

Combinatorial optimization and graphical models

Submodular functions

Optimization - 10725

Carlos Guestrin

Carnegie Mellon University

April 28th, 2008

©2008 Carlos Guestrin

1

Most probable explanation (MPE) in a Markov network

$$\psi_i(x_i) = e^{f_i(x_i)}$$

$$\psi_{ij}(x_i, x_j) = e^{f_{ij}(x_i, x_j)}$$

- Markov net:



$$P(x) = \frac{1}{Z} \prod_i \psi_i(x_i) \prod_{i,j} \psi_{ij}(x_i, x_j)$$

$$= \frac{1}{Z} e^{\sum_i f_i(x_i) + \sum_{i,j} f_{ij}(x_i, x_j)}$$

$$Z = \sum_x e^{\sum_i f_i(x_i) + \sum_{i,j} f_{ij}(x_i, x_j)}$$

- Most probable explanation:

$$\arg \max_x P(x) = \arg \max_x \log P(x) = \arg \max_x \left(\sum_i f_i(x_i) + \sum_{i,j} f_{ij}(x_i, x_j) \right) - \log Z$$

- In general, NP-complete problem, and hard to approximate

©2008 Carlos Guestrin

2

MPE for attractive MNs – 2 classes

■ Attractive MN: $\alpha_{i0} \geq 0, \alpha_{i1} \geq 0$
 □ E.g., image classification

$f_i(x_i) \leftarrow \text{anything}$ (local measurement)
 $f_{ij}(x_i, x_j) = \begin{cases} \lambda_{ij}, & x_i = x_j \\ 0, & \text{otherwise} \end{cases}$ (smoothness "constraint")

Finding most probable explanation: $\max_x \sum_i f_i(x_i) + \sum_{ij} f_{ij}(x_i, x_j)$

Can be solved by min cut


Goal: little edge cut as possible

OPT $x_1=1, x_2=1, x_3=0, x_4=0$
 Val OPT: $\lambda_{12} + \lambda_{34} + \alpha_{10} + \alpha_{20} + \alpha_{31} + \alpha_{41}$

Variables: $x_1, \lambda_{12}, \alpha_{10}, \alpha_{20}, \alpha_{31}, \alpha_{41}$

MPE, Attractive MNs, k classes

- MPE for k classes:



- Multiway cut:

- Graph G, edge weights w_{ij}
- Finding minimum cut, separate s_1, \dots, s_k



- Multiway cut problem is NP-complete

Multiway cut – combinatorial algorithm

Very simple alg:

- For each $i=1 \dots k$
 - Find cut C_i that separates s_i from rest
- Discard $\arg \max_i w(C_i)$, return union of rest

$$C_1 \leftarrow w(C_1) = 5$$

$$C_2 \leftarrow w(C_2) = 8$$



$$C_3 \leftarrow w(C_3) = 6$$

return $C_1 \cup C_3$

Algorithm achieves $2-2/k$ approximation

- OPT cut A^* separates graph into k components
 - No advantage in more than k
- From A^* form A_1^*, \dots, A_k^* , where A_i^* separates s_i from rest

- Each edge in A^* appears in $\sum_i A_i^*$ exactly 2 times
 - Thus $2w(A^*) = \sum_i w(A_i^*)$



$$w(C_i) \leq w(A_i^*)$$

Multiway cut proof

- Thus, for OPT cut A^* we have that:

$$2w(A^*) = \sum_i w(A_i^*)$$

- Each A_i^* separates s_i from rest, thus

$$w(C_i) \leq w(A_i^*)$$

but C_i is OPT cut of s_i from rest

- But, can do better, because


$$\text{throw away } \max_i w(C_i) \geq \frac{1}{k} \sum_i w(C_i)$$

$$\sum_i w(C_i) - \max_i w(C_i) \leq \left(1 - \frac{1}{k}\right) \sum_i w(C_i)$$

$$\leq \left(1 - \frac{1}{k}\right) \sum_i w(A_i^*) = 2\left(1 - \frac{1}{k}\right) w(A^*)$$

General solutions to combinatorial problems

- Nonlinear optimization
 - Extremely general
 - NP-Hard to solve
- Covex problems:
 - Special problem structure
 - General efficient solution algorithms
 - Applies to many continuous problems
- Combinatorial optimization
 - Integer programming
 - Very general, but NP-hard
 - Convex relaxations
 - Approximation guarantees for many problems
 - Usually solution is problem specific
 - Though there are general principles
- Is there general problem structure in combinatorial problems?
 - Analogously to convexity?
 - Recognize structure and use general algorithms?



$$\theta f(x) + (1-\theta)f(y) \geq f(\theta x + (1-\theta)y)$$

e.g., totally unimodular matrices

©2008 Carlos Guestrin


7

Water distribution networks

- Water distribution in a city → very complex system
- Pathogens in water can affect thousands (or millions) of people
- Currently: Add chlorine to the source and hope for the best

ATTACK!
could deliberately introduce pathogen

Chlorine



Simulator from EPA


©2008 Carlos Guestrin

8

Monitoring water networks

[Krause, Leskovec, G., Faloutsos, VanBriesen '08]

- Contamination of drinking water could affect millions of people



$$A \subseteq V$$



Hach Sensor
~\$14K

- Simulator from EPA
- Place sensors to detect contaminations
- “Battle of the Water Sensor Networks” competition

Where should we place sensors to detect contaminations quickly?


©2008 Carlos Guestrin

9

Model-based sensing


- Model predicts impact of contaminations
 - For water networks: Water flow simulator from EPA
 - For lake monitoring: Learn probabilistic models from data (later)
- For each subset $A \subseteq V$ compute “sensing quality” $F(A)$

Model predicts
High impact
Contamination



High sensing quality $F(A) = 0.9$

©2008 Carlos Guestrin



Low sensing quality $F(A)=0.01$


10

Sensor placement

- Given: finite set V of locations, sensing quality F
- Want: $A^* \subseteq V$ such that

$$A^* = \underset{|\mathcal{A}| \leq k}{\operatorname{argmax}} F(\mathcal{A})$$

Typically NP-hard!



Greedy algorithm:

Start with $A = \emptyset$

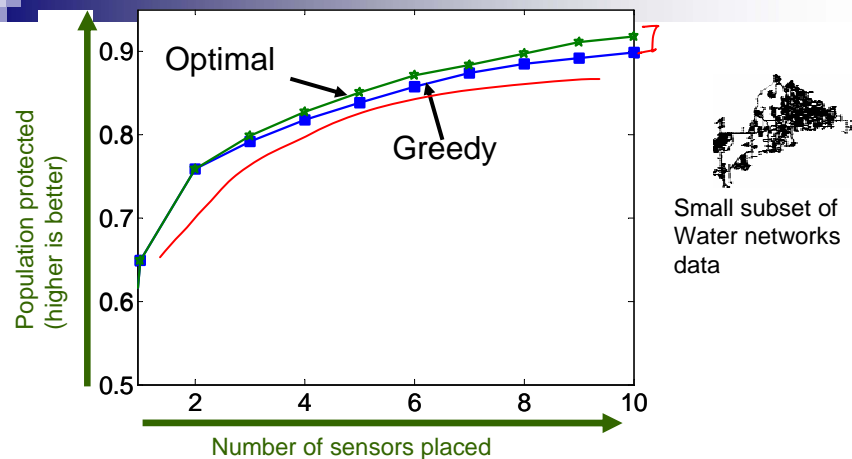
For $i = 1$ to k

$s^* := \underset{s}{\operatorname{argmax}} F(A \cup \{s\})$

$A := A \cup \{s^*\}$

How well can this simple heuristic do?

Performance of greedy algorithm




Greedy score empirically close to optimal. Why?

Key property: Diminishing returns

submodularity


Placement A = $\{S_1, S_2\}$

Placement B = $\{S_1, S_2, S_3, S_4\}$



Adding S' will help a lot!

New sensor S'



Adding S' doesn't help much

Theorem [Krause, Leskovec, G., Faloutsos, VanBriesen '08]:

Sensing quality $F(A)$ in water networks is submodular!

©2008 Carlos Guestrin

13

Submodularity


supermodular $F(A \cup S) - F(A) \leq F(B \cup S) - F(B)$

modular $F(A \cup \{x\}) - F(A) = F(B \cup \{x\}) - F(B)$
 $(\Rightarrow) F(A) = \sum_{A_i \in A} F(A_i)$

- Formalizes notion of diminishing returns

$$\forall A, B \quad F(A \cup \{x\}) - F(A) \geq F(B \cup \{x\}) - F(B)$$

$$A \subseteq B$$




- Equivalent definition:

$$\forall A, B \quad F(A) + F(B) \geq F(A \cup B) + F(A \cap B)$$

©2008 Carlos Guestrin

14

What do we get from submodularity

max $f(x)$
 x
 e.g. $x \in$ 
 Hand 

- Submodularity is a general property of set functions
- Submodular function can be minimized in polynomial time!

$$\min_{A \subseteq V} F(A)$$

F is submodular
 polytime


- But our problem is water problem

$$\max_{\substack{A \\ |A| \leq k}} F(A)$$

F is submodular

Another example...

- Maximum cover
 - Ground elements $v \in V$
 - Set of sets $S_i \subseteq V$
 - Pick k sets, maximize number of covered elements



$$F(A) = \sum_v I(v \text{ is } \bigcup_{i \in A} S_i)$$

A is a set of sets

$$A = \{S_1, S_2, \dots, S_k\}$$

$$F(A \cup S_i) - F(A) \geq F(B \cup S_i) - F(B) \quad A \subseteq B$$

consider element v

\hookrightarrow $\begin{cases} \text{covered } A \\ \text{covered } S_i, \text{ but not } B \text{ or } A \\ \text{covered } B, \text{ but not } A \end{cases}$

Maximizing submodular functions – cardinality constraint

■ Given

□ Submodular function $F(A)$

□ Normalized $F(\emptyset) = 0$

□ Non-decreasing $F(A) \leq F(B) \quad A \subseteq B$

e.g.)
max cover

■ Greedy algorithm guarantees

$$F(A_{\text{greedy}}) \geq \underbrace{(1 - \frac{1}{e})}_{63\%} F(A_{\text{opt}})$$

$\uparrow \max_{A: |A| \leq k} F(A)$

■ Can you get better algorithm?

NO, unless $P=NP$, if you could solve max-cover better than $(1 - \frac{1}{e})$ then $P=NP$

Online bounds

■ Submodularity provides bounds on the quality of the solution A obtained by any algorithm

□ For normalized, non-decreasing functions

■ Advantage of adding elements to A:

■ Bound on the quality of any set A:

■ Tighter bound:

Battle of the Water Sensor Networks Competition

- Real metropolitan area network (12,527 nodes)
- Water flow simulator provided by EPA
- 3.6 million contamination events
- Multiple objectives: Detection time, affected population, ...
- Place sensors that detect well “on average”



©2008 Carlos Guestrin

19

BWSN Competition results


- 13 participants
- Performance measured in 30 different criteria

G: Genetic algorithm

D: Domain knowledge

H: Other heuristic

E: “Exact” method (MIP)




©2008 Carlos Guestrin

20

What was the trick?

Simulated alB.6M contaminations on 2 weeks / 40 processors
 152 GB data on disk / 16 GB in main memory (compressed)

→ Very accurate sensing quality ☺ Very slow evaluation of $F(A)$ ☹



30 hours/20 sensors

6 weeks for all

30 settings ☹



submodularity
to the rescue:

Using "lazy evaluations":

1 hour/20 sensors

Done after 2 days! ☺

©2008 Carlos Guestrin

21

Lazy evaluations

- Naïve implementation of greedy:
 - Advantage of an element never increases:
 - Advantage:
 - What if you already picked a larger set:
 - Set after picking i elements: A_i
- Lazy evaluations:
 - Keep a priority queue over elements:
 - Initialize with advantage of each element
 - Pick element on top, recompute priority
 - If element remains on top

©2008 Carlos Guestrin

22

Other maximization settings

- Non-monotone submodular functions
- Non-unit costs
- Complex constraints
 - Paths
 - Spanning trees
- Worst-case optimization