The Erdos Selfridge Strategy

Given a set system $(X, \mathcal{R})$ with $n = |X|$, $m = |\mathcal{R}|$ and each set of $\mathcal{R}$ having size exactly k, two players, say named ROW PLAYER and COLUMN PLAYER, play the following game:

> ROW PLAYER starts, and picks an element of X.
>
> Then COLUMN PLAYER picks one of the remaining elements of X.
>
> Then it is again ROW PLAYER's turn to pick a remaining element and so on.
>
> ROW PLAYER wins if it is able to collect *all* the elements of *some* set of $\mathcal{R}$. COLUMN PLAYER wins otherwise (that is, COLUMN PLAYER has a hitting set for $\mathcal{R}$).

Clearly, the larger the number of sets, the easier it is for ROW PLAYER to win.

> For example, if $\mathcal{R}$ is the powerset of X and $k \leq \frac{n}{2}$, then clearly *any* strategy for ROW PLAYER wins after k steps, since ROW PLAYER would have collected k elements of X, which is also a set in $\mathcal{R}$.

In fact, there is even a set system with $m = 2^{k-1}$ sets such that the ROW PLAYER has a winning strategy:

> take X to be the vertices of a perfectly balanced binary tree of depth $k-1$ and let $\mathcal{R}$ consist of all root-to-leaf paths (each of length k). Then ROW PLAYER picks the root vertex, and then in the next turn, ROW PLAYER picks the root vertex of the left subtree if COLUMN PLAYER had picked a vertex in the right subtree, and vice versa. Continuing like this, ROW PLAYER would eventually reach a leaf, completing that set.

Theorem 0.1. *If $m < 2^{k-1}$, then there is a winning strategy for* COLUMN PLAYER.

Proof. COLUMN PLAYER will assign a weight to each set of $\mathcal{R}$, and in its turn, pick an element which is contained in the sets of maximum total weight.

The weight of all sets from which COLUMN PLAYER has already picked an element is set to 0. Otherwise, here is the weight function:

$$\omega(S) = 2^{\#\text{ of elements of } S \text{ picked by ROW PLAYER so far}}$$

Seting $\omega^1(S) = 1$ for all $S \in \mathcal{R}$ and $\Omega^1(\mathcal{R}) = \sum_{S \in \mathcal{R}} \omega^1(S)$. Then here is the algorithm for iterations $i = 1, \ldots$:

1. ROW PLAYER picks an element $x_i \in X$ from the set of elements not already picked.
2. Update the weight function: for each $S \in \mathcal{R}$ containing x_i, set

$$\omega^{i+1}(S) = 2 \cdot \omega^i(S).$$

3. COLUMN PLAYER picks an element $y_i \in X \setminus \{x_1, \dots, x_i, y_1, \dots, y_{i-1}\}$ that hits the sets of maximum weight. That is, maximizing

$$\sum_{\substack{S \in \mathcal{R}: \\ y_i \in S}} \omega^{i+1}(S).$$

4. For all $S \in \mathcal{R}$ hit by y_i, set $\omega^{i+1}(S) = 0$. We can just remove such sets from $\mathcal{R}$.

Here is the key claim.

Claim 0.2. Apart from the first iteration*, the total weight of all sets are monotonically non-increasing.*

Proof. At the $(i+1)$-th iteration, the weight increase due to the choice of ROW PLAYER is:

$$\Omega^{i+2}(\mathcal{R}) = \Omega^{i+1}(\mathcal{R}) + \sum_{\substack{S \in \mathcal{R}: \\ x_{i+1} \in S}} \omega^{i+1}(S). \tag{0.3}$$

Note that the element x_{i+1} was available to be picked by the COLUMN PLAYER in the *previous* iteration i—with the goal of maximizing precisely the same summation as the second term in the R.H.S. of Equation (0.3). Therefore it must be that the weight of sets of $\mathcal{R}$ set to 0 in iteration i was at least as large as the increase in the total weight at iteration $i+1$. Thus

Wt decrease (COLUMN PLAYER) at iteration i + Wt increase (ROW PLAYER) at iteration $i+1 \leq 0$.

Summing up over all iterations i completes the claim. □

Now, since all weights are non-negative, as long as the total weight of the sets does not reach 2^k in any iteration, the COLUMN PLAYER wins. To this end:

At the start, we have $\Omega^1(\mathcal{R}) = m$.

After the ROW PLAYER chooses x_1, in the worst case, the weight of all sets could have doubled and so we have $\Omega^2(\mathcal{R}) \leq 2m$.

From now onwards, the total weight of the sets do not increase due to Claim 0.2.

Thus COLUMN PLAYER wins if

$$2m < 2^k \quad \Longrightarrow \quad m < 2^{k-1}.$$

□

Bibliography and discussion. The algorithm and analysis is from [ES73].

[ES73] P. Erdős and J. L. Selfridge. “On a combinatorial game”. In: *Journal of Combinatorial Theory, Series A* 14.3 (1973), pp. 298–301.

Rounding via Duality and Epsilon-Nets

Let V be a set of n elements, and $\mathcal{S}$ a collection of m subsets of V. Further, let $\alpha > 0$ be a parameter such that for *any* weight function $w \colon \mathcal{S} \to \mathbb{R}^+$, there exists an element $v \in V$ hitting sets of total weight at least an α-th fraction of the total weight of $\mathcal{S}$. That is,

$$\sum_{S \in \mathcal{S} \colon v \in S} w(S) \geq \alpha \cdot \left(\sum_{S \in \mathcal{S}} w(S) \right).$$

Then what is the size of a minimum hitting set of $\mathcal{S}$?

A natural greedy algorithm constructs one of size $O\left(\frac{\ln m}{\alpha}\right)$:

> add a point hitting the maximum *number* of sets in $\mathcal{S}$ to our hitting set, and reiterate on the remaining unhit elements in $\mathcal{S}$. After having added the first point, the number of unhit sets in $\mathcal{S}$ is at most
>
> $$|\mathcal{S}| - \alpha |\mathcal{S}| = |\mathcal{S}| \cdot (1 - \alpha).$$
>
> And after adding the second point, the number of unhit sets in $\mathcal{S}$ is at most
>
> $$|\mathcal{S}| (1 - \alpha) - \alpha \cdot |\mathcal{S}| (1 - \alpha) = |\mathcal{S}| \cdot (1 - \alpha)^2.$$
>
> And after adding t points, the number of unhit sets in $\mathcal{S}$ is at most
>
> $$|\mathcal{S}| (1 - \alpha)^t.$$
>
> Setting $t = \frac{\ln(2|\mathcal{S}|)}{\alpha}$, the number of unhit sets in $\mathcal{S}$ after t iterations is at most
>
> $$|\mathcal{S}| \cdot (1 - \alpha)^t \leq |\mathcal{S}| \cdot \exp(-\alpha t) = |\mathcal{S}| \cdot \exp\left(-\alpha \frac{\ln 2|\mathcal{S}|}{\alpha}\right) = \frac{1}{2} < 1.$$
>
> That is, all sets are hit and thus the chosen points form a hitting set, of size $t = O\left(\frac{\ln m}{\alpha}\right)$.

We now present a technique where, if an additional structural property is true, one can construct a hitting set of size depending *only* on α. Here is the key insight, that follows immediately from LP duality.

Lemma 0.4. *Let V be a finite set of n elements, $\mathcal{S}$ a finite collection of subsets of V, and $\alpha \in (0, 1]$ a parameter such that the following is true:*

> *for any weight function $w \colon \mathcal{S} \to \mathbb{R}^+$, there exists an element $v \in V$ such that*
>
> $$\sum_{S \in \mathcal{S} \colon v \in S} w(S) \geq \alpha \cdot \left(\sum_{S \in \mathcal{S}} w(S) \right).$$

Then there exists a weight function $w_V \colon V \to \mathbb{R}^+$ such that for any $S \in \mathcal{S}$,

$$\sum_{v \in S} w_V(v) \geq \alpha \cdot \left(\sum_{v \in V} w_V(v) \right).$$

Proof. We will use a LP to construct the weight function w_V, and then prove that it satisfies the required property using LP duality.

The *covering* LP below has $|V|$ variables, while its dual *packing* LP has $|\mathcal{S}|$ variables. All variables are non-negative.

Covering LP (Primal)		**Packing LP (Dual)**
Minimize $\sum_{v \in V} x_v$	$\Longleftrightarrow$	Maximize $\sum_{S \in \mathcal{S}} y_S$
subject to		subject to
for each $S \in \mathcal{S}$: $\sum_{u \in S} x_u \geq 1.$		for each $v \in V$: $\sum_{S \in \mathcal{S}:\, v \in S} y_S \leq 1.$

Claim 0.5. *The dual LP—and thus the primal LP—has value at most* $\frac{1}{\alpha}$.

Proof. Let $W = \sum_{S \in \mathcal{S}} y_S$ be the value of the dual LP.

By our input assumption applied to the weights given by the y_S variables, there exists an element $v' \in V$ hitting sets of $\mathcal{S}$ with total weight at least αW.

On the other hand, the dual LP constraint implies that this is at most 1 for any $v \in V$, and thus

$$\alpha \cdot W \;\leq \sum_{S \in \mathcal{S} : v' \in S} y_S \;\leq\; 1, \qquad \text{implying that} \quad W \leq \frac{1}{\alpha}.$$

By LP duality, the value of primal LP is also at most $\frac{1}{\alpha}$. □

Now we can set the weight function w_V:

$$w_V(v) = x_v \qquad \text{for each } v \in V.$$

The primal LP has value at most $1/\alpha$ by Claim 0.5, and every $S \in \mathcal{S}$ contains elements of total weight at least 1. Thus each set of $\mathcal{S}$ has total weight at least an α-th fraction of the total weight under w_V, as desired. □

In game theory, the above lemma is typically stated in the setting of a 2-player game. From a graph-theoretic point of view, V and $\mathcal{S}$ are vertices of the bipartite incidence graph.

Lemma 0.4 converts our initial problem into, in the language of sampling theory, an ϵ-net problem: our initial goal is reduced to computing a hitting set for a set system $(V, \mathcal{S})$ where each set of $\mathcal{S}$ contains an α-th fraction of the total weight.

The following sampling approach, while only giving the same bound as the greedy algorithm, points us in the right direction:

> Independently, with replacement, and with respect to the weights w_V's, choose a random element of V repeatedly t times, to get a random sample R of size at most t. Then
>
> $$\Pr\left[\text{a fixed } S \in \mathcal{S} \text{ is not hit by } R\right] \le (1-\alpha)^t \le \exp\left(-\alpha\, t\right).$$
>
> By the union bound,
>
> $$\Pr\left[\text{some } S \in \mathcal{S} \text{ is not hit by } R\right] \le |\mathcal{S}| \cdot \exp\left(-\alpha\, t\right).$$
>
> Now setting $t = \frac{\ln(2|\mathcal{S}|)}{\alpha}$, the above probability becomes at most $\frac{1}{2}$, and so there exists a hitting set of $\mathcal{S}$ of size $O\left(\frac{\ln m}{\alpha}\right)$.

We next present three statements and their proofs, all of which follow the above formula:

1. Show the existence of a point hitting many sets.
2. Use LP duality to assign weights to points such that each set has high weight.
3. Show, using properties for the given scenario, that one can 'round' the weights to integer ones to get a hitting set.

APPLICATION 1: THE (p, q)-THEOREM FOR CONVEX SETS

Theorem 0.6. *Let $\mathcal{C} = \{C_1, \ldots, C_n\}$ be a set of n compact convex objects in $\mathbb{R}^d$, satisfying the following (p, q)-property for given integers $p \ge q \ge d+1$:*

> *for any $\mathcal{C}' \subseteq \mathcal{C}$ of size p, there exists a subset of $\mathcal{C}'$ of size q that can be hit by a point in $\mathbb{R}^d$.*

Then there exists a $Q \subseteq \mathbb{R}^d$, $|Q| = \tilde{O}\left(\left(p^{\frac{q-1}{q-d}}\right)^d\right)$, that is a hitting set for all sets in $\mathcal{C}$.

Proof. As $\mathcal{C}$ is finite, let P be a finite set of points, one from each distinct cell in the arrangement of $\mathcal{C}$ in $\mathbb{R}^d$. It suffices to restrict our hitting set for $\mathcal{C}$ to a subset of P.

The first step is to show, given the (p, q)-property, the following statement (stated without proof):

Lemma 0.7. *For any weight distribution $w\colon \mathcal{C} \to \mathbb{R}^+$ with $W = \sum_{C \in \mathcal{C}} w(C)$, there exists a point $r \in P$ such that*

$$\sum_{C \in \mathcal{C}\colon r \in C} w(C) \geq \alpha_{p,q,d} \cdot W, \qquad \textit{where} \qquad \alpha_{p,q,d} = \Omega\left(\frac{1}{p^{\frac{q-1}{q-d}}}\right).$$

Now apply Lemma 0.4 with $V = P$, $\mathcal{S} = \mathcal{C}$ and $\alpha = \alpha_{p,q,d}$ as given by Lemma 0.7 to get the following:

> there exists a function $w_P\colon P \to \mathbb{R}^+$ such that each $C \in \mathcal{C}$ contains points with total weight at least an $\alpha_{p,q,d}$-th fraction of the total weight of P under w_P.

Finally we need the following theorem on *weak ϵ-nets* (stated without proof):

Theorem 0.8. *Given a finite set P of points in $\mathbb{R}^d$ and a parameter $\epsilon > 0$, there exists a set $Q \subseteq \mathbb{R}^d$ such that any convex set in $\mathbb{R}^d$ containing at least $\epsilon |P|$ points of P contains at least one point of Q, where*

$$|Q| = \tilde{O}\left(\frac{1}{\epsilon^d}\right).$$

Applying Theorem 0.8 to P with weights given by w_P and with $\epsilon = \alpha_{p,q,d}$, we get a set $Q \subseteq \mathbb{R}^d$ that hits all sets of $\mathcal{C}$, with

$$|Q| = \tilde{O}\left(\frac{1}{(\alpha_{p,q,d})^d}\right) = \tilde{O}\left(\left(p^{\frac{q-1}{q-d}}\right)^d\right).$$

□

Application 2: Separating Points From a Convex Set

Given a finite set $P \subseteq \mathbb{R}^d$ and a convex set C, Carathéodory's Theorem implies the following: if the convex-hull of every $(d+1)$-sized subset of P is disjoint from C, then P can be separated from C by *one* hyperplane. We now show that by using more than one hyperplane to separate C from P, one can extend this to low-dimensional intersections.

Theorem 0.9. *Let C be a convex set in $\mathbb{R}^d$ and P a finite set of points in $\mathbb{R}^d$ satisfying the following property:*

the convex-hull of any $\left(\lfloor \frac{d}{2} \rfloor + 1\right)$*-sized subset of* P *is disjoint from* C.

Then C *can be separated from* P *by a set* Q *of* $O\left(d^4 \log d\right)$ *hyperplanes; that is, for each* $p \in P$*, there exists a hyperplane in* Q *separating* p *from* C.

Proof. Given P and C, let $\mathcal{H}$ be a finite set of hyperplanes, all disjoint from C, such that for any $P' \subseteq P$ where $\operatorname{conv}(P')$ is separable from C by a hyperplane, there exists a $h \in \mathcal{H}$ separating $\operatorname{conv}(P')$ from C. Thus it suffices to construct Q from the hyperplanes in $\mathcal{H}$.

For each $p \in P$, let $\mathcal{H}_p \subseteq \mathcal{H}$ be the hyperplanes separating p from C. Further, for each $h \in \mathcal{H}$, let h^+ denote the half-space defined by h that does *not* contain C.

The first step is to prove the existence of a hyperplane separating many points of P from C.

Lemma 0.10. *Given a rational weight function* $w\colon P \to \mathbb{R}^+$ *with* $W = \sum_{p \in P} w(p)$*, there exists a hyperplane* $h \in \mathcal{H}$ *such that*

$$\sum_{p \in P\colon h \in \mathcal{H}_p} w(p) \geq \alpha_d \cdot W, \qquad \text{where} \quad \alpha_d = \Omega\left(\frac{1}{d^3}\right).$$

Stated equivalently, letting $\mathcal{S} = \{\mathcal{H}_p \colon p \in P\}$ *with weights* $w(\mathcal{H}_p) = w(p)$*, there exists a* $h \in \mathcal{H}$ *hitting sets of* $\mathcal{S}$ *of total weight at least an* α_d*-th fraction of the total weight of the sets of* $\mathcal{S}$.

Proof. Assume $w(p) = \hat{w}(p)/D$, where $\hat{w}(p)$ and D are integers. Let P' be the point set constructed by replacing each $p \in P$ with $\hat{w}(p)$ copies of p. We need the following (stated without proof).

Theorem 0.11. *Given any finite set* P *in* $\mathbb{R}^d$*, there exists a set* $Q \subseteq P$ *of size* $\lfloor \frac{d}{2} \rfloor + 1$ *such that any half-space containing* Q *contains at least* $\alpha_d \cdot |P|$ *points of* P.

Apply Theorem 0.11 to P' to get a $\lfloor \frac{d}{2} \rfloor + 1$-sized set $Q \subseteq P'$ such that *any* half-space containing Q contains at least $\alpha_d |P'|$ points of P'. As this applies to the half-space h^+ corresponding to $h \in \mathcal{H}$ separating $\operatorname{conv}(Q)$ from C, we get

$$\sum_{p \in P\colon h \in \mathcal{H}_p} w(p) = \frac{1}{D} \sum_{p \in P\colon h \in \mathcal{H}_p} \hat{w}(p) \geq \frac{1}{D} \cdot \alpha_d \cdot |P'| = \alpha_d \cdot W.$$

□

Apply Lemma 0.4 with $V = \mathcal{H}$, $\mathcal{S} = \{\mathcal{H}_p \colon p \in P\}$, and $\alpha = \alpha_d$ from Lemma 0.10, to get:

there exists a function $w_{\mathcal{H}}\colon \mathcal{H} \to \mathbb{R}^+$ such that for each $p \in P$, the total weight of all the hyperplanes in $\mathcal{H}$ separating p from C is at least an α_d-th fraction of the total weight of $\mathcal{H}$.

Finally, we need the following theorem on ϵ-*nets* for half-spaces:

Theorem 0.12. *Given any finite set $\mathcal{H}$ of half-spaces in $\mathbb{R}^d$ and an $\epsilon > 0$, there exists a set $Q \subseteq \mathcal{H}$ of size $O\left(\frac{d}{\epsilon}\log\frac{1}{\epsilon}\right)$ such that any point of $\mathbb{R}^d$ contained in at least $\epsilon|\mathcal{H}|$ half-spaces of $\mathcal{H}$ is contained in some half-space of Q.*

By applying Theorem 0.12 to the half-spaces $\{h^+\colon h \in \mathcal{H}\}$, $\epsilon = 1/\alpha_d$, and weights given by $w_{\mathcal{H}}$, there exists a set $Q \subseteq \mathcal{H}$ of

$$O\left(\frac{d}{\epsilon}\log\frac{1}{\epsilon}\right) = O\left(d \cdot \alpha_d \cdot \log \alpha_d\right) = O\left(d^4 \log d\right)$$

hyperplanes such that each point of p is separated by some hyperplane in Q. This completes the proof. □

Bibliography and discussion.

LP Duality

Consider the following primal LP:

Primal LP

$$\max_{x \in \mathbb{R}^n} c \cdot x$$

$$\text{s.t.}$$

$$Ax \leq b,$$

$$x \geq 0.$$

Here $c \in \mathbb{R}^n$, A is a $m \times n$ matrix, and $b \in \mathbb{R}^m$.

Let $r_1, \ldots, r_m$ be the m rows of A, and $c_1, \ldots, c_n$ be the n columns of A.

The feasible region in $\mathbb{R}^n$, denoted by $\mathcal{C}$, is defined by the intersection of the m half-spaces (called constraints), where the j-th half-space, $j \in [m]$, is defined by the inequality

$$r_j \cdot x \leq b_j.$$

Any $x \in \mathcal{C}$ satisfies all the above m inequalities.

We will assume that $\mathcal{C}$ is non-empty and compact, $o \in \mathcal{C}$, and that $m \geq n$.

Fact 0.13. *Any positive linear combination of the constraints is also a half-space. Further, it contains the feasible region $\mathcal{C}$.*

Proof. Let $y_1 \geq 0, \ldots, y_m \geq 0$ be m non-negative coefficients for the m constraints. Then, by definition, for any $p \in \mathcal{C}$, we have for all $j \in [m]$:

$$r_j \cdot p \leq b_j \qquad \Longleftrightarrow \qquad y_j\,(r_j \cdot p) \leq y_j\,b_j.$$

Adding them up, we get that for any $p \in \mathcal{C}$,

$$\sum_{j \in [m]} y_j\,(r_j \cdot p) \leq \sum_{j \in [m]} y_j\,b_j,$$

or equivalently,

$$\left(\sum_{j \in [m]} y_j r_j\right) \cdot p \leq \sum_{j \in [m]} y_j\,b_j. \tag{0.14}$$

In words, as a $p \in \mathcal{C}$ satisfies all the m constraints, it also satisfies the sum of their *non-negative* linear combination given by Equation (0.14), which is just the half-space with the equation:

$$\underbrace{\left(\sum_{j\in[m]} y_j r_j \right)}_{\text{a vector } r'\in\mathbb{R}^n} \cdot x \le \sum_{j\in[m]} y_j\, b_j.$$

Geometrically, this is the half-space with the normal vector r'.

□

The next, intuitively obvious, geometric claim is at the heart of the argument.

Fact 0.15. *Let $v \in \mathbb{R}^n$ be a vertex of $\mathcal{C}$ that is the common intersection of n constraints, where the k-th constraint is denoted by the half-space*

$$\boldsymbol{h_k^-} : \quad r_k \cdot x \le b_k.$$

*Then **any** half-space tangent to $\mathcal{C}$ at v and containing $\mathcal{C}$, can be derived by a **convex** combination of the above n constraints.*

Proof.

Broadly, we take the *geometric* dual of each of these n hyperplanes with respect to the origin o (recall that $o \in \mathcal{C}$). Then,

$a)$ the dual of the convex-hull of these n dual points is exactly the space of all hyperplanes tangent to $\mathcal{C}$ at v, and

$b)$ a convex combination of the dual points corresponds to a convex combination of the half-space constraints.

This is essentially where Farkas' lemma comes into play—though it is applied to the geometric dual space where the n constraints defining v become points.

For a hyperplane $h \subset \mathbb{R}^n$, let h^- be the half-space defined by h and containing the origin o. Let $f(\cdot)$ be the standard geometric duality function with respect to the origin o. That is, $f(\cdot)$ maps

$a)$ hyperplanes in $\mathbb{R}^n$ not containing $o \longrightarrow$ points in $\mathbb{R}^n \setminus \{o\}$, and

$b)$ points in $\mathbb{R}^n \setminus \{o\} \longrightarrow$ hyperplanes in $\mathbb{R}^n$,

such that this duality function has two properties:

Preserves sidedness. A half-space h^- contains a point p if and only if the half-space $f(p)^-$ contains the point $f(h)$.

Preserves convex combinations. For any hyperplane h and hyperplanes $h_i \colon r_i \cdot x = b_i$, $b_i > 0$ and $i \in [t]$:

if $f(h)$ is a convex combination of the $f(h_i)$'s, then the hyperplane h is a convex combination of the hyperplane h_i's. That is,

$$\exists\, \alpha_1, \ldots, \alpha_t \geq 0 \colon \sum_{i \in [t]} \alpha_i = 1 \text{ and } f(h) = \sum_{i \in [t]} \alpha_i f(h_i)$$

$$\Updownarrow$$

$$\exists\, \alpha'_1, \ldots, \alpha'_t \geq 0 \colon \sum_{i \in [t]} \alpha'_i = 1 \text{ and } h \text{ can be written as } \left(\sum_{i \in [t]} \alpha'_i r_i \right) \cdot x = \sum_{i \in [t]} \alpha'_i b_i.$$

Fix a hyperplane h containing the given vertex v. We now show that h intersects $\operatorname{int}(\mathcal{C})$ if and only if the point $f(h)$ does not lie in $\operatorname{conv}(\{f(h_1), \ldots, f(h_n)\})$.

—$f(h)$ does not lie in $\operatorname{conv}\big(\{f(h_1), \ldots, f(h_n)\}\big) \implies h$ intersects $\operatorname{int}(\mathcal{C})$:

1. By the dual property of sidedness preservation, the points $f(h), f(h_1), \ldots, f(h_n)$ lie on the hyperplane $f(v)$.
2. Applying Farkas' lemma, in $\mathbb{R}^{n-1}$, to

   $$f(h), f(h_1), \ldots, f(h_n) \subset f(v),$$

   there is a $(n-2)$-dimensional hyperplane $h' \subset f(v)$ separating $f(h)$ from

   $$\operatorname{conv}\big(\{f(h_1), \ldots, f(h_n)\}\big).$$

3. We can extend h' to a $(n-1)$-dimensional hyperplane, say denoted by $f(q)$, that separates $f(h)$ from

   $$\operatorname{conv}\left(o \bigcup \{f(h_1), \ldots, f(h_n)\}\right).$$

4. That is, $f(q)^-$ contains $\big\{f(h_1), \ldots, f(h_n)\big\}$ and does not contain $f(h)$.

   By the dual property of sidedness preservation,

   (a) $h_1^-, \ldots, h_n^-$ contain q, and
   (b) h^- does not contain q.

5. Thus h must intersect $\operatorname{int}(\mathcal{C})$.

—h intersects $\operatorname{int}(\mathcal{C}) \implies f(h)$ does not lie in $\operatorname{conv}(\{f(h_1), \ldots, f(h_n)\})$:

Then there is a point $q \in \mathcal{C}$ that is not contained in h^-. That is, the half-space $f(q)^-$ contains $f(h_1), \ldots, f(h_n)$ but it does not contain $f(h)$. Therefore $f(h)$ cannot lie in $\operatorname{conv}(\{f(h_1), \ldots, f(h_n)\})$.

Finally, the dual property of convex combination preservation completes the proof.

$\square$

Remark: Rather than just using the n constraints around the vertex v, one can also do the above proof, without any change, by using all the m half-space constraints defining $\mathcal{C}$. This gives the following (weaker) statement:

h intersects $\operatorname{int}(\mathcal{C})$ if and only if the point $f(h)$ does not lie in $\operatorname{conv}(\{f(h_1), \ldots, f(h_m)\})$.

Note that then one uses Farkas' Lemma in $\mathbb{R}^n$ instead of $\mathbb{R}^{n-1}$.

Returning to our LP, let v_c be the vertex of $\mathcal{C}$ that is extreme in the given direction c. Denote the n constraints defining v, for $k = 1, \ldots, n$, by:

$$r_{j_k} \cdot x \leq b_{j_k}.$$

Then we have $r_{j_k} \cdot v_c = b_{j_k}$ for all $k \in [n]$.

Let h^- be the half-space containing $\mathcal{C}$ that is tangent to $\mathcal{C}$ at v, and orthogonal to c. By Fact 0.15, h^- can be written, for some $y_{j_1} \geq 0, \ldots, y_{j_n} \geq 0$, as:

$$\left(\sum_{k \in [n]} y_{j_k} \, r_{j_k} \right) \cdot x \leq \sum_{k \in [n]} y_{j_k} b_{j_k}.$$

The normal vector of this half-space—that is, $\sum_{k \in [n]} y_{j_k} \, r_{j_k}$—is in direction c; by scaling the y_{j_k}'s, we can assume that it is precisely c:

$$c = \sum_{k \in [n]} y_{j_k} \, r_{j_k}.$$

After scaling, the coefficients y_{j_k}'s are no longer convex—that is, their sum is no longer 1—but are non-negative.

This proves strong duality: since v_c lies on the bounding hyperplane of h^-, we have

$$\left(\sum_{k \in [n]} y_{j_k} \, r_{j_k} \right) \cdot v_c = \sum_{k \in [n]} y_{j_k} b_{j_k} \qquad \Longleftrightarrow \qquad c \cdot v_c = \sum_{k \in [n]} y_{j_k} b_{j_k}.$$

Formally, we have to also consider the coefficients of the other $m - n$ constraints, which are set to 0. That is, let $y \in \mathbb{R}^m$ be a vector with the above y_{j_k}'s at their proper indices, and set the remaining $m - n$ indices to 0. Then we have

$$c \cdot v_c = \left(\sum_{k \in [n]} y_{j_k} \, r_{j_k} \right) \cdot v_c = \sum_{k \in [n]} y_{j_k} \, (r_{j_k} \cdot v_c) = \sum_{k \in [n]} y_{j_k} b_{j_k} = \sum_{j \in [\boldsymbol{m}]} y_j b_j = b \cdot y.$$

Also, it is easy to see that y is a feasible dual solution.

Dual LP

$$\min_{y \in \mathbb{R}^m} \; b \cdot y$$

$$\text{s.t.}$$

$$y^T A \geq c,$$

$$y \geq 0.$$

Remark: In our proof above, we are getting the stronger property $y^T A = c$ instead of $y^T A \geq c \ldots$ and this is not correct.

Indeed, the proof given above is incomplete: it ignores the case where the extremal vertex in direction c is realized by constraints that include the coordinate planes, due to our condition that $x \geq 0$.

To fix this, include these additional n constraints: for $i = 1, \ldots, n$,

$$-e_i \cdot x \leq 0.$$

The i-th constraint requires the i-th coordinate of x to be positive.

Then, as before, we can write c as a non-negative linear combination of constraint half-spaces, this time including the non-negativity constraints. That is, h^- can be written as

$$\underbrace{\left(\sum_{j \in [m]} y_j r_j + \sum_{i \in [n]} y_i (-e_i) \right)}_{c} \cdot x \leq \sum_{j \in [m]} y_j \, b_j + \sum_{i \in [n]} y_i \, 0.$$

Now removing the terms on the L.H.S. due to the non-negativity constraints gives the required form:

$$\underbrace{\left(\sum_{j\in[m]} y_j r_j\right)}_{\geq c} \cdot x \leq \sum_{j\in[m]} y_j\, b_j.$$

We note the primal and dual slackness conditions, and their geometric interpretations.

Let $x \in \mathbb{R}^n$ and $y \in \mathbb{R}^m$ be *feasible* primal and dual solutions. Then they are optimal if and only if all these conditions hold:

Dual complementary slackness conditions. For each $j \in [m]$:

$$\text{either } y_j = 0 \text{ or } \; r_j \cdot x = b_j.$$

> Given y, keep the constraints whose corresponding y value is non-zero. Then x must lie on each of these constraints. Usually there are n such constraints, representing a vertex of $\mathcal{C}$ which should be the optimal solution x.
>
> Note that these conditions alone do not say that x must be maximum point of $\mathcal{C}$ in the direction c—just that x, *which lies in $\mathcal{C}$ by the feasibility assumption*, must lie on the common intersection of some (at most n) constraints.

Primal complementary slackness conditions. For each $i \in [n]$:

$$\text{either } x_i = 0 \text{ or } \left(\sum_{j\in[m]} y_j r_j\right)_i = c_i.$$

> Assume, for the moment, that none of the x_i's are 0.
>
> Then the linear combination—as given by the y_j's—of the normal vector r_j's of the constraints (really, we just care here about those constraints with $y_i \neq 0$) must equal the vector c.
>
> That is, the linear combination of the constraints of $\mathcal{C}$ with $y_j \neq 0$ gives a half-space with normal vector c.

Together, these two conditions fix the optimal x and y solutions:

the first set of conditions tells us that x must lie on the boundary of $\mathcal{C}$, and the second set of conditions tells us that there is a tangent half-space to this boundary point x which is normal to the direction c.

Thus x must be the optimal bound in direction c, and y the precise coefficients that give us the extremal half-space tangent to $\mathcal{C}$ that contains x.

Purely mechanically, each constraint in **Primal** corresponds to a variable in the **Dual**, and each variable in the **Primal** corresponds to a constraint in the **Dual**. Then x and y, feasible solutions, are optimal if and only if:

whenever a primal variable is not 0, its corresponding dual constraint is tight, and

whenever a dual variable is not 0, its corresponding primal constraint is tight.

Bibliography and discussion. See Schrijver's A First Course in Combinatorial Optimization, section 2.4.

Linear Classifiers: Winnow Algorithm

We are given m vectors $V = \{v_1, \ldots, v_m\}$ in $\mathbb{R}^n$, with the property that there exists a hyperplane through the origin that contains all of them on one side. In fact, we assume something slightly stronger:

there exists a parameter $\epsilon > 0$ and a vector $u^* \geq 0$ with $\sum_{i=1}^n (u^*)_i = 1$, such that

$$\text{for all } v \in V: \quad u^* \cdot v \geq \epsilon.$$

Then our goal is to find a vector $u \in \mathbb{R}^n$ such that $u \cdot v_i \geq 0$, for all $i \in [m]$.

Let ρ be the absolute value of the largest coordinate in V. That is,

$$\rho = \max_{v \in V} \max_{i \in [n]} |(v)_i|.$$

In this section, we show that one can find a separating vector u with the MWU technique.

Overview of ideas. To find an approximation to u^*, we start with the candidate vector $\omega^1 = (1, \ldots, 1) \in \mathbb{R}^n$. Iteratively, at step t, we find a violating vector $v^t \in V$ with $\omega^t \cdot v^t < 0$ and compute the next point by modifying ω^t, somewhat surprisingly, *multiplicatively coordinate-wise*. That is, for each coordinate $i \in [m]$,

$$(\omega^{t+1})_i = (\omega^t)_i + \omega_i^t \cdot \eta \frac{(v^t)_i}{\rho}.$$

The key point is that now our standard weight function, $\Omega^t = \sum_{i \in [m]} (\omega^t)_i$, is decreasing:

$$\Omega^{t+1} = \sum_{i \in [m]} (\omega^{t+1})_i = \sum_{i \in [m]} (\omega^t)_i + \frac{\eta}{\rho} \sum_{i \in [m]} (\omega^t)_i \cdot (v^t)_i < \sum_{i \in [m]} (\omega^t)_i = \Omega^t,$$

since $\omega^t \cdot v^t < 0$.

As usual, computing upper and lower bounds for Ω^T, we will arrive—ignoring secondary quadratic terms for the moment—at the conclusion that for each coordinate $i \in [m]$,

$$\sum_{t=1}^T \frac{(v^t)_i}{T} \leq \frac{\rho \ln n}{T \eta}. \tag{0.16}$$

In other words, each coordinate of the average point $\tilde{v} = \sum_{t=1}^T v^t / T$ is upper bounded by the R. H. S. above, as a decreasing function of T.

However, along the direction of the vector $u^* \geq 0$, the projection lengths of all $t \in [T]$ satisfy $u^* \cdot v^t \geq \epsilon$, and so by pigeonhole principle, one coordinate of $\tilde{v}$ would become too large, for T sufficiently large, going along direction u^*. But this would contradict to Equation (0.16). That is,

$$\epsilon \leq \sum_{t=1}^{T} \boldsymbol{u^* \cdot \frac{v^t}{T}} = \sum_{t=1}^T \sum_{i \in [m]} (u^*)_i \cdot \frac{(v^t)_i}{T} = \sum_{i \in [m]} (u^*)_i \cdot \left(\sum_{t=1}^T \frac{(v^t)_i}{T} \right) \leq \frac{\rho \ln n}{T \eta}.$$

The MWU algorithm is as follows.

Initialize $\omega^1 \in \mathbb{R}^n$ to be the all 1's vector, and let η be a scaling parameter to be set later.

Then for iteration $t = 1, \ldots$:

1. let $v^t \in V$ be a vector violating our constraints. That is,

$$v^t \cdot \omega^t < 0.$$

   If there does not exist such a violating vector in V, we're done:

   – return ω^t as the desired vector.

2. update each coordinate of ω^t: for $i = 1, \ldots, n$, set

$$\omega_i^{t+1} = \omega_i^t \left(1 + \eta \frac{(v^t)_i}{\rho}\right),$$

   for a small-enough parameter η to be set later.

We claim that there exists a large-enough value T such that the above algorithm succeeds in at most T iterations.

Let's say the algorithm continues for T steps.

As usual, we upper and lower bound the total weight $\Omega^t = \sum_{i=1}^n \omega_i^t$.

Upper bound. By the choice of v^t at each step t, we have

$$\Omega^{t+1} = \Omega^t + \sum_{i=1}^n \eta \frac{(v^t)_i}{\rho} \cdot \omega_i^t = \Omega^t + \frac{\eta}{\rho} \left(v^t \cdot \omega^t\right) \leq \Omega^t.$$

Lower bound.

$$\omega_i^{T+1} = \prod_{t=1}^T \left(1 + \eta \frac{(v^t)_i}{\rho}\right).$$

Thus we have, for each $i = 1, \ldots, n$:

$$\prod_{t=1}^T \left(1 + \eta \frac{(v^t)_i}{\rho}\right) \leq \Omega^{T+1} \leq n.$$

Taking logarithms and dividing by η,

$$\frac{1}{\eta}\ln\left(\prod_{t=1}^{T}\left(1+\eta\frac{(v^t)_i}{\rho}\right)\right)\leq\frac{\ln n}{\eta}.$$

Applying Lemma 0.54 with $a_t=\frac{(v^t)_i}{\rho}$ implies that, for each $i\in[1,n]$,

$$\sum_{t=1}^{T}\frac{(v^t)_i}{\rho}-\eta\sum_{t=1}^{T}\frac{|(v^t)_i|}{\rho}\leq\frac{\ln n}{\eta}.$$

Taking a convex combination of these n inequalities, where we weigh the i-th inequality with $(u^*)_i$, we get:

$$\sum_{i=1}^{n}\left((u^*)_i\cdot\left(\sum_{t=1}^{T}\frac{(v^t)_i}{\rho}-\eta\sum_{t=1}^{T}\frac{|(v^t)_i|}{\rho}\right)\right)\leq\frac{\ln n}{\eta}.$$

Writing in terms of dot products, we arrive at the final statement:

$$\sum_{t=1}^{T}u^*\cdot\frac{v^t}{\rho}-\eta\sum_{t=1}^{T}\sum_{i=1}^{n}(u^*)_i\frac{|(v^t)_i|}{\rho}\leq\frac{\ln n}{\eta}.$$

Using $u^*\cdot v^t\geq\epsilon$, $\sum_{i=1}^{n}(u^*)_i=1$, and $|(v^t)_i|\leq\rho$, it follows that

$$\frac{\epsilon}{\rho}\cdot T-\eta T\leq\frac{\ln n}{\eta}.$$

This gives the required bound:

$$T\leq\frac{\ln n}{\eta\left(\frac{\epsilon}{\rho}-\eta\right)}.$$

Setting $\eta=\frac{\epsilon}{2\rho}$ gives that

$$T=O\left(\frac{\rho^2}{\epsilon^2}\ln n\right).$$

Discrepancy: General Case

We give a MWU algorithm to compute a two-coloring with small discrepancy.

Theorem 0.41. *Let $(X, \mathcal{R})$ be a finite set system with $X = \{v_1, \ldots, v_n\}$ and $m = |\mathcal{R}|$. Then there is a deterministic MWU algorithm that computes a two-coloring of X with discrepancy $O\left(\sqrt{n \ln m}\right)$.*

Overview of ideas. We will color the elements of X sequentially, in the order $v_1, \ldots, v_n$, with a $+1$ or a -1 color. The elements that are so far uncolored will have color 0.

The idea is to maintain a weight for each set, where this weight depends *exponentially* on the current discrepancy of that set.

Let $\eta > 0$ be a parameter to be set later.

Define the weight of any $S \in \mathcal{R}$ as:

$$W(S) = \exp(\eta \cdot \mathrm{disc}(S)),$$

where $\mathrm{disc}(S)$ denotes the current discrepancy of S. That is, with the so-far uncolored elements having color 0.

Set

$$W(\mathcal{R}) = \sum_{S \in \mathcal{R}} W(S) = \sum_{S \in \mathcal{R}} \exp(\eta \cdot \mathrm{disc}(S)).$$

As with the MWU technique, when coloring element v_i, we will assign it a color that minimizes $W(\mathcal{R})$.

The key technical lemma is to show that, at each iteration, there is a choice of color for v_k such that the sum $W(\mathcal{R})$ grows slowly.

This then implies that no set can have too large a discrepancy.

Assume we have colored the elements $v_1, \ldots, v_{k-1}$ and now have to assign a color to v_k.

Let $\mathrm{disc}_k(\cdot)$ be the discrepancy and $W_k(\cdot)$ the weights, at the start of the k-th iteration. Note that for all $S \in \mathcal{R}$,

$$\mathrm{disc}_1(S) = 0 \qquad \text{and} \qquad W_1(S) = 1.$$

Claim 0.42. *At the start of the k-th iteration, let $\mathcal{R}' \subseteq \mathcal{R}$ be the sets with $\mathrm{disc}_k(\cdot) \neq 0$:*

$$\mathcal{R}' = \{S \in \mathcal{R} \colon \mathrm{disc}_k(S) \neq 0\}.$$

Then we can assign a color to v_k—that is, a $+1$ or a -1 value—such that

$$\sum_{S\in\mathcal{R}'} W_{k+1}(S) \leq \left(\frac{e^{\eta}+e^{-\eta}}{2}\right)\cdot \sum_{S\in\mathcal{R}'} W_k(S).$$

Proof. Set the color of v_k to $+1$ or -1 with equal probability. Then note that for any $S\in\mathcal{R}$ with $\operatorname{disc}(S)\neq 0$, the discrepancy of S increases by 1 or decreases by 1 with equal probability. Thus for any $S\in\mathcal{R}'$,

$$\begin{aligned}\mathrm{E}\left[W_{k+1}(S)\right] &= \frac{1}{2}\cdot e^{\eta(\operatorname{disc}_k(S)+1)} + \frac{1}{2}\cdot e^{\eta(\operatorname{disc}_k(S)-1)}\\ &= e^{\eta \operatorname{disc}_k(S)}\cdot\left(\frac{e^{\eta}+e^{-\eta}}{2}\right).\end{aligned}$$

By linearity of expectation, we have

$$\begin{aligned}\mathrm{E}\left[\sum_{S\in\mathcal{R}'} W_{k+1}(S)\right] &= \sum_{S\in\mathcal{R}'}\mathrm{E}\left[W_{k+1}(S)\right] = \sum_{S\in\mathcal{R}'} e^{\eta\operatorname{disc}_k(S)}\left(\frac{e^{\eta}+e^{-\eta}}{2}\right)\\ &= \left(\frac{e^{\eta}+e^{-\eta}}{2}\right)\cdot\sum_{S\in\mathcal{R}'} W_k(S).\end{aligned}$$

Thus for one of the two choices for the color of v_k, the desired statement holds.

□

Remark: The use of probability in the above proof is purely for 'implementing' an averaging argument. Essentially, we showed that

$$\begin{aligned}&\underbrace{\sum_{S\in\mathcal{R}'} W\big(S \mid \operatorname{color}(v_k)=+1\big)}_{W_{k+1}(\mathcal{R}')\text{ assuming }\operatorname{color}(v_k)=+1} + \underbrace{\sum_{S\in\mathcal{R}'} W\big(S \mid \operatorname{color}(v_k)=-1\big)}_{W_{k+1}(\mathcal{R}')\text{ assuming }\operatorname{color}(v_k)=-1}\\ &= \sum_{S\in\mathcal{R}'} W\big(S \mid \operatorname{color}(v_k)=+1\big) + W\big(S \mid \operatorname{color}(v_k)=-1\big)\\ &= \sum_{S\in\mathcal{R}'} e^{\eta}\cdot W_k(S) + e^{-\eta}\cdot W_k(S)\\ &= \left(e^{\eta}+e^{-\eta}\right)\cdot\sum_{S\in\mathcal{R}'} W_k(S),\end{aligned}$$

and so one of the two sums must be at most $\frac{1}{2}$ of the R.H.S. above.

For the moment, assume that for all $S\in\mathcal{R}$ and $k>0$, we always have $\operatorname{disc}_k(S)\neq 0$. Then we're done:

Upper and lower bounding the total weight, we get

$$\max_{S\in\mathcal{R}} W_{n+1}(S) \;\le\; W_{n+1}(\mathcal{R}) \;\le\; W_1(\mathcal{R})\cdot\left(\frac{e^{\eta}+e^{-\eta}}{2}\right)^{n}.$$

Using the inequality $e^{\eta}+e^{-\eta}\le 2e^{\eta^2/2}$ (Fact 0.43 below), and that $W_1(\mathcal{R})=m$,

$$\exp\left(\eta\cdot\max_{S\in\mathcal{R}}\operatorname{disc}_{n+1}(S)\right) \;\le\; W_{n+1}(\mathcal{R}) \;\le\; m\cdot e^{n\eta^2/2}.$$

Taking logarithms,

$$\max_{S\in\mathcal{R}}\operatorname{disc}_{n+1}(S)\le\frac{\ln m}{\eta}+\frac{n\eta}{2}.$$

The above is minimized by setting $\eta=\Theta\left(\sqrt{\frac{\ln m}{n}}\right)$, giving the desired upper bound on the discrepancy for each set.

Fact 0.43.

$$\frac{e^{\eta}+e^{-\eta}}{2}\le e^{\eta^2/2}.$$

Proof. Using Taylor series at 0 gives, for any $\eta>0$,

$$e^{\eta}=1+\frac{\eta}{1!}+\frac{\eta^2}{2!}+\frac{\eta^3}{3!}+\frac{\eta^4}{4!}+\cdots$$

$$e^{-\eta}=1-\frac{\eta}{1!}+\frac{\eta^2}{2!}-\frac{\eta^3}{3!}+\frac{\eta^4}{4!}+\cdots$$

Adding them up cancels the linear term—so the quadratic term becomes the dominant one for $\eta<1$—and we get

$$e^{\eta}+e^{-\eta}=2\left(1+\frac{\eta^2}{2!}+\frac{\eta^4}{4!}+\frac{\eta^6}{6!}+\frac{\eta^8}{8!}+\cdots\right)$$

Using the fact that $(2i)!\ge 2^i i!$,

$$<2\left(1+\frac{(\eta^2)}{2^1\cdot 1!}+\frac{(\eta^2)^2}{2^2\cdot 2!}+\frac{(\eta^2)^3}{2^3\cdot 3!}+\frac{(\eta^2)^4}{2^4\cdot 4!}+\cdots\right)$$

$$=2\left(1+\frac{(\eta^2/2)}{1!}+\frac{(\eta^2/2)^2}{2!}+\frac{(\eta^2/2)^3}{3!}+\frac{(\eta^2/2)^4}{4!}+\cdots\right)$$

$$= 2\, e^{\eta^2/2}.$$

□

The above does not *quite* work—we used Claim 0.42 which only applies to sets of $\mathcal{R}$ with $\mathrm{disc}_k(\cdot) \neq 0$. Indeed, the restriction to sets with $\mathrm{disc}_k(\cdot) \neq 0$ is necessary:

> The key property in Claim 0.42 is that the discrepancy of each $S \in \mathcal{R}'$ can *both* increase or decrease by 1. This is what allows us to upper bound the *average* multiplicative factor increase in the weight of each $S \in \mathcal{R}'$ by $(e^{\eta} + e^{-\eta})/2$.
>
> However, this is not true when $\mathrm{disc}(S) = 0$—then the discrepancy of S can only increase by 1, no matter what color is given to v_k, and so the multiplicative factor becomes e^{η}, which is too big by a factor of roughly 2.

We now present two ways to get around this problem:

Bounding total increase in weights. The weight function is the same as earlier:

$$W(S) = \exp(\eta \cdot \mathrm{disc}(S)).$$

As before:

1. we choose the color of v_k only by considering $S \in \mathcal{R}$ with $\mathrm{disc}(S) > 0$, and then applying Claim 0.42.
2. The total weight of the sets with $\mathrm{disc}(S) > 0$ increases by a multiplicative factor of at most $\left(\frac{e^{\eta}+e^{-\eta}}{2}\right)$.

However, additionally, the weight of each set with $\mathrm{disc}(S) = 0$ goes from 1 to e^{η}. But this is not really a problem:

> the weight of S is already small when $\mathrm{disc}(S) = 0$—it is $e^{\eta 0} = 1$, and will become e^{η}. These small weights can be incorporated in the calculation without significantly changing the upper bound.

Taking both into account, we have

$$W_{k+1}(\mathcal{R}) \leq m \cdot e^{\eta} + W_k \cdot \left(\frac{e^{\eta} + e^{-\eta}}{2}\right)$$

Opening it up inductively,

$$
\begin{aligned}
W_{n+1}(\mathcal{R}) &\le m \cdot e^{\eta} + \left(m \cdot e^{\eta} + W_n \cdot \left(\frac{e^{\eta} + e^{-\eta}}{2} \right) \right) \cdot \left(\frac{e^{\eta} + e^{-\eta}}{2} \right) \\
&= m \cdot e^{\eta} + m \cdot e^{\eta} \cdot \left(\frac{e^{\eta} + e^{-\eta}}{2} \right) + W_n \cdot \left(\frac{e^{\eta} + e^{-\eta}}{2} \right)^2 \\
&\ \ \vdots \\
&\le \left(\sum_{i=0}^{n-1} m \cdot e^{\eta} \cdot \left(\frac{e^{\eta} + e^{-\eta}}{2} \right)^i \right) + m \cdot \left(\frac{e^{\eta} + e^{-\eta}}{2} \right)^n \\
&\le m \cdot n \cdot e^{\eta} \cdot \left(\frac{e^{\eta} + e^{-\eta}}{2} \right)^n + m \cdot \left(\frac{e^{\eta} + e^{-\eta}}{2} \right)^n .
\end{aligned}
$$

Now the previous double-counting argument finishes the proof as before:

$$
\exp\left(\eta \cdot \max_{S \in \mathcal{R}} \operatorname{disc}_{n+1}(S) \right) \ \le\ W_{n+1}(\mathcal{R})\ \le 2m \cdot n \cdot e^{\eta} \cdot e^{n\eta^2/2}.
$$

Taking logarithms,

$$
\max_{S \in \mathcal{R}} \operatorname{disc}_n(S) = O\left(\frac{\ln mn}{\eta} + n\eta \right).
$$

Setting $\eta = \Theta\left(\sqrt{\ln m/n}\right)$ gives an upper bound of $O\left(\sqrt{n \ln m}\right)$, assuming $m \ge n$.

Using a different weight function. The trick here—on seeing the multiplicative factor of $\left(\frac{e^{\eta}+e^{-\eta}}{2}\right)$—is to slightly modify the weight function so that even when $\operatorname{disc}(S) = 0$, the weight increases by a smaller multiplicative factor.

We set the new weight function, denoted by $\omega(\cdot)$, to be:

$$
\omega(S) = \frac{\exp(\eta \cdot \operatorname{disc}(S)) + \exp(-\eta \cdot \operatorname{disc}(S))}{2}. \tag{0.44}
$$

Now note that even when $\operatorname{disc}_k(S) = 0$ with $\omega_k(S) = 1$, we have

$$
\omega_{k+1}(S) = \frac{\exp(\eta) + \exp(-\eta)}{2},
$$

which is the precise multiplicative increase we wanted.

Further, the general upper bound on the multiplicative weight increase continues to hold, as before, for the case $\mathrm{disc}(S) \neq 0$:

$$\begin{aligned}
\mathrm{E}\left[\omega_{k+1}(S)\right] &= \frac{1}{2}\left(\frac{e^{\eta\cdot(\mathrm{disc}_k(S)+1)} + e^{-\eta\cdot(\mathrm{disc}_k(S)+1)}}{2}\right) + \frac{1}{2}\left(\frac{e^{\eta\cdot(\mathrm{disc}_k(S)-1)} + e^{-\eta\cdot(\mathrm{disc}_k(S)-1)}}{2}\right) \\
&= \frac{e^{\eta\,\mathrm{disc}_k(S)}\cdot e^{\eta}}{4} + \frac{e^{-\eta\,\mathrm{disc}_k(S)}\cdot e^{-\eta}}{4} + \frac{e^{\eta\,\mathrm{disc}_k(S)}\cdot e^{-\eta}}{4} + \frac{e^{-\eta\,\mathrm{disc}_k(S)}\cdot e^{\eta}}{4} \\
&= \frac{e^{\eta\,\mathrm{disc}_k(S)}}{2}\left(\frac{e^{\eta}+e^{-\eta}}{2}\right) + \frac{e^{-\eta\,\mathrm{disc}_k(S)}}{2}\left(\frac{e^{\eta}+e^{-\eta}}{2}\right) \\
&= \left(\frac{e^{\eta\,\mathrm{disc}_k(S)} + e^{-\eta\,\mathrm{disc}_k(S)}}{2}\right)\cdot\left(\frac{e^{\eta}+e^{-\eta}}{2}\right) \\
&= \omega_k\left(S\right)\cdot\left(\frac{e^{\eta}+e^{-\eta}}{2}\right).
\end{aligned}$$

Now the previous double-counting argument finishes the proof.

> **Remark:** Here is one way to naturally derive the weight function given in Equation (0.44).
>
> Our goal is to minimize $\mathrm{disc}(S)$—in other words, for each $S \in \mathcal{R}$, the number of '$+1$' colors should not be too large, and neither should the number of '-1' colors.
>
> Our earlier weight function, $\exp\left(\eta\,\mathrm{disc}(S)\right)$, was capturing this compactly using the absolute value function. But the drawback of this is that it made it insensitive to the case when $\mathrm{disc}(S) = 0$.
>
> We can fix this by *separately* adding the two exponential constraints—one prohibiting too many '$+1$' colors, and the other prohibiting too many '-1' colors:
>
> > For each $S \in \mathcal{R}$, let P_S be the number of elements of color '$+1$', and N_S the number of elements of color '-1'.
> >
> > Then we minimize the weight function
> >
> > $$\exp\left(\eta\left(P_S - N_S\right)\right) + \exp\left(\eta(N_S - P_S)\right).$$
>
> This is exactly Equation (0.44) scaled by a factor of 2! The constant 2 is not important and could have been omitted—the calculation without it gives the same bound.

Bibliography and discussion. Another way one can arrive at the function $\frac{1}{2}\left(e^{\eta}+e^{-\eta}\right)$ is via the proof of the tail bound used to prove the $O\left(\sqrt{n \ln m}\right)$ bound for discrepancy via a random coloring (see [You95]).

[You95] N. E. Young. "Randomized Rounding Without Solving the Linear Program". In: *Proceedings of the Sixth Annual Symposium on Discrete Algorithms (SODA)*. 1995, pp. 170–178.