



AFRL-RI-RS-TR-2018-148

SCALABLE RAPIDLY DEPLOYABLE CONVEX OPTIMIZATION FOR DATA ANALYTICS

STANFORD UNIVERSITY

JUNE 2018

FINAL TECHNICAL REPORT

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED

STINFO COPY

**AIR FORCE RESEARCH LABORATORY
INFORMATION DIRECTORATE**

NOTICE AND SIGNATURE PAGE

Using Government drawings, specifications, or other data included in this document for any purpose other than Government procurement does not in any way obligate the U.S. Government. The fact that the Government formulated or supplied the drawings, specifications, or other data does not license the holder or any other person or corporation; or convey any rights or permission to manufacture, use, or sell any patented invention that may relate to them.

This report is the result of contracted fundamental research deemed exempt from public affairs security and policy review in accordance with SAF/AQR memorandum dated 10 Dec 08 and AFRL/CA policy clarification memorandum dated 16 Jan 09. This report is available to the general public, including foreign nations. Copies may be obtained from the Defense Technical Information Center (DTIC) (<http://www.dtic.mil>).

AFRL-RI-RS-TR-2018-148 HAS BEEN REVIEWED AND IS APPROVED FOR PUBLICATION IN ACCORDANCE WITH ASSIGNED DISTRIBUTION STATEMENT.

FOR THE CHIEF ENGINEER:

/ S /

JOHN SPINA
Work Unit Manager

/ S /

JON S. JONES
Technical Advisor, Information Intelligence
Systems and Analysis Division
Information Directorate

This report is published in the interest of scientific and technical information exchange, and its publication does not constitute the Government's approval or disapproval of its ideas or findings.

REPORT DOCUMENTATION PAGE				Form Approved OMB No. 0704-0188	
The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.					
1. REPORT DATE (DD-MM-YYYY) JUN 2018		2. REPORT TYPE FINAL TECHNICAL REPORT		3. DATES COVERED (From - To) SEP 2012 – DEC 2017	
4. TITLE AND SUBTITLE SCALABLE RAPIDLY DEPLOYABLE CONVEX OPTIMIZATION FOR DATA ANALYTICS				5a. CONTRACT NUMBER N/A	
				5b. GRANT NUMBER FA8750-12-2-0306	
				5c. PROGRAM ELEMENT NUMBER 62702E	
6. AUTHOR(S) Stephen Boyd				5d. PROJECT NUMBER XDAT	
				5e. TASK NUMBER A0	
				5f. WORK UNIT NUMBER 02	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Board of Trustees of the Leland Stanford Junior University 3160 Porter Drive, Suite 100 Palo Alto, CA 94304				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Air Force Research Laboratory/RIEA 525 Brooks Road Rome NY 13441-4505				10. SPONSOR/MONITOR'S ACRONYM(S) AFRL/RI	
				11. SPONSOR/MONITOR'S REPORT NUMBER AFRL-RI-RS-TR-2018-148	
12. DISTRIBUTION AVAILABILITY STATEMENT Approved for Public Release; Distribution Unlimited. This report is the result of contracted fundamental research deemed exempt from public affairs security and policy review in accordance with SAF/AQR memorandum dated 10 Dec 08 and AFRL/CA policy clarification memorandum dated 16 Jan 09					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT Over the period of the contract we have developed the full stack for wide use of convex optimization, in machine learning and many other areas. We have developed a new open-source domain specific language (DSL) for convex optimization, in Python, Julia, and recently, in R. Each of these has been published, and each are widely used and cited. We have developed a novel open-source solver, SCS, that is bundled with the DSLs, and solves any combination of linear programs, SOCPs, SDPs, exponential cone programs, and power cone programs. CVXPY supports basic methods for distributed optimization, on multiple heterogenous platforms. We have also done basic research in various application areas, using CVXPY, to demonstrate its usefulness. See attached report for publication information.					
15. SUBJECT TERMS CVXPY, Xdata, Convex Optimization					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT UU	18. NUMBER OF PAGES 16	19a. NAME OF RESPONSIBLE PERSON JOHN SPINA
a. REPORT U	b. ABSTRACT U	c. THIS PAGE U			19b. TELEPHONE NUMBER (Include area code)

TABLE OF CONTENTS

1	SUMMARY	1
2	INTRODUCTION.....	1
3	METHODS, ASSUMPTIONS, AND PROCEDURES	1
3.1	Software for convex optimization.....	1
4	RESULTS AND DISCUSSION	1
4.1	Project Overview and Milestones.....	1
4.2	Solver for large-scale convex optimization.....	2
4.3	Dynamic estimation and management applications	5
5	CONCLUSIONS	9
6	NOTES	9
6.1	Awards.....	9
	SYMBOLS, ABBREVIATIONS, AND ACRONYMS	12

1 SUMMARY

The objective of the XDATA program was to develop computational techniques and software tools for analyzing and visualizing large volumes of data, including semi-structured data such as tabular, relational, categorical, or meta-data and unstructured data such as text documents or message traffic.

2 INTRODUCTION

Over the period of the contract we have developed the full stack for wide use of convex optimization, from open-source software and solver to applications in machine learning and many other areas. We categorize our main result in three ways: software for convex optimization, solver for large-scale optimization, and dynamic estimation applications. During the period of the grant we have published more than thirty papers related to the main thrust of the grant, and various applications. Professor Stephen Boyd's Google citation count has risen from around 80000 to around 130000 over the period of the grant. All of these papers, and associated code, are available at Professor Boyd's website.

3 METHODS, ASSUMPTIONS, AND PROCEDURES

3.1 Software for convex optimization

Using convex optimization in an application requires either developing a custom solver or converting the problem into a standard form. Both of these tasks require expertise, and are time consuming and error prone. We used a domain-specific language (DSL), which would allow non-experts to write short, easily readable scripts that can apply a wide range of convex optimization and machine learning methods to huge datasets. Like CVX (Grant and Boyd, 2014), this language could dramatically shorten the time between inception and deployment of a given optimization-based application.

4 RESULTS AND DISCUSSION

4.1 Project Overview and Milestones

We have developed a new open-source DSL for convex optimization, in Python, Julia, and recently, in R. Each of these has been published, and each are widely used and cited. This allows the user to specify the problem in a natural way that follows the math; this specification is then automatically converted into the standard form required by generic solvers. For example, CVXPY converts problems into a standard form known as conic form, a generalization of a linear program. The conversion is done using graph implementations of convex functions (Grant and Boyd, 2008). The resulting cone program is equivalent to the original problem, so by solving it we obtain a solution of the original problem. Solvers that handle conic form are known as cone solvers; each one can handle combinations of several types of cones. CVXPY interfaces with the open-source cone solvers CVXOPT, ECOS (Domahidi et al., 2013), and SCS (ODonoghue et al., 2016), which

are implemented in combinations of Python and C. These solvers have different characteristics, such as the types of cones they can handle and the type of algorithms employed.

We also developed many tools on top of CVXPY. Disciplined convex stochastic programming (DCSP) (Ali et al., 2015) allows modelers to specify -in a straightforward way- and solve convex optimization problems that include expectations of arbitrary expressions, partial optimizations, optimizations over (only) a subset of the optimization variables, which additionally pave the way for the specification of multi-stage stochastic programs, and chance constraints that are required to hold with high probability. These three building blocks can be used to express a wide variety of stochastic optimization problems. And we introduced disciplined convex-concave programming (DCCP), which combines the ideas of disciplined convex programming (DCP) with convex-concave programming (CCP) (Shen et al., 2016). Selected publications are listed below:

- CVXPY: A python-embedded modeling language for convex optimization
S. Diamond and S. Boyd
 - Convex optimization in Julia
M. Udell, K. Mohan, D. Zeng, J. Hong, S. Diamond, and S. Boyd
 - CVXR: An R Package for disciplined convex optimization
A. Fu, B. Narasimhan, and S. Boyd
 - Disciplined convex stochastic programming: A new framework for stochastic optimization
A. Ali, Z. Kolter, S. Diamond, and S. Boyd
 - Disciplined multi-convex programming
X. Shen, S. Diamond, M. Udell, Y. Gu, and S. Boyd
 - Disciplined convex-concave programming
X. Shen, S. Diamond, Y. Gu, and S. Boyd
 - A new architecture for optimization modeling frameworks
M. Wytock, S. Diamond, F. Heide, and S. Boyd

4.2 Solver for large-scale convex optimization

Due to the explosion in size and complexity of modern datasets, it has been increasingly important to be able to solve such large-scale problems across multiple machines. Thus, either the decentralized collection or storage of these datasets as well as accompanying distributed solution methods are highly desirable. We first surveyed several important classes of such algorithms including proximal algorithm (Patrik and Boyd, 2016), which is compatible with non-smooth, constrained, large-scale problem in a distributed fashion. They are very generally applicable, but are especially well-suited to problems of substantial recent interest involving large or high-dimensional datasets.

In addition, we also presented more fundamental theoretical analysis for convex optimization. We derived a second-order ordinary differential equation (ODE) (Su et al., 2016), which exhibits approximate equivalence to Nesterov's scheme and thus can serve as a tool for analysis. We show

that the continuous time ODE allows for a better understanding of Nesterov’s scheme. As a by-product, we obtain a family of schemes with similar convergence rates. The ODE interpretation also suggests restarting Nesterov’s scheme, leading to an algorithm which can be rigorously proven to converge at a linear rate whenever the objective is strongly convex. Other theoretical works such as the convergence analysis of ADMM (Giselsson and Boyd, 2015) or duality gap analysis (Udell and Boyd, 2016) were derived and several practical preconditioning and matrix-free methods (Giselsson and Boyd, 2016) (Diamond and Boyd, 2017) were also suggested.

Moreover, we have developed a novel open-source solver including SCS (O’Donoghue et al., 2016), which uses an operator splitting method, the alternating directions method of multipliers, to solve the homogeneous self-dual embedding. SCS is bundled with the DSLs, and solves any combination of linear programs, SOCPs, SDPs, exponential cone programs and power cone programs. Compared to interior-point methods, first-order methods scale to very large problems, at the cost of requiring more time to reach very high accuracy. Compared to other first-order methods for cone programs, our approach finds either primal and dual solutions when available or a certificate of infeasibility or unboundedness otherwise, is parameter free, and the per-iteration cost of the method is the same as applying a splitting method to the primal or dual alone. We also report numerical results that show speedups over interior-point cone solvers for large problems, and scaling to very large general cone programs.

We have additionally developed several solvers such as SnapVX (Hallac et al, 2017), and the nonconvex objective or constraint (Diamond et al., 2017) (Takapoui et al., 2016). SnapVX is a high-performance solver that combines the capabilities of two open source software packages: a large scale graph processing library, Snap.py, and our CVXPY. General nonconvex quadratically constrained quadratic programs (QCQPs) (Park and Boyd, 2017), which is based on Suggest-and-Improve framework, generalizes a number of known methods and provides heuristics to get approximate solutions for which no specialized methods are available. We also investigated the convex-concave procedure (CCP) (Lipp and Boyd, 2016), a local heuristic that utilizes the tools of convex optimization to find local optima of difference of convex (DC) programming problems. The class of DC problems is very general, and includes difficult problems such as the traveling salesman problem. We extend the standard procedure in two major ways and describe several variations. First, we describe a method that allows the algorithm to be initialized without a feasible point. Second, we generalize the algorithm to include vector inequalities. Selected publications are listed below:

- Block splitting for distributed optimization
N. Parikh and S. Boyd
- Conic optimization via operator splitting and homogeneous self-dual embedding
B. O’Donoghue, E. Chu, N. Parikh, and S. Boyd
- SnapVX: A network-based convex optimization solver
D. Hallac, C. Wong, S. Diamond, A. Sharang, R. Sasic, S. Boyd, and J. Leskovec.
- Variations and extensions of the convex-concave procedure
T. Lipp and S. Boyd

- A general system for heuristic solution of convex problems over nonconvex sets
S. Diamond, R. Takapoui, and S. Boyd
- OSQP: An operator splitting solver for quadratic programs
B. Stellato, G. Banjac, P. Goulart, A. Bemporad, and S. Boyd
- Embedded code generation using the OSQP solver
G. Banjac, B. Stellato, N. Moehle, P. Goulart, A. Bemporad and S. Boyd
- A semidefinite programming method for integer convex quadratic minimization
J. Park and S. Boyd
- A simple effective heuristic for embedded mixed-integer quadratic programming
R. Takapoui, N. Moehle, S. Boyd, and A. Bemporad
- General heuristics for nonconvex quadratically constrained quadratic programming
J. Park and S. Boyd
- Linear programming heuristics for the graph isomorphism problem
R. Takapoui and S. Boyd
- Concave quadratic cuts for mixed-integer quadratic problems
J. Park and S. Boyd
- Bounding duality gap for problems with separable objective
M. Udell and S. Boyd
- Diagonal scaling in Douglas-Rachford splitting and ADMM
P. Giselsson and S. Boyd
- Metric selection in fast dual forward backward splitting
P. Giselsson and S. Boyd
- Linear convergence and metric selection in Douglas-Rachford splitting and ADMM
P. Giselsson and S. Boyd
- Matrix-free convex optimization modeling,
S. Diamond and S. Boyd
- A differential equation for modeling Nesterovs accelerated gradient method
W. Su, S. Boyd, and E. Candes
- Proximal algorithms
N. Parikh and S. Boyd

- A primer on monotone operator methods
E. Ryu and S. Boyd

4.3 Dynamic estimation and management applications

The framework of convex optimization gives us a uniform method for handling dynamic estimation problems in finance, control, and radiation treatment planning. We developed a method for time-series analysis via convex optimization and dynamic programming. We introduced the time-varying graphical lasso (TVGL) (Hallac et al., 2017), a method of inferring time-varying networks from raw time series data and identifying changes in the model over time accordingly. We cast the problem in terms of estimating a sparse time-varying inverse covariance matrix, which reveals a dynamic network of interdependencies between the entities. We derive a scalable message-passing algorithm based on the Alternating Direction Method of Multipliers (ADMM) to solve this problem in an efficient way. TVGL obtained interpretable results and outperformed state-of-the-art baselines in terms of both accuracy and scalability.

For discovering repeated patterns in temporal time-series data, we proposed a new method of model-based clustering, which we call Toeplitz Inverse Covariance-Based Clustering (TICC) (Hallac et al., 2017). Each cluster in the TICC method is defined by a correlation network, or Markov random field (MRF), characterizing the interdependencies between different observations in a typical subsequence of that cluster. Based on this graphical representation, TICC simultaneously segments and clusters the time series data. We solve the TICC problem through alternating minimization, using a variation of the expectation maximization (EM) algorithm. And we demonstrated on an automobile sensor dataset how TICC can be used to learn interpretable clusters in real-world scenarios.

We developed a method, greedy Gaussian segmentation (GGS) (Hallac et al., 2016), breaking a multivariate time-series into segments over which the data is well explained as independent samples from a Gaussian distribution. We formulate this as a covariance-regularized maximum likelihood problem, which can be reduced to a combinatorial optimization problem of searching over the possible breakpoints, or segment boundaries. GGS is a heuristic method with linear complexity in the time series length that approximately solves the problem, and always yields a locally optimal choice, in the sense that no change of any one breakpoint improves the objective. It is quite efficient and easily scales to problems with vectors of dimension over 1000 and time series of arbitrary length. Finally, we illustrate the approach on various financial time series.

In addition to time-series analysis, based on convex predictive model, we also suggested several efficient strategies of dynamic resource allocation problem. We considered this to mitigate multi-user interference in the downlink Digital Subscriber Line (DSL) transmission (Zhang et al., 2017): given the real-time demands, determine the optimal transmission scheme: the optimal NOI and DOI size in each data frame as well as the optimal grouping strategy in the DOI, and optimally adjust the transmission scheme. We formulate these optimal dynamic resource allocation problems and propose efficient real-time algorithms to solve them to global optimality.

Also, we developed a decentralized method for solving this problem called proximal message passing (Kraning et al., 2012). The method is iterative: At each step, each device exchanges simple messages with its neighbors in the network and then solves its own optimization problem, minimizing its own objective function, augmented by a term determined by the messages it has received. The method is completely decentralized, and needs no global coordination other than synchronizing iterations; the problems to be solved by each device can typically be solved extremely efficiently and in parallel. This method is fast enough that even a serial implementation can solve substantial problems in reasonable time frames.

Lastly, we applied convex optimization to plan radiation treatment (Ungun et al., 2016) and also described the single-period and multi-period trading methods (Boyd et al., 2017) in one simple framework, giving a clear description of the development and the approximations made. The relevant publications are listed below.

- Toeplitz inverse covariance-based clustering of multivariate time series data
D. Hallac, S. Vire, S. Boyd, and J. Leskovec
- Multi-period trading via convex optimization
S. Boyd, E. Busseti, S. Diamond, R. Kahn, K. Koh, P. Nystrup, and J. Speth
- Performance bounds and suboptimal policies for multi-period investment
S. Boyd, M. Mueller, B. O'Donoghue, and Y. Wang
- Dynamic network energy management via proximal message passing
M. Kraning, E. Chu, J. Lavaei, and S. Boyd
- Network inference via the time-varying graphical lasso
D. Hallac, Y. Park, S. Boyd, and J. Leskovec
- A convex optimization approach to radiation treatment planning with dose constraints
A. Fu, B. Ungun, L. Xing, and S. Boyd
- Dynamic resource allocation for energy efficient transmission in digital subscriber lines
N. Zhang, Z. Yao, Y. Liu, S. Boyd, and Z.-Q. Luo
- Dynamic energy management with scenario-based robust MPC
M. Wytock, N. Moehle, and S. Boyd
- Greedy Gaussian segmentation of multivariate time series
D. Hallac, P. Nystrup, and S. Boyd
- Fitting jump models
A. Bemporad, V. Breschi, D. Piga, and S. Boyd
- Real time radiation treatment planning with optimality guarantees via cluster and bound methods
B. Ungun, L. Xing, and S. Boyd

4.4 Other publications Other than three categorized accomplishments, we researched various theoretical aspects of convex optimization and developed several heuristic algorithms (for nonconvex problems). We also took an advantage of convex optimization and dynamic programming in a wide range of applications such as feature selection (Boyd et al., 2018), generalized low rank model (Udell and Boyd, 2016), importance sampling (Ryu and Boyd, 2014), value function approximation (Moehle and Boyd, 2017) (O’Donoghue et al., 2013), MIMO proportional-integral-derivative (PID) controller design (Boyd et al, 2015), modeling heterogeneous data (Park et al., 2017), etc. All relevant papers are listed below.

- Saturating splines and feature selection
N. Boyd, T. Hastie, S. Boyd, S. Recht, and M. Jordan
- Learning the network structure of heterogeneous data via pairwise exponential Markov random fields
Y. Park, D. Hallac, S. Boyd, and J. Leskovec
- A distributed method for optimal capacity reservation
N. Moehle, X. Shen, Z.-Q. Luo, and S. Boyd
- Approximate dynamic programming via iterated Bellman inequalities
Y. Wang, B. O’Donoghue, and S. Boyd
- Value function approximation for direct control of switched power converters
N. Moehle and S. Boyd
- Stochastic matrix-free equilibration
S. Diamond and S. Boyd
- Maximum torque-per-current waveform design for induction motors via semidefinite programming
N. Moehle and S. Boyd
- Antagonistic control
T. Lipp and S. Boyd
- Optimization of rotational arc station parameter optimized radiation therapy
P. Dong, B. Ungun, S. Boyd, and L. Xing
- Generalized low rank models
M. Udell, C. Horn, R. Zadeh, and S. Boyd

- Line search for averaged operator iteration
P. Giselsson, M. Filt, and S. Boyd
- Risk-constrained Kelly gambling
E. Busseti, E. Ryu, and S. Boyd
- Optimal current waveforms for switched reluctance motors
N. Moehle and S. Boyd
- MIMO PID tuning via iterated LMI restriction
S. Boyd, M. Hast, and K. J. Astrom
- Model predictive control for wind power gradients
T. Hovgaard, S. Boyd, and J. Jrgensen
- A perspective-based convex relaxation for switched-affine optimal control
N. Moehle and S. Boyd
- Volume weighted average price optimal execution
E. Busseti and S. Boyd
- Extensions of Gauss quadrature via linear programming
E. Ryu and S. Boyd
- Optimal current waveforms for brushless permanent magnet motors
N. Moehle and S. Boyd
- Linear models based on noisy data and the Frisch scheme
L. Ning, T. Georgiou, A. Tannenbaum, and S. Boyd
- Network lasso: Clustering and optimization in large graphs
D. Hallac, J. Leskovec, and S. Boyd
- Adaptive importance sampling via stochastic convex programming
E. Ryu and S. Boyd
- Optimal crowd-powered rating and filtering algorithms

A. Parameswaran, S. Boyd, H. Garcia-Molina, A. Gupta, N. Polyzotis, and J. Widom

- Security constrained optimal power flow via proximal message passing
S. Chakrabarti, M. Kraning, E. Chu, R. Baldick, and S. Boyd
- Quadratic approximate dynamic programming for input-affine systems
A. Keshavarz and S. Boyd
- Minimum-time speed optimization over a fixed path
T. Lipp and S. Boyd
- Preconditioning in fast dual gradient methods
P. Giselsson and S. Boyd
- Monotonicity and restart in fast gradient methods
P. Giselsson and S. Boyd
- Infeasibility detection in the alternating direction method of multipliers for convex optimization
F. Banjac, P. Goulart, B. Stellato, and S. Boyd

5 Conclusions

Using convex optimization in an application requires either developing a custom solver or converting the problem into a standard form. Both of these tasks require expertise, and are time consuming and error prone. Like CVX (Grant and Boyd, 2014), this language could dramatically shorten the time between inception and deployment of a given optimization-based application. We categorize our main results in three ways: software for convex optimization, solver for large-scale optimization, and dynamic estimation applications.

6 Notes

6.1 Awards

During the period of the grant, Boyd has received a number of awards and recognitions:

2017 *Honorary Doctorate*, Université Catholique de Louvain (UCL).

2016 Selected as *INFORMS Fellow*.

2015 Chaim Weizmann Memorial lectures, Weizmann Institute of Science.

2015 Election as *SIAM Fellow*, with citation: “For fundamental contributions to the development, teaching, and practice of optimization in engineering.”

2015 Russell Severance Springer Professorship, University of California at Berkeley. **2014** Appointed to *Information Science and Technology Study Group* (DARPA ISAT). **2014** Elected to the *National Academy of Engineering*.

He and his group have won several best paper awards:

2017 *Runner-up, SIGKDD Best Paper Award*, for “Toeplitz Inverse Covariance-Based Clustering of Multivariate Time Series Data,” with coauthors David Hallac, Sagar Vare, and Jure Leskovec.

2017 *Chinese Control and Decision Conference Best Student Paper Award*, for “Disciplined Multi-Convex Programming,” with student co-author Xinyue Shen.

2015 *IEEE Signal Processing Letters Best Paper Award*, for paper “Compressed Sensing with Quantized Measurements”, co-authored with A. Zymnis and E. Candes.

2014 *IFAC Best Paper Prize*, for paper “Smoothed State Estimates Under Abrupt Changes Using Sum-of-Norms Regularization,” co-authored with H. Ohlsson, F. Gustafsson, and L. Ljung.

He has won several major teaching and writing awards:

2017 *Stanford Tau Beta Pi Teaching Honor Roll*.

2017 *IEEE James H. Mulligan, Jr. Education Medal*, IEEE’s highest education award, for a career of outstanding contributions to education in the fields of interest of IEEE, with citation: “For inspirational education of students and researchers in the theory and application of optimization.”

2016 *Walter J. Gores Award for Excellence in Teaching*, Stanford University’s highest award for excellence in teaching, with citation: “For his signature course, Convex Optimization, which attracts more than 300 Stanford students each year, is taught at more than 100 universities and over the past 20 years has had a profound influence on how researchers and engineers think about convex models to solve their problems; For revolutionizing the way mathematical optimization is taught and applied in engineering and the social and natural sciences worldwide; For the brilliance, clarity and humor with which he presents mathematically advanced topics, make them accessible and interesting to students in many fields.”

2014 *INFORMS Saul Gass Expository Writing Award*.

He has been invited to give, and has given, many plenary and keynote lectures:
List does not include many other distinguished lectures and colloquia presentations.

2018 Ørsted lectures, *Denmark Technical University*.

2018 Simons Mathematics lectures, *MIT*.

2017 Keynote lecture, *H2O World*.

2017 Robert T. Chien Distinguished lecture, *University of Illinois*.

2017 Distinguished lecture, *Hong Kong Polytechnic University*.

2017 Distinguished lecture, *Chinese University of Hong Kong*.

2017 Keynote lecture, *Chinese Control and Decision Conference*, Chongqing.

2017 Institute for Advanced Studies Distinguished lecture, *City University of Hong Kong*. **2016** Plenary lecture, *4th European Conference on Computational Optimization*, Leuven. **2016** Plenary lecture, *9th IEEE Sensor Array and Multichannel Signal Processing Workshop*, Rio de Janeiro.

2016 Plenary public lecture, *American Control Conference*, Boston.

2016 Plenary lecture, *International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*, Shanghai.

2015 Plenary lecture, *International Conference on Computer Vision (ICCV)*, Santiago, Chile.

2015 Erik Jonsson Distinguished lecture, *UT Dallas*.

2015 Plenary lecture, *5th International Conference on Control and Optimization with Industrial Applications*, Baku, Azerbaijan.

2015 Plenary lecture, *5th IFAC Workshop on Distributed Estimation and Control in Networked Systems*, Philadelphia.

2015 Keynote lecture, *ShanghaiTech Symposium on Data Science*.

2015 *55th Annual Weizmann Memorial lectures*.

2014 Keynote lecture, *Stanford-Berkeley Robotics Symposium*.

2014 Plenary lecture, *International Conference on Operations Research*, Aachen.

2014 Plenary lecture, *Convex Optimization and Beyond*, Edinburgh.

2014 Plenary lecture, *Mediterranean Conference on Control and Automation*, Palermo.

2014 Plenary lecture, *48th Annual Conference on Information Sciences and Systems (CISS)*, Princeton.

SYMBOLS, ABBREVIATIONS, AND ACRONYMS

ADMM.....	Alternating Direction Method of Multipliers
AFRL	Air Force Research Lab
CCP	Convex-Concave Programming
CVX	Convex Optimization
CVXPY	Convex Optimization for Python
DC	Difference of Convex
DCCP	Disciplined Convex-Concave Programming
DCP	Disciplined Convex Programming
DCSP.....	Disciplined Convex Stochastic Programming
DSL	Domain-Specific Language
DSL	Digital Subscriber Line
ECOS	Embedded Configurable Operating System
EM.....	Expectation Maximization
GGs.....	Greedy Gaussian Segmentation
MRF	Markov Random Field
ODE	Ordinary Differential Equation
QCQP	Quadratically Constrained Quadratic Programs
SCS	Splitting Conic Solver
SDP	Semidefinite Programming
SOCp.....	Second-Order Cone Program
TICC	Toeplitz Inverse Covariance-Based Clustering
TVGL.....	Time-Varying Graphical Lasso