MATH 233B: Polyhedral techniques in combinatorial optimization
Instructor: Jan Vondrák
Lecture date: 3/9/2017

1 Submodular maximization subject to a matroid constraint

Given: Monotone submodular function $f : 2^E \to \mathbb{R}_+$, matroid $\mathcal{M} = (E, \mathcal{I})$.
Goal: Maximize $f(S)$ subject to $S \in \mathcal{I}$.

We have seen that the Submodular Welfare problem is a special case of this, and the greedy algorithm provides a 1/2-approximation. Our goal here is to present a $(1 - 1/e)$-approximation algorithm, which is best possible.

1.1 The continuous greedy algorithm

We consider the following relaxation of the problem:

$$\max\{F(\mathbf{x}) : \mathbf{x} \in P(\mathcal{M})\}$$

where $F(\mathbf{x}) = \mathbb{E}[f(\hat{\mathbf{x}})]$ is the multilinear relaxation of f, and $P(\mathcal{M})$ is the matroid polytope corresponding to $\mathcal{M}$. Our first task is to solve this problem (at least approximately). In this part of the algorithm, we do not need to assume much about P, except that it is a *solvable polytope.*

Definition 1 *A class of polytopes is called* solvable *if there is an algorithm which for any polytope $P \subset \mathbb{R}^n$ in the class and any $\mathbf{w} \in \mathbb{R}^n$ solves the problem* $\max\{\mathbf{w} \cdot \mathbf{x} : \mathbf{x} \in P\}$ *in time polynomial in n and the bit-size of* $\mathbf{w}$.

In other words, we assume that we know how to optimize linear functions over P. (This is certainly true for our matroid polytope $P(\mathcal{M})$.) However, F is not linear, or even concave, so standard optimization techniques do not apply here. We describe a *continuous greedy algorithm* which is designed specifically for the problem $\max\{F(\mathbf{x}) : \mathbf{x} \in P\}$. The continuous greedy algorithm attacks this problem in a way similar to the discrete greedy algorithm: in each step, it tries to make a small improvement of maximum possible benefit. However, in contrast, it achieves a $(1 - 1/e)$-approximation for *any solvable polytope* P, i.e. in a much more general setting than a cardinality constraint. In a sense this is analogous to the fact that linear programming can be solved optimally for any (solvable) polytope, while discrete optimization problems can be solved optimally only for very special combinatorial structures.

In a compact form, the algorithm is described as follows.

Algorithm **ContinuousGreedy**(F, P):
  Initialize $\mathbf{x}(0) := \mathbf{0} \in \mathbb{R}^E$;
  Define $\mathbf{v}_{max}(\mathbf{x}) = \text{argmax}_{\mathbf{v} \in P}(\mathbf{v} \cdot \nabla F(\mathbf{x}))$;
  For $t \in [0, 1]$ solve
    $\frac{d\mathbf{x}}{dt} = \mathbf{v}_{max}(\mathbf{x}(t))$;
  Return $\mathbf{x}(1)$.

In words, the continuous greedy algorithm at time t moves in a direction $\mathbf{v}_{max}(\mathbf{x}(t))$ which maximizes the local gain $\mathbf{v} \cdot \nabla F(\mathbf{x})$ over all directions $\mathbf{v} \in P$. Importantly, we are able to find $\mathbf{v}_{max}(\mathbf{x}) = \text{argmax}_{\mathbf{v} \in P}(\mathbf{v} \cdot \nabla F(\mathbf{x}(t))$ efficiently in each step, since this is just a linear optimization problem over P. The actual algorithm will proceed in finite steps, increasing t by some fixed small increment $\delta > 0$. However, we view the algorithm as truly continuous, which makes the analysis easier.

We remark that it is instructive to assume that P is a *down-monotone polytope*, i.e. $0 \leq \mathbf{x} \leq \mathbf{y} \in P \Rightarrow \mathbf{x} \in P$. This is without loss of generality, as F is monotone and hence P can be replaced by its down-monotone closure $P^{\downarrow} = \{\mathbf{y} \geq 0 : \exists \mathbf{x} \in P, \mathbf{y} \leq \mathbf{x}\}$ without affecting the optimum. Then, the initial point $\mathbf{0}$ is in $P^{\downarrow}$ and the trajectory $\mathbf{x}(t)$ is contained in $P^{\downarrow}$ as well. Nevertheless, the assumption of down-monotonicity is not used anywhere in the analysis. First, let us prove the following lemma.

Lemma 2 *For any monotone submodular function $f : 2^E \to \mathbb{R}_+$, its multilinear extension $F : [0,1]^E \to \mathbb{R}_+$, and a polytope $P \subseteq [0,1]^E$, let $\mathbf{x} \in [0,1]^E$ be an arbitrary point and $OPT = \max\{F(\mathbf{y}) : \mathbf{y} \in P\}$. Then there is $\mathbf{v} \in P$ such that*

$$\mathbf{v} \cdot \nabla F(\mathbf{x}) \geq OPT - F(\mathbf{x}).$$

Proof: We claim that any $\mathbf{v} \in P$ such that $OPT = F(\mathbf{v})$ satisfies the conclusion of the lemma. Consider a line from $\mathbf{x}$ towards $\mathbf{v} \vee \mathbf{x}$; this line has direction $\mathbf{d} = (\mathbf{v} \vee \mathbf{x}) - \mathbf{x} = (\mathbf{v} - \mathbf{x}) \vee 0$, which satisfies $0 \leq \mathbf{d} \leq \mathbf{v}$. Let $\phi(\xi) = F(\mathbf{x} + \xi \mathbf{d})$ denote the objective function along this line. The starting point has value $\phi(0) = F(\mathbf{x})$. Now consider the point corresponding to $\xi = 1$; its value is $\phi(1) = F(\mathbf{x} + \mathbf{d}) \geq OPT$, because $\mathbf{x} + \mathbf{d} = \mathbf{v} \vee \mathbf{x} \geq \mathbf{v}$ and F is non-decreasing in each coordinate. As we saw in the last lecture, $\phi(\xi)$ is concave. This means that

$$\phi(1) - \phi(0) \leq \phi'(0) = \mathbf{d} \cdot \nabla F(\mathbf{x}).$$

Using $\mathbf{d} \leq \mathbf{v}$ and the nonnegativity of ∇F (due to monotonicity), we obtain

$$\mathbf{v} \cdot \nabla F(\mathbf{x}) \geq \mathbf{d} \cdot \nabla F(\mathbf{x}) \geq \phi(1) - \phi(0) \geq OPT - F(\mathbf{x}).$$

$\square$

Recall that $\mathbf{v} \cdot \nabla F(\mathbf{x})$ is the derivative of F along a line of direction $\mathbf{v}$ at the point $\mathbf{x}$. The continuous greedy algorithm chooses a direction maximizing this quantity over $\mathbf{v} \in P$. I.e., what we proved is that at each time t, the local gain in terms of F is at least $OPT - F(\mathbf{x}(t))$. This leads to a differential equation whose solution yields the factor of $1 - 1/e$.

Theorem 3 *For any monotone submodular function $f : 2^E \to \mathbb{R}_+$, its multilinear extension $F : [0,1]^E \to \mathbb{R}_+$, and any solvable polytope $P \subseteq [0,1]^E$, the continuous greedy algorithm finds a point $\mathbf{x}(1) \in P$ such that*

$$F(\mathbf{x}(1)) \geq (1 - 1/e)OPT$$

where $OPT = \max\{F(\mathbf{x}) : \mathbf{x} \in P\}$.

Proof: By design of the continuous greedy algorithm, we obtain

$$\mathbf{x}(1) = \int_0^1 \mathbf{v}_{max}(\mathbf{x}(t))dt.$$

Therefore, $\mathbf{x}(1)$ is a convex linear combination of vectors $\mathbf{v}_{max}(\mathbf{x}(t)) \in P$, and hence $\mathbf{x}(1) \in P$.

To prove the approximation guarantee, let us analyze $F(\mathbf{x}(t))$. By the chain rule and the design of the continuous greedy algorithm,

$$\frac{d}{dt}F(\mathbf{x}(t)) = \frac{d\mathbf{x}}{dt} \cdot \nabla F(\mathbf{x}(t)) = \mathbf{v}_{max}(\mathbf{x}(t)) \cdot \nabla F(\mathbf{x}(t)) = \max_{\mathbf{v} \in P} \mathbf{v} \cdot \nabla F(\mathbf{x}(t)).$$

By Lemma 2, we know that this quantity is at least $OPT - F(\mathbf{x}(t))$. Hence, we obtain

$$\frac{d}{dt}F(\mathbf{x}(t)) + F(\mathbf{x}(t)) \geq OPT.$$

This differential inequality can be solved as follows. Multiplying by e^t, we obtain:

$$\frac{d}{dt}(e^t F(\mathbf{x}(t))) = e^t \frac{d}{dt}F(\mathbf{x}(t)) + e^t F(\mathbf{x}(t)) \geq e^t OPT.$$

The initial condition is $F(\mathbf{x}(0)) = 0$ which implies

$$e^t F(\mathbf{x}(t)) \geq \int_0^t e^\tau \, OPT \, d\tau = (e^t - 1)OPT.$$

Therefore, the solution at time t satisfies $F(\mathbf{x}(t)) \geq (1 - e^{-t})OPT$ and at time $t = 1$ we get $F(\mathbf{x}(1)) \geq (1 - 1/e)OPT$. $\square$

As we noted, an actual implementation needs to emulate the continuous greedy algorithm in a finite number of steps, in some sense numerically solving the differential equation $\frac{d\mathbf{x}}{dt} = \mathbf{v}_{max}(\mathbf{x}(t))$. This can be done by standard techniques, for example by choosing a fixed time increment $\delta > 0$ and approximating the process by $\mathbf{x}(t+\delta) = \mathbf{x}(t) + \delta \mathbf{v}_{max}(\mathbf{x}(t))$. It can be shown that this deviates from the analysis above by an error which tends to zero as $\delta \to 0$.

1.2 Rounding in the matroid polytope

The problem that remains is how to round the fractional solution, $\mathbf{x}_0 = \mathbf{x}(1)$ found by the continuous greedy algorithm. How can we convert $\mathbf{x}_0$ into an integral solution of comparable value? Observe that if f were a linear function (and hence F as well), this would be a trivial problem: Every point $\mathbf{x}_0 \in P(\mathcal{M})$ can be decomposed as a convex combination $\mathbf{x}_0 = \sum_{I \in \mathcal{I}} \alpha_I \chi_I$ where $\alpha_I \geq 0$ and $\sum \alpha_I = 1$. If F is a linear function, then $F(\mathbf{x}_0) = \sum_{I \in \mathcal{I}} \alpha_I f(I)$ and hence at least one of the sets I satisfies $f(I) \geq F(\mathbf{x}_0)$. Nevertheless, this is not true for submodular functions.

We will see that there is a rounding procedure which does not lose in terms of the objective value even if the objective function is submodular. In other words, the multilinear optimization problem $\max\{F(\mathbf{x}) : \mathbf{x} \in P(\mathcal{M})\}$ has integrality gap 1 for any submodular function and any matroid polytope. More precisely, we show the following.

Theorem 4 *There is an (efficient) rounding procedure such that given any submodular function $f : 2^E \to \mathbb{R}$, matroid $\mathcal{M} = (N, \mathcal{I})$ and a point $\mathbf{x}_0 \in P(\mathcal{M})$, it outputs an independent set $S \in \mathcal{I}$ such that $f(S) \geq F(\mathbf{x}_0)$.*

We remark that monotonicity does not play any role in this theorem. If we prove this theorem, we obtain together with the continuous greedy algorithm a $(1-1/e)$-approximation for maximizing any monotone submodular function subject to a matroid constraint.

Corollary 5 *For any monotone submodular function* $f : 2^E \to \mathbb{R}_+$ *and a matroid* $\mathcal{M} = (N, \mathcal{I})$, *there is a* $(1-1/e)$*-approximation for the problem* $\max\{f(S) : S \in \mathcal{I}\}$.

To prove Theorem 4, we proceed as follows. Let us assume that the starting point is $\mathbf{x}_0 = \sum_{i=1}^{\ell} \alpha_i \chi_{B_i}$ where B_i is a base of the matroid and $\sum_{i=1}^{\ell} \alpha_i = 1, \alpha_i \geq 0$. Since we can assume that $\mathbf{x}_0$ is in the matroid base polytope, this is certainly true. In fact, our continuous greedy algorithm (or rather, a suitable discretization thereof), provides $\mathbf{x}_0$ conveniently in this form. Moreover, the algorithm gives $\mathbf{x}_0 = \frac{1}{N}\sum_{i=1}^{N} \chi_{B_i}$ so let's assume that this is the case and for simplicity $N = 2^q$.

We will proceed in q phases, where in each phase, pairs of bases are merged. The merge procedure relies on the strong exchange property and works as follows.

Procedure **MergeBases**(B_1, B_2):
  While $B_1 \neq B_2$ do
    Find $i \in B_1 \setminus B_2, j \in B_2 \setminus B_1$ such that $B_1 - i + j$ and $B_2 - j + i$ are bases.
    Randomly with equal probability,
      either replace $B_1 \leftarrow B_1 - i + j$,
      or replace $B_2 \leftarrow B_2 - j + i$.
  EndWhile
  Output B_1.

We prove the following lemma.

Lemma 6 *If* B_1, B_2 *are two bases participating in a convex combination* $\mathbf{x} = \frac{1}{N}\sum_{i=1}^{N} \chi_{B_i}$, *and we replace* $\frac{1}{N}(\chi_{B_1} + \chi_{B_2})$ *by* $\frac{2}{N}\chi_B$, *where* B *is the output of* **MergeBases**(B_1, B_2), *then*

$$\mathbb{E}\left[F\left(\mathbf{x} - \frac{1}{N}(\chi_{B_1} + \chi_{B_2}) + \frac{2}{N}\chi_B\right)\right] \geq F(\mathbf{x}).$$

Proof: We use the fact that F is convex along any line of direction $\mathbf{e}_i - \mathbf{e}_j$. In each step of **MergeBases**, conditioned on the current point being $\mathbf{x}$, we pick two elements i, j and we either add or subtract $\frac{1}{N}(\mathbf{e}_i - \mathbf{e}_j)$. Therefore, in expectation the value after this step is

$$\frac{1}{2}(F(\mathbf{x} + \frac{1}{N}(\mathbf{e}_i - \mathbf{e}_j)) + F(\mathbf{x} - \frac{1}{N}(\mathbf{e}_i - \mathbf{e}_j))) \geq F(\mathbf{x})$$

by the convexity of F. Hence $F(\mathbf{x})$ can only increase in expectation. □

By applying this procedure to $2^{\ell-1}$ disjoint pairs of bases, we obtain a new point $\mathbf{x}' = \frac{2}{N}\sum_{i=1}^{N/2} \chi_{B'_i}$ such that $\mathbb{E}[F(\mathbf{x}')] \geq F(\mathbf{x}_0)$. The number of bases is reduced from $N = 2^q$ to 2^{q-1}. We repeat this procedure until we obtain a single (random) base, B^*, such that $\mathbb{E}[f(B^*)] = \mathbb{E}[F(\chi_{B^*})] \geq F(\mathbf{x}_0)$. This is the output of the rounding procedure, which proves Theorem 4.