RC25395 (WAT1307-029) July 15, 2013
Mathematics

IBM Research Report

Computational and Algebraic Aspects of Convexity

Amir Ali Ahmadi
IBM Research Division
Thomas J. Watson Research Center
P.O. Box 208
Yorktown Heights, NY 10598
USA

Computational and Algebraic Aspects of Convexity[1]

Amir Ali Ahmadi
Goldstine Fellow, IBM Watson Research Center
aaa@us.ibm.com

We would like to take this opportunity to express our gratitude to the 2012 Informs Computing Society Prize Committee, which was chaired by Pascal Van Hentenryck and included Daniel Bienstock and Dorit Hochbaum. Thank you Pascal, Dan, and Dorit for this very kind distinction. It is a great honor to receive this prize.

The ICS Prize was awarded for our work in [4], [5], [6] on the study of some very basic questions about *convexity*. The first paper [4] studies the computational complexity of recognizing convexity of functions and sets in polynomial optimization. The second and third papers [5], [6] are on an algebraic relaxation for convexity, known as *sum-of-squares-convexity* (sos-convexity), which has links to semidefinite programming and plays a notable role in the emerging field of *convex algebraic geometry* [9]. Some of the results, we believe, turned out to be interesting—our complexity paper answered an open question of Naum Shor from 1992; our study of sos-convexity revealed an unforeseen connection with a classical result of Hilbert in real algebraic geometry. We welcome your feedback and hope that you also find some aspects of the work interesting. A brief description of the main contributions follows with more details to be found in references mentioned above or in [1].

1 Complexity of Deciding Convexity

Over the last century, the notion of convexity has established itself as a central concept in the theory of optimization and operations research. Extensive and greatly successful research in the applications of convex optimization has shown that surprisingly many problems of practical importance can be cast as convex optimization problems. Moreover, we have a fair number of rules based on the calculus of convex functions that allow us to design—whenever we have the freedom to do so—problems that are by construction convex. Nevertheless, in order to be able to exploit the potential of convexity in optimization in full, a very basic question is to understand whether we are even able to *recognize* the presence of convexity in optimization problems. In other words, can we have an efficient algorithm that tests whether a given optimization problem is convex?

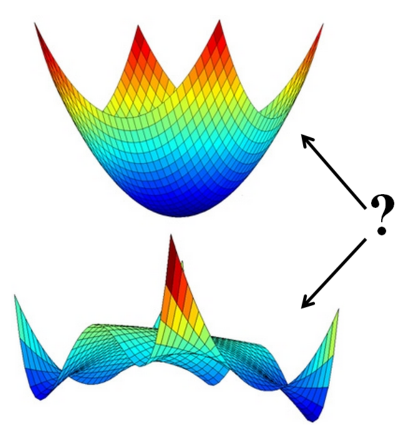

A class of optimization problems that allow for a rigorous study of this question from a computational complexity viewpoint is the class of polynomial optimization problems. These are optimization problems where the objective is given by a polynomial function and the feasible set is described by polynomial

[1]This note is to appear in the newsletter of the Informs Computing Society (ICS), in a column on the 2012 ICS Prize for best series of papers at the interface of operations research and computer science, awarded to A.A. Ahmadi, A. Olshevsky, P.A. Parrilo, and J.N. Tsitsiklis for papers listed in [4], [5], [6].

inequalities. Our research in this direction was motivated by a concrete question of N. Z. Shor that appeared as one of seven open problems in complexity theory for numerical optimization put together by Pardalos and Vavasis in 1992 [17]:

> "Given a degree-4 polynomial in n variables, what is the complexity of determining whether this polynomial describes a convex function?"

The reason why Shor's question is specifically about degree 4 polynomials is that deciding convexity of odd degree polynomials is trivial[2] and deciding convexity of degree 2 (quadratic) polynomials can be reduced to the simple task of checking whether a constant matrix is positive semidefinite. So, the first interesting case really occurs for degree 4 (quartic) polynomials. The main contribution of our first paper [4] is to show that deciding convexity of polynomials is strongly NP-hard already for polynomials of degree 4.

The implication of strong NP-hardness of this problem is that unless P=NP, there exists no algorithm that can take as input the (rational) coefficients of a quartic polynomial, have running time bounded by a polynomial in the numeric value of these coefficients (let alone in the number of bits needed to represent the coefficients), and output correctly on every instance whether or not the polynomial is convex. The reduction that establishes our result is purely algebraic. It shows that the NP-hard problem of testing global nonnegativity of so-called *biquadratic forms* (a special subclass of quartic polynomials) can be turned into the question of checking convexity by doubling the number of variables and without changing the degree.

Although the implications of convexity are very significant in optimization theory, our results suggest that unless additional structure is present, ensuring the mere presence of convexity is likely an intractable task. It is therefore natural to wonder whether there are other properties of optimization problems that share some of the attractive consequences of convexity, but are easier to check.

1.1 Complexity of deciding variants of convexity

In the same paper, we also study the complexity of recognizing some well-known variants of convexity, namely, the problems of deciding *strict convexity*, *strong convexity*, *pseudoconvexity*, and *quasiconvexity* of polynomials. The relationship between these notions is as follows (with none of the converse implications being true in general):

$$\text{strong convexity} \Longrightarrow \text{strict convexity} \Longrightarrow \text{convexity} \Longrightarrow \text{pseudoconvexity} \Longrightarrow \text{quasiconvexity}.$$

Table 1: Summary of our complexity results. A yes (no) entry means that the question is trivial for that particular entry because the answer is always yes (no) independent of the input. By P, we mean that the problem can be solved in polynomial time.

property vs. degree	1	2	odd ≥ 3	even ≥ 4
strong convexity	no	P	no	strongly NP-hard
strict convexity	no	P	no	strongly NP-hard
convexity	yes	P	no	strongly NP-hard
pseudoconvexity	yes	P	P	strongly NP-hard
quasiconvexity	yes	P	P	strongly NP-hard

Strict convexity is a property that is often useful to check because it guarantees uniqueness of the optimal solution in optimization problems. The notion of strong convexity is a common assumption in

[2] Independent of their coefficients, polynomials of odd degree are never (globally) convex, unless they are linear and hence always convex.

convergence analysis of many iterative Newton-type algorithms in optimization theory; see, e.g., [10, Chaps. 9–11]. So, in order to ensure the theoretical convergence rates promised by many of these algorithms, one needs to first make sure that the objective function is strongly convex. The problem of checking quasiconvexity (convexity of sublevel sets) of polynomials also arises frequently in practice. For instance, if the feasible set of an optimization problem is defined by polynomial inequalities, by certifying quasiconvexity of the defining polynomials we can ensure that the feasible set is convex. In several statistics and clustering problems, we are interested in finding minimum volume convex sets that contain a set of data points in space. This problem can be tackled by searching over the set of quasiconvex polynomials [14]. In economics also, quasiconcave functions are prevalent as desirable utility functions [7]. Finally, the notion of pseudoconvexity is a natural generalization of convexity that inherits many of the attractive properties of convex functions. For example, every stationary point or every local minimum of a pseudoconvex function must be a global minimum. Because of these nice features, pseudoconvex programs have been studied extensively in nonlinear programming [15], [11].

Our complexity results as a function of the degree of the polynomial are listed in Table 1.1. As you can see, all of these properties are easy to decide for quadratics but hard for polynomials of even degree 4 or higher. Somewhat surprisingly, we were able to show that testing quasiconvexity and pseudoconvexity can be done in polynomial time if the degree is odd.

2 Convexity and SOS-Convexity

Of course, NP-hardness of a problem does not stop us from studying it, but on the contrary, stresses the need for finding good approximation algorithms that can deal with a large number of instances efficiently. Towards this goal, we study in [5], [6] a *semidefinite relaxation* for convexity of polynomials known as *sos-convexity.*

Definition 2.1. *A polynomial $p(x) = p(x_1, \ldots, x_n)$ is sos-convex if its Hessian $H(x)$ can be factored as $H(x) = M^T(x)M(x)$ with a possibly nonsquare polynomial matrix*[3] *$M(x)$.*

It is easy to see that sos-convexity is a sufficient condition for convexity of polynomials. Indeed, if we have the factorization $H(x) = M^T(x)M(x)$, then $H(x)$ must be positive semidefinite for all x. Moreover, one can show that sos-convexity of a polynomial p can be decided by solving a single semidefinite program whose size is polynomial in the size of the coefficients of p; see e.g. [5].

The idea behind sos-convexity is related to the concept of representing *nonnegative polynomials* as *sums of squares*[4]—a deep-rooted subject in real algebraic geometry that has found widespread recent applications in optimization theory. Just like a sum of squares (sos) decomposition produces an algebraic certificate for nonnegativity, sos-convexity can be thought of as an algebraic certificate for convexity. In [6, Thm. 3.1], we showed that if one applies the sos relaxation to the standard definition of convexity or its first order characterization (as opposed to Definition 2.1, which appeals to the second order characterization of convexity), then one ends up with conditions that are equivalent to sos-convexity. This is reassuring in that it demonstrates the legitimacy of sos-convexity as *the* right semidefinite relaxation for convexity.

2.1 The first example of a convex polynomial that is not sos-convex

The connection of sos-convexity to semidefinite programming has motivated its use in many application domains, such as statistics (convex regression, minimum volume shape fitting, etc., [14]) and control theory (stability of hybrid systems [3]). Aside from its computational implications, sos-convexity is a

[3] A polynomial matrix is simply a matrix whose entries are (multivariate) polynomials.

[4] A polynomial p is nonnegative if $p(x) \geq 0$ for all $x \in \mathbb{R}^n$, and it is a sum of squares if it can be written as $p = \sum q_i^2$, for some polynomials q_i.

concept of interest in the field of *convex algebraic geometry* [9], which is devoted to the study of convex sets with algebraic structure. In particular, one of the early results in this area, due to Helton and Nie [12], states that any subset of $\mathbb{R}^n$ defined as $\{x|\ g_i(x) \leq 0\}$, with each g_i an sos-convex polynomial, can be represented as the projection of the feasible set of a semidefinite program.

Motivated by results of this type, it had been speculated whether sos-convexity of a polynomial was in fact equivalent to its convexity. We showed in [5], via a concrete counterexample (a polynomial in 3 variables and degree 8), that the answer was negative. This example came before our NP-hardness result in [4]. Indeed, complexity considerations alone suggest (assuming P≠NP) that convex but not sos-convex polynomials should exist, at least when the number of variables goes to infinity.

For the curious reader, we include here the first example of a convex but not sos-convex polynomial, published in [5]:

$$\begin{aligned} p(x) \;=\;& 32x_1^8 + 118x_1^6x_2^2 + 40x_1^6x_3^2 + 25x_1^4x_2^4 - 43x_1^4x_2^2x_3^2 - 35x_1^4x_3^4 + 3x_1^2x_2^4x_3^2 \\ & -16x_1^2x_2^2x_3^4 + 24x_1^2x_3^6 + 16x_2^8 + 44x_2^6x_3^2 + 70x_2^4x_3^4 + 60x_2^2x_3^6 + 30x_3^8. \end{aligned}$$

We found this polynomial with the help of a computer and as a solution of a carefully-designed semidefinite program; see [5, Sect. 4] for details. In general, finding polynomials with such properties is a nontrivial task. For example, the following closely related problem is still open.

Open problem.

Find an explicit example of a convex, nonnegative polynomial that is not a sum of squares.

Blekherman [8] has shown via volume arguments that for degree $d \geq 4$ and asymptotically for large n such polynomials must exist, although no examples are known. It is known, however, that such a convex polynomial must necessarily be not sos-convex [12],[6]. The question is particularly interesting from an optimization viewpoint since it implies that the well-known sum of squares relaxation for minimizing polynomials [19], [18] is not always exact, even in the easy case of minimizing convex polynomials.

2.2 A full characterization of cases where convexity equals sos-convexity

One of the cornerstones of real algebraic geometry is Hilbert's seminal paper in 1888 [13], where he gives a complete characterization of the degrees and dimensions for which nonnegative polynomials can be written as sums of squares of polynomials. In particular, Hilbert proves in [13] that there exist nonnegative polynomials that are not sums of squares, although explicit examples of such polynomials appeared only about 80 years later [16] and the study of the gap between nonnegative and sums of squares polynomials continues to be an active area of research to this day.

Once we produced the first example of a convex but not sos-convex polynomial in [5], it was natural to aim for a characterization of the dimensions and degrees for which such polynomials can exist, similar to the characterization that Hilbert provided for nonnegativity and sum of squares.

The contribution of our final paper [6] is to provide such a characterization for the inclusion relationship between convexity and sos-convexity. The results are summarized in Figure 1 and cover both the case of polynomials and forms (homogeneous polynomials). The entry $(n, d) = (3, 4)$ in the table on the right is particularly challenging and is joint work with G. Blekherman [2].

convex=sos-convex?

Polynomials

n,d	2	4	≥6
1	yes	yes	yes
2	yes	yes	no
3	yes	no	no
≥4	yes	no	no

Forms

n,d	2	4	≥6
1	yes	yes	yes
2	yes	yes	yes
3	yes	yes	no
≥4	yes	no	no

Figure 1: The tables answer whether every convex polynomial (form) in n variables and of degree d is sos-convex [6]. A similar characterization for nonnegativity and sum of squares was done by Hilbert in 1888 [13].

The remarkable overall outcome of this research is that convex polynomials (resp. forms) were shown to to be sos-convex *precisely* in the cases where nonnegative polynomials (resp. forms) are sums of squares, as shown by Hilbert. The proofs given in our work certainly do not use or depend on the results of Hilbert and it is not clear to us at this point whether this similarity is just a coincidence. It remains to be seen whether there is a more satisfying explanation of the resemblance of the two results, perhaps revealing a more fundamental connection between two basic concepts, nonnegativity and convexity.

References

[1] A. A. Ahmadi. *Algebraic relaxations and hardness results in polynomial optimization and Lyapunov analysis*. PhD thesis, Massachusetts Institute of Technology, September 2011. Available at `http://aaa.lids.mit.edu/publications`.

[2] A. A. Ahmadi, G. Blekherman, and P. A. Parrilo. Convex ternary quartics are sos-convex. In preparation, 2013.

[3] A. A. Ahmadi and R. Jungers. SOS-convex Lyapunov functions with applications to nonlinear switched systems. Submitted to the *2013 IEEE Conference on Decision and Control*, 2013.

[4] 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(1-2):453–476, 2013.

[5] A. A. Ahmadi and P. A. Parrilo. A convex polynomial that is not sos-convex. *Mathematical Programming*, 135(1-2):275–292, 2012.

[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] K. J. Arrow and A. C. Enthoven. Quasi-concave programming. *Econometrica*, 29(4):779–800, 1961.

[8] G. Blekherman. Convex forms that are not sums of squares. arXiv:0910.0656, 2009.

[9] G. Blekherman, P. A. Parrilo, and R. Thomas. *Semidefinite optimization and convex algebraic geometry*. SIAM Series on Optimization, 2013.

[10] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.

[11] R. W. Cottle and J. A. Ferland. On pseudo-convex functions of nonnegative variables. *Math. Programming*, 1(1):95–101, 1971.

[12] J. W. Helton and J. Nie. Semidefinite representation of convex sets. *Mathematical Programming*, 122(1, Ser. A):21–64, 2010.

[13] D. Hilbert. Über die Darstellung Definiter Formen als Summe von Formenquadraten. *Math. Ann.*, 32, 1888.

[14] A. Magnani, S. Lall, and S. Boyd. Tractable fitting with convex polynomials via sum of squares. In *Proceedings of the 44^{th} IEEE Conference on Decision and Control*, 2005.

[15] O. L. Mangasarian. Pseudo-convex functions. *J. Soc. Indust. Appl. Math. Ser. A Control*, 3:281–290, 1965.

[16] T. S. Motzkin. The arithmetic-geometric inequality. In *Inequalities (Proc. Sympos. Wright-Patterson Air Force Base, Ohio, 1965)*, pages 205–224. Academic Press, New York, 1967.

[17] P. M. Pardalos and S. A. Vavasis. Open questions in complexity theory for numerical optimization. *Mathematical Programming*, 57(2):337–339, 1992.

[18] P. A. Parrilo and B. Sturmfels. Minimizing polynomial functions. *Algorithmic and Quantitative Real Algebraic Geometry, DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, 60:83–99, 2003.

[19] N. Z. Shor. Class of global minimum bounds of polynomial functions. *Cybernetics*, 23(6):731–734, 1987. (Russian orig.: Kibernetika, No. 6, (1987), 9–11).