Stochastic Optimization

L-Shaped Method

CSci 688 — Frans Schalekamp (frans@wm.edu) — Spring 2015

Stochastic Program

Consider the Stochastic Program

$$\begin{aligned}
\text{minimize } & c^T x + E_\xi(q(\omega_i)^T y(\omega_i)) & \\
\text{s.t. } & Ax = b & \\
& Wy(\omega_i) + T(\omega_i)x = h(\omega_i) & \text{almost surely} \\
& x \geq 0 & \\
& y(\omega_i) \geq 0 & \text{almost surely.}
\end{aligned}$$

The L-shaped method is a way of solving a Stochastic Program in the special case that the random variable ξ takes on finitely many values, say $\{\omega_1, \omega_2, \ldots, \omega_n\}$. In this case we can write the Stochastic Program as a huge Linear Program:

$$\begin{aligned}
\text{minimize } & c^T x + \sum_i p(\omega_i) q(\omega_i)^T y(\omega_i) & \\
\text{s.t. } & Ax = b & \\
& Wy(\omega_i) + T(\omega_i)x = h(\omega_i) & \text{for all } \omega_i \\
& x \geq 0 & \\
& y(\omega_i) \geq 0 & \text{for all } \omega_i.
\end{aligned}$$

where $p(\omega_i) = P(\xi = \omega_i)$, the probability that the random variable ξ is ω_i.

The key insight underlying the L-shaped Method is that if x is fixed, then we could find each $y(\omega_i, x)$ by just solving the following small LP for each ω_i separately:

$$\begin{aligned}
\text{minimize } & q(\omega_i)^T y(\omega_i, x) \\
\text{s.t. } & Wy(\omega_i, x) = h(\omega_i) - T(\omega_i)x \\
& y(\omega_i, x) \geq 0.
\end{aligned}$$

We are going to replace the second stage cost $q(\omega_i)^T y(\omega_i)$ in the objective of the Linear Program by a variable ϑ_i for each ω_i, and use duality to get bounds on these ϑ_i. The dual of the second stage LP for a given ω_i is

$$\begin{aligned}
\text{maximize } & (h(\omega_i) - T(\omega_i)x)^T \pi(\omega_i) \\
\text{s.t. } & W^T \pi(\omega_i) \leq q(\omega_i) \\
& \pi(\omega_i) \text{ unrestricted.}
\end{aligned}$$

Weak duality tells us that (for a given x), for any primal feasible $y(\omega_i, x)$ and dual feasible $\pi(\omega_i)$ we have

$$q(\omega_i)^T y(\omega_i, x) \geq (h(\omega_i) - T(\omega_i)x)^T \pi(\omega_i). \quad (1)$$

So the constraints that we want to add to the LP when we are replacing the second stage cost $q(\omega_i)^T y(\omega_i)$ in the objective of the Linear Program by a variable ϑ_i for each ω_i, are

$$\vartheta_i \geq (h(\omega_i) - T(\omega_i)x)^T \pi(\omega_i)$$

for all feasible dual solutions $\pi(\omega_i)$. In other words, every feasible solution $\pi(\omega_i)$ gives rise to an inequality for the variables ϑ_i and x, as is emphasized by rewriting the inequality as:

$$\vartheta_i \geq (h(\omega_i) - T(\omega_i)x)^T \pi(\omega_i) = \pi(\omega_i)^T (h(\omega_i) - T(\omega_i)x) = \pi(\omega_i)^T h(\omega_i) - \pi(\omega_i)^T T(\omega_i)x$$

or, putting all variables on the left hand side of the inequality:

$$\vartheta_i + \pi(\omega_i)^T T(\omega_i)x \geq \pi(\omega_i)^T h(\omega_i)$$

for all feasible dual solutions $\pi(\omega_i)$ to the second-stage LP for scenario ω_i. Note that we have such a family of constraints for every scenario ω_i.

We have thus derived the following LP, which I will refer to as the "LP with the Ridiculous Number of Constraints".

$$\begin{aligned}
\text{minimize } & c^T x + \sum_i p(\omega_i)\vartheta_i \\
\text{s.t. } & Ax = b \\
& \vartheta_i + \pi(\omega_i)^T T(\omega_i)x \geq \pi(\omega_i)^T h(\omega_i) && \text{for all feasible dual solutions } \pi(\omega_i) \\
& && \text{to the second-stage LP for scenario } \omega_i, \\
& && \text{for all scenarios } \omega_i \\
& x \geq 0.
\end{aligned}$$

L-Shaped Method, Multicut Version

Because the "LP with the Ridiculous Number of Constraints" has a ridiculous number of constraints, we want to find a way so that we only have to consider LPs that have only a reasonably sized subset of these constraints. This is what the L-shaped Method is about: it iteratively adds these type of constraints, namely exactly those that correspond to the optimal dual solutions to the second-stage LP for all scenarios ω_i for a given x.

So the method is the following: it starts with the following LP to which we will add constraints, and which we will refer to as the "Master LP":

$$\begin{aligned}
\text{minimize } & c^T x \\
\text{s.t. } & Ax = b \\
& x \geq 0.
\end{aligned}$$

(Note that in this first iteration there are no bounds on the ϑ_i variables, so we drop these variables from the objective in this iteration.) We find an optimal $\bar{x}$. We then find the optimal dual solutions to the second-stage LP for all scenarios ω_i for this $\bar{x}$ — let us call these $\tilde{\pi}^1(\omega_i)$. All these are feasible dual solutions, and so define a new constraint for ϑ and x. We add all of these constraints to the "Master LP", which now looks like

$$\begin{aligned}
\text{minimize } & c^T x + \sum_i p(\omega_i)\vartheta_i \\
\text{s.t. } & Ax = b \\
& \vartheta_i + \tilde{\pi}^1(\omega_i)^T T(\omega_i)x \geq \tilde{\pi}^1(\omega_i)^T h(\omega_i) && \text{for all scenarios } \omega_i \\
& x \geq 0.
\end{aligned}$$

We solve this LP with the new constraints, giving an optimal $\bar{x}$ and $\bar{\vartheta}_i$ for all i. We again find the optimal dual solutions to the second-stage LP for all scenarios ω_i for this $\bar{x}$ — let us call these $\tilde{\pi}^2(\omega_i)$.

All these are feasible dual solutions, and so define a new constraint for ϑ and x. We add all of these constraints to the "Master LP", which now looks like

$$
\begin{aligned}
\text{minimize } & c^T x + \sum_i p(\omega_i)\vartheta_i \\
\text{s.t. } & Ax = b \\
& \vartheta_i + \tilde{\pi}^1(\omega_i)^T T(\omega_i)x \geq \tilde{\pi}^1(\omega_i)^T h(\omega_i) && \text{for all scenarios } \omega_i \\
& \vartheta_i + \tilde{\pi}^2(\omega_i)^T T(\omega_i)x \geq \tilde{\pi}^2(\omega_i)^T h(\omega_i) && \text{for all scenarios } \omega_i \\
& x \geq 0.
\end{aligned}
$$

The method repeats these steps (solving Master LP, finding an optimal $\bar{x}$ and $\bar{\vartheta}_i$ for all i, solving all second-stage LPs for this $\bar{x}$, adding the constraints that correspond to the optimal dual solutions to the Master LP), until it finds that $\bar{x}$ and $\bar{\vartheta}_i$ satisfy all newly added constraints..

Proof of Optimality

Theorem 1. *If the L-shaped Method ends then the final $\bar{x}$ is an optimal first-stage decision for the Stochastic Program.*

Lemma 1. *Optimal objective value of the "LP with the Ridiculous Number of Constraints" is at most the optimal objective value of the Stochastic Program.*

Proof. Let $\bar{x}$, $\bar{y}(\omega_i)$ be an optimal solution to the Stochastic Program. Define $\bar{\vartheta}_i = q(\omega_i)^T \bar{y}(\omega_i)$ for all i. Note that $\bar{x}$, $\bar{\vartheta}_i$ is a feasible solution to the "LP with the Ridiculous Number of Constraints" by weak duality (see equation (1)), with the same objective. In other words, for any solution to the Stochastic Program, we can define a solution to the "LP with the Ridiculous Number of Constraints" with the same objective, so the optimal objective value "LP with the Ridiculous Number of Constraints" can not be more than the optimal objective value of the Stochastic Program. □

Lemma 2. *Optimal objective of the Master LP is at most the optimal objective of the "LP with the Ridiculous Number of Constraints".*

Proof. The Master LP only includes a subset of the constraints of the "LP with the Ridiculous Number of Constraints". Therefore any feasible solution to "LP with the Ridiculous Number of Constraints" is also feasible to the Master LP (in any iteration). □

Lemma 3. *If $\bar{x}$ and $\bar{\vartheta}_i$ satisfy all newly added constraints in an iteration of the L-shaped Method, then there exist $\bar{y}(\omega_i)$ so that $\bar{x}, \bar{y}(\omega_i)$ is an optimal solution to the Stochastic Program.*

Proof. Let $\tilde{\pi}(\omega_i)$ be the optimal dual second-stage solutions that were used to generate the new constraints, for all i. By strong duality, there exist feasible primal second-stage solutions $\bar{y}(\omega_i, \bar{x})$ so that $q(\omega_i)^T \bar{y}(\omega_i, \bar{x}) = (h(\omega_i) - T(\omega_i)\bar{x})^T \tilde{\pi}(\omega_i)$ for all i.

The inequalities that were added in the last iteration are $\vartheta_i \geq (h(\omega_i) - T(\omega_i)x)^T \tilde{\pi}(\omega_i)$, and since $\bar{\vartheta}_i$ and $\bar{x}$ satisfy these inequalities we therefore conclude that there exist feasible primal second-stage solutions $\bar{y}(\omega_i, \bar{x})$ so that $\bar{\vartheta}_i \geq q(\omega_i)^T \bar{y}(\omega_i, \bar{x})$ for all i.

The objective of the Stochastic Program for this feasible solution $\bar{x}$, $\bar{y}(\omega_i)$ is $c^T\bar{x} + \sum_i p(\omega_i) q(\omega_i)^T \bar{y}(\omega_i, \bar{x}) \leq c^T\bar{x} + \sum_i p(\omega_i)\bar{\vartheta}_i$ = optimal objective of the Master LP $\leq$ optimal objective of the "LP with the Ridiculous Number of Constraints" $\leq$ optimal objective value of the Stochastic Program.

In other words, we have found a feasible solution to the Stochastic Program with objective at least as good as the objective of an optimal solution, which means this solution must be an optimal solution. □

Finite Running Time

Note that if we can choose the optimal second-stage dual solutions to be extreme points in the L-shaped method, then the L-shaped Method is guaranteed to be finite. This is because the number of extreme points of each of the LPs is finite, and therefore the number of constraints that can be added is finite. After a finite number of iterations the newly added constraints were already presented at the start of the iteration, and $\bar{x}$ and $\bar{\vartheta}$ satisfy all constraints.

Infeasible Second-Stage LPs

Note that we can always choose the optimal second-stage dual solutions to be extreme points in the L-shaped method if the optimal solutions of the second-stage dual LPs are bounded (finite). The only thing to worry about is when these LPs are not bounded. By duality we know that if the second-stage dual LP is unbounded, the second-stage primal LP is infeasible. In other words, we actually want to exclude first-stage decisions x for which this is the case. This means it suffices to generate a constraint for the Master LP which cuts off $\bar{x}$, in case we detect that some second-stage dual LP is unbounded.

If we find that, given an $\bar{x}$, the second-stage dual LP for some ω_i is unbounded, then we can find a ray $\bar{r}(\omega_i)$ so that $(h(\omega_i) - T(\omega_i)\bar{x})^T \bar{r}(\omega_i) > 0$. This means that to exclude this first-stage decision $\bar{x}$ from consideration in the Master LP, it suffices to add the constraint

$$(h(\omega_i) - T(\omega_i)x)^T \bar{r}(\omega_i) \leq 0.$$

To emphasize that this is an equality for the variables x we rewrite the inequality as

$$(h(\omega_i) - T(\omega_i)x)^T \bar{r}(\omega_i) = \bar{r}(\omega_i)^T (h(\omega_i) - T(\omega_i)x) = \bar{r}(\omega_i)^T h(\omega_i) - \bar{r}(\omega_i)^T T(\omega_i)x \leq 0,$$

or

$$\bar{r}(\omega_i)^T T(\omega_i)x \geq \bar{r}(\omega_i)^T h(\omega_i).$$

Note that if we modify the L-shaped method so that it adds this constraint for an extreme ray when it finds an unbounded solution to a dual of a second-stage LP (if there exists a ray, then there exists an extreme ray), and it uses an extreme point for the second-stage dual solution if it is bounded, then this (modified) L-shaped method is guaranteed to terminate, because there are only finitely many extreme rays and extreme points.

Normal (One Cut) L-Shaped Method

Instead of having a variable ϑ_i for every scenario ω_i in the Master LP, we could alternatively think about having only one variable ϑ representing the convex combination of these (with weights equal to the probabilities $p(\omega_i)$). Instead of adding a constraint for every second-stage LP in an iteration of the L-shaped Method, you now add either a constraint that follows from the infeasiblity of a second-stage LP, or one constraint if all second-stage LPs are feasible, namely the same convex combination of the constraints that you would add in the multicut version, i.e.

$$\sum_i p(\omega_i)\vartheta_i + \sum_i p(\omega_i)\tilde{\pi}(\omega_i)^T T(\omega_i)x \geq \sum_i p(\omega_i)\tilde{\pi}(\omega_i)^T h(\omega_i)$$

and thus (because you are replacing $\sum_i p(\omega_i)\vartheta_i$ by one variable, namely ϑ)

$$\vartheta + \sum_i p(\omega_i)\tilde{\pi}(\omega_i)^T T(\omega_i)x \geq \sum_i p(\omega_i)\tilde{\pi}(\omega_i)^T h(\omega_i).$$

The proof of optimality and proof of termination for this procedure are analogous to the proofs for the multicut version.

Exercises

1. Discuss the difference between the multicut and the normal L-shaped methods, in terms of (1) the number of variables, (2) a bound on the potential number of iterations, and (3) a bound on the number of iterations of the methods, such that it will surely outperform just solving the Linear Program

$$\begin{aligned} \text{minimize } & c^T x + \sum_i p(\omega_i) q(\omega_i)^T y(\omega_i) \\ \text{s.t. } & Ax = b \\ & W y(\omega_i) + T(\omega_i) x = h(\omega_i) && \text{for all } \omega_i \\ & x \geq 0 \\ & y(\omega_i) \geq 0 && \text{for all } \omega_i. \end{aligned}$$

Assume that solving a linear program with n variables takes $n^{3.5}$ time.

Example

The following are illustrations to accompany the example of the Normal (One Cut) L-shaped Method, that is in the textbook of Birge & Loveaux (pages 184 and beyond).

In the following figures the new constraint that is created in every iteration of the L-shaped Method is depicted with a blue plane. Any feasible combination of x_1, x_2 and ϑ has a ϑ value that is at least as large as the corresponding value on the plane (i.e., the feasible x_1, x_2 and ϑ lie above the plane). The constraints are then combined to one (red) boundary, and the Master Problem is solved. The solution to the Master Problem is indicated by a green box.

The first figure illustrates the inequality that is generated after the first iteration. This then is the boundary of the feasible region of the Master Problem in the second iteration, and the green box ($\bar{x}_1 = 40$, $\bar{x}_2 = 80$, $\bar{\vartheta} = -18299.2$) is an optimal solution to this LP.

The constraint that is created in this second iteration of the L-shaped Method is depicted with a blue plane in the next figure. (You can see that the optimal solution is cut off by this blue plane.)

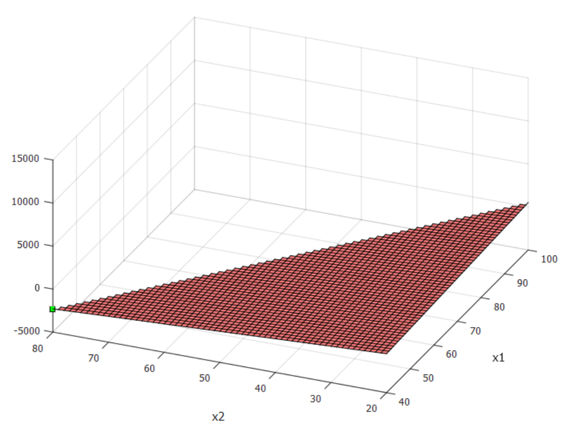

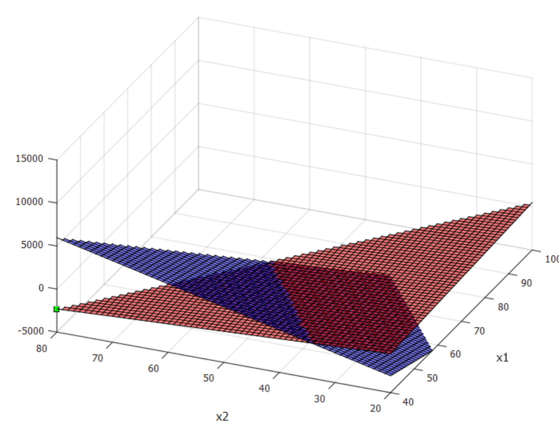

The constraints are then combined to one (red) boundary, and the Master Problem is solved. The solution to the Master Problem (in this third iteration) is indicated by a green box.

Solving the second-stage problem gives new extreme points for the dual LPs for the second-stage problems, which define a new constraint on x_1, x_2 and ϑ (drawn again in blue).

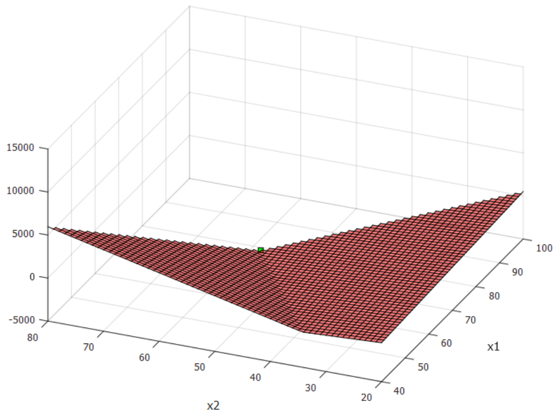

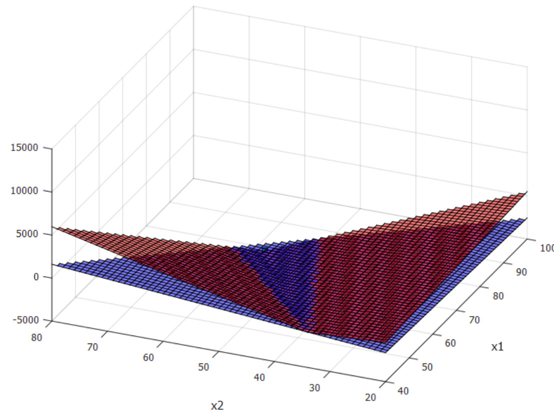

The method continues until the newly created constraint does not cut off the optimal solution of the Master Program.

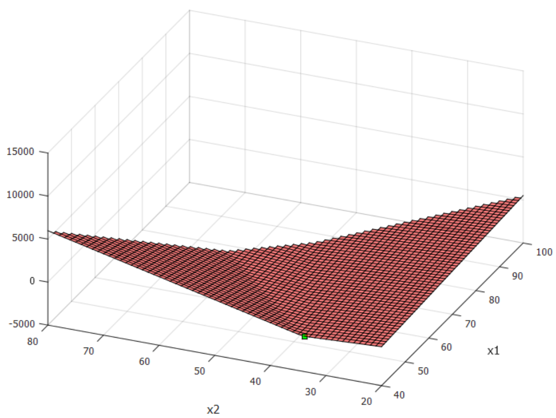

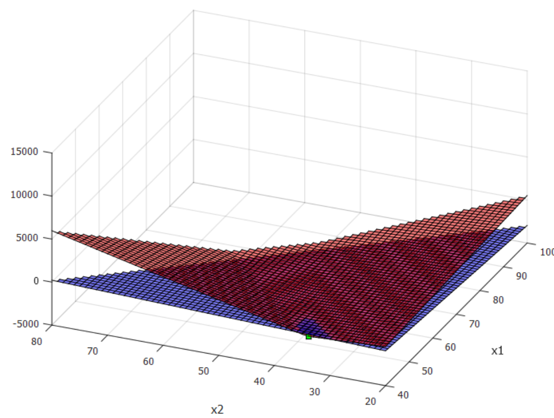

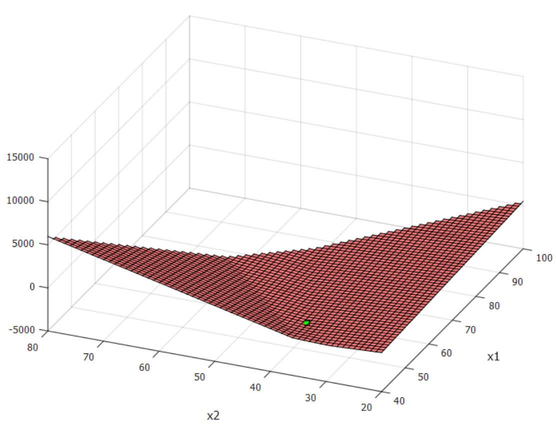

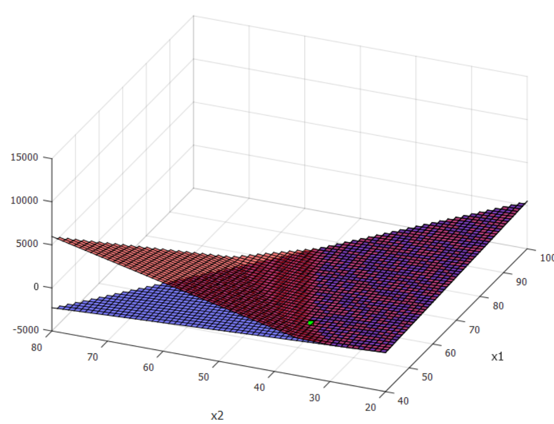

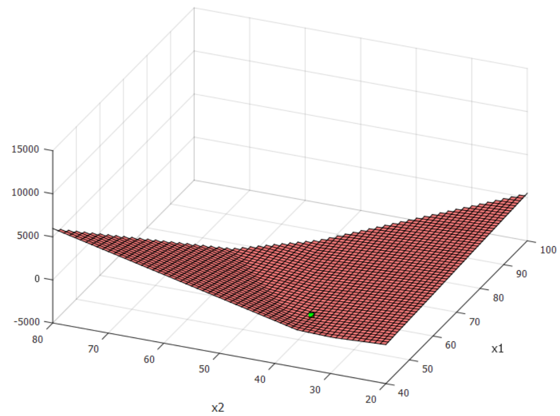
15000
10000
5000
0
-5000
80
70
60
50
40
30
20
x2
40
50
60
70
80
90
100
x1