Generative Adversarial Networks

STEPHEN MCALEER

Outline

1. Discriminative vs. generative models
2. Deep generative models
3. Generative adversarial networks
4. Applications and further work

Discriminative vs. Generative Models

Training set
x, y

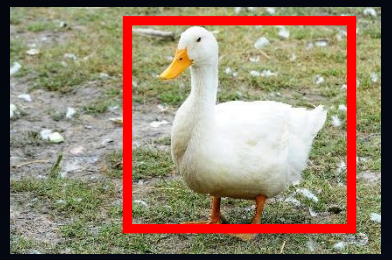

duck

car

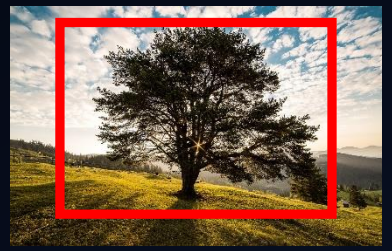

tree

Discriminative Model
$p(y|x)$

pickup truck

x_2

Generative Model
$p(x, y)$

$z \sim N(0,1)$

horse

Deep Learning

- Learn hierarchical model of data
- Higher-level features derived from lower-level features
- Has achieved much success in discriminative tasks

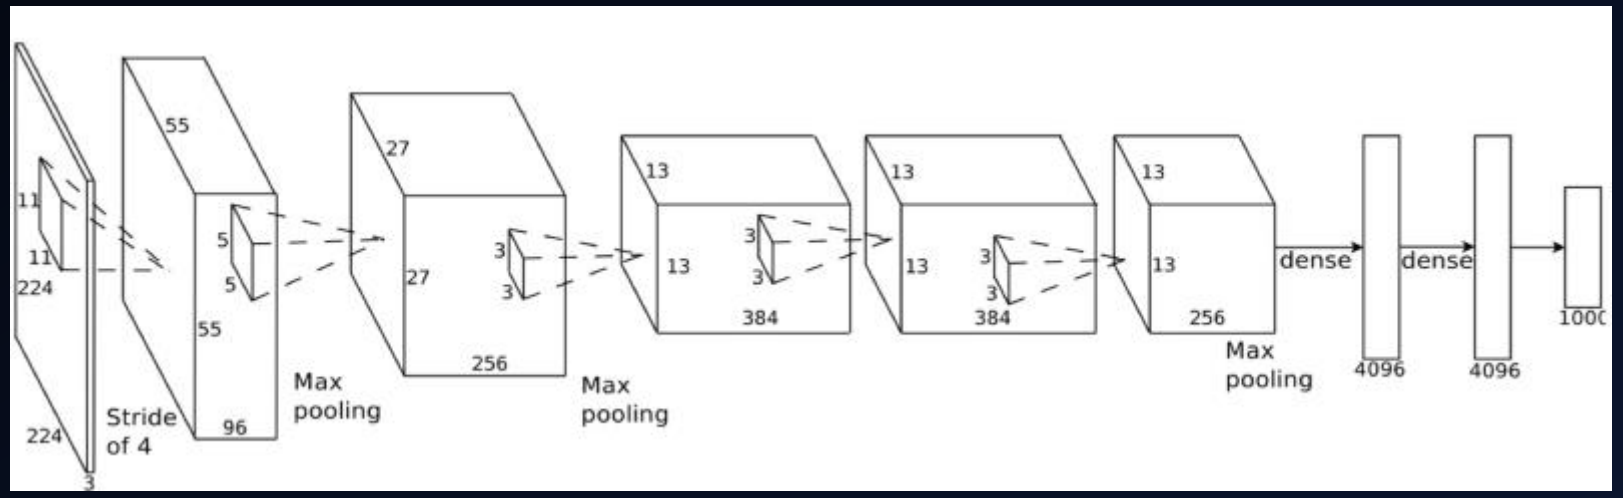

Generating an Image

- Opposite of convolutional neural nets
- How to train it?

x_2

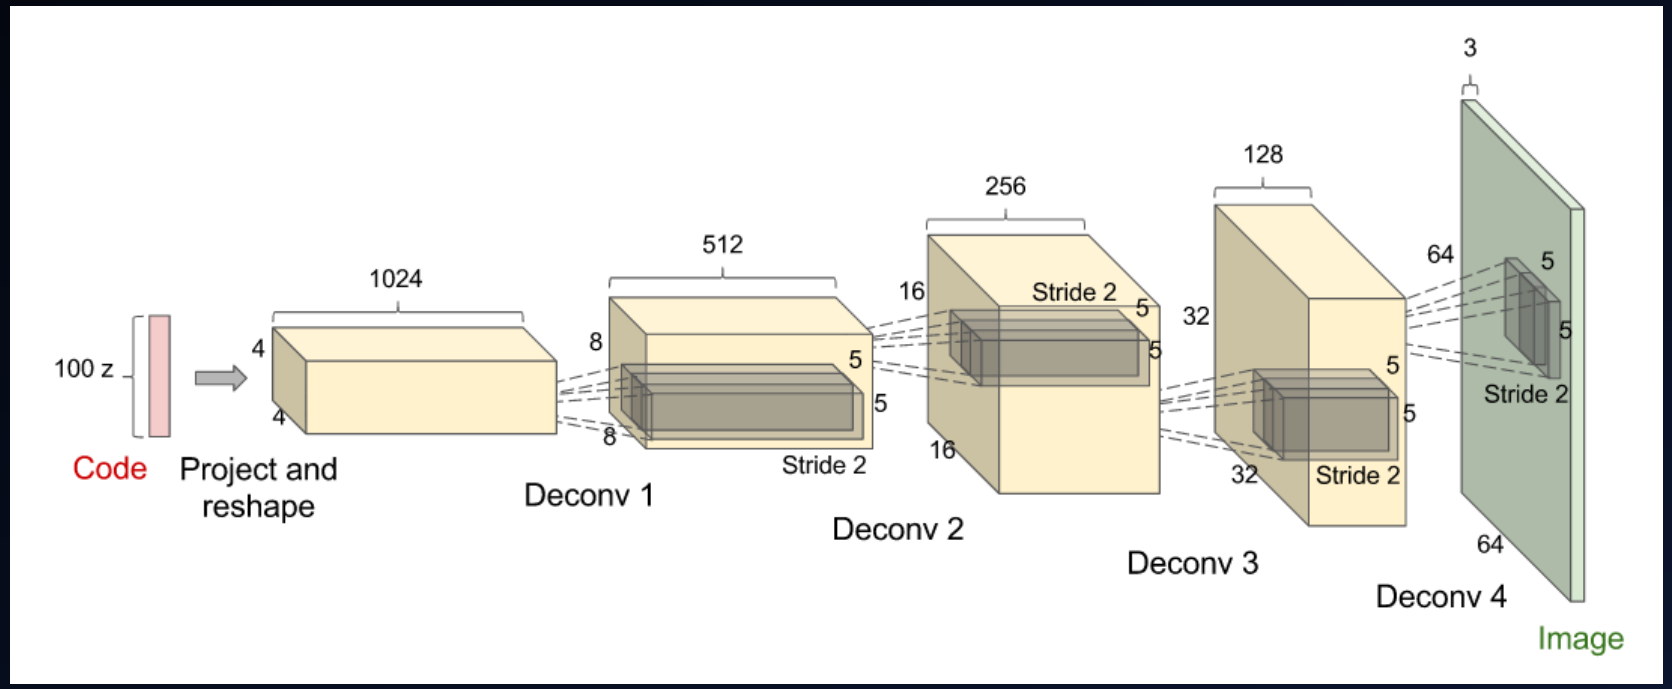

Deep Generative Models

- Variational Autoencoders[1]
- Generative Adversarial Networks[2]

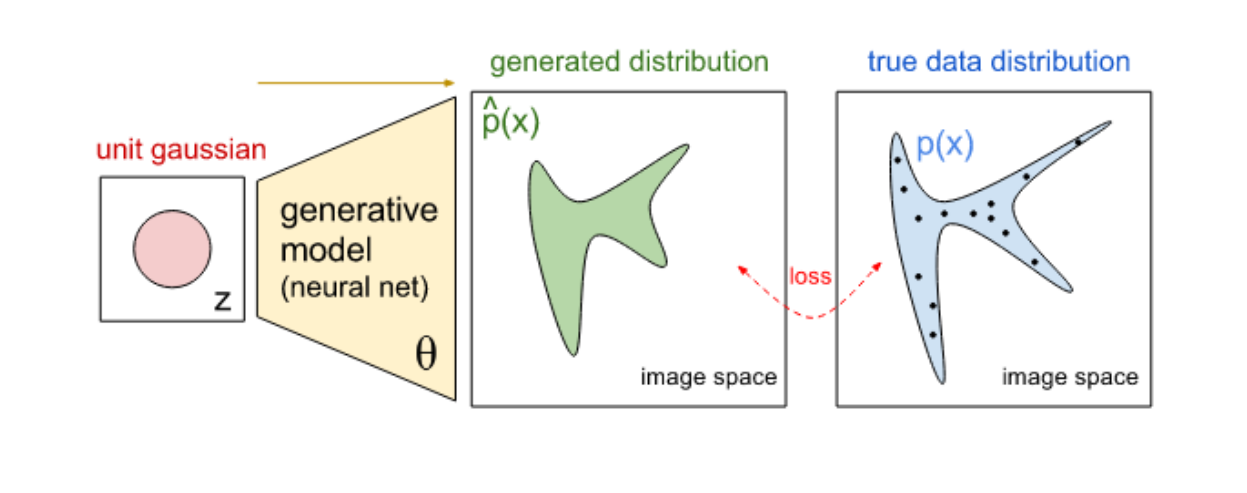

1. Kingma, Diederik P and Welling, Max. **Auto-Encoding Variational Bayes**. In The 2nd International Conference on Learning Representations (ICLR), 2013.
2. I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. **Generative adversarial nets**. In NIPS, pages 2672–2680. 2014.

Generative Adversarial Networks[1]

- Two networks compete with one another
- The generator generates imitations of data
- The discriminator distinguishes generated data from real data
- Backwards-differentiable(!)

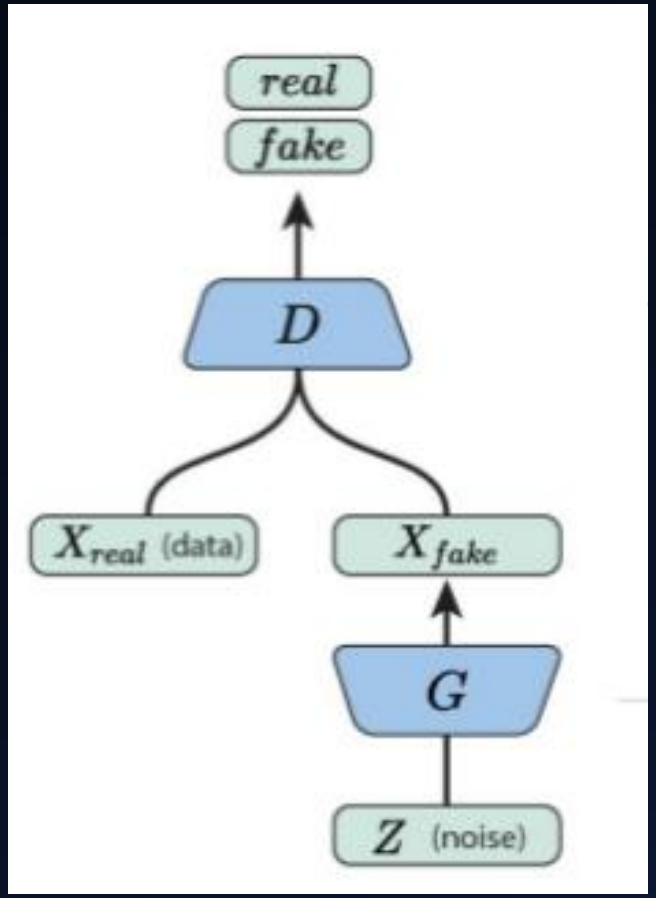

1. I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. **Generative adversarial nets**. In NIPS, pages 2672–2680. 2014.

Generative Adversarial Networks[1]

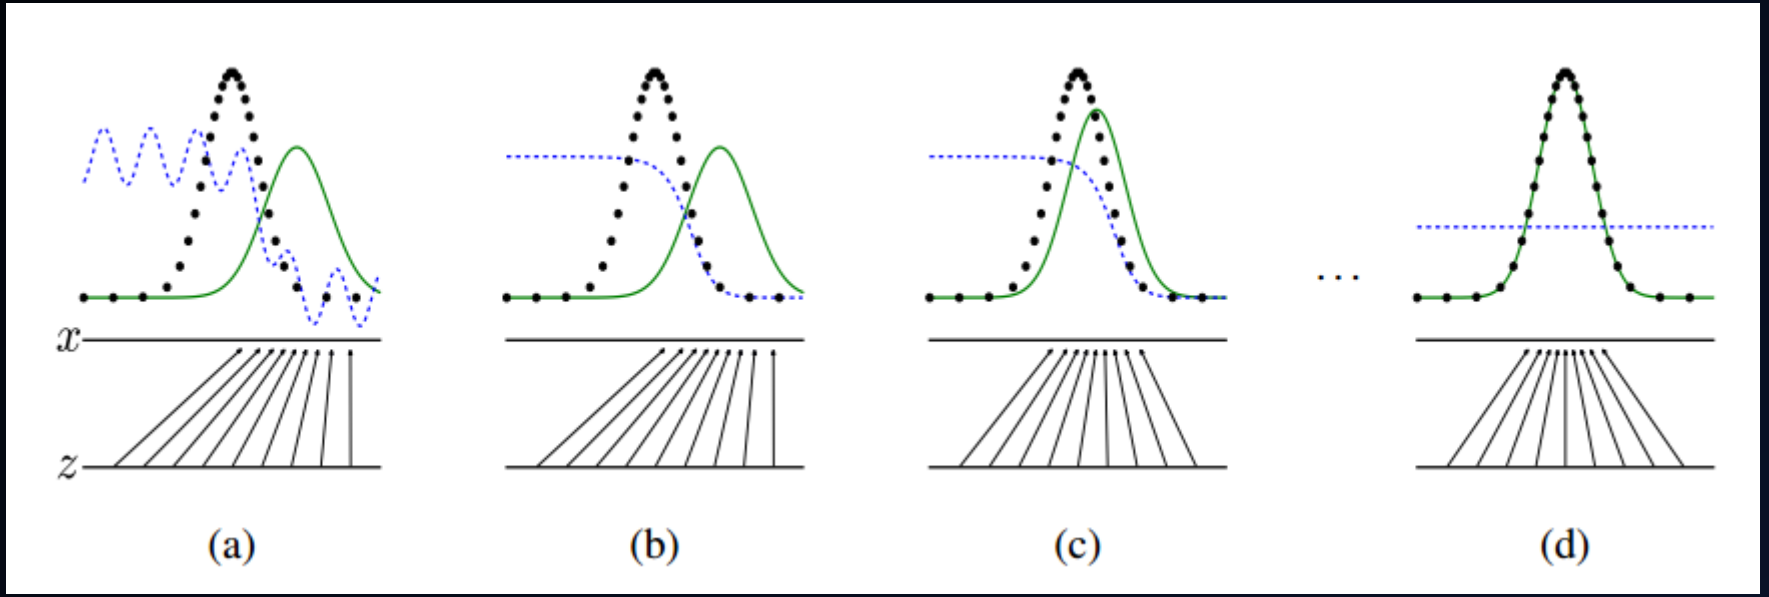

$$\min_G \max_D V(D,G) = \mathbb{E}_{\boldsymbol{x}\sim p_{\text{data}}(\boldsymbol{x})}[\log D(\boldsymbol{x})] + \mathbb{E}_{\boldsymbol{z}\sim p_{\boldsymbol{z}}(\boldsymbol{z})}[\log(1 - D(G(\boldsymbol{z})))].$$

1. I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. **Generative adversarial nets**. In NIPS, pages 2672–2680. 2014.

Results

Alec Radford, Luke Metz, and Soumith Chintala. **Unsupervised representation learning with deep convolutional generative adversarial networks**. arXiv preprint arXiv:1511.06434, 2015.

Results

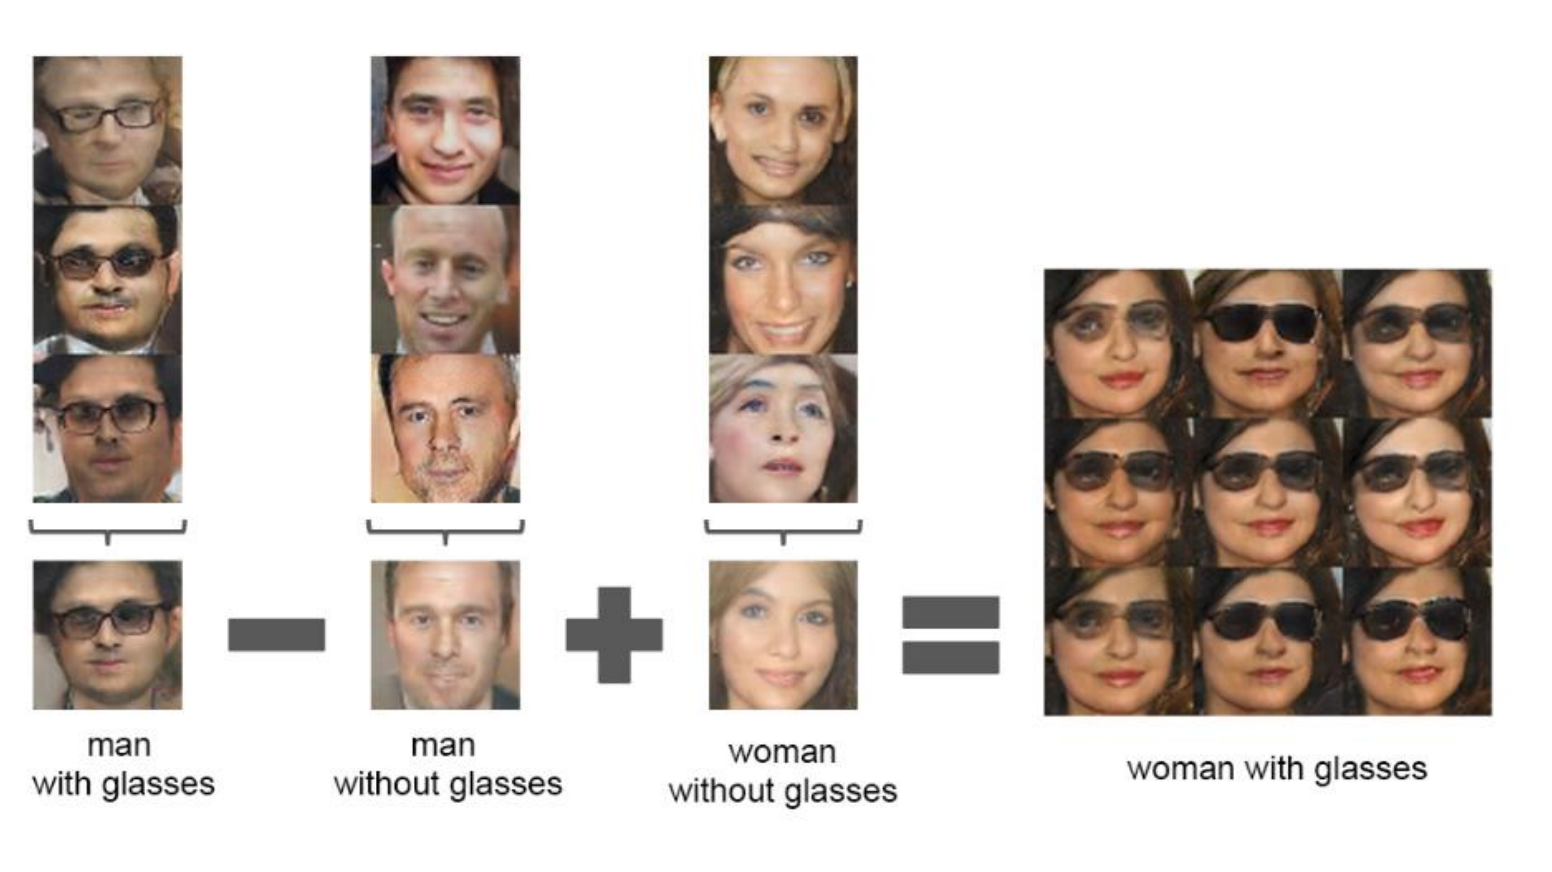

Alec Radford, Luke Metz, and Soumith Chintala. **Unsupervised representation learning with deep convolutional generative adversarial networks**. arXiv preprint arXiv:1511.06434, 2015.

Results

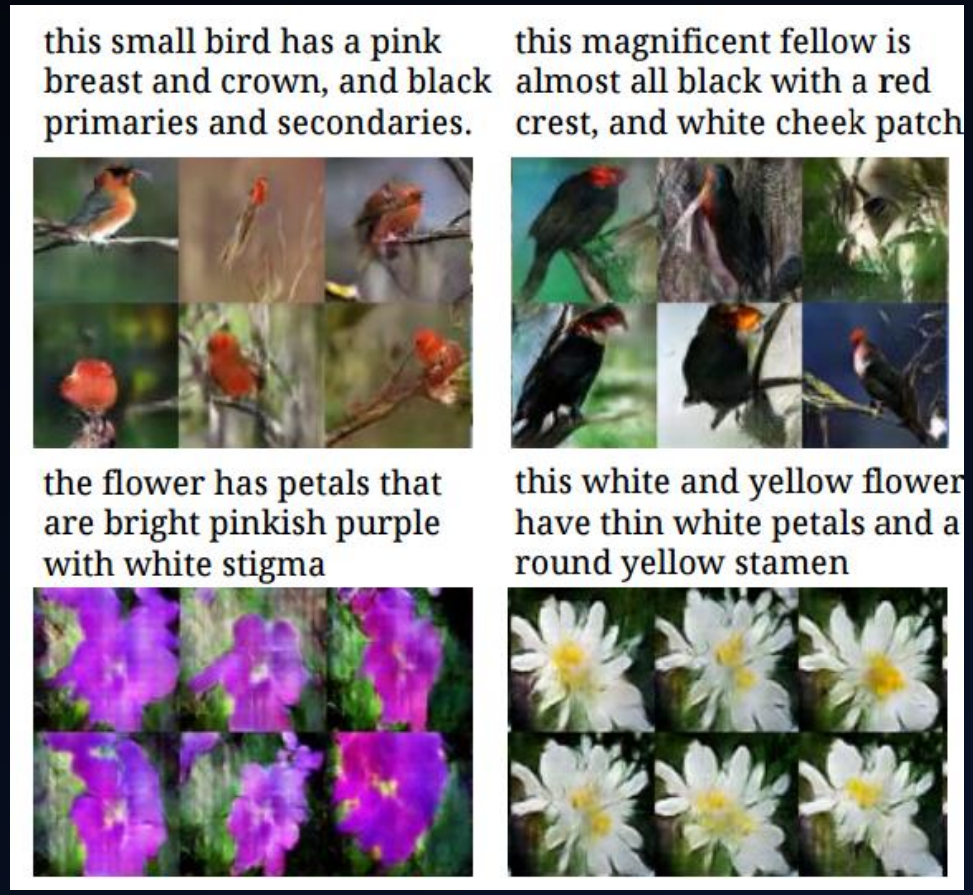

S. Reed, Z. Akata, X. Yan, L. Logeswaran, B. Schiele, and H. Lee. **Generative adversarial text-to-image synthesis**. In ICML, 2016b.

Further Research

- Disentangled representations
- Different architectures
- Combining GANs with other models
- Finding minima in non-convex games

Sources

- I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. In NIPS, pages 2672–2680. 2014.
- Kingma, Diederik P and Welling, Max. Auto-Encoding Variational Bayes. In The 2nd International Conference on Learning Representations (ICLR), 2013.
- Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. arXiv preprint arXiv:1511.06434, 2015.
- S. Reed, Z. Akata, X. Yan, L. Logeswaran, B. Schiele, and H. Lee. Generative adversarial text-to-image synthesis. In ICML, 2016b.
- A. van den Oord, N. Kalchbrenner, and K. Kavukcuoglu. Pixel recurrent neural networks. arXiv preprint arXiv:1601.06759, 2016.