

blendtorch: Stochastic Optimization of Distributional Parameters Involving Non-Differentiable Render Functions

Christoph Heindl¹

¹Computational Perception, JKU
Visual Computing, PROFACTOR GmbH
christoph.heindl@gmail.com

February 3, 2021

Abstract

We consider a gradient based parameter optimization of a stochastic computational graph consisting of random scene properties and a deterministic, non-differentiable render function. Our approach leverages ideas from Generative Adversarial Networks (GANs) and gradient estimators from reinforcement learning to jointly optimize all distributional and structural parameters based on generated images. This document should be regarded as an unfinished notebook to emphasize our main idea.

1 Optimization

Consider the objective

$$\arg \min_{\Omega_Z} \arg \max_{\Omega_D} \mathbb{E}_{\mathbf{x} \sim p(\mathbf{x})} [\log D(\mathbf{x}; \Omega_D)] + \mathbb{E}_{\mathbf{z} \sim p(\mathbf{z}|\Omega_Z)} [\log(1 - D(G(\mathbf{z}); \Omega_D))], \quad (1)$$

a variant of the objective function of GANs [1] modified as follows: The scene configuration variables \mathbf{z} are samples from structured probabilistic model governed by distributional parameters Ω_Z . The generator (render function) G transforms the scene configuration \mathbf{z} into a synthetic image and is assumed to be non-differentiable and without parameters.

The discriminator D remains unchanged compared to GANs and hence we do not consider it in the remainder of this discussion. Optimizing Ω_Z is not so straight forward, since G is non-differentiable and the parameters of the optimization are distributional. Consider optimal discriminator parameters Ω_D^* ,

then optimizing Ω_Z reduces to minimizing

$$\arg \min_{\Omega_Z} S(\Omega_Z) = \arg \min_{\Omega_Z} \mathbb{E}_{\mathbf{z} \sim p(\mathbf{z}|\Omega_Z)} \underbrace{[\log(1 - D(G(\mathbf{z}); \Omega_D^*))]}_{f(\mathbf{z})}. \quad (2)$$

In the following we derive the score-function gradient estimator [3, 2] that enables us to apply the ideas of stochastic gradient descent to Ω_Z as follows

$$\Omega_Z^{t+1} = \Omega_Z^t - \alpha \nabla_{\Omega_Z} S(\Omega_Z^t), \quad (3)$$

where $\nabla_{\Omega_Z} S(\Omega_Z^t)$ is given by

$$\begin{aligned} \nabla_{\Omega_Z} S(\Omega_Z) &= \nabla_{\Omega_Z} \mathbb{E}_{\mathbf{z} \sim p(\mathbf{z}|\Omega_Z)} f(\mathbf{z}) \\ &= \nabla_{\Omega_Z} \int p(\mathbf{z} | \Omega_Z) f(\mathbf{z}) d\mathbf{z} \\ &= \int \nabla_{\Omega_Z} p(\mathbf{z} | \Omega_Z) f(\mathbf{z}) d\mathbf{z} \\ &= \int \frac{p(\mathbf{z} | \Omega_Z)}{p(\mathbf{z} | \Omega_Z)} \nabla_{\Omega_Z} p(\mathbf{z} | \Omega_Z) f(\mathbf{z}) d\mathbf{z} \\ &= \int p(\mathbf{z} | \Omega_Z) \nabla_{\Omega_Z} \log p(\mathbf{z} | \Omega_Z) f(\mathbf{z}) d\mathbf{z} \\ &= \mathbb{E}_{\mathbf{z} \sim p(\mathbf{z}|\Omega_Z)} [\nabla_{\Omega_Z} \log p(\mathbf{z} | \Omega_Z) f(\mathbf{z})], \end{aligned} \quad (4)$$

The expectation in Equation 4 can be approximated by the following unbiased estimator

$$\nabla_{\Omega_Z} S(\Omega_Z) \approx \frac{1}{N} \sum_{\hat{\mathbf{z}}} \nabla_{\Omega_Z} \log p(\hat{\mathbf{z}} | \Omega_Z) f(\hat{\mathbf{z}}), \quad (5)$$

with $\hat{\mathbf{z}} \sim p(\hat{\mathbf{z}} | \Omega_Z)$. Note that in the above derivation, $G(\mathbf{z})$ only appears in the weighting term $f(\mathbf{z})$ for which no gradients are required. Depending on the probabilistic model $\log p(\hat{\mathbf{z}} | \Omega_Z)$ usually decomposes into simpler terms.

In an oscillating fashion we update the structural parameters Ω_D of discriminator and distributional parameters Ω_Z until an equilibrium is reached in which the discriminator cannot distinguish between samples of the target and the generator distribution.

References

- [1] Ian Goodfellow et al. “Generative adversarial nets”. In: *Advances in neural information processing systems*. 2014, pp. 2672–2680.
- [2] John Schulman et al. “Gradient estimation using stochastic computation graphs”. In: *Advances in Neural Information Processing Systems*. 2015, pp. 3528–3536.
- [3] Ronald J Williams. “Simple statistical gradient-following algorithms for connectionist reinforcement learning”. In: *Machine learning* 8.3-4 (1992), pp. 229–256.