The Mechanics of n-Player Differentiable Games

Sergey Pozdniakov, Ekaterina Trofimova,
Ivan Bulygin, Alexander Lyzhov, Kaziy Botashev

Problem statement

Hessian of a *n*-player game with parameter w:

$$H(w) = \nabla_w \xi(w)^T,$$

where ξ refers to the dynamics of the game (or simultaneous gradient)

Decomposition of the Hessian:

$$\underbrace{\begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}}_{H} = \underbrace{\begin{pmatrix} 0 & 0 \\ 0 & 0 \end{pmatrix}}_{S} + \underbrace{\begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}}_{A}.$$

SGA

Idea of the game dynamics adjustment by

$$\xi + \lambda A^T \xi$$

The dynamic ξ is a Hamiltonian vector field, since it conserves the level-sets of the Hamiltonian H

Quadratic case

All losses are quadratic $\implies \xi = Gw$, Nash equilibrium is $w = 0$ (if exist)

$w_{n+1} = w_n + lr\, Gw_n$ converges, but what about $w_{n+1} = w_n + lr(Gw_n + \varepsilon)$?

Convergence means, that $\lim_{n\to\infty} E[w_n] = 0$ and $\Sigma[w_n]$ is finite (equivalently $E[w_n^2]$ is finite)

Stationary noise

$M[\varepsilon]$ and $\Sigma[\varepsilon]$ doesn't depend on w

deterministic case $w_{n+1} = Qw_n$ $(Q = I + lr\,G)$ converges $(\rho(Q) < 1)$

what about $w_{n+1} = Qw_n + \varepsilon$?

1) $E[w_{n+1}] = QE[w_n] + lrE[\varepsilon] = QE[w_n]$ converges to zero

2) $\Sigma[w_{n+1}] = Q\Sigma[w_n]Q^T + lr^2\Sigma[\varepsilon]$

$$\Sigma[w_n] = \sum_{i=0}^{n-1} Q^i lr^2 \Sigma[\varepsilon] Q^{iT}$$

$Q^i lr^2 \Sigma[\varepsilon] Q^{iT} < const\ \rho^{2i}(Q)$ pointwise $\implies \Sigma[w_n]$ is finite

Non-stationary noise

Our example is one-dimensional:

$$w_{n+1} = (1 + lr\ g)w_n + lr\ \varepsilon\ w_n \qquad E[\varepsilon w_n] = 0$$

$$E[w_{n+1}^2] = E[w_n^2]((1 + lr\ g)^2 + lr^2\ E[\varepsilon^2])$$

$lr \in \left[0, \frac{-2g}{g^2+E[\varepsilon^2]}\right]$ both dynamics converges

$lr \in \left[\frac{-2g}{g^2+E[\varepsilon^2]}, \frac{-2}{g}\right]$ deterministic dynamic converges, stochastic diverges

$lr > \frac{-2}{g}$ both dynamic diverges

We can calculate final covariation analytically!

$$\Sigma_\infty = \sum_{i=0}^{\infty} Q^i lr^2 \Sigma[\varepsilon] Q^{iT}$$

$$Q = CJC^{-1} \qquad \Sigma_\infty = lr^2 \sum_{i=0}^{\infty} CJ^iC^{-1}\Sigma[\varepsilon]C^{-T}J^{iT}C^T$$

$$Q = \begin{bmatrix} \frac{1}{2} & -\frac{1}{5} \\ -\frac{1}{5} & \frac{1}{3} \end{bmatrix} \quad \Sigma[\varepsilon] = \begin{bmatrix} 5 & 1 \\ 1 & 5 \end{bmatrix} \quad C = \begin{bmatrix} \frac{2}{3} & -\frac{3}{2} \\ 1 & 1 \end{bmatrix} J = \begin{bmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{bmatrix} \quad J^i = \begin{bmatrix} \lambda_1^i & 0 \\ 0 & \lambda_2^i \end{bmatrix} \quad \lambda_1 = \frac{1}{5} \quad \lambda_2 = \frac{19}{30}$$

$$\Sigma_\infty = lr^2 \sum_{i=0}^{\infty} \begin{bmatrix} \frac{80}{169}(\lambda_1^2)^i + \frac{180}{169}(\lambda_1\lambda_2)^i & -\frac{36}{169}(\lambda_1^2)^i + \frac{36}{169}(\lambda_1\lambda_2)^i \\ -\frac{81}{169}(\lambda_1\lambda_2)^i + \frac{81}{169}(\lambda_2)^i & \frac{180}{169}(\lambda_1\lambda_2)^i + \frac{80}{169}(\lambda_2^2)^i \end{bmatrix} \qquad \sum_{i=0}^{\infty} q^i = \frac{1}{1-q}$$

Done!

Can be applied for arbitrarily dimensionalities.

Example of final distribution

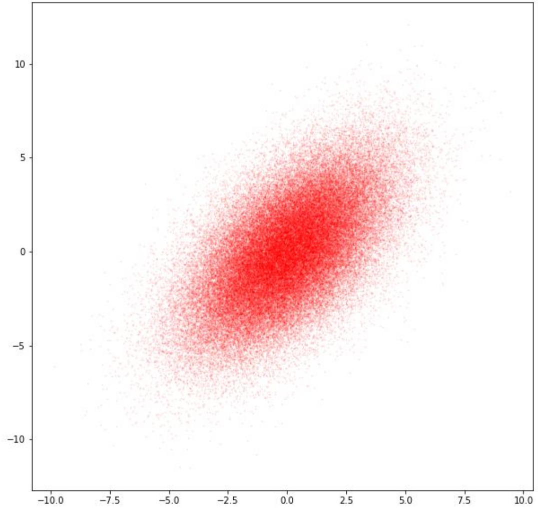

Some properties:

$$w_{n+1} = Qw_n + lr\,\varepsilon$$
$$Q = I + lr\,G$$
$$\Sigma_\infty \sim \frac{lr^2}{1-\rho^2(Q)}$$
$$\rho_Q \sim 1 - lr$$
$$\Sigma_\infty \sim lr$$
$$\Delta w \sim \sqrt{lr}$$
$$\Sigma_\infty \sim (Amplitude\ of\ noise)^2$$
$$\Delta w \sim Amplitude\ of\ noise$$

Simple two-player game

Lets consider Player 1 and Player 2 with the respective loss functions:

$$l_1(x, y) = \frac{1}{2}x^2 + 10xy, \quad l_2(x, y) = -\frac{1}{2}y^2 + 10xy$$

In this game Player 1 controls the variable x and Player 2 controls y respectively.

The simultaneous gradient is given by:

$$\xi = (\frac{\delta l_1}{\delta x}, \frac{\delta l_2}{\delta y}) = (x + 10y, -y + 10x)$$

SGA convergency

GD convergency

Learning rate = 0.005

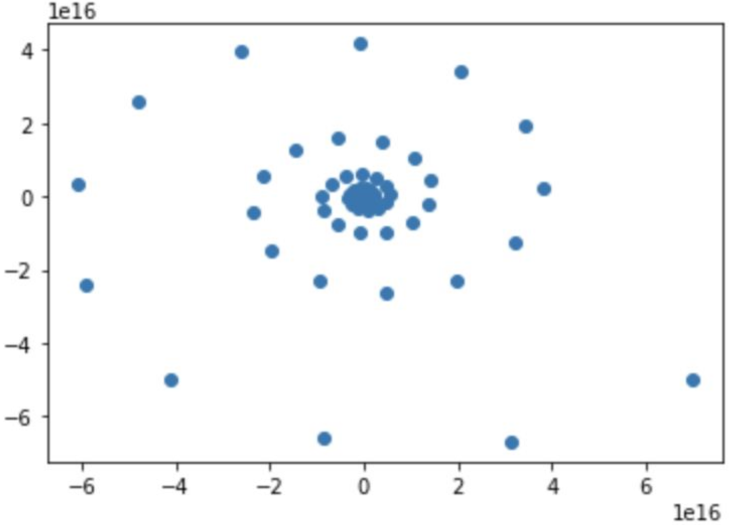

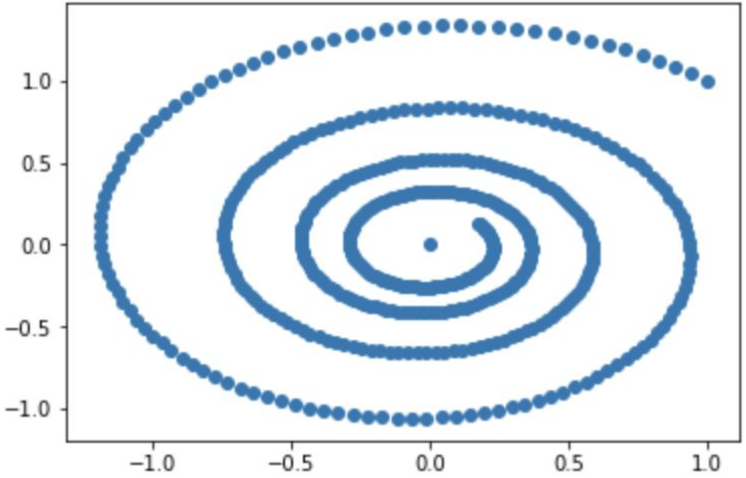

Simple four-player game

$$\ell_1(w,x,y,z) = \frac{\epsilon}{2}w^2 + wx + wy + wz$$
$$\ell_2(w,x,y,z) = -wx + \frac{\epsilon}{2}x^2 + xy + xz$$
$$\ell_3(w,x,y,z) = -wy - xy + \frac{\epsilon}{2}y^2 + yz$$
$$\ell_4(w,x,y,z) = -wz - xz - yz + \frac{\epsilon}{2}z^2,$$

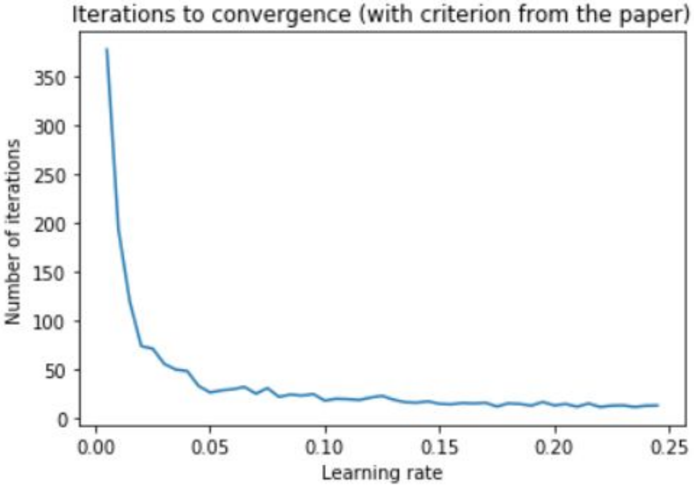

Simultaneous GD

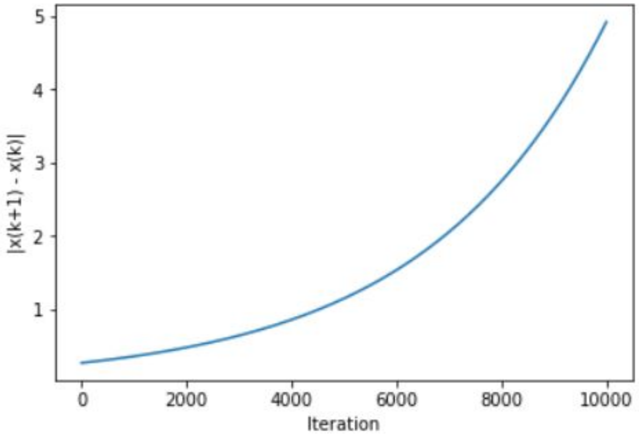

Symplectically adjucted GD

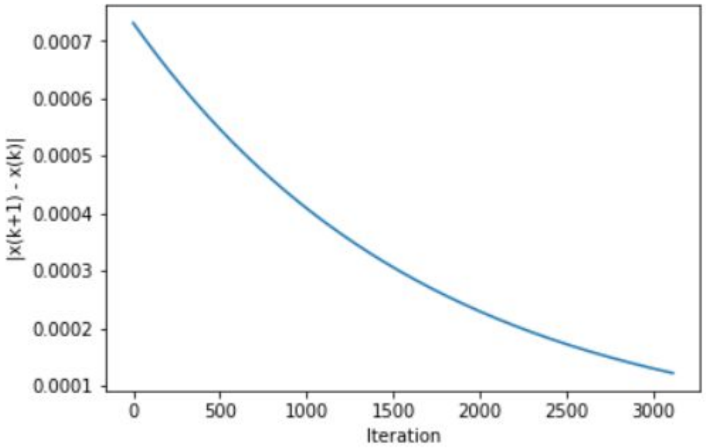

Rock-paper-scissors experiment

1-round game with simultaneous GD

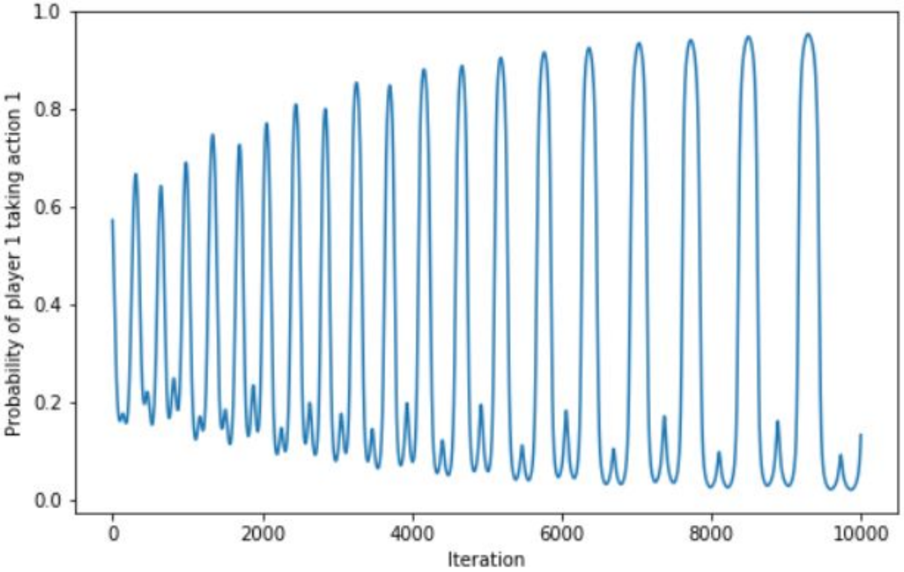

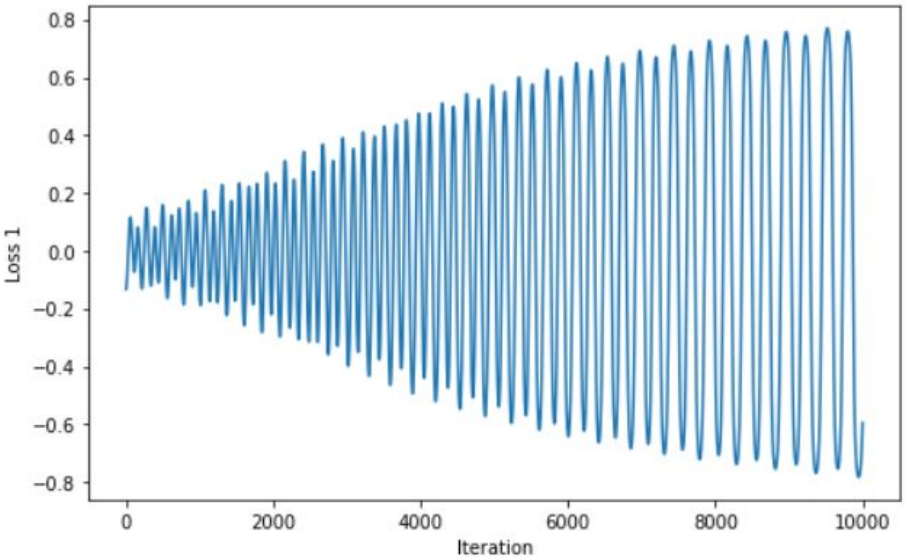

Rock-paper-scissors experiment

1-round game with symplectically adjucted GD

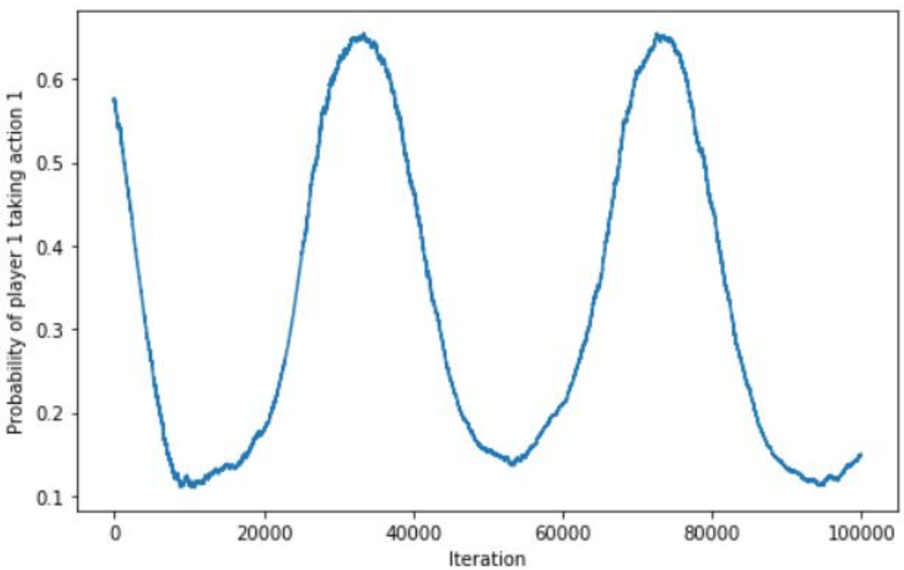

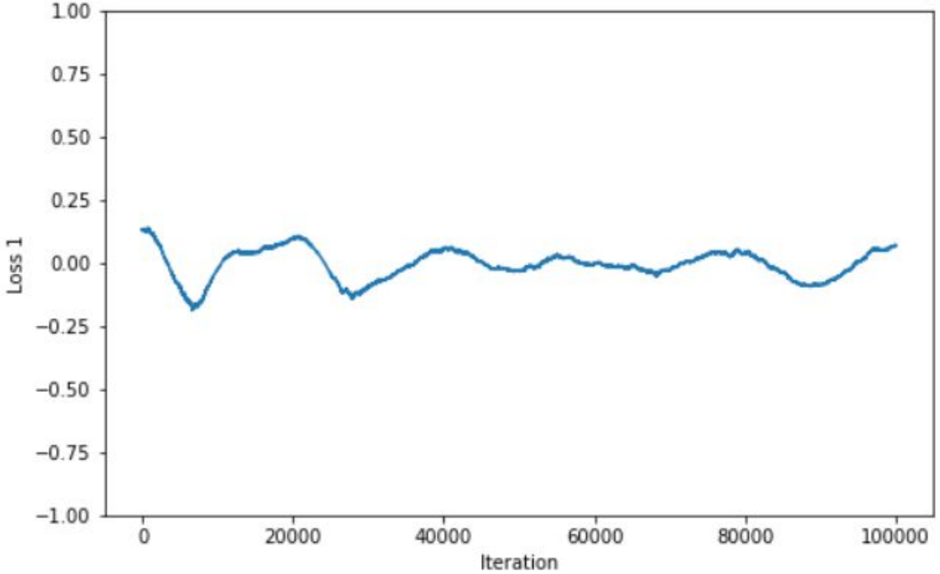

Rock-paper-scissors experiment: eye candy

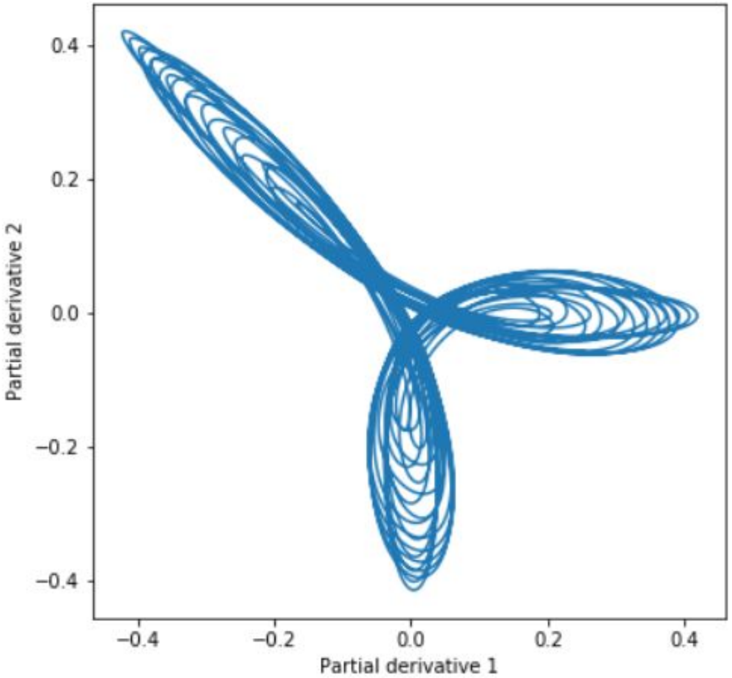

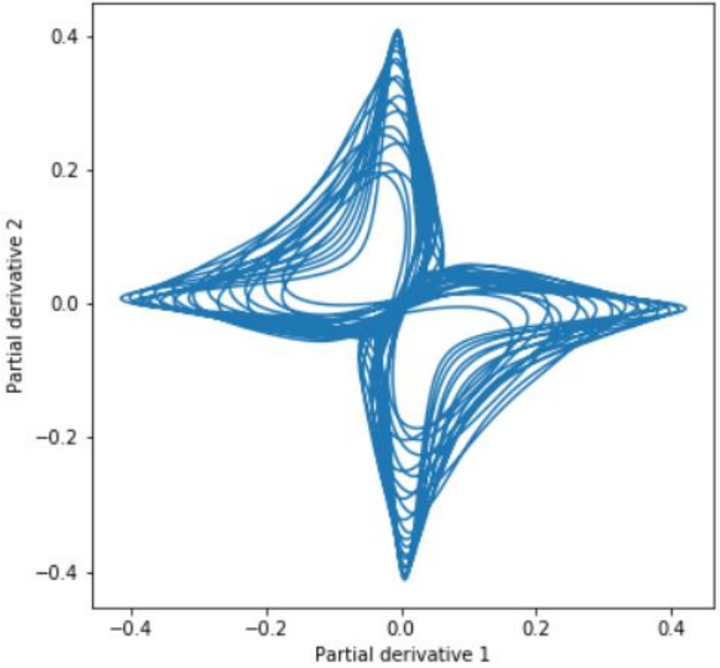

Experiments: GAN

Formulation

The goal is to learn Generator approximate 4x4 grid of Gaussians distributions:

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

Generator & **Discriminator** are MLPs

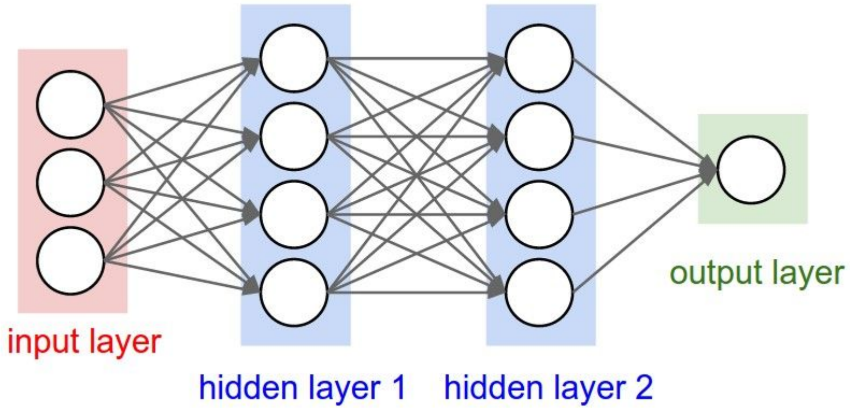

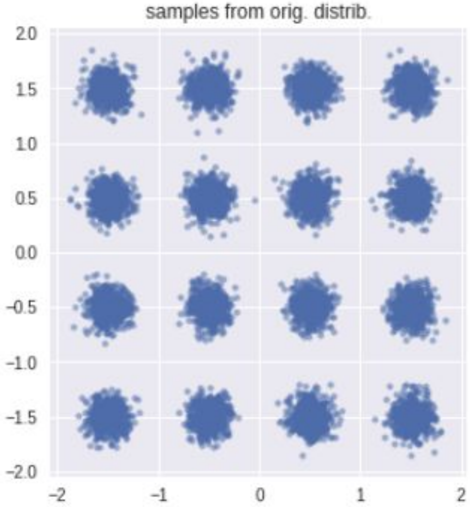

Hypothesis: SGA accelerate convergence and eliminate mode-collapse

Experiments: GAN

Results

Without SGA

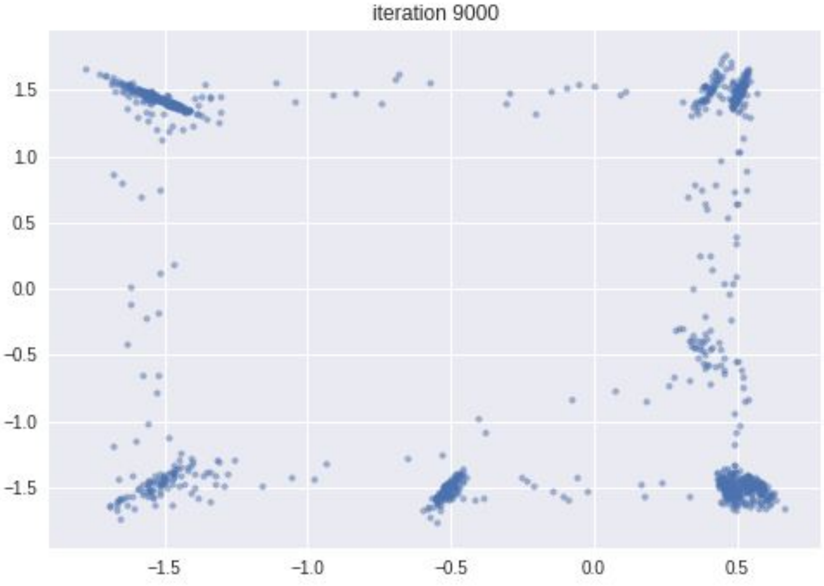

With SGA

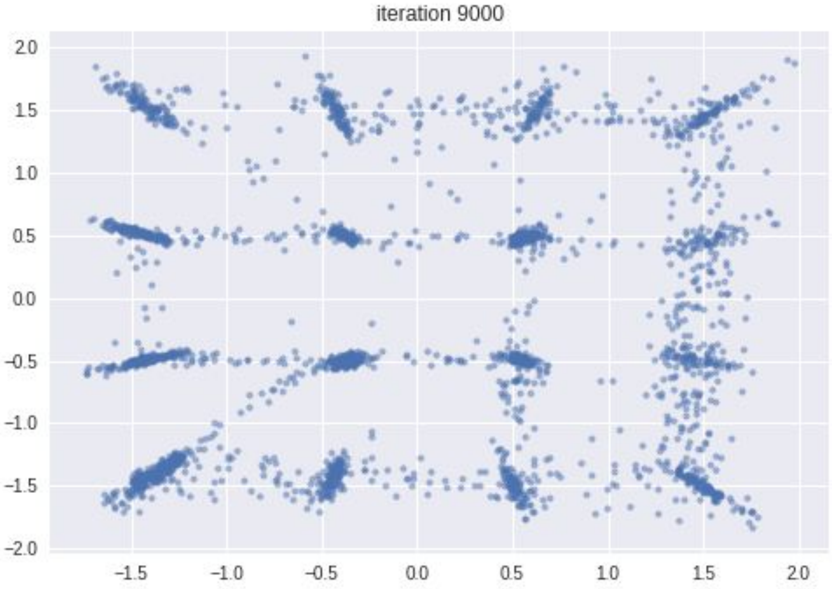

Job to be done

There is a list of things we found interesting to investigate further:

- Perform experiments with real stochasticity in policy gradient methods
- Compare SGA and Unrolled GAN for mode-collapse elimination

The Mechanics of n-Player Differentiable Games

Sergey Pozdniakov, Ekaterina Trofimova,
Ivan Bulygin, Alexander Lyzhov, Kaziy Botashev