

Higher-Order Newton Methods with Polynomial Work per Iteration

Amir Ali Ahmadi*, Abraar Chaudhry*, Jeffrey Zhang†

Abstract

We present generalizations of Newton’s method that incorporate derivatives of an arbitrary order d but maintain a polynomial dependence on dimension in their cost per iteration. At each step, our d^{th} -order method uses semidefinite programming to construct and minimize a sum of squares-convex approximation to the d^{th} -order Taylor expansion of the function we wish to minimize. We prove that our d^{th} -order method has local convergence of order d . This results in lower oracle complexity compared to the classical Newton method. We show on numerical examples that basins of attraction around local minima can get larger as d increases. Under additional assumptions, we present a modified algorithm, again with polynomial cost per iteration, which is globally convergent and has local convergence of order d .

Keywords. Newton’s method, tensor methods, semidefinite programming, sum of squares methods, convergence analysis.

1 Introduction

Newton’s method is perhaps one of the most well-known and prominent algorithms in optimization. In its attempt to minimize a function $f : \mathbb{R}^n \rightarrow \mathbb{R}$, this algorithm replaces f with its second-order Taylor expansion at an iterate $x_k \in \mathbb{R}^n$ and defines the next iterate x_{k+1} to be a critical point of this quadratic approximation. This critical point coincides with a minimizer of the quadratic approximation in the case where the Hessian of f at x_k is positive semidefinite.

The work required in each iteration of Newton’s method consists of solving a system of linear equations which arises from setting the gradient of the quadratic approximation to zero. This can be carried out in time that grows polynomially with the dimension n . Perhaps the most well-known theorem about the performance of Newton’s method is its *local quadratic convergence*. More precisely, under the assumptions that the second derivative of f is locally Lipschitz around a local minimizer x^* , and that the Hessian at x^* is positive definite, there exists

*Princeton University, Operations Research and Financial Engineering. AAA and AC were partially supported by the MURI award of the AFOSR and the Sloan Fellowship.

†Yale University, Department of Biomedical Informatics and Data Science.

a full-dimensional basin around x^* and a constant c , such that if x_0 is in this basin, one has

$$\|x_{k+1} - x^*\| \leq c\|x_k - x^*\|^2$$

for all $k \geq 0$. We note however that Newton's method is in general not globally convergent. Lack of global convergence can occur even when in addition to the previous assumptions, f is assumed to be strongly convex (see, e.g., Example 5.1 in Section 5).

As higher-order Taylor expansions provide closer local approximations to the function f , it is natural to ask why Newton's method limits the order of Taylor approximation to 2. The main barrier to higher-order Newton methods is the computational burden associated with minimizing polynomials of degree larger than 2 which would arise from higher-order Taylor expansions. For instance, any of the following tasks that one could consider for each iteration of a higher-order Newton method are in general NP-hard:

- (i) finding a global minimum of polynomials of degree even¹ and at least 4 (see, e.g., [39]),
- (ii) finding a local minimum of polynomials of degree at least 4 (see [8, Theorem 2.1]),
- (iii) finding a second-order point (i.e., a point where the gradient vanishes and the Hessian is positive semidefinite) of polynomials of degree at least 4 (see [7, Theorem 2.2]),
- (iv) finding a critical point (i.e., a point where the gradient vanishes) of polynomials of degree at least 3 (see [7, Theorem 2.1]).

In addition to matters related to computation, there are geometric distinctions between Newton's method and higher-order analogues of it. For example, even when the function f is strongly convex and the starting iterate is arbitrarily close to its minimizer, Taylor expansions of even degree and larger than 2 may not be bounded below. One can see this by examining the strongly convex univariate function $f(x) = x^2 - x^4 + x^6$ and its 4th order Taylor expansion near the origin.

Despite these barriers, the question of whether one can make higher-order Newton methods tractable and in some way superior to Newton's method has been considered at least since the work of Chebyshev [20] (see Section 1.1 for more recent literature). More specifically, the question that is of interest to us is whether it is possible to design a higher-order Newton method (i.e., a method which utilizes a Taylor expansion of degree $d > 2$ in each iteration) in such a way that (i) the work per iteration grows polynomially with the dimension, and (ii) the local order of convergence grows with d , hence requiring fewer function evaluations as d increases. In this paper, we show that this is indeed possible (Algorithm 1 and Theorem 4).

¹Note that odd-degree polynomials are unbounded below.

Our algorithm relies on sum of squares techniques in optimization [44], [30] and semidefinite programming and does not require the function f to be convex. For any fixed degree d , our approach is to approximate the d -th order Taylor expansion of f with an “sos-convex” polynomial (see Section 2 for a definition). Sos-convex polynomials form a subclass of convex polynomials whose convexity has an explicit algebraic proof. One can then use a first-order sum of squares relaxation to minimize this sos-convex polynomial. It turns out that both the task of finding a suitable sos-convex polynomial and that of minimizing it can be carried out by solving two semidefinite programs whose sizes are polynomial in the dimension n (in fact of the same order as the number of terms in the Taylor expansion). As is well known, semidefinite programs can be solved to arbitrary accuracy in polynomial time; see [48] and references therein.

We work with sos-convex polynomials instead of general convex polynomials since the latter set lacks a tractable description [5], and the former, as we show, turns out to be sufficient for achieving an algorithm with superlinear local convergence. Our sum of squares based algorithm works for higher-order Newton methods of any order d and can be easily implemented using any sum of squares parser (e.g., YALMIP [35] or SOSTOOLS [45]). This is in contrast to previous work where implementable algorithms have been worked out only for $d = 3$; see [40], [23, Sect. 1.5], [25, Sect. 5]. While we present our algorithms in the unconstrained case, they can be readily implemented in the presence of sos-convex constraints (such as linear constraints or convex quadratic constraints). We note, however, that our interest in this paper is only on generalizing Newton’s method in terms of its convergence order and polynomial work per iteration, and not on the practical aspects of implementation. Designing more scalable algorithms for semidefinite programs is an active area of research [36, 49]. In addition, we believe that there are promising future research directions which could make our algorithms more practical at larger scale (see Section 7).

1.1 Related Work

Over the years, there have been many adaptations of and extensions to Newton’s method. A primary example is the pioneering work of Nesterov and Polyak [41], where the idea of Newton’s method with cubic regularization was introduced. We do not review the large literature that emerged from this work since the order d of Taylor expansion in this line of work is still equal to 2, and hence these methods are not considered “higher-order” (i.e., $d > 2$). However, the framework that we propose, similar to most of the literature, follows the structure of [41] (and [33, 37]) in terms of minimizing, in each iteration, a Taylor expansion of a certain order plus an appropriate regularization term. Recently, there has been a body of work following this structure with Taylor expansions of order higher than two [40, 9, 12, 28, 29, 25]. Unlike our paper, these works are in the setting of convex optimization, do not study the complexity of minimizing the regularized Taylor expansion in each iteration (except in the case of $d = 3$ for a subset of these papers), and derive sublinear rates of global convergence. There has also been work on lower bounds on the rates of convergence

for such methods [11, 1, 13, 40]. These lower bounds are nearly achieved by the algorithms in the aforementioned papers. The recent textbook [17] provides an accessible summary of this literature and its broader scope. See also [14, 16, 15] and references therein.

In terms of work per iteration of higher-order Newton methods, Nesterov presents a polynomial-time algorithm in [40] for minimizing a quartically-regularized third-order Taylor expansion. This problem is revisited recently in [18], where an algorithm for recovering an approximate second-order point for a possibly non-convex quartically-regularized third-order Taylor expansion is presented. In [47], a different third-order Newton method is presented which has polynomial work per iteration. In each iteration, this algorithm moves to a local minimum of the third-order Taylor expansion. It turns out that local minima of cubic polynomials can be found by semidefinite programs of polynomial size [7]. To the best of our knowledge, no efficient algorithm for higher-order Newton methods of degree $d > 3$ has been presented. In fact, designing such an algorithm is referred to as an open problem in [23, Sec. 1.5] and [25, Sec. 5]. Interestingly, Nesterov asks in [40, Sec. 6] whether it is possible to tackle this problem using “some tools from algebraic geometry and the related technique of sums of squares”. This is precisely the approach that we take in this paper.

To our knowledge, the only works that establish superlinear rates of local convergence for higher-order Newton methods are [47] and [24] (and the related PhD thesis [23]), the latter of which came to our attention at the time of writing this paper. In [47], the authors establish third-order local convergence rate for an unregularized third-order Newton method applied to a strongly convex function. In [24], the authors establish superlinear local convergence for higher-order Newton methods applied to convex optimization problems with composite objective. When the smooth part of the objective function is strongly convex, the authors show local convergence of order d in function value and norm of the subgradient for their proposed d^{th} -order Newton method. An algorithm carrying out the work per iteration of this method, however, is available only in the case of $d = 3$ (and is the same as that in [40]). Moreover, similar to much of the literature, the regularization term that is added to the Taylor expansion in this method requires knowledge of the Lipschitz constant of the d^{th} derivative of f . Our proof technique for local superlinear convergence is different than [24] both in the parts where the sum of squares programming aspects come in and in the parts that they do not. Furthermore, our method has polynomial work per iteration for *any* degree d . It also does not rely on knowledge of any Lipschitz constants. Our regularization term is instead derived from the optimal value of a semidefinite program which can be written down from the coefficients of the Taylor expansion alone. This optimized approach can potentially lead to smaller deviations from the Taylor expansion and therefore an improved convergence factor. Finally, we note that in our work, assumptions on convexity of f and knowledge of the Lipschitz constant of its d^{th} derivative are made only in Section 6, where global convergence is established. Our approach in Section 6 is based on incorporating sum of squares methods into the framework of Nesterov in [40], though in theory this can also be done with other globally

convergent higher-order Newton methods. In fact, at the time of revision, there has already been interesting follow-up work to our paper which combines our sum of squares framework with adaptive regularization techniques for tensor methods and analyzes the complexity of the resulting algorithm for finding an approximate stationary point of a nonconvex function [19].

1.2 Organization and Contributions

In Section 2, we review preliminaries on sos-convexity, sos-convex polynomial optimization, and error rates of derivatives of Taylor expansions. In Section 3, we present our main algorithm (Algorithm 1). In Section 4, we prove that our algorithm is well-defined in the sense that the semidefinite programs it executes are always feasible and that the next iterate is always uniquely defined (Theorem 3). We then prove that our semidefinite programming-based d^{th} -order Newton scheme has local convergence of order d (Theorem 4). Compared to the classical Newton method, this leads to fewer calls to the Taylor expansion oracle (a common oracle in this literature; see e.g., [12], [29, Sect. 2.2], [1, Sect. 1.1], [11, Sect. 2], [17, Chap. 1.2]) at the price of requiring higher-order derivatives. The proof of Theorem 4 is more involved than the proof of local quadratic convergence of Newton’s method. This is in part because the expression for the next iterate of Newton’s method is explicit, whereas our next iterate comes from the solution to two semidefinite programs. We also remark that our proof framework is applicable to a broader class of higher-order Newton methods that may not necessarily use sum of squares techniques.

In Section 5, we present three numerical examples. We give an explicit expression and a geometric interpretation of our third-order Newton method in dimension one. We compare the basins of attraction of local minima for our higher-order methods to those of the classical Newton method. In Section 6, we present a slightly modified higher-order Newton method which is globally convergent under additional convexity and Lipschitzness assumptions similar to those in [40]. This modified algorithm works in the case of d being an odd integer and still has polynomial work per iteration and local convergence of order d . Finally, in Section 7, we present a few directions for future research.

2 Preliminaries

2.1 SOS-Convex Polynomial Optimization

In each iteration of the higher-order Newton methods that we propose, two semidefinite programs (SDPs) need to be solved. These SDPs arise from the notion of sos-convexity, which is reviewed in this subsection.

Definition 1. *A polynomial $p : \mathbb{R}^n \mapsto \mathbb{R}$ is said to be a sum of squares (sos) if there exist polynomials $q_1, \dots, q_r : \mathbb{R}^n \mapsto \mathbb{R}$ such that $p = \sum_{i=1}^r q_i^2$.*

As is well known, one can check if a polynomial is sos by solving an SDP. The next theorem establishes this link. We denote that a symmetric matrix

A is positive semidefinite (i.e., has nonnegative eigenvalues) with the standard notation $A \succeq 0$.

Theorem 1 (see, e.g., [44]). *For a variable $x \in \mathbb{R}^n$ and an even integer d , let $\phi_{\frac{d}{2}}(x)$ denote the vector of all monomials of degree at most $\frac{d}{2}$ in x . A polynomial $p : \mathbb{R}^n \mapsto \mathbb{R}$ of degree d is sos if and only if there exists a symmetric matrix Q such that (i) $p(x) = \phi_{\frac{d}{2}}(x)^T Q \phi_{\frac{d}{2}}(x)$ for all $x \in \mathbb{R}^n$, and, (ii) $Q \succeq 0$.*

The first constraint above can be written as a finite number of linear equations by coefficient matching. Therefore, the two constraints together represent the intersection of an affine subspace with the cone of positive semidefinite matrices. Thus, as polynomials can be encoded as an ordered vector of coefficients, the set of sos polynomials of a given degree has a description as the feasible region of a semidefinite program. Furthermore, the size of this SDP grows polynomially in n when d is fixed.

Throughout this paper, we denote the gradient vector (resp. Hessian matrix) of a function $g : \mathbb{R}^n \mapsto \mathbb{R}$ with the standard notation ∇g (resp. $\nabla^2 g$).

Definition 2 (SOS-Convex). *A polynomial $p : \mathbb{R}^n \mapsto \mathbb{R}$ is said to be sos-convex if the polynomial $q : \mathbb{R}^n \times \mathbb{R}^n \mapsto \mathbb{R}$ defined as $q(x, y) := y^T \nabla^2 p(x) y$ is sos.*

Note that any sos-convex polynomial is convex. The converse statement is not true, except for certain dimensions and degrees (see [6]). By Theorem 1 above, the set of sos-convex polynomials of a given degree also form the feasible region of a semidefinite program. Because the polynomial $q(x, y)$ is quadratic in y , one can reduce the size of the underlying SDP. More specifically, a polynomial $p : \mathbb{R}^n \mapsto \mathbb{R}$ of degree d is sos-convex² if and only if there exists a symmetric matrix $Q \succeq 0$ such that $y^T \nabla^2 p(x) y = (\phi_{\frac{d}{2}-1}(x) \otimes y)^T Q (\phi_{\frac{d}{2}-1}(x) \otimes y)$. (Here, \otimes denotes the Kronecker product.) We see that the size of the SDP that represents sos-convex polynomials of degree d in n variables grows polynomially in n when d is fixed.

We next explain why sos-convex polynomial optimization problems can be solved with the first level of the so-called Lasserre hierarchy. A polynomial optimization problem is a problem of the form

$$\begin{aligned} \inf_{x \in \mathbb{R}^n} \quad & g_0(x) \\ \text{s.t.} \quad & g_j(x) \leq 0 \quad j = 1, \dots, m, \end{aligned} \tag{1}$$

where $g_j(x)$ are real-valued polynomial functions of a variable $x \in \mathbb{R}^n$. The first-level Lasserre relaxation (see [30]) corresponding to problem (1) takes the form

$$\begin{aligned} \sup_{\gamma \in \mathbb{R}, \lambda \in \mathbb{R}^m} \quad & \gamma \\ \text{s.t.} \quad & g_0(x) - \gamma + \sum_{j=1}^m \lambda_j g_j(x) \text{ is sos} \\ & \lambda_j \geq 0 \quad j = 1, \dots, m. \end{aligned} \tag{2}$$

²Note that an odd-degree polynomial can never be convex, except for the trivial case of affine polynomials.

The reader can check that the optimal value of (2) is always a lower bound on that of (1). The next theorem establishes that this lower-bound is tight when the defining polynomials of (1) are sos-convex.

Theorem 2 (See Corollary 2.5 from [31], and Theorem 3.3 from [32]). *Suppose that the polynomials g_0, \dots, g_m in (1) are sos-convex, the optimal value of (1) is finite, and that the Slater condition holds³. Then, the optimal values of (1) and (2) are the same. Moreover, an optimal solution to (1) can be readily recovered from a solution to the semidefinite program that is dual to (2).*

This result is already proven by Lasserre in [32] using a lemma of Helton and Nie from [26]. For completeness and for the benefit of the reader, we give an alternative short proof of the first claim.

Proof. Recalling that an sos polynomial is nonnegative and that $\lambda_j \geq 0$ for $j = 1, \dots, m$, it is easy to see that the optimal value of (1) is larger than or equal to the optimal value of (2). To show the opposite inequality, let γ^* be the optimal value of (1). Then, the convex function $x \mapsto g_0(x) - \gamma^*$ is nonnegative over the set $\{x \mid g_j(x) \leq 0, j = 1, \dots, m\}$. By the convex Farkas lemma (see, e.g., [46, Theorem 2.1]), there exists a nonnegative vector $\lambda^* \in \mathbb{R}^m$ such that $p(x) := g_0(x) - \gamma^* + \sum_{j=1}^m \lambda_j^* g_j(x) \geq 0$ for all $x \in \mathbb{R}^n$. Notice that $p(x)$ is sos-convex since it is a conic combination of sos-convex polynomials. Thus, by [6, Theorem 3.1], the polynomial $q(x, y) := p(y) - p(x) - \nabla p(x)^T (y - x)$ is sos. Let x^* be an optimal solution to (1) (such a vector must exist [10]). Observe that the polynomial $y \mapsto q(x^*, y)$ is also sos (since it is the restriction of $q(x, y)$ to $x = x^*$). By optimality of x^* to (1), we have $p(x^*) \leq 0$. Since p is nonnegative, we have $p(x^*) = 0$ and $\nabla p(x^*) = 0$. Thus, $p(y) = q(x^*, y)$, and hence $p(y)$ must be sos. Therefore, γ^*, λ^* is feasible to (2), and hence the optimal value of (2) is at least γ^* ; i.e., the optimal value of (1). □

For a proof of the second claim and an explicit expression of the dual of (2), see Theorem 3.3 from [32].

2.2 Error rates of Taylor remainders

In this subsection, we review certain error rates of multivariate Taylor expansions that will be used in our arguments. We denote by $\nabla^d f$ the d^{th} order symmetric tensor of order- d partial derivatives of the function f . We denote the tensor product of a set of vectors $x_1, \dots, x_d \in \mathbb{R}^n$ with $x_1 \boxtimes x_2 \boxtimes \dots \boxtimes x_d$.⁴ We use the notation $x^{\boxtimes d}$ to denote the tensor product of a vector $x \in \mathbb{R}^n$ with itself d times. With this notation, we can define the d^{th} -order Taylor expansion of a d -times differentiable function f at a point \bar{x} as

$$T_{\bar{x}, d}(x) := f(\bar{x}) + \sum_{i=1}^d \frac{1}{i!} \langle \nabla^i f(\bar{x}), (x - \bar{x})^{\boxtimes i} \rangle,$$

³That is, there exists some $\bar{x} \in \mathbb{R}^n$ such that $g_j(\bar{x}) < 0$ for all $j = 1, \dots, m$.

⁴We use this slightly nonstandard notation to avoid confusion with the Kronecker product.

where $\langle \cdot, \cdot \rangle$ denotes the standard tensor inner product. The *remainder* or *error* term of the Taylor expansion is

$$R_{\bar{x},d}(x) := f(x) - T_{\bar{x},d}(x).$$

For a d^{th} -order tensor D , let us define the following norm

$$\|D\| := \max_{\|x_1\|, \dots, \|x_d\| \leq 1} \langle D, x_1 \boxtimes x_2 \boxtimes \dots \boxtimes x_d \rangle,$$

where $\|x_i\|$ denotes the Euclidean 2-norm of the vector $x_i \in \mathbb{R}^n$. Note that for cases of $d = 1$ and $d = 2$, this expression reduces to the standard Euclidean norm and the spectral norm, respectively.

We will need the following lemma in Section 4.

Lemma 1 (see, e.g., inequality (11) in [9]). *Fix a vector $\bar{x} \in \mathbb{R}^n$. Suppose $\nabla^d f$ has a Lipschitz constant L over a convex set C containing \bar{x} , i.e.,*

$$\|\nabla^d f(x) - \nabla^d f(y)\| \leq L\|x - y\|$$

for all $x, y \in C$. Then, for any $x \in C$, we have

$$\|\nabla R_{\bar{x},d}(x)\| \leq \frac{L}{d!} \|x - \bar{x}\|^d.$$

and

$$\|\nabla^2 R_{\bar{x},d}(x)\| \leq \frac{L}{(d-1)!} \|x - \bar{x}\|^{d-1}.$$

3 Algorithm Definition

For a given integer $d \geq 3$, we consider the task of minimizing a function f which is assumed to have derivatives up to order d , and a local minimum x^* satisfying $\nabla^2 f(x^*) \succ 0$. We also assume that the d^{th} derivative of f is locally Lipschitz around the point x^* , i.e., there is a radius $r_L > 0$, and a scalar $L \geq 0$, such that for points x, y in the set $\{z \in \mathbb{R}^n \mid \|z - x^*\| \leq r_L\}$, we have

$$\|\nabla^d f(x) - \nabla^d f(y)\| \leq L\|x - y\|.$$

Note that the latter assumption is always satisfied if the $d + 1^{\text{th}}$ derivative of f exists and is continuous. Our goal is to minimize f by iteratively minimizing a surrogate function of the type

$$T_{x_k,d}(x) + t\|x - x_k\|^{d'},$$

where x_k is our current iterate, $T_{x_k,d}$ is the Taylor expansion of f of order d at x_k , d' is the smallest even integer greater than d (as we require the surrogate to be a polynomial), and t is chosen according to the following sum of squares program:

$$\begin{aligned}
& \min_{t \in \mathbb{R}} t \\
& \text{s.t. } T_{x_k, d}(x) + t \|x - x_k\|^{d'} \quad \text{sos-convex} \\
& t \geq 0.
\end{aligned} \tag{3}$$

In view of Theorem 1 and the remarks after Definition 2, this program can be reformulated as an SDP of size polynomial in n . Letting $t(x_k)$ denote the optimal value of (3) for a given x_k , we define our surrogate function to be

$$\psi_{x_k, d}(x) := T_{x_k, d}(x) + t(x_k) \|x - x_k\|^{d'}. \tag{4}$$

In our algorithm, we choose x_{k+1} to be the minimizer of $\psi_{x_k, d}$ (which exists and is unique; see Theorem 3 below). By Theorem 2, since $\psi_{x_k, d}$ is sos-convex, we can find its minimizer via another SDP of size polynomial in n .

If x_k is far from x^* so that $\nabla^2 f(x_k)$ is not positive definite, it may occur that (3) is infeasible. If this occurs, we fix a positive scalar⁵ ε and instead solve the SDP:

$$\begin{aligned}
& \min_{\bar{t} \in \mathbb{R}} \bar{t} \\
& \text{s.t. } T_{x_k, d}(x) + \frac{1}{2} \left(\varepsilon - \lambda_{\min} \nabla^2 f(x_k) \right) \|x - x_k\|^2 + \bar{t} \|x - x_k\|^{d'} \quad \text{sos-convex} \\
& \bar{t} \geq 0.
\end{aligned} \tag{5}$$

Let $\bar{t}(x_k)$ denote the optimal value of (5) and define

$$\bar{\psi}_{x_k, d}(x) := T_{x_k, d}(x) + \frac{1}{2} \left(\varepsilon - \lambda_{\min} \nabla^2 f(x_k) \right) \|x - x_k\|^2 + \bar{t}(x_k) \|x - x_k\|^{d'}. \tag{6}$$

We then let x_{k+1} be the minimizer $\bar{\psi}_{x_k, d}$ (which again exists and is unique; see Theorem 3 below). As before, we can find a minimizer of $\bar{\psi}_{x_k, d}$ by solving an SDP of size polynomial in n ; see Theorem 2.

Our overall algorithm is summarized below:

Algorithm 1: d^{th} -order Newton method

Parameter: $\varepsilon > 0$
Input: $x_0 \in \mathbb{R}^n$

- 1 **for** $k = 0, \dots$ **do**
- 2 **if** $\nabla^2 f(x_k) \succ 0$ **then**
- 3 Solve (3) to find $t(x_k)$
- 4 Let x_{k+1} be the minimizer of $\psi_{x_k, d}$ (see (4))
- 5 **else**
- 6 Solve (5) to find $\bar{t}(x_k)$
- 7 Let x_{k+1} be the minimizer of $\bar{\psi}_{x_k, d}$ (see (6))
- 8 **end**
- 9 **end**

⁵Our analysis applies to any positive value of ε .

4 Algorithm Analysis and Convergence

In this section, we present our main technical results. Theorem 3 shows that our algorithm is well-defined for all initial conditions. Theorem 4 gives our convergence result. We remind the reader that the assumptions made on f are described in the first paragraph of Section 3. In particular, the function f is not required to be convex, and the d^{th} derivatives of f are not required to be globally Lipschitz.

Theorem 3. *Algorithm 1 is well-defined in the sense that*

- (i) *the problems (3) and (5) are always feasible when required at Lines 3 and 6, and*
- (ii) *the functions $\psi_{x_k,d}$ and $\bar{\psi}_{x_k,d}$ (see (4) and (6)) always possess a unique minimizer when required at Lines 4 and 7.*

Theorem 4. *There exist constants $r, c > 0$ such that if $\|x_0 - x^*\| \leq r$, then the sequence $\{x_k\}$ generated by Algorithm 1 satisfies*

$$\|x_{k+1} - x^*\| \leq c \|x_k - x^*\|^d$$

for all k .

The power d in this theorem is referred to as the *order* of convergence and the constant c is referred to as the *factor* of convergence. We note that the factor of convergence arising from our proof is explicit.

To prove Theorems 3 and 4, we first establish some technical lemmas. Lemmas 2 and 3 are used to prove the first claim of Theorem 3; Lemmas 4 and 5 are for the second claim; and Lemmas 3, 4, and 6 are employed in the proof of Theorem 4.

In Lemma 2, we show that a particular polynomial is in the interior of the cone of sos-convex polynomials. This is used in Lemma 3 to show that we can always make our surrogate functions defined in (3) and (5) sos-convex.

Lemma 2. *Let $x := (x_1, \dots, x_n)$. The polynomial*

$$p(x) = x^T x + (x^T x)^d$$

is in the interior of the cone of sos-convex polynomials in n variables and of degree at most $2d$.

Proof. We first establish the following claim:

Claim 0. For all $d \geq 0$, the polynomial

$$\tilde{p}_d(x) = 1 + (d+1)(x^T x)^d$$

can be written as $\phi_d(x)^T Q \phi_d(x)$, where ϕ_d is the standard basis of monomials of degree up to d with the monomials appearing in ascending order of degree, and Q is a positive definite matrix.

To prove Claim 0, it suffices to show that for all $d \geq 0$, there exists a constant $\alpha_d > 0$ and a positive definite matrix \hat{Q}_d such that $1 + \alpha_d(x^T x)^d = \phi_d(x)^T \hat{Q}_d \phi_d(x)$. Indeed, if $\alpha_d < d + 1$, we can observe that

$$\tilde{p}_d(x) = 1 + \alpha_d(x^T x)^d + ((d + 1) - \alpha_d)(x^T x)^d = \phi_d(x)^T \left(\hat{Q}_d + Q' \right) \phi_d(x),$$

where Q' can be taken to be positive semidefinite since $((d + 1) - \alpha_d)(x^T x)^d$ is sos. If $\alpha_d > d + 1$, we can observe that

$$\tilde{p}_d(x) = \frac{d + 1}{\alpha_d} + (d + 1)(x^T x)^d + \left(1 - \frac{d + 1}{\alpha_d}\right) = \phi_d(x)^T \left(\frac{d + 1}{\alpha_d} \hat{Q}_d + Q' \right) \phi_d(x),$$

where Q' can be taken to be positive semidefinite since $(1 - \frac{d+1}{\alpha_d})$ is sos.

Let us now proceed by induction on d to prove the claim made in the previous paragraph. The case of $d = 0$ is clear since we can take any $\alpha_0 > 0$ and the associated matrix \hat{Q}_0 is simply a 1×1 matrix containing the scalar $1 + \alpha_0$. Now suppose that the induction hypothesis holds for $d = k$. To construct α_{k+1} and \hat{Q}_{k+1} , we will add matrices associated with the polynomials $1 + \alpha_k(x^T x)^k$ and $\alpha(x^T x)^{k+1} - \alpha_k(x^T x)^k$, where α is an arbitrary scalar. From the induction hypothesis, there exist a scalar $\alpha_k > 0$ and a matrix $\hat{Q}_k \succ 0$ of size $\binom{n+k}{k} \times \binom{n+k}{k}$ that satisfy

$$1 + \alpha_k(x^T x)^k = \phi_k(x)^T \hat{Q}_k \phi_k(x) = \phi_{k+1}(x)^T \begin{bmatrix} \hat{Q}_k & 0 \\ 0 & 0 \end{bmatrix} \phi_{k+1}(x).$$

Meanwhile, observe that we can write

$$\alpha(x^T x)^{k+1} - \alpha_k(x^T x)^k = \phi_{k+1}(x)^T \begin{bmatrix} 0 & A \\ A^T & \alpha P \end{bmatrix} \phi_{k+1}(x)$$

for some matrices A and $P \succ 0$, where the zero block is of size $\binom{n+k}{k} \times \binom{n+k}{k}$. Indeed, we can take the matrix P to be diagonal with its diagonal entries equalling the coefficients of $(x^T x)^{k+1}$ and move the coefficients of $\alpha_k(x^T x)^k$ to the matrix A . Adding the two identities, we observe that:

$$1 + \alpha(x^T x)^{k+1} = \phi_{k+1}(x)^T \begin{bmatrix} \hat{Q}_k & A \\ A^T & \alpha P \end{bmatrix} \phi_{k+1}(x).$$

Since \hat{Q}_k and P are both positive definite matrices, by the Schur complement condition, whenever $\alpha P - A^T \hat{Q}_k^{-1} A \succ 0$, the matrix on the right-hand side of the above expression will be positive definite. One can therefore choose α_{k+1} to be any large enough value of α that satisfies the previous condition and let $\hat{Q}_{k+1} := \begin{bmatrix} \hat{Q}_k & A \\ A^T & \alpha_{k+1} P \end{bmatrix}$. We have thus proved Claim 0.

By Claim 0 (with d replaced by $d - 1$), we can fix a positive definite matrix

Q such that $1 + d(x^T x)^{d-1} = \phi(x)^T_{d-1} Q \phi_{d-1}(x)$ for all x . One can check that

$$\begin{aligned} y^T \nabla^2 p(x) y &= y^T \left(2I + 2d(x^T x)^{d-1} I + 4d(d-1)(x^T x)^{d-2} x x^T \right) y \\ &= 2(y^T y)(1 + d(x^T x)^{d-1}) + 4d(d-1)(x^T x)^{d-2} (x^T y)^2 \\ &= 2(y^T y) \phi(x)^T_{d-1} Q \phi_{d-1}(x) + 4d(d-1)(x^T x)^{d-2} (x^T y)^2 \\ &= (\phi_{d-1}(x) \otimes y)^T (Q \otimes 2I + Q') (\phi_{d-1}(x) \otimes y), \end{aligned}$$

where Q' can be taken to be positive semidefinite since $4d(d-1)(x^T x)^{d-2} (x^T y)^2$ is sos. Since the matrix $Q \otimes 2I + Q'$ is positive definite, it follows that p is in the interior of the cone of sos-convex polynomials of degree at most $2d$. \square

Lemma 3. *Suppose $f : \mathbb{R}^n \mapsto \mathbb{R}$ has continuous derivatives up to order d over a compact set $B \subseteq \mathbb{R}^n$. If $\nabla^2 f(x) \succ 0$ for all $x \in B$, then $t(x)$ (i.e., the optimal value of (3)) is uniformly bounded from above over B .*

Proof. Let δ be a positive scalar such that $\lambda_{\min} \nabla^2 f(x) \geq \delta$ for all $x \in B$. Let x' be any vector in B , and define

$$F_{x'}(x) := \frac{2}{\delta} T_{x',d}(x' + x).$$

Since $\nabla^2 f(x') \succeq \delta I$, we have $\nabla^2 F_{x'}(0) \succeq 2I$. Let $Q_{x'}$ (resp. $C_{x'}$) be the sum of the quadratic and higher (resp. cubic and higher) terms of $F_{x'}$. For a polynomial p , define $\|p\|_\infty$ as the infinity norm of the coefficients of p when expressed in the standard monomial basis. By Lemma 2, we can fix a positive scalar R such that for any polynomial p of degree at most d with $\|p\|_\infty \leq R$, we have that the polynomial $\|x\|^2 + \|x\|^{d'} + p$ is sos-convex. Fix a scalar M such that $\|C_{x'}\|_\infty < M$ for all $x' \in B$. Define $\alpha := \min\{1, \frac{R}{M}\}$. We have $\|x \mapsto C_{x'}(\alpha x)\|_\infty \leq \alpha^3 \|C_{x'}\|_\infty$ since all terms of $C_{x'}$ are of cubic or higher order. Then we can write

$$\begin{aligned} \frac{1}{\alpha^2} Q_{x'}(\alpha x) + \|x\|^{d'} &= \frac{1}{2} x^T \nabla^2 F_{x'}(0) x + \frac{1}{\alpha^2} C_{x'}(\alpha x) + \|x\|^{d'} \\ &= \frac{1}{2} x^T (\nabla^2 F_{x'}(0) - 2I) x + \frac{1}{\alpha^2} C_{x'}(\alpha x) + (\|x\|^2 + \|x\|^{d'}). \end{aligned}$$

Since $\nabla^2 F_{x'}(0) \succeq 2I$, the first term is sos-convex. We can bound the second term as follows: $\|x \mapsto \frac{1}{\alpha^2} C_{x'}(\alpha x)\|_\infty \leq \alpha \|C_{x'}\|_\infty \leq \alpha M \leq R$. Thus, the sum of the second and the third term is sos-convex by the definition of R . It follows that the polynomial

$$\frac{1}{\alpha^2} Q_{x'}(\alpha x) + \|x\|^{d'} \text{ is sos-convex.}$$

We can then conclude the sos-convexity of the polynomials

$$(a) \quad Q_{x'}(\alpha x) + \alpha^2 \|x\|^{d'},$$

- (b) $F_{x'}(\alpha x) + \alpha^2 \|x\|^{d'}$,
- (c) $F_{x'}(\alpha x) + \alpha^{2-d'} \|\alpha x\|^{d'}$,
- (d) $F_{x'}(x) + \alpha^{2-d'} \|x\|^{d'}$,
- (e) $T_{x',d}(x' + x) + \frac{\delta}{2} \alpha^{2-d'} \|x\|^{d'}$, and
- (f) $T_{x',d}(x) + \frac{\delta}{2} \alpha^{2-d'} \|x - x'\|^{d'}$,

respectively (a) by scaling, (b) by the observation that the affine terms do not affect sos-convexity, (c) by rewriting, (d) by a linear change of coordinates, (e) by another rescaling, and (f) by an affine change of coordinates. Thus, we have $t(x) \leq \frac{\delta}{2} \alpha^{2-d'}$ for $x \in B$. □

We next use a quadrature rule for integration to establish a technical lemma that is needed for the remainder of this section. By a polynomial matrix, we mean a matrix whose entries are polynomial functions.

Lemma 4. *Let $M : \mathbb{R} \mapsto \mathbb{S}^{n \times n}$ be univariate polynomial matrix whose entries have degree at most d , where d is even. Suppose $M(s) \succeq 0$ for all $s \in [0, 1]$. Then,*

$$\int_0^1 M(s) ds \succeq \frac{1}{2(d^2 - 1)} M(\alpha)$$

for $\alpha \in \{0, 1\}$.

Proof. Using a quadrature rule for integration proposed in [21] and analyzed in [27], there exist a set of weights $w_0, \dots, w_d \geq 0$, with $w_0 = \frac{1}{d^2 - 1}$, and a set of points $s_0, \dots, s_d \in [-1, 1]$, with $s_0 = 1$, such that for any polynomial p of degree at most d we have

$$\int_{-1}^1 p(s) ds = \sum_{i=0}^d w_i p(s_i).$$

Now we can write

$$\begin{aligned} \int_0^1 M(s) ds &= \frac{1}{2} \int_{-1}^1 M\left(\frac{1-s}{2}\right) ds \\ &= \frac{1}{2} \sum_{i=0}^d w_i M\left(\frac{1-s_i}{2}\right) \\ &\succeq \frac{1}{2} w_0 M\left(\frac{1-s_0}{2}\right) = \frac{1}{2(d^2 - 1)} M(0). \end{aligned}$$

By replacing s with $1 - s$, the claim with $\alpha = 1$ follows. □

The next lemma directly proves the second claim of Theorem 3 and is possibly of independent interest.

Lemma 5. *If a convex polynomial $p : \mathbb{R}^n \mapsto \mathbb{R}$ satisfies $\nabla^2 p(x_0) \succ 0$ for any point $x_0 \in \mathbb{R}^n$, then p has a unique minimizer.*

Proof. Without loss of generality, assume $x_0 = 0$. Let d be an even integer greater than the degree of the Hessian of p . For any $x \in \mathbb{R}^n$,

$$\begin{aligned} p(x) &= p(0) + x^T \nabla p(0) + x^T \left(\int_0^1 \int_0^t \nabla^2 p(sx) ds dt \right) x \\ &= p(0) + x^T \nabla p(0) + x^T \left(\int_0^1 t \int_0^1 \nabla^2 p(stx) ds dt \right) x \\ &\geq p(0) + x^T \nabla p(0) + x^T \left(\int_0^1 t \left(\frac{1}{2(d^2 - 1)} \nabla^2 p(0) \right) dt \right) x \\ &= p(0) + x^T \nabla p(0) + \frac{1}{4} x^T \left(\frac{1}{d^2 - 1} \nabla^2 p(0) \right) x, \end{aligned}$$

where the inequality follows from Lemma 4. Thus, p is lower bounded by a coercive⁶ quadratic function, and hence p is coercive itself. A coercive function that is convex (and hence continuous) has at least one minimizer.

Suppose for the sake of contradiction that p had two minimizers \bar{x}, \bar{y} . Then, by convexity, any point on the line segment connecting \bar{x} and \bar{y} would also be a minimizer. Since p is a polynomial, it follows that p must be constant along the line passing through \bar{x} and \bar{y} . This contradicts coercivity. \square

We remark that the statement of Lemma 5 does not hold for non-polynomial convex functions (consider, e.g., the univariate function $\max\{0, x^2 - 1\}$).

The next lemma is used in the proof of Theorem 4.

Lemma 6. *There exists a constant $r > 0$ such that if $\|x_k - x^*\| \leq r$, then $\lambda_{\min} \nabla^2 \psi_{x_k, d}(x^*) \geq \frac{1}{2} \lambda_{\min} \nabla^2 f(x^*)$.*

Proof. We show that we can take

$$r = \min \left\{ r_L, \left(\frac{(d-1)! \lambda_{\min} \nabla^2 f(x^*)}{2L} \right)^{\frac{1}{d-1}} \right\}$$

where r_L and L are as in the first paragraph of Section 3. By Lemma 1, For every x satisfying $\|x - x_k\| \leq r_L$, we have

$$\|\nabla^2 f(x) - \nabla^2 T_{x_k, d}(x)\| \leq \frac{L}{(d-1)!} \|x - x_k\|^{d-1}.$$

Thus, if $\|x^* - x_k\| \leq r$, we have

$$\|\nabla^2 f(x^*) - \nabla^2 T_{x_k, d}(x^*)\| \leq \frac{1}{2} \lambda_{\min} \nabla^2 f(x^*).$$

⁶We recall that a function $g : \mathbb{R}^n \mapsto \mathbb{R}$ is coercive if $g(x) \rightarrow \infty$ as $\|x\| \rightarrow \infty$.

It follows that

$$\lambda_{\min} \nabla^2 T_{x_k, d}(x^*) \geq \frac{1}{2} \lambda_{\min} \nabla^2 f(x^*).$$

Indeed, if there was a unit vector y such that if $y^T \nabla^2 T_{x_k, d}(x^*) y < \frac{1}{2} \lambda_{\min} \nabla^2 f(x^*)$, the previous inequality would be violated.

Recall from (4) that $\psi_{x_k, d}$ is obtained by adding to $T_{x_k, d}$ the convex function $t(x_k) \|x - x_k\|^{d'}$. Therefore, we have $\nabla^2 \psi_{x_k, d}(x^*) \succeq \nabla^2 T_{x_k, d}(x^*)$, which gives the claim. \square

We now have all the ingredients to prove Theorems 3 and 4.

Proof of Theorem 3.

(i) When $\nabla^2 f(x_k) \succ 0$, the proof of Lemma 3 with $B = \{x_k\}$ demonstrates a feasible solution to (3). This argument also extends to show feasibility of (5) since the polynomial $T_{x_k, d}(x) + \frac{1}{2}(\varepsilon - \lambda_{\min} \nabla^2 f(x_k)) \|x - x_k\|^2$ has a positive definite Hessian at x_k .

(ii) At Line 4 (resp. Line 7), $\psi_{x_k, d}$ (resp. $\bar{\psi}_{x_k, d}$) has a positive definite Hessian at x_k . Moreover, the polynomial $\psi_{x_k, d}$ (resp. $\bar{\psi}_{x_k, d}$) is sos-convex and therefore convex. Thus, by Lemma 5, $\psi_{x_k, d}$ (resp. $\bar{\psi}_{x_k, d}$) has a unique minimizer. \square

Proof of Theorem 4. Since $d > 1$, it suffices to show that there exist constants $r', c' > 0$ such that if $\|x_0 - x^*\| \leq r'$, then $\|x_1 - x^*\| \leq c' \|x_0 - x^*\|^d$.

By continuity of the map $x \mapsto \lambda_{\min} \nabla^2 f(x)$, there exists a scalar $r_1 > 0$ such that $\lambda_{\min} \nabla^2 f(x) \geq \frac{1}{2} \lambda_{\min} \nabla^2 f(x^*) > 0$ for all x with $\|x - x^*\| \leq r_1$.

Let $r_2 > 0$ be the constant needed for the conclusion of Lemma 6 to hold. Define

$$r' := \min\{r_L, r_1, r_2\}$$

and $\Omega := \{x \in \mathbb{R}^n \mid \|x - x^*\| \leq r'\}$. Suppose $x_0 \in \Omega$. Note that in this case, Algorithm 1 finds the next iterate x_1 by minimizing the polynomial $\psi_{x_0, d}$ defined in (4).

By the fundamental theorem of calculus, we have

$$\nabla \psi_{x_0, d}(x^*) - \nabla \psi_{x_0, d}(x_1) = \left(\int_0^1 \nabla^2 \psi_{x_0, d}(x_1 + s(x^* - x_1)) ds \right) (x^* - x_1).$$

Since x_1 minimizes $\psi_{x_0, d}$, we have $\nabla \psi_{x_0, d}(x_1) = 0$, and thus

$$\nabla \psi_{x_0, d}(x^*) = \left(\int_0^1 \nabla^2 \psi_{x_0, d}(x_1 + s(x^* - x_1)) ds \right) (x^* - x_1).$$

We can bound the norm of this vector from below:

$$\|\nabla \psi_{x_0, d}(x^*)\| \geq \lambda_{\min} \left(\int_0^1 \nabla^2 \psi_{x_0, d}(x_1 + s(x^* - x_1)) ds \right) \|x^* - x_1\|. \quad (7)$$

Applying first Lemma 4 and then Lemma 6, we have

$$\begin{aligned}\lambda_{\min} \left(\int_0^1 \nabla^2 \psi_{x_0, d}(x_1 + s(x^* - x_1)) ds \right) &\geq \frac{\lambda_{\min} \nabla^2 \psi_{x_0, d}(x^*)}{2((d' - 2)^2 - 1)} \\ &\geq \frac{\lambda_{\min} \nabla^2 f(x^*)}{4((d' - 2)^2 - 1)}.\end{aligned}$$

Substituting this into (7) and rearranging yields

$$\|x_1 - x^*\| \leq \frac{4((d' - 2)^2 - 1)}{\lambda_{\min} \nabla^2 f(x^*)} \|\nabla \psi_{x_0, d}(x^*)\|. \quad (8)$$

Expanding $\nabla \psi_{x_0, d}(x^*)$, we have

$$\begin{aligned}\|\nabla \psi_{x_0, d}(x^*)\| &= \left\| \nabla T_{x_0, d}(x^*) + \nabla(t(x_0) \|x - x_0\|^{d'}) \Big|_{x^*} \right\| \\ &= \left\| \nabla T_{x_0, d}(x^*) + t(x_0) d' \|x^* - x_0\|^{d'-2} (x^* - x_0) \right\| \\ &\leq \|\nabla T_{x_0, d}(x^*)\| + t(x_0) d' \|x^* - x_0\|^{d'-1}.\end{aligned}$$

Applying Lemma 1 and noting that $\nabla f(x^*) = 0$, we have

$$\|\nabla \psi_{x_0, d}(x^*)\| \leq \frac{L}{d!} \|x^* - x_0\|^d + t(x_0) d' \|x^* - x_0\|^{d'-1}.$$

Using Lemma 3 and the fact that $\|x^* - x_0\| \leq r'$, we get

$$\begin{aligned}\|\nabla \psi_{x_0, d}(x^*)\| &\leq \frac{L}{d!} \|x^* - x_0\|^d + (\sup_{x \in \Omega} t(x)) d' \max\{r', 1\} \|x^* - x_0\|^d \\ &= \left(\frac{L}{d!} + (\sup_{x \in \Omega} t(x)) d' \max\{r', 1\} \right) \|x^* - x_0\|^d.\end{aligned}$$

Substituting into (8), we have

$$\|x_1 - x^*\| \leq \left(\frac{4((d' - 2)^2 - 1)}{\lambda_{\min} \nabla^2 f(x^*)} \left(\frac{L}{d!} + (\sup_{x \in \Omega} t(x)) d' \max\{r', 1\} \right) \right) \|x^* - x_0\|^d$$

as desired. \square

5 Numerical Examples

We present three examples to compare the performance of our d^{th} -order Newton methods and the classical Newton method.

5.1 The Univariate Case

In the univariate case, the iterations of the classical Newton method read

$$x_{k+1} = x_k - \frac{f'(x_k)}{f''(x_k)}.$$

In terms of finding a root of f' , this iteration can be interpreted as first computing the first-order Taylor expansion of f' at x_k , and then finding the root of this affine function to define x_{k+1} .

We derive a similar explicit formula for our higher-order Newton method in the case where $n = 1$, $d = 3$, and f'' is positive. Since convex univariate polynomials are sos-convex, finding explicit solutions to the two SDPs involved in each iteration of our algorithm reduces to arguments about roots of univariate polynomials.

Proposition 1. *In the univariate case, when $f''(x_k) > 0$ and $f'''(x_k) \neq 0$, the next iterate of the 3rd-order version of Algorithm 1 is given by⁷*

$$x_{k+1} = x_k - 2 \frac{f''(x_k)}{f'''(x_k)} - \sqrt[3]{\frac{f'(x_k) - \frac{2}{3} \frac{(f''(x_k))^2}{f'''(x_k)}}{\frac{(f'''(x_k))^2}{12f''(x_k)}}}.$$

Proof. To simplify notation, we let $T := T_{x_k,3}$ and $\psi := \psi_{x_k,3}$. By translation, we may assume $x_k = 0$. Then $T(x) = f(x_k) + x f'(x_k) + \frac{1}{2} x^2 f''(x_k) + \frac{1}{6} x^3 f'''(x_k)$, $\psi(x) = T(x) + t x^4$, where t is the smallest constant that makes ψ convex. We have $\psi''(x) = f''(x_k) + x f'''(x_k) + 12t x^2$. The discriminant of ψ'' is $(f'''(x_k))^2 - 48t f''(x_k)$, which tells us that $t = \frac{(f'''(x_k))^2}{48f''(x_k)}$.

To find x_{k+1} , we look for the root of ψ' . One can write the expression for ψ' in the following form:

$$\psi'(x) = \frac{(f'''(x_k))^2}{12f''(x_k)} \left(x + 2 \frac{f''(x_k)}{f'''(x_k)} \right)^3 + f'(x_k) - \frac{2}{3} \frac{(f''(x_k))^2}{f'''(x_k)}.$$

Observe that a univariate cubic polynomial of the form $a(x-b)^3 + c$, with $a \neq 0$, has a unique root at $x = b - \sqrt[3]{\frac{c}{a}}$. Therefore, after a translation back by x_k , we have

$$x_{k+1} = x_k - 2 \frac{f''(x_k)}{f'''(x_k)} - \sqrt[3]{\frac{f'(x_k) - \frac{2}{3} \frac{(f''(x_k))^2}{f'''(x_k)}}{\frac{(f'''(x_k))^2}{12f''(x_k)}}}.$$

□

As in the case of the classical Newton method, the expression in Proposition 1 can be interpreted geometrically in terms of finding a root of f' . This iteration computes the second-order Taylor expansion of f' at x_k , adds a sufficiently large cubic term to enforce monotonicity, and then finds the root of this monotone cubic function to define x_{k+1} .

⁷Note that when $f''(x_k) > 0$ and $f'''(x_k) = 0$, the third-order Taylor series is convex and coincides with the second-order Taylor series. Therefore, the next iterates of the third-order and the classical Newton method coincide.

Example 1 In this example, we apply our method to the univariate function

$$f(x) = \sqrt{x^2 + 1} - 1. \quad (9)$$

This is a strictly convex function with its unique minimizer at $x^* = 0$. One can check that the classical Newton method converges to this minimizer if and only if $|x_0| < 1$. Using Proposition 1, we can calculate the exact basin of convergence of our third-order Newton method to be $(-\beta, \beta)$, where

$$\beta = \sqrt{\frac{1}{3} \left(11 + \frac{142}{\sqrt[3]{1691 + 9i\sqrt{47}}} + \sqrt[3]{1691 + 9i\sqrt{47}} \right)} \sim 3.407.$$

This is strictly larger than the basin of convergence of the classical method.

Figure 1 demonstrates the difference between one iteration of the classical and our third-order Newton method starting at the point $x_0 = 1.5$. We display the quadratic and quartic polynomials $T_{x_0,2}$ and $\psi_{x_0,3}$. The minimizers of these polynomials are denoted by x_1^{Newton} and $x_1^{3\text{ON}}$, which are respectively the next iterates of the classical and our third-order Newton method. Since the third-order Taylor expansion of f provides a more accurate approximation, we see that the next iterate of our method is closer to x^* , while that of the classical Newton method moves farther away from x^* .




Figure 1: A comparison of one iteration of the classical Newton method and our third-order Newton method applied to the function in (9) starting at $x_0 = 1.5$.

For our d^{th} -order Newton methods with $d > 3$, we calculate the radii of convergence numerically. These radii increase with degree as the following table demonstrates:

Degree d	Radius of Convergence
2 (Classical Newton)	1
3	~ 3.4
4	~ 4.5
5	~ 5.9

We can visualize the speed of convergence of the fifth-order method, for example, in Figure 2. In this figure, we plot the absolute value of $|x_k - x^*|$ starting at $x_0 = 5.9$, which is close to the boundary of the basin. In just five iterations, the method reaches a point with absolute value approximately 10^{-15} .




Figure 2: 5th-order Newton iterates applied to the function in (9).

Example 2 In this example, we compare our third-order method to the classical Newton method when applied to the function

$$f(x) = 2x \arctan(x) - \log(1 + x^2) + \frac{1}{10}x^2. \quad (10)$$

This is a strongly convex function with its unique minimizer at $x^* = 0$.

In Figure 3, N_2 (resp. N_3) is the map that takes a point to the corresponding next iterate of the classical (resp. third-order) Newton method. In this example, the third-order method satisfies $|N_3(x)| < |x|$ for all nonzero x , implying global convergence of the method. Meanwhile, the classical Newton method oscillates between ± 13.494 when x_0 is outside of the range $[-\alpha, \alpha]$, where $\alpha \sim 1.712$ is point of intersection of the functions N_2 and $-x$.

In Figure 4, we can see a comparison of the iterates of the third-order and the classical Newton method starting from the initial condition $x_0 = 1.7$. While both methods converge to the minimizer, the third-order method converges much faster.




Figure 3: Comparison of the classical Newton map N_2 and our third-order Newton map N_3 applied to the function in (10). Subfigure (a) implies that the third-order method is globally convergent, while the classical method is not. Subfigure (b) zooms in on the behavior of these maps near the origin to show that the basin of attraction for the classical method is approximately $(-1.712, 1.712)$.




Figure 4: Iterates of our third-order and the classical Newton method applied to the function in (10) starting from a point in the basin of attraction of both methods.

5.2 A Multivariate Example

In our last example, we compare the classical and the third-order Newton methods applied to a standard test function in nonlinear optimization called the Beale function:

$$f(x_1, x_2) = (1.5 - x_1 + x_1x_2)^2 + (2.25 - x_1 + x_1x_2^2)^2 + (2.625 - x_1 + x_1x_2^3)^2.$$

This nonconvex function has a single global minimum at $x^* = (3, 0.5)^T$ and no other local minima. In Figure 5, we explore the behavior of both methods with initial conditions in the region $\{x \in \mathbb{R}^2 \mid \|x\|_\infty \leq 4\}$. We initialize the classical method and our third-order method at a fine grid of points in this box and run both methods for 350 iterations. For our third-order method, we take the parameter ε in Algorithm 1 to be equal to 0.01. In Figure 5, the color yellow corresponds to initial points that converge to x^* , and the color blue corresponds to any other behavior including divergence or convergence to a point which is not a local minimum. In this example, the two basins are incomparable, but that of the third-order method is more contiguous and larger in volume.

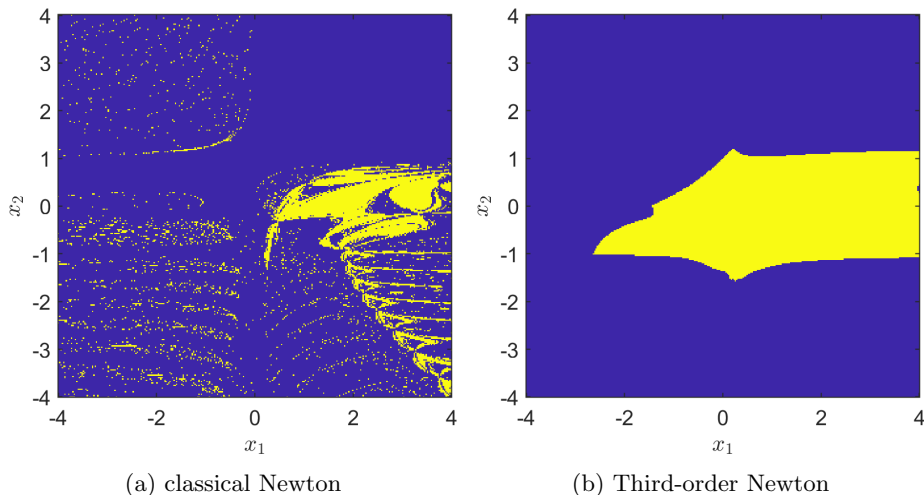


Figure 5: The basins of attraction for the classical and the third-order Newton methods for the minimizer of the Beale function. The basin for the classical method has fractal structure, demonstrating more sensitivity to initialization.

6 Global convergence

In this section, we present a slightly modified algorithm which has global convergence under additional assumptions. There is a vast literature on modifications to Newton’s method that lead to global convergence in special circumstances: see, e.g., [41, 43, 38, 22]. In the setting of our work, it turns out that we can use a result of Nesterov from [40] to show that a simple modification to our

algorithm that still has polynomial work per iteration is globally convergent when the Taylor expansion is made to an odd order.⁸ This modified algorithm (Algorithm 2 below) also inherits the local convergence order of Algorithm 1.

As in [40], suppose the d^{th} derivative of the function $f : \mathbb{R}^n \mapsto \mathbb{R}$ that we wish to minimize has a Lipschitz constant L_d , and that an upper bound M on L_d is known. In this setting, consider the following algorithm:

Algorithm 2: d^{th} -order globally convergent Newton method (d odd)

Input: $x_0 \in \mathbb{R}^n$
1 for $k = 0, \dots$ **do**
2 Solve (3) to find $t(x_k)$
3 Let x_{k+1} be the minimizer of
 $T_{x_k, d}(x) + \max\{\frac{dM}{(d+1)!}, t(x_k)\} \|x - x_k\|^{d+1}$
4 end

Using the same arguments as those in the proof of Theorem 3, one can see that the next iterate x_{k+1} produced by this algorithm is well-defined whenever $\nabla^2 f(x_k) \succ 0$. Also as before, problem (3) can be solved as a semidefinite program of size polynomial in the dimension. This claim also holds for the problem of finding the (unique) minimizer of the degree $d + 1$ polynomial

$$T_{x_k, d} + \max\left\{\frac{dM}{(d+1)!}, t(x_k)\right\} \|x - x_k\|^{d+1}.$$

This is because the polynomials $\|x - x_k\|^{d+1}$ and $T_{x_k, d} + t(x_k)\|x - x_k\|^{d+1}$ are sos-convex and a conic combination of two sos-convex polynomials is sos-convex, making Theorem 2 applicable.

Theorem 5. *Suppose $f : \mathbb{R}^n \rightarrow \mathbb{R}$ has bounded level sets, a positive definite Hessian everywhere, and the Lipschitz constant of its d^{th} derivative bounded above by M .⁹ Then, the iterates of Algorithm 2 starting from any $x_0 \in \mathbb{R}^n$ converge to the (unique) minimizer of f . Furthermore, Algorithm 2 has local convergence rate of order d .*

Proof. Since the Hessian of f is positive definite everywhere, the function f is strictly convex. This, along with boundedness of the level sets, implies that f has a unique (global) minimizer which we call x^* .

Define $\psi_{x_k, d}(x) := T_{x_k, d}(x) + \max\{\frac{dM}{(d+1)!}, t(x_k)\} \|x - x_k\|^{d+1}$. By Theorem 1 from [40], we have $\psi_{x_k, d}(x) \geq f(x)$ for all $x \in \mathbb{R}^n$, thus the method is monotone; i.e., $f(x_{k+1}) \leq f(x_k)$. Let $M_k := \max\left\{M, \frac{(d+1)!t(x_k)}{d}\right\}$ and $\delta_k := f(x_k) - f(x^*)$.

⁸The reason we need the Taylor expansion order to be odd is that in the work of Nesterov, the Taylor polynomial is regularized by a term of degree one larger. We need this new term to be a polynomial function for sum of squares methods to be readily applicable.

⁹The assumptions that we make here are the same as those in [40] except that our assumption of positive definiteness of the Hessian is stronger than the assumption of positive semidefiniteness of the Hessian made in [40].

Since the set $\{x \in \mathbb{R}^n \mid f(x) \leq f(x_0)\}$ is compact and the method is monotone, there exists a scalar D such that $\|x_k - x^*\| \leq D$ for all k . By the arguments in the proof of Theorem 2 from [40], we can conclude that

$$\delta_k - \delta_{k+1} \geq C_k \delta_k^{\frac{d+1}{d}},$$

where $C_k := \frac{d}{d+1} \left(\frac{d!}{(dM_k + L_d)D^{d+1}} \right)^{\frac{1}{d}}$.

By Lemma 3, we know that

$$t_{\max} := \sup_{\|x - x^*\| \leq D} t(x)$$

is finite. Letting $M_{\max} := \max\{M, \frac{(d+1)t_{\max}}{d}\}$, we have $M_k \leq M_{\max}$, and therefore $C_k \geq \frac{d}{d+1} \left(\frac{d!}{(dM_{\max} + L_d)D^{d+1}} \right)^{\frac{1}{d}}$ for all k . Continuing the argument from the proof of Theorem 2 from [40], we can conclude that

$$f(x_k) - f(x^*) \leq \frac{(dM_{\max} + L_d)D^{d+1}}{d!} \left(\frac{d+1}{k} \right)^d.$$

Thus, we have $f(x_k) - f(x^*) \rightarrow 0$ and therefore $x_k \rightarrow x^*$.

For the local superlinear convergence rate, it suffices to show that for x_k close enough to x^* , we have

$$\|x_{k+1} - x^*\| \leq c' \|x_k - x^*\|^d$$

for some constant c' . Let r_1 and r_2 be as in the proof of Theorem 4, $r' := \min\{r_1, r_2\}$, and $\Omega := \{x \in \mathbb{R}^n \mid \|x - x^*\| \leq r'\}$. By the arguments in the proof of Theorem 4, for every $x_k \in \Omega$, we have

$$\begin{aligned} \|\nabla \psi_{x_k, d}(x^*)\| &\leq \frac{L_d}{d!} \|x^* - x_k\|^d + \max \left\{ \frac{dM}{(d+1)!}, t(x_k) \right\} (d+1) \|x^* - x_k\|^d \\ &\leq \frac{L_d}{d!} \|x^* - x_k\|^d + \max \left\{ \frac{dM}{(d+1)!}, \sup_{x \in \Omega} t(x) \right\} (d+1) \|x^* - x_k\|^d. \end{aligned}$$

Substituting into (8) (with x_0 replaced with x_k), we have

$$\|x_{k+1} - x^*\| \leq c' \|x^* - x_k\|^d,$$

where

$$c' := \frac{4((d-1)^2 - 1)}{\lambda_{\min} \nabla^2 f(x^*)} \left(\frac{L_d}{d!} + \max \left\{ \frac{dM}{(d+1)!}, \sup_{x \in \Omega} t(x) \right\} (d+1) \right).$$

We note that by Lemma 3, $\sup_{x \in \Omega} t(x)$ is finite. \square

7 Future directions

Besides the question of extending the results of Section 6 to the case of d even, there are a few other potential directions for future research that we wish to highlight:

- Can we replace the SDPs used in Algorithm 1 with more scalable conic programs such as linear programs (LPs) or second-order cone programs (SOCPs)? There has been work (see, e.g., [4]) on replacing methods based on sos programming with LP or SOCP-based approaches that rely on more tractable subsets of sos polynomials, such as the so-called *diagonally dominant sum of squares* (dsos) or *scaled diagonally dominant sum of squares* (sdsos) polynomials. In our setting, we might wish to replace the constraint in (3) (or (5)) that a polynomial is sos-convex with a constraint that it is “dsos-convex” or “sdsos-convex” (see, e.g., [3]). The results in [3] on the difference of dsos-convex decompositions of arbitrary polynomials could be explored to potentially replace the first SDP in each iteration of Algorithm 1 with an LP or SOCP. One would then need to establish an appropriate dsos or sdsos version of Theorem 2 to replace our second SDP with an LP or SOCP. It would be interesting to compare the factor of convergence of such an algorithm to that of the SDP-based approach.
- Can we create a method that uses a sparse subset of higher-order derivatives of the function f and that perhaps approximates the remaining derivatives in order to speed up each iteration? Such a method would be a higher-order analogue to the so-called “quasi-Newton” methods which rely on approximations of the Hessian of f (see, e.g., [42, Chap. 6]). An example of such a higher-order quasi-Newton method which results in semidefinite programs of small size in each iteration has been proposed in [2], but its convergence properties are currently unknown.
- Can we use our method or a modification thereof to solve systems of nonlinear equations (in a way that is superior to simply minimizing the sum of the squares of the equations)? The classical Newton method and its variants can be used for this purpose (see, e.g., [42, Sect. 11.1]). What are the right higher-order analogues of these approaches?
- Each iteration of the algorithms that we have presented in this paper can be interpreted as running just one iteration of the so-called “convex-concave procedure” (see, e.g., [34]) to a particular difference of convex decomposition of the Taylor expansion of f . Are there benefits of working with alternative difference of convex decompositions (see, e.g., [3]) of the Taylor expansion, or running more iterations of the convex-concave procedure before the Taylor polynomial is updated?

Acknowledgements

We would like to thank Jean-Bernard Lasserre for insightful discussions around the results in [32].

References

- [1] N. Agarwal and E. Hazan. Lower bounds for higher-order convex optimization. In *Proceedings of the 31st Conference On Learning Theory*, volume 75 of *Proceedings of Machine Learning Research*, pages 774–792, 2018.
- [2] A. A. Ahmadi, C. Dibek, and G. Hall. Sums of separable and quadratic polynomials. *Mathematics of Operations Research*, 48, 2022.
- [3] A. A. Ahmadi and G. Hall. DC decomposition of nonconvex polynomials with algebraic techniques. *Mathematical Programming*, 169(1):69–94, 2018.
- [4] A. A. Ahmadi and A. Majumdar. DSOS and SDSOS optimization: More tractable alternatives to sum of squares and semidefinite optimization. *SIAM Journal on Applied Algebra and Geometry*, 3(2):193–230, 2019.
- [5] A. A. Ahmadi, A. Olshevsky, P. A. Parrilo, and J. N. Tsitsiklis. NP-hardness of deciding convexity of quartic polynomials and related problems. *Mathematical Programming*, 137:453–476, 2013.
- [6] A. A. Ahmadi and P. A. Parrilo. A complete characterization of the gap between convexity and sos-convexity. *SIAM Journal on Optimization*, 23(2):811–833, 2013.
- [7] A. A. Ahmadi and J. Zhang. Complexity aspects of local minima and related notions. *Advances in Mathematics*, 397:108119, 2022.
- [8] A. A. Ahmadi and J. Zhang. On the complexity of finding a local minimizer of a quadratic function over a polytope. *Mathematical Programming*, 195(1-2):783–792, 2022.
- [9] M. Baes. Estimate sequence methods: extensions and approximations. *Institute for Operations Research, ETH, Zürich, Switzerland*, 2(1), 2009.
- [10] E. G. Belousov and D. Klatte. A Frank–Wolfe type theorem for convex polynomial programs. *Computational Optimization and Applications*, 22(1):37–48, 2002.
- [11] E. G. Birgin, J. L. Gardenghi, J. M. Martínez, S. A. Santos, and P. L. Toint. Worst-case evaluation complexity for unconstrained nonlinear optimization using high-order regularized models. *Mathematical Programming*, 163(1):359–368, 2017.

- [12] S. Bubeck, Q. Jiang, Y. T. Lee, Y. Li, and A. Sidford. Near-optimal method for highly smooth convex optimization. In *Conference on Learning Theory*, pages 492–507. Proceedings of Machine Learning Research, 2019.
- [13] Y. Carmon, J. C. Duchi, O. Hinder, and A. Sidford. Lower bounds for finding stationary points i. *Mathematical Programming*, 184(1):71–120, 2020.
- [14] C. Cartis, N. I. Gould, and P. L. Toint. Universal regularization methods: varying the power, the smoothness and the accuracy. *SIAM Journal on Optimization*, 29(1):595–615, 2019.
- [15] C. Cartis, N. I. Gould, and P. L. Toint. A concise second-order complexity analysis for unconstrained optimization using high-order regularized models. *Optimization Methods and Software*, 35(2):243–256, 2020.
- [16] C. Cartis, N. I. Gould, and P. L. Toint. Sharp worst-case evaluation complexity bounds for arbitrary-order nonconvex optimization with inexpensive constraints. *SIAM Journal on Optimization*, 30(1):513–541, 2020.
- [17] C. Cartis, N. I. Gould, and P. L. Toint. *Evaluation Complexity of Algorithms for Nonconvex Optimization: Theory, Computation and Perspectives*. SIAM, 2022.
- [18] C. Cartis and W. Zhu. Second-order methods for quartically-regularised cubic polynomials, with applications to high-order tensor methods. *arXiv preprint arXiv:2308.15336*, 2023.
- [19] C. Cartis and W. Zhu. Global convergence of high-order regularization methods with sums-of-squares Taylor models. *arXiv preprint arXiv:2404.03035*, 2024.
- [20] P. L. Chebyshev. Polnoe Sobranie Sochinenii. *Izd. Akad. Nauk SSSR*, 5:7–25, 1951.
- [21] C. W. Clenshaw and A. R. Curtis. A method for numerical integration on an automatic computer. *Numerische Mathematik*, 2(1):197–205, 1960.
- [22] A. Conn, N. Gould, and P. Toint. *Trust Region Methods*. MPS-SIAM Series on Optimization. Society for Industrial and Applied Mathematics, 2000.
- [23] N. Doikov. *New second-order and tensor methods in convex optimization*. PhD thesis, Université catholique de Louvain, 2021.
- [24] N. Doikov and Y. Nesterov. Local convergence of tensor methods. *Mathematical Programming*, 193(1):315–336, 2022.
- [25] G. N. Grapiglia and Y. Nesterov. Tensor methods for finding approximate stationary points of convex functions. *Optimization Methods and Software*, 37(2):605–638, 2022.

- [26] J. W. Helton and J. Nie. Semidefinite representation of convex sets. *Mathematical Programming*, 122:21–64, 2010.
- [27] J. P. Imhof. On the method for numerical integration of Clenshaw and Curtis. *Numerische Mathematik*, 5(1):138–141, 1963.
- [28] B. Jiang, T. Lin, and S. Zhang. A unified adaptive tensor approximation scheme to accelerate composite convex optimization. *SIAM Journal on Optimization*, 30(4):2897–2926, 2020.
- [29] B. Jiang, H. Wang, and S. Zhang. An optimal high-order tensor method for convex optimization. *Mathematics of Operations Research*, 46(4):1390–1412, 2021.
- [30] J.-B. Lasserre. Global optimization with polynomials and the problem of moments. *SIAM Journal on Optimization*, 11:796–817, 2000.
- [31] J.-B. Lasserre. Representation of nonnegative convex polynomials. *Archiv der Mathematik*, 91(2):126–130, 2008.
- [32] J.-B. Lasserre. Convexity in semialgebraic geometry and polynomial optimization. *SIAM Journal on Optimization*, 19:1995–2014, 2009.
- [33] K. Levenberg. Method for the solution of certain problems in least squares. *J Numer Anal*, 16:588–A604, 1944.
- [34] T. Lipp and S. Boyd. Variations and extension of the convex–concave procedure. *Optimization and Engineering*, 17(2):263–287, 2016.
- [35] J. Löfberg. YALMIP: A toolbox for modeling and optimization in MATLAB. In *IEEE International Conference on Robotics and Automation*, pages 284–289, 2004.
- [36] A. Majumdar, G. Hall, and A. A. Ahmadi. Recent scalability improvements for semidefinite programming with applications in machine learning, control, and robotics. *Annual Review of Control, Robotics, and Autonomous Systems*, 3:331–360, 2020.
- [37] D. W. Marquardt. An algorithm for least-squares estimation of nonlinear parameters. *Journal of the Society for Industrial and Applied Mathematics*, 11(2):431–441, 1963.
- [38] J. J. Moré. *Recent Developments in Algorithms and Software for Trust Region Methods*, pages 258–287. Springer Berlin Heidelberg, 1983.
- [39] K. G. Murty and S. N. Kabadi. Some NP-complete problems in quadratic and nonlinear programming. *Mathematical Programming*, 39(2):117–129, 1987.
- [40] Y. Nesterov. Implementable tensor methods in unconstrained convex optimization. *Mathematical Programming*, 186(1):157–183, 2021.

- [41] Y. Nesterov and B. T. Polyak. Cubic regularization of Newton method and its global performance. *Mathematical Programming*, 108(1):177–205, 2006.
- [42] J. Nocedal and S. J. Wright. *Numerical Optimization*. Springer, 2006.
- [43] J. M. Ortega and W. C. Rheinboldt. *Iterative Solution of Nonlinear Equations in Several Variables*. SIAM, 2000.
- [44] P. A. Parrilo. *Structured semidefinite programs and semialgebraic geometry methods in robustness and optimization*. PhD thesis, California Institute of Technology, 2000.
- [45] S. Prajna, A. Papachristodoulou, and P. A. Parrilo. Introducing SOS-TOOLS: A general purpose sum of squares programming solver. In *Proceedings of the 41st IEEE Conference on Decision and Control*, volume 1, pages 741–746, 2002.
- [46] I. Pólik and T. Terlaky. A survey of the S-lemma. *SIAM Review*, 49(3):371–418, 2007.
- [47] O. Silina and J. Zhang. An unregularized third order Newton method. *arXiv preprint arXiv:2209.10051*, 2022.
- [48] L. Vandenberghe and S. Boyd. Semidefinite programming. *SIAM Review*, 38(1):49–95, 1996.
- [49] A. Yurtsever, J. A. Tropp, O. Fercoq, M. Udell, and V. Cevher. Scalable semidefinite programming. *SIAM Journal on Mathematics of Data Science*, 3(1):171–200, 2021.