Learning-Augmentation for Online Convex Covering and Concave Packing

Maoyuan 'Raymond' Song

Computer Science Department, Purdue University

Online Covering and Packing Linear Programs

We consider the following primal covering minimization problem:

$$\min c^T x \text{ over } x \in \mathbb{R}^n_{\geq 0} \text{ subject to } Ax \geq 1 \tag{1}$$

and the corresponding dual packing maximization problem:

$$\max \mathbf{1}^T y \text{ over } y \in \mathbb{R}^m_{\geq 0} \text{ subject to } A^T y \leq c \tag{2}$$

where:

- $A \in \mathbb{R}^{m \times n}_{\geq 0}$ is the constraint matrix;
- $\mathbf{1}$ is the vector of all ones;
- $c \in \mathbb{R}^n_{>0}$ is the linear coefficients of the cost function.

In the online setting, the rows of the constraint matrix A arrives online, and the algorithm must update x or y in a non-decreasing manner.

Primal-Dual Learning-Augmented (PDLA) Algorithms for Online Covering

We adapt and extend the *primal-dual method* to incorporate learning-augmentation and advices.

- The algorithm is given an *advice* $x' \in \mathbb{R}^n_{\geq 0}$, suggesting a solution to LP (1).
- The user chooses a *confidence parameter* $\lambda \in [0, 1]$ that controls the desired consistency-robustness tradeoff.

Performance Measures

- **Consistency**: The ratio between the value of the algorithm's solution and the value of the advice.
- **Robustness**: The ratio between the value of the algorithm's solution and the value of optimal offline solution.

Our PDLA algorithm solves the primal and the dual simultaneously. Whenever constraint $i \in [n]$ arrives online:

- Increase y_i with growth rate 1;
- Increase x_j with growth rate

$$\frac{A_{ij}}{c_j}\left(x_j + \frac{\lambda}{A_i \mathbf{1}} + \frac{(1-\lambda)x'_j \mathbf{1}_{x_j < x'_j}}{A_i x'_c}\right)$$

where x'_c is the advice x' restricted to entries j where $x_j < x'_j$ holds.

Theorem 1. *Our algorithm is $O(\frac{1}{1-\lambda})$-consistent, $O(\log \frac{\kappa n}{\lambda})$-robust, where κ is the condition number of A.*

Simple Switching Strategies

Private communication with Roie Levin pointed out that the following algorithm beats our PDLA algorithms *for all choices of* λ.

Switching Algorithm for Online Covering

Input: problem instance A, c, advice x', online algorithm $\mathcal{O}$;
Output: solution x.
When constraint i arrives:
  Run $\mathcal{O}$ for round i to obtain $x^{(i)}$.
  if $c^T x' \geq c^T x^{(i)}$ **then** $x \leftarrow x^{(i)}$.
  else for $j \in [n]$ **do**
    $x_j \leftarrow \max\{x'_j, x_j^{(i-1)}\}$.

Follow-up work devised another similar algorithm based on simple switching ideas for online packing, and can be extended to concave objective functions.

Switching Algorithm for Online Packing

Input: problem instance A^T, c, advice y', online algorithm $\mathcal{O}$;
Output: solution y.
When column i arrives:
  Run $\mathcal{O}$ for round i to obtain $y_i^{\mathcal{O}}$.
  if $A^T y' \leq c$ **then** $y_i \leftarrow \frac{1}{2}(y'_i + y_i^{\mathcal{O}})$.
  else $y_i \leftarrow y_i^{\mathcal{O}}$.

Theorem 2. *The switching algorithm for online covering LPs and the switching algorithm for online (concave) packing are both 2-consistent and 2α-robust, where α is the competitiveness of $\mathcal{O}$.*

Important Conceptual Questions

- What structural properties enable switching strategies?
- Characterizing space of problems that allows black-boxes?
- The fundamental power of learning-augmentation?

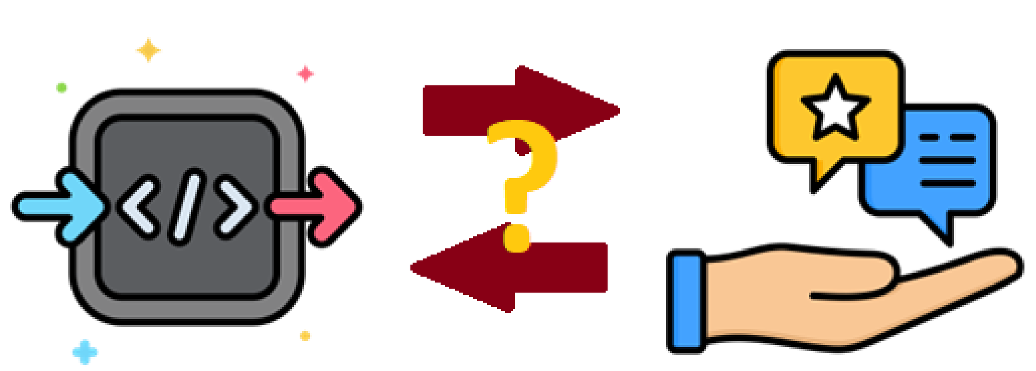

PDLA Algorithms for Online Convex Covering

We study online covering with convex objectives:

$$\min f(x) \text{ over } x \in \mathbb{R}^n_{\geq 0} \text{ subject to } Ax \geq \mathbf{1} \tag{3}$$

where $f : \mathbb{R}^n_{\geq 0} \mapsto \mathbb{R}_{\geq 0}$ is monotone, convex, and differentiable. We also assume:

- $f(0) = 0$, and ∇f is monotone;
- d is the row sparsity of A;
- There exists some $p := \sup_x \frac{\langle x, \nabla f(x) \rangle}{f(x)}$.

Here, the switching strategy **do not work**!

- The objective $f(x)$ is not necessarily sub-additive.
- Additional cost of switching between solutions is unbounded.

PDLA Algorithm for Online Convex Covering

Input: problem instance A, c, advice x', confidence parameter λ;
Output: solution x.
When constraint i arrives:
  while $A_i x \leq 1$ **do**
    for $j \in [n]$ **do**
      Increment x_j at rate

$$\frac{a_{ij}}{\nabla_j f(x)}\left(x_j + \frac{\lambda}{A_{ij} d} + \frac{(1-\lambda)x' \mathbf{1}_{x_j < x'_j}}{A_i x'_c}\right)$$

    Increment y_i at rate (for some δ chosen later)

$$r := \frac{\delta}{\log(1 + \frac{2}{\lambda} d^2)}$$

    for each $j \in [n]$ s.t. dual constraint j is tight **do**
      $m_j^* \leftarrow \arg\max_{t=1}^{i}\{A_{tj} | y_t > 0\}$.
      Decrease $y_{m_j^*}$ at rate $\frac{A_{ij}}{A_{(m_j^*)j}} \cdot r$.

Theorem 3. *Our PDLA algorithm is $O(\frac{1}{1-\lambda})$-consistent, and $O((p \log \frac{d}{\lambda})^p)$-robust, matching prior work up to the confidence parameter δ, and improves upon Theorem 1 when f is linear.*

Contact Information

- *Email:* `MaoyuanRS@gmail.com`
- *Webpage:* `maoyuans.github.io`
- *Twitter:* `@maoyuans`

PURDUE UNIVERSITY