Introduction to Mathematical Optimization

Nick Henderson, AJ Friend (Stanford University)
Kevin Carlberg (Sandia National Laboratories)

August 14, 2018

Kevin Carlberg

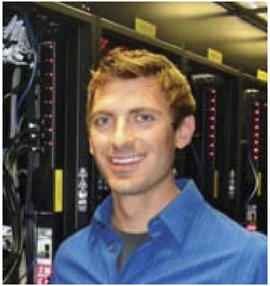

- *Background*
  - PhD Aeronautics & Astronautics, 2011, Stanford University (PhD Minor in CME)
  - Harry S Truman Postdoctoral Fellow, Sandia National Laboratories, 2011–2014
  - Principal Member of Technical Staff, Sandia National Laboratories, 2014–
- *Research interests*: nonlinear model reduction, computational mechanics, machine learning, uncertainty quantification, numerical optimization, Krylov-subspace methods, time-parallel methods

Pre-course feedback

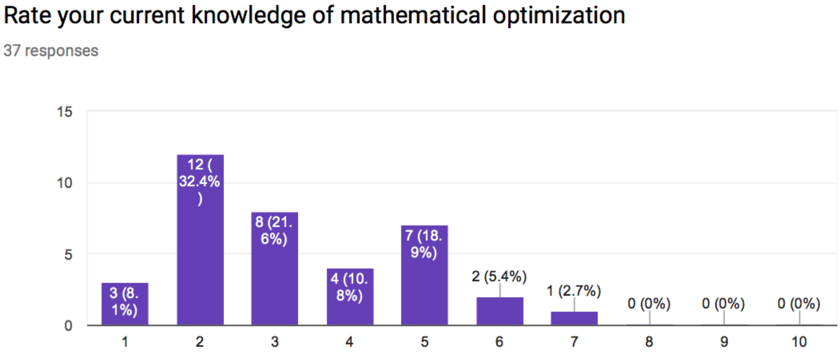

- I will not assume you have any previous knowledge
- Yet, we will cover some advanced topics and tools for the 5-7 people

Pre-course feedback

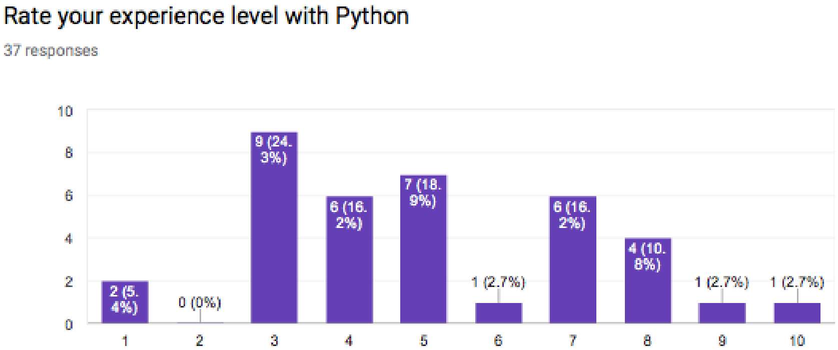

- Provided ipython notebooks that can be run without experience with Python
- Those with Python experience can quickly figure out how to use optimization tools

Pre-course feedback

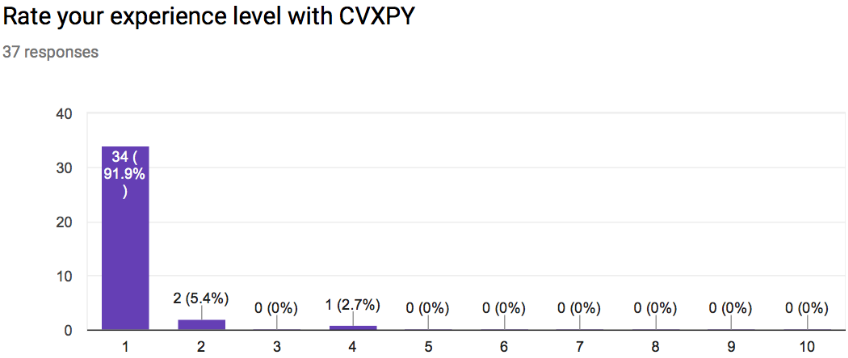

- No previous knowledge is needed
- CVXPY has a shallow learning curve: designed to allow you to ''express your problem in a natural way''

Pre-course feedback

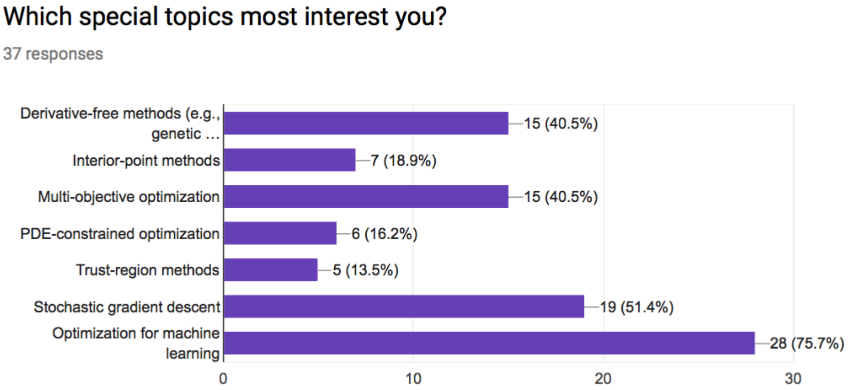

- We will briefly cover optimization aspects of machine learning
    - stochastic gradient descent
    - neural networks
    - distributed optimization

Pre-course feedback

- Lots of interest in machine learning, Python, general knowledge

Course materials

- You may download the course materials from `https://goo.gl/CbuBwm`
- These include:
  1. Slides
  2. iPython notebooks
  3. html exports of iPython notebooks
- html exports are useful for following along if you don't want to run code interactively, weren't able to install CVXPY, etc.

Rough outline

Session 1: 9:00am–10:30am

- Introduction to optimization (2_introduction.pdf)
- Unconstrained optimization (3_unconstrained.pdf)

Session 2: 10:45am–12:00pm

- Optimization in Python (4_optimization-in-python.pdf)
- Constrained optimization (5_constrained.pdf)

Session 3: 2:00pm–3:15pm

- Optimization for machine learning (6_optimization-for-ml.pdf)

Session 4: 3:15pm–4:45pm

- Convex optimization (7_convexity.pdf)
- Convex-optimization examples (8_modeling.pdf)

ICME at Stanford

- Home at Stanford for a critical multidisciplinary field that uses advanced mathematical and computing capabilities to understand and solve big, complex problems
- Over 50 faculty affiliated from 20+ departments
- MS and PhD students apply studies to wide range of domains
- MS tracks in Data Science, Computational Finance, Geoscience, Imaging Sciences
- World-class Facilities, such as the HIVE Immersive Visualization Environment, and High Performance GPU Cluster
- Offering MS and PhD degrees in computational mathematics for over 30 years: first as part of CS; then 14 years ago ICME spun off into separate institute
- Over 60 courses for graduate students as well as some undergraduates

ICME at Stanford: external partners

- These Fundamentals of Data Science workshops are part of the ICME program for **external partners.**
- We have events throughout the year around computational math, data science, machine learning and related fields.
- If you want your company to be involved in ICME collaboration, let us know!
- My organization (Sandia National Laboratories) has close connections with Stanford ICME

References

- J. Nocedal and S. J. Wright. *Numerical Optimization*, Springer, 1999.
- S. Boyd and L. Vadenberghe. *Convex Optimization*, Cambridge University Press, 2004. (available online)
  - Excellent lectures by S. Boyd online
  - Class notes and lectures for EE364a, EE364b online
  - CVX101 MOOC
- P.E. Gill, W. Murray, and M.H. Wright, *Practical Optimization*, London, Academic Press, 1981.
- Bottou, L., Curtis, F.E. and Nocedal, J., 2018. Optimization methods for large-scale machine learning. SIAM Review, 60(2), pp.223-311. (available on the arXiv)
  - Advanced recent review article