Stochastic Optimization: Making Complex Design, Planning, and Operation Decisions in the Face of Uncertainty

Jim Luedtke

Industrial & Systems Engineering
University of Wisconsin-Madison
jim.luedtke@wisc.edu

Chaos and Complex Systems Seminar, January 30, 2018

Goal of This Talk

Overview of stochastic optimization

- What types of problems might it be useful for?
- Different "flavors" of stochastic optimization models
- Some sense of how they are solved

Stochastic optimization is a branch of mathematical optimization, so we'll start with that

Will avoid this!

$$\min \mathbb{E}\Big[\sum_{t\in[T]} c_t(\xi^t)^\top x_t(\xi^t)\Big]$$

$$A_t(\xi^t)x_t(\xi^t) + B_t(\xi^t)x_{t-1}(\xi^{t-1}) = b_t(\xi^t), \quad \forall t \in [T], \mathbb{P}\text{-a.s.}$$

$$C_t(\xi^t)x \geq d_t(\xi^t), \quad \forall t \in [T], \mathbb{P}\text{-a.s.}$$

But, it's also not going to be (primarily) an application talk

Outline

Mathematical Optimization

Two-stage Stochastic Optimization

Limiting Risk

Multi-stage Stochastic Optimization

Solution Methods

Mathematical Optimization

1. Decision variables: Values/decisions to be determined by the model
2. Objective: Minimize/maximize profit, time, energy, cost... (function of decision variables)
3. Constraints: Restrictions on what values decision variables can take (inequalities/equations)

Mathematical Optimization Problem

Find values for the decision variables that satisfy all the constraints and achieve the best possible objective value.

- Optimization $\Rightarrow$ REALLY the best value (or bounds on error)
- Solution methods vary greatly, depending on structure

Historically (and often still used): Mathematical Programming, Linear Programming, Nonlinear Programming, Stochastic Programming.. ("Programming" $\equiv$ "Planning")

Example 1: Power Grid Economic Dispatch

Problem solved by Independent System Operators every 5 minutes

- Given power demands and renewable energy inputs at points in grid
- Determine generation amounts at gas/coal plants, amount to buy from spot market
- Minimize cost
- Do not exceed line limits, generation limits, etc.

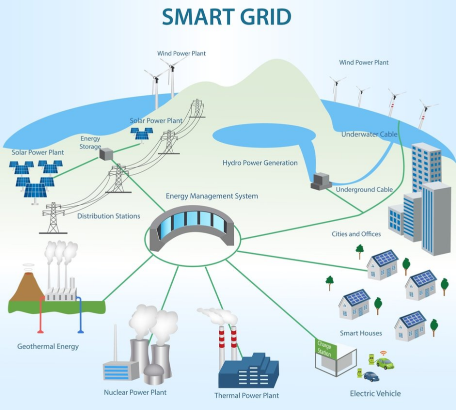

Example 2: Scheduling Service System Employees

Service Systems: Call centers, hospital departments, repair shops, etc.

- Given estimated hourly service demands throughout next week
- Determine employee schedules
- Minimize overtime cost, unmet customer demands
- Limited by number of employees having different skill sets, schedules must meet certain rules, etc.

Example 3: Wildfire Initial Response Planning

Pre-positioning firefighting equipment (dozers, etc.) to be ready for "initial response"

- Given available equipment and current locations
- Determine where equipment should be placed
- Maximize ability to respond to fires "fast enough"
- Limited by budget, amount of equipment, space at locations

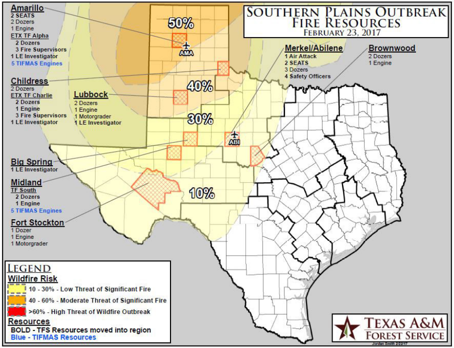

Classes of Optimization Models

Continuous

- Decision variables can take on any real values
- Infinitely many solutions
- Methods based on iterative updates to solution, e.g., using function derivatives (calculus)

Discrete/integer

- Decision variables restricted to be integer $\Rightarrow$ Can model yes/no with $0/1$ variables
- May be finite set of solutions, but too many to enumerate
- Methods based on continuous relaxations and "smart search"

Wide variation in difficulty!

- Some problems are "well-solved" (polynomial-time algorithms)
- Some problems are "theoretically hard" (NP-hard, etc.)
- Even if a problem is "hard", in many cases optimization algorithms can *optimally* solve most practical instances
  - Pet peave: "Problem is NP-hard, so there is no hope to solve it optimally."

Uncertainty in Optimization Models

Often data in a model is not perfectly known when solving the model

- Measurement errors
- Future events

Examples:

- Energy demand and wind/solar outputs
- Customer volume in service systems
- Forest fire locations and size
- Investment returns
- Cell response to enzyme changes

Ignore Uncertainty?

The "Flaw" of Averages

- The flaw of averages occurs when uncertainties are replaced by "single average numbers" planning.
- Joke: Did you hear the one about the statistician who drowned fording a river with an average depth of three feet.

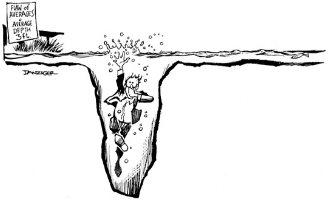

Uncertainty in Optimization Models

How to incorporate uncertainty into the model?

- Assume uncertain outcomes are random variables $\Rightarrow$ Stochastic optimization
- Assume uncertain outcomes lie within some known set and protect against worst-case $\Rightarrow$ Robust optimization

Outline

Mathematical Optimization

Two-stage Stochastic Optimization

Limiting Risk

Multi-stage Stochastic Optimization

Solution Methods

Two-stage Stochastic Optimization

Classic two-stage framework

1. Choose "here and now" decisions

$\Rightarrow$ Observe random variables

2. Make "recourse" decisions (in response to observed random variables)

Goal: Choose current decisions to minimize immediate cost plus expected value of cost of "best response" decisions

- Or, maximize expected profit, etc.
- Later: Goals other than expected value

Power Grid Unit Commitment

Daily/Weekly problem for independent system operators

- Many generators require significant time/cost to "turn on" and "turn off"
- Need to schedule the on/off status of these in advance (e.g., on hourly basis, for next day or week) $\Rightarrow$ "Commitment decisions"

Two-stage stochastic optimization model

- Here and now decisions: Which generators to "turn on/off" and when
- Random variables: Electric load and renewable generation, in each time period and each location in grid
- Recourse decisions: Economic dispatch, i.e., decide generation levels (for generators that are on)
- Minimize total expected cost
- Note: Recourse has multiple periods, but if they are independent they can be considered a single decision stage

Service System Scheduling

Service systems with multiple employee and customer "types"

- Employee "type" based on which customers they can serve
- Schedules must be made in advance, but assignment of servers to customers can be done real-time

Two-stage stochastic optimization model

- Here and now decisions: Employee schedules (e.g., for next week)
- Random variables: Number of each customer type arriving in each period
- Recourse decisions: Assign customers to available employees, determine "lost" customers
- Minimize expected number of lost customers
- Again: Recourse has multiple periods, but if they are independent, can be considered a single decision stage

Wildfire Initial Response Planning

Two-stage stochastic optimization model

- Here and now decisions: Where to place firefighting resources
- Random variables: Location and size of fires
- Recourse decisions: Determine which equipment to use to respond to fires, measure size of uncontained fire
- Minimize expected amount of uncontained fires

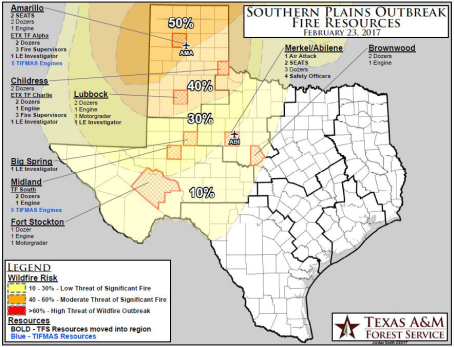

Outline

Mathematical Optimization

Two-stage Stochastic Optimization

Limiting Risk

Multi-stage Stochastic Optimization

Solution Methods

Risk-Averse?

Expected cost/profit is often appropriate objective

- E.g., when optimizing operational decisions that will be repeated
- Typically leads to "safer" solutions than ignoring uncertainty!

But, solutions good on average may still have undesirable risk of "bad outcome"

- E.g., Two options for investing \$1000, each with two equally likely outcomes

1. Stock: Lose \$200, or gain \$300 (expected gain = \$50)
2. Bitcoin: Lose \$900, or gain \$2000 (expected gain = \$550)

Risk Measures

Random outcome $\Leftrightarrow$ Distribution of possible values

- Expected value summarizes distribution with a single number, by averaging over all the outcomes

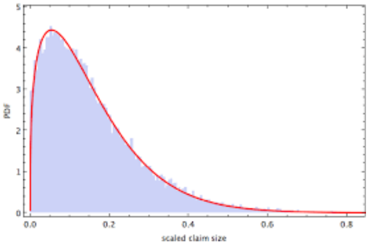

Risk measures

- Alternative ways to summarize distribution
- E.g., measure distribution spread (variance), or focus on the "bad events"

Risk Measures

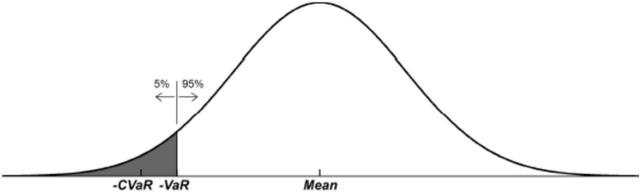

Example: Conditional value-at-risk

- Average over the 5% of worst outcomes (e.g., losses)
- Idea: Good outcomes are anyway good, more important to choose solution that is better in the bad cases

Risk Measures

Typical use in stochastic optimization model:

- Minimize expected cost, with bound on risk: Repeat with many values
- Construct an "efficient frontier" of pareto optimal solutions
- Idea dates to Markowitz' classic mean/variance portfolio model

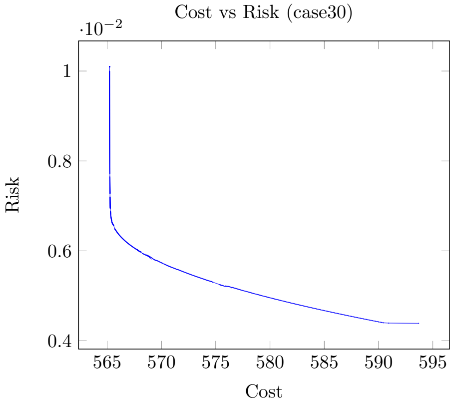

Chance Constraints

Sometimes difficult to quantify "cost" of a bad outcome

- Power line limit exceeded: May not fail at all, may lead to cascading failure
- Service system scheduling: Poor service coverage $\Rightarrow$ Lose customer "goodwill"
- Wildfire initial response planning: Difficult to predict magnitude of fires for which initial response failed

Chance constraint

Restrict the probability of undesirable event to be below a limit, ϵ

- Probability(line limit exceeded) ≤ 0.01
- Probability(any customers unserved) ≤ 0.05
- Probability(any fire not contained) ≤ 0.10

How to choose limit ϵ?

- Try many and construct efficient frontier of pareto optimal solutions

Outline

Mathematical Optimization

Two-stage Stochastic Optimization

Limiting Risk

Multi-stage Stochastic Optimization

Solution Methods

Multi-stage Stochastic Optimization Problem

Finite-horizon sequential decision making problems under uncertainty

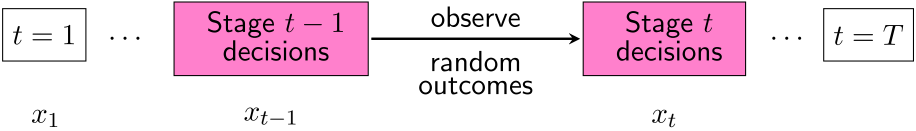

When making decisions in stage t:

- Must anticipate full sequence of future random events and optimal responses to those

Examples: When would they be multi-stage?

- Power grid unit commitment: If can adjust future commitment decisions every hour (or day)
- Employee scheduling: Recourse is multi-stage if customers willing to wait from one period to next
- Wildfire initial response planning: Move equipment periodicallly based on evolving availability

Outline

Mathematical Optimization

Two-stage Stochastic Optimization

Limiting Risk

Multi-stage Stochastic Optimization

Solution Methods

First Challenge: Evaluating Expected Value

Stochastic optimization models typically have many random variables

- Need estimate of expected value of objective *as function of decision variables*
- Exact evaluation impossible even for a single set of decision variable values
- Similar challenge for chance constraints (calculate probability of event) or risk measures

Typical appraoch: "Sample average approximation"

- Approximate vector of random variables with finite set of "scenarios"
- Expected value $\Rightarrow$ Weighted sum
- Often scenarios from a Monte Carlo sample, but many more advanced approaches

Sample average approximation $\Rightarrow$ Deterministic, but very large-scale optimization model

Approximating Expected Value

Key question: How many scenarios required for "good approximation"?

- Significant research into this for variety of problems (two-stage, chance constraints, risk measures, multi-stage)
- Good news: Surprisingly, required number grows "mildly" with number of decision variables and random variables
- Bad news: Required number grows fast with desired accuracy

Conclusion

In many cases sampling enables solving stochastic optimization problems to "modest accuracy"

- Exception: For multi-stage problems, sample size grows exponentially with number of stages

Next challenge: How to solve the very large-scale optimization model defined by sample average approximation?

Solving Sample Average Approximation: Two-Stage Problems

Large-scale because model must account for actions in every scenario

- Structure: With first-stage decisions fixed, each scenario can be considered independently
- Single problem of size $N \times m \Rightarrow N$ separate problems of size m

Algorithms exploit this structure via decomposition

1. Choose first-stage decisions by solving a "master problem"
2. Solve recourse problem for each scenario to evaluate first-stage decisions
3. Collect information from recourse problems to update master problem, and repeat

Similar idea for problems with risk measures, chance constraints

Other Approaches: Two-Stage Problems

Alternative decomposition strategy:

- Solve full problem (first and second-stage) separately for each scenario
- Issue: Different here-and-now decisions in different scenarios
- Average the here-and-now decisions $\Rightarrow$ "Consensus" decision
- Re-solve separate subproblem, but with penalty for straying from consensus

Stochastic approximation (stochastic gradient descent):

- Sampling is integrated in algorithm

Multi-Stage Problems

Frequent simplifying assumption: Random variables in different stages are independent

- Often can vbe satisfied with appropriate modeling
- Many methods based on recursive approximation of “cost-to-go” function
- Related approach: (Deep) reinforcement learning (AlphaGo)

Approximation by restricting flexibility in decisions

- Require all (or some) decisions to be linear function of observed random variables
- Restricted problem is one or two-stage problem
- Similar techniques in “dual” problem yields bounds on solution quality

Open Challenges

Many areas for ongoing work

- Stochastic + discrete
- Multi-stage
- Chance constraints with small risk tolerance
- High-impact, rare events
- Integrating stochastic optimization and machine learning/prediction models (data-driven)
- Use it for...

Questions?

jim.luedtke@wisc.edu