Riemannian Optimization on Semialgebraic Sets

Matthias Himmelmann

Technische Universität Braunschweig
matthias.himmelmann@tu-braunschweig.de

Riemannian optimization uses local methods to solve optimization problems constrained to smooth manifolds. A linear step along some descent direction usually leaves the constraints, and hence retraction maps are applied to approximate the exponential map and return to the manifold. By developing algorithms based on homotopy continuation, we are able to compute the orthogonal projection onto the closest point on any submanifold of $\mathbb{R}^n$, also referred to as the Euclidean distance retraction.

We extend these tools to semialgebraic sets by adapting the predictor-corrector scheme central to homotopy continuation. Specifically, we replace Lagrange multipliers by Karush-Kuhn-Tucker conditions and introduce heuristics for escaping singularities and saddle points. Under the condition that each tangent descent step remains sufficiently close to the constraint set, we establish convergence guarantees for our method.

This work has applications in two main areas. First, the Euclidean distance retraction makes the effective approximation of physically meaningful deformation paths of geometric constraint systems possible and as such, it has been used to study mechanisms of geometric materials such as cylinder packings. Second, the generalization to semialgebraic sets lets us apply the Riemannian optimization framework to constrained polynomial optimization, providing a novel tool for finding upper bounds on the global minimum.