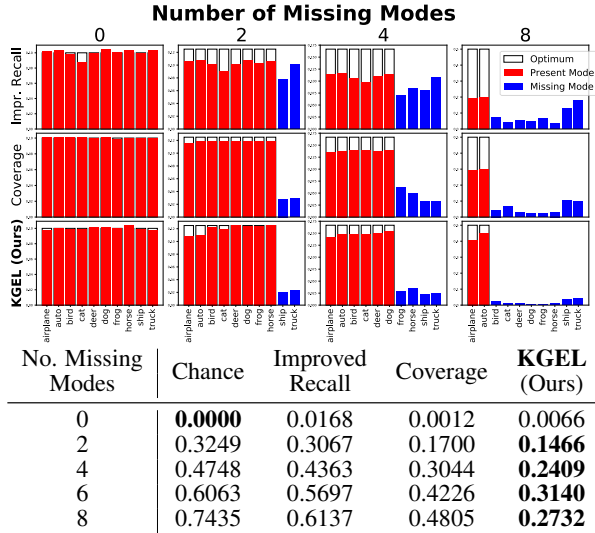
The two-sample test finds, among distributions that satisfy the moment conditions, the two that are closest to their respective empirical distributions. We call the ME version the **Kernel Two-sample Generalized Empirical Likelihood (KGEL2)**. This variant improves upon the one-sample test in two ways. First, it opens up more diagnostic tools: we can find DGM samples that are not in the support of the data distribution (Sec. 4.2). Second, it provides finite scores for a wider range of generative models: the intersection of the two convex hulls only needs to be non-empty.

The intersection of the convex hulls may be empty, in which case the empirical likelihood is  $-\infty$ . In practice this only occurs with particularly poorly performing generative models, in which case the per-sample probabilities are unlikely to be helpful.


**Calculation** We solve GEL objectives using Newton’s method on the dual problem. The computational complexity

<sup>10</sup>For mean tests, this is geometrically equivalent to the model mean not lying in the convex hull of test points. Again, due to space constraints, we defer discussion of the convex hull condition to Appendix A.2.

(a) Mode Dropping



(b) Mode Imbalance



Mode 1 Probability	Chance	Improved Recall	Coverage	KGEL (Ours)
0.1	0.3249	0.3021	0.2808	<b>0.1206</b>
0.3	0.1452	0.1313	0.1354	<b>0.0486</b>
0.5	<b>0.0000</b>	0.0080	0.0007	0.0030
0.7	0.1452	0.1506	0.1323	<b>0.0434</b>
0.9	0.3249	0.3238	0.2706	<b>0.1207</b>

Figure 3. Comparison of Evaluation Metrics on mode dropping (panel (a)), and mode imbalance (panel (b)). In the mode dropping experiment, up to 8 modes are dropped, and the table on the bottom left is the Hellinger distance between the oracle probability and the calculated one. In the mode imbalance experiment, we calculate the Hellinger distance between oracle and calculated probabilities, and the mode probability changes.

is  $\mathcal{O}(nd^3)$ , where  $n$  and  $d$  are the number of samples and dimensionality, respectively. For detailed derivations of the dual and for code, we refer the reader to Appendix C.

## 4. Experiments

### 4.1. Validation

**Detecting Mode Dropping** We repeat the experiment in Sec. 3.1 with the KGEL test. For our “generative model” we use 40,000 images from the CIFAR-10 training set, with 4,000 images per class. We use the remaining 10,000 images in the training set to draw 1,024 witness points. Otherwise, the experimental setup is identical.

The results in Fig. 3(a) show improved results of the KGEL test relative to the GEL test and other baselines. In Appendix E.1, we also show similar performance when using StyleGAN2+ADA [34] samples instead of CIFAR-10 training data.

**Detecting Mode Imbalance** In more realistic scenarios, a generative model may not drop a mode entirely, but instead undersample it relative to other modes. To better understand how accurately the proposed method predicts the degree of mode imbalance, we run a controlled experiment varying the number of examples from two modes. We use the first five categories — airplane, automobile, bird, cat, and deer — as the first mode, and the remaining categories as the second. We vary the proportion of examples in the first mode from 0.0 to 1.0, and test KGEL, a normalized improved recall, and coverage to see how well the metrics recover the probabilities. Note that at probabilities 0 and 1, this is equivalent to five modes being dropped.

The results in Figure 3(b) show that our proposed method significantly outperforms competitors. At equal probability, all methods perform similarly. For any degree of unequal probability, KGEL significantly outperforms competitors: the second best method is 29% to 212% worse.

### Identifying Improper Conditioning with Label Information

When a generative model synthesizes an image from an incorrect class (say a dog instead of butterfly), a test of only images would not identify this issue. This is true of the standard metrics such as FID, IS, and improved precision and recall; and less standard ones such as density coverage.

With KGEL tests that include label information, we can identify the degree of mislabeling. In this test, we take the training set as the generative model distribution. To simulate incorrect conditioning, we change the labels for 10% to 60% of the training set examples on CIFAR-10. We then use the kernel test, with kernel  $k(x, y) = k_i(x^{(i)}, y^{(i)})k_l(x^{(l)}, y^{(l)})$ , where  $k_i$  is the exponential kernel, and  $k_l$  the delta kernel.

The results in Fig. 4 compare the proposed KGEL test to improved precision and density, both calculated on a per-label basis.<sup>11</sup> For improved precision, we vary the value of  $k$  to obtain the precision-recall curve, while for density, we vary  $k$  to optimize for performance of the area under the precision recall curve.

We find that our proposed method significantly outperforms improved precision, and density using  $k$  so that expected coverage is greater than 0.95. One can improve results by artificially increasing  $k$ . Here we increase it

<sup>11</sup>For this test, we run 10 tests, one for labels airplane, automobile, etc. KGEL runs only once.




Figure 4. Performance of various metrics under label corruption. Left: Precision-Recall Curve at 30% label error. Middle: Precision-Recall Curve at 60% label error. Right: Area Under Precision-Recall Curve for different label errors.

to 100, over a factor of 30 of the suggested rate. Even with careful tuning of the baselines, however, KGEL still outperforms these methods without any tuning of its own, suggesting that the proposed method is more capable of identifying improper label conditioning.

## 4.2. Real-World Applications

**Assessing Within-Class Distributions using Two-Sample Tests** We use the Kernel Two-sample GEL (KGEL2) test to better understand how well generative models capture within-class distributions. We use two DGMs in this study — BigGAN-deep [4] and the Cascaded Diffusion Model (CDM) [28] — as these models represent among the best examples of their respective model classes. The KGEL2 test uses Pool3 features, 256 witness points from the ImageNet v2 dataset [60], and an exponential kernel. The test compares per-class samples from the ImageNet training set and an equivalent number from the DGM. The results in Fig. 5 show Monarch butterfly samples from the model least representative of the data distribution (model probability 0.0), and samples from data distribution least representative of the model (data probability 0.0). BigGAN-deep is least able to represent butterflies that comprise a large portion of the image and swarms of butterflies, while CDMs are least able to represent swarms of butterflies. We show more examples in Appendix F.1.

When we select the 50 samples per class with the lowest model probability (for a total of 50,000 examples), Inception Score of the CDM decreases from 166.2 to 36.82, and of BigGAN-deep from 218.1 to 88.37.

**Understanding Classifier Guidance and the Truncation Trick** DGMs often employ a mechanism for trading off sample diversity and sample quality. Two of the most widely used are classifier guidance [69, 72] for diffusion models, and the truncation trick [4] for GANs. In this section, we use the labeled KGEL2 test to help us better understand the effect of these methods. We use the Ablated Diffusion Model (ADM) [13] to measure the effect of classifier guidance, and BigGAN-deep [4] to measure the effect of the truncation trick. Fig. 6 show the results when the guidance parameter varies from 0.0 to 10.0, and the truncation parameter  $\tau$  varies from 0.2 to 1.0. As the truncation parameter decreases, unsurprisingly, the number of examples with 0 data probability increases dramatically

Table 2. GEL for Different Models on ImageNet  $256 \times 256$ . Numbers reported as  $2^{D(\hat{P}_n \| P_{\pi^*})}$ . For two-sample tests, the numbers reported are model/test  $2^{D(\hat{P}_n \| P_{\pi^*})}$ . Table 4 contains a complete set of results. VQGAN\* denotes parameters  $k = 600$ ,  $t = 1.0$ ,  $p = 0.92$  and VQGAN\*\* parameters  $k = 600$ ,  $a = 0.05$ ,  $p = 1.0$ .

PAPER MODEL	KGEL	KGEL	KGEL2	KGEL2
	EL	ET	EL	ET
- THEOR. OPT	1.0	1.0	1.0/1.0	1.0/1.0
- TRAINING SET	1.138	1.164	1.020/1.018	1.021/1.017
[4] BIGGAN-DEEP- $\tau=1.0$	<b>1.735</b>	<b>2.075</b>	<b>1.150/1.166</b>	1.166/1.222
[4] BIGGAN-DEEP- $\tau=0.6$	2.316	3.017	1.224/1.271	1.260/1.404
[59] VQ-VAE2	$+\infty$	40.55	2.933/2.730	5.851/5.723
[17] VQGAN*	4.443	9.034	1.295/1.495	1.487/1.717
[17] VQGAN**	1.772	2.219	1.175/1.202	<b>1.148/1.158</b>
[13] ADM	2.035	2.994	1.180/1.208	1.255/1.244
[13] ADM-G (1.0)	1.786	2.289	1.155/ <b>1.151</b>	1.185/1.188
[28] CDM	1.857	2.467	1.161/1.166	1.210/1.204

(from 14 at  $\tau = 1.0$  to 3420 at  $\tau = 0.2$ ). For classifier guidance, we find that optimal performance occurs at scale 1.0, while increasing the scale beyond this point increases the number of examples with 0 data probability to 650.

**Model Comparison** We perform a comparison of many popular ImageNet  $256 \times 256$  models. Tab. 2 shows the result for some representative models, while we include a full set of results in Tab. 4 (we also include results for CIFAR-10 in Appendix E.2). To help delineate between the performance of different models, we report results as  $2^{D_{KL}(\hat{P}_n \| P_{\pi})}$  and  $2^{D_{KL}(P_{\pi} \| \hat{P}_n)}$  for empirical likelihood, and exponentially tilting GEL, respectively. We find that performance of BigGAN-deep, ADM with classifier guidance, and VQGAN with 0.05 acceptance rate perform similarly.

## 5. Related Work

In contrast to the amount of research focused on improving generative models, comparatively little is focused on generative model evaluation metrics. There are, however, some notable exceptions. Inception Score (IS) [65] was likely the first broadly used evaluation metric, and measured sample quality and class diversity. Followup work such as the Modified Inception Score [24] sought to address IS's inability

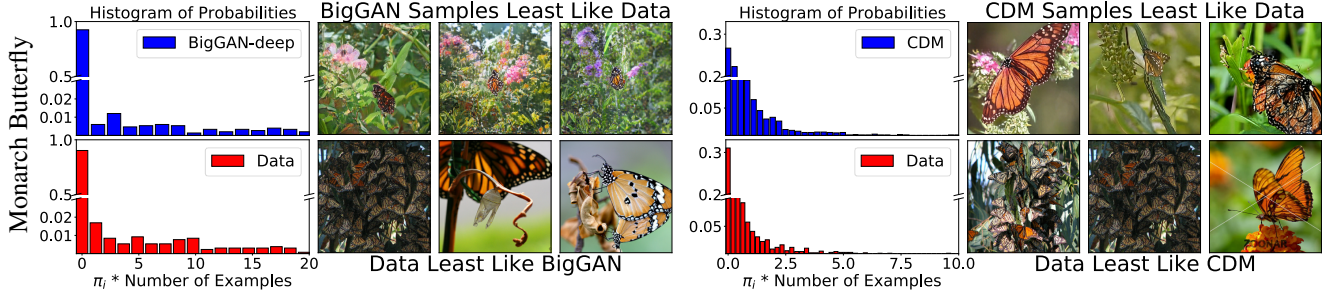


Figure 5. Model and test probabilities of the KGEL2 test can be used to identify data that the model cannot represent, and samples outside the data distribution. We identify examples with 0 model and data probabilities for BigGAN-deep (left) and Cascaded Diffusion Models (right). The blue and red histograms are those for model and data probabilities, respectively. The three top-right images are model samples least like the data (0 model probability), and the bottom-right are examples from the data the model cannot represent (0 data probability).

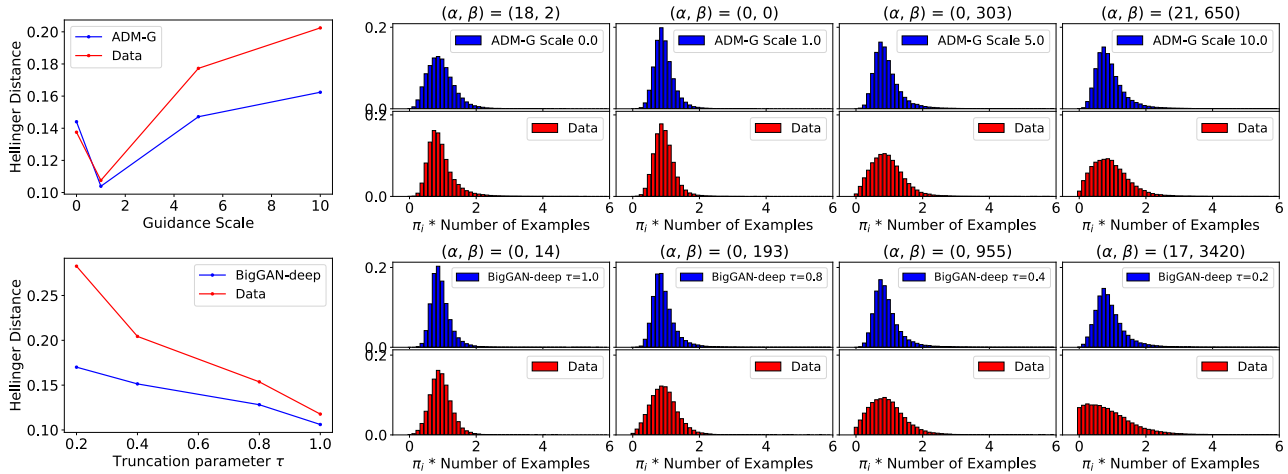


Figure 6. Right: Histogram of two-sample probabilities for different ADM classifier Guidance Scales and BigGAN-deep truncation probabilities. We use the labeled KGEL2 test with 1000-dim label-balanced ImageNet-v2 witness. The histograms show how increasing the classifier guidance scale and decreasing the truncation parameter increases the number of examples with 0.0 data probability.  $\alpha$  is the number of DGM samples with 0.0 model probability, and  $\beta$  is the number of examples with 0.0 data probability. Left: Hellinger distance between weighted DGM and empirical distributions (blue); and weighted data and empirical distributions (left).

to measure within-class diversity. This class of metrics has largely been replaced by Fréchet Inception Distance [26], now the most popular metric for comparing DGMs. While FID addressed many of the issues of IS, it suffered from statistical bias. Later work [2,9] proposed unbiased alternatives.

As deep generative models have improved, a number of authors have proposed more nuanced metrics. A popular approach has been to adopt precision and recall (PR) metrics [15,42,47,64]. These approaches first estimate approximate manifolds of the data and model distributions, and then determine how many data and model samples lie in the model and data manifolds, respectively. The manifold estimation step, however, is highly sensitive to hyperparameters and number of samples. Other “manifold” methods include the Geometry Score [36] and the Intrinsic Multi-Scale Distance [76].

Metrics that do not fit the above categories include a metric on human evaluation of sample quality [82], accuracy (on real data) of classifiers trained on DGM data [18,57,66,67,81], and for GANs, metrics based on its latent space [35].

Even less literature exists in the machine learning community on Empirical Likelihood methods. Authors proposed an EL test with linear-time MMD constraints [14] for testing simple distributions (such as the Normal). GEL approaches have been proposed for distributionally-robust optimization [16,43] and for off-policy evaluation [12,32]. Finally, authors introduced lower bounds on GEL objectives with functional moment restrictions for parameter estimation [40].

## 6. Discussion

In this work, we proposed generalized empirical likelihood methods as a tool for a better evaluation of DGMs. We propose a set of interpretable tests that allow us to diagnose deficiencies such as mode dropping and improper label conditioning. Current results are promising, and the generality of the approach may lead to new tests. In particular, we are interested in using features of new modalities, such as text, to better evaluate models such as text-conditioned DGMs.



## References

- [1] Michael S Albergo, Gurtej Kanwar, Sébastien Racanière, Danilo J Rezende, Julian M Urban, Denis Boyda, Kyle Cranmer, Daniel C Hackett, and Phiala E Shanahan. Flow-based sampling for fermionic lattice field theories. *Physical Review D*, 104:114507, 2021. [1](#)
- [2] Miłkołaj Bińkowski, Danica J Sutherland, Michael Arbel, and Arthur Gretton. Demystifying MMD GANs. In *Proceedings of the International Conference on Learning Representations*, 2018. [3](#), [8](#)
- [3] Denis Boyda, Gurtej Kanwar, Sébastien Racanière, Danilo J Rezende, Michael S Albergo, Kyle Cranmer, Daniel C Hackett, and Phiala E Shanahan. Sampling using  $SU(N)$  gauge equivariant flows. *Physical Review D*, 103(7):074504, 2021. [1](#)
- [4] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale GAN training for high fidelity natural image synthesis. In *Proceedings of the International Conference on Learning Representations*, 2019. [7](#), [15](#), [16](#)
- [5] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners. In *Advances in Neural Information Processing Systems*, 2020. [1](#)
- [6] George Casella and Roger Berger. *Statistical Inference*. Duxbury Resource Center, 2001. [12](#)
- [7] Rewon Child. Very deep VAEs generalize autoregressive models and can outperform them on images. In *Proceedings of the International Conference on Learning Representations*, 2021. [16](#)
- [8] Rewon Child, Scott Gray, Alec Radford, and Ilya Sutskever. Generating long sequences with sparse transformers, 2019. [1](#)
- [9] Min Jin Chong and David Forsyth. Effectively unbiased FID and inception score and where to find them. In *Proceedings of the Conference on Computer Vision and Pattern Recognition*, 2020. [8](#)
- [10] Kacper P Chwialkowski, Aaditya Ramdas, Dino Sejdinovic, and Arthur Gretton. Fast two-sample testing with analytic representations of probability measures. In *Advances in Neural Information Processing Systems*, 2015. [3](#)
- [11] Noel Cressie and Timothy RC Read. Multinomial goodness-of-fit tests. *Journal of the Royal Statistical Society Series B (Methodological)*, pages 440–464, 1984. [2](#)
- [12] Bo Dai, Ofir Nachum, Yinlam Chow, Lihong Li, Csaba Szepesvári, and Dale Schuurmans. CoinDICE: Off-policy confidence interval estimation. In *Advances in Neural Information Processing Systems*, 2020. [8](#)
- [13] Prafulla Dhariwal and Alex Nichol. Diffusion models beat GANs on image synthesis. In *Advances in Neural Information Processing Systems*, 2021. [7](#), [15](#)
- [14] Lizhong Ding, Zhi Liu, Yu Li, Shizhong Liao, Yong Liu, Peng Yang, Ge Yu, Ling Shao, and Xin Gao. Linear kernel tests via empirical likelihood for high-dimensional data. In *Proceedings of the Conference on Artificial Intelligence*, 2019. [8](#)
- [15] Josip Djolonga, Mario Lucic, Marco Cuturi, Olivier Bachem, Olivier Bousquet, and Sylvain Gelly. Precision-recall curves using information divergence frontiers. In *Proceedings of the International Conference on Artificial Intelligence and Statistics*, 2020. [8](#)
- [16] John Duchi, Peter Glynn, and Hongseok Namkoong. Statistics of robust optimization: A generalized empirical likelihood approach. *Mathematics of Operations Research*, 46:946–969, 2021. [8](#)
- [17] Patrick Esser, Robin Rombach, and Bjorn Ommer. Taming transformers for high-resolution image synthesis. In *Proceedings of the Conference on Computer Vision and Pattern Recognition*, 2021. [7](#), [15](#)
- [18] Cristóbal Esteban, Stephanie L Hyland, and Gunnar Rätsch. Real-valued (medical) time series generation with recurrent conditional GANs. *arXiv preprint arXiv:1706.02633*, 2017. [8](#)
- [19] Kenji Fukumizu, Francis R Bach, and Michael I Jordan. Dimensionality reduction for supervised learning with reproducing kernel Hilbert spaces. *Journal of Machine Learning Research*, 5:73–99, 2004. [3](#)
- [20] Karan Goel, Albert Gu, Chris Donahue, and Christopher Re. It’s raw! Audio generation with state-space models. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvári, Gang Niu, and Sivan Sabato, editors, *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 7616–7633. PMLR, 17–23 Jul 2022. [1](#)
- [21] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in Neural Information Processing Systems*, 2014. [15](#)
- [22] Arthur Gretton, Karsten M Borgwardt, Malte J Rasch, Bernhard Schölkopf, and Alexander Smola. A kernel two-sample test. *Journal of Machine Learning Research*, 13:723–773, 2012. [2](#), [3](#)
- [23] Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Guo, Mohammad Gheshlaghi Azar, Bilal Piot, Koray Kavukcuoglu, Remi Munos, and Michal Valko. Bootstrap your own latent: A new approach to self-supervised learning. In *Advances in Neural Information Processing Systems*, 2020. [4](#)
- [24] Swaminathan Gurumurthy, Ravi Kiran Sarvadevabhatla, and Venkatesh Babu Radhakrishnan. Generative adversarial networks for diverse and limited data. In *Proceedings of the Conference on Computer Vision and Pattern Recognition*, 2017. [7](#)
- [25] Alistair R Hall. *Generalized method of moments*. Advanced Texts in Econometrics. Oxford University Press, 2005. [2](#)
- [26] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. GANs trained by a two

- time-scale update rule converge to a local Nash equilibrium. In *Advances in Neural Information Processing Systems*, 2017. [1](#), [8](#)
- [27] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *Advances in Neural Information Processing Systems*, 2020. [16](#)
- [28] Jonathan Ho, Chitwan Saharia, William Chan, David J Fleet, Mohammad Norouzi, and Tim Salimans. Cascaded diffusion models for high fidelity image generation. *Journal of Machine Learning Research*, 23(47):1–33, 2022. [7](#), [15](#)
- [29] H Hotelling. The generalization of student’s ratio. *Annals of Mathematical Statistics*, 1931. [2](#)
- [30] Wittawat Jitkrittum, Zoltán Szabó, Kacper P Chwialkowski, and Arthur Gretton. Interpretable distribution features with maximum testing power. *Advances in Neural Information Processing Systems*, 29, 2016. [3](#)
- [31] Bahman Kalantari. Randomized triangle algorithms for convex hull membership. *arXiv preprint arXiv:1410.3564*, 2014. [13](#), [14](#)
- [32] Nikos Karampatziakis, John Langford, and Paul Mineiro. Empirical likelihood for contextual bandits. In *Advances in Neural Information Processing Systems*, 2020. [8](#)
- [33] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. Progressive growing of GANs for improved quality, stability, and variation. In *Proceedings of the International Conference on Learning Representations*, 2018. [16](#)
- [34] Tero Karras, Miika Aittala, Janne Hellsten, Samuli Laine, Jaakko Lehtinen, and Timo Aila. Training generative adversarial networks with limited data. In *Advances in Neural Information Processing Systems*, 2020. [6](#), [15](#), [16](#)
- [35] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of the Conference on Computer Vision and Pattern Recognition*, 2018. [8](#)
- [36] Valentin Khruikov and Ivan Oseledets. Geometry score: A method for comparing generative adversarial networks. In *Proceedings of the International Conference on Machine Learning*, 2018. [8](#)
- [37] Jack Kiefer and Jacob Wolfowitz. Consistency of the maximum likelihood estimator in the presence of infinitely many incidental parameters. *The Annals of Mathematical Statistics*, pages 887–906, 1956. [2](#)
- [38] Diederik P Kingma and Max Welling. Auto-encoding variational Bayes. In *Proceedings of the International Conference on Learning Representations*, 2014. [15](#)
- [39] Yuichi Kitamura. Empirical likelihood methods in econometrics: Theory and practice. *Cowles Foundation Discussion Paper No. 1569*, 2006. [13](#)
- [40] Heiner Kremer, Jia-Jie Zhu, Krikamol Muandet, and Bernhard Schölkopf. Functional generalized empirical likelihood estimation for conditional moment restrictions. In *Proceedings of the International Conference on Machine Learning*, 2022. [8](#)
- [41] Alex Krizhevsky and Geoffrey Hinton. Learning Multiple Layers of Features from Tiny Images. Technical report, University of Toronto, 2009. [4](#)
- [42] Tuomas Kynkäänniemi, Tero Karras, Samuli Laine, Jaakko Lehtinen, and Timo Aila. Improved precision and recall metric for assessing generative models. In *Advances in Neural Information Processing Systems*, 2019. [1](#), [4](#), [8](#), [21](#)
- [43] Henry Lam and Enlu Zhou. The empirical likelihood approach to quantifying uncertainty in sample average approximation. *Operations Research Letters*, 45(4):301–307, 2017. [8](#)
- [44] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *Proceedings of the International Conference on Computer Vision*, 2015. [16](#)
- [45] Takeru Miyato, Toshiki Kataoka, Masanori Koyama, and Yuichi Yoshida. Spectral normalization for generative adversarial networks. In *Proceedings of the International Conference on Learning Representations*, 2018. [16](#)
- [46] Krikamol Muandet, Kenji Fukumizu, Bharath Sriperumbudur, and Bernhard Schölkopf. *Kernel Mean Embedding of Distributions: A Review and Beyond*. Foundations and Trends® in Machine Learning. Now Publishers, Inc., 2017. [5](#)
- [47] Muhammad Ferjad Naeem, Seong Joon Oh, Youngjung Uh, Yunjey Choi, and Jaejun Yoo. Reliable fidelity and diversity metrics for generative models. In *Proceedings of the International Conference on Machine Learning*, 2020. [4](#), [8](#)
- [48] Frank Noé, Simon Olsson, Jonas Köhler, and Hao Wu. Boltzmann generators: Sampling equilibrium states of many-body systems with deep learning. *Science*, 365(6457), 2019. [1](#)
- [49] Aäron van den Oord, Sander Dieleman, Heiga Zen, Karen Simonyan, Oriol Vinyals, Alex Graves, Nal Kalchbrenner, Andrew Senior, and Koray Kavukcuoglu. Wavenet: A generative model for raw audio. *arXiv preprint arXiv:1609.03499*, 2016. [1](#)
- [50] OpenAI. Gpt-4 technical report, 2023. [1](#)
- [51] Art Owen. Empirical likelihood ratio confidence regions. *The Annals of Statistics*, pages 90–120, 1990. [2](#)
- [52] Art B Owen. *Empirical Likelihood*. Chapman and Hall/CRC, 2001. [2](#), [12](#), [13](#)
- [53] Jing Qin. Semi-empirical likelihood ratio confidence intervals for the difference of two sample means. *Annals of the Institute of Statistical Mathematics*, 46(1):117–126, 1994. [5](#)
- [54] Jin Qin and Jerry Lawless. Empirical likelihood and general estimating equations. *The Annals of Statistics*, pages 300–325, 1994. [2](#)
- [55] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with CLIP latents. *arXiv preprint arXiv:2204.06125*, 2022. [1](#)
- [56] Suman Ravuri, Shakir Mohamed, Mihaela Rosca, and Oriol Vinyals. Learning implicit generative models with the method of learned moments. In *Proceedings of the International Conference on Machine Learning*, 2018. [16](#)
- [57] Suman Ravuri and Oriol Vinyals. Classification accuracy score for conditional generative models. In *Advances in Neural Information Processing Systems*, 2019. [8](#)
- [58] Ali Razavi, Aäron van den Oord, Ben Poole, and Oriol Vinyals. Preventing posterior collapse with delta-VAEs. In *Proceedings of the International Conference on Learning Representations*, 2019. [16](#)

- [59] Ali Razavi, Aaron van den Oord, and Oriol Vinyals. Generating diverse high-fidelity images with VQ-VAE-2. In *Advances in Neural Information Processing Systems*, 2019. 7, 15
- [60] Benjamin Recht, Rebecca Roelofs, Ludwig Schmidt, and Vaishaal Shankar. Do ImageNet classifiers generalize to ImageNet? In *Proceedings of the International Conference on Machine Learning*, 2019. 7
- [61] Danilo J Rezende, Shakir Mohamed, and Daan Wierstra. Stochastic backpropagation and approximate inference in deep generative models. In *Proceedings of the International Conference on Machine Learning*, 2014. 15
- [62] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the Conference on Computer Vision and Pattern Recognition*, 2022. 1
- [63] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily Denton, Seyed Kamyar Seyed Ghasemipour, Raphael Gontijo-Lopes, Burcu Karagol Ayan, Tim Salimans, Jonathan Ho, David J Fleet, and Mohammad Norouzi. Photo-realistic text-to-image diffusion models with deep language understanding. In *Advances in Neural Information Processing Systems*, 2022. 1
- [64] Mehdi SM Sajjadi, Olivier Bachem, Mario Lucic, Olivier Bousquet, and Sylvain Gelly. Assessing generative models via precision and recall. In *Advances in Neural Information Processing Systems*, 2018. 8, 21
- [65] Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training GANs. In *Advances in Neural Information Processing Systems*, 2016. 7
- [66] Shibani Santurkar, Ludwig Schmidt, and Aleksander Madry. A classification-based study of covariate shift in GAN distributions. In *Proceedings of the International Conference on Machine Learning*, 2018. 8
- [67] Konstantin Shmelkov, Cordelia Schmid, and Karteek Alahari. How good is my GAN? In *Proceedings of the European Conference on Computer Vision*, 2018. 8
- [68] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *Proceedings of the International Conference on Learning Representations*, 2015. 4
- [69] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *Proceedings of the International Conference on Machine Learning*, 2015. 7, 15
- [70] Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. In *Advances in Neural Information Processing Systems*, 2019. 16
- [71] Yang Song and Stefano Ermon. Improved techniques for training score-based generative models. In *Advances in Neural Information Processing Systems*, 2020. 16
- [72] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *Proceedings of the International Conference on Learning Representations*, 2021. 7
- [73] Zoltán Szabó and Bharath K Sriperumbudur. Characteristic and universal tensor product kernels. *Journal of Machine Learning Research*, 18:1–29, 2018. 5
- [74] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *Proceedings of the Conference on Computer Vision and Pattern Recognition*, 2016. 4
- [75] Min Tsao. Bounds on coverage probabilities of the empirical likelihood ratio confidence regions. *Annals of Statistics*, pages 1215–1221, 2004. 12
- [76] Anton Tsitsulin, Marina Munkhoeva, Davide Mottin, Panagiotis Karras, Alex Bronstein, Ivan Oseledets, and Emmanuel Mueller. The shape of data: Intrinsic distance for data distributions. In *Proceedings of the International Conference on Learning Representations*, 2020. 8
- [77] Arash Vahdat and Jan Kautz. NVAE: A deep hierarchical variational autoencoder. In *Advances in Neural Information Processing Systems*, 2020. 16
- [78] Jean-Philippe Vert, Hiroto Saigo, and Tatsuya Akutsu. Local alignment kernels for biological sequences. *Kernel Methods in Computational Biology*, pages 131–154, 2004. 5, 15
- [79] Martin J Wainwright, Michael I Jordan, et al. Graphical models, exponential families, and variational inference. *Foundations and Trends® in Machine Learning*, 1(1–2):1–305, 2008. 13
- [80] Peter Wirnsberger, George Papamakarios, Borja Ibarz, Sébastien Racanière, Andrew J Ballard, Alexander Pritzel, and Charles Blundell. Normalizing flows for atomic solids. *Machine Learning: Science and Technology*, 3(2):025009, 2022. 1
- [81] Jianwei Yang, Anitha Kannan, Dhruv Batra, and Devi Parikh. LR-GAN: Layered recursive generative adversarial networks for image generation. In *Proceedings of the International Conference on Learning Representations*, 2017. 8
- [82] Sharon Zhou, Mitchell Gordon, Ranjay Krishna, Austin Narcomey, Durim Morina, and Michael S Bernstein. HYPE: human-eye perceptual evaluation of generative models. In *Advances in Neural Information Processing Systems*, 2019. 8