SEAMLETS: CONTENT-AWARE NONLINEAR WAVELET TRANSFORM

David D. Conger, Hayder Radha, Fellow, IEEE
Michigan State University
Dept. of Electrical & Computer Engineering
East Lansing, MI 48824-1226

Mrityunjay Kumar
Eastman Kodak Company
Rochester, NY 14650-2118

ABSTRACT

With the rise of mobile media devices, resizing an image or video to fit a screen of arbitrary size has become an important topic. In general, arbitrary *resizing* does not preserve the original image aspect ratio, and hence, can introduce significant visual distortion. Content-aware *retargeting* methods have been proposed to preserve salient features of an image or video ([1][3][5][6]) even under arbitrary changing of the aspect ratio. While these are useful tools, little attention has been given to the development of an efficient retargeting-driven representation of images. Such representation, for example, can lead to efficient distribution of media destined for retargeting at a variety of devices with different aspect ratios. We propose a new framework, *seamlets*, which utilizes the discrete wavelet transform (DWT) to perform content-aware arbitrary media resizing. In essence, seamlets provide an efficient multi-resolution representation for retargeting applications. The *seamlet transform* elegantly generalizes the DWT and *seam carving*, which is an intriguing retargeting method. The result is a new image framework, *seamlets*, that inherits the benefits of retargeting and wavelets.

Index Terms— Wavelet transforms, Image retargeting, Seam carving, Image resizing

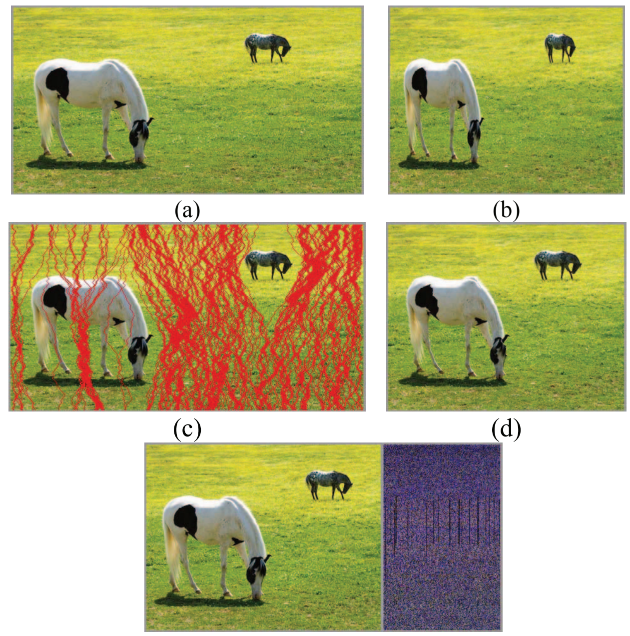

Figure 1: (a) Original image, (b) uniform resizing to 60% width, (c) seams (red) found using seam carving and (d) result of removing seams to achieve 60% width, (e) Vertical seamlet decomposition using Haar wavelet. [Image courtesy of mr.bmonroe on flickr.com]

1. INTRODUCTION

Mobile devices such as cell phones and personal media players are quickly changing what is meant by typical image sizes. With an increased demand for media content, the ability to make full use of their displays becomes a real challenge. The screen sizes are generally dictated by what is considered mobile, and are therefore no bigger than what might fit in one's pocket. Furthermore, there are no standards on screen sizes and so generally every manufacturer produces different sized screens. This initially was not a problem since displaying media on these devices was never a primary function. However, with increased memory and processing power, distributing media to mobile devices has become a mainstay and, more often than not, a primary selling point.

Media can be cropped or stretched to fit different device displays, but generally this introduces undesirable distortion (cf. Figure 1(b)). Adding borders to an image or video to preserve the aspect ratio is another option, but this leaves an already-small mobile device display with an even smaller content area.

This motivates what is known as *retargeting*, which is the process of changing the aspect ratio of an image or video while maintaining the content and appearance of the original media. One rather elegant solution called seam carving (SC) – developed by S. Avidan and A. Shamir in [1][4] – provides a way to systematically remove pixels from visually "unimportant" paths, or *seams*, effectively reducing the height or width by one pixel at a time, in a relatively unnoticeable way (cf. Figure 1(c,d)). Similarly, pixels can be added to these paths to achieve an increase in the dimension. By repeating seam removal or insertion, essentially any dimension can be achieved, although, too many iterations eventually introduce very noticeable distortion. This of course depends on the signal and, specifically, how much "important" information an image or video contains.

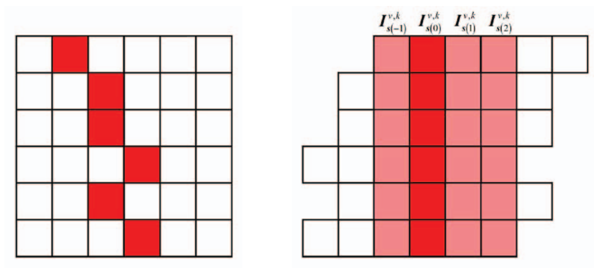

(a) (b)

Figure 2: (a) Image with seam shown in red, (b) Image rows realigned at seam with seam image shown in pink for a=2, b=1.

It is worth mentioning that several global retargeting solutions have been proposed ([3][5][6]). Although they prove to be effective options, such techniques require processing an entire image or video at once which may be too complex and too costly for mobile devices. SC on the other hand is more scalable since it decreases or increases the dimension by one at each step, effectively achieving every target size between the original and desired target sizes. This is especially beneficial when the target size is not known ahead of time. In this case, the SC algorithm proceeds by finding all the seams in the image or video and pre-storing the locations of these paths. Then, later, any size can be attained on-demand using the path data and the original media [1][4].

In this paper we present the groundwork for a general framework that provides an efficient multi-resolution image representation tailored for retargeting applications. For example, the proposed framework enables efficient distribution of images that are destined for retargeting at arbitrary and heterogeneous mobile devices. As generating this framework involves applying wavelets along seam paths, we call this new decomposition, *seamlets*. Methods such as SC can already be implemented in mobile devices, but large-scale distribution may require either wasteful transmission or disproportionate processing. We believe that seamlets provide a viable representation for retargeting.

The remainder of the paper is organized as follows. Section 2 describes the seamlet transform including the analysis and synthesis stages for one-dimensional signals. Extension of seamlets to 2D is outlined in Section 3. Section 4 discusses applications of seamlets for retargeting of images over networks. Section 5 provides some simulation results.

2. THE SEAMLET TRANSFORM

As mentioned above, image retargeting is essentially changing the aspect ratio of an image. Using seam carving, this can be achieved in a couple different ways. To make an image narrower, for instance, we can either remove vertical seams or add horizontal seams. While each approach will produce different looking images and different image dimensions, the resulting images can be rescaled to match the desired target size. Alternatively, the target size can be met by removing or adding both horizontal and vertical seams. This avoids the need for rescaling, however, there are limitations to how many seams can be added or removed.

Determining which way to proceed is an image dependent decision and a whole other topic for discussion. While our algorithm can be used for all of these cases, for the sake of brevity we will focus on the description of the one-dimensional seamlet transform for the case of "removing" vertical seams to meet a particular aspect ratio.

2.1. Overview

Here we present an overview of the vertical seamlet decomposition (refer to Figure 3(a)). At step k the image $I^{(k-1)}$ is modified as follows:

- the k-th seam path is found using seam carving
- the pixels near and on the seam are filtered using the DWT (as explained further below)
- the approximation coefficients replace those pixels to produce $I^{(k)}$
- the detail coefficients are stored as $y^{v,k}$

The starting image is $I^{(0)}$ and the final retargeted image is $I^{(K)}$. For convenience we will refer to the k-th instance of the retargeted image, namely $I^{(k-1)}$, as the k-th base image.

2.2. The Seam Path

A seam is an 8-connected path of pixels from one side of an image to the opposite side, containing exactly one pixel per row (vertical seam) or one pixel per column (horizontal seam). Based on [1], we define the k-th vertical seam path of an m×n image as

$$\mathbf{s}^{v,k} = \left\{ s_i^{v,k} \right\}_{i=1}^{m} = \left\{ \left(i, x^k\left(i\right) \right) \right\}_{i=1}^{m}$$

where x is a single-valued function such that

$$\left| x\left(i\right) - x\left(i-1\right) \right| \leq 1 \quad \forall i \in \left[2, m\right].$$

Although seamlet analysis does not require any particular path finding algorithm, we use seam carving because we are interested in content preservation. It should be noted that there are two seam criteria defined in [2]: forward energy and backward energy. In general, the forward energy criterion produces better results at the cost of longer compute times. For our purposes it does not matter which is used since either criteria finds seam paths as defined above and consequently either one can be used.

Having found the seam path, the SC algorithm would proceed by simply discarding the pixels along the path and move on to the next step. Instead, we will employ the DWT to reduce the spatial resolution near the seam path by filtering each row (or column for horizontal seams). To formulate the neighborhood near the seam, we would also need to know the seams adjacent to the k-th seam. We define a seam path shifted by p as

$$\boldsymbol{s}^{v,k}(p)=\left\{s_i^{v,k}(p)\right\}_{i=1}^{m}=\left\{\left(i,x^k(i)+p\right)\right\}_{i=1}^{m}$$

where $\boldsymbol{s}^{v,k}(0)$ is the k-th optimum seam path as defined above. Let the pixels of the base image associated with the k-th seam path be defined as

$$\boldsymbol{I}_{\boldsymbol{s}(p)}^{v,k}=I^{(k-1)}\left(\boldsymbol{s}^{v,k}(p)\right)$$

where $I^{(k-1)}$ is the base image at the start of step k.

2.3. Seamlet Analysis

A key difference from traditional wavelet decomposition of an image is that seamlet decomposition occurs iteratively. Rather than filtering the entire image at once, we only filter the pixels near and along a single seam at each step k. Formally, this pixel area around and including a vertical seam can be redefined as the $m\times(a+b+1)$ seam image

$$I_{v,k}=\left[\boldsymbol{I}_{\boldsymbol{s}(-b)}^{v,k}\quad\cdots\quad\boldsymbol{I}_{\boldsymbol{s}(0)}^{v,k}\quad\cdots\quad\boldsymbol{I}_{\boldsymbol{s}(a)}^{v,k}\right]$$

where a and b are nonnegative integers chosen based on the type and support size of the wavelet being used. Of course, there could be other ways to define the neighborhood of pixels around a seam, such as designing a way to account for the directionality. But for this brief introduction to the seamlet transform, we will use this straightforward definition. Note, however, that this is *not* simply a rectangular subset of the original image (see Figure 2(b)).

Because we want to reduce the dimension of the image, we are exploiting the fact that the DWT decimates the signal. This will invariably put certain conditions on acceptable values for a and b, and consequently on the size of the seam image. For example, using the basic Haar wavelet, the minimum possible size seam-image is $m\times 2$.

The next step is to apply the DWT to each row (or column for horizontal seams) using lowpass filter h and highpass filter g:

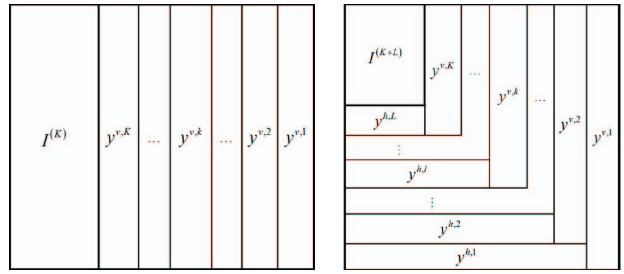

(a) (b)

Figure 3: (a) Vertical seamlet decomposition, (b) 2D seamlet decomposition.

$$I'_{v,k}(i,j)=\sum_{t=-b}^{a}I_{v,k}(i,t)h(2j-t)$$

$$y^{v,k}(i,j)=\sum_{t=-b}^{a}I_{v,k}(i,t)g(2j-t).$$

Note that the columns of the seam-image are indexed from $-b$ to a. The lowpass output $I'_{v,k}(i,j)$ replaces the seam image pixels corresponding to the starting image $I^{(k-1)}$; the highpass output $y^{v,k}(i,j)$ is stored beside the image (see Figure 3(a)).

Additionally, the seam path locations are stored in an index map that is used for reconstruction:

$$T(i,j)=\left(x^{K+1-j}(i)\right)_{i,j}$$

where x is found using seam carving as defined above.

This process is repeated until some minimum-size image $I^{(K)}$ is reached. The result is the seamlet decomposition (cf. Figure 1(e)).

2.4. Seamlet Synthesis

Reconstructing the image works by recombining the highpass and lowapss seamlet coefficients in the reverse order that they were decomposed from $I^{(K)}$ to $I^{(k_0)}$ where $I^{(k_0)}$ is either the original image or meets some desired target size.

Assuming the data is stored in a lossless manner, the quality of the reconstruction depends on the wavelet being implemented. If the wavelet provides perfect reconstruction, the seamlet reconstruction will also be perfect.

3. EXTENSION TO 2D

Seamlet decomposition can be extended to 2D by switching between vertical and horizontal seam processing. The framework for this is shown in Figure 3(b), however, note that it is *not* required to switch directions at every step as is shown; nor does the number of horizontal seams need to equal the number of vertical seams. The order of directions as well as the stopping point for either direction is image dependent and, consequently, difficult to automate.

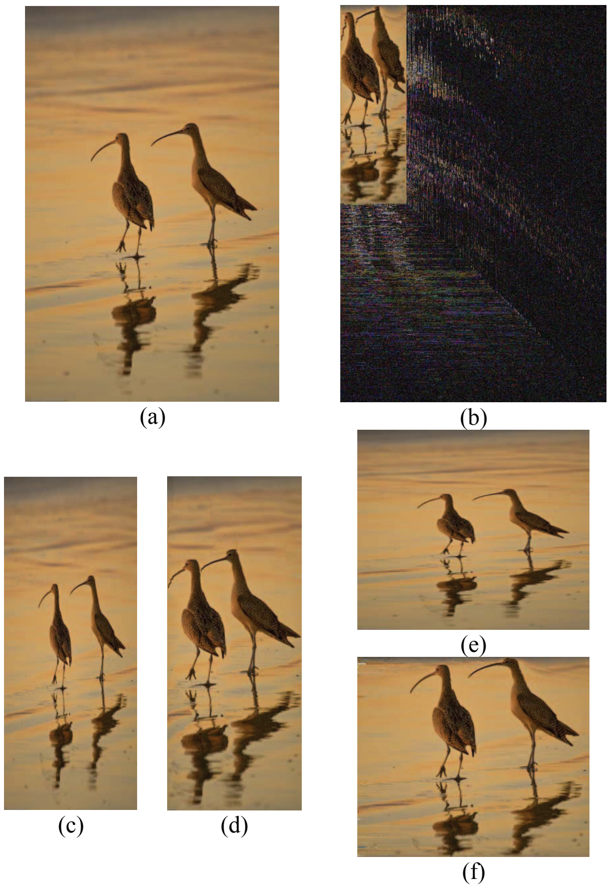

Figure 4: (a) Original image, (b) 2D seamlet decomposition, (c) resizing and (d) seamlet reconstruction to 501×200, (e) resizing and (f) seamlet reconstruction to 301×350. [Image courtesy of mikebaird on flickr.com]

Synthesis for this case would proceed just as for the 1D case by reconstructing the seam images in the reverse order they were decomposed. Of course, this might not immediately produce an image of the desired target size for the reason that not every possible image size is passed as an image is retargeted from $I^{(0)}$ to $I^{(K)}$. A simple solution to this problem is to reconstruct the image to $I^{(k_0)}$ such that both target dimensions are met or exceeded for the first time during the reconstruction process. From there, the image $I^{(k_0)}$ can be retargeted to the desired size.

Another solution is possible by considering the intersections of seams and modifying the seamlet coefficients to artificially reorder when each seam was processed. Although we have omitted these details for the sake of space, we have used this in our simulations.

4. SIMULATIONS

While we have experimented with various types of wavelets, we found that the Haar wavelet is particularly well-suited for retargeting due to its small support width. This easily enables reductions by one dimension in the height or width, which is consistent with the seam carving algorithm. For this reason, we have shown simulation examples using the Haar wavelet.

To support our claim of arbitrary media retargeting, we have taken an image represented using the 2D seamlet decomposition and reconstructed it to various sizes (see Figure 4). The original 600×400 image was decomposed to a 300x100 base image by alternating between vertical and horizontal seams. Starting from there, we were able to reconstruct to any arbitrary size using the methods discussed in Section 4. For comparison, we have also used resizing to obtain the same sizes. Due to space constraints, we are unable to show the results using only seam carving, however, there is no significant difference from our results.

5. CONCLUSIONS

We have presented a wavelet-based seam carving algorithm in the light of an efficient multi-resolution representation. Yet, we believe there is more significance to this algorithm in terms of wavelet theory. In future we plan to look further into these details and discover other applications.

6. REFERENCES

[1] S. Avidan, and A. Shamir, “Seam carving for content-aware image resizing,” *ACM Transactions on Graphics*, Vol. 26, No. 3, July 2007.

[2] S. Avidan, and A. Shamir, “Seam carving for media retargeting,” *Communications of the ACM*, Vol 52, No. 1, January 2009.

[3] Y. Guo, F. Liu, J. Shi, Z.-H. Zhou, and M. Gleicher, “Image retargeting using mesh parametrization,” *IEEE Transactions on Multimedia*, Vol. 11, No. 4, 2009.

[4] M. Rubinstein, A. Shamir, and S. Avidan, “Improved Seam Carving for Video Retargeting,” *ACM Transactions on Graphics*, Vol. 27, No. 3, 2008.

[5] D. Simakov, Y. Caspi, E. Shechtman, and M. Irani, “Summarizing visual data using bidirectional similarity,” *In Proceedings of CVPR*, 2008.

[6] L. Wolf, M. Guttmann, and D. Cohen-Or, “Non-homogeneous content-driven video-retargeting,” *IEEE International Conference on Computer Vision (ICCV)*, 2007.