

Maximum Satisfiability

- Given a set of clauses, MAXSAT seeks the truth assignment that satisfies the most.
- MAX2SAT is already NP-complete (p. 333), so MAXSAT is NP-complete.
- Consider the more general k -MAXGSAT for constant k .
 - Let $\Phi = \{ \phi_1, \phi_2, \dots, \phi_m \}$ be a set of boolean expressions in n variables.
 - Each ϕ_i is a *general* expression involving up to k variables.
 - k -MAXGSAT seeks the truth assignment that satisfies the most expressions.

A Probabilistic Interpretation of an Algorithm

- Let ϕ_i involve $k_i \leq k$ variables and be satisfied by s_i of the 2^{k_i} truth assignments.
- A random truth assignment $\in \{0, 1\}^n$ satisfies ϕ_i with probability $p(\phi_i) = s_i/2^{k_i}$.
 - $p(\phi_i)$ is easy to calculate as k is a constant.
- Hence a random truth assignment satisfies an average of

$$p(\Phi) = \sum_{i=1}^m p(\phi_i)$$

expressions ϕ_i .

The Search Procedure

- Clearly

$$p(\Phi) = \frac{1}{2} \{ p(\Phi[x_1 = \mathbf{true}]) + p(\Phi[x_1 = \mathbf{false}]) \}.$$

- Select the $t_1 \in \{ \mathbf{true}, \mathbf{false} \}$ such that $p(\Phi[x_1 = t_1])$ is the larger one.
- Note that $p(\Phi[x_1 = t_1]) \geq p(\Phi)$.
- Repeat the procedure with expression $\Phi[x_1 = t_1]$ until all variables x_i have been given truth values t_i and all ϕ_i are either true or false.

The Search Procedure (continued)

- By our hill-climbing procedure,

$$\begin{aligned} & p(\Phi) \\ & \leq p(\Phi[x_1 = t_1]) \\ & \leq p(\Phi[x_1 = t_1, x_2 = t_2]) \\ & \leq \dots \\ & \leq p(\Phi[x_1 = t_1, x_2 = t_2, \dots, x_n = t_n]). \end{aligned}$$

- So at least $p(\Phi)$ expressions are satisfied by truth assignment (t_1, t_2, \dots, t_n) .

The Search Procedure (concluded)

- Note that the algorithm is *deterministic*!
- It is called **the method of conditional expectations**.^a

^aErdős and Selfridge (1973); Spencer (1987).

Approximation Analysis

- The optimum is at most the number of satisfiable ϕ_i —i.e., those with $p(\phi_i) > 0$.
- Hence the ratio of algorithm's output vs. the optimum is^a

$$\geq \frac{p(\Phi)}{\sum_{p(\phi_i) > 0} 1} = \frac{\sum_i p(\phi_i)}{\sum_{p(\phi_i) > 0} 1} \geq \min_{p(\phi_i) > 0} p(\phi_i).$$

- This is a polynomial-time ϵ -approximation algorithm with $\epsilon = 1 - \min_{p(\phi_i) > 0} p(\phi_i)$ by Eq. (19) on p. 702.
- Because $p(\phi_i) \geq 2^{-k}$ for a satisfiable ϕ_i , the heuristic is a polynomial-time ϵ -approximation algorithm with $\epsilon = 1 - 2^{-k}$.

^aRecall that $\sum_i a_i / \sum_i b_i \geq \min_i (a_i / b_i)$.

Back to MAXSAT

- In MAXSAT, the ϕ_i 's are clauses (like $x \vee y \vee \neg z$).
- Hence $p(\phi_i) \geq 1/2$ (why?).
- The heuristic becomes a polynomial-time ϵ -approximation algorithm with $\epsilon = 1/2$.^a
- Suppose we set each boolean variable to true with probability $(\sqrt{5} - 1)/2$, the golden ratio.
- Then follow through the method of conditional expectations to **derandomize** it.

^aJohnson (1974).

Back to MAXSAT (concluded)

- We will obtain a $\lfloor (3 - \sqrt{5}) \rfloor / 2$ -approximation algorithm.^a

- Note $\lfloor (3 - \sqrt{5}) \rfloor / 2 \approx 0.382$.

- If the clauses have k *distinct* literals,

$$p(\phi_i) = 1 - 2^{-k}.$$

- The heuristic becomes a polynomial-time ϵ -approximation algorithm with $\epsilon = 2^{-k}$.

- This is the best possible for $k \geq 3$ unless $P = NP$.

- All the results hold even if clauses are weighted.

^aLieberherr and Specker (1981).


MAX CUT Revisited

- MAX CUT seeks to partition the nodes of graph $G = (V, E)$ into $(S, V - S)$ so that there are as many edges as possible between S and $V - S$.
- It is NP-complete (p. 368).
- **Local search** starts from a feasible solution and performs “local” improvements until none are possible.
- Next we present a local-search algorithm for MAX CUT.

A 0.5-Approximation Algorithm for MAX CUT

- 1: $S := \emptyset$;
- 2: **while** $\exists v \in V$ whose switching sides results in a larger cut **do**
- 3: Switch the side of v ;
- 4: **end while**
- 5: **return** S ;

Analysis



Analysis (continued)

- Partition $V = V_1 \cup V_2 \cup V_3 \cup V_4$, where
 - Our algorithm returns $(V_1 \cup V_2, V_3 \cup V_4)$.
 - The optimum cut is $(V_1 \cup V_3, V_2 \cup V_4)$.
- Let e_{ij} be the number of edges between V_i and V_j .
- Our algorithm returns a cut of size

$$e_{13} + e_{14} + e_{23} + e_{24}.$$

- The optimum cut size is

$$e_{12} + e_{34} + e_{14} + e_{23}.$$

Analysis (continued)

- For each node $v \in V_1$, its edges to $V_1 \cup V_2$ are outnumbered by those to $V_3 \cup V_4$.
 - Otherwise, v would have been moved to $V_3 \cup V_4$ to improve the cut.
- Considering all nodes in V_1 together, we have

$$2e_{11} + e_{12} \leq e_{13} + e_{14}.$$

- $2e_{11}$, because each edge in V_1 is counted twice.
- The above inequality implies

$$e_{12} \leq e_{13} + e_{14}.$$

Analysis (concluded)

- Similarly,

$$e_{12} \leq e_{23} + e_{24}$$

$$e_{34} \leq e_{23} + e_{13}$$

$$e_{34} \leq e_{14} + e_{24}$$

- Add all four inequalities, divide both sides by 2, and add the inequality $e_{14} + e_{23} \leq e_{14} + e_{23} + e_{13} + e_{24}$ to obtain

$$e_{12} + e_{34} + e_{14} + e_{23} \leq 2(e_{13} + e_{14} + e_{23} + e_{24}).$$

- The above says our solution is at least half the optimum.

Remarks

- A 0.12-approximation algorithm exists.^a
- 0.059-approximation algorithms do not exist unless $NP = ZPP$.^b

^aGoemans & Williamson (1995).

^bHåstad (1997).

Approximability, Unapproximability, and Between

- Some problems have approximation thresholds less than 1.
 - KNAPSACK has a threshold of 0 (p. 749).
 - NODE COVER (p. 708), BIN PACKING, and MAXSAT have a threshold larger than 0.
- The situation is maximally pessimistic for TSP (p. 727) and INDEPENDENT SET,^a which cannot be approximated
 - Their approximation threshold is 1.

^aSee the textbook.

Unapproximability of TSP^a

Theorem 79 *The approximation threshold of TSP is 1 unless $P = NP$.*

- Suppose there is a polynomial-time ϵ -approximation algorithm for TSP for some $\epsilon < 1$.
- We shall construct a polynomial-time algorithm to solve the NP-complete HAMILTONIAN CYCLE.
- Given any graph $G = (V, E)$, construct a TSP with $|V|$ cities with distances

$$d_{ij} = \begin{cases} 1, & \text{if } [i, j] \in E, \\ \frac{|V|}{1-\epsilon}, & \text{otherwise.} \end{cases}$$

^aSahni & Gonzales (1976).

The Proof (concluded)

- Run the alleged approximation algorithm on this TSP.
- Suppose a tour of cost $|V|$ is returned.
 - This tour must be a Hamiltonian cycle.
- Suppose a tour that includes an edge of length $\frac{|V|}{1-\epsilon}$ is returned.
 - The total length of this tour is $> \frac{|V|}{1-\epsilon}$.^a
 - Because the algorithm is ϵ -approximate, the optimum is at least $1 - \epsilon$ times the returned tour's length.
 - The optimum tour has a cost exceeding $|V|$.
 - Hence G has no Hamiltonian cycles.

^aSo this reduction is **gap introducing**.

METRIC TSP

- METRIC TSP is similar to TSP.
- But the distances must satisfy the triangular inequality:

$$d_{ij} \leq d_{ik} + d_{kj}$$

for all i, j, k .

A 0.5-Approximation Algorithm for METRIC TSP^a

- It suffices to present an algorithm with the approximation ratio of

$$\frac{c(M(x))}{\text{OPT}(x)} \leq 2$$

(see p. 703).

^aChoukhmane (1978); Iwainsky, Canuto, Taraszow, & Villa (1986); Kou, Markowsky, & Berman (1981); Plesník (1981).

A 0.5-Approximation Algorithm for METRIC TSP (concluded)

- 1: $T :=$ a minimum spanning tree of G ;
- 2: $T' :=$ double the edges of T ; {Note: T' is an Eulerian *multigraph*.}
- 3: $C :=$ an Euler cycle of T' ;
- 4: Remove repeated nodes of C ; {"Shortcutting."}
- 5: **return** C ;

Analysis

- Let C_{opt} be an optimal TSP tour.
- Note first that

$$c(T) \leq c(C_{\text{opt}}). \quad (20)$$

- C_{opt} is a spanning tree after the removal of one edge.
- But T is a minimum spanning tree.
- Because T' doubles the edges of T ,

$$c(T') = 2c(T).$$

Analysis (concluded)

- Because of the triangular inequality, “shortcutting” does not increase the cost.
 - $(1, 2, 3, 2, 1, 4, \dots) \rightarrow (1, 2, 3, 4, \dots)$, a Hamiltonian cycle.

- Thus



$$c(C) \leq c(T').$$

- Combine all the inequalities to yield

$$c(C) \leq c(T') = 2c(T) \leq 2c(C_{\text{opt}}),$$


as desired.

A 100-Node Example




The cost is 7.72877.

A 100-Node Example (continued)




The minimum spanning tree T .

A 100-Node Example (continued)



“Shortcutting” the repeated nodes on the Euler cycle C .

A 100-Node Example (concluded)



The cost is $10.5718 \leq 2 \times 7.72877 = 15.4576$.

A $(1/3)$ -Approximation Algorithm for METRIC TSP^a

- It suffices to present an algorithm with the approximation ratio of

$$\frac{c(M(x))}{\text{OPT}(x)} \leq \frac{3}{2}$$

(see p. 703).

- This is the best approximation ratio for METRIC TSP as of 2016!

^aChristofides (1976).

A $(1/3)$ -Approximation Algorithm for METRIC TSP (concluded)

- 1: $T :=$ a minimum spanning tree of G ;
- 2: $V' :=$ the set of nodes with an odd degree *in* T ; $\{|V'|$
must be even.}
- 3: $G' :=$ the induced subgraph of G by V' ; $\{G'$ is a
complete graph on V' .}
- 4: $M :=$ a minimum-cost perfect matching of G' ;
- 5: $G'' := T \cup M$; $\{G''$ is an Eulerian *multigraph*.}
- 6: $C :=$ an Euler cycle of G'' ;
- 7: Remove repeated nodes of C ; {"Shortcutting."}
- 8: **return** C ;

Analysis

- Let C_{opt} be an optimal TSP tour.
- By Eq. (20) on p. 732, $c(T) \leq c(C_{\text{opt}})$.
- Let C' be C_{opt} on V' by “shortcutting.”
 - C_{opt} is a Hamiltonian cycle on V .
 - Replace any path (v_1, v_2, \dots, v_k) on C_{opt} with (v_1, v_k) , where $v_1, v_k \in V'$ but $v_2, \dots, v_{k-1} \notin V'$.
- By the triangular inequality,

$$c(C') \leq c(C_{\text{opt}}).$$

- C' is now a Hamiltonian cycle on V' .

Analysis (continued)

- C' consists of two perfect matchings on G' .^a
 - The first, third, ... edges constitute one.
 - The second, fourth, ... edges constitute the other.
- The cheaper perfect matching has cost

$$\frac{c(C')}{2} \leq \frac{c(C_{\text{opt}})}{2}.$$

- As a result, the minimum-cost one M must satisfy

$$c(M) \leq \frac{c(C')}{2} \leq \frac{c(C_{\text{opt}})}{2}.$$

^aNote that G' is a complete graph.

Analysis (concluded)



- Minimum-cost perfect matching can be solved in polynomial time.^a
- Finally, by combining the two earlier inequalities, the Euler cycle C has a cost of

$$\begin{aligned}c(C) &\leq c(T) + c(M) \\ &\leq c(C_{\text{opt}}) + \frac{c(C_{\text{opt}})}{2} \\ &= \frac{3}{2} c(C_{\text{opt}}),\end{aligned}$$

as desired.




^aEdmonds (1965); Micali & V. Vazirani (1980).

A 100-Node Example




The cost is 7.72877.

A 100-Node Example (continued)




A 100-Node Example (continued)




A perfect matching M (not necessarily optimal, however).

A 100-Node Example (continued)






The Euler cycle C of $G'' = T \cup M$.

A 100-Node Example (continued)



“Shortcutting” the repeated nodes on the Euler cycle C .

A 100-Node Example (concluded)



The cost is $8.74583 \leq (3/2) \times 7.72877 = 11.5932$.

KNAPSACK Has an Approximation Threshold of Zero^a

Theorem 80 *For any ϵ , there is a polynomial-time ϵ -approximation algorithm for KNAPSACK.*

- We have n weights $w_1, w_2, \dots, w_n \in \mathbb{Z}^+$, a weight limit W , and n values $v_1, v_2, \dots, v_n \in \mathbb{Z}^+$.^b
- We must find an $I \subseteq \{1, 2, \dots, n\}$ such that $\sum_{i \in I} w_i \leq W$ and $\sum_{i \in I} v_i$ is the largest possible.

^aIbarra & Kim (1975).

^bIf the values are fractional, the result is slightly messier, but the main conclusion remains correct. Contributed by Mr. Jr-Ben Tian (B89902011, R93922045) on December 29, 2004.

The Proof (continued)

- Let

$$V = \max\{v_1, v_2, \dots, v_n\}.$$

- Clearly, $\sum_{i \in I} v_i \leq nV$.
- Let $0 \leq i \leq n$ and $0 \leq v \leq nV$.
- $W(i, v)$ is the minimum weight attainable by selecting only from the first i items and with a total value of v .
 - It is an $(n + 1) \times (nV + 1)$ table.

The Proof (continued)

- Set $W(0, v) = \infty$ for $v \in \{1, 2, \dots, nV\}$ and $W(i, 0) = 0$ for $i = 0, 1, \dots, n$.^a
- Then, for $0 \leq i < n$ and $1 \leq v \leq nV$,^b

$$W(i+1, v) = \begin{cases} \min\{W(i, v), W(i, v - v_{i+1}) + w_{i+1}\}, & \text{if } v \geq v_{i+1}, \\ W(i, v), & \text{otherwise.} \end{cases}$$

- Finally, pick the largest v such that $W(n, v) \leq W$.^c

^aContributed by Mr. Ren-Shuo Liu (D98922016) and Mr. Yen-Wei Wu (D98922013) on December 28, 2009.

^bThe textbook's formula has an error.

^cLawler (1979).

0 v nV

	$\leq W$	

The Proof (continued)

With 6 items, weights $(3, 3, 1, 3, 2, 1)$, values $(4, 3, 3, 3, 2, 3)$, and $W = 12$, maximum total value 16 is achieved with $I = \{1, 2, 3, 4, 6\}$ and total weight 11.

0	∞	∞	∞	∞	∞	∞	∞	∞	∞	∞	∞	∞	∞	∞	∞	∞	∞	∞
0	∞	∞	∞	3	∞	∞	∞	∞	∞	∞	∞	∞	∞	∞	∞	∞	∞	∞
0	∞	∞	3	3	∞	∞	6	∞	∞	∞	∞	∞	∞	∞	∞	∞	∞	∞
0	∞	∞	1	3	∞	4	4	∞	∞	7	∞	∞	∞	∞	∞	∞	∞	∞
0	∞	∞	1	3	∞	4	4	∞	7	7	∞	∞	10	∞	∞	∞	∞	∞
0	∞	2	1	3	3	4	4	6	6	7	9	9	10	∞	12	∞	∞	∞
0	∞	2	1	3	3	2	4	4	5	5	7	7	8	10	10	11	∞	13

The Proof (continued)

- The running time $O(n^2V)$ is not polynomial.
- Call the problem instance

$$x = (w_1, \dots, w_n, W, v_1, \dots, v_n).$$

- Additional idea: Limit the number of precision bits.
- Define

$$v'_i = \left\lfloor \frac{v_i}{2^b} \right\rfloor.$$

- Note that

$$v_i \geq 2^b v'_i > v_i - 2^b.$$

The Proof (continued)

- Call the approximate instance

$$x' = (w_1, \dots, w_n, W, v'_1, \dots, v'_n).$$

- Solving x' takes time $O(n^2V/2^b)$.
 - Use $v'_i = \lfloor v_i/2^b \rfloor$ and $V' = \max(v'_1, v'_2, \dots, v'_n)$ in the dynamic programming.
 - It is now an $(n + 1) \times (nV + 1)/2^b$ table.
- The selection I' is optimal for x' .
- But I' may not be optimal for x , although it still satisfies the weight budget W .

The Proof (continued)

With the same parameters as p. 753 and $b = 1$: Values are now $(2, 1, 1, 1, 1, 1)$ and a smaller total maximum value $4 + 3 + 3 + 2 + 3 = 15$ is achieved with $I' = \{1, 2, 3, 5, 6\}$ and total weight 10.^a

0	∞						
0	∞	3	∞	∞	∞	∞	∞
0	3	3	6	∞	∞	∞	∞
0	1	3	4	7	∞	∞	∞
0	1	3	4	7	10	∞	∞
0	1	3	4	6	9	12	∞
0	1	2	4	5	7	10	13

^aThe original optimal $I = \{1, 2, 3, 4, 6\}$ has value 6 and weight 11 for x' , whereas I' has the same total value 6 but smaller total weight 10.

The Proof (continued)

- The value of I' for x is close to that of the optimal I :

$$\begin{aligned}\sum_{i \in I'} v_i &\geq \sum_{i \in I'} 2^b v'_i = 2^b \sum_{i \in I'} v'_i \\ &\geq 2^b \sum_{i \in I} v'_i = \sum_{i \in I} 2^b v'_i \\ &\geq \sum_{i \in I} (v_i - 2^b) \\ &\geq \left(\sum_{i \in I} v_i \right) - n2^b.\end{aligned}$$

The Proof (continued)

- In summary,

$$\sum_{i \in I'} v_i \geq \left(\sum_{i \in I} v_i \right) - n2^b.$$

- Without loss of generality, assume $w_i \leq W$ for all i .
 - Otherwise, item i is redundant.
- V is a lower bound on OPT.
 - Picking an item with value V is a legitimate choice.

The Proof (concluded)

- The relative error from the optimum is:

$$\frac{\sum_{i \in I} v_i - \sum_{i \in I'} v_i}{\sum_{i \in I} v_i} \leq \frac{\sum_{i \in I} v_i - \sum_{i \in I'} v_i}{V} \leq \frac{n2^b}{V}.$$

- Suppose we pick $b = \lfloor \log_2 \frac{\epsilon V}{n} \rfloor$.
- The algorithm becomes ϵ -approximate.^a
- The running time is then $O(n^2 V / 2^b) = O(n^3 / \epsilon)$, a polynomial in n and $1/\epsilon$.^b

^aSee Eq. (16) on p. 697.

^bIt hence depends on the *value* of $1/\epsilon$. Thanks to a lively class discussion on December 20, 2006. If we fix ϵ and let the problem size increase, then the complexity is cubic. Contributed by Mr. Ren-Shan Luoh (D97922014) on December 23, 2008.

Comments

- INDEPENDENT SET and NODE COVER are reducible to each other (Corollary 41, p. 360).
- NODE COVER has an approximation threshold at most 0.5 (p. 710).
- But INDEPENDENT SET is unapproximable (see the textbook).
- INDEPENDENT SET limited to graphs with degree $\leq k$ is called k -DEGREE INDEPENDENT SET.
- k -DEGREE INDEPENDENT SET is approximable (see the textbook).

Finis