Research Statement — Rising Stars in EECS 2017

My research focuses on a particular class of optimization problems, those where one optimizes *over the set of nonnegative polynomials*. In these problems, the coefficients of the polynomial are the unknowns and the objective and constraints are affine in these coefficients, bar one constraint which requires that the polynomial be nonnegative (either globally, or over a subset of the Euclidean space defined by polynomial inequalities). Problems of this type appear ubiquitously in applications: in dynamics and control, they are used, e.g., to find Lyapunov functions; in game theory and economics, they are used to compute Nash equilibria; in power engineering, they can be used to solve the optimal power flow problem; in combinatorial optimization and integer programming, they can be used to provide optimality certificates. My interest in this area stems from the fact that it closely connects two different fields: computational algebra and convex optimization. Optimizing over the set of nonnegative polynomials is an NP-hard problem, whose study requires the use of tools from approximation algorithms and, more generally, theoretical computer science. One of these topics of study is to provide tractable approximations to the initial problem using convex optimization. A well-known sufficient condition for nonnegativity, dating back to the 1900s and the work of David Hilbert[6], is via the concept of *sum of squares* polynomials. A polynomial p is said to be a sum of squares if it can be written as the sum of squares of other polynomials, i.e., $p(x) = \sum_i q_i(x)^2$, where $\{q_i\}$ are polynomials. It is easy to see that if a polynomial p is a sum of squares then it is nonnegative, i.e., $p(x) \geq 0, \forall x$. As a consequence, one can replace the nonnegativity condition in the previous optimization problem by a sum of squares condition and optimize over a subset of the set of nonnegative polynomials. The benefit of doing so is that the resulting optimization problem (also called a *sum of squares program*) can be recast as a semidefinite program, which can be solved in polynomial time. The research I have done as a graduate student straddles both computational algebra and optimization theory. A significant part of my work has been on the optimization side, to develop algorithms that can solve large-scale semidefinite programs. Another part of my work has been to combine tools from computational algebra and optimization to study sets of polynomials with certain structural properties, such as convexity or monotonocity. The structural properties I consider are of practical interest in many applications, from statistics to computer vision. In the sections below, I expand on these two research directions.

Algorithms for large-scale semidefinite programs. Semidefinite programming (SDP) is a subfield of convex optimization, containing problems of the form $\min_{X \in S^n} \mathrm{Tr}(CX)$ such that $\mathrm{Tr}(A_iX) = b_i$ for $i = 1, \ldots, m$ and $X \succeq 0$, where S^n is the set of symmetric $n \times n$ matrices, C and $A_i \in S^n$, and $b_i \in \mathbb{R}$. SDP has many applications of its own, but as mentioned previously, it is of particular interest to me due to its connection to sum of squares optimization. It is well-known that semidefinite programs can be solved in polynomial time via, for example, interior point methods. Their solving time scales poorly however with problem dimension, and even when the dimension of the program is small, they arguably remain the most expensive class of convex optimization problems to solve. This is particularly detrimental to semidefinite programs that arise from sum of squares programs. Indeed, for a sum of squares program involving polynomials of degree $2d$ and in n variables, the resulting SDP will be of size $\binom{n+d}{d}$. Consequently, a large portion of my work has been to develop algorithms that replace the semidefinite program at hand with optimization problems that are more tractable and scalable. More specifically, in joint work with Ahmadi and Dash [1, 4], I developped two different methods that produce iterative sequences of linear and second order cone programs, whose objective values approximate the objective value of the semidefinite program that they aim to replace. Future work in this area includes a shift from theory to practice, with the development of customized solvers for SDP using these ideas.

Convex polynomials: theory, applications, and algorithms. Containing 3D point clouds using sets of compact description, optimizing over the set of norms, speeding-up solving time for difference of convex programs: these topics may seem unrelated at first glance, but at their core, they all share one common characteristic—*convex polynomials*. Convexity, analogously to nonnegativity, is a property that is hard to even test and much work has been done to provide tractable relaxations for it. A better understanding of the set of convex polynomials and its algorithmic applications in diverse domains has been another focal point of my research. In [2] for example, my area of interest was norms. In the paper, jointly with Ahmadi and de Klerk, I show that any norm can be approximated to arbitrary accuracy by a *polynomial norm*, i.e., the d^{th} root of a degree-d homogeneous strictly convex polynomial. The paper provides a comprehensive study of such norms, including efficient ways of testing and optimizing over the set. This opens up the door, e.g., to data-tailored regularizers in statistics and automated Lyapunov analysis of certain classes of control systems. In [5], motivated by applications in computer vision and robotics, I show (in joint work with Ahmadi, Makadia, and Sindhwani) how one can use sublevel sets of convex polynomials to contain 3D objects as tightly as possible (see the figure below).

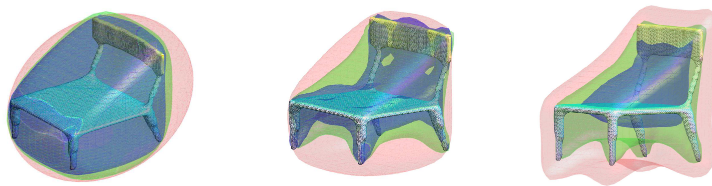

Describing bounding volumes using polynomial sublevel sets also facilitates tasks such as distance or penetration computation among 3D objects. Finally, in joint work with Ahmadi [3], I studied the set of polynomials that could be written as the difference of two convex polynomials. This is of interest in the context of difference of convex optimization, a subclass of optimization problems where both the objective and the constraints are given as the difference of two convex functions. The paper shows that any polynomial can be written as a difference of two convex polynomials, but that finding such a decomposition is a hard problem. It proposes relaxations based on conic programming and explains how these can be used to speed-up existing heuristics for solving difference of convex optimization problems. In the future, I aim to continue my work on studying sets of polynomials which are shape-constrained. My current interest is in the set of monotonic polynomials.

Career goals. I conclude with a few words regarding my career goals. When I was growing up, and all the way through my undergraduate studies, I never considered going into academia. This was in part due to the fact that academia as I was exposed to it lacked role models and people I could identify with. Hence, when the time came to consider graduate studies, I applied for a Master's degree, rather than a PhD at Princeton, until the university convinced me to switch my application to a PhD. I am truly grateful to have had this opportunity, as I have met some exceptional researchers since my arrival here who have led me to revise my first impression of academia. I realize though that many students would not be so lucky. As a consequence, one of the things that I most aspire to be is the kind of professor who encourages and inspires their students to step out of their comfort zones when it comes to their long-term ambitions. This workshop would be extremely valuable to me as a first step to achieving this goal.

References

[1] A. A. Ahmadi, S. Dash, and G. Hall. Optimization over structured subsets of positive semidefinite matrices via column generation. *Discrete Optimization*, 2016.

[2] A. A. Ahmadi, E. de Klerk, and G. Hall. Polynomial norms. *arXiv preprint arXiv:1704.07462*, 2017.

[3] A. A. Ahmadi and G. Hall. DC decomposition of nonconvex polynomials with algebraic techniques. *Mathematical Programming*, 2017.

[4] A. A. Ahmadi and G. Hall. Sum of squares basis pursuit with linear and second order cone programming. *Contemporary Mathematics*, 2017.

[5] A. A. Ahmadi, G. Hall, A. Makadia, and V. Sindhwani. Geometry of 3D environments and sum of squares polynomials. *RSS*, 2017.

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