Approximate Secular Equations for the Cubic Regularization Subproblem

Yihang Gao, Man-Chung Yue and Micahel K. Ng

The University of Hong Kong

November 8, 2022

Overview

Problem

$$\min_{\boldsymbol{x}\in\mathbb{R}^n} f(\boldsymbol{x}),$$

where $f(\boldsymbol{x})$ is the non-convex objective function.

- ϵ-approximate stationary point.

$$\|\nabla f(\boldsymbol{x})\|_2 \leq \epsilon$$

  and $\boldsymbol{x}$ satisfies the second-order necessary condition (e.g., $\lambda_{\min}(\nabla^2 f(\boldsymbol{x})) \geq -\sqrt{\epsilon}$).
- Second-order methods: Cubic Regularization (CR) [Nesterov et al., 2006] and Trust Region (TR) [Conn et al., 2000] methods, etc.
- Local convergence properties (e.g., superlinear, linear, and sublinear convergence) under mild assumptions [Yue et al., 2019].

Cubic Regularization

Each iteration of CR and its variants involve solving the following form, called cubic regularization subproblem (CRS):

$$\min_{\mathbf{x}\in\mathbb{R}^n} f_{\mathbf{A},\mathbf{b},\rho}(\mathbf{x}) := \mathbf{b}^{\mathrm{T}}\mathbf{x} + \frac{1}{2}\mathbf{x}^{\mathrm{T}}\mathbf{A}\mathbf{x} + \frac{\rho}{3}\|\mathbf{x}\|^3, \tag{1}$$

- In CR, we have the following update scheme:

$$\mathbf{x}_{k+1} \in \arg\min_{\mathbf{x}} \; f(\mathbf{x}_k) + \mathbf{g}_k^{\mathrm{T}}(\mathbf{x}-\mathbf{x}_k) + \frac{1}{2}(\mathbf{x}-\mathbf{x}_k)^{\mathrm{T}}\mathbf{H}_k(\mathbf{x}-\mathbf{x}_k) + \frac{\rho}{3}\|\mathbf{x}-\mathbf{x}_k\|^3,$$

  where $\mathbf{g}_k = \nabla f(\mathbf{x}_k)$ and $\mathbf{H}_k = \nabla^2 f(\mathbf{x}_k)$.
- **A** is not necessarily positive definite.

Secular Equation

- We denote by $\lambda_1 \leq \cdots \leq \lambda_n$ the eigenvalues of $\mathbf{A}$ and by $\mathbf{v}_1, \cdots, \mathbf{v}_n$ the corresponding eigenvectors. In other words, we have the eigendecomposition $\mathbf{A} = \sum_{i=1}^n \lambda_i \mathbf{v}_i \mathbf{v}_i^T = \mathbf{V}\mathbf{\Lambda}\mathbf{V}^{\mathrm{T}}$, where $\mathbf{\Lambda} = \mathrm{diag}(\lambda_1, \ldots, \lambda_n)$ and $\mathbf{V} = [\mathbf{v}_1, \cdots, \mathbf{v}_n]$.

Proposition ([Nesterov et al., 2006])

A vector $\mathbf{x}^*$ *solves the CRS* (1) *if and only if it satisfies the system*

$$\begin{cases} (\mathbf{A} + \rho\|\mathbf{x}^*\|\mathbf{I})\mathbf{x}^* + \mathbf{b} = \mathbf{0}, & (2) \\ \mathbf{A} + \rho\|\mathbf{x}^*\|\mathbf{I} \succeq \mathbf{0}. & (3) \end{cases}$$

Moreover, if $\mathbf{A} + \rho\|\mathbf{x}^*\|\mathbf{I} \succ \mathbf{0}$*, then* $\mathbf{x}^*$ *is the unique solution (and hence a critical point).*

- If $\mathbf{b}^{\mathrm{T}}\mathbf{v}_1 \neq 0$, then $\mathbf{A} + \rho\|\mathbf{x}^*\|\mathbf{I} \succ \mathbf{0}$ and the solution $\mathbf{x}^*$ is the unique (second-order) critical point (and hence the unique solution).

Secular Equation

- Conditions (2) and (3) can be written as
$$\begin{cases} (\Lambda + \sigma \mathbf{I}) \cdot \mathbf{y}^* = \mathbf{c}, \\ \lambda_1 + \sigma > 0. \end{cases}$$
where $\sigma =: \rho \|\mathbf{x}^*\|$, $[y_1^*, \cdots, y_n^*]^{\mathrm{T}} := \mathbf{y}^* = \mathbf{V}^{\mathrm{T}} \mathbf{x}^*$ and $[c_1, \cdots, c_n]^{\mathrm{T}} := \mathbf{c} = -\mathbf{V}^{\mathrm{T}} \mathbf{b}$.
- Since the Euclidean norm is invariant to orthogonal transformation, we have
$$\frac{\sigma^2}{\rho^2} = \|\mathbf{x}^*\|^2 = \|\mathbf{y}^*\|^2 = \sum_{i=1}^{n} \frac{c_i^2}{(\lambda_i + \sigma)^2}.$$
- We first find the (unique) root $\sigma > \max\{-\lambda_1, 0\}$ of the equation
$$w(\sigma) = \sum_{i=1}^{n} \frac{c_i^2}{(\lambda_i + \sigma)^2} - \frac{\sigma^2}{\rho^2}, \tag{4}$$
called the secular equation, and then solves the linear system $(\mathbf{A} + \sigma \mathbf{I})\mathbf{x} = -\mathbf{b}$.

The First-Order Truncated Secular Equation

- We define the first-order truncated secular equation by

$$w_1(\sigma;\mu) = \sum_{i=1}^{m} \frac{c_i^2}{(\lambda_i+\sigma)^2} + \sum_{i=m+1}^{n} \frac{c_i^2}{(\mu+\sigma)^2} - \frac{\sigma^2}{\rho^2}, \tag{5}$$

where $\mu \geq \lambda_m$.

Lemma

For any $\mu \geq \lambda_m$, the function $w_1(\cdot;\mu)$ as defined in (5) admits a unique root.

The First-Order Truncated Secular Equation

Theorem

Let σ_1^* and σ^* be the unique roots of $w_1(\sigma;\mu)$ and $w(\sigma)$, respectively. Then

$$|\sigma_1^* - \sigma^*| \leq C_m \cdot \max_{m+1 \leq i \leq n} |\lambda_i - \mu|, \tag{6}$$

where $C_m > 0$ is a constant, upper bounded by $\frac{2\|\mathbf{b}\|^2}{(\lambda_m - \lambda_1)^3} \cdot \min\left\{\frac{(\lambda_n + B_1)^3}{2\|\mathbf{b}\|^2}, \frac{\rho^2}{2B_1}\right\}$ with $B_1 = \frac{-\lambda_1 + \sqrt{\lambda_1^2 + 4\rho \cdot \|\mathbf{b}\|}}{2}$ being an upper bound for $|\sigma_1^*|$.

Proposition

Let $\mathbf{x}^*$ and $\tilde{\mathbf{x}}$ be solutions to the equations $(\mathbf{A} + \sigma^* \mathbf{I})\,\mathbf{x}^* = -\mathbf{b}$ and $(\mathbf{A} + \sigma_1^* \mathbf{I})\,\tilde{\mathbf{x}} = -\mathbf{b}$, respectively. Then, $\|\tilde{\mathbf{x}} - \mathbf{x}^*\| = \mathcal{O}\left(|\sigma_1^* - \sigma^*|\right)$.

The First-Order Truncated Secular Equation

- An intuitive choice of μ that works well in practice and is computationally cheap is the average of unknown eigenvalues, i.e.,

$$\mu_1 = \frac{\sum_{i=m+1}^{n} \lambda_i}{n-m} = \frac{\text{tr}(\mathbf{A}) - \sum_{i=1}^{m} \lambda_i}{n-m}. \tag{7}$$

- An example on Random Gaussian matrices.
Suppose that $\mathbf{A} = \widetilde{\mathbf{A}}/\sqrt{2n}$, where $\widetilde{\mathbf{A}}$ is a symmetric random matrix with i.i.d. entries on and above the diagonal. By the Wigner semicircle law, as $n \to \infty$, the eigenvalues of $\mathbf{A}$ distribute according to a density of a semi-circle shape. In particular, we can deduce that with a probability of $1 - o(1)$,

$$\max_{m+1 \le i \le n} |\lambda_i - \mu| \le \mathcal{O}\left(\left(1 - \frac{m+1}{n}\right)^{2/3}\right) \approx \left(\frac{3\pi}{4\sqrt{2}}\right)^{2/3} \cdot \left(1 - \frac{m+1}{n}\right)^{2/3} \tag{8}$$

The Second-Order Truncated Secular Equation

- With the second-order Taylor approximation, we define the second-order truncated secular equation by

$$w_2(\sigma;\mu) = \sum_{i=1}^{m} \frac{c_i^2}{(\lambda_i+\sigma)^2} + \sum_{i=m+1}^{n} \frac{c_i^2}{(\mu+\sigma)^2} - 2\sum_{i=m+1}^{n} \frac{c_i^2\cdot(\lambda_i-\mu)}{(\mu+\sigma)^3} - \frac{\sigma^2}{\rho^2}, \quad (9)$$

  where $\mu \geq \lambda_m$.

Lemma

With

$$\mu = \frac{\sum_{i=m+1}^{n} c_i^2\cdot\lambda_i}{\sum_{i=m+1}^{n} c_i^2}, \quad (10)$$

the function $w_2(\cdot;\mu)$ as defined in (9) admits a unique root.

The Second-Order Truncated Secular Equation

Theorem

Let σ_2^* *and* σ^* *be the unique root of* $w_2(\sigma;\mu)$ *and* $w(\sigma)$*, respectively, and*

$$\mu = \frac{\sum_{i=m+1}^{n} c_i^2 \cdot \lambda_i}{\sum_{i=m+1}^{n} c_i^2}.$$

Then,

$$|\sigma_2^* - \sigma^*| \leq C_m \cdot \max_{m+1 \leq i \leq n} (\lambda_i - \mu)^2, \tag{11}$$

where $C_m > 0$ *is a constant bounded by* $\frac{3\|\mathbf{b}\|^2}{(\lambda_m - \lambda_1)^4} \cdot \min\left\{\frac{(\lambda_n + B_1)^3}{2\|\mathbf{b}\|^2}, \frac{\rho^2}{2B_1}\right\}$ *with* $B_1 = \frac{-\lambda_1 + \sqrt{\lambda_1^2 + 4\rho \cdot \|\mathbf{b}\|}}{2}$ *being an upper bound for* $|\sigma_2^*|$.

Implementation Details

- The resulting CRS solver, namely the approximate secular equation method (ASEM), is summarized as follows:
  Step 1: obtaining the partial eigen information $\{\lambda_1, \cdots, \lambda_m\}$ and $\{\mathbf{v}_1, \cdots, \mathbf{v}_m\}$ of $\mathbf{A}$.
  Step 2: solving the secular equation (5) with μ defined in (7) or (10); we get σ^*.
  Step 3: iteratively solving the linear system $(\mathbf{A} + \sigma^* \mathbf{I})\mathbf{x} + \mathbf{b} = \mathbf{0}$.
  Output: the solution $\mathbf{x}$.
- Krylov subspace (Lanczos) method for partial eigen information. Matlab (eigs function) and Python (Scipy package) etc.
- Bisection method for finding roots of approximate secular equations $w_1(\cdot, \mu)$ and $w_2(\cdot, \mu)$.
- Krylov subpace (Lanczos) method for solving the linear system $(\mathbf{A} + \sigma \mathbf{I})\mathbf{x} = -\mathbf{b}$.

Synthetic Examples

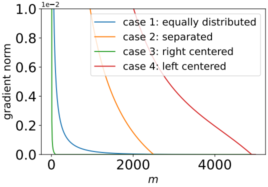

Figure: Trajectories of suboptimality (gradient norm $\|\nabla f_{\mathbf{A},\mathbf{b},\rho}(\mathbf{x})\|$) with different distributions for eigenvalues in Experiment 1.

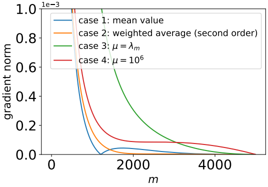

Figure: Trajectories of suboptimality (gradient norm $\|\nabla f_{\mathbf{A},\mathbf{b},\rho}(\mathbf{x})\|$) with different μ in Experiment 2.

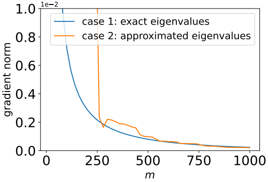

Figure: Trajectories of suboptimality (gradient norm $\|\nabla f_{\mathbf{A},\mathbf{b},\rho}(\mathbf{x})\|$) with exact and approximated eigenvalues and eigenvectors in Experiment 3.

CUTEst Examples

Problem	Method	$f(\mathbf{x}_{\text{out}})$	$\|\nabla f(\mathbf{x}_{\text{out}})\|$	$\lambda_1(\nabla^2 f(\mathbf{x}_{\text{out}}))$	iter	time(s)
TOINTGSS ($n = 1000$)	ARC-CP	3.60E+14	4.12E-05	1.40E-16	1000	6.08
	ARC-GD	3.60E+14	1.42E-06	3.89E-16	100	6.98
	ARC-Krylov(1)	3.60E+14	4.12E-05	1.20E-15	300	6.75
	ARC-Krylov(10)	3.60E+14	2.20E-08	1.29E-15	19	**1.87**
	ARC-ASEM(1)	3.60E+14	**8.01E-10**	-1.63E-15	19	2.17
	ARC-ASEM(10)	3.60E+14	8.17E-10	-7.67E-16	19	2.74
BRYBAND ($n = 2000$)	ARC-CP	7.49E+14	1.10E-03	5.40E+00	1000	8.27
	ARC-GD	**1.25E+05**	4.93E+03	4.40E+02	100	11.05
	ARC-Krylov(10)	7.49E+14	6.60E-06	5.40E+00	100	9.85
	ARC-Krylov(30)	7.49E+14	1.14E-07	5.40E+00	14	2.37
	ARC-ASEM(1)	7.49E+14	1.02E-07	5.40E+00	14	**2.24**
	ARC-ASEM(10)	7.49E+14	**1.01E-07**	5.40E+00	14	3.83
DIXMAANG ($n = 3000$)	ARC-CP	1.00E+00	3.13E-04	6.67E-04	2000	35.16
	ARC-GD	1.00E+00	9.24E-05	6.67E-04	200	33.83
	ARC-Krylov(10)	1.00E+00	3.44E-05	6.67E-04	500	32.18
	ARC-Krylov(30)	1.00E+00	9.06E-09	6.67E-04	46	**6.65**
	ARC-ASEM(1)	1.00E+00	5.53E-09	6.67E-04	**30**	7.51
	ARC-ASEM(10)	1.00E+00	**4.85E-09**	6.67E-04	42	18.74
TQUARTIC ($n = 5000$)	ARC-CP	8.04E-01	6.10E-02	-5.41E-05	500	71.92
	ARC-GD	8.05E-01	2.77E-02	-4.80E-05	100	98.14
	ARC-Krylov(1)	8.05E-01	2.76E-02	-4.71E-05	100	29.43
	ARC-Krylov(10)	**5.05E-14**	**8.48E-09**	4.00E-04	46	16.09
	ARC-ASEM(1)	7.43E-14	9.62E-09	4.00E-04	46	**15.47**
	ARC-ASEM(10)	7.43E-14	9.62E-09	4.00E-04	46	16.18

Figure: Results on CUTEst problems in Experiment 5 (ARC [Cartis et al., 2011]).

References I

Yurii Nesterov, Boris T Polyak (2006)
Cubic regularization of Newton method and its global performance,
Mathematical Programming.

Andrew Conn, Nicholas IM Gould, Philippe L Toint (2000)
Trust region methods,
SIAM.

Man-Chung Yue, Zirui Zhou, Man-Cho So (2019)
On the quadratic convergence of the cubic regularization method under a local error bound condition,
SIAM Journal on Optimization.

Yair Carmon, John Duchi (2019)
Gradient descent finds the cubic-regularized nonconvex Newton step,
SIAM Journal on Optimization.

Coralia Cartis, Nicholas IM Gould, Philippe L Toint (2011)
Adaptive cubic regularisation methods for unconstrained optimization. Part I: motivation, convergence and numerical results,
Mathematical programming.

The End

Thanks!