
Conditional Image Synthesis with Auxiliary Classifier GANs: Appendix

Augustus Odena¹ Christopher Olah¹ Jonathon Shlens¹

1. Hyperparameters

We summarize hyperparameters used for the ImageNet model in Table 1 and for the CIFAR-10 model in Table 2.

Conditional Image Synthesis with Auxiliary Classifier GANs: Appendix

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 (?).

	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 (?).