QCSC: Quadratic Cuts for Strongly Convex Optimization and its Stochastic Extension SQDP

Adriana Washington Henarejos*
School of Applied Mathematics (FGV EMAp)
E-mail: adriana.washington07@gmail.com

Vincent Guigues
School of Applied Mathematics (FGV EMAp)
E-mail: vincent.guigues@fgv.br

ABSTRACT

Convex optimization plays a central role in applied mathematics and computational sciences, providing the foundation for efficient algorithms in areas such as control, data science, and operations research. Within this field, methods that exploit the **structure of strong convexity**, that is, the presence of curvature in the objective function, can achieve faster convergence and improved numerical behavior. Motivated by this observation, this work introduces the **Quadratic Cuts for Strongly Convex Optimization (QCSC)** algorithm, a novel approach for solving deterministic strongly convex optimization problems through the construction of piecewise quadratic lower approximations of the objective function.

In the classical Kelley's method, each iteration builds a polyhedral approximation of the objective based on linear cuts of the form

$$\ell_{\tilde{f}}(x; x_k) = \tilde{f}(x_k) + \langle \tilde{f}'(x_k), x - x_k \rangle, \tag{1}$$

which underestimate any convex function. QCSC, however, incorporates curvature by adding a quadratic term, with the strong convexity ensuring that

$$q_{\tilde{f}}(x; x_k) = \tilde{f}(x_k) + \langle \tilde{f}'(x_k), x - x_k \rangle + \frac{\mu}{2}\|x - x_k\|^2 \leq \tilde{f}(x), \tag{2}$$

and producing a tighter and smoother model of the objective function and leading to faster convergence in practice.

We consider the strongly convex minimization problem

$$\varphi^* = \min_{x \in X} \tilde{f}(x), \tag{3}$$

where $\tilde{f} : \mathbb{R}^n \to \mathbb{R} \cup \{+\infty\}$ is a proper, lower semicontinuous, μ-strongly convex function, and $X \subset \mathbb{R}^n$ is a nonempty compact convex set.

At each iteration k, QCSC defines the model $\Gamma_k(x)$ as the maximum of all past quadratic cuts,

$$\Gamma_k(x) = \max_{j<k} q_{\tilde{f}}(x; x_j), \tag{4}$$

and determines the next iterate x_k as the minimizer of $\Gamma_k(x)$ over X. The best function value among the visited points defines y_k, and the optimality gap $t_k = \tilde{f}(y_k) - \Gamma_k(x_k)$ is monitored until it falls below a given tolerance $\bar{\varepsilon}$.

*PhD student in Applied Mathematics and Data Science

A theoretical analysis establishes the validity of the quadratic cuts and provides convergence and complexity guarantees, showing similarities with proximal bundle-type methods while keeping a simpler structure. A composite reformulation further clarifies the role of the global quadratic term. Also, numerical experiments on strongly convex problems indicate promising performance, confirming the practical viability of the deterministic approach and motivating its extension to stochastic settings.

Building upon this deterministic foundation, a stochastic generalization called **SQDP (Stochastic Quadratic Dynamic Programming)** is proposed to solve multistage stochastic optimization problems of the form

$$\begin{aligned} &\inf_{x_1,\ldots,x_T} \mathbb{E}_{\xi_2,\ldots,\xi_T}\left[\sum_{t=1}^{T} f_t(x_t(\xi_1,\xi_2,\ldots,\xi_t), x_{t-1}(\xi_1,\xi_2,\ldots,\xi_{t-1}),\xi_t)\right] \\ &x_t(\xi_1,\xi_2,\ldots,\xi_t) \in X_t(x_{t-1}(\xi_1,\xi_2,\ldots,\xi_{t-1}),\xi_t) \text{ a.s.}, t=1,\ldots,T, \end{aligned} \tag{5}$$

where x_0 is given, ξ_1 is deterministic, $(\xi_t)_{t=2}^T$ is a stochastic process, $X_t\left(x_{t-1},\xi_t\right) = \{x_t \in \mathbb{R}^n : x_t \in \mathcal{X}_t, g_t\left(x_t, x_{t-1},\xi_t\right) \leq 0, A_t x_t + B_t x_{t-1} = b_t\}, t = 1,\ldots,T$, and the function $f_t\left(\cdot,\cdot,\xi_{tj}\right)$ is strongly convex on $\mathcal{X}_t \times \mathcal{X}_{t-1}$. Classical SDDP (Stochastic Dual Dynamic Programming) approximates these recourse functions using affine cuts, but such cuts may poorly capture curvature in high-curvature or ill-conditioned problems. SQDP overcomes this limitation by constructing **piecewise quadratic lower approximations** of the recourse functions, obtained by adding a stage-dependent quadratic term to each backward-pass cut.

This modification preserves the stage-wise decomposition structure of SDDP while exploiting strong convexity to generate tighter value function approximations. As a result, each backward pass produces more informative cuts, and the forward simulation benefits from more accurate decision rules. From an algorithmic viewpoint, SQDP follows the same forward–backward scheme as SDDP: the forward pass samples scenarios and produces trial points, and the backward pass updates the quadratic cut models for each stage. The main difference basically lies entirely in the construction and accumulation of quadratic cuts.

Numerical experiments demonstrate substantial performance gains over SDDP. SQDP converges faster, stabilizes its value function approximations in fewer iterations, and performs particularly well when the curvature of the recourse functions is large. These results highlight that incorporating curvature through quadratic cuts significantly improves the classical dynamic programming approximation, reinforcing the potential of quadratic-cut methods for large-scale stochastic optimization.

Keywords: *Optimization, Strong Convexity, Cutting-plane methods, Stochastic programming*

Referências

[1] V. Guigues, A. Washington, *Stochastic Quadratic Dynamic Programming*, arXiv:2506.07314, 2025.