Thin Models for Big Data

Matteo Fischetti, University of Padova



Big Data

- Big data is a broad term for data sets so large and complex that traditional data processing applications are inadequate
- Multimodal (incomplete) data from any kind of sources (physical/social systems etc.)
- Not just a matter of **size** ...



 We (OR/Mathematical Optimization people) are already proud if we can solve problems with few MB's of perfectly clean input data... #TooBigForUs?

Big Data in the realm of Analytics





•

17

★ 17

...

Jean-François Puget @JFPuget · 16 ago The Future of #Analytics Is Prescriptive, Not Predictive clickz.com/2416723 #orms Well said.

#WeArePrescriptive



Jean-François Puget

Following

Ö

Prescriptive Analytics Is Easier And More Profitable Than Predictive #Analytics ibm.com/developerworks ... #orms

S Visualizza traduzione



03:15 - 6 ago 2015

It turns out that

Prescriptive Analytics

(large scale) Mathematical Optimization

we are

no longer old-fashioned

but actually have a LOT to say about Big Data

But ... are we "scalable enough"?

- Prescriptive Analytics calls for **scalable** methods to deal with **larger and larger problems**
- Operations Research / Mathematical Optimization models and algorithms can be inherently rather **sophisticated**
- Challenge: improve scalability of OR heuristic and exact methods
- Simplicity is one of keys to scalability → focus (as much as possible) on simple models and solution schemes
- Warning: **simple** does not mean trivial ...



Occam's razor

• **Occam's razor**, or law of parsimony (lex parsimoniae): a problem-solving principle devised by the English philosopher William of Ockham (1287–1347).



- Among competing hypotheses, the one with the fewest assumptions is more likely be true and should be preferred—the fewer assumptions that are made, the better.
- The simpler (the model, the algorithm) the better
- Used as a heuristic guide in the development of theoretical models in physics (Albert Einstein, Max Planck, Werner Heisenberg, etc.)
- Not to misinterpreted and used as an excuse to address oversimplified models: "Everything should be kept as simple as possible, but no simpler" (Albert Einstein)

Thinning out optimization models

• The practical difficulty in solving hard problems sometimes comes for **overmodelling**:

Too many vars.s and constr.s just

suffocate the model

(and the cure is <u>not</u> to complicate it even more!)

Let your model breathe!



- Simpler and more effective models can sometimes be obtained by:
 - 1. Choosing a model that better fits the **instances** of interest
 - 2. Removing variables that play a little role for the **problem class** of interest
 - 3. Using **decomposition** to break the problem into smaller pieces

Example 1: QAP

- Quadratic Assignment Problem (QAP): extremely hard to solve
- Unsolved esc* instances from QAPLIB (attempted on constellations of thousand computers around the world for many CPU years, with no success)
- The thin out approach: *esc* instances are
 - very large \rightarrow use **slim MILP** models with high node throughput
 - decomposable \rightarrow solve pieces **separately**
 - very symmetrical → find a cure and simplify the model through Orbital Shrinking to actually reduce the size of the instances

Fischetti, L. Liberti, "Orbital shrinking", Lecture Notes in Computer Science, Vol. 7422, 48-58, 2012.

Outcome:

- a. all esc* but two instances solved in minutes on a notebook
- b. esc128 (by far the largest QAP ever attempted) solved in just seconds
- M. Fischetti, M. Monaci, D. Salvagnin, "Three ideas for the Quadratic Assignment Problem", Operations Research 60 (4), 954-964, 2012.

Example 2: Steiner Trees

- Recent **DIMACS 11 (2014)** challenge on Steiner Tree Problems: various versions and categories (exact/heuristic/parallel/...) and scores (avg/formula 1/ ...)
- Standard MILP models use x var.s (arcs) and y var.s (nodes)
- Many very hard (unsolved) instances available on STEINLIB
- **Observation**: many hard instances have uniform arc costs
- **Thin out**: remove *x* var.s and work on the *y*-space (kind of Benders' projection)
- Heuristics based on the **blur** principle: initially forget about details...
- Vienna-Padua team → MozartBalls code
- Outcome:
 - Some open instances solved in a few **seconds**
 - MozartBalls ranked first in most DIMACS categories



M. Fischetti, M. Leitner, I. Ljubic, M. Luipersbeck, M. Monaci, M. Resch, D. Salvagnin, M. Sinnl, "Thinning out Steiner trees: a node-based model for uniform edge costs", Tech.Rep., 2014

Example 3: facility location

- Uncapacitated facility location with linear (UFL) and quadratic (qUFL) costs
- **Huge** MILP models involving y var.s (selection) and x var.s (assignment)
- **Thin out**: assignment var.s *x* **suffocate** the model, just remove them...
- A perfect fit with **Benders decomposition** (more later!)
- Outcome:
 - Many hard UFL instances solved in just **seconds** on a notebook
 - Seven open instances solved to optimality, 22 best-known improved
 - Speedup of **4 orders of magnitude** for qUFL up to size 150x150
 - Solved qUFL instances up to 2,000 locations and 10,000 clients in just 5 minutes (MIQCP's with 20M SOC constraints and 40M var.s)

M. Fischetti, I. Ljubic, M. Sinnl, "Thinning out facilities: a Benders decomposition approach for the uncapacitated facility location problem with separable convex costs", TR 2015.



Thin out your favorite model call Benders toll free

Benders decomposition well known ... but not so many MIPeople actually use it

We will next give a brief tutorial on Modern Benders

hoping to convince young researchers to test it...

...and not to be #TooBoring



Benders in a nutshell

CLASSICAL BENDERS ('601) EFFAS y's Benders' cats for conver pr. (beofrion) $g(x,y) \leq 0 \qquad (< BenDers' cuts y A: torue:$ $h(y) \leq 0 \qquad (min f(x,y)) = 0 \qquad g(x,y) \leq 0 \qquad g(x,y) \leq 0 \qquad g(x,y) \leq 0 \qquad g(x,y) \leq 0 \qquad g' \leq y \leq y^{*}$ min $\{w: w \geq \beta_{i}^{i} + \beta_{j}^{i} y, i \geq 3, ..., n\} \qquad (a ortimal value: w(y^{*})) = 0 \qquad (y' \leq y' \leq y') \qquad (b ortimal value: w(y^{*})) = 0 \qquad (y' \leq y' \leq y') \qquad (y' \leq y') \qquad (y' \leq y') \qquad (y' \leq y') = 0 \qquad (y' \leq y') \qquad (y' = y') \qquad (y$ How to solve -> BUNDLE METHODS -> IN-OUT (SPECIALIZED)

Modern Benders

Consider the original **convex** MINLP $\min f(x, y)$ $g(x, y) \le 0$ $Ay \le b$ y integer

and assume for the sake of simplicity

 $S := \{y : Ay \le b\}$ nonempty and bounded

 $X(y):=\{x:g(x,y)\leq 0\}$ nonempty, closed and bounded for all $y\in S$

so the convex function

$$\Phi(y) := \min_{x \in X(y)} f(x, y)$$

is well-defined for all $y \in S$.

OR 2015, Vienna, Sept. 2015

13

Working on the y-space (projection)

$\min f(x,y)$	$\Phi(y) := \min_{x \in X(y)} f(x, y)$	$\min \Phi(y)$
$g(x,y) \leq 0$		$Ay \leq b$
$Ay \leq b$		$y { m integer}$
y integer		

Original MINLP in the (x,y) space \rightarrow

Master problem in the y space

Warning: projection changes the objective function shape!





Life of P(H)I

- Solving Benders' master problem calls for the minimization of a nonlinear function (even if you start from a linear problem!)
- Branch-and-cut MINLP solvers generate a sequence of linear cuts to approximate this function from below (outer-approximation)







$$w \ge \Phi(y) \ge \Phi(y^*) + \xi(y^*)^T (y - y^*)$$

Benders cut computation

• **Benders** (for linear) and **Geoffrion** (general convex) told us how to compute a **(sub)gradient** to be used in the cut derivation, by using the optimal primal-dual solution (x^*, u^*) available after computing $\Phi(y^*)$

$$\xi(y^*) = \nabla_y f(x^*, y^*) + u^* \nabla_y g(x^*, y^*)$$

- This formula is **problem-specific** and perhaps **#scaring**
- By rewriting

$$\Phi(y^*) = \min\{f(x, \mathbf{q}) \mid g(x, \mathbf{q}) \le 0, \, y^* \le \mathbf{q} \le y^*\}$$

we obtain a much **simpler recipe** to derive the same Benders cut:

- 1) solve the original convex problem with new var. bounds $y^* \le y \le y^*$
- 2) take opt_val and reduced costs r_j 's
- 3) write $w \ge opt_val + \sum_j r_j(y_j y_j^*)$

#TheCurseOfKelley

• Master problem is typically solved by a **cutting plane method** where primal (fractional) solutions *y*^{*} and Benders cuts are generated on the fly

• A main reason for Benders' slow convergence is the use of **Kelley's** cutting plane recipe "**Always cut the optimal solution of the previous master**"

• In the first iterations, the master can contain too few constraints (sometimes, only variable bounds) \rightarrow **zig-zagging** in the *y* space (lower bound stalling)

→ Stabilization required as in

Column Generation and Lagrangian Relaxation

e.g. through bundle methods



Escaping the #CurseOfKelley

- Root node LP bound very critical → many ships sank here!
- Kelley's cutting plane can be desperately slow, bundle/interior points methods required
- Stabilization using "interior points"



(Ben-Ameur and Neto 2007, Fischetti and Salvagnin 2010, Naoum-Sawaya and Elhedhli 2013).

- For facility location problems, we implemented a very simple "chase the carrot" heuristic to determine an internal path towards the optimal y
- Our very first implementation worked so well that we did not have an incentive to try and improve it **#OccamPrinciple**

Our #ChaseTheCarrot heuristic



• We (the donkey) start with y = (1, 1, ..., 1) and optimize the master LP as in Kelley, to get optimal y^* (the carrot on the stick).

• We move y half-way towards y*. We then separate a point y' in the segment y-y* close to y. The generated Benders cut is added to the master LP, which is reoptimizied to get the new optimal y* (carrot moves).

• Repeat until bound improves, then switch to Kelley for final bound refinement (kind of cross-over)

• Warning: adaptations needed if feasibility Benders cuts can be generated...

Effect of the improved cut-loop



- Comparing Kelley cut loop at the root node with Kelley+ (add epsilon to y*) and with our chase-the-carrot method (inout)
- Koerkel-Ghosh **qUFL** instance gs250a-1 (250x250, quadratic costs)
- *nc = n. of Benders cuts generated at the end of the root node
- times in logarithmic scale

Conclusions

- Operations Research (OPS! Prescriptive Analytics) can play an important role in the Datazoic era ...
 - ... giving us an incentive to design simpler and more scalable solvers
 - ... possibly based on clever decomposition
- Benders is a good boy, if you do not mistreat him!

Thanks for your attention!

Slides available at http://www.dei.unipd.it/~fisch/papers/slides/



(not too) Big Data and OR as good friends