# On the Complexity of Finding a Local Minimizer of a Quadratic Function over a Polytope

Amir Ali Ahmadi[*] and Jeffrey Zhang[†‡]

**Abstract**

We show that unless P=NP, there cannot be a polynomial-time algorithm that finds a point within Euclidean distance $c^n$ (for any constant $c \geq 0$) of a local minimizer of an $n$-variate quadratic function over a polytope. This result (even with $c = 0$) answers a question of Pardalos and Vavasis that appeared in 1992 on a list of seven open problems in complexity theory for numerical optimization. Our proof technique also implies that the problem of deciding whether a quadratic function has a local minimizer over an (unbounded) polyhedron, and that of deciding if a quartic polynomial has a local minimizer are NP-hard.

**Keywords:** Local minimizers, quadratic programs, computational complexity, polynomial optimization.

## 1 Introduction

Recall that a *local minimizer* of a function $f : \mathbb{R}^n \to \mathbb{R}$ over a set $\Omega \subseteq \mathbb{R}^n$ is a point $\bar{x} \in \Omega$ for which there exists a scalar $\epsilon > 0$ such that $p(\bar{x}) \leq p(x)$ for all $x \in \Omega$ with $\|x - \bar{x}\| \leq \epsilon$. In the case where $f$ is a linear function and $\Omega$ is polyhedral (i.e., the case of linear programming), it is well known that a local minimizer (which also has to be a global minimizer) can be found in polynomial time in the Turing model of computation [7, 6]. Perhaps the next simplest constrained optimization problem to consider is one where $f$ is a quadratic function and $\Omega$ is polyhedral. Such an optimization problem is known as a *quadratic program* and can be written as

$$\begin{array}{ll} \min\limits_{x \in \mathbb{R}^n} & x^T Q x + c^T x \\ \text{subject to} & a_i^T x \leq b_i, \forall i \in \{1, \ldots, m\}, \end{array} \tag{1}$$

where $Q \in \mathbb{R}^{n \times n}$, $c, a_1, \ldots, a_m \in \mathbb{R}^n$, and $b_1, \ldots, b_m \in \mathbb{R}$. The matrix $Q$ is taken without loss of generality to be symmetric. When complexity questions about quadratic programs are studied in the Turing model of computation, all these data are rational and the input size is the total number of bits required to write them down. When the matrix $Q$ is positive semidefinite, then we recover the case of convex quadratic programming, where it is well known that finding a global minimizer can be done in polynomial time [8]. However, when $Q$ has even a single negative eigenvalue, finding a global minimizer is NP-hard [13]. It is therefore natural to ask whether one can instead find a local minimizer of a general quadratic program efficiently. In fact, this precise question appeared in 1992 on a list of seven open problems in complexity theory for numerical optimization [14]:

[*] Amir Ali Ahmadi is with the department of Operations Research and Financial Engineering at Princeton University. Email: aaa@princeton.edu.

[†] Jeffrey Zhang is with the department of Mathematical Sciences at Carnegie Mellon University. Email: jeffz@cmu.edu.

[‡] This work was partially supported by an AFOSR MURI award, the DARPA Young Faculty Award, the Princeton SEAS Innovation Award, the NSF CAREER Award, the Google Faculty Award, and the Sloan Fellowship.

*"What is the complexity of finding even a local minimizer for nonconvex quadratic programming, assuming the feasible set is compact? Murty and Kabadi (1987, [11]) and Pardalos and Schnitger (1988, [12]) have shown that it is NP-hard to test whether a given point for such a problem is a local minimizer, but that does not rule out the possibility that another point can be found that is easily verified as a local minimizer."*

A few remarks on the phrasing of this problem are in order. First, note that in this question, the feasible set of the quadratic program is assumed to be compact (i.e., a polytope). Therefore, there is no need to focus on the related and often prerequisite problem of deciding the *existence* of a local minimizer (since any global minimizer e.g. is a local minimizer). The latter question makes sense in the case where the feasible set of the quadratic program is unbounded; the complexity of this question is also addressed in our paper (Theorem 2.5). Second, as the quote points out, the question of finding a local minimizer is also separate from a complexity viewpoint from that of testing if a given point is a local minimizer. This related question has been studied more extensively and its complexity has already been settled for optimization problems whose objective and constraints are given by polynomial functions of any degree; see [11, 12, 2].

To point out some of the subtle differences between these variations of the problem more specifically, we briefly review the reduction of Murty and Kabadi [11], which shows the NP-hardness of deciding if a given point is a local minimizer of a quadratic program. In [11], the authors show that the problem of deciding if a symmetric matrix $Q$ is *copositive*—i.e. whether $x^TQx \geq 0$ for all vectors $x$ in the nonnegative orthant—is NP-hard. From this, it is straightforward to observe that the problem of testing whether a given point is a local minimizer of a quadratic function over a polyhedron is NP-hard: Indeed, the origin is a local minimizer of $x^TQx$ over the nonnegative orthant if and only if the matrix $Q$ is copositive. However, it is not true that $x^TQx$ has a local minimizer over the nonnegative orthant if and only if $Q$ is copositive. Although the "if" direction holds, the "only if" direction does not. For example, consider the matrix

$$Q = \begin{bmatrix} 0 & 1 \\ 1 & -2 \end{bmatrix},$$

which is clearly not copositive, even though the point $(1,0)^T$ is a local minimizer of $x^TQx$ over the nonnegative orthant.

Our main results in this paper are as follows. We show that unless P=NP, no polynomial-time algorithm can find a point within Euclidean distance $c^n$ (for any constant $c \geq 0$) of a local minimizer of an $n$-variate quadratic program with a compact feasible set (Theorem 2.6). See also Corollaries 2.7 and 2.8. To prove this, we show as an intermediate step that deciding whether a quartic polynomial or a quadratic program has a local minimizer is strongly NP-hard[1] (Theorems 2.1 and 2.5). Finally, we show that unless P=NP, there cannot be a polynomial-time algorithm that decides if a quadratic program with a compact feasible set has a unique local minimizer and if so returns this minimizer (Theorem 2.9).

Overall, our results suggest that without additional problem structure, questions related to finding local minimizers of quadratic programs are not easier (at least from a complexity viewpoint) than those related to finding global minimizers. It also suggests that any efficient heuristic that aims to find a local minimizer of a quadratic program must necessarily fail on a "significant portion" of instances; see e.g. Corollary 2.2 of [5] for a more formal complexity theoretic statement.

[1]This implies that these problems remain NP-hard even if the bitsize of all numerical data are $O(\log n)$, where $n$ is the number of variables. For a strongly NP-hard problem, even a pseudo-polynomial time algorithm—i.e., an algorithm whose running time is polynomial in the magnitude of the numerical data of the problem but not necessarily in their bitsize—cannot exist unless P=NP. See [4] or [1, Section 2] for more details.

### 1.1 Notation and Basic Definitions

For a vector $x \in \mathbb{R}^n$, the notation $x^2$ denotes the vector in $\mathbb{R}^n$ whose $i$-th entry is $x_i^2$, and $Diag(x)$ denotes the diagonal $n \times n$ matrix whose $i$-th diagonal entry is $x_i$. The notation $x \geq 0$ denotes that the vector $x$ belongs to the nonnegative orthant, and for such a vector, $\sqrt{x}$ denotes the vector in $\mathbb{R}^n$ whose $i$-th entry is $\sqrt{x_i}$. For two matrices $X, Y \in \mathbb{R}^{m\times n}$, we denote by $X \odot Y$ the matrix in $\mathbb{R}^{m\times n}$ whose $(i,j)$-th entry is $X_{ij}Y_{ij}$. For vectors $x, y \in \mathbb{R}^n$, the notation $y_x$ (sometimes $(y)_x$ if there is room for confusion with other indices) denotes the vector containing the entries of $y$ where $x_i$ is nonzero in the same order (the length of $y_x$ is hence equal to the number of nonzero entries in $x$). Similarly, for a vector $x \in \mathbb{R}^n$ and a matrix $Y \in \mathbb{R}^{n\times n}$, the notation $Y_x$ (sometimes $(Y)_x$ if there is room for confusion with other indices) denotes the principal submatrix of $Y$ consisting of rows and columns of $Y$ whose indices correspond to indices of nonzero entries of $x$. The notation $I$ (resp. $J$) refers to the identity matrix (resp. the matrix of all ones); the dimension will be clear from context. For a symmetric matrix $M \in \mathbb{R}^{n\times n}$, the notation $M \succeq 0$ (resp. $M \succ 0$) denotes that $M$ is *positive semidefinite* (resp. *positive definite*), i.e. that it has nonnegative (resp. positive) eigenvalues. As mentioned already, we say that $M$ is *copositive* if $x^T Mx \geq 0, \forall x \geq 0$. The *simplex* in $\mathbb{R}^n$ is denoted by $\Delta_n := \{x \in \mathbb{R}^n \mid x \geq 0, \sum_{i=1}^n x_i = 1\}$. Finally, for a scalar $c$, the notation $\lceil c \rceil$ denotes the ceiling of $c$, i.e. the smallest integer greater than or equal to $c$.

We recall that a *form* is a homogeneous polynomial; i.e. a polynomial whose monomials all have the same degree. A form $p : \mathbb{R}^n \to \mathbb{R}$ is said to be *nonnegative* if $p(x) \geq 0, \forall x \in \mathbb{R}^n$, and *positive definite* if $p(x) > 0, \forall x \neq 0$. A *critical point* of a differentiable function $f : \mathbb{R}^n \to \mathbb{R}$ is a point $x \in \mathbb{R}^n$ at which the gradient $\nabla f(x)$ is zero. A *second-order point* of a twice-differentiable function $f : \mathbb{R}^n \to \mathbb{R}$ is a critical point $x$ at which the Hessian matrix $\nabla^2 f(x)$ is positive semidefinite.

All graphs in this paper are undirected, unweighted, and have no self-loops. The *adjacency matrix* of a graph $G$ on $n$ vertices is the $n \times n$ symmetric matrix whose $(i,j)$-th entry equals one if vertices $i$ and $j$ share an edge in $G$ and zero otherwise. The *complement* of a graph $G$, denoted by $\bar{G}$, is the graph with the same vertex set as $G$ and such that two distinct vertices are adjacent if and only if they are not adjacent in $G$. An *induced subgraph* of $G$ is a graph containing a subset of the vertices of $G$ and all edges connecting pairs of vertices in that subset.

## 2 The Main Result

### 2.1 Complexity of Deciding Existence of Local Minimizers

To show that a polynomial-time algorithm for finding a local minimizer of a quadratic function over a polytope (i.e., a bounded polyhedron) implies P = NP (Theorem 2.6), we first show that it is NP-hard to decide whether a quartic polynomial has a local minimizer (Theorem 2.1). The main idea of this proof is used to show the same claim about deciding whether a quadratic program has a local minimizer (Theorem 2.5) and to then establish our main result (Theorem 2.6).

**Theorem 2.1.** *It is strongly NP-hard to decide if a degree-4 polynomial has a local minimizer.*

We will prove this theorem by presenting a polynomial-time reduction from the STABLESET problem, which is known to be (strongly) NP-hard [4]. Recall that in the STABLESET problem, we are given as input a graph $G$ on $n$ vertices and a positive integer $r \leq n$. We are then asked to decide whether $G$ has a *stable set* of size $r$, i.e. a set of $r$ pairwise non-adjacent vertices. We denote the size of the largest stable set in a graph $G$ by the standard notation $\alpha(G)$. We also recall that a *clique* in a graph $G$ is a set of pairwise adjacent vertices. The size of the largest clique in $G$ is denoted by $\omega(G)$. The following theorem of Motzkin and Straus [10] relates $\omega(G)$ to the optimal value of a quadratic program.

**Theorem 2.2** ([10]). *Let $G$ be a graph on $n$ vertices with adjacency matrix $A$ and clique number $\omega$. The optimal value of the quadratic program*

$$\begin{aligned} \max_{x\in\mathbb{R}^n} \quad & x^T A x \\ \text{subject to} \quad & x \geq 0, \\ & \sum_{i=1}^n x_i = 1 \end{aligned} \tag{2}$$

*is* $1 - \frac{1}{\omega}$.

For a scalar $k$ and a symmetric matrix $A$ (which will always be an adjacency matrix), the following notation will be used repeatedly in our proofs:

$$M_{A,k} := kA + kI - J, \tag{3}$$

$$q_{A,k}(x) := x^T M_{A,k} x, \tag{4}$$

and

$$p_{A,k}(x) := (x^2)^T M_{A,k} x^2. \tag{5}$$

Note that nonnegativity of the quadratic form $q_{A,k}$ over the nonnegative orthant is equivalent to (global) nonnegativity of the quartic form $p_{A,k}$ and to copositivity of the matrix $M_{A,k}$.

The following corollary of Theorem 2.2 will be of more direct relevance to our proofs. The first statement in the corollary has been observed e.g. by de Klerk and Pashechnik [3], but its proof is included here for completeness. The second statement, which will also be needed in the proof of Theorem 2.1, follows straightforwardly.

**Corollary 2.3.** *For a scalar $k > 0$ and a graph $G$ with adjacency matrix $A$, the matrix $M_{A,k}$ in (3) is copositive if and only if $\alpha(G) \leq k$. Furthermore, if $\alpha(G) < k$, the quartic form $p_{A,k}$ in (5) is positive definite.*[2]

*Proof.* First observe that $\alpha(G) = \omega(\bar{G})$, and that the adjacency matrix of $\bar{G}$ is $J - A - I$. Thus from Theorem 2.2, the maximum value of $x^T(J - A - I)x$ over $\Delta_n$ is $1 - \frac{1}{\alpha(G)}$, and hence the minimum value of $x^T(A + I)x$ over $\Delta_n$ is $\frac{1}{\alpha(G)}$. Therefore, for any $k > 0$, $\alpha(G) \leq k$ if and only if $x^T(A + I)x \geq \frac{1}{k}$ for all $x \in \Delta_n$, which holds if and only if $x^T(k(A + I) - J)x \geq 0$ for all $x \in \Delta_n$. The first statement of the corollary then follows from the homogeneity of $x^T(k(A + I) - J)x$.

To show that $p_{A,k}$ is positive definite when $\alpha(G) < k$, observe that

$$kA + kI - J = (\alpha(G)(A + I) - J) + (k - \alpha(G))(A + I).$$

Considering the two terms on the right separately, we observe that $(x^2)^T(\alpha(G)(A+I) - J)x^2$ (i.e., $p_{A,\alpha(G)}$) is nonnegative since $M_{A,\alpha(G)}$ is copositive, and that $(x^2)^T(k - \alpha(G))(A + I)x^2$ is positive definite. Therefore, their sum $(x^2)^T(kA + kI - J)x^2$ is positive definite. □

We now present the proof of Theorem 2.1. While the statement of the theorem is given for degree-4 polynomials, it is straightforward to extend the result to higher-degree polynomials. We note that degree four is the smallest degree for which deciding existence of local minimizers is intractable. For degree-3 polynomials, it turns out that this question can be answered by solving semidefinite programs of polynomial size [2].

[2]The converse of this statement also holds, but we do not need it for the proof of Theorem 2.1.

*Proof (of Theorem 2.1).* We present a polynomial-time reduction from the STABLESET problem. Let a graph $G$ on $n$ vertices with adjacency matrix $A$ and a positive integer $r \leq n$ be given. We show that $G$ has a stable set of size $r$ if and only if the quartic form $p_{A,r-0.5}$ defined in (5) has no local minimizer. This is a consequence of the following more general fact that we prove below: For a noninteger scalar $k$, the quartic form $p_{A,k}$ has no local minimizer if and only if $\alpha(G) \geq k$.

We first observe that if $\alpha(G) < k$, then $p_{A,k}$ has a local minimizer. Indeed, recall from the second claim of Corollary 2.3 that under this assumption, $p_{A,k}$ is positive definite. Since $p_{A,k}$ vanishes at the origin, it follows that the origin is a local minimizer. Suppose now that $\alpha(G) \geq k$. Since $k$ is noninteger, this implies that $\alpha(G) > k$. We show that in this case, $p_{A,k}$ has no local minimizer by showing that the origin must be the only second-order point of $p_{A,k}$. Since any local minimizer of a polynomial is a second-order point, only the origin can be a candidate local minimizer for $p_{A,k}$. However, by the first claim of Corollary 2.3, the matrix $M_{A,k}$ is not copositive and hence $p_{A,k}$ is not nonnegative. As $p_{A,k}$ is homogeneous, this implies that $p_{A,k}$ takes negative values arbitrarily close to the origin, ruling out the possibility of the origin being a local minimizer.

To show that when $\alpha(G) > k$, the origin is the only second-order point of $p_{A,k}$, we compute the gradient and Hessian of $p_{A,k}$. We have

$$\nabla p_{A,k}(x) = 4x \odot M_{A,k}x^2,$$

and

$$\nabla^2 p_{A,k}(x) = 8M_{A,k} \odot xx^T + 4Diag(M_{A,k}x^2).$$

Suppose for the sake of contradiction that $p_{A,k}$ has a nonzero second-order point $\bar{x}$. Since $\bar{x}$ is a critical point, $\nabla p_{A,k}(\bar{x}) = 0$ and thus $(M_{A,k}\bar{x}^2)_{\bar{x}} = 0$. It then follows that

$$(\nabla^2 p_{A,k}(\bar{x}))_{\bar{x}} = 8(M_{A,k})_{\bar{x}} \odot \bar{x}_{\bar{x}}\bar{x}_{\bar{x}}^T.$$

Because $\nabla^2 p_{A,k}(\bar{x}) \succeq 0$ and thus all its principal submatrices are positive semidefinite, we have $8(M_{A,k})_{\bar{x}} \odot \bar{x}_{\bar{x}}\bar{x}_{\bar{x}}^T \succeq 0$. Since

$$(M_{A,k})_{\bar{x}} \odot \bar{x}_{\bar{x}}\bar{x}_{\bar{x}}^T = Diag(\bar{x}_{\bar{x}})(M_{A,k})_{\bar{x}}Diag(\bar{x}_{\bar{x}}),$$

and since $Diag(\bar{x}_{\bar{x}})$ is an invertible matrix, it follows that $(M_{A,k})_{\bar{x}} \succeq 0$.

We now consider the induced subgraph $G_{\bar{x}}$ of $G$ with vertices corresponding to the indices of the nonzero entries of $\bar{x}$. Note that the adjacency matrix of $G_{\bar{x}}$ is $A_{\bar{x}}$. Furthermore, observe that $M_{A_{\bar{x}},k} = (M_{A,k})_{\bar{x}}$, and therefore $M_{A_{\bar{x}},k}$ is positive semidefinite and thus copositive. We conclude from the first claim of Corollary 2.3 that $\alpha(G_{\bar{x}}) \leq k$. We now claim that

$$M_{A_{\bar{x}},k}\bar{x}_{\bar{x}}^2 = (M_{A,k})_{\bar{x}}\bar{x}_{\bar{x}}^2 = (M_{A,k}\bar{x}^2)_{\bar{x}} = 0.$$

The first equality follows from that $M_{A_{\bar{x}},k} = (M_{A,k})_{\bar{x}}$, the second from that the indices of the nonzero entries of $\bar{x}$ are the same as those of $\bar{x}^2$, and the third from that $\nabla p_{A,k}(\bar{x}) = 0$. Hence, $p_{A_{\bar{x}},k}(\bar{x}_{\bar{x}}) = 0$. Since $\bar{x}_{\bar{x}}$ is nonzero, $p_{A_{\bar{x}},k}$ is not positive definite. By the second claim of Corollary 2.3, we must have $\alpha(G_{\bar{x}}) \geq k$. Therefore, $\alpha(G_{\bar{x}}) = k$. However, because $k$ was assumed to be noninteger, we have a contradiction. □

It turns out that the proof of Theorem 2.1 also shows that it is NP-hard to decide if a quartic polynomial has a strict local minimizer. Recall that a *strict local minimizer* of a function $f : \mathbb{R}^n \to \mathbb{R}$ over a set $\Omega \subseteq \mathbb{R}^n$ is a point $\bar{x} \in \Omega$ for which there exists a scalar $\epsilon > 0$ such that $p(\bar{x}) < p(x)$ for all $x \in \Omega \backslash \bar{x}$ with $\|x - \bar{x}\| \leq \epsilon$.

**Corollary 2.4.** *It is strongly NP-hard to decide if a degree-4 polynomial has a strict local minimizer.*

*Proof.* Observe from the proof of Theorem 2.1 that for a graph $G$ and a noninteger scalar $k$, the quartic form $p_{A,k}$ has a local minimizer if and only if $\alpha(G) < k$. In the case where $p_{A,k}$ does have a local minimizer, we showed that $p_{A,k}$ is positive definite, and thus the local minimizer (the origin) must be a strict local minimizer. □

We now turn our attention to local minimizers of quadratic programs.

**Theorem 2.5.** *It is strongly NP-hard to decide if a quadratic function has a local minimizer over a polyhedron. The same is true for deciding if a quadratic function has a strict local minimizer over a polyhedron.*

*Proof.* We present a polynomial-time reduction from the STABLESET problem to the problem of deciding if a quadratic function has a local minimizer over a polyhedron. The reader can check that same reduction is valid for the case of strict local minimizers.

Let a graph $G$ on $n$ vertices with adjacency matrix $A$ and a positive integer $r \leq n$ be given. Let $k = r - 0.5$, $q_{A,k}$ be the quadratic form defined in (4), and consider the optimization problem

$$\begin{array}{ll} \min\limits_{x \in \mathbb{R}^n} & q_{A,k}(x) \\ \text{subject to} & x \geq 0. \end{array} \tag{6}$$

We show that a point $x \in \mathbb{R}^n$ is a local minimizer of (6) if and only if $\sqrt{x}$ is a local minimizer of the quartic form $p_{A,k}$ defined in (5). By the arguments in the proof of Theorem 2.1, we would have that (6) has no local minimizer if and only if $G$ has a stable set of size $r$.

Indeed, if $x$ is not a local minimizer of (6), there exists a sequence $\{y_j\} \subseteq \mathbb{R}^n$ with $y_j \to x$ and such that for all $j$, $y_j \geq 0$ and $q_{A,k}(y_j) < q_{A,k}(x)$. The sequence $\{\sqrt{y_j}\}$ would then satisfy $p_{A,k}(\sqrt{y_j}) < p_{A,k}(\sqrt{x})$ and $\sqrt{y_j} \to \sqrt{x}$, proving that $\sqrt{x}$ is not a local minimizer of $p_{A,k}$. Similarly, if $x$ is not a local minimizer of $p_{A,k}$, there exists a sequence $\{z_j\} \subseteq \mathbb{R}^n$ such that $z_j \to x$ and $p_{A,k}(z_j) < p_{A,k}(x)$ for all $j$. The sequence $\{z_j^2\}$ would then prove that $x^2$ is not a local minimizer of (6).

□

*Remark* 2.1. The decision problems that Theorems 2.1 and 2.5 prove to be NP-hard are unlikely to be in NP. Indeed, membership of these problems in NP would imply (via our reductions) that STABLESET is in co-NP, which would further imply that NP=co-NP.

## 2.2 Complexity of Finding a Local Minimizer of a Quadratic Function Over a Polytope

We now address the original question posed by Pardalos and Vavasis concerning the complexity of finding a local minimizer of a quadratic program with a compact feasible set. Note again that if the feasible set is compact, the existence of a local minimizer is guaranteed. In fact, there will always be a local minimizer that has rational entries with polynomial bitsize [15].

**Theorem 2.6.** *If there is a polynomial-time algorithm that finds a point within Euclidean distance $c^n$ (for any constant $c \geq 0$) of a local minimizer of an $n$-variate quadratic function over a polytope, then $P = NP$.*

*Proof.* Fix any constant $c \geq 0$. We show that if an algorithm could take as input a quadratic program with a bounded feasible set and in polynomial time return a point within distance $c^n$ of any local minimizer, then this algorithm would solve the STABLESET problem in polynomial time.

Let a graph $G$ on $n$ vertices with adjacency matrix $A$ and a positive integer $r \leq n$ be given. Let $k = r - 0.5$, $q_{A,k}$ be the quadratic form defined in (4), and consider the quadratic program

$$\begin{aligned} \min_{x \in \mathbb{R}^n} \quad & q_{A,k}(x) \\ \text{subject to} \quad & x \geq 0, \\ & \sum_{i=1}^{n} x_i \leq \lceil 3c^n \sqrt{n} \rceil. \end{aligned} \tag{7}$$

Note that the feasible set of this problem is bounded. Moreover, the number of bits required to write down this quadratic program is polynomial in $n$. Indeed, the scalar $\lceil 3c^n \sqrt{n} \rceil$ takes at most $2 + n\lceil \log_2(c+1) \rceil + \frac{1}{2}\lceil \log_2(n+1) \rceil$ bits to write down, and the remaining $O(n^2)$ numbers in the problem data are bounded in magnitude by $n$, so they each take $O(\log_2(n))$ bits to write down.

We will show that if $\alpha(G) < k$, the origin is the unique local minimizer of (7), and that if $\alpha(G) \geq k$ (or equivalently $\alpha(G) > k$), any local minimizer $\bar{x}$ of (7) satisfies $\sum_{i=1}^{n} \bar{x}_i = \lceil 3c^n \sqrt{n} \rceil$. Since the (Euclidean) distance from the origin to the hyperplane $\{x \in \mathbb{R}^n | \sum_{i=1}^{n} x_i = \lceil 3c^n \sqrt{n} \rceil\}$ is at least $3c^n$, there is no point that is within distance $c^n$ of both the origin and this hyperplane. Thus, the graph $G$ has no stable set of size $r$ (or equivalently $\alpha(G) < k$) if and only if the Euclidean norm of all points within distance $c^n$ of any local minimizer of (7) is less than or equal to $c^n$.

To see why $\alpha(G) < k$ implies that the origin is the unique local minimizer of (7), recall from the second claim of Corollary 2.3 that the quartic form $p_{A,k}$ defined in (5) must be positive definite. Thus, for any nonzero vector $x \geq 0$, we have $q_{A,k}(x) > 0$. This implies that the origin is a local minimizer of (7). Moreover, since $q_{A,k}$ is homogeneous, we have that no other feasible point can be a local minimizer. Indeed, for any nonzero vector $x \geq 0$ and any nonnegative scalar $\epsilon < 1$, $q_{A,k}(\epsilon x) < q_{A,k}(x)$.

To see why when $\alpha(G) > k$, the last constraint of (7) must be tight at all local minimizers, recall from the proof of Theorem 2.5 that when $\alpha(G) > k$, the optimization problem in (6) has no local minimizer. Therefore, for any vector $x$ that is feasible to (7) and satisfies $\sum_{i=1}^{n} x_i < \lceil 3c^n \sqrt{n} \rceil$, there exists a sequence $\{y_i\} \subseteq \mathbb{R}^n$ with $y_i \to x$, and satisfying

$$y_i \geq 0, \sum_{i=1}^{n} y_i < \lceil 3c^n \sqrt{n} \rceil, q_{A,k}(y_i) < q_{A,k}(x), \forall i.$$

As the points $y_i$ are feasible to (7), any vector $x$ satisfying $\sum_{i=1}^{n} x_i < \lceil 3c^n \sqrt{n} \rceil$ cannot be a local minimizer of (7). Thus, if $\alpha(G) > k$, any local minimizer $\bar{x}$ of (7) satisfies $\sum_{i=1}^{n} \bar{x}_i = \lceil 3c^n \sqrt{n} \rceil$. □

By replacing the quantity $\lceil 3c^n \sqrt{n} \rceil$ in the proof of Theorem 2.6 with $\lceil 3n^{c+0.5} \rceil$ and $2n$ respectively, we get the following two corollaries.

**Corollary 2.7.** *If there is a pseudo-polynomial time algorithm that finds a point within Euclidean distance $n^c$ (for any constant $c \geq 0$) of a local minimizer of an $n$-variate quadratic function over a polytope, then $P = NP$.*

**Corollary 2.8.** *If there is a polynomial-time algorithm that finds a point within Euclidean distance $\epsilon\sqrt{n}$ (for any constant $\epsilon \in [0, 1)$) of a local minimizer of a restricted set of quadratic programs over $n$ variables whose numerical data are integers bounded in magnitude by $2n$, then $P = NP$.*

In [14], Pardalos and Vavasis also propose two follow-up questions about quadratic programs with compact feasible sets. The first is about the complexity of finding a "KKT point". As is, our proof does not have any implications for this question since the origin is always a KKT point of the quadratic programs that arise from our reductions. The second question asks whether finding a local minimizer is easier in the special case where the problem only has one local minimizer (which is thus also the global minimizer). Related to this question, we can prove the following claim.

**Theorem 2.9.** *If there is a polynomial-time algorithm which decides whether a quadratic program with a bounded feasible set has a unique local minimizer, and if so returns this minimizer*[3]*, then P=NP.*

*Proof.* Suppose there was such an algorithm (call it Algorithm U). We show that Algorithm U would solve the STABLESET problem in polynomial time. Let a graph $G$ on $n$ vertices with adjacency matrix $A$ and a positive integer $r \leq n$ be given, and input the quadratic program (7), with $k = r - 0.5$, into Algorithm U. Observe from the proof of Theorem 2.6 that there are three possibilities for this quadratic program: (i) the origin is the unique local minimizer, (ii) there is a unique local minimizer and it is on the hyperplane $\{x \in \mathbb{R}^n | \sum_{i=1}^n x_i = \lceil 3c^n \sqrt{n} \rceil\}$, and (iii) there are multiple local minimizers and they are all on the hyperplane $\{x \in \mathbb{R}^n | \sum_{i=1}^n x_i = \lceil 3c^n \sqrt{n} \rceil\}$. Case (i) indicates that $\alpha(G) < k$, and the output of Algorithm U in this case would be the origin. Cases (ii) and (iii) both indicate that $\alpha(G) > k$. The output of Algorithm U is a point away from the origin in case (ii), and the declaration that the local minimizer is not unique in case (iii). Thus Algorithm U would reveal which case we are in, and that would allow us to decide if $G$ has a stable set of size $r$ in polynomial time. □

To conclude, we have established intractability of several problems related to local minima of quadratic programs. We hope our results motivate more research on identifying classes of quadratic programs where local minimizers can be found more efficiently than global minimizers. One interesting example is the case of the concave knapsack problem, where Moré and Vavasis [9] show that a local minimizer can be found in polynomial time even though unless P=NP, a global minimizer cannot.

## References


[1] A. A. Ahmadi and G. Hall. On the complexity of detecting convexity over a box. *Mathematical Programming*, pages 1–15, 2019.

[2] A. A. Ahmadi and J. Zhang. Complexity aspects of local minima and related notions. *Available at https://arxiv.org/pdf/2008.06148.pdf.*

[3] E. De Klerk and D. V. Pasechnik. Approximation of the stability number of a graph via copositive programming. *SIAM Journal on Optimization*, 12(4):875–892, 2002.

[4] M. R. Garey and D. S. Johnson. *Computers and Intractability*, volume 29. WH Freeman New York, 2002.

[5] L. A. Hemaspaandra and R. Williams. SIGACT News Complexity Theory Column 76: an atypical survey of typical-case heuristic algorithms. *ACM SIGACT News*, 43(4):70–89, 2012.


[3] This unique local (and therefore global) minimizer is guaranteed to have rational entries with polynomial bitsize; see [15].

[6] N. Karmarkar. A new polynomial-time algorithm for linear programming. In *Proceedings of the Sixteenth Annual ACM Symposium on Theory of Computing*, pages 302–311, 1984.

[7] L. G. Khachiyan. A polynomial algorithm in linear programming. In *Doklady Akademii Nauk*, volume 244, pages 1093–1096. Russian Academy of Sciences, 1979.

[8] M. K. Kozlov, S. P. Tarasov, and L. G. Khachiyan. The polynomial solvability of convex quadratic programming. *USSR Computational Mathematics and Mathematical Physics*, 20(5):223–228, 1980.

[9] J. J. Moré and S. A. Vavasis. On the solution of concave knapsack problems. *Mathematical Programming*, 49(1-3):397–411, 1990.

[10] T. S. Motzkin and E. G. Straus. Maxima for graphs and a new proof of a theorem of Turán. *Canadian Journal of Mathematics*, 17:533–540, 1965.

[11] K. G. Murty and S. N. Kabadi. Some NP-complete problems in quadratic and nonlinear programming. *Mathematical Programming*, 39(2):117–129, 1987.

[12] P. M. Pardalos and G. Schnitger. Checking local optimality in constrained quadratic programming is NP-hard. *Operations Research Letters*, 7(1):33–35, 1988.

[13] P. M. Pardalos and S. A. Vavasis. Quadratic programming with one negative eigenvalue is NP-hard. *Journal of Global optimization*, 1(1):15–22, 1991.

[14] P. M. Pardalos and S. A. Vavasis. Open questions in complexity theory for numerical optimization. *Mathematical Programming*, 57(1-3):337–339, 1992.

[15] S. A. Vavasis. Quadratic programming is in NP. *Information Processing Letters*, 36(2):73–77, 1990.