Black-Box Optimization with Local Generative Surrogates (L-GSO)

RealML workshop, ICML 2020

Sergey Shirobokov*, Vladislav Belavin*,
Michael Kagan, Andrey Ustyuzhanin,
Atilim Gunes Baydin

Contact:

s.shirobokov17@imperial.ac.uk (Twitter: @SergeyShir994),
vbelavin@hse.ru

Problem statement

TL;DR: We approximate a stochastic black-box function with generative neural networks to enable gradient based optimisation

Intractable

$$\mathrm{argmin}_{\psi} E_y[R(y_{\psi})] = \int R(y)\,\boxed{p(y|x;\psi)}\,q(x)\,dx\,dy \approx \sum_x R(F(x,\psi))$$

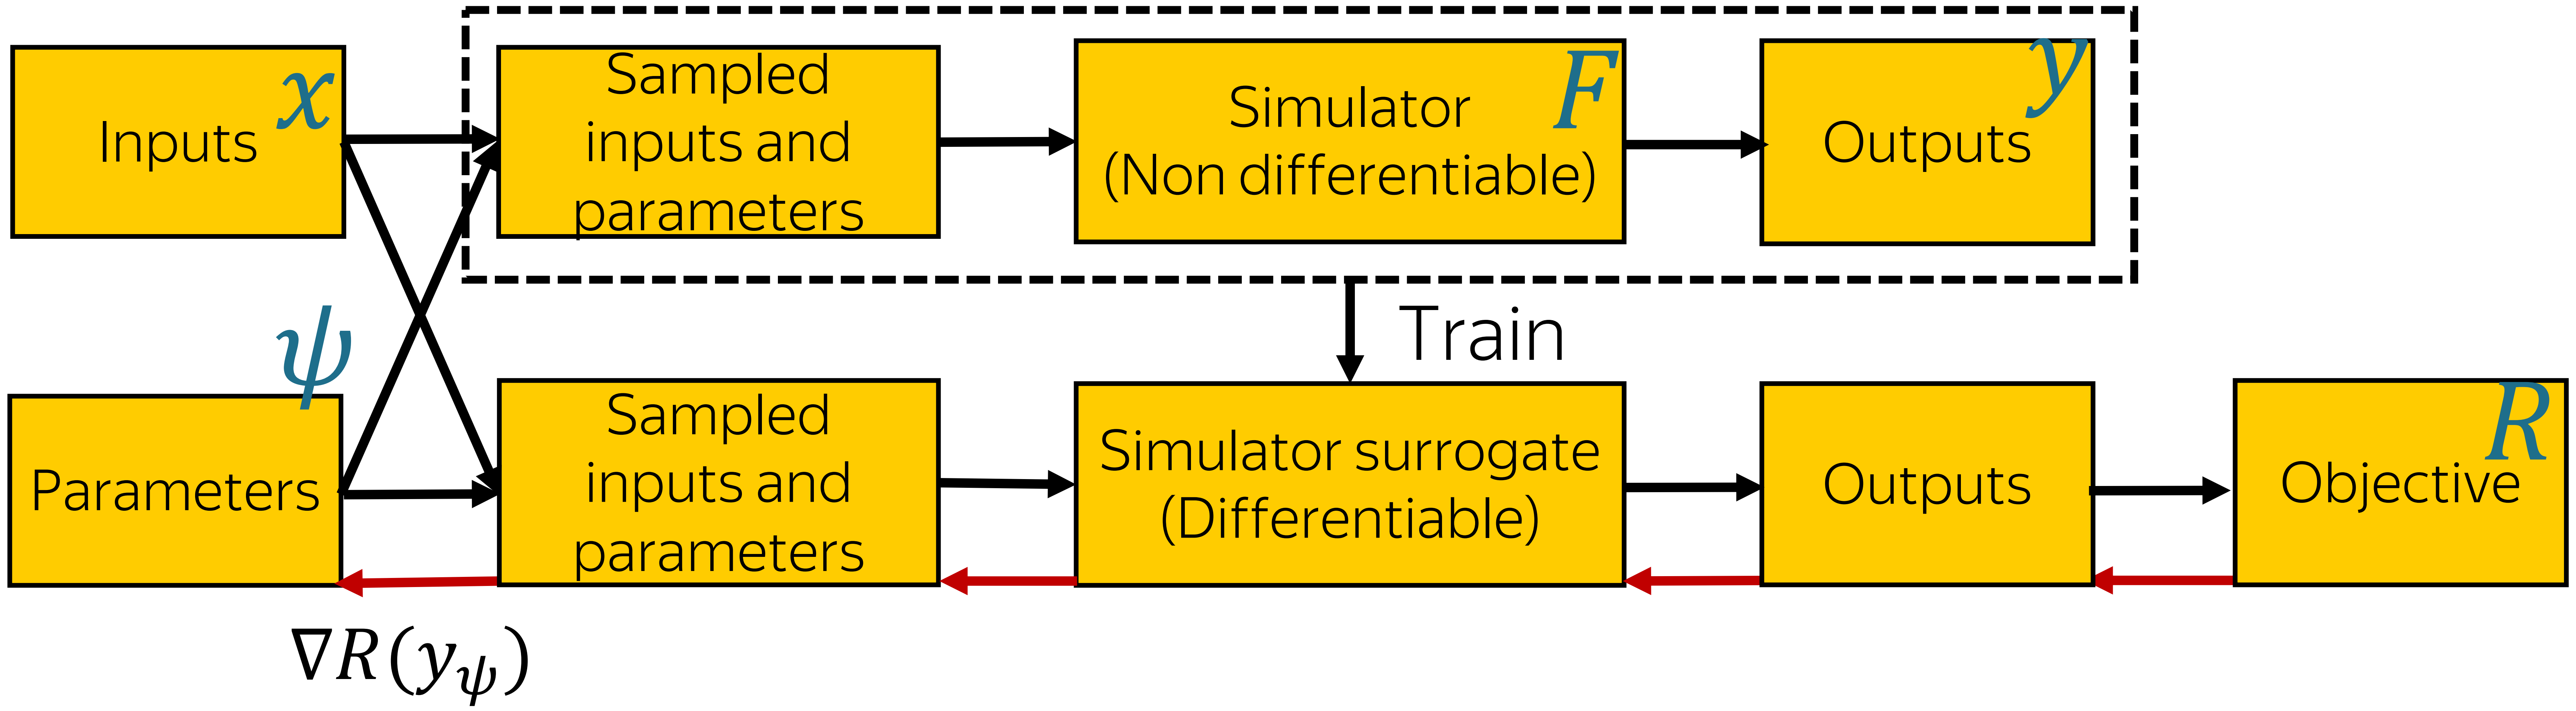

Our Method (L-GSO)

Approximate: $\nabla_\psi E_y[R(y_\psi)] \approx \sum_x \nabla_\psi R(S_\theta(z,x,\psi))$

- S_θ is a Deep Generative model (GAN, FFJORD, etc)

- To fight ”curse of dimensionality”: We train surrogate in a local neighbourhood of parameter space and perform optimisation step

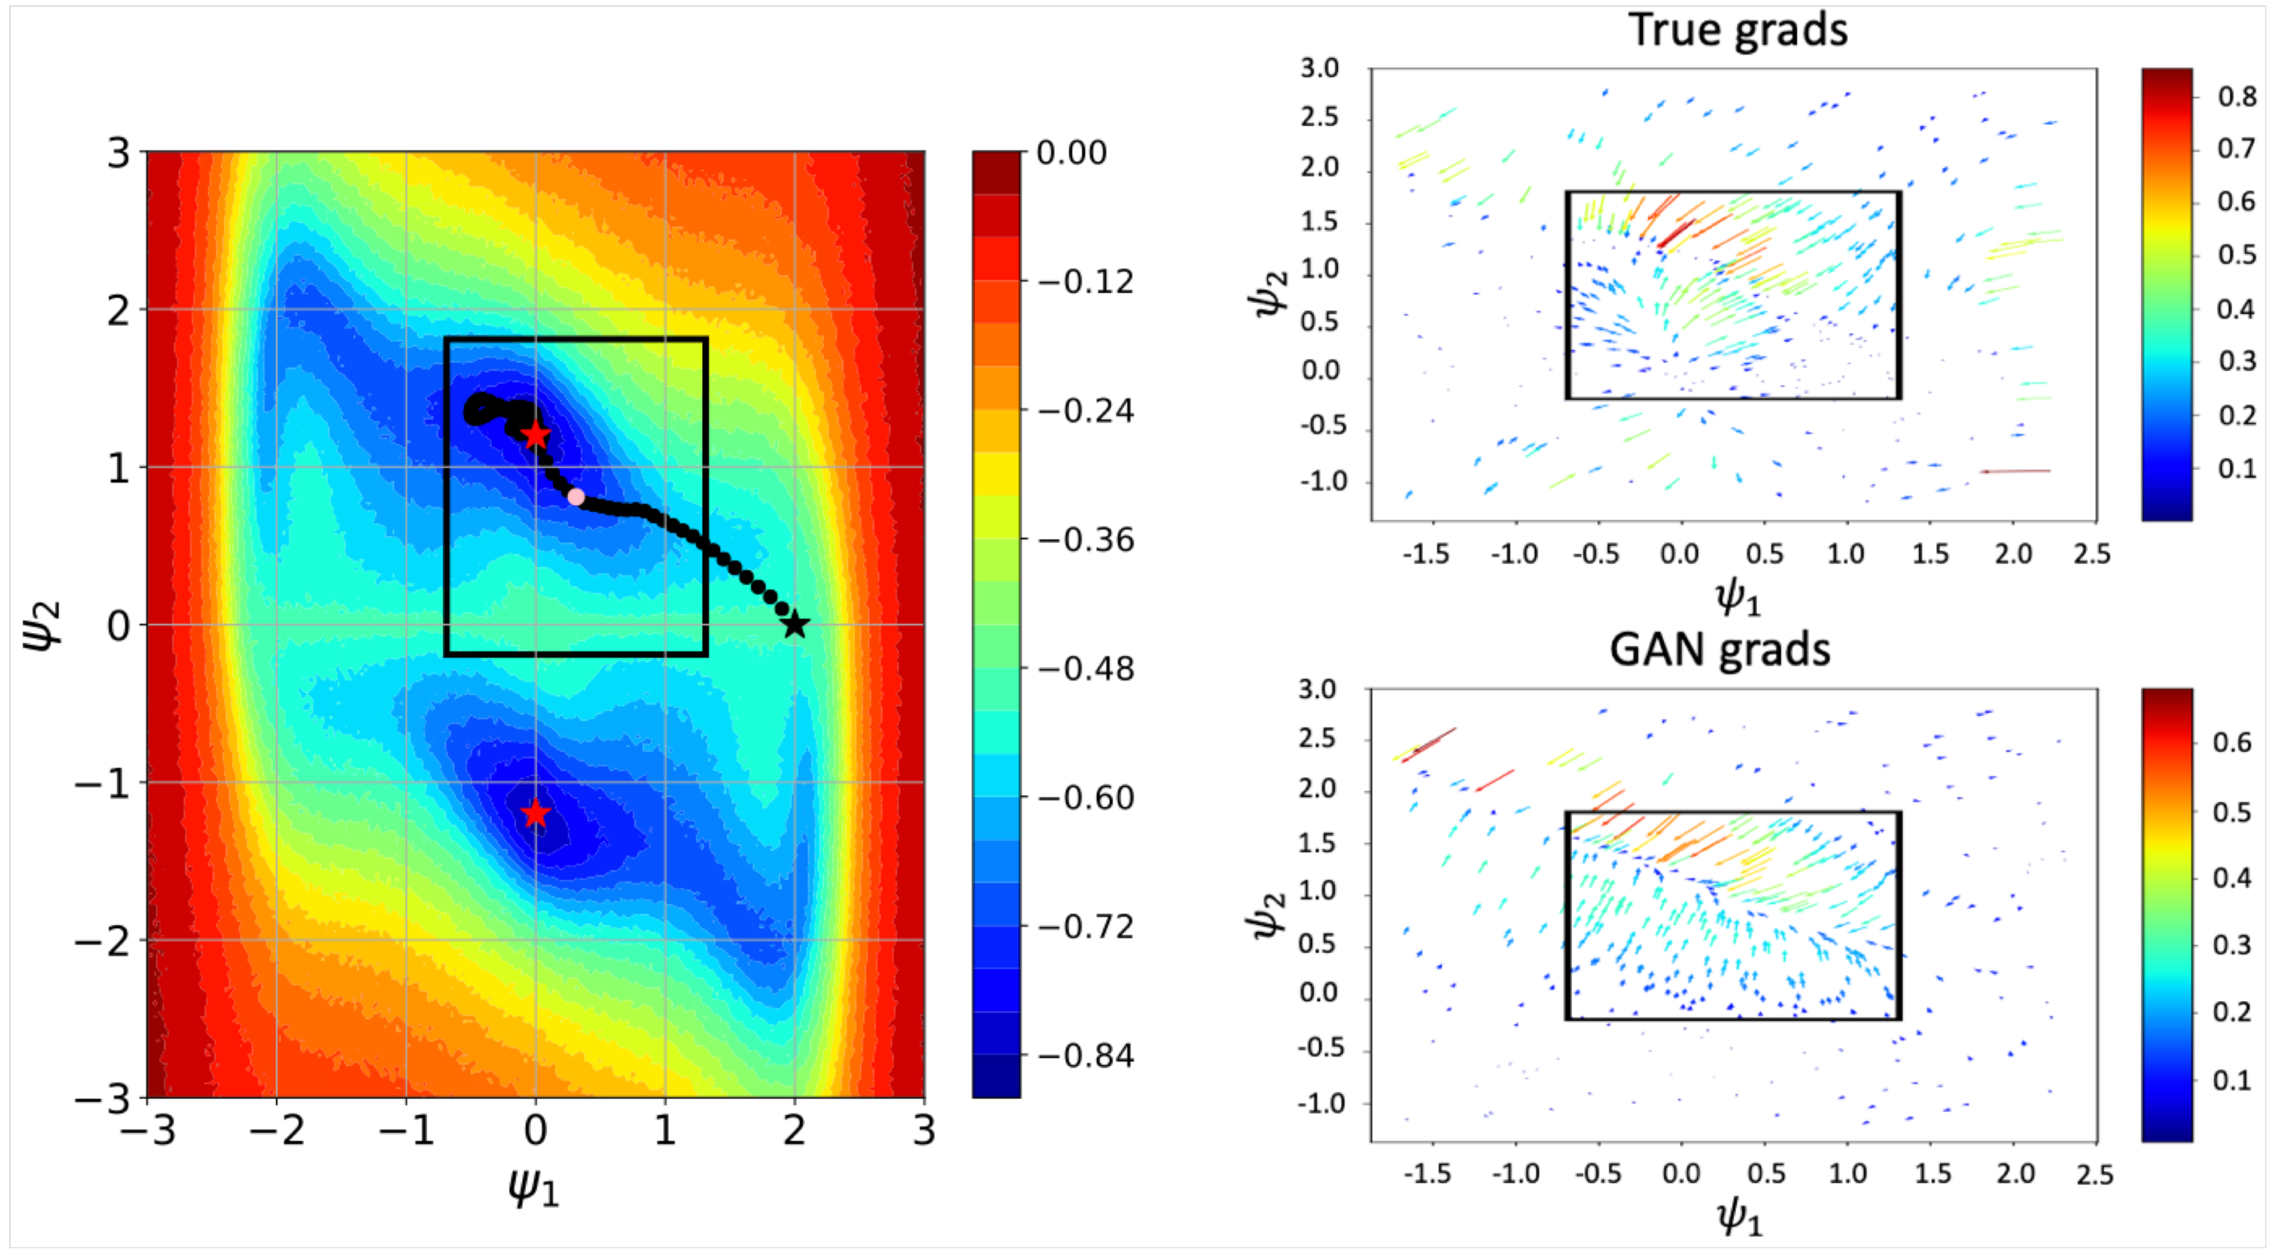

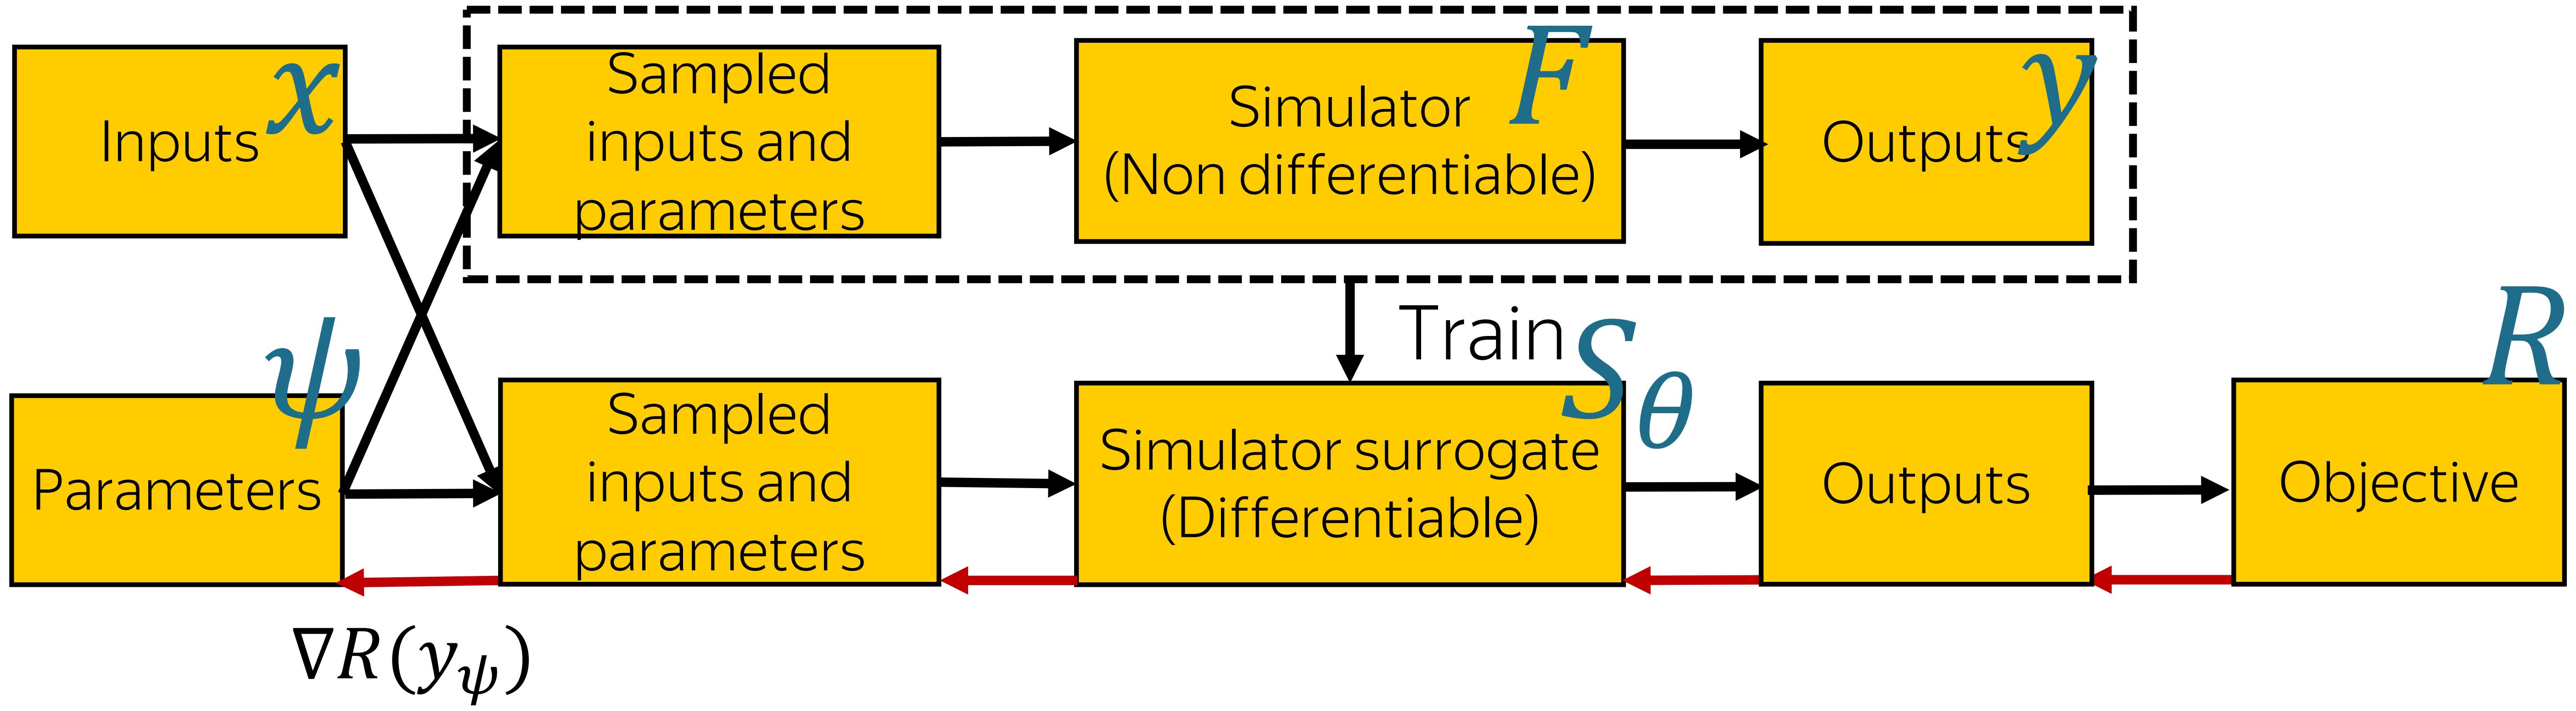

Results

Toy problems: Rosenbrock function, Degenerate Rosenbrock, Neural Network weights optimisation

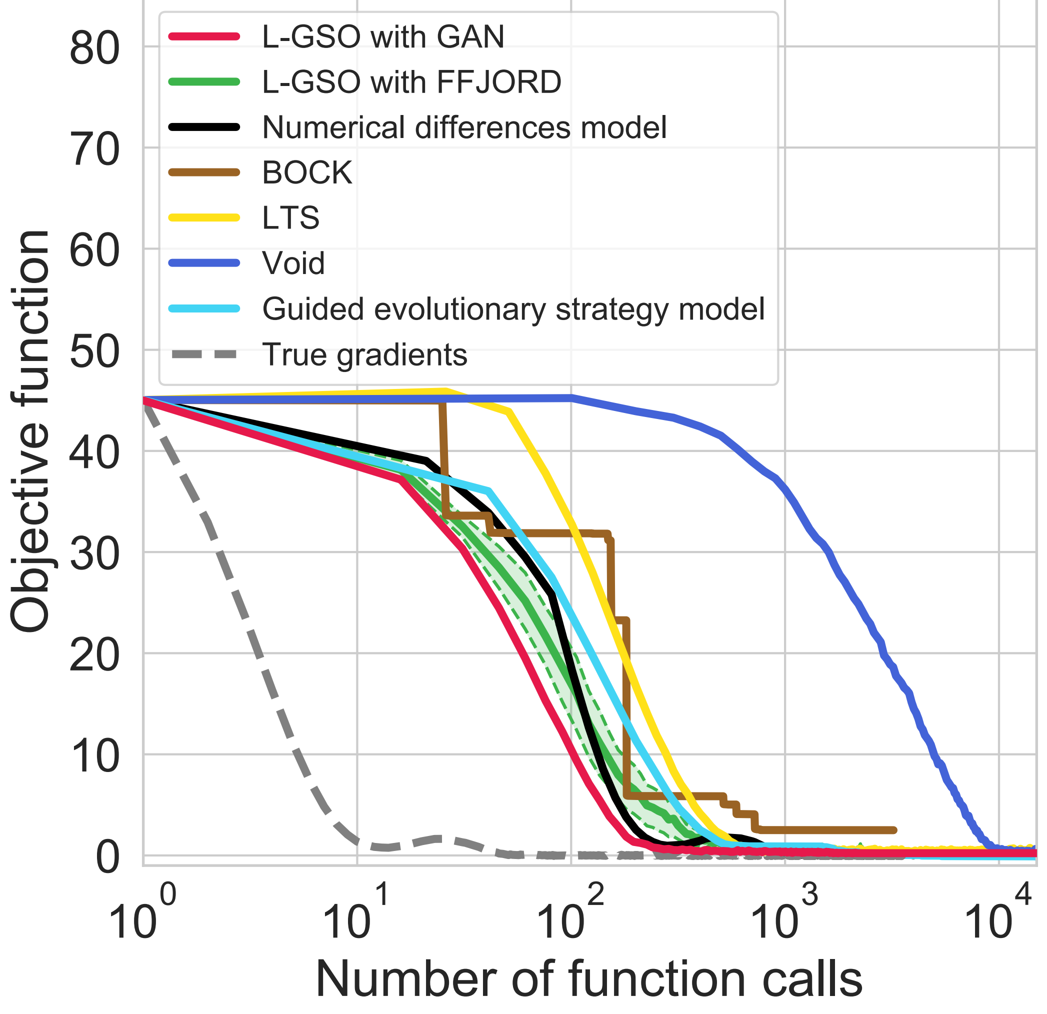

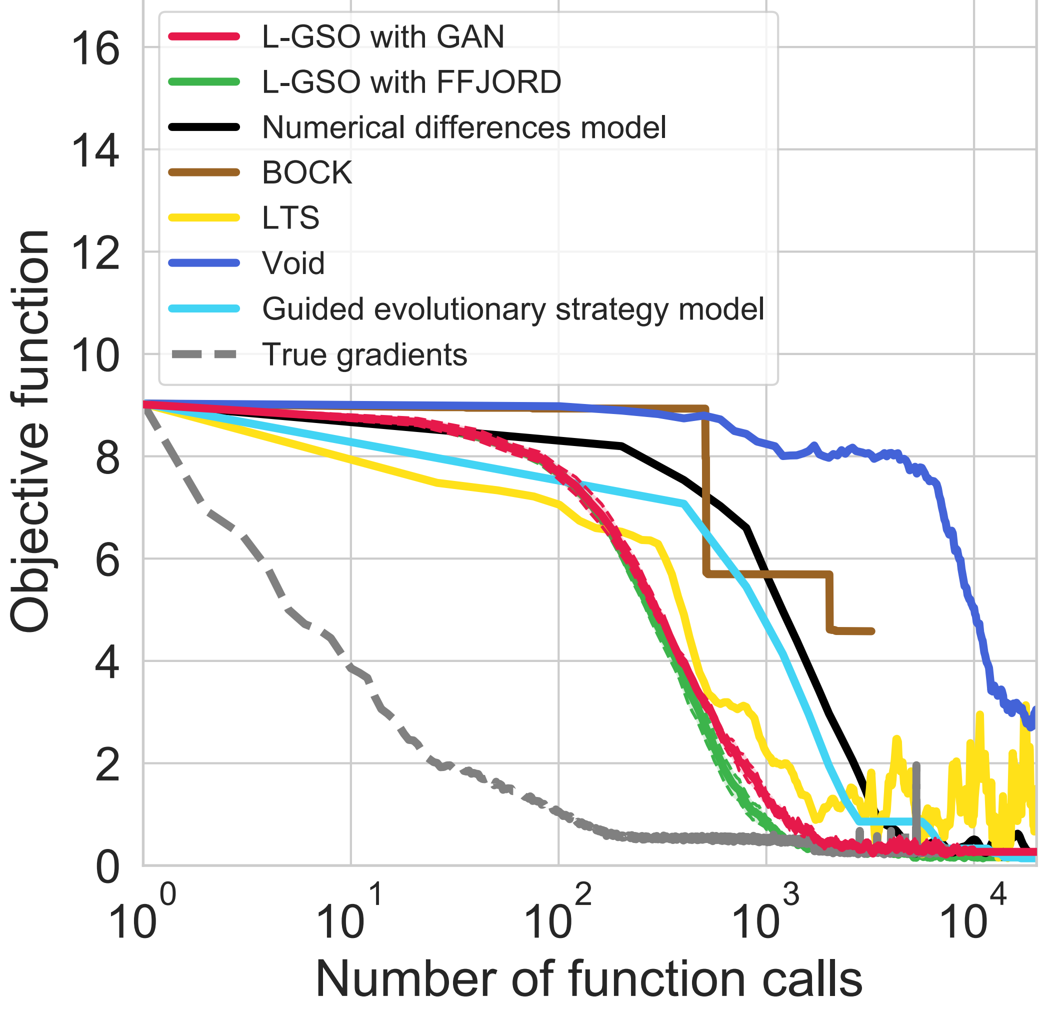

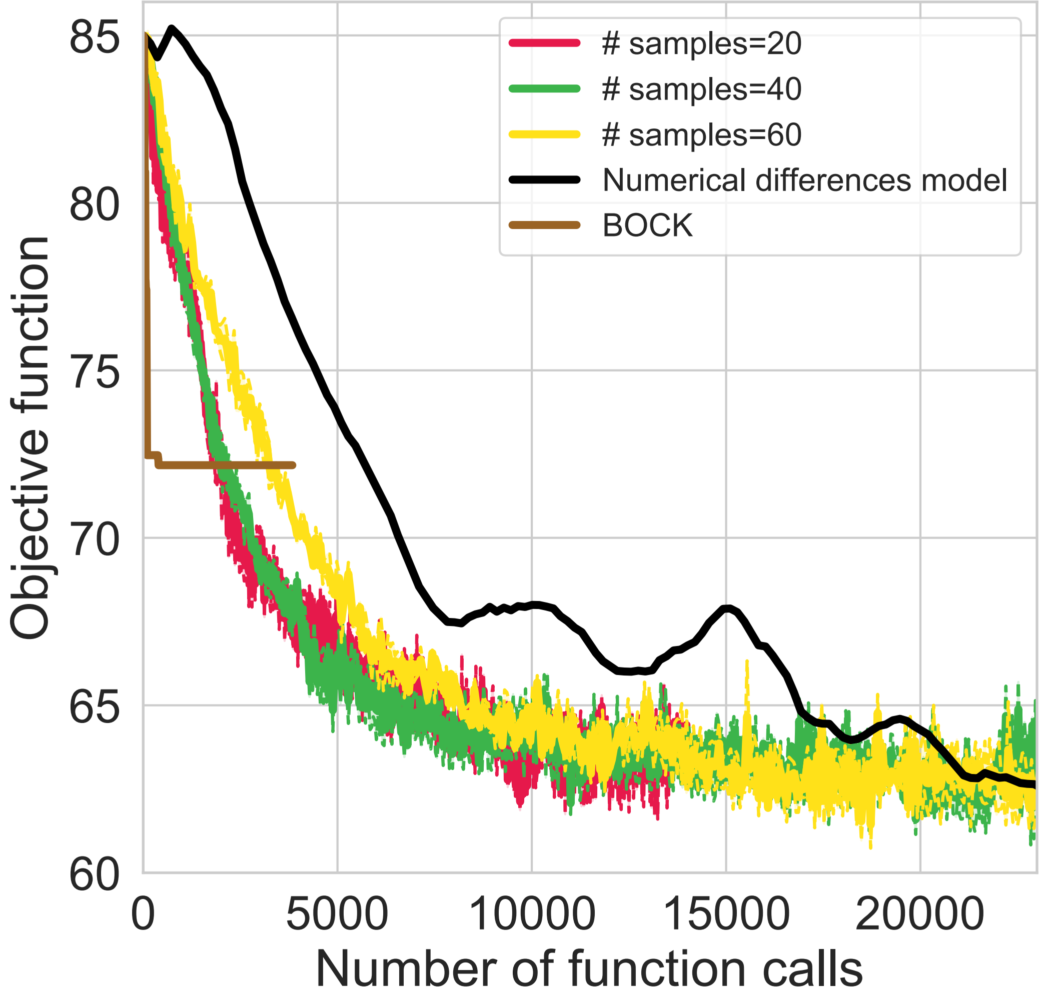

Optimisation in 42 dimensional space of physics simulator

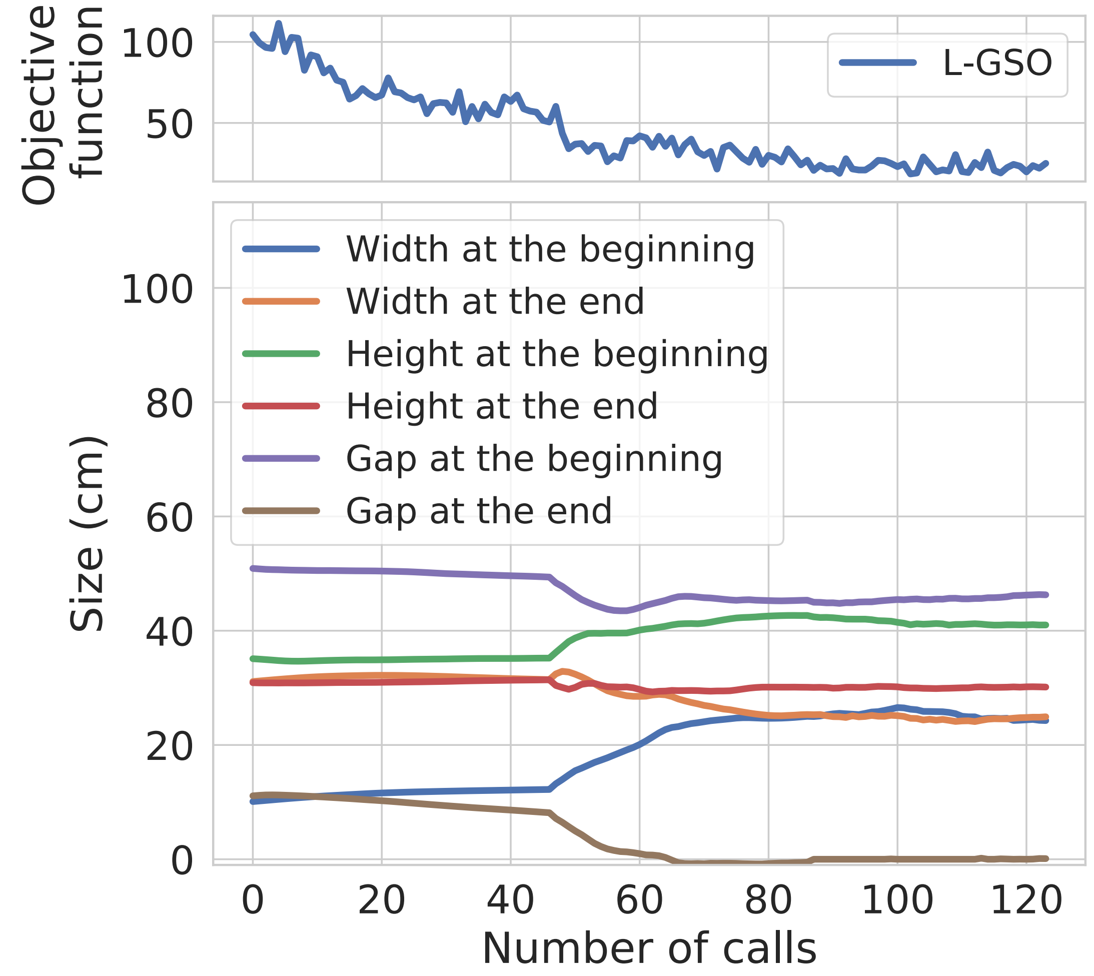

- Our algorithm is comparable with **all** baselines and outperforms REINFORCE- and evolutionary-based algorithms in speed of convergence
- Our algorithm outperforms all algorithms in high-dimensional setting when parameters are constrained to a lower dimension manifold
- We do not observe bias in gradients

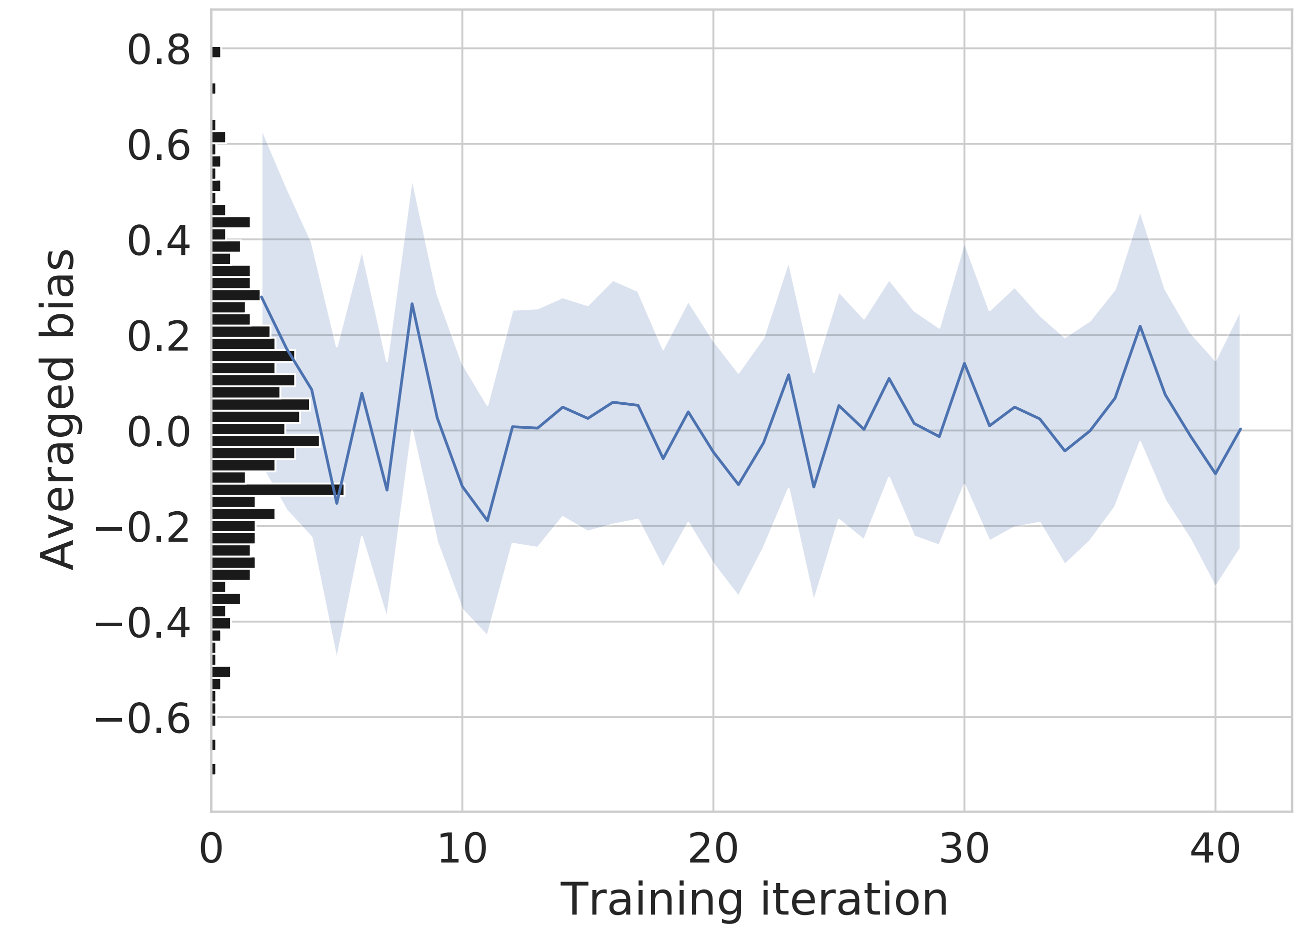

Contacts:

- ArXiv: 2002.04632 (link)

- s.shirobokov17@imperial.ac.uk
- vbelavin@hse.ru
- @SergeySHir994 (link)