Graph cuts

Reminder:

Markov Random Fields (MRF)

1D: HMM (Hidden Markov Models)
Markov chain

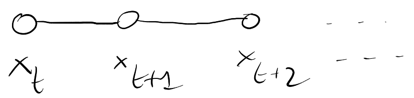

minimize a criterion:

$$\sum_t V_t(x_t) + \sum_t D_t(x_t, x_{t+1})$$

→ solvable with dynamic programming

2D

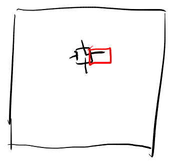

ex: image = pixel grid + neighborhood → image segmentation in 2 classes

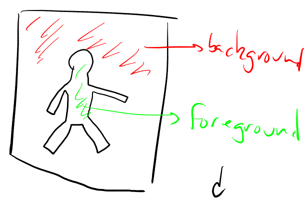

For each pixel, choose one of the 2 possible labels

local prior
For each pixel p
(based on its color, the texture around ... local descriptor)

criterion to optimize:

$$\underbrace{\sum_p V_p(x_p) + \sum_{p\sim q} D_{pq}(x_p, x_q)}_{\text{MRF}}$$

label chosen for pixel p (binary)

neighboring pixels

↳ general case: NP-hard

I Max Flow

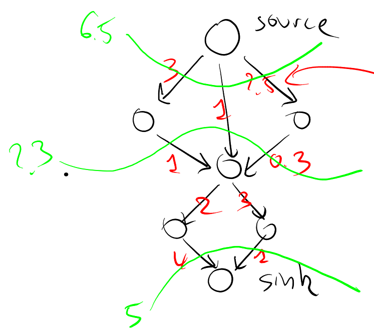

each edge (≡ pipe), maximum capacity (unit: L/s)

→ maximum flow that can pour from "source" to "sink"?

First Algorithm: Ford-Fulkerson

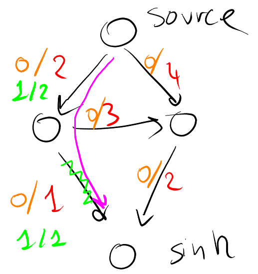

Flow (current)

Maximum capacity

→ path through edges not full yet

Residual graph

↳ with only non-saturated edges

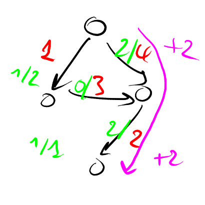

→

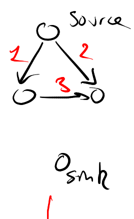

Algo: iteratively:
- Find a path connecting source & sink in the residual graph

– saturate this path
(i.e. add flow along this path
until not possible to add more)

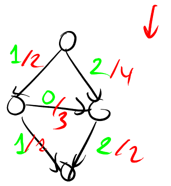

Problem:

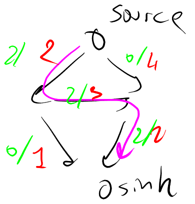

→ Residual

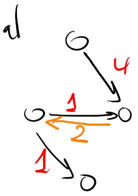

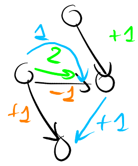

Corrected algorithm: residual graph: { edges not saturated,
and reverse of edges with non-0 flow }

residual capacity = flow in the true direction

not necessary for directed edges from the source or the sink

Iteratively,
– find a path connecting the source & the sink
– saturate it

Issue: complexity (in worst case)

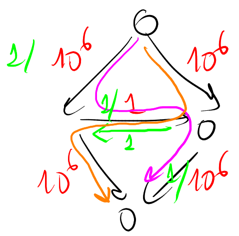

Issue: un-saturate an already-saturated edge
many many times
if not looking at "optimal paths"

Variations: { Dinitz (Dinic)
Edmonds – Karp

Idea: – don't pick "a" path
– pick the shortest path

⇒ better worst-case complexity
$O(\#\text{nodes}^3)$

Many other algorithms
→ Push-Relabel
→ in practice: Boykov-Kolmogorov: heuristics for 2D image segmentation problem

II Link between maximum flow & minimum cut

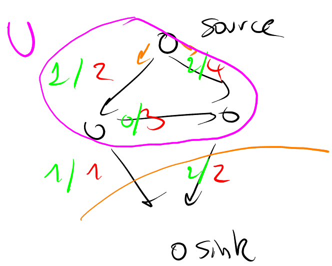

Residual graph at the end →

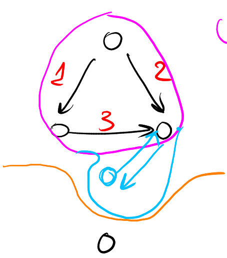

U: all nodes reachable from the source in the residual graph

maximum flow (reading from source outputs): 3

1) All edges going out of U are saturated

2) Flows on edges going to U are 0

→ partition of the graph; U and ¬U
↳ cut of the graph

Total flow
maximum flow $= \sum_{\text{edges out of } U}$ capacity(edge)
= what pours from the source

because these edges are saturated

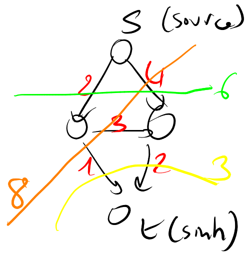

Theorem :

Best s-t cut is given by the maximum flow solution

↳ cost of the best cut = maximum flow

cost of a cut = $\sum_{\text{cut edges}}$ capacity (of these edges) ← s-t min-cut : Find a cut in the graph that separates s from t such that the cost of the cut is minimized

III Image segmentation as a s-t min cut problem

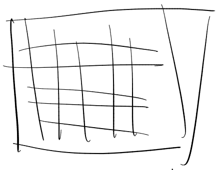
image grid

- local preference for each pixel p : $V_p(x_p)$ (labels : 0 / 1) $\sum_p$
- interaction between neighbors (for spatial coherency) : $D_{pq}(x_p, x_q)$ $\sum_{pq}$

each pixel is represented by a node

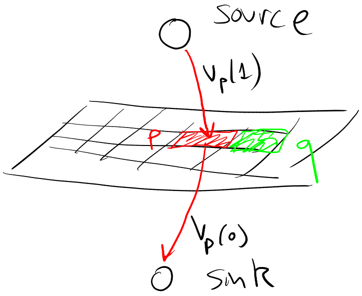

zoom

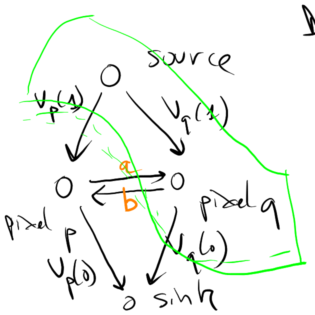

D_{pq}

$x_p \backslash x_q$	0	1
0	0	a
1	b	0

$0+0 \le a+b$

=> interaction between pixels ⟷ these values a, b

↳ positivity of all capacities : required by MaxFlow

=> $a, b \ge 0$

Conditions

$V_p : \begin{pmatrix} V_p(0) \\ V_p(1) \end{pmatrix} : \begin{pmatrix} -3 \\ -2 \end{pmatrix}$ +10 +10 only important thing : the difference

same solution

=> Free to choose any values (in $\mathbb{R}$) for individual potentials

=> sub-modularity :

$M :$

M_{11}	M_{12}
M_{21}	M_{22}

+10 +10 +10 +10

⇓ I can suppose that all coefficient of D_{pq} are ≥ 0

M is submodular ⟺ $M_{11} + M_{22} \le M_{21} + M_{12}$

it costs more to be off-diagonal ⟺ to choose different labels for neighboring pixel

- naive construction => interaction terms are submodular.
- Reverse ?

H0 : submodular matrix

M_{11}	M_{12}
M_{21}	M_{22}

$M_{11}+M_{22} \le M_{21}+M_{12}$

→ $-M_{11}$ on each element

0	$M_{12}-M_{11}$
$M_{21}-M_{11}$	$M_{22}-M_{11}$

interaction matrix

$x_p \backslash x_q$	0	1
0	$y+t$	$z+a+y$
1	$x+b+t$	$x+z$

V.S.

total cost of the cut

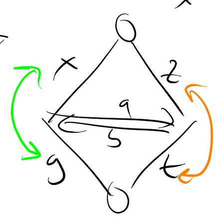

⇓ $-y-t$

0	$\boxed{z-t}$ $+a$
$\boxed{x-y}$ $+b$	$\boxed{x-y}$ $+\boxed{z-t}$

binary variable

=> Any MRF with submodular interaction matrices (2x2) can be solved exactly by a graph cut

while : if not submodular : NP-hard in the general case

$\prod_p \sum_{q \in N(p)} D_{pq}(x_p, x_q)$

↳ as many interaction matrices as edges

$\boxed{p|q|r}$ ↳ $E(p,q,r)$ → "high order" terms

IV Multi-label energies

- proposed labelling: x (x_p)
- criterion to optimize:

$$\sum_p V_p(x_p) + \sum_{pq} D_{pq}(x_p, x_q)$$

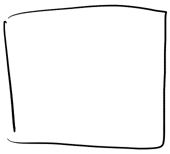

Image segmentation task

3 classes (or more ≥ 2) $\Rightarrow$ can't apply the technique above

1) α-expansion:

Sequence of binary problems

Iteratively:
- select one class c $(=\alpha)$
- binary problem:
  given the current solution:
  - For each pixel: either keep the current choice
    or move to that class c

$\Rightarrow$ yields an exact solution to each sub-problem (binary ones)

$\Rightarrow$ local solution to the original problem but with quality guarantees

→ The binary problem obtained needs to be submodular

$\Rightarrow$ interaction matrix between 2 pixels

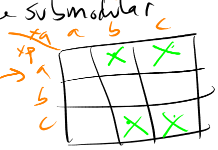

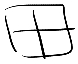

2x2 sub-matrix

currently: $x_p = a$, $x_q = b$

all 4 corners of rectangular submatrices → 2x2  are submodular

↳ supplementary H0: $D_{pq}(\ell,\ell) = 0 \quad \forall \ell$

$$\begin{array}{|c|c|c|} \hline 0 & & \\ \hline & 0 & \\ \hline & & 0 \\ \hline \end{array}$$

$\Rightarrow D_{pq}$ is a distance between labels
- $D_{pq}(a,b) \geq 0 \quad \forall a,b$
- $D_{pq}(a,b) + D_{pq}(b,c) \geq D_{pq}(a,c)$

2) α-β swap

Similar idea: sequence of binary problems

Iteratively:
- pick 2 classes α, β
- binary problem: allow pixels with one of these 2 labels to move to the other one

Conditions: binary pb has to be submodular

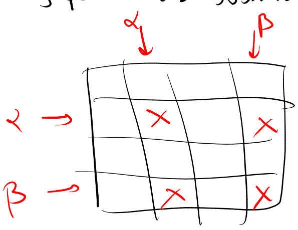

2x2 submatrix
→ 2 corners on the diagonal

$\Rightarrow$ constraints: weaker than for α-expansion

Example of interaction matrix: Potts model

$$D_{pq}(a,b) = \delta_{a=b} = \begin{cases} 0 & \text{if } a=b \\ 1 & \text{otherwise} \end{cases}$$

$\Rightarrow$ distance between labels

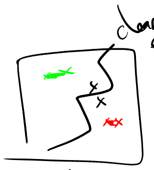

wish: $\neq$ labels costs less } not possible with submodular interactions!

How to set the problem

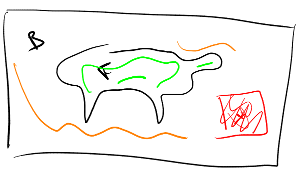

Image of a cow on grass

→ statistics on pixel colors

↳ histogram of background pixels in a set of segmented images

of Foreground pixels

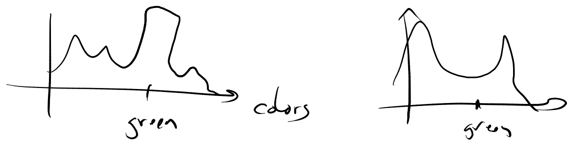

individual potentials V_p

interactive segmentation by adding more clues

→ interaction between pixels

→ estimate the probability of a cut there

cost of cutting that edge : $e^{-\gamma \| color(p) - color(q) \|^2}$

Infinite links

↳ Ishikawa's construction

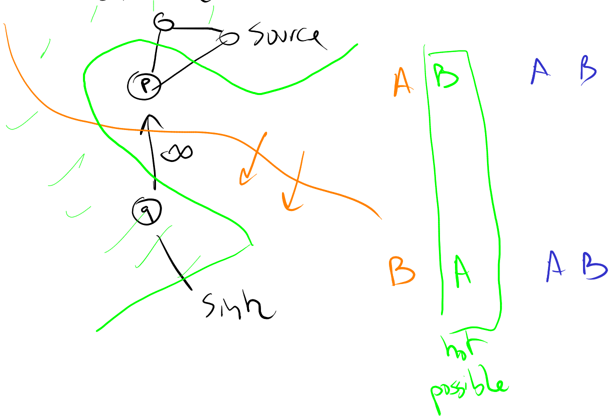

$$
\begin{array}{c|cc}
D_{pq} & A & B \\
\hline
A & & \\
B & \infty &
\end{array}
$$

Ex of icebergs sequence segmentation

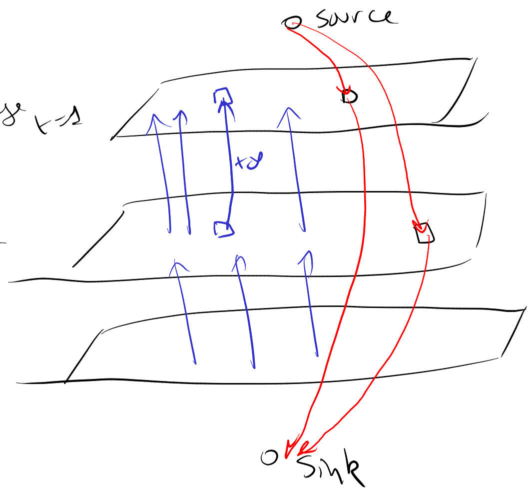

≡ constraint that Foreground can only shrink with time

$$
D_{t,t+1} : \begin{array}{c|cc}
{}_{t+1} \backslash {}^{t} & F & B \\
\hline
F & & \\
B & +\infty &
\end{array}
$$

Official worst-case complexity : cubic

↳ using heuristics by B&K → fast

↳ O(#frames)

⇒ globally-optimal solution (1 only cut)

Infinite links between any pixel p at time t and the same location at time t-1

NB: needs pixel-wise registration of all images

Application:

image → 5 levels of grey {0, 50, 100, 200, 250}

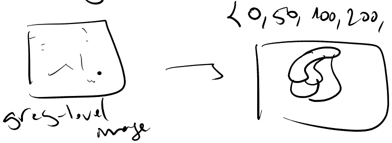

~ looks like a cartoon

Renaissance