


Complexity guarantees and polling strategies for Riemannian direct-search methods



Bastien Cavarretta, F. Goyens, C. W. Royer, F. Yger

LAMSADE, Université Paris Dauphine-PSL, (bastien.cavarretta@dauphine.psl.eu)

Motivation: Adversarial Attacks of Neural Networks



- **Goal:** Fool a classifier *without* knowing weights, architecture.
- **Strategy:** Maximize cross-entropy loss w.r.t. inputs,

$$\max_{\substack{\delta \in \mathbb{R}^{1024 \times 3} \\ \|\delta\|_2 = 0.5}} -\log(p_{\text{horse}}(x + \delta))$$

Derivative-free Optimization on Riemannian Manifolds

Derivative-free optimization algorithms rely *only* on function values.

$$\min_{x \in \mathcal{M}} f$$

- No gradients. They must *limit* number of queries.
- Ex: Nelder-Mead, Evolutionary algorithms, Model-based.
- Direct-search methods: complexity in dimension squared.
- **Our work:** 1) Analyze a Riemannian direct-search method. 2) Apply it to adversarial attacks.

Riemannian direct search

Input: $x_0 \in \mathcal{M}$, $\alpha_0 > 0$, $c > 0$

for $k = 0, 1, \dots$ **do**

Generate polling set $D_k \subset T_{x_k} \mathcal{M}$ of tangent vectors (PSS)

If $f(\text{Exp}_{x_k}(\alpha_k d)) < f(x_k) - \frac{c}{2} \alpha_k^2 \|d\|_{x_k}^2$, for some $d \in D_k$,

$$x_{k+1} := \text{Exp}_{x_k}(\alpha_k d) \text{ and } \alpha_{k+1} := \min\{2\alpha_k, 1\}.$$

Else $x_{k+1} := x_k$ and $\alpha_{k+1} := \alpha_k/2$.

PSS, cosine measure

A *Positive Spanning Set* (PSS) of $T_x \mathcal{M}$ is a set $D_x \subset T_x \mathcal{M}$ s.t.

$$\text{cone}(D_x) = T_x \mathcal{M}.$$

The *cosine measure* of D_x is

$$\text{cm}_x(D_x) := \min_{v \in T_x \mathcal{M} \setminus \{0\}} \max_{d \in D_x} \frac{\langle v, d \rangle_x}{\|v\|_x \|d\|_x}$$


Complexity

Riemannian direct-search algorithms reach $\|\text{grad} f(x)\|_x \leq \epsilon$ in


$$O(\underbrace{\sup_x |D_x|}_{\text{size}} (\underbrace{\inf_x \text{cm}_x(D_x)}_{\text{cosine measure}})^{-2} \epsilon^{-2})$$

→ Trade-off between size and cosine measure of the Positive Spanning Set.

PSS, $\mathcal{M} = \mathbb{R}^m$




$$\text{cm}(D^{\oplus}) = \frac{1}{\sqrt{m}}$$



$$\text{cm}(D^{\odot}) = \frac{1}{m}$$

A high cosine measure means a well-spread PSS.

Polling strategies




Projected PSS

- ✓ Simple projection.
- ✗ Redundancy. Varying cosine measure.

Intrinsic PSS

- ✓ Consistent with $\mathbb{R}^m$. $O(m^2 \epsilon^{-2})$
- ✗ Requires orthogonal basis of $T_x \mathcal{M}$.


Numerical results



Data profiles on 400 quadratic problems constrained to linear subspace and sphere in a linear subspace of dimension m in $\mathbb{R}^n$.

- Projected (dashed) *vs* Intrinsic (plain) PSS.
- Left: $m = 4$. Codimension $n - m$ between 0 and 32. Right: converse.

When $n - m \gg m$, pick intrinsic PSSs!



Robust accuracy of a NN

We attack a WRN28-10 on 75 CIFAR10 images with 30720 queries.

- Perturbations: low-rank matrices $\delta = UV$, where $U \in \mathbb{R}^{32 \times r}$ and $V \in \mathbb{R}^{r \times 32}$, with $\|\delta\|_F = 0.5$.
- Attacks: Direct search methods in $\mathbb{R}^{32 \times r} \times \mathbb{R}^{r \times 32}$, and probabilistic descent variants (PD), with 2 or $2r$ random polling directions (DS-PD(2), or DS-PD($2r$)).

attack \ r	3	9	16	22	32
DS	0.033	0.033	0.022	0.011	0.011
DS-PD(2)	0.067	0.022	0.033	0.033	0.022
DS-PD($2r$)	0.056	0.044	0.011	0.022	0.033

References

- V. Kungurtsev, F. Rinaldi, D. Zeffiro *Retraction-Based Direct Search Methods for Derivative-Free Riemannian Optimization*, JOTA, 2024.
- BC. , F. Goyens, C. W. Royer, F. Yger, *Complexity guarantees and polling strategies for Riemannian direct-search methods*, arXiv:2511.15360, 2025.
- L. Roberts, C. W. Royer, *Direct Search Based on Probabilistic Descent in Reduced Spaces*, SIAM Journal on Optimization, 2023.
- D. Savostianova, E. Zangrando, F. Tudisco, *Low-Rank Adversarial PGD Attack*, arXiv:2410.12607, 2024.