Matroids and Submodular Optimization	September 9, 2026

Lecture 5

Lecturer: Siyue Liu — *Scribe: ChatGPT*

1 Submodular Maximization

Let E be a finite ground set. For a set function $f : 2^E \to \mathbb{R}$, an element $e \in E$, and a set $S \subseteq E$, the *marginal value* of e with respect to S is

$$f(e \mid S) := f(S \cup \{e\}) - f(S).$$

The function f is *submodular* if it has diminishing returns:

$$f(e \mid S) \geq f(e \mid T) \qquad \text{whenever } S \subseteq T \subseteq E \text{ and } e \notin T.$$

It is *monotone* if $f(S) \leq f(T)$ whenever $S \subseteq T$, and it is *normalized* if $f(\varnothing) = 0$.

We study two basic maximization settings. First, f is monotone and the chosen set must satisfy a cardinality or matroid constraint. We then drop monotonicity and consider unconstrained maximization of a nonnegative submodular function.

1.1 A cardinality constraint

Assume throughout this subsection that $f : 2^E \to \mathbb{R}_{\geq 0}$ is normalized, monotone, and submodular. Given an integer $1 \leq k \leq |E|$, consider

$$\max\{f(S) : S \subseteq E,\ |S| \leq k\}.$$

Example 1.1 (Maximum k-coverage)**.** Let U be a universe, and associate a subset $C_e \subseteq U$ with each $e \in E$. Selecting $S \subseteq E$ covers $\bigcup_{e \in S} C_e$, so

$$f(S) = \left| \bigcup_{e \in S} C_e \right|.$$

This is a normalized, monotone, submodular function. The maximum k-coverage problem asks for at most k sets whose union is as large as possible.

Greedy algorithm. Set $A_0 = \varnothing$. For $i = 1, \ldots, k$, choose

$$a_i \in \operatorname*{arg\,max}_{e \in E \setminus A_{i-1}} f(e \mid A_{i-1}), \qquad A_i = A_{i-1} \cup \{a_i\}.$$

Return A_k.

Theorem 1.2 (Cardinality-constrained greedy)**.** *Let* OPT *be an optimal solution. The greedy solution satisfies*

$$f(A_k) \geq \left(1 - \left(1 - \frac{1}{k}\right)^k\right) f(\mathrm{OPT}) \geq \left(1 - \frac{1}{e}\right) f(\mathrm{OPT}).$$

Thus greedy is a $(1 - 1/e)$-approximation [1].

Proof. At iteration i, the greedy choice, submodularity, and monotonicity give

$$\begin{aligned} f(a_i \mid A_{i-1}) &\geq \frac{1}{k} \sum_{o \in \mathrm{OPT}} f(o \mid A_{i-1}) \\ &\geq \frac{1}{k} (f(A_{i-1} \cup \mathrm{OPT}) - f(A_{i-1})) \\ &\geq \frac{1}{k} (f(\mathrm{OPT}) - f(A_{i-1})). \end{aligned}$$

Define the residual gap $\Delta_i = f(\mathrm{OPT}) - f(A_i)$. The preceding inequality is equivalent to

$$\Delta_i \leq \left(1 - \frac{1}{k}\right) \Delta_{i-1}.$$

Since $f(A_0) = f(\varnothing) = 0$, iterating the recurrence yields

$$f(\mathrm{OPT}) - f(A_k) = \Delta_k \leq \left(1 - \frac{1}{k}\right)^k f(\mathrm{OPT}) \leq e^{-1} f(\mathrm{OPT}).$$

Rearranging proves the result. □

Remark (Continuous intuition). Put $y_i = f(A_i)$ and regard one iteration as a time step of length $1/k$. The recurrence resembles the differential inequality

$$\frac{dy}{dt} \geq f(\mathrm{OPT}) - y, \qquad y(0) = 0.$$

The corresponding equality has solution $y(t) = (1 - e^{-t}) f(\mathrm{OPT})$, which gives $y(1) = (1 - 1/e) f(\mathrm{OPT})$. The discrete recurrence above is the rigorous version of this calculation.

1.2 A matroid constraint

Let $M = (E, \mathcal{I})$ be a matroid of rank r, and let $f : 2^E \to \mathbb{R}_{\geq 0}$ again be normalized, monotone, and submodular. We now consider

$$\max\{f(S) : S \in \mathcal{I}\}.$$

Greedy algorithm. Starting with $A_0 = \varnothing$, for $i = 1, \ldots, r$ choose

$$a_i \in \arg\max\{f(e \mid A_{i-1}) : e \in E \setminus A_{i-1},\ A_{i-1} \cup \{e\} \in \mathcal{I}\}, \qquad A_i = A_{i-1} \cup \{a_i\}.$$

Because the objective is monotone, greedy may be continued until A_r is a basis. An optimal independent set may likewise be extended to an optimal basis.

Theorem 1.3 (Matroid-constrained greedy)**.** *The greedy solution A_r satisfies*

$$f(A_r) \geq \frac{1}{2} f(\mathrm{OPT}).$$

Thus greedy is a 1/2-approximation for monotone submodular maximization under a matroid constraint [1].

Proof. Write $A_i = \{a_1, \dots, a_i\}$. By repeated basis exchange, the elements of the optimal basis can be ordered as $\mathrm{OPT} = \{o_1, \dots, o_r\}$ so that, with $O_i = \{o_1, \dots, o_i\}$,

$$A_i \cup (\mathrm{OPT} \setminus O_i) \in \mathcal{B}(M) \qquad (i = 0, \dots, r).$$

In particular, $A_{i-1} \cup \{o_i\} \in \mathcal{I}$, so o_i is a feasible candidate when greedy chooses a_i. Therefore,

$$\begin{aligned} f(A_r) &= \sum_{i=1}^{r} f(a_i \mid A_{i-1}) \\ &\geq \sum_{i=1}^{r} f(o_i \mid A_{i-1}) \\ &\geq \sum_{i=1}^{r} f(o_i \mid A_r) \\ &\geq f(A_r \cup \mathrm{OPT}) - f(A_r). \end{aligned}$$

The second line uses the greedy choice; the next two use submodularity. Monotonicity gives $f(A_r \cup \mathrm{OPT}) \geq f(\mathrm{OPT})$, and hence

$$f(A_r) \geq f(\mathrm{OPT}) - f(A_r).$$

The claim follows. □

Example 1.4 (Tightness of the greedy analysis)**.** Consider the partition matroid with blocks

$$P_1 = \{1, 2\}, \qquad P_2 = \{3, 4\},$$

and capacity one in each block. Let $U = \{x, y\}$ and consider the coverage sets

$$C_1 = \{x\}, \qquad C_2 = \{y\}, \qquad C_3 = C_4 = \{x\}.$$

With tie-breaking in favor of element 1, greedy chooses 1 first and gains nothing from its later choice in P_2, so its value is 1. The independent set $\{2, 3\}$ covers both x and y and has value 2.

The tie can be removed while keeping the ratio arbitrarily close to 1/2. Add an element z of weight $\delta > 0$ to the universe, replace C_1 by $\{x, z\}$, and give x and y weight 1. Greedy now uniquely chooses 1 first and obtains $1 + \delta$, whereas $\{2, 3\}$ has value 2. Letting $\delta \downarrow 0$ proves that the factor 1/2 is tight for this greedy algorithm.

More generally, if $w : U \to \mathbb{R}_{\geq 0}$ assigns weights to the universe, the associated weighted coverage function is

$$f(T) = \sum_{u \in \bigcup_{e \in T} C_e} w(u), \qquad T \subseteq E.$$

2 Unconstrained Nonmonotone Maximization

We now assume only that $f : 2^E \to \mathbb{R}_{\geq 0}$ is submodular. In particular, f need not be monotone or normalized. The unconstrained problem is

$$\max\{f(S) : S \subseteq E\}.$$

Example 2.1 (Cut functions)**.** For an undirected graph $G = (V, F)$, the cut function

$$f(S) = |\delta(S)|, \qquad S \subseteq V,$$

is nonnegative, symmetric, and submodular. If every vertex is placed in S independently with probability $1/2$, each edge is cut with probability $1/2$. Thus a uniform random cut has expected value $|F|/2$ and is a $1/2$-approximation to the maximum cut.

For a digraph $D = (V, F)$, the directed-cut function

$$f(S) = |\delta^+(S)|$$

is nonnegative and submodular, but generally not symmetric. Each arc belongs to $\delta^+(S)$ with probability $1/4$, so a uniform random set gives a $1/4$-approximation to the maximum directed cut.

Random-half-sampling algorithm. Include each element of E independently with probability $1/2$, and return the resulting set R.

Theorem 2.2 (Random sampling)**.** *For every nonnegative submodular function f,*

$$\mathbb{E}[f(R)] \geq \frac{1}{4} f(\mathrm{OPT}).$$

If f is symmetric*, meaning that $f(S) = f(E \setminus S)$ for every $S \subseteq E$, then*

$$\mathbb{E}[f(R)] \geq \frac{1}{2} f(\mathrm{OPT}).$$

These are guarantees for the random-half-sampling algorithm [2].

We prove the theorem through the multilinear extension of f.

2.1 The multilinear extension

For $x \in [0,1]^E$, let $R(x)$ be the random subset of E containing each element e independently with probability x_e.

Definition 2.3 (Multilinear extension)**.** The multilinear extension $F : [0,1]^E \to \mathbb{R}$ of f is

$$\begin{aligned} F(x) &:= \mathbb{E}[f(R(x))] \\ &= \sum_{S \subseteq E} f(S) \prod_{e \in S} x_e \prod_{e \notin S} (1 - x_e). \end{aligned}$$

Although F need not be globally concave, submodularity gives concavity along directions whose coordinates have one sign.

Lemma 2.4 (Concavity along nonnegative directions)**.** *Let $d \in \mathbb{R}^E_{\geq 0}$. For every interval on which $x + td \in [0,1]^E$, the function*

$$g(t) = F(x + td)$$

is concave. The same conclusion holds when $d \leq 0$.

Proof. Fix distinct $i, j \in E$, and let T denote a random subset of $E \setminus \{i, j\}$ drawn according to the remaining coordinates of x. Multilinearity gives

$$\frac{\partial^2 F}{\partial x_i^2}(x) = 0,$$

while

$$\frac{\partial^2 F}{\partial x_i \partial x_j}(x) = \mathbb{E}\big[f(T \cup \{i, j\}) - f(T \cup \{i\}) \\ - f(T \cup \{j\}) + f(T)\big] \leq 0$$

by submodularity. Consequently,

$$g''(t) = d^\mathsf{T} \nabla^2 F(x + td) d = \sum_{i \neq j} d_i d_j \frac{\partial^2 F}{\partial x_i \partial x_j}(x + td) \leq 0$$

whenever the entries of d have the same sign. $\square$

For completeness, multilinearity also gives the useful first-derivative identity

$$\frac{\partial F}{\partial x_i}(x) = \mathbb{E}\big[f(R(x) \cup \{i\}) - f(R(x) \setminus \{i\})\big].$$

The distribution of coordinates other than i is the same in both terms.

2.2 Two sampling inequalities

For $A \subseteq E$ and $p \in [0, 1]$, let $A(p)$ be the random subset obtained by retaining each element of A independently with probability p.

Claim 2.5. *For every $A \subseteq E$ and $p \in [0, 1]$,*

$$\mathbb{E}[f(A(p))] = F(p\mathbf{1}_A) \geq pf(A) + (1 - p)f(\varnothing).$$

Proof. The map $p \mapsto F(p\mathbf{1}_A)$ is concave by the preceding lemma. It therefore lies above the chord joining $F(0) = f(\varnothing)$ and $F(\mathbf{1}_A) = f(A)$. $\square$

Claim 2.6. *For $A, B \subseteq E$ and $p, q \in [0, 1]$, sample $A(p)$ and $B(q)$ independently. Then*

$$\mathbb{E}[f(A(p) \cup B(q))] \geq (1 - p)(1 - q)f(\varnothing) + p(1 - q)f(A) \\ + (1 - p)qf(B) + pqf(A \cup B).$$

Proof. Condition on $A(p)$ and apply the first claim to the submodular function $T \mapsto f(A(p) \cup T)$:

$$\mathbb{E}[f(A(p) \cup B(q)) \mid A(p)] \geq qf(A(p) \cup B) + (1 - q)f(A(p)).$$

Taking expectations and applying the first claim twice more, now to the random set $A(p)$, gives

$$\mathbb{E}[f(A(p) \cup B(q))] \geq q\big(pf(A \cup B) + (1 - p)f(B)\big) \\ + (1 - q)\big(pf(A) + (1 - p)f(\varnothing)\big),$$

which is the stated inequality. $\square$

Proof of the random-sampling theorem. Take

$$A = \mathrm{OPT}, \qquad B = E \setminus \mathrm{OPT}, \qquad p = q = \frac{1}{2}.$$

Since A and B partition E, the set $A(1/2) \cup B(1/2)$ has exactly the same distribution as the output R of random-half sampling. The second claim therefore gives

$$\mathbb{E}[f(R)] \geq \frac{1}{4}\big(f(\varnothing) + f(\mathrm{OPT}) + f(E \setminus \mathrm{OPT}) + f(E)\big).$$

Because f is nonnegative, this is at least $f(\mathrm{OPT})/4$.

If f is symmetric, then $f(E \setminus \mathrm{OPT}) = f(\mathrm{OPT})$. Keeping these two terms in the preceding bound gives

$$\mathbb{E}[f(R)] \geq \frac{1}{2} f(\mathrm{OPT}),$$

as required. $\square$

Summary. Greedy loses only the factor created by the constraint: it achieves $1 - 1/e$ for a cardinality bound and $1/2$ for a general matroid. Without monotonicity, a uniform random subset already retains one quarter of the optimal value, and symmetry doubles this guarantee. The multilinear extension turns these discrete sampling statements into concavity inequalities along nonnegative directions.

References

[1] M. L. Fisher, G. L. Nemhauser, and L. A. Wolsey, "An analysis of approximations for maximizing submodular set functions—II," in *Polyhedral Combinatorics*, Mathematical Programming Studies 8, Springer, 1978, pp. 73–87.

[2] U. Feige, V. S. Mirrokni, and J. Vondrák, "Maximizing non-monotone submodular functions," *SIAM Journal on Computing* 40 (2011), 1133–1153. doi:10.1137/090779346.