Approximate Energy Minimization via Graph Cuts

Filip Malmberg

Recall: Local search optimization

- Local search algorithms are a class of combinatorial optimization algorithms.
- Starting from a feasible solution, they iteratively try to obtain a better solution by searching the "neighborhood" of the current solution.
- A critical issue is the manner in which the neighborhood is defined. Larger neighborhoods generally improve the quality of the solution, but increase the computation time.

Very large-scale neighborhood search

- We now turn our attention to a special class of local search methods, where
  1. the size of the neighborhood is "very large" with respect to the size of the input data.
  2. the neighborhood can be searched in an efficient manner.
- A number of such methods have been proposed in the literature [1].
- In this lecture, we will consider an approach using minimal graph cuts. [2].

Move-making algorithms

- Local search methods are also often referred to as *move-making* algorithms.
- The solutions adjacent to a solution S are those that can be obtained from S in a single *move*.

Standard moves

- A *standard move* consists of changing the label of a single vertex. (Algorithm known as *iterated conditional modes*, *ICM* – this is what we did in the first lecture!)
- A local minimum with respect to standard moves thus means that we cannot decrease the energy by changing the lable of a single vertex. This is a rather weak optimality condition
- In order to get to a *global* optimum (or even a "good" local optimum) from a particular starting point, we must be able to reach the desired optimum by changing the label of one vertex at a time *and each such operation must improve the solution*.

A "canonical" labeling problem

- We seek a label assignment configuration x that minimizes a given objective function E, written as follows:

$$E(\mathbf{x}) = \sum_{i \in \mathcal{V}} \phi_i(x_i) + \sum_{i,j \in \mathcal{E}} \phi_{ij}(x_i, x_j) \; . \tag{1}$$

  where x_i denotes the label of vertex $v \in \mathcal{V}$ which must belong to a finite set of integers $\{0, 1 \ldots, K-1\}$
- Bad news: NP-hard in the general case.
- Good news: For $K = 2$ labels, a large class of these problems can be solved by minimal graph cuts.

Solving binary labeling problems with graph cuts

- A minimal graph cut divides the graph so the every node is connected either to the source or the sink – a binary labeling!
- For the binary labeling problem, with $K = 2$, a globally optimal solution can be computed by solving a max-flow/min-cut problem on a suitably constructed graph, provided that all pairwise terms are *submodular* [2, 6].
- A pairwise term ϕ_{ij} is said to be submodular if

$$\phi_{ij}(0,0) + \phi_{ij}(1,1) \leq \phi_{ij}(0,1) + \phi_{ij}(1,0) \; . \tag{2}$$

Why submodular?

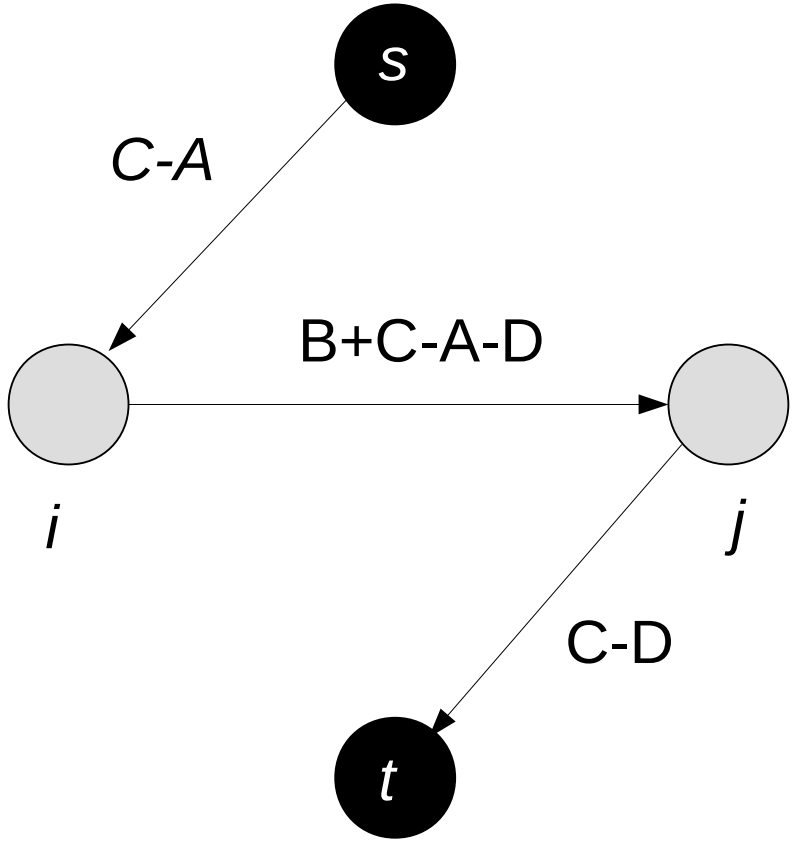

Figure 1: Example of graph construction. Submodularity ensures positive edge weights.

Graph cuts with more than two terminals?

- Unfortunately, computing globally minimal graph cuts for more than two terminals is NP-hard. Therefore, we can not directly apply the graph cut approach to solve optimization with more than two labels.

Approximate energy minimization with graph cuts

- In a paper from 2001 [2], Boykov et al. proposed two types of large moves based on minimal graph cuts:
  - α-β-swap moves.
  - α-expansion moves.
- In contrast to standard moves, both these moves allow a large number of vertices to change their labels simultaneously.
- This paper has more than 9000 citations (600 since 2020) according to google scholar!

α-β-swap moves.

- Given a pair of labels α, β, a move from a labeling L_1 to a labeling L_2 is called an α-β-swap move if the only difference between L_1 and L_2 is that
  - some vertices that were labeled α in L_1 are labeled β in L_2, and
  - some vertices that were labeled β in L_1 are labeled α in L_2.

α-β-swap moves.

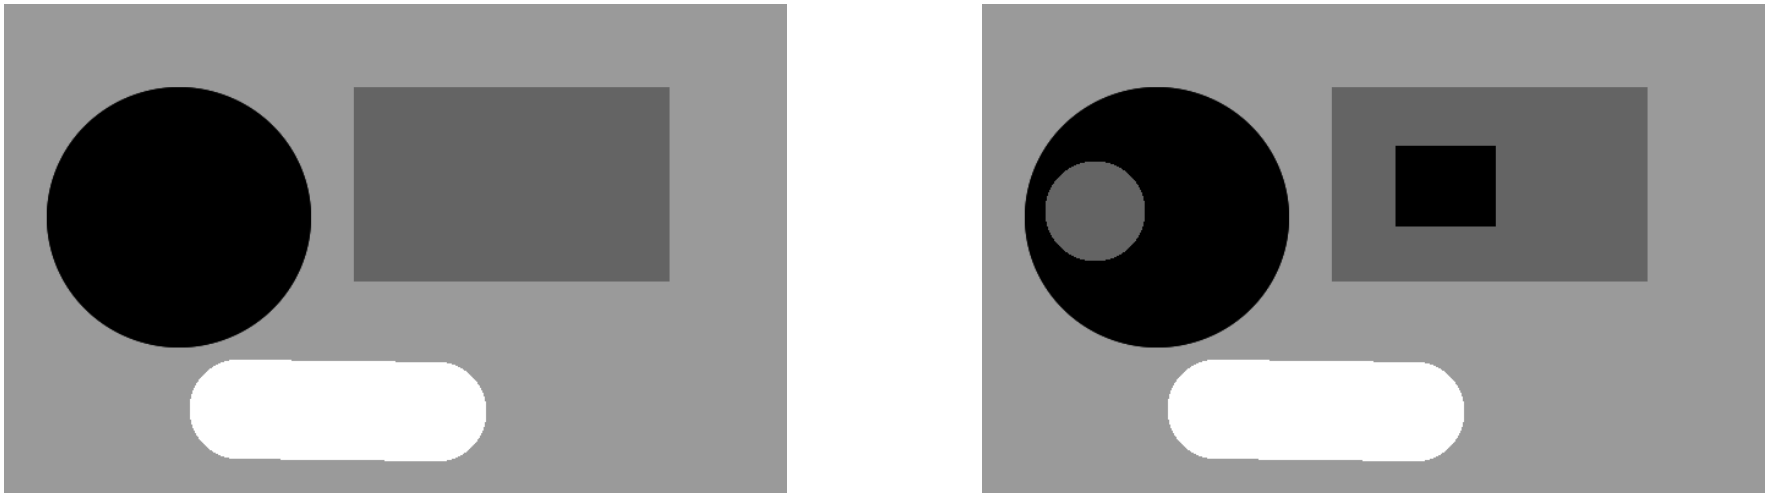

Figure 2: Example of an α-β-swap move. The labeling on the right is a "dark gray"-"black"-swap move from the labeling on the left.

α-expansion moves.

- Given a label α, a move from a labeling L_1 to a labeling L_2 is called an α-expansion move if the only difference between L_1 and L_2 is that some vertices that were not labeled α in L_1 are labeled α in L_2.

α-expansion moves.

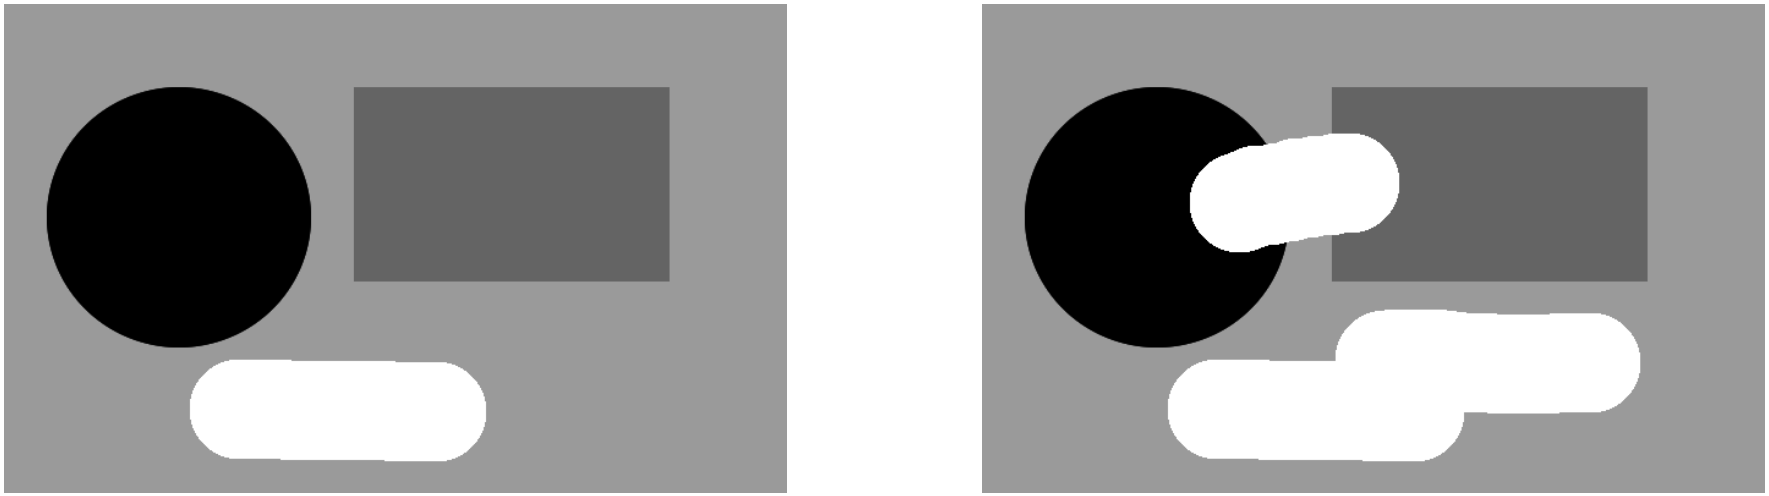

Figure 3: Example of an α-β-swap move. The labeling on the right is a "white"-expansion move from the labeling on the left.

Optimization algorithm

- Given a labeling L, there is an exponential number of possible swap and expansion moves.
- Checking these moves naively requires exponential time if perfomed naively.
- Instead Boykov et al. propose efficient methods to find the optimal α-β-swap or α-expansion move given a current labeling, using minimal graph cuts.
- The two algorithms are similar in structure.

Swap move algorithm

Start with an arbitrary labeling L.
Set done $\leftarrow$ false.
while *not done* **do**
  Set done $\leftarrow$ true.
  for *each pair of labels* α *and* β **do**
    Find, among all labelings within one α-β swap from L, the labeling L' with the lowest energy.
    if $f(L') < f(L)$ **then**
      Set $L \leftarrow L'$.
      Set done $\leftarrow$ false.
    end
  end
end

Expansion move algorithm

Start with an arbitrary labeling L.
Set done $\leftarrow$ false.
while *not done* **do**
  Set done $\leftarrow$ true.
  for *each label* α **do**
    Find, among all labelings within one α expansion from L, the labeling L' with the lowest energy.
    if $f(L') < f(L)$ **then**
      Set $L \leftarrow L'$.
      Set done $\leftarrow$ false.
    end
  end
end

Comparison between the algorithms

As shown by Boykov et al. [2] both algorithms can be used to find strong local minima of a fairly general class of objective functions:

- The swap-move algorithm can optimize any objective function that is a *semi-metric*.
- The expansion move-algorithm can optimize any objective function that is a *metric*.
- When applicable, the expansion move algorithm is guaranteed to produce results that are within a known factor of the global minimum.

Example results

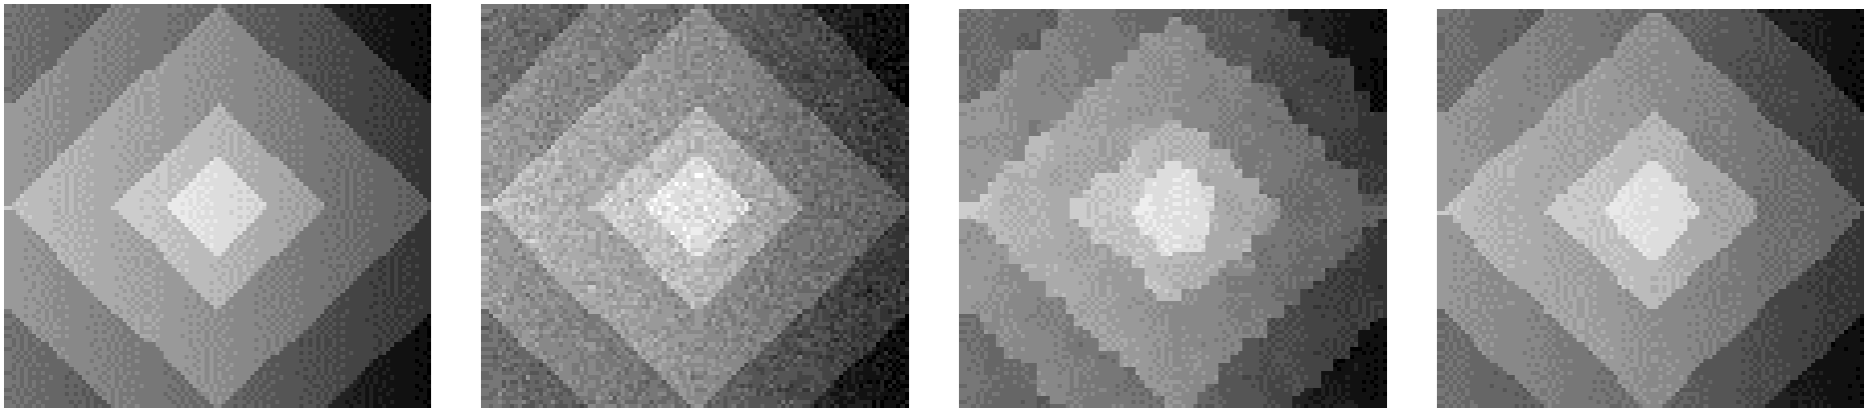

Figure 4: Restoration example by Boykov et al. (a) Original image. (b) Original image corrupted by noise. (c) Local minimum with respect to standard moves. (d) Local minimum with respect to expansion moves.

Example results

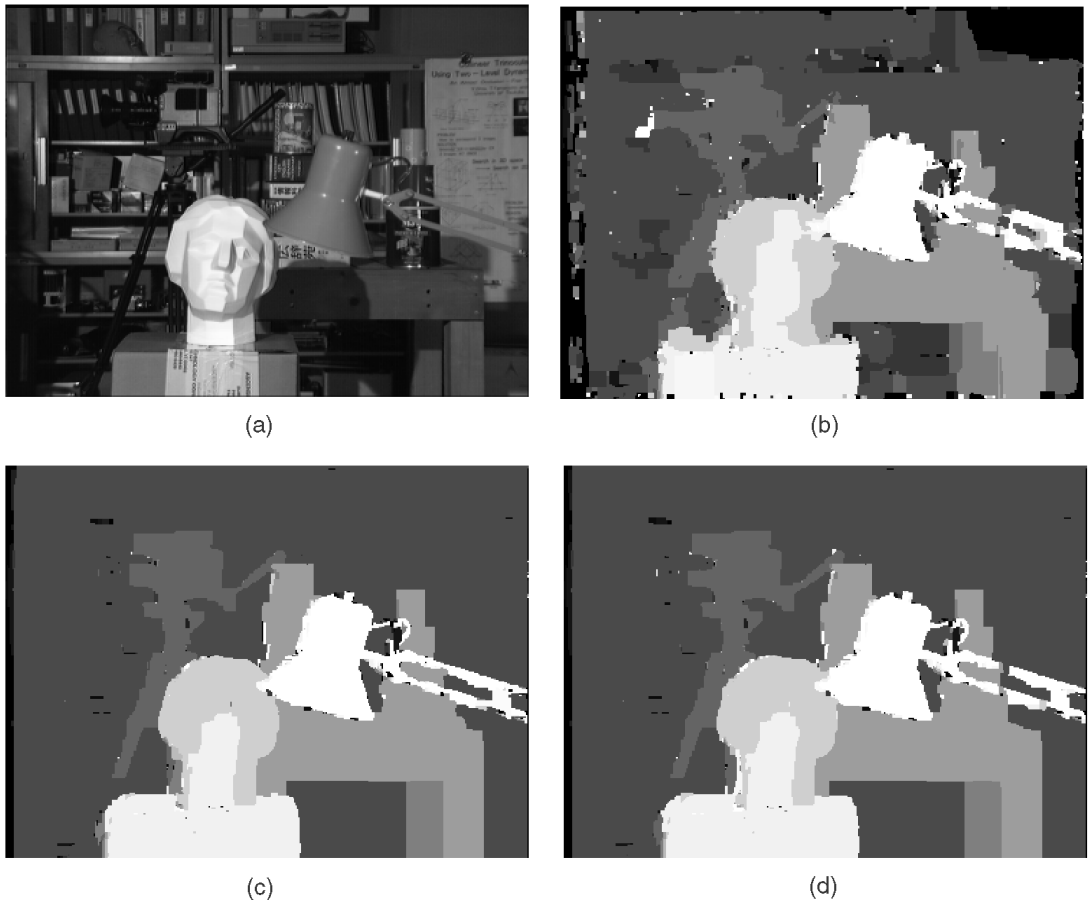

Figure 5: Stereo disparity example by Boykov et al. (a) One image from a stereo pair. (b) Disparity result, simulated annealing. (c) Disparity result, swap algorithm. (d) Disparity result, expansion algorithm.

What if our function is not sub-modular?

Many approximate methods have been proposed for minimizing non-submodular objective functions. Some good examples:

- Quadratic pseudo-boolean optimization (QPBO)[5]. Code: `https://pub.ist.ac.at/~vnk/software.html`
- "Local submodularization for binary pairwise energies" [3].
- "A comparative study of modern inference techniques for structured discrete energy minimization problems" [4] (Excellent survey paper, with implementations)

Summary

- Via the swap-move and expansion-move algorithms presented here, minimum graph cuts can be used to solve a fairly wide class of combinatorial optimization problems typically occuring in image processing.
- The solutions are guaranteed to be locally optimal, in a strong sense.
- In contrast to "standard" minimum cuts, these techniques can solve labeling problems involving more than two labels. The number of labels must still be finite.

References

[1] Ravindra K Ahuja, Özlem Ergun, James B Orlin, and Abraham P Punnen.

A survey of very large-scale neighborhood search techniques.

Discrete Applied Mathematics, 123(1):75–102, 2002.

[2] Yuri Boykov, Olga Veksler, and Ramin Zabih.

Fast approximate energy minimization via graph cuts.

Pattern Analysis and Machine Intelligence, IEEE Transactions on, 23(11):1222–1239, 2001.

[3] Lena Gorelick, Yuri Boykov, Olga Veksler, Ismail Ben Ayed, and Andrew Delong.

Local submodularization for binary pairwise energies.

IEEE transactions on pattern analysis and machine intelligence, 39(10):1985–1999, 2016.

[4] Jörg H Kappes, Bjoern Andres, Fred A Hamprecht, Christoph Schnörr, Sebastian Nowozin, Dhruv Batra, Sungwoong Kim, Bernhard X Kausler, Thorben Kröger, Jan Lellmann, et al.

A comparative study of modern inference techniques for structured discrete energy minimization problems.

International Journal of Computer Vision, 115(2):155–184, 2015.

[5] Vladimir Kolmogorov and Carsten Rother.

Minimizing nonsubmodular functions with graph cuts-a review.

IEEE transactions on pattern analysis and machine intelligence, 29(7):1274–1279, 2007.

[6] Vladimir Kolmogorov and Ramin Zabin.

What energy functions can be minimized via graph cuts?

Pattern Analysis and Machine Intelligence, IEEE Transactions on, 26(2):147–159, 2004.