

Introduction to Algorithms and Data Structures

Lecture 19: “Seam Carving” (via Dynamic Programming)

Mary Cryan

School of Informatics
University of Edinburgh

Seam Carving

We need to fit images into new dimensions (relevant for tablet/mobile layouts).

Naïve approaches to adapting to varying dimensions include cropping and scaling. Both have their flaws - see video by Shai Avidan and Ariel Shamir: https://www.youtube.com/watch?v=6NcIJXTlugc&feature=emb_logo

A better, more flexible, approach is to search for **seams** in the image, a **seam** being a connected sequence of pixels running from top-to-bottom (*vertical*) or from left-to-right (*horizontal*).

- ▶ More general than deleting a column.
- ▶ Seams can be deleted (or duplicated) and the rows and columns of the altered image will have uniform lengths 1 less (or more) than before.
- ▶ For re-sizing images, we will want to find seams of low-energy (where there is little difference between the seam pixels and their surrounding pixels).

Seam Carving

We are given an image $\mathbf{I} : [m] \times [n]$ where each pixel is a colour (maybe RGB). Our dimensions are $m \times n$ but we want to fit to different dimensions $m' \times n'$.

Definition

In a image \mathbf{I} of dimensions $m \times n$, we define a **vertical seam** to be any sequence

$$\mathbf{s} = j_1, \dots, j_m \in [n]$$

such that for every $i \in [m] \setminus \{1\}$, we have $|j_i - j_{i-1}| \leq 1$ and $2 \leq j_i \leq n - 1$ (don't touch left/right sides).

A **horizontal seam** is any sequence

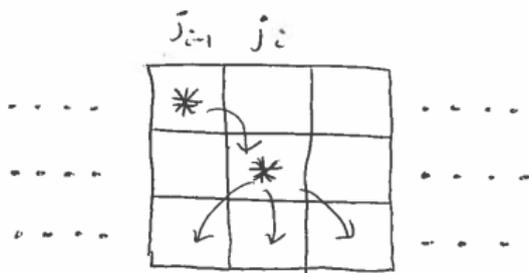
$$\mathbf{s} = i_1, \dots, i_n \in [m]$$

such that for every $j \in [n] \setminus \{1\}$, we have $|i_j - i_{j-1}| \leq 1$ and $2 \leq i_j \leq m - 1$.

(recall $[k]$ is the set of values $\{1, \dots, k\}$)

Seams and Energy

In building a (vertical) seam we are allowed to move (straight) down, 1-pixel left, or 1-pixel right.



"8-connected window for constructing a seam"

We need to evaluate the **energy** of a pixel, as we prefer low-energy seams.

We assume some energy function $e = e_{\mathbf{I}}$ applied to pixels of that image.

Then

$$e(\mathbf{s}) =_{\text{def}} \begin{cases} \sum_{i=1}^m e_{\mathbf{I}}(i, j_i) & \mathbf{s} \text{ is a vertical seam} \\ \sum_{j=1}^n e_{\mathbf{I}}(l_j, j) & \mathbf{s} \text{ is a horizontal seam} \end{cases}$$

Energy functions

Many options for (pixel) energy function, often a (local) **gradient** score.

- ▶ L_1 gradient scoring (for pixel (i, j)) can be written as

$$e_{\mathbf{I}}(i, j) =_{\text{def}} \left| \frac{\partial}{\partial x} \mathbf{I} \right|_{i, j} + \left| \frac{\partial}{\partial y} \mathbf{I} \right|_{i, j}.$$

- ▶ $\frac{\partial}{\partial x}, \frac{\partial}{\partial y}$ are defined in the image processing context:

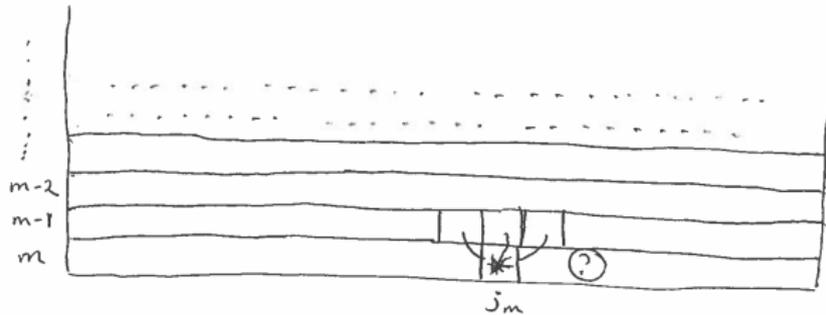
- ▶ For example the **Sobel operators** can be used to calculate $\frac{\partial}{\partial x}$ and $\frac{\partial}{\partial y}$:

$$\begin{bmatrix} -1 & 0 & +1 \\ -2 & 0 & +2 \\ -1 & 0 & +1 \end{bmatrix} \quad \begin{bmatrix} -1 & -2 & -1 \\ 0 & 0 & 0 \\ +1 & +2 & +1 \end{bmatrix}$$

- ▶ The first (vertical) would calculate $\frac{\partial}{\partial x} |_{i, j}$ as $(I_{i-1, j+1} - I_{i-1, j-1}) + 2(I_{i, j+1} - I_{i, j-1}) + (I_{i+1, j+1} - I_{i+1, j-1})$.
- ▶ A colour image will have 3 color channels, so 3 $\frac{\partial}{\partial x}$ and 3 $\frac{\partial}{\partial y}$ scores to sum.

Computing an optimal seam

We assume we are looking for a vertical seam (without loss of generality).
Take a recursive view and understand the optimal seam in terms of (slightly) shorter seam



Suppose an optimal vertical seam ends at pixel (m, j_m) .

- then we know that j_{m-1} was either j_m , j_{m-1} , or j_m+1 .
- depends on which of $(m-1, j_m)$, $(m-1, j_{m-1})$ and $(m-1, j_m+1)$ is the bottom endpoint of the best $(m-1)$ -length vertical seam.
- three subproblems of length $(m-1)$.

Recurrence for optimal (vertical) seam

Assume that we have precomputed $e_{\mathbf{I}}(i,j)$ for every pixel $1 \leq i \leq m$, every $1 \leq j \leq n$ ($\Theta(1)$ time for each (i,j)).

($e_{\mathbf{I}}(i,j)$ slightly different for top/bottom rows)

Definition

For every $i,j, 1 \leq i \leq m, 1 \leq j \leq n$, we define $opt_{\mathbf{I}}(i,j)$ to be the cost of the minimum-cost vertical seam from (somewhere in) row 1 to pixel (i,j) , where cost is scored as at the end of slide 4.

We have the following recurrence:

$$opt_{\mathbf{I}}(i,j) = e_{\mathbf{I}}(i,j) + \begin{cases} 0 & \text{if } i = 1 \\ \min\{opt_{\mathbf{I}}(i-1,j-1), \\ opt_{\mathbf{I}}(i-1,j), \\ opt_{\mathbf{I}}(i-1,j+1)\} & \text{if } i > 1 \end{cases}$$

(we will set $e_{\mathbf{I}}(i,j) \leftarrow \infty$ if $j = 1$ or n)

Dynamic programming implementation

We will need a table/array of size $m \cdot n$, let the table be opt .
(I assume indexing starts at 1 for this algorithm)

- ▶ Entry $opt[i, j]$ will store the value of $opt_{\mathbf{I}}(i, j)$ (when we have computed it).
- ▶ We will need to have the local $e_{\mathbf{I}}(i, j)$ energy values pre-computed (for each $1 \leq i \leq m, 1 \leq j \leq n$) and stored in a table e (of dimensions $m \times n$). We will set $e[i, 1] \leftarrow \infty, e[i, n] \leftarrow \infty$ for all $i \in [m]$ to make sure the seam avoids the sides.
- ▶ For image processing applications we definitely need to know the pixel sequence for the *actual* seam of optimal score.
Define another table/array p of dimensions $m \times n$ to hold $-1, 0, 1$ values (indicating whether j_i was $j_{i-1} - 1, j_{i-1}, j_{i-1} + 1$ for the good seam).

Dynamic programming implementation

Algorithm Vertical-Seam(I, m, n)

1. **for** $j \leftarrow 1$ **to** n
2. **for** $i \leftarrow 1$ **to** m
3. $e[i, j] \leftarrow$ "compute $e_I(i, j)$ " // $\Theta(1)$ time
4. $opt[1, j] \leftarrow e[1, j], p[1, j] \leftarrow 0$ //Base case
5. **for** $i \leftarrow 1$ **to** m
6. **for** $j \leftarrow 1$ **to** n
7. $opt[i, j] \leftarrow opt[i - 1, j], p[i, j] \leftarrow 0$ //default case
8. **if** $opt[i - 1, j - 1] < opt[i, j]$ **then**
9. $opt[i, j] \leftarrow opt[i - 1, j - 1], p[i, j] \leftarrow -1$
10. **if** $opt[i - 1, j + 1] < opt[i, j]$ **then**
11. $opt[i, j] \leftarrow opt[i - 1, j + 1], p[i, j] \leftarrow +1$
12. $opt[i, j] \leftarrow opt[i, j] + e[i, j]$ //Always add $e[i, j]$
13. $j^* \leftarrow 2$
14. **for** $j \leftarrow 1$ **to** n
15. **if** $opt[m, j] < opt[m, j^*]$ **then** $j^* \leftarrow j$
16. Print("Best vertical seam ends at cell (m, j^*) ").

Wrapping up

- ▶ After the algorithm has terminated, we find the optimal vertical seam by searching row m for the minimum $opt[m, j]$ value (one final loop).
- ▶ The double-loop between lines 5-12 does the main work, computing the opt values using the recurrence. There are $m \cdot n$ iterations of lines 7-12, and we can check that for a specific (i, j) , lines 7-12 take $O(1)$ time. So the algorithm has worst-case running-time $O(mn)$.
(ok, should also say lines 1-4 take $O(mn)$ and lines 13-16 take $O(n)$)
- ▶ The values in the p array make it very easy to reconstruct the actual sequence of pixels forming the seam (even easier than edit distance, etc).
- ▶ We specify how to compute $e[i, j]$ as the specific energy function can vary (see Avidan-Shamir paper). These are functions of local (to (i, j)) pixels and hence will always be computable in $O(1)$ -time (per (i, j) entry).

As discussed in the video, the seam will either be deleted (if we are aiming to **reduce the width**) or alternatively duplicated (if aiming to **increase width**).

There is a similar algorithm for Horizontal-Seam.

Reading Materials

Seam Carving:

- ▶ “Seam-Carving for Content-Aware Image Resizing” (clickable). I have used slightly different notation.
- ▶ There are many Seam-Carving implementations in Python available on the Internet, worth Googling them and taking a look. They are usually using `numpy` for file i/o and other functionality, plus some image processing resource too.

