
Autoregressive Quantile Networks for Generative Modeling

Georg Ostrovski^{*1} Will Dabney^{*1} Rémi Munos¹

Abstract

We introduce autoregressive implicit quantile networks (AIQN), a fundamentally different approach to generative modeling than those commonly used, that implicitly captures the distribution using quantile regression. AIQN is able to achieve superior perceptual quality and improvements in evaluation metrics, without incurring a loss of sample diversity. The method can be applied to many existing models and architectures. In this work we extend the PixelCNN model with AIQN and demonstrate results on CIFAR-10 and ImageNet using Inception score, FID, non-cherry-picked samples, and inpainting results. We consistently observe that AIQN yields a highly stable algorithm that improves perceptual quality while maintaining a highly diverse distribution.

1. Introduction

There has been a staggering increase in progress on generative modeling in recent years, built largely upon fundamental advances such as generative adversarial networks (Goodfellow et al., 2014), variational inference (Kingma & Welling, 2013), and autoregressive density estimation (van den Oord et al., 2016c). These have led to breakthroughs in state-of-the-art generation of natural images (Karras et al., 2017) and audio (van den Oord et al., 2016a), and even been used for unsupervised learning of disentangled representations (Higgins et al., 2017; Chen et al., 2016). These domains often have real-valued distributions with underlying metrics; that is, there is a domain-specific notion of similarity between data points. This similarity is ignored by the predominant work-horse of generative modeling, the Kullback-Leibler (KL) divergence. Progress is now being made towards algorithms that optimize with respect to these underlying metrics (Arjovsky et al., 2017; Bousquet et al., 2017).

^{*}Equal contribution ¹DeepMind, London, UK. Correspondence to: Georg Ostrovski <ostrovski@google.com>, Will Dabney <wdabney@google.com>.

In this paper, we present a novel approach to generative modeling, that, while strikingly different from existing methods, is grounded in the well-understood statistical methods of *quantile regression*. Unlike the majority of recent work, we approach generative modeling without the use of the KL divergence, and without explicitly approximating a likelihood model. Like GANs, in this way we produce an implicitly defined model, but unlike GANs our optimization procedure is inherently stable and lacks degenerate solutions which cause loss of diversity and mode collapse.

Much of the recent research on GANs has been focused on improving stability (Radford et al., 2015; Arjovsky et al., 2017; Daskalakis et al., 2017) and sample diversity (Gulrajani et al., 2017; Salimans et al., 2016; 2018). By stark contrast, methods such as PixelCNN (van den Oord et al., 2016b) readily produce high diversity, but due to their use of KL divergence are unable to make reasonable trade-offs between likelihood and perceptual similarity (Theis et al., 2015; Bellemare et al., 2017; Bousquet et al., 2017).

Our proposed method, *autoregressive implicit quantile networks* (AIQN), combines the benefits of both: a loss function that respects the underlying metric of the data leading to improved perceptual quality, and a stable optimization process leading to highly diverse samples. While there has been an increasing tendency towards complex architectures (Chen et al., 2017; Salimans et al., 2017) and multiple objective loss functions to overcome these challenges, AIQN is conceptually simple and does not rely on any special architecture or optimization techniques. Empirically it proves to be robust to hyperparameter variations and easy to optimize.

Our work is motivated by the recent advances achieved by reframing GANs in terms of optimal transport, leading to the Wasserstein GAN algorithm (Arjovsky et al., 2017), as well as work towards understanding the relationship between optimal transport and both GANs and VAEs (Bousquet et al., 2017). In agreement with these results, we focus on loss functions grounded in perceptually meaningful metrics. We build upon recent work in distributional reinforcement learning (Dabney et al., 2018a), which has begun to bridge the gap between approaches in reinforcement learning and unsupervised learning. Towards a practical algorithm we base our experimental results on Gated PixelCNN (van den Oord et al., 2016b), and show that using AIQN significantly im-

proves objective performance on CIFAR-10 and ImageNet 32x32 in terms of Fréchet Inception Distance (FID) and Inception score, as well as subjective perceptual quality in image samples and inpainting.

2. Background

We begin by establishing some notation, before turning to a review of three of the most prevalent methods for generative modeling. Calligraphic letters (e.g. \mathcal{X}) denote sets or spaces, capital letters (e.g. X) denote random variables, and lower case letters (e.g. x) indicate values. A probability distribution with random variable $X \in \mathcal{X}$ is denoted $p_X \in \mathcal{P}(\mathcal{X})$, its cumulative distribution function (c.d.f.) F_X , and inverse c.d.f. or quantile function $Q_X = F_X^{-1}$. When probability distributions or quantile functions are parameterized by some θ we will write p_θ or Q_θ recognizing that here we do not view θ as a random variable.

Perhaps the simplest way to approach generative modeling of a random variable $X \in \mathcal{X}$ is by fixing some discretization of \mathcal{X} into n separate values, say $x_1, \dots, x_n \in \mathcal{X}$, and parameterize the approximate distribution with $p_\theta(x_i) \propto \exp(\theta_i)$. This type of categorical parameterization is widely used, only slightly less commonly when \mathcal{X} does not lend itself naturally to such a partitioning. Typically, the parameters θ are optimized to minimize the Kullback-Leibler (KL) divergence between observed values of X and the model p_θ , $\theta^* = \arg \min_\theta D_{KL}(p_X \| p_\theta)$.

However, this is only tractable when \mathcal{X} is a small discrete set or at best low-dimensional. A common method for extending a generative model or density estimator to multivariate distributions is to factor the density as a product of scalar-valued conditional distributions. Let $X = (X_1, \dots, X_n)$, then for any permutation of the dimensions $\sigma: \mathbb{N}_n \rightarrow \mathbb{N}_n$,

$$p_X(x) = \prod_{i=1}^n p_{X_{\sigma(i)}}(x_{\sigma(i)} | x_{\sigma(1)}, \dots, x_{\sigma(i-1)}). \quad (1)$$

When the conditional density is modeled by a simple (e.g. Gaussian) base distribution, the ordering of the dimensions can be crucial (Papamakarios et al., 2017). However, it is common practice to choose an arbitrary ordering and rely upon a more powerful conditional model to avoid these problems. This class of models includes PixelRNN and PixelCNN (van den Oord et al., 2016c;b), MAF (Papamakarios et al., 2017), MADE (Germain et al., 2015), and many others. Fundamentally, all these approaches use the KL divergence as their loss function.

Another class of methods, generally known as *latent variable methods*, can bypass the need for autoregressive models using a different modeling assumption. Specifically, consider the Variational Autoencoder (VAE) (Kingma & Welling, 2013; Rezende et al., 2014), which represents

p_θ as the marginalization over a latent random variable $Z \in \mathcal{Z}$. The VAE is trained to maximize an approximate lower bound of the log-likelihood of the observations:

$$\log p_\theta(x) \geq -D_{KL}(q_\theta(z|x) \| p(z)) + \mathbb{E} [\log p_\theta(x|z)].$$

Although VAEs are straightforward to implement and optimize, and effective at capturing structure in high-dimensional spaces, they often miss fine-grained detail, resulting in blurry images.

Generative Adversarial Networks (GANs) (Goodfellow et al., 2014) pose the problem of learning a generative model as a two-player zero-sum game between a discriminator D , attempting to distinguish between $x \sim p_X$ (real data) and $x \sim p_\theta$ (generated data), and a generator G , attempting to generate data indistinguishable from real data. The generator is an implicit latent variable model that reparameterizes samples, typically from an isotropic Gaussian distribution, into values in \mathcal{X} . The original formulation of GANs,

$$\arg \min_G \sup_D \left[\mathbb{E}_X \log(D(X)) + \mathbb{E}_Z \log(1 - D(G(Z))) \right],$$

can be seen as minimizing a lower-bound on the Jensen-Shannon divergence (Goodfellow et al., 2014; Bousquet et al., 2017). That is, even in the case of GANs we are often minimizing functions of the KL divergence¹.

Many recent advances have come from principled combinations of these three fundamental methods (Makhzani et al., 2015; Dumoulin et al., 2016; Rosca et al., 2017).

2.1. Distance Metrics and Loss Functions

A common perspective in generative modeling is that the choice of model should encode existing metric assumptions about the domain, combined with a generic likelihood-focused loss such as the KL divergence. Under this view, the KL's general applicability and robust optimization properties make it a natural choice, and most implementations of the methods we reviewed in the previous section attempt to, at least indirectly, minimize a version of the KL.

On the other hand, as every model inevitably makes trade-offs when constrained by capacity or limited training, it is desirable for its optimization goal to incentivize trade-offs prioritizing approximately correct solutions, when the data space is endowed with a metric supporting a meaningful (albeit potentially subjective) notion of approximation. It has been argued (Theis et al., 2015; Bousquet et al., 2017; Arjovsky et al., 2017; Bellemare et al., 2017) that the KL may not always be appropriate from this perspective, by making sub-optimal trade-offs between likelihood and similarity.

¹The Jensen-Shannon divergence is the sum of KLs between distributions P, Q and their uniform mixture $M = 0.5(P + Q)$: $JSD(P||Q) = 0.5(D_{KL}(P||M) + D_{KL}(Q||M))$.




Figure 1. Categorization of generative models by discretization vs. reparameterization and loss functions by likelihood-based vs. quantile-regression-based.

Indeed, many limitations of existing models can be traced back to the use of KL, and the resulting trade-offs in approximate solutions it implies. For instance, its use appears to play a central role in one of the primary failure modes of VAEs, that of blurry samples. Zhao et al. (2017) argue that the Gaussian posterior $p_\theta(x|z)$ implies an overly simple model, which, when unable to perfectly fit the data, is forced to average (thus creating blur), and is not incentivized by the KL towards an alternative notion of approximate solution. Theis et al. (2015) emphasized that an improvement of log-likelihood does not necessarily translate to higher perceptual quality, and that the KL loss is more likely to produce atypical samples than some other training criteria.

We offer an alternative perspective: a good model should encode assumptions about the data distribution, whereas a good loss should encode the notion of similarity, that is, the underlying metric on the data space. From this point of view, the KL corresponds to an actual absence of explicit underlying metric, with complete focus on probability.

The optimal transport metrics W_c , for underlying metric $c(x, x')$, and in particular the p -Wasserstein distance, when c is an L_p metric, have frequently been proposed as being well-suited replacements to KL (Bousquet et al., 2017; Genevay et al., 2017). Briefly, the advantages are (1) avoidance of mode collapse (no need to choose between spreading over modes or collapsing to a single mode as in KL), and (2) the ability to trade off errors and incentivize approximations that respect the underlying metric.

Recently, Arjovsky et al. (2017) introduced the Wasserstein

GAN, reposing the two-player game as the estimation of the gradient of the 1-Wasserstein distance between the data and generator distributions. They reframe this in terms of the dual form of the 1-Wasserstein, with the critic estimating a function f which maximally separates the two distributions. While this is an exciting line of work, it still faces limitations when the critic solution is approximate, i.e. when f^* is not found before each update. In this case, due to insufficient training of the critic (Bellemare et al., 2017) or limitations of the function approximator, the gradient direction produced can be arbitrarily bad (Bousquet et al., 2017).

Thus, we are left with the question of how to minimize a distribution loss respecting an underlying metric. Recent work in distributional reinforcement learning has proposed the use of *quantile regression* as a method for minimizing the 1-Wasserstein in the univariate case when approximating using a mixture of Dirac functions (Dabney et al., 2018b).

2.2. Quantile Regression

In this section, we review quantile regression as a method for estimating the quantile function of a distribution at specific points, i.e. its inverse cumulative distribution function. This leads to recent work on approximating a distribution by a neural network approximation of its quantile function, acting as a reparameterization of a random sample from the uniform distribution.

The quantile regression loss (Koenker & Hallock, 2001) for a quantile at $\tau \in [0, 1]$ and error u (positive for underestimation and negative for overestimation) is given by $\rho_\tau(u) = (\tau - \mathbb{I}\{u \leq 0\})u$. It is an asymmetric loss function penalizing underestimation by weight τ and overestimation by weight $1 - \tau$. For a given scalar distribution Z with c.d.f. F_Z and a quantile τ , the inverse c.d.f. $q = F_Z^{-1}(\tau)$ minimizes the expected quantile regression loss $\mathbb{E}_{z \sim Z} [\rho_\tau(z - q)]$.

Using this loss allows one to train a neural network to approximate a scalar distribution represented by its inverse c.d.f. For this, the network can output a fixed grid of quantiles (Dabney et al., 2018b), with the respective quantile regression losses being applied to each output independently. A more effective approach is to provide the desired quantile τ as an additional input to the network, and train it to output the corresponding value of $F_Z^{-1}(\tau)$. The *implicit quantile network* (IQN) model (Dabney et al., 2018a) reparameterizes a sample $\tau \sim \mathcal{U}([0, 1])$ through a deterministic function to produce samples from the underlying data distribution. These two methods can be seen to belong to the top-right and bottom-right categories in Figure 1. An IQN Q_θ can be trained by stochastic gradient descent on the quantile regression loss, with $u = z - Q_\theta(\tau)$ and training samples (z, τ) drawn from $z \sim Z$ and $\tau \sim \mathcal{U}([0, 1])$.

One drawback to the quantile regression loss is that gradients do not scale with the magnitude of the error, but instead with the sign of the error and the quantile weight τ . This increases gradient variance and can negatively impact the final model’s sample quality. Increasing the batch size, and thus averaging over more values of τ , would have the effect of lowering this variance. Alternatively, we can smooth the gradients as the model converges by allowing errors, under some threshold κ , to be scaled with their magnitude, reverting to an *expectile* loss. This results in the Huber quantile loss (Huber, 1964; Dabney et al., 2018b):

$$\rho_\tau^\kappa(u) = \begin{cases} \frac{|\tau - \mathbb{I}\{u \leq 0\}|}{2\kappa} u^2, & \text{if } |u| \leq \kappa, \\ |\tau - \mathbb{I}\{u \leq 0\}|(|u| - \frac{1}{2}\kappa), & \text{otherwise.} \end{cases} \quad (2)$$

3. Autoregressive Implicit Quantiles

Let $X = (X_1, \dots, X_n) \in \mathcal{X}_1 \times \dots \times \mathcal{X}_n = \mathcal{X}$ be an n -dimensional random variable. We begin by analyzing the effect of two naive applications of IQN to modeling the distribution of X .

First, suppose we use the same quantile target, $\tau \in [0, 1]$, for every output dimension. The only modification to IQN would be to output n dimensions instead of 1, the loss being applied to each output dimension independently. This is equivalent to assuming that the dimensions of X are *comonotonic*. Two random variables are comonotonic if and only if they can be expressed as non-decreasing (deterministic) functions of a single random variable (Dhaene et al., 2006). Thus a joint quantile function for a comonotonic X can be written as $F_X^{-1}(\tau) = (F_{X_1}^{-1}(\tau), F_{X_2}^{-1}(\tau), \dots, F_{X_n}^{-1}(\tau))$. While there are many interesting uses for comonotonic random variables, we believe this assumption is too strong to be useful more broadly.

Second, one could use a separate value $\tau_i \in [0, 1]$ for each X_i , with the IQN being unchanged from the first case. This corresponds to making an independence assumption on the dimensions of X . Again we would expect this to be an unreasonably restrictive modeling assumption for many domains, such as the case of natural images.

Now, we turn to our proposed approach of extending IQN to multivariate distributions. We fix an ordering of the n dimensions. If the density function p_X is expressed as a product of conditional likelihoods, as in Equation 1, then the joint c.d.f. can be written as

$$\begin{aligned} F_X(x) &= \mathbb{P}(X_1 \leq x_1, \dots, X_n \leq x_n), \\ &= \prod_{i=1}^n F_{X_i | X_{i-1}, \dots, X_1}(x_i). \end{aligned}$$

Furthermore, for $\tau_{joint} = \prod_{i=1}^n \tau_i$, we can write the *joint-quantile function* of X as

$$F_X^{-1}(\tau_{joint}) = (F_{X_1}^{-1}(\tau_1), \dots, F_{X_n | X_{n-1}, \dots}^{-1}(\tau_n)).$$

This approach has been used previously by Koenker & Xiao (2006), who introduced a quantile autoregression model for quantile regression on time-series.

We propose to extend IQN to an autoregressive model of the above conditional form of a joint-quantile function. Denoting $\mathcal{X}_{1:i} = \mathcal{X}_1 \times \dots \times \mathcal{X}_i$, let $\tilde{\mathcal{X}} := \bigcup_{i=0}^n \mathcal{X}_{1:i}$ be the space of ‘partial’ data points. We can define the autoregressive IQN as a deterministic function $Q_\theta: \tilde{\mathcal{X}} \times [0, 1]^n \rightarrow \tilde{\mathcal{X}}$, mapping partial samples $\tilde{x} \in \tilde{\mathcal{X}}$ and quantile targets $\tau_i \in [0, 1]$ to estimates of F_X^{-1} . We can then train Q_θ using a quantile regression loss (Equation 2). For generation, one can iterate $x_{1:i} = Q_\theta(x_{1:i-1}, \tau_i)$, on a sequence of growing partial samples² $x_{1:i-1}$ and independently sampled $\tau_i \sim \mathcal{U}([0, 1])$, for $i = 1, \dots, n$, to finally obtain a sample $x = x_{1:n}$.

3.1. Quantile Regression and the Wasserstein

As previously mentioned, for the restricted model class of a uniform mixture of Diracs, quantile regression can be shown to minimize the 1-Wasserstein metric (Dabney et al., 2018b). We extend this analysis for the case of arbitrary approximate quantile functions, and find that quantile regression minimizes a closely related divergence which we call *quantile divergence*, defined, for any distributions P and Q , as

$$q(P, Q) := \int_0^1 \left[\int_{F_P^{-1}(\tau)}^{F_Q^{-1}(\tau)} (F_P(x) - \tau) dx \right] d\tau.$$

Indeed, the expected quantile loss of any parameterized quantile function \bar{Q}_θ equals, up to a constant, the quantile divergence between P and the distribution Q_θ implicitly defined by \bar{Q}_θ :

$$\mathbb{E}_{\tau \sim \mathcal{U}([0, 1])} \left[\mathbb{E}_{z \sim P} [\rho_\tau(z - \bar{Q}_\theta(\tau))] \right] = q(P, Q_\theta) + h(P),$$

where $h(P)$ does not depend on Q_θ . Thus quantile regression minimizes the quantile divergence $q(P, Q_\theta)$ and the sample gradient $\nabla_\theta \rho_\tau(z - \bar{Q}_\theta(\tau))$ (for $\tau \sim \mathcal{U}([0, 1])$ and $z \sim P$) is an unbiased estimate of $\nabla_\theta q(P, Q_\theta)$. See Appendix for proofs.

3.2. Quantile Density Function

Although IQN does not directly model the log-likelihood of the data distribution, observe that we can still query the implied density at a point (Jones, 1992):

$$\frac{\partial}{\partial \tau} F_X^{-1}(\tau) = \frac{1}{p_X(F_X^{-1}(\tau))}.$$

Indeed, this quantity, known as the *sparsity function* (Tukey, 1965) or the *quantile-density function* (Parzen, 1979) plays

²Throughout we understand $x_0 = x_{1:0} \in \mathcal{X}_{1:0}$ to denote the ‘empty tuple’, and the function Q_θ to map this to a single unconditional sample $x_1 = x_{1:1} = Q_\theta(x_0, \tau_1)$.

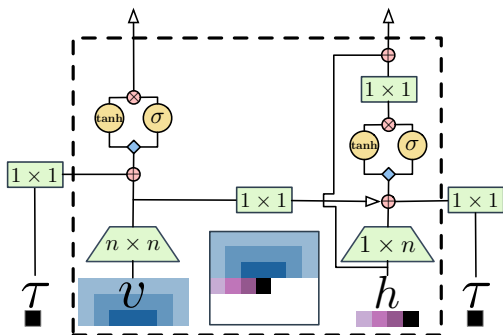


Figure 2. Illustration of Gated PixelCNN layer block for PixelIQN. Dashed line shows boundary of standard Gated PixelCNN, with v the vertical and h the horizontal stack. Conditioning on τ is identical to the location-dependent conditioning in Gated PixelCNN.

a central role in the analysis of quantile regression models (Koenker, 1994). A common approach involves choosing a bandwidth parameter h and estimating this quantity through finite-differences around the value of interest as $(F_X^{-1}(\tau+h) - F_X^{-1}(\tau-h))/2h$ (Siddiqui, 1960). However, as we have the full quantile function, the quantile-density function can be computed exactly using a single step of back-propagation to compute $\frac{\partial F_X^{-1}(\tau)}{\partial \tau}$. As this only allows querying the density given the value of τ , application to general likelihoods would require finding the value of τ that produces the closest approximation to the query point x . Though arguably too inefficient for training, this could potentially be used to interrogate the model.

4. PixelIQN

To test our proposed method, which is architecturally compatible with many generative model approaches, we wanted to compare and contrast IQN, that is quantile regression and quantile reparameterization, with a method trained with an explicit parameterization to minimize KL divergence. A natural choice for this was PixelCNN, specifically we build upon the Gated PixelCNN of van den Oord et al. (2016b).

The Gated PixelCNN takes as input an image $x \sim X$, sampled from the training distribution at training time, and potentially all zeros or partially generated at generation time, as well as a location-dependent context s . The model consists of a number of residual layer blocks, whose structure is chosen to allow each output pixel to be a function of all preceding input pixels (in a raster-scan order). At its core, each layer block computes two gated activations of the form

$$y = \tanh(W_{k,f} * x + V_{k,f} * s) \odot \sigma(W_{k,g} * x + V_{k,g} * s),$$

with k the layer index, $*$ denoting convolution, and $V_{k,f}$ and $V_{k,g}$ being 1×1 convolution kernels. See Figure 2 for a full schematic depiction of a Gated PixelCNN layer

block. After a number of such layer blocks, the PixelCNN produces a final output layer with shape $(n, n, 3, 256)$, with a softmax across the final dimension, corresponding to the approximate conditional likelihood for the value of each pixel-channel. That is, the conditional likelihood is the product of these individual autoregressive models,

$$p(x|s) = \prod_{i=1}^{3n^2} p(x_i|x_1, \dots, x_{i-1}, s_i).$$

Typically the location-dependent conditioning term was used to condition on class labels, but here, we will use it to condition on the sample point³ $\tau \in [0, 1]^{3n^2}$. Thus, in addition to the input image x we input, in place of s , the sample points $\tau = (\tau_1, \dots, \tau_{3n^2})$, to be reparameterized, with each $\tau_i \sim \mathcal{U}([0, 1])$. Finally, our network outputs only the full sample image of shape $(n, n, 3)$, without the need for an additional softmax layer. Note that the number of τ values generated exactly corresponds to the number of random draws from softmax distributions in the original PixelCNN. We are simply changing the role of the randomness, from a draw at the output to a part of the input.

Architecturally, our proposed model, PixelIQN, is exactly the network given by van den Oord et al. (2016b), with the one exception that we output only a single value per pixel-channel and do not require the softmax activations.

In PixelCNN training is done by passing the training image through the network, and training each output softmax distribution using the KL divergence between the training image and the approximate distribution,

$$\sum_i D_{KL}(\delta_{x_i}, p(\cdot|x_1, \dots, x_{i-1})).$$

For PixelIQN, the input is the training image x and a sample point $\tau \sim \mathcal{U}([0, 1]^{3n^2})$. The output values $Q_x(\tau) \in \mathbb{R}^{3n^2}$ are interpreted as the approximate quantile function at τ , $Q_x(\tau)_i = Q_X(\tau_i|x_{i-1}, \dots)$, trained with a single step of quantile regression towards the observed sample x :

$$\sum_i \rho_{\tau_i}^k(x_i - Q_X(\tau_i|x_{i-1}, \dots)).$$

4.1. CIFAR-10

We begin by demonstrating PixelIQN on CIFAR-10 (Krizhevsky & Hinton, 2009). For comparison, we train both a baseline Gated PixelCNN and a PixelIQN. Both models correspond to the 15-layer network variant in (van den Oord et al., 2016b), see Appendix for detailed hyperparameters and training procedure. The two methods have substantially different loss functions, so we performed a

³Conditioning on labels remains possible (see Section 4.2).




Figure 3. CIFAR-10: Real example images (left), samples generated by PixelCNN (center), and samples generated by PixelIQN (right).




Figure 4. Evaluations by Inception score (higher is better) and FID (lower is better) on CIFAR-10 and ImageNet 32x32. Dotted lines correspond to models trained with class-label conditioning.

hyperparameter search using a short training run, with the same number (500) of hyperparameter configurations evaluated for both models. For all results, we report full training runs using the best found hyperparameters in each case. The evaluation metric used for the hyperparameter search was the Fréchet Inception Distance (FID) (Heusel et al., 2017), see Appendix for details. In addition to FID, we report Inception score (Salimans et al., 2016) for both models.

Figure 4 (left) shows Inception score and FID for both models evaluated at several points throughout training. The fully trained PixelCNN achieves an Inception score and FID of 4.6 and 65.9 respectively, while PixelIQN substantially outperforms it with an Inception score of 5.3 and FID of 49.5. This also compares favorably with e.g. WGAN (Arjovsky et al., 2017), which reaches an Inception score of 3.8. For subjective evaluations, we give samples from both models in Figure 3. Samples coming from PixelIQN are much more visually coherent. Of note, the PixelIQN model achieves

Method	CIFAR-10		ImageNet (32x32)	
	Inception	FID	Inception	FID
WGAN	3.82	-	-	-
WGAN-GP	6.5	36.4	-	-
DC-GAN	6.4	37.11	7.89	-
PixelCNN	4.60	65.93	7.16	40.51
PixelIQN	5.29	49.46	8.68	26.56
PixelIQN(l)	-	-	7.29	37.62
PixelCNN*	-	-	8.33	33.27
PixelIQN*	-	-	10.18	22.99

Table 1. Inception score and FID for CIFAR-10 and ImageNet. WGAN and DC-GAN results taken from (Arjovsky et al., 2017; Radford et al., 2015). PixelIQN(l) is the small 15-layer version of the model. Models marked * refer to class-conditional training.

a performance level comparable to that of the fully trained PixelCNN with only about one third the number of training updates (and about one third of the wall-clock time).

4.2. ImageNet 32x32

Next, we turn to the small ImageNet dataset (Russakovsky et al., 2015), first used for generative modeling in the PixelRNN work (van den Oord et al., 2016c). Again, we evaluate using FID and Inception score. For this much harder dataset, we base our PixelCNN and PixelIQN models on the larger 20-layer variant used in (van den Oord et al., 2016b). Due to substantially longer training time for this model, we did not perform additional hyperparameter tuning, and mostly used the same hyperparameter values as in the previous sections for both models; details can be found in the Appendix.

Figure 4 shows Inception score and FID throughout training of PixelCNN and PixelIQN. Again, PixelIQN substan-




Figure 5. ImageNet 32x32: Real example images (left), samples generated by PixelCNN (center), and samples generated by PixelIQN (right). Neither of the sampled image sets were cherry-picked. More samples by PixelIQN in the Appendix.

tially outperforms the baseline in terms of final performance and sample complexity. For final scores and a comparison to state-of-the-art GAN models, see Table 1. Figure 5 shows random (non-cherry-picked) samples from both models. Compared to PixelCNN, PixelIQN samples appear to have superior quality with more global consistency and less ‘high-frequency noise’.

In Figure 6, we show the inpainting performance of PixelIQN, by fixing the top half of a validation set image as input and sampling repeatedly from the model to generate different completions. We note that the model consistently generates plausible completions with significant diversity between different completion samples for the same input image. Meanwhile, WGAN-GP has been seen to produce deterministic completions (Bellemare et al., 2017).

Following (van den Oord et al., 2016b), we also trained a class-conditional PixelIQN variant, providing to the model the one-hot class label corresponding to a training image (in addition to a τ sample). Samples from a class-conditional model can be expected to have higher visual quality, as the class label provides $\log_2(1000) \approx 10$ bits of information, see Figure 7. As seen in Figure 4 and Table 1, class conditioning also further improves Inception score and FID. To generate each sample for the computation of these scores, we sample one of 1000 class labels randomly, then generate an image conditioned on this label via the trained model.

Finally, motivated by the very long training time for the large PixelCNN model (approximately 1 day per 100K training steps, on 16 NVIDIA Tesla P100 GPUs), we also trained smaller 15-layer versions of the models (same as the ones used on CIFAR-10) on the small ImageNet dataset. For comparison, these take approximately 12 hours for 100K training steps on a single P100 GPU, or less than 3 hours on 8 P100 GPUs. As expected, little PixelCNN, while suitable




Figure 6. Small ImageNet inpainting examples. Left image is the input provided to the network at the beginning of sampling, right is the original image, columns in between show different completions. More examples in the Appendix.

for the CIFAR-10 dataset, fails to achieve competitive scores on the ImageNet dataset, achieving Inception score 5.1 and FID 66.4. Astonishingly, little PixelIQN on this dataset reaches Inception score 7.3 and FID 38.5, see Figure 4 (right). It thereby not only outperforms the little PixelCNN, but also the larger 20-layer version! This strongly supports the hypothesis that PixelCNN, and potentially many other models, are constrained not only by their model capacity, but crucially also by the sub-optimal trade-offs made by their log-likelihood training criterion, failing to align with perceptual or evaluation metrics.

5. Discussion and Conclusions

Most existing generative models for images belong to one of two classes. The first are likelihood-based models, trained with an elementwise KL reconstruction loss, which,

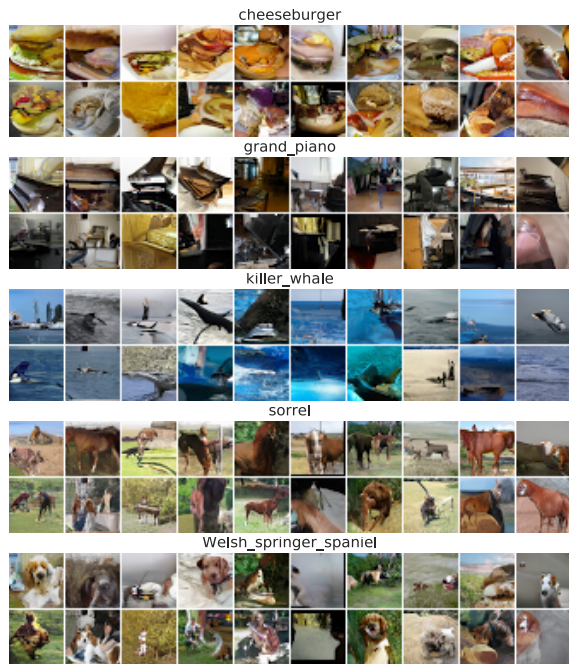


Figure 7. Class-conditional samples from PixelIQN. More samples of each class and more classes in the Appendix.

while perceptually meaningless, provides robust optimization properties and high sample diversity. The second are GANs, trained based on a discriminator loss, typically better aligned with a perceptual metric and enabling the generator to produce realistic, globally consistent samples. Their advantages come at the cost of a harder optimization problem, high parameter sensitivity, and most importantly, a tendency to collapse modes of the data distribution.

AIQNs are a new, fundamentally different, technique for generative modeling. By using a quantile regression loss instead of KL divergence, they combine some of the best properties of the two model classes. By their nature, they preserve modes of the learned distribution, while producing perceptually appealing high-quality samples. The inevitable approximation trade-offs a generative model makes when constrained by capacity or insufficient training can vary significantly depending on the loss used. We argue that the proposed quantile regression loss aligns more effectively with a given metric and therefore makes subjectively more advantageous trade-offs.

Devising methods for quantile regression over multidimensional outputs is an active area of research. New methods are continuing to be investigated (Carlier et al., 2016; Hallin & Miroslav, 2016), and a promising direction for future work is to find ways to use these to replace autoregressive models. One approach to reducing the computational burden of such models is to apply AIQN to the latent dimensions

of a VAE. Similar in spirit to Rosca et al. (2017), this would use the VAE to reduce the dimensionality of the problem and the AIQN to sample from the true latent distribution. In the Appendix we give preliminary results using such a technique, on CelebA 64×64 (Liu et al., 2015).

We have shown that IQN, computationally cheap and technically simple, can be readily applied to existing architectures, PixelCNN and VAE (Appendix), improving robustness and sampling quality of the underlying model. We demonstrated that PixelIQN produces more realistic, globally coherent samples, and improves Inception score and FID.

We further point out that many recent advances in generative models could be easily combined with our proposed method. Recent algorithmic improvements to GANs such as mini-batch discrimination and progressive growing (Salimans et al., 2016; Karras et al., 2017), while not strictly necessary in our work, could be applied to further improve performance. PixelCNN++ (Salimans et al., 2017) is an architectural improvement of PixelCNN, with several beneficial modifications supported by experimental evidence. Although we have built upon the original Gated PixelCNN in this work, we believe all of these modifications to be compatible with our work, except for the use of a mixture of logistics in place of PixelCNN’s softmax. As we have entirely replaced this model component, this change does not map onto our model. Of note, the motivation behind this change closely mirrors our own, in looking for a loss that respects the underlying metric between examples. The recent PixelSNAIL model (Chen et al., 2017) achieves state-of-the-art modeling performance by enhancing PixelCNN with ELU nonlinearities, modified block structure, and an attention mechanism. Again, all of these are fully compatible with our work and should improve results further.

Finally, the implicit quantile formulation lifts a number of architectural restrictions of previous generative models. Most importantly, the reparameterization as an inverse c.d.f. allows to learn distributions over continuous ranges without pre-specified boundaries or quantization. This enables modeling continuous-valued variables, for example for generation of sound (van den Oord et al., 2016a), opening multiple interesting avenues for further investigation.

Acknowledgements

We would like to acknowledge the important role many of our colleagues at DeepMind played for this work. We especially thank Aäron van den Oord and Sander Dieleman for invaluable advice on the PixelCNN model; Ivo Danihelka and Danilo J. Rezende for careful reading and insightful comments on an earlier version of the paper; Igor Babuschkin, Alexandre Galashov, Dominik Grewe, Jacob Menick, and Mihaela Rosca for technical help.

References

- Arjovsky, M., Chintala, S., and Bottou, L. Wasserstein generative adversarial networks. In *Proceedings of the 34th International Conference on Machine Learning (ICML)*, 2017.
- Bellemare, M. G., Danihelka, I., Dabney, W., Mohamed, S., Lakshminarayanan, B., Hoyer, S., and Munos, R. The Cramer distance as a solution to biased Wasserstein gradients. *arXiv preprint arXiv:1705.10743*, 2017.
- Bousquet, O., Gelly, S., Tolstikhin, I., Simon-Gabriel, C.-J., and Schoelkopf, B. From optimal transport to generative modeling: the vegan cookbook. *arXiv preprint arXiv:1705.07642*, 2017.
- Carlier, G., Chernozhukov, V., and Galichon, A. Vector quantile regression: an optimal transport approach. *Annals of Statistics*, 44(3):1165–1192, 06 2016.
- Chen, X., Duan, Y., Houthoofd, R., Schulman, J., Sutskever, I., and Abbeel, P. Infogan: Interpretable representation learning by information maximizing generative adversarial nets. In *Advances in Neural Information Processing Systems*, pp. 2172–2180, 2016.
- Chen, X., Mishra, N., Rohaninejad, M., and Abbeel, P. PixelSNAIL: An improved autoregressive generative model. *arXiv preprint arXiv:1712.09763*, 2017.
- Dabney, W., Ostrovski, G., Silver, D., and Munos, R. Implicit quantile networks for distributional reinforcement learning. *Proceedings of the 34th International Conference on Machine Learning (ICML)*, 2018a.
- Dabney, W., Rowland, M., Bellemare, M. G., and Munos, R. Distributional reinforcement learning with quantile regression. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2018b.
- Daskalakis, C., Ilyas, A., Syrgkanis, V., and Zeng, H. Training GANs with optimism. *arXiv preprint arXiv:1711.00141*, 2017.
- Dhaene, J., Vanduffel, S., Goovaerts, M. J., Kaas, R., Tang, Q., and Vyncke, D. Risk measures and comonotonicity: a review. *Stochastic Models*, 22(4):573–606, 2006.
- Dumoulin, V., Belghazi, I., Poole, B., Lamb, A., Arjovsky, M., Mastropietro, O., and Courville, A. Adversarially learned inference. *arXiv preprint arXiv:1606.00704*, 2016.
- Genevay, A., Peyré, G., and Cuturi, M. GAN and VAE from an optimal transport point of view. *arXiv preprint arXiv:1706.01807*, 2017.
- Germain, M., Gregor, K., Murray, I., and Larochelle, H. MADE: Masked autoencoder for distribution estimation. In *International Conference on Machine Learning*, pp. 881–889, 2015.
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., and Bengio, Y. Generative adversarial nets. In *Advances in Neural Information Processing Systems*, pp. 2672–2680, 2014.
- Gulrajani, I., Ahmed, F., Arjovsky, M., Dumoulin, V., and Courville, A. C. Improved training of Wasserstein GANs. In *Advances in Neural Information Processing Systems*, pp. 5769–5779, 2017.
- Hallin, M. and Miroslav, Š. Multiple-output quantile regression. *Handbook of Quantile Regression*, pp. 185–208, 2016.
- Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., and Hochreiter, S. GANs trained by a two time-scale update rule converge to a local nash equilibrium. In *Advances in Neural Information Processing Systems*, pp. 6629–6640, 2017.
- Higgins, I., Matthey, L., Pal, A., Burgess, C., Glorot, X., Botvinick, M., Mohamed, S., and Lerchner, A. Beta-VAE: Learning basic visual concepts with a constrained variational framework. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2017.
- Huber, P. J. Robust estimation of a location parameter. *Annals of Mathematical Statistics*, 35(1):73–101, 1964.
- Jones, M. C. Estimating densities, quantiles, quantile densities and density quantiles. *Annals of the Institute of Statistical Mathematics*, 44(4):721–727, 1992.
- Karras, T., Aila, T., Laine, S., and Lehtinen, J. Progressive growing of GANs for improved quality, stability, and variation. *arXiv preprint arXiv:1710.10196*, 2017.
- Kingma, D. P. and Welling, M. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- Koenker, R. Confidence intervals for regression quantiles. In *Asymptotic Statistics*, pp. 349–359. Springer, 1994.
- Koenker, R. and Hallock, K. Quantile regression: an introduction. *Journal of Economic Perspectives*, 15(4):43–56, 2001.
- Koenker, R. and Xiao, Z. Quantile autoregression. *Journal of the American Statistical Association*, 101(475):980–990, 2006.
- Krizhevsky, A. and Hinton, G. Learning multiple layers of features from tiny images. Technical report, 2009.

- Liu, Z., Luo, P., Wang, X., and Tang, X. Deep learning face attributes in the wild. In *Proceedings of International Conference on Computer Vision (ICCV)*, 2015.
- Makhzani, A., Shlens, J., Jaitly, N., Goodfellow, I., and Frey, B. Adversarial autoencoders. *arXiv preprint arXiv:1511.05644*, 2015.
- Papamakarios, G., Murray, I., and Pavlakou, T. Masked autoregressive flow for density estimation. In *Advances in Neural Information Processing Systems*, pp. 2335–2344, 2017.
- Parzen, E. Nonparametric statistical data modeling. *Journal of the American Statistical Association*, 74(365):105–121, 1979.
- Radford, A., Metz, L., and Chintala, S. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434*, 2015.
- Rezende, D. J., Mohamed, S., and Wierstra, D. Stochastic backpropagation and approximate inference in deep generative models. In *International Conference on Machine Learning*, pp. 1278–1286, 2014.
- Rosca, M., Lakshminarayanan, B., Warde-Farley, D., and Mohamed, S. Variational approaches for autoencoding generative adversarial networks. *arXiv preprint arXiv:1706.04987*, 2017.
- Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., et al. Imagenet large scale visual recognition challenge. *International Journal of Computer Vision*, 115(3): 211–252, 2015.
- Salimans, T., Goodfellow, I., Zaremba, W., Cheung, V., Radford, A., and Chen, X. Improved techniques for training GANs. In *Advances in Neural Information Processing Systems*, pp. 2234–2242, 2016.
- Salimans, T., Karpathy, A., Chen, X., and Kingma, D. P. PixelCNN++: Improving the pixelcnn with discretized logistic mixture likelihood and other modifications. *arXiv preprint arXiv:1701.05517*, 2017.
- Salimans, T., Zhang, H., Radford, A., and Metaxas, D. Improving gans using optimal transport. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2018.
- Siddiqui, M. M. Distribution of quantiles in samples from a bivariate population. *J. Res. Nat. Bur. Standards B*, 64: 145–150, 1960.
- Theis, L., van den Oord, A., and Bethge, M. A note on the evaluation of generative models. *arXiv preprint arXiv:1511.01844*, 2015.
- Tukey, J. W. Which part of the sample contains the information? *Proceedings of the National Academy of Sciences*, 53(1):127–134, 1965.
- van den Oord, A., Dieleman, S., Zen, H., Simonyan, K., Vinyals, O., Graves, A., Kalchbrenner, N., Senior, A., and Kavukcuoglu, K. Wavenet: A generative model for raw audio. *arXiv preprint arXiv:1609.03499*, 2016a.
- van den Oord, A., Kalchbrenner, N., Espeholt, L., Vinyals, O., Graves, A., and Kavukcuoglu, K. Conditional image generation with PixelCNN decoders. In *Advances in Neural Information Processing Systems*, pp. 4790–4798, 2016b.
- van den Oord, A., Kalchbrenner, N., and Kavukcuoglu, K. Pixel recurrent neural networks. In *Proceedings of the International Conference on Machine Learning*, 2016c.
- Zhao, S., Song, J., and Ermon, S. Towards deeper understanding of variational autoencoding models. *arXiv preprint arXiv:1702.08658*, 2017.