Fractional Set Cover LP - An Application of MWU Method

Anil Maheshwari

anil@scs.carleton.ca
School of Computer Science
Carleton University
Canada

Outline

LP for Set Cover

ρ-bounded oracle

MWU Method

LP for Set Cover

Set Cover Problem

Input: A universe U consisting of n elements.
A set of m subsets $\mathcal{S} = \{S_1, S_2, \ldots, S_m\}$ of U such that $\cup_{i=1}^{m} S_i = U$.

Output: Find a minimum number of subsets of $\mathcal{S}$ such that their union covers all elements of U.

Integer Linear Program for Set Cover

Define a $0-1$ indicator variable x_S for each set $S \in \mathcal{S}$, where $x_S = 1$ if and only if S is included in the set cover.

The ILP formulation is:

$$\begin{aligned} \min \sum_{S \in \mathcal{S}} x_S \\ \forall u \in U : \sum_{u \in S} x_S \geq 1 \\ \forall S \in \mathcal{S} : x_S \in \{0,1\} \end{aligned}$$

The relaxed Linear Program (LP) is where we replace the integrality constraint $x_S \in \{0,1\}$ by $0 \leq x_S \leq 1$.

Moreover, we can replace $0 \leq x_S \leq 1$ by $x_S \geq 0$ as this is a minimization problem.

Optimization Problem → Feasibility Problem

- The value of the objective function in ILP formulation is one of the numbers $\{1, \ldots, m\}$.

- Suppose we guess that the size of the set cover is $\beta \in \{1, \ldots, m\}$.

- If the following feasibility inequality can be satisfied, we know that the optimal value is at most β.

$$\begin{aligned} \forall u \in U : \sum_{u \in S} x_S &\geq 1 \\ \sum_{S \in \mathcal{S}} x_S &\leq \beta \\ \forall S \in \mathcal{S} : x_S &\geq 0 \end{aligned}$$

- We can perform a binary search to find the true optimal value.

Fractional Set Cover Feasibility Problem

Let $\mathcal{P} = \{x \in \Re^m | \sum_{i=1}^{m} x_i \leq \beta \bigwedge \forall i \in \{1, \ldots, m\}, x_i \geq 0\}$.

$\mathcal{P}$ is a convex polytope that defines the feasible solutions for our Linear Program.

We can express the feasibility problem succinctly as follows:

Fractional Set Cover Feasibility Problem:

Report $x \in \mathcal{P}$ such that $\forall u \in U : \sum_{u \in S} x_S \geq 1$,

Otherwise report infeasibility.

An Approximate Abstract Feasibility Problem

Let A be a $n \times m$ matrix, b is a vector of length m,
$\mathcal{P}$ is the convex feasibility region.

Approximate Abstract Feasibility Problem:
Let $\epsilon \geq 0$ be an error parameter.
If $z \in \mathcal{P}$ and $Az \geq b$ is feasible then
report $x \in \mathcal{P}$ such that $A_i x \geq b_i - \epsilon$, $\forall i \in \{1, \ldots, n\}$,
Otherwise report infeasibility.

Approximate Set Cover Feasibility Problem

The matrix A is a $0-1$ matrix of size $n \times m$. It represents elements of $\mathcal{U} = \{u_1, \ldots, u_n\}$ as rows and subsets in $\mathcal{S} = \{S_1, \ldots, S_m\}$ as columns.

$$A[i,j] = \begin{cases} 1, & \text{if } u_i \in S_j \\ 0, & \text{otherwise} \end{cases}$$

Approximate Set Cover Feasibility Problem:
For a universe U of size n and m-subsets of U, we have

- characteristic matrix A of size $n \times m$,
- vector b of length n consisting of 1's,
- $\mathcal{P} = \{x \in \Re^m | \sum_{i=1}^{m} x_i \leq \beta \wedge \forall i \in \{1, \ldots, m\}, x_i \geq 0\}$.
- Error parameter $\epsilon \geq 0$.

If $z \in \mathcal{P}$ and $Az \geq b$ is feasible then
report $x \in \mathcal{P}$ such that $A_i x \geq 1 - \epsilon$, $\forall i \in \{1, \ldots, n\}$,
Otherwise report infeasibility.

Note: A_i represents the i-th row of matrix A.

Approximate Feasibility $\rightarrow$ Feasibility

- Suppose we have an instance of the set cover feasibility problem with a given choice of β that is feasible.

- Approximate set cover feasibility problem returns us an $x \in \mathcal{P}$ such that $\forall u \in U : \sum_{u \in S} x_S \geq 1 - \epsilon$.

- Set $\bar{x} = \frac{x}{1-\epsilon}$.

- $\bar{x}$ satisfies the constraints $\bar{x}_i \geq 0$ for $i = 1, \ldots, m$ and $\forall u \in U : \sum_{u \in S} \bar{x}_S \geq 1$.

- $\bar{x}$ results in the objective value of at most $(1 + O(\epsilon))\beta$.

Next: We use the MWU method to solve the Approximate Set Cover Feasibility Problem. But first ρ-bounded oracle for set cover.

ρ-bounded oracle

ρ-bounded oracle

The ρ-bounded oracle takes as input a probability distribution $p = (p_1, \ldots, p_n)$, where $\sum_{i=1}^{m} p_i = 1$, on the rows of A (i.e. on the elements of U) and returns the following.

ρ-bounded oracle

If $x \in \mathcal{P}$ and $p^T Ax \geq p^T b$ is feasible,
return $x^* \in \Re^m$ such that $\forall i : |A_i x^* - b_i| \leq \rho$.
Otherwise, return that the system is infeasible.

Note: $p^T Ax \geq p^T b$ is a single inequality.
Its a linear combination of rows of A given by the vector p.
Finding $x \in \mathcal{P}$ that satisfies $p^T Ax \geq p^T b$ may be easier than satisfying n constraints of the fractional set cover feasibility problem.

Next: Construction of a ρ-bounded oracle for the set cover.

ρ-bounded oracle for set cover

For set cover $p^T b = 1$, as b is a vector of all 1s, and p is a vector of probabilities that add to 1.

What do we want?

We want $x \geq 0$, $\sum_{S \in \mathcal{S}} x_S \leq \beta$, and

$p^T Ax = \sum_{u \in U} p_u \left(\sum_{u \in S} x_S \right) = \sum_{S \in \mathcal{S}} x_S p(S) \geq 1$, where $p(S)$ denotes the sum of the probabilities associated to the elements in S.

How to find x?
- Find the set $S \in \mathcal{S}$ that maximizes $p(S)$ for the given vector p.
- Suppose the set $S^* \in \mathcal{S}$ maximizes this value.
- Set $x_{S^*} = \beta$ and for every other set $S \neq S^*$ set $x_S = 0$.
- Observe that the vector x^* has 0's in all the coordinates except the coordinate corresponding to S^* where it is equal to β.

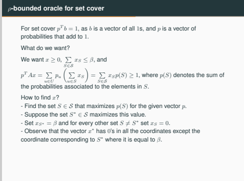

- What are the two equalities in $p^T Ax = \sum_{u \in U} p_u \left(\sum_{u \in S} x_S \right) = \sum_{S \in \mathcal{S}} x_S p(S)$?
- We will interpret the product $p^T Ax$ in two different ways
- Think of each element $u_i \in \mathcal{U}$ has an associated probability p_i.

1st interpretation: Product Ax is a vector of dimension n, where its i-th entry is the number of sets in $\mathcal{S}$ that contain the element u_i.
- $p^T Ax$ is the dot-product of vectors p^T and Ax, where the i-th entry in Ax is multiplied by the probability p_i.
- Thus, $p^T Ax$ is the sum of the products of the probability p_i of element u_i times the number of occurrences of u_i in $\mathcal{S}$.

2nd Interpretation: For each set $S \in \mathcal{S}$ sum up the probabilities associated to each element in that set (this is the quantity $p(S)$).
- We take the sum $p(S)$ over all sets so that for each element we take into account the number of times it occurs in the sets of $\mathcal{S}$.

-Therefore, $p^T Ax = \sum_{u \in U} p_u \left(\sum_{u \in S} x_S \right) = \sum_{S \in \mathcal{S}} x_S p(S)$.

ρ-bounded oracle for set cover (contd.)

x^* is feasible

The vector $x^* = (0, 0, \ldots, \beta, 0, \ldots, 0) \in \mathcal{P}$, as each of its coordinates is ≥ 0 and the sum of the coordinates is $\leq \beta$.

Consider $\sum\limits_{S \in \mathcal{S}} x^*_S p(S)$.

If $\sum\limits_{S \in \mathcal{S}} x^*_S p(S) \geq 1$, we have the x^* that we are looking for.

Claim

If $\sum_{S \in \mathcal{S}} x^*_S p(S) < 1$, then no other $x \in \mathcal{P}$ can satisfy the inequality $\sum\limits_{S \in \mathcal{S}} x_S p(S) \geq 1$.

Proof.

Note that under the constraints ($x \geq 0$, $\sum\limits_{S \in \mathcal{S}} x_S \leq \beta$) the choice of x (=x^*) that maximizes the expression $\sum\limits_{S \in \mathcal{S}} x_S p(S)$ didn't satisfy the inequality. Any other assignment will have a value at most the max value. $\square$

ρ-bounded oracle for set cover (contd.)

What should be the value of ρ?

Find the smallest value ρ such that for all $i \in \{1, \ldots, n\}$, $|A_i x^* - b_i| \leq \rho$.
- Due to the choice of x^* and the matrix A being a $0-1$ matrix, the product $A_i x^*$ is either 0 or β.

- Thus $|A_i x^* - b_i| = |A_i x^* - 1| \leq |\beta - 1| \leq \beta \leq m$.

- Note that $\beta \geq 1$ as the set cover consists of one or more sets to cover $\mathcal{U}$.

- Set $\rho = \min\{\beta, m\}$.

Now we have the required ρ-bounded oracle for the set cover.

ρ-bounded oracle
If $x \in \mathcal{P}$ and $p^T Ax \geq p^T b$ is feasible,
we return $x^* = (0, 0, \ldots, \beta, 0, \ldots, 0) \in \Re^m$ such that $\forall i : |A_i x^* - b_i| \leq \rho$.
Otherwise, we return that the system is infeasible.

MWU Method

MWU Method for Approximate Fractional Set Cover

Step 1: Fix an $\eta \in [0, \min(\frac{1}{2}, \frac{\epsilon}{2\rho})]$ and set $w^1 = (1, \ldots, 1)$.

Step 2: For $t = 1$ to $T = 4\frac{\rho^2 \ln n}{\epsilon^2}$ days do:

1. Compute the probability vector $p^t = (\frac{w_1^t}{\Phi^t}, \ldots, \frac{w_n^t}{\Phi^t})$, where $\Phi^t = \sum_{i=1}^{n} w_i^t$.
2. Execute the ρ-bounded oracle for p^t. It either returns that the system is infeasible and we STOP or returns the vector x^t.
3. Compute the costs of each expert i by evaluating $m_i^t = \frac{1}{\rho}(A_i x^t - b_i)$. (Observe that $m_i^t \in [-1, 1]$.)
4. Update weights for the next day for each expert i by executing $w_i^{t+1} = w_i^t(1 - \eta m_i^t)$.

Step 3: If we didn't report infeasibility during the T days of execution, return $\bar{x} = \frac{1}{T}\sum_{t=1}^{T} x^t$ as the answer to the Approximate Set Cover Feasibility Problem.

Observations

1. Suppose we didn't report infeasibility on any of the T days of execution of MWU Method. We claim that $\bar{x} = \frac{1}{T} \sum_{t=1}^{T} x^t$ is a feasible solution for the approximate fractional set cover problem as it is a convex combination (average) of T feasible vectors. Hence it is within the polytope $\mathcal{P}$.
2. If $A_i x^t \geq b_i$, $m_i^t \geq 0$, and the i-th constraint is satisfied. But if $A_i x^t < b_i$, then $m_i^t < 0 \implies$ for the rows of A for which the constraints are satisfied their weights will be smaller compared to the rows for which the constraints are not satisfied.
3. In the next round the unsatisfied rows (experts) will get higher probabilities compared to the satisfied rows. The more unsatisfied the row is higher is its probability. Hence it is likely that it may get satisfied in future rounds.

Analysis

From the analysis of MWU Method we know that the expected cost of this algorithm is bounded with respect to the cost of any expert i by

$$\sum_{t=1}^{T} M^t = \sum_{t=1}^{T} p^t \cdot m^t \leq \frac{\ln n}{\eta} + \eta T + \sum_{t=1}^{T} m_i^t$$

Claim

$\sum_{t=1}^{T} M^t \geq 0.$

Proof.

Since $m^t = \frac{1}{\rho}(Ax^t - b)$, we have

$$M^t = p^t \cdot m^t = \frac{1}{\rho}\left(p^t \cdot (Ax^t - b)\right) = \frac{1}{\rho}(p^t \cdot Ax^t - p^t \cdot b) \geq 0$$

The last inequality holds as the system is satisfied, i.e. $p^t Ax^t \geq p^t b$. $\square$

Analysis (contd.)

Claim

For each $i = 1, \ldots, n$, $A_i \bar{x} \geq b_i - \epsilon$, where $\bar{x} = \frac{1}{T} \sum_{t=1}^{T} x^t$.

Proof.

- For $t = 1, \ldots, T$, each $M^t \geq 0 \implies \frac{\ln n}{\eta} + \eta T + \sum_{t=1}^{T} m_i^t \geq \sum_{t=1}^{T} M^t \geq 0$.

- Substitute $m_i^t = \frac{1}{\rho}(A_i x^t - b_i)$, we obtain: $\frac{\ln n}{\eta} + \eta T + \frac{1}{\rho} \sum_{t=1}^{T} (A_i x^t - b_i) \geq 0$.

- This is equivalent to $\frac{\ln n}{\eta} + \eta T + \frac{1}{\rho} \sum_{t=1}^{T} A_i x^t - \frac{T}{\rho} b_i \geq 0$.

- Multiply by $\frac{\rho}{T}$ and use $\bar{x} = \frac{1}{T} \sum_{t=1}^{T} x^t$:

$$\frac{\rho \ln n}{T\eta} + \rho\eta + A_i \bar{x} - b_i \geq 0$$

- Substitute $T = 4\frac{\rho^2 \ln n}{\epsilon^2}$ and $\eta \in [0, \min(\frac{1}{2}, \frac{\epsilon}{2\rho})]$ we obtain

$$\epsilon + A_i \bar{x} - b_i \geq 0$$

□

Computational Complexity

- We run the algorithm for T days.

- For each day we make a call to the ρ-bounded oracle.

- The overall time complexity is bounded by the time it takes to run $O(\frac{\rho^2 \ln n}{\epsilon^2})$ calls to the oracle.

Summary

1. Set cover problem $\rightarrow$ ILP formulation $\rightarrow$ Relaxed LP that may use fractional values.
2. LP formulation $\rightarrow$ Fractional set cover feasibility problem (binary search to find the minimum set cover).
3. Feasibility Problem $\rightarrow$ Approximate feasibility problem.
4. Answer to approximate feasibility problem yields an assignment to variables (x_S's) that satisfies all the constraints and is within a $1 + O(\epsilon)$ factor of optimal.
5. To answer the approximate feasibility problem, we take help of a ρ-bounded oracle.
6. An easier problem as it has one inequality.
7. Use ρ-bounded oracle in the MWU method to find successive x's that are within the convex region $\mathcal{P}$.
8. If the approximate feasibility problem has a solution, the average of x^t's yields a good approximation, using the MWU guarantees. Otherwise, we report infeasibility and adjust the guess on the optimal value and restart.

Reference

- Arora, Hazan and Kale, The multiplicative weights update method: a meta-algorithm and applications, Theory of Computing 8(1): 121-164, 2012.

- Several Lecture Notes