Computer Vision for Visual Effects

CVFX 2015

Seam Carving

- *Seam Carving for Content-Aware Image Resizing*
  - Avidan & Shamir
  - SIGGRAPH 2007

Retargeting

crop? scaling?
fisheye-view warping?

content-aware

Seam Carving

- Energy function with respect to image content
- Carving-out or inserting a seam
  - a connected path of low energy pixels crossing the image from top to bottom or from left to right
- Successively removing or inserting seams = reduce or enlarge the size of an image

Energy Functions?

- Gradient magnitude
- Entropy
- Visual saliency
- Eye-gaze movement
- User-guided
- ...

Background

- Feature-Aware Texturing
- Digital Photomontage
- Drag-and-Drop Pasting
- Image Quilting
- ...

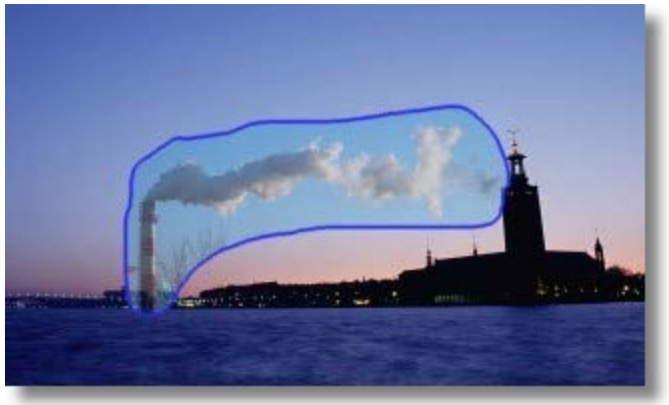

The Basic Energy Function

- Remove unnoticeable pixels that blend with their surroundings

$$e_1(\mathbf{I}) = |\frac{\partial}{\partial x}\mathbf{I}| + |\frac{\partial}{\partial y}\mathbf{I}|$$

- How to remove pixels?

Seams

- Vertical seam

image $\mathbf{I}$: $n \times m$

$$\mathbf{s^x} = \{s_i^x\}_{i=1}^n = \{(x(i), i)\}_{i=1}^n, \text{ s.t. } \forall i, |x(i) - x(i-1)| \leq 1$$

$$x : [1, \ldots, n] \rightarrow [1, \ldots, m]$$

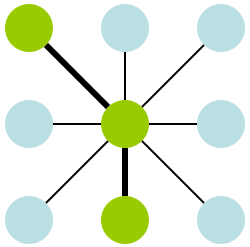

$$\mathbf{I_s} = \{\mathbf{I}(s_i)\}_{i=1}^n = \{\mathbf{I}(x(i), i)\}_{i=1}^n$$

- Horizontal seam

$$\mathbf{s^y} = \{s_j^y\}_{j=1}^m = \{(j, y(j))\}_{j=1}^m, \text{ s.t. } \forall j |y(j) - y(j-1)| \leq 1$$

$$y : [1, \ldots, m] \rightarrow [1, \ldots, n]$$

Optimization

$$E(\mathbf{s}) = E(\mathbf{I}_\mathbf{s}) = \sum_{i=1}^{n} e(\mathbf{I}(s_i))$$

$$s^* = \min_{\mathbf{s}} E(\mathbf{s}) = \min_{\mathbf{s}} \sum_{i=1}^{n} e(\mathbf{I}(s_i))$$

- Dynamic programming

$$\begin{aligned} M(i,j) \quad = \quad & e(i,j) + \\ & \min(M(i-1,j-1), M(i-1,j), M(i-1,j+1)) \end{aligned}$$

(a) Original

(b) Crop

(c) Column

(d) Seam

(e) Pixel

(f) Optimal

$$\frac{1}{|\mathbf{I}|}\sum_{p\in I} e(p)$$

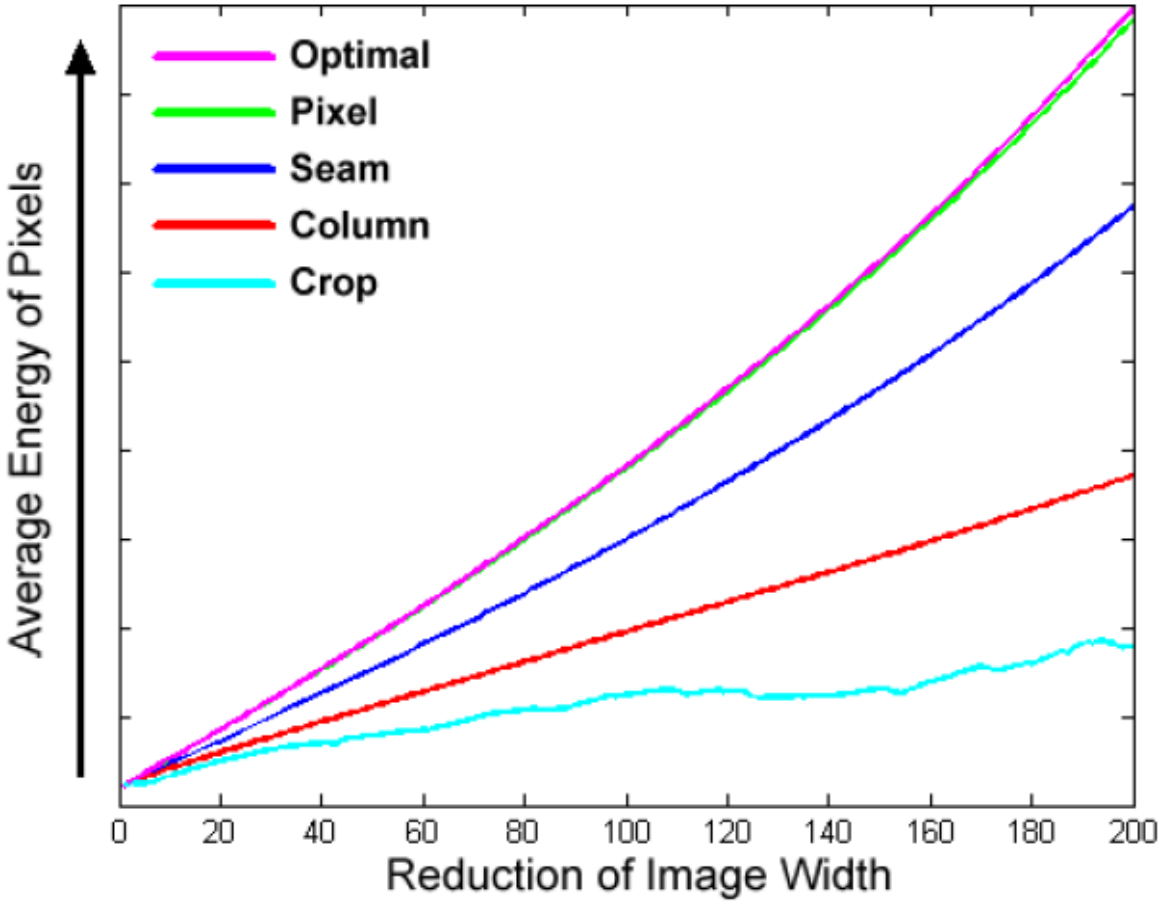

not a good energy function?

Other Energy Functions

- Entropy
  - How?
  - 9x9 window

$$H(X) = \mathrm{E}(I(X)) = -\sum_{i=1}^{n} p(x_i) \log_2 p(x_i)$$

- Segmentation
  - How? (segmentation+e_1) What's the effect?
- Histogram of Oriented Gradients (HOG)
  - 8 bins, 11x11 window

$$e_{HoG}(\mathbf{I}) = \frac{|\frac{\partial}{\partial x}\mathbf{I}| + |\frac{\partial}{\partial y}\mathbf{I}|}{\max\left(HoG(\mathbf{I}(x,y))\right)}$$

follow the edge and do not cross it

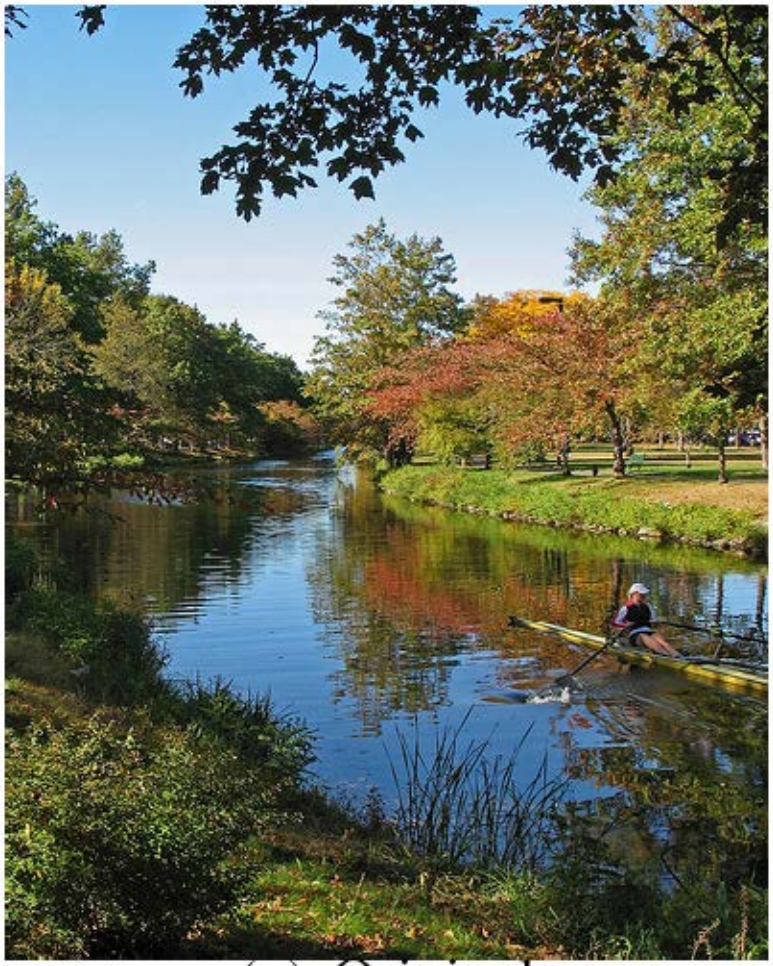

(a) Original

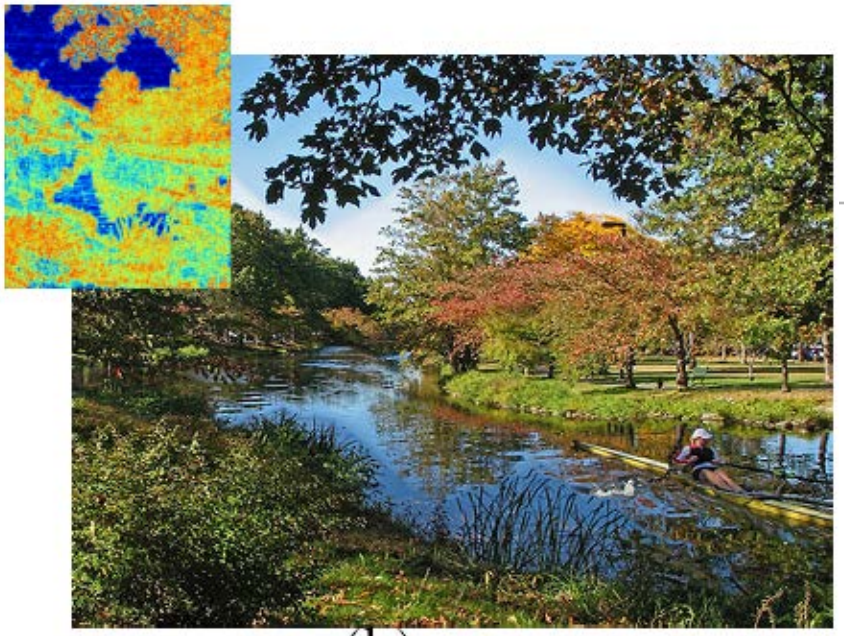

(b) e_1

(c) $e_{Entropy}$

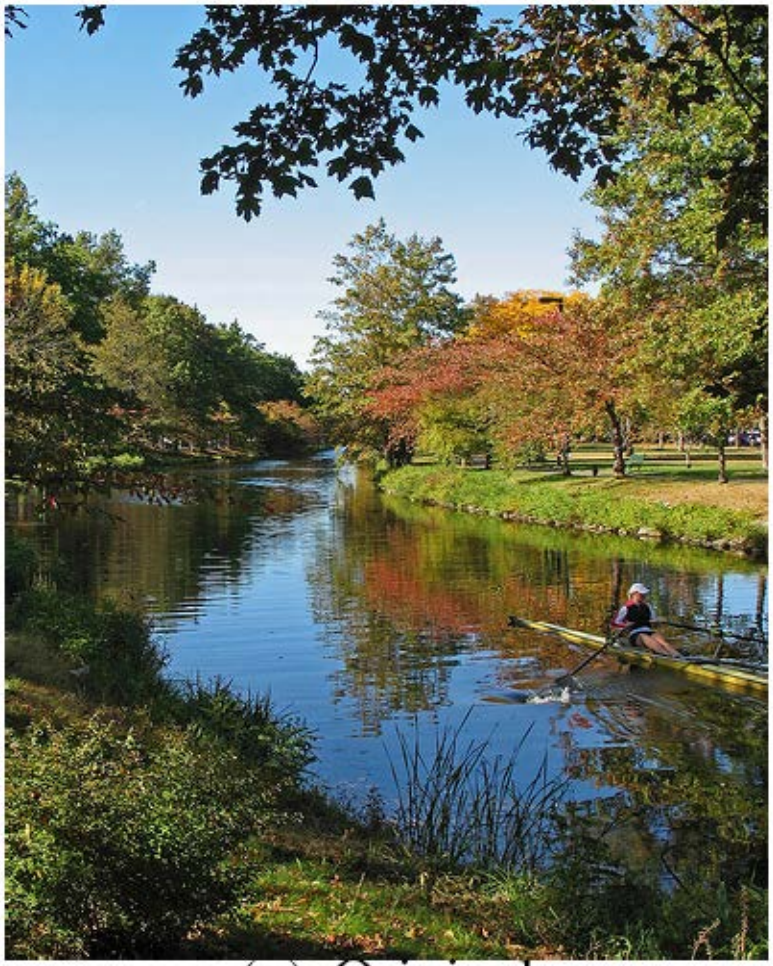

(a) Original

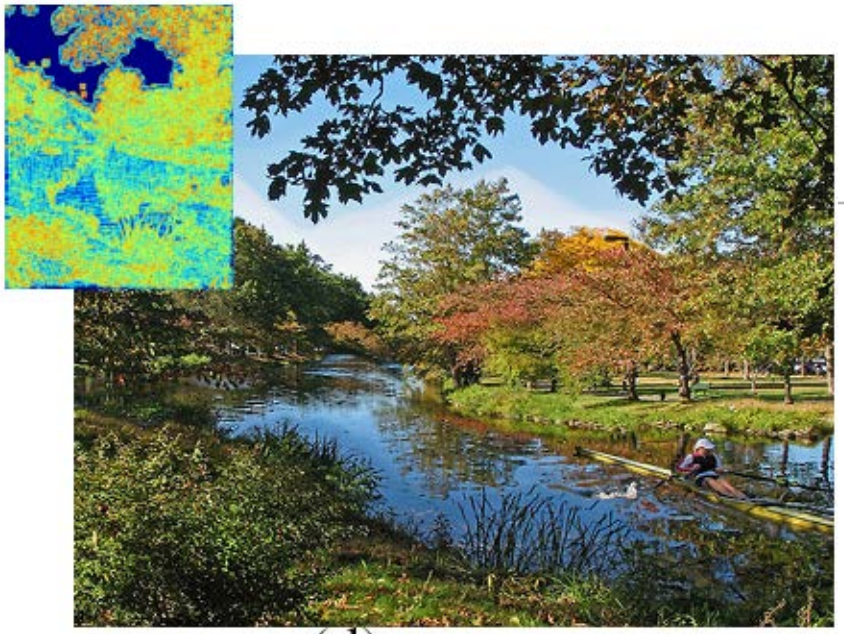

(d) e_{HoG}

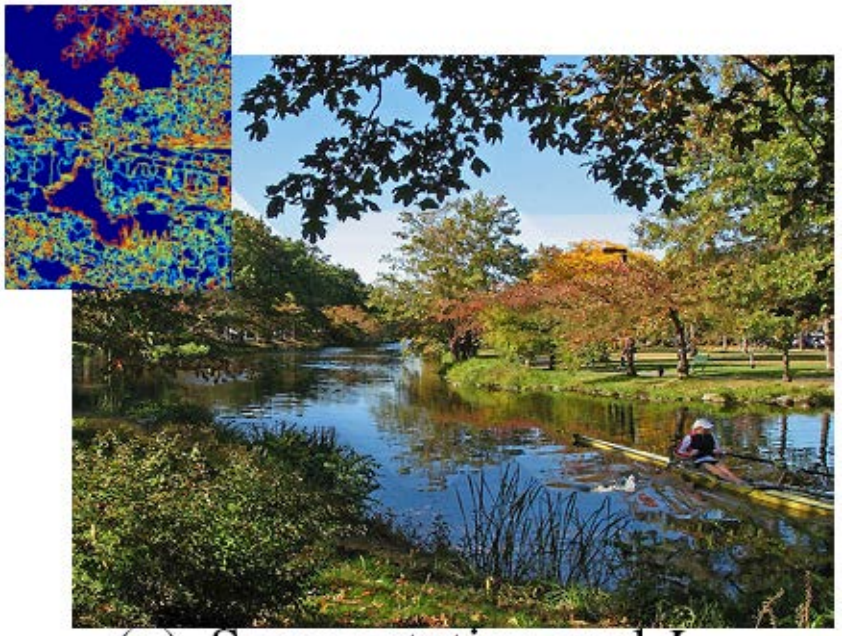

(e) Segmentation and L_1

Aspect Ratio Change

seam removal
scaling
cropping




$n \times m$ to $n \times m'$ where $m - m' = c$

Retargeting with Optimal Seams-Order

$$n \times m \longrightarrow n' \times m' \qquad m' < m \text{ and } n' < n$$

$$\min_{\mathbf{s^x},\mathbf{s^y},\alpha} \sum_{i=1}^{k} E(\alpha_i \mathbf{s_i^x} + (1-\alpha_i)\mathbf{s_i^y})$$

$$k = r + c \qquad r = (n - n') \qquad c = (m - m')$$

$$\alpha_i \in \{0,1\}\ , \sum_{i=1}^{k} \alpha_i = c\ , \sum_{i=1}^{k} (1-\alpha_i) = r$$

Dynamic Programming

transport map

$$\mathbf{T}(0,0) = 0 \qquad n-r+1 \times m-c \qquad n-r \times m-c+1$$

$$\begin{aligned}\mathbf{T}(r,c) \quad = \quad & min(\mathbf{T}(r-1,c) + E(\mathbf{s^x}(\mathbf{I_{n-r-1 \times m-c}})), \\ & \mathbf{T}(r,c-1) + E(\mathbf{s^y}(\mathbf{I_{n-r \times m-c-1}})))\end{aligned}$$

Enlarging

- Compute the optimal vertical seam on the current image
- Duplicate the pixels of the seam by averaging them with their left and right neighbors

- How to prevent from choosing the same seam repeatedly?
  - How large the step size should be?
  - How to handle excessive image enlarging

(a) (b)

(c) (d)

(e) (f)

Content Amplification

scaling followed by seam carving

Gradient Domain

Object Removal

- Remove user-marked target
- Change the whole image

Object Removal

Multi-size Images

- Real-time retargeting
- Compute the index maps **V** and **H**
  - Keep the order of removal

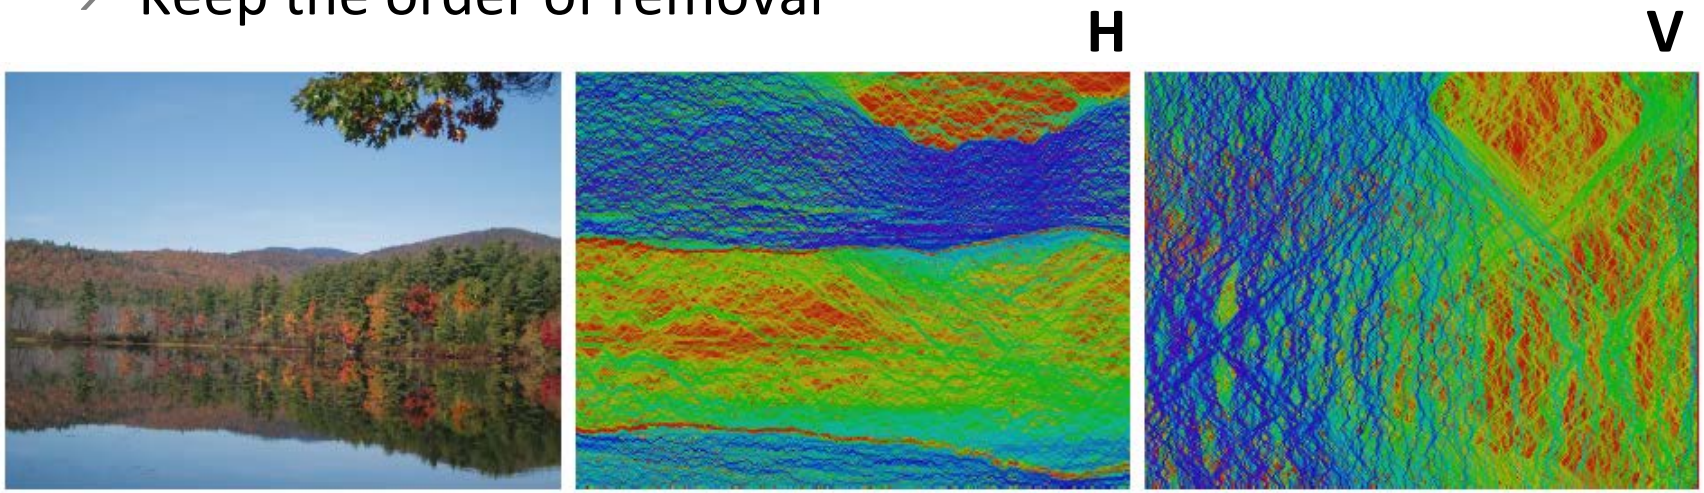

- How to support both dimension resizing using the index maps?

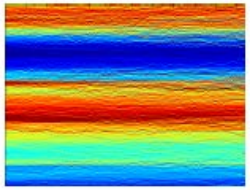

Consistent Index Maps

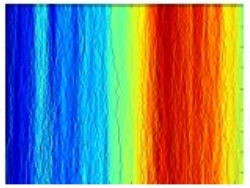

- **H** and **V** are consistent if every horizontal seam intersects (or touches) all the vertical seam indexes and every vertical seam intersects all horizontal seam indexes.

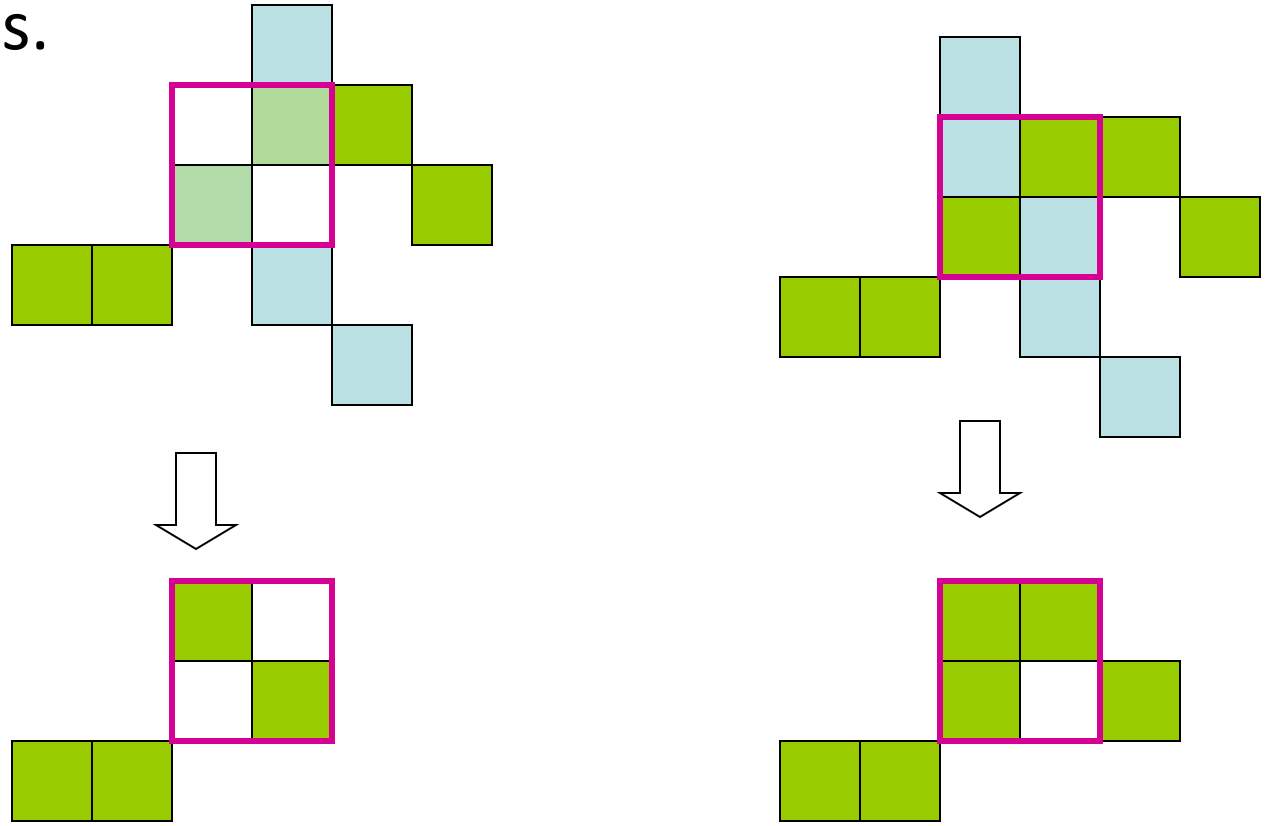

Conclusion