

CONDITIONAL IMAGE SYNTHESIS WITH AUXILIARY CLASSIFIER GANS

Augustus Odena*, Christopher Olah & Jonathon Shlens

Google Brain

{augustusodena, colah, shlens}@google.com

1 INTRODUCTION

Characterizing the structure of natural images has been a rich research endeavor. Natural images obey intrinsic invariances and exhibit multi-scale statistical structures that have historically been difficult to quantify (Simoncelli & Olshausen, 2001). Recent advances in machine learning offer an opportunity to substantially improve the quality of image models. Improved image models advance the state-of-the-art in image denoising (Ballé et al., 2015), compression (Toderici et al., 2016), in-painting (van den Oord et al., 2016a), and super-resolution (Ledig et al., 2016). Better models of natural images also improve performance in semi-supervised learning tasks (Kingma et al., 2014; Springenberg, 2015; Odena, 2016; Salimans et al., 2016) and reinforcement learning problems (Blundell et al., 2016).

Generative adversarial networks (GANs) offer a distinct and promising approach that focuses on a game-theoretic formulation for training an image synthesis model (Goodfellow et al., 2014). Recent work has shown that GANs can produce convincing image samples on datasets with low variability and low resolution (Denton et al., 2015; Radford et al., 2015). However, GANs struggle to generate globally coherent, high resolution samples - particularly from datasets with high variability. Moreover, a theoretical understanding of GANs is an on-going research topic (Uehara et al., 2016; Mohamed & Lakshminarayanan, 2016).




Figure 1: 128×128 resolution samples from 5 classes taken from an AC-GAN trained on the ImageNet dataset. Note that the classes shown have been selected to highlight the success of the model and are not representative. Samples from all ImageNet classes are in the Appendix.

In this work we demonstrate that that adding more structure to the GAN latent space along with a specialized cost function results in higher quality samples. We exhibit 128×128 pixel samples from all classes of the ImageNet dataset (Russakovsky et al., 2015) with increased global coherence (Figure 1). Importantly, we demonstrate quantitatively that our high resolution samples are not just naive resizings of low resolution samples. In particular, downsampling our 128×128 samples to 32×32 leads to a 50% decrease in visual discriminability. We also introduce a new metric for assessing the variability across image samples and employ this metric to demonstrate that our synthesized images exhibit diversity comparable to training data for a large fraction (84.7%) of ImageNet classes.

*Work completed as a participant in the 2016-2017 Google Brain Residency program.

2 BACKGROUND

A generative adversarial network (GAN) consists of two neural networks trained in opposition to one another. The generator G takes as input a random noise vector z and outputs an image $X_{fake} = G(z)$. The discriminator D receives as input either a training image or a synthesized image from the generator and outputs a probability distribution $P(S | X) = D(X)$ over possible image sources. The discriminator is trained to maximize the log-likelihood it assigns to the correct source:

$$L = E[\log P(S = real | X_{real})] + E[\log P(S = fake | X_{fake})]$$

The generator is trained to minimize that same quantity.

The basic GAN framework can be augmented using side information. One strategy is to supply both the generator and discriminator with class labels in order to produce class conditional samples (Mirza & Osindero, 2014). Class conditional synthesis can significantly improve the quality of generated samples (van den Oord et al., 2016b). Richer side information such as image captions and bounding box localizations may improve sample quality further (Reed et al., 2016a;b).

Instead of feeding side information to the discriminator, one can task the discriminator with reconstructing side information. This is done by modifying the discriminator to contain an auxiliary decoder network¹ that outputs the class label for the training data (Odena, 2016; Salimans et al., 2016) or a subset of the latent variables from which the samples are generated (Chen et al., 2016). Forcing a model to perform additional tasks is known to improve performance on the original task (e.g. Sutskever et al. (2014); Szegedy et al. (2014); Ramsundar et al. (2016)). In addition, an auxiliary decoder could leverage pre-trained discriminators (e.g. image classifiers) for further improving the synthesized images (Nguyen et al., 2016). Motivated by these considerations, we introduce a model that combines both strategies for leveraging side information. That is, the model proposed below is class conditional, but with an auxiliary decoder that is tasked with reconstructing class labels.

3 AC-GANS

We propose a variant of the GAN architecture which we call an auxiliary classifier GAN (or AC-GAN). In the AC-GAN, every generated sample has a corresponding class label, $c \sim p_c$ in addition to the noise z . G uses both to generate images $X_{fake} = G(c, z)$. The discriminator gives both a probability distribution over sources and a probability distribution over the class labels, $P(S | X)$, $P(C | X) = D(X)$. The objective function has two parts: the log-likelihood of the correct source, L_S , and the log-likelihood of the correct class, L_C .

$$L_S = E[\log P(S = real | X_{real})] + E[\log P(S = fake | X_{fake})]$$

$$L_C = E[\log P(C = c | X_{real})] + E[\log P(C = c | X_{fake})]$$

D is trained to maximize $L_S + L_C$ while G is trained to maximize $L_C - L_S$. AC-GANs learn a representation for z that is independent of class label (e.g. Kingma et al. (2014)).

Early experiments demonstrated that increasing the number of classes trained on while holding the model fixed decreased the quality of the model outputs (Appendix E). The structure of the AC-GAN model permits separating large datasets into subsets by class and training a generator and discriminator for each subset. We exploit this property in our experiments to train across the entire ImageNet data set.

4 RESULTS

We train several AC-GAN models on the ImageNet data set (Russakovsky et al., 2015). Broadly speaking, the architecture of the generator G is a series of ‘deconvolution’ layers that transform the noise z and class c into an image (Odena et al., 2016). We train two variants of the model architecture

¹ Alternatively, one can force the discriminator to work with the joint distribution (X, z) and train a separate inference network that computes $q(z|X)$ (Dumoulin et al., 2016; Donahue et al., 2016).

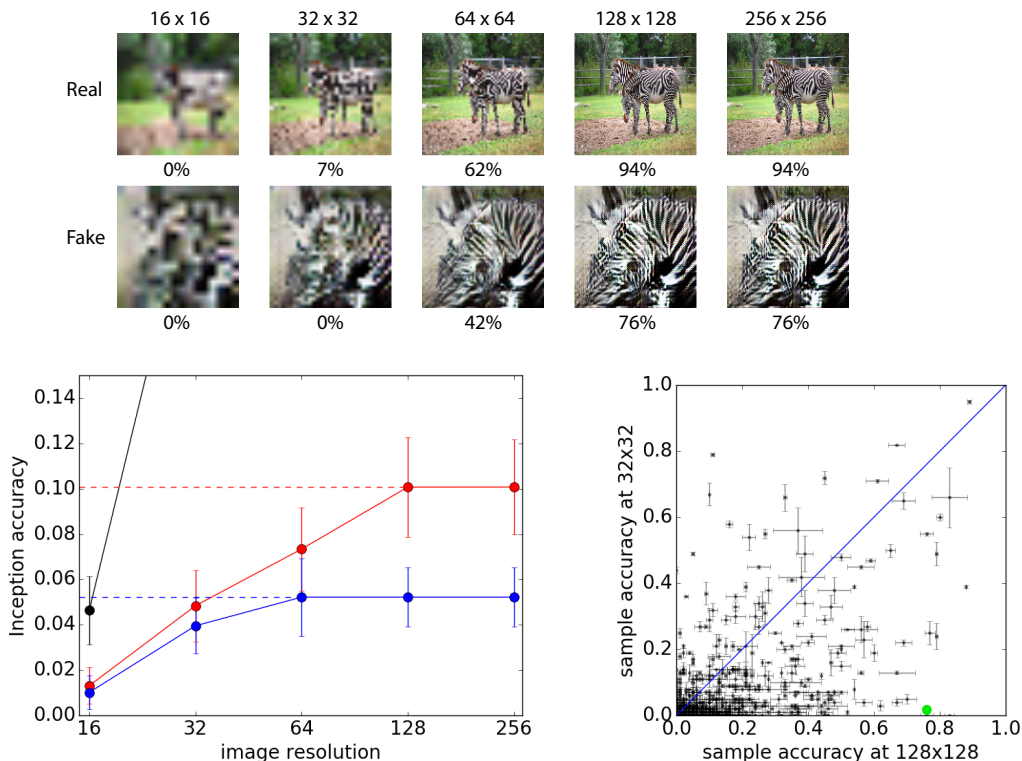


Figure 2: Generating high resolution images improves discriminability. Top: Training data and synthesized images from the *zebra* class resized to a lower spatial resolution (indicated above) and subsequently artificially resized to the original resolution. Inception accuracy is shown below the corresponding images. Bottom Left: Summary of accuracies across varying spatial resolutions for training data and image samples from 64×64 and 128×128 models. Error bar measures standard deviation across 10 subsets of images. Dashed lines highlight the accuracy at the output spatial resolution of the model. The training data (clipped) achieves accuracies of 24%, 54%, 81% and 81% at resolutions of 32, 64, 128, and 256 respectively. Bottom Right: Comparison of accuracy scores at 128×128 and 32×32 spatial resolutions (x and y axis, respectively). Each point represents an ImageNet class. 84.4% of the classes are below the line of equality. The green dot corresponds to the *zebra* class. We also artificially resized 128×128 and 64×64 images to 256×256 as a sanity check to demonstrate that simply increasing the number of pixels will not increase discriminability.

for generating images at 128×128 and 64×64 spatial resolutions. The discriminator D is a deep convolutional neural network with a Leaky ReLU nonlinearity (Maas et al., 2013). See Appendix B for more details. As mentioned earlier, we find that reducing the variability introduced by all 1000 classes of ImageNet significantly improves the quality of training. We train 100 AC-GAN models – each on images from just 10 classes – for 50000 mini-batches of size 100.

Evaluating the quality of image synthesis models is challenging due to the variety of probabilistic criteria (Theis et al., 2015) and the lack of a perceptually meaningful image similarity metric. Nonetheless, in subsequent sections we attempt to measure the quality of the AC-GAN by building several *ad-hoc* measures for image sample discriminability and diversity. Our hope is that this work might provide quantitative measures that may be used to aid training and subsequent development of image synthesis models.

Figure 2 contains our attempt to measure whether the model is making ‘full use’ of its output resolution. See Appendix A for more details and further experiments.

ACKNOWLEDGMENTS

We thank the developers of TensorFlow (Abadi et al., 2016). We thank Luke Metz and Vincent Dumoulin for extensive and helpful comments on drafts. We also thank Ben Poole, Sam Schoenholz, Barret Zoph, Martín Abadi, Manjunath Kudlur and Jascha Sohl-Dickstein for helpful discussions.

REFERENCES

- Abadi et al. Tensorflow: Large-scale machine learning on heterogeneous distributed systems. *CoRR*, abs/1603.04467, 2016. URL <http://arxiv.org/abs/1603.04467>.
- Johannes Ballé, Valero Laparra, and Eero P. Simoncelli. Density modeling of images using a generalized normalization transformation. *CoRR*, abs/1511.06281, 2015. URL <http://arxiv.org/abs/1511.06281>.
- Yoshua Bengio, Grégoire Mesnil, Yann Dauphin, and Salah Rifai. Better mixing via deep representations. *CoRR*, abs/1207.4404, 2012. URL <http://arxiv.org/abs/1207.4404>.
- C. Blundell, B. Uria, A. Pritzel, Y. Li, A. Ruderman, J. Z Leibo, J. Rae, D. Wierstra, and D. Hassabis. Model-Free Episodic Control. *ArXiv e-prints*, June 2016.
- X. Chen, Y. Duan, R. Houthoofd, J. Schulman, I. Sutskever, and P. Abbeel. InfoGAN: Interpretable Representation Learning by Information Maximizing Generative Adversarial Nets. *ArXiv e-prints*, June 2016.
- Emily L. Denton, Soumith Chintala, Arthur Szlam, and Robert Fergus. Deep generative image models using a laplacian pyramid of adversarial networks. *CoRR*, abs/1506.05751, 2015. URL <http://arxiv.org/abs/1506.05751>.
- Laurent Dinh, Jascha Sohl-Dickstein, and Samy Bengio. Density estimation using real NVP. *CoRR*, abs/1605.08803, 2016. URL <http://arxiv.org/abs/1605.08803>.
- J. Donahue, P. Krähenbühl, and T. Darrell. Adversarial Feature Learning. *ArXiv e-prints*, May 2016.
- V. Dumoulin, I. Belghazi, B. Poole, A. Lamb, M. Arjovsky, O. Mastropietro, and A. Courville. Adversarially Learned Inference. *ArXiv e-prints*, June 2016.
- I. J. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative Adversarial Networks. *ArXiv e-prints*, June 2014.
- D. P Kingma and M. Welling. Auto-Encoding Variational Bayes. *ArXiv e-prints*, December 2013.
- Diederik P. Kingma, Danilo Jimenez Rezende, Shakir Mohamed, and Max Welling. Semi-supervised learning with deep generative models. *CoRR*, abs/1406.5298, 2014. URL <http://arxiv.org/abs/1406.5298>.
- Diederik P. Kingma, Tim Salimans, and Max Welling. Improving variational inference with inverse autoregressive flow. *CoRR*, abs/1606.04934, 2016. URL <http://arxiv.org/abs/1606.04934>.
- C. Ledig, L. Theis, F. Huszar, J. Caballero, A. Aitken, A. Tejani, J. Totz, Z. Wang, and W. Shi. Photo-Realistic Single Image Super-Resolution Using a Generative Adversarial Network. *ArXiv e-prints*, September 2016.
- Kede Ma, Qingbo Wu, Zhou Wang, Zhengfang Duanmu, Hongwei Yong, Hongliang Li, and Lei Zhang. Group mad competition - a new methodology to compare objective image quality models. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016.
- Andrew Maas, Awni Hannun, and Andrew Ng. Rectifier nonlinearities improve neural network acoustic models. In *Proceedings of The 33rd International Conference on Machine Learning*, 2013.
- Mehdi Mirza and Simon Osindero. Conditional generative adversarial nets. *CoRR*, abs/1411.1784, 2014. URL <http://arxiv.org/abs/1411.1784>.

- Shakir Mohamed and Balaji Lakshminarayanan. Learning in implicit generative models. *arXiv preprint arXiv:1610.03483*, 2016.
- Anh Mai Nguyen, Alexey Dosovitskiy, Jason Yosinski, Thomas Brox, and Jeff Clune. Synthesizing the preferred inputs for neurons in neural networks via deep generator networks. *CoRR*, abs/1605.09304, 2016. URL <http://arxiv.org/abs/1605.09304>.
- A. Odena. Semi-Supervised Learning with Generative Adversarial Networks. *ArXiv e-prints*, June 2016.
- Augustus Odena, Vincent Dumoulin, and Chris Olah. Deconvolution and checkerboard artifacts. <http://distill.pub/2016/deconv-checkerboard/>, 2016.
- Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. *CoRR*, abs/1511.06434, 2015. URL <http://arxiv.org/abs/1511.06434>.
- Bharath Ramsundar, Steven Kearnes, Patrick Riley, Dale Webster, David Konerding, and Vijay Pande. Massively multitask networks for drug discovery. In *Proceedings of The 33rd International Conference on Machine Learning*, 2016.
- Scott Reed, Zeynep Akata, Santosh Mohan, Samuel Tenka, Bernt Schiele, and Honglak Lee. Learning what and where to draw. *arXiv preprint arXiv:1610.02454*, 2016a.
- Scott Reed, Zeynep Akata, Xinchun Yan, Lajanugen Logeswaran, Bernt Schiele, and Honglak Lee. Generative adversarial text-to-image synthesis. In *Proceedings of The 33rd International Conference on Machine Learning*, 2016b.
- D. Rezende and S. Mohamed. Variational Inference with Normalizing Flows. *ArXiv e-prints*, May 2015.
- D. Rezende, S. Mohamed, and D. Wierstra. Stochastic Backpropagation and Approximate Inference in Deep Generative Models. *ArXiv e-prints*, January 2014.
- Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. *CoRR*, abs/1505.04597, 2015. URL <http://arxiv.org/abs/1505.04597>.
- Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)*, 115(3):211–252, 2015. doi: 10.1007/s11263-015-0816-y.
- T. Salimans, I. Goodfellow, W. Zaremba, V. Cheung, A. Radford, and X. Chen. Improved Techniques for Training GANs. *ArXiv e-prints*, June 2016.
- Eero Simoncelli and Bruno Olshausen. Natural image statistics and neural representation. *Annual Review of Neuroscience*, 24:1193–1216, 2001.
- J. T. Springenberg. Unsupervised and Semi-supervised Learning with Categorical Generative Adversarial Networks. *ArXiv e-prints*, November 2015.
- Ilya Sutskever, Oriol Vinyals, and Le Quoc V. Sequence to sequence learning with neural networks. In *Neural Information Processing Systems*, 2014.
- Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott E. Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. *CoRR*, abs/1409.4842, 2014. URL <http://arxiv.org/abs/1409.4842>.
- Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jonathon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. *CoRR*, abs/1512.00567, 2015. URL <http://arxiv.org/abs/1512.00567>.
- L. Theis, A. van den Oord, and M. Bethge. A note on the evaluation of generative models. *ArXiv e-prints*, November 2015.

- George Toderici, Damien Vincent, Nick Johnston, Sung Jin Hwang, David Minnen, Joel Shor, and Michele Covell. Full resolution image compression with recurrent neural networks. *CoRR*, abs/1608.05148, 2016. URL <http://arxiv.org/abs/1608.05148>.
- M. Uehara, I. Sato, M. Suzuki, K. Nakayama, and Y. Matsuo. Generative Adversarial Nets from a Density Ratio Estimation Perspective. *ArXiv e-prints*, October 2016.
- Aäron van den Oord, Nal Kalchbrenner, and Koray Kavukcuoglu. Pixel recurrent neural networks. *CoRR*, abs/1601.06759, 2016a. URL <http://arxiv.org/abs/1601.06759>.
- Aäron van den Oord, Nal Kalchbrenner, Oriol Vinyals, Lasse Espeholt, Alex Graves, and Koray Kavukcuoglu. Conditional image generation with pixelenn decoders. *CoRR*, abs/1606.05328, 2016b. URL <http://arxiv.org/abs/1606.05328>.
- Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004a.
- Zhou Wang, Eero P Simoncelli, and Alan C Bovik. Multiscale structural similarity for image quality assessment. In *Signals, Systems and Computers, 2004. Conference Record of the Thirty-Seventh Asilomar Conference on*, volume 2, pp. 1398–1402. Ieee, 2004b.

A FURTHER EXPERIMENTS

Here we present further experimental results that did not fit into the main paper.

A.1 GENERATING HIGH RESOLUTION IMAGES IMPROVES DISCRIMINABILITY

Building a class-conditional image synthesis model necessitates measuring the extent to which synthesized images appear to belong to the intended class. In particular, we would like to know that a high resolution sample is not just a naive resizing of a low resolution sample. Consider a simple experiment: pretend there exists a model that synthesizes 32×32 images. One can trivially increase the resolution of synthesized images by performing bilinear interpolation. This would yield higher resolution images, but these images would just be blurry versions of the low resolution images that are not discriminable. Hence, the goal of an image synthesis model is not simply to produce high resolution images, but to produce high resolution images that are more discriminable than low resolution images.

To measure discriminability, we feed synthesized images to a pre-trained Inception network (Szegedy et al., 2015) and report the fraction of the samples for which the Inception network assigned the correct label². We calculate this accuracy measure on a series of real and synthesized images which have had their spatial resolution artificially decreased by bilinear interpolation (Figure 2, top panels). Note that as the spatial resolution is decreased, the accuracy decreases - indicating that resulting images contain less class information (Figure 2, scores below top panels). We summarized this finding across all 1000 ImageNet classes for the ImageNet training data (black), a 128×128 resolution AC-GAN (red) and a 64×64 resolution AC-GAN (blue) in Figure 2 (bottom, left). The black curve (clipped) provides an upper-bound on the discriminability of real images.

The goal of this analysis is to show that synthesizing higher resolution images leads to increased discriminability. The 128×128 model achieves an accuracy of $10.1\% \pm 2.0\%$ versus $7.0\% \pm 2.0\%$ with samples resized to 64×64 and $5.0\% \pm 2.0\%$ with samples resized to 32×32 . In other words, downsizing the outputs of the AC-GAN to 32×32 and 64×64 decreases visual discriminability by 50% and 38% respectively. Furthermore, 84.4% of the ImageNet classes have higher accuracy at 128×128 than at 32×32 (Figure 2, bottom left).

We performed the same analysis on an AC-GAN trained to 64×64 spatial resolution. This model achieved less discriminability than a 128×128 AC-GAN model. Accuracies from the 64×64 model plateau at a 64×64 spatial resolution consistent with previous results. Finally, the 64×64 resolution model achieves less discriminability at 64 spatial resolution than the 128×128 model.

A.2 MEASURING THE DIVERSITY OF GENERATED IMAGES

An image synthesis model is not very interesting if it only outputs one image. Indeed, a well-known failure mode of GANs is that the generator will collapse and output a single prototype that maximally fools the discriminator (Goodfellow et al., 2014; Salimans et al., 2016). A class-conditional model of images is not very interesting if it only outputs one image per class. The Inception accuracy can not measure whether a model has collapsed. A model that simply memorized one example from each ImageNet class would do very well by this metric. Thus, we seek a complementary metric to explicitly evaluate the intra-class diversity of samples generated by the AC-GAN.

Several methods exist for quantitatively evaluating image similarity by attempting to predict human perceptual similarity judgements. The most successful of these is multi-scale structural similarity (MS-SSIM) (Wang et al., 2004b; Ma et al., 2016). MS-SSIM is a multi-scale variant of a well-characterized perceptual similarity metric that attempts to discount aspects of an image that are not important for human perception (Wang et al., 2004a). MS-SSIM values range between 0.0 and 1.0; higher MS-SSIM values correspond to perceptually more similar images. As a proxy for image diversity, we measure the MS-SSIM scores between randomly chosen pairs of images within a

² One could also use the Inception score (Salimans et al., 2016), but our method has several advantages: accuracy figures are easier to interpret than exponentiated KL-divergences; accuracy may be assessed for individual classes; accuracy measures whether a class-conditional model generated samples from the intended class. To compute the Inception accuracy, we modified a version of Inception-v3 supplied in <https://github.com/openai/improved-gan/>.

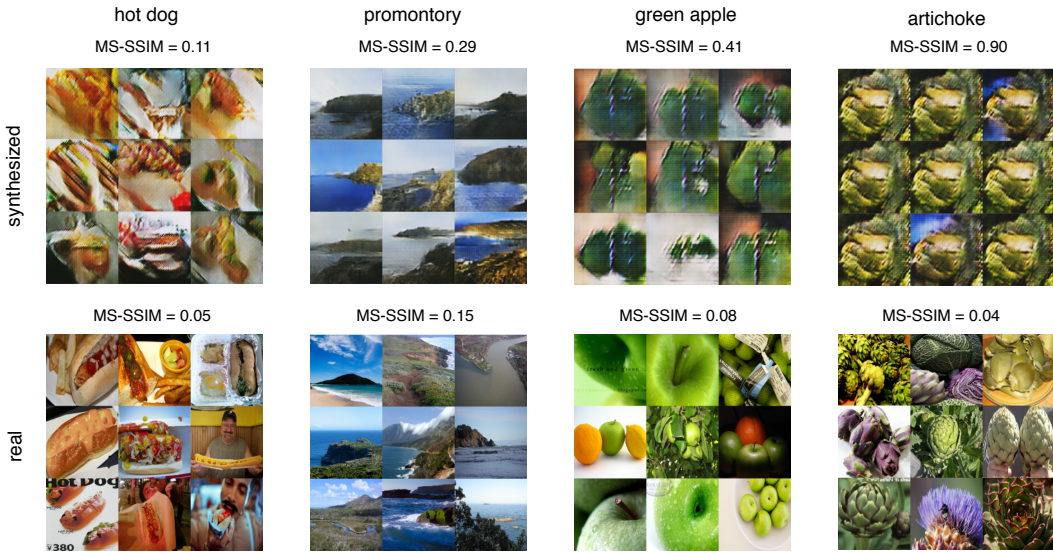


Figure 3: Examples of different MS-SSIM scores. The top and bottom rows contain AC-GAN samples and training data, respectively.

given class. Samples from classes that have higher diversity result in lower mean MS-SSIM scores (Figure 3, left columns); samples from classes with lower diversity have higher mean MS-SSIM scores (Figure 3, right columns). Training images from the ImageNet training data contain a variety of mean MS-SSIM scores across the classes indicating the variability of image diversity in ImageNet classes (Figure 4, left panel, x-axis). Note that the highest mean MS-SSIM score (indicating the least variability) is 0.25 for the training data.

We calculate the mean MS-SSIM score for all 1000 ImageNet classes generated by the AC-GAN model. We track this value during training to identify whether the generator has collapsed (Figure 4, right panel, red curve). We also employ this metric to compare the diversity of the training images to the samples from the GAN model after training has completed. Figure 4 (left) plots the mean MS-SSIM values for image samples and training data broken up by class. The blue line is the line of equality. Out of the 1000 classes, we find that 847 have mean sample MS-SSIM scores below that of the maximum MS-SSIM for the training data. In other words, 84.7% of classes have sample variability that exceeds that of the least variable class from the ImageNet training data.

Note that SSIM is not intended as a generic distance measure that will measure any distance in semantic space. Instead, we are focused on the real potential for a GAN to output a set of images that are only differ across brightness, illumination and contrast. The variability across these dimensions are not perceptually relevant for measuring the diversity of the GAN’s output and we want our distance metric to treat these variants of an image as negligible, hence our selection of the SSIM metric.

A.3 GENERATED IMAGES ARE BOTH DIVERSE AND DISCRIMINABLE

We have presented quantitative metrics demonstrating that AC-GAN samples may be diverse and discriminable but we have yet to examine how these metrics interact. Figure 5 shows the joint distribution of Inception accuracies and MS-SSIM scores across all classes. Inception accuracy and MS-SSIM are anti-correlated ($r^2 = -0.16$). In fact, 74% of the classes with low diversity ($MS-SSIM \geq 0.25$) contain Inception accuracies $\leq 1\%$. These results suggest that GANs that drop modes are most likely to produce low quality images. Conversely, 78% of classes with high diversity ($MS-SSIM < 0.25$) have Inception accuracies that exceed 1%. In comparison, the Inception-v3 model achieves 78.8% accuracy on average across all 1000 classes (Szegedy et al., 2015). A fraction of the classes AC-GAN samples reach this level of accuracy. This indicates opportunity for future image synthesis models.




Figure 4: (Left) Comparison of the mean MS-SSIM scores between pairs of images within a given class for ImageNet training data and samples from the GAN (blue line is equality). The horizontal red line marks the maximum MS-SSIM value across all ImageNet classes. Each point is an individual class. The mean standard deviation of scores across the training data and the samples was 0.06 and 0.08 respectively. Scores below the red line (84.7% of classes) arise from classes where GAN training largely succeeded. (Right) Intra-class MS-SSIM for selected ImageNet classes throughout a training run. Classes that successfully train tend to have decreasing mean MS-SSIM scores, to a point.




Figure 5: Inception accuracy vs MS-SSIM for all 1000 ImageNet classes ($r^2 = -0.16$). Samples from AC-GAN models do not achieve variability at the expense of discriminability.

A.4 COMPARISON TO PREVIOUS RESULTS

Previous quantitative results for image synthesis models trained on ImageNet are reported in terms of log-likelihood (van den Oord et al., 2016a;b). Log-likelihood is a coarse and potentially inaccurate measure of sample quality (Theis et al., 2015). Additionally, log-likelihood is intractable to compute for GANs. Instead we compare with previous state-of-the-art results on CIFAR-10 using a lower spatial resolution (32×32). Following the procedure in Salimans et al. (2016), we compute the Inception score³ for 50000 samples from an AC-GAN with resolution (32×32), split into 10 groups at random. We also compute the Inception score for 25000 extra samples, split into 5 groups at random. We select the best model based on the first score and report the second score. Performing a grid search across 27 hyperparameter configurations, we are able to achieve a score of 8.25 ± 0.07 compared to state of the art 8.09 ± 0.07 (Salimans et al., 2016). Moreover, we accomplish this without employing any of the new techniques introduced in that work (i.e. virtual batch normalization, minibatch discrimination, and label smoothing). This provides additional evidence that AC-GANs are effective even without the benefit of class splitting (Appendix E).

³ The Inception score is given by $\exp(E_x[D_{KL}(p(y|x} || p(y))])$ where x is a particular image, $p(y|x)$ is the conditional output distribution over the classes in a pre-trained Inception network (Szegedy et al., 2014) given x , and $p(y)$ is the marginal distribution over the classes.

A.5 SEARCHING FOR SIGNATURES OF OVERFITTING

One possibility that must be investigated is that the AC-GAN has overfit on the training data. As a first check that the network does not memorize the training data, we identify the nearest neighbors of image samples in the training data measured by L1 distance in pixel space (Figure 6). The nearest neighbors from the training data do not resemble the corresponding samples. This provides evidence that the AC-GAN is not merely memorizing the training data.




Figure 6: Nearest neighbor analysis. (Left) Samples from a single ImageNet class. (Right) Corresponding nearest neighbor (L1 distance) in training data for each sample.

A more sophisticated method for understanding the degree of overfitting in a model is to explore that model’s latent space by interpolation. In an overfit model one might observe discrete transitions in the interpolated images and regions in latent space that do not correspond to meaningful images (Bengio et al., 2012; Radford et al., 2015; Dinh et al., 2016). Figure 7 (left) highlights interpolations in the latent space between several image samples. Notably, the generator learned that certain combinations of dimensions correspond to semantically meaningful features (e.g. size of the arch, length of a bird’s beak) and there are no discrete transitions or ‘holes’ in the latent space. A second method for exploring the latent space of the AC-GAN is to exploit the structure of the model. The AC-GAN factorizes its representation into class information and a class-independent latent representation z . Sampling the AC-GAN with z fixed but altering the class label corresponds to generating samples with the same ‘style’ across multiple classes (Kingma et al., 2014). Figure 7 (right) shows samples from 8 bird classes. Elements of the same row have the same z . Although the class changes for each column, elements of the global structure (e.g. position, layout, background) are preserved, indicating that AC-GAN can represent certain types of ‘compositionality’.



Figure 7: (Left) Latent space interpolations for selected ImageNet classes. Left-most and right-columns show three pairs of image samples - each pair from a distinct class. Intermediate columns highlight linear interpolations in the latent space between these three pairs of images. (Right) Class-independent information contains global structure about the synthesized image. Each column is a distinct bird class while each row corresponds to a fixed latent code z .

B HYPERPARAMETERS

	Operation	Kernel	Strides	Feature maps	BN?	Dropout	Nonlinearity
$G_x(z)$ – $110 \times 1 \times 1$ input							
	Linear	N/A	N/A	768	×	0.0	ReLU
	Transposed Convolution	5×5	2×2	384	✓	0.0	ReLU
	Transposed Convolution	5×5	2×2	256	✓	0.0	ReLU
	Transposed Convolution	5×5	2×2	192	✓	0.0	ReLU
	Transposed Convolution	5×5	2×2	3	×	0.0	Tanh
$D(x)$ – $128 \times 128 \times 3$ input							
	Convolution	3×3	2×2	16	×	0.5	Leaky ReLU
	Convolution	3×3	1×1	32	✓	0.5	Leaky ReLU
	Convolution	3×3	2×2	64	✓	0.5	Leaky ReLU
	Convolution	3×3	1×1	128	✓	0.5	Leaky ReLU
	Convolution	3×3	2×2	256	✓	0.5	Leaky ReLU
	Convolution	3×3	1×1	512	✓	0.5	Leaky ReLU
	Linear	N/A	N/A	11	×	0.0	Soft-Sigmoid
	Optimizer	Adam ($\alpha = 0.0002, \beta_1 = 0.5, \beta_2 = 0.999$)					
	Batch size	100					
	Iterations	50000					
	Leaky ReLU slope	0.2					
	Weight, bias initialization	Isotropic gaussian ($\mu = 0, \sigma = 0.02$), Constant(0)					

Table 1: ImageNet hyperparameters. A Soft-Sigmoid refers to an operation over $K + 1$ output units where we apply a Softmax activation to K of the units and a Sigmoid activation to the remaining unit. We also use activation noise in the discriminator as suggested in Salimans et al. (2016).

Operation	Kernel	Strides	Feature maps	BN?	Dropout	Nonlinearity
$G_x(z) - 110 \times 1 \times 1$ input						
Linear	N/A	N/A	384	×	0.0	ReLU
Transposed Convolution	5×5	2×2	192	✓	0.0	ReLU
Transposed Convolution	5×5	2×2	96	✓	0.0	ReLU
Transposed Convolution	5×5	2×2	3	×	0.0	Tanh
$D(x) - 32 \times 32 \times 3$ input						
Convolution	3×3	2×2	16	×	0.5	Leaky ReLU
Convolution	3×3	1×1	32	✓	0.5	Leaky ReLU
Convolution	3×3	2×2	64	✓	0.5	Leaky ReLU
Convolution	3×3	1×1	128	✓	0.5	Leaky ReLU
Convolution	3×3	2×2	256	✓	0.5	Leaky ReLU
Convolution	3×3	1×1	512	✓	0.5	Leaky ReLU
Linear	N/A	N/A	11	×	0.0	Soft-Sigmoid
Generator Optimizer	Adam ($\alpha = [0.0001, 0.0002, 0.0003]$, $\beta_1 = 0.5$, $\beta_2 = 0.999$)					
Discriminator Optimizer	Adam ($\alpha = [0.0001, 0.0002, 0.0003]$, $\beta_1 = 0.5$, $\beta_2 = 0.999$)					
Batch size	100					
Iterations	50000					
Leaky ReLU slope	0.2					
Activation noise standard deviation	[0, 0.1, 0.2]					
Weight, bias initialization	Isotropic gaussian ($\mu = 0$, $\sigma = 0.02$), Constant(0)					

Table 2: CIFAR-10 hyperparameters. When a list is given for a hyperparameter it means that we performed a grid search using the values in the list. For each set of hyperparameters, a single AC-GAN was trained on the whole CIFAR-10 dataset. For each AC-GAN that was trained, we split up the samples into groups so that we could give some sense of the variance of the Inception Score. To the best of our knowledge, this is identical to the analysis performed in Salimans et al. (2016).

C QUALITATIVE COMPARISON OF SAMPLES



Figure 8: Samples generated from the ImageNet dataset. (Left) Samples generated from the model in Salimans et al. (2016). (Right) Samples generated from an AC-GAN. AC-GAN samples possess global coherence absent from the samples of the earlier model. It's relatively clear that - clockwise from top-left - the images are of a flower, a bird in a tree, a landscape, and a spider. The spider even seems to have around 8 legs, depending on how the counting is done.

D USING MS-SSIM TO FIND NEAREST NEIGHBORS



Figure 9: Nearest neighbor analysis using MS-SSIM instead of L1. Neighbors appear more structurally similar than when L1 is used, but there is still no meaningful evidence of memorization. (Left) Samples from a single ImageNet class. (Right) Corresponding nearest neighbor in training data for each sample.

E MEASURING THE EFFECT OF CLASS SPLITS ON IMAGE SAMPLE QUALITY.

Class conditional image synthesis affords the opportunity to divide up a dataset based on image label. In our final model we divide 1000 ImageNet classes across 100 AC-GAN models. In this section we describe early experiments that highlight the benefit of cutting down the diversity of classes for training an AC-GAN. We employed an ordering of the labels and divided it into contiguous groups of 10. This ordering can be seen in the following section, where we display samples from all 1000 classes. Two aspects of the split merit discussion: the number of classes per split and the intra-split diversity.

We find that training a fixed model on more classes harms the model’s ability to produce compelling samples (Figure 10). Performance on larger splits can be improved by giving the model more parameters. However, using a small split is not sufficient to achieve good performance. We were unable to train a GAN (Goodfellow et al., 2014) to converge reliably even for a split size of 1.




Figure 10: Mean pairwise MS-SSIM values for 10 ImageNet classes plotted against the number of ImageNet classes used during training. We fix everything except the number of classes trained on, using values from 10 to 100. We only report the MS-SSIM values for the first 10 classes to keep the MS-SSIM comparable. MS-SSIM quickly goes above 0.25 (the red line) as the class count increases. These scores were computed using 9 random restarts per class count, using the same number of training steps for each model. Since we have observed that generators do not recover from the collapse phase, the use of a fixed number of training steps seems justified in this case.

This raises the question of whether it is easier to train a model on a diverse set of classes than on a similar set of classes. We were unable to find conclusive evidence that the selection of classes in a split significantly affects sample quality.

We don’t have a hypothesis about what causes this sensitivity to class count that is well-supported experimentally. We can only note that, since the failure case that occurs when the class count is increased is ‘generator collapse’, it seems plausible that general methods for addressing ‘generator collapse’ could also address this sensitivity.

F SAMPLES FROM ALL 1000 IMAGENET CLASSES

The following is a link to 10 samples from each of the 1000 ImageNet classes:
<https://goo.gl/photos/8bgHBkCwDEVTXAPaA>