

# Visibility Maps for Improving Seam Carving

Alex Mansfield<sup>1</sup>, Peter Gehler<sup>1</sup>, Luc Van Gool<sup>1,2</sup> and Carsten Rother<sup>3</sup>

<sup>1</sup> Computer Vision Laboratory, ETH Zürich, Switzerland

<sup>2</sup> ESAT-PSI, KU Leuven, Belgium

<sup>3</sup> Microsoft Research Ltd, Cambridge, UK

{mansfield, pgehler, vangool}@vision.ee.ethz.ch, carrot@microsoft.com

**Abstract.** In this paper, we present a new, improved seam carving algorithm. Seam carving efficiently removes pixels from an image to produce a retargeted image. It has proved popular with users and has been used as a component in many retargeting algorithms. We introduce the visibility map, a new framework for pixel removing image editing methods. This allows us to cast retargeting as a binary graph labelling problem. We derive a general algorithm which uses seam carving operations for efficient greedy optimization of a well defined energy, and compare this with forward energy seam carving and shift map image editing. We test this method with varying parameters on a large number of images, and present an improved seam carving algorithm which can demonstrably produce better results. We draw general conclusions about pixel removing methods for retargeting and motivate future directions of research.

## 1 Introduction


Image retargeting aims to generate effective visualizations of images from different sources on different displays. Given the increasing variation of sources and displays, from more traditional cameras and monitors to time-of-flight webcams and smartphones, there has been great interest in this application in recent years.

Seam carving [2, 15] is one of the most popular image retargeting methods. This simple algorithm removes a set of pixels from the input image to generate the output. Since its introduction by Avidan and Shamir in 2007, many seam carving implementations have become available, including in Adobe Photoshop<sup>4</sup>, the Liquid Rescale plugin for GIMP<sup>5</sup> and online at [rsizr.com](http://rsizr.com). Seam carving has also been built on in many academic works [4, 6–8, 11, 16, 20]. However, most of this work uses seam carving as a complete algorithm, without modification.

The goal of this paper is to analyze, extend and improve seam carving itself. To this end we cast the problem in a new framework, the *visibility map*, illustrated in Fig. 1. This map provides a natural description of methods that remove pixels from the input image to generate the output, allowing us to describe retargeting as a binary graph labelling problem. We define an energy over a visibility map that can still be optimized using seam carving operations. We

<sup>4</sup> See <http://www.adobe.com/products/photoshop/photoshopextended/features/>

<sup>5</sup> Available at <http://liquidrescale.wikidot.com/>



**Fig. 1.** The visibility map shows which pixels are visible in the output image after an image editing operation. Visible pixels are labelled 1 (shown in white), non-visible pixels labelled 0 (shown in black). As with all figures, best viewed in colour

explore different versions of the visibility map energy and also optimization options that open up due to this new viewpoint. Results for numerous parameters are generated on a large set of images to determine an improved seam carving algorithm.

Our improved seam carving has a number of advantages. Most importantly, it optimizes an energy defined directly between the input and output images, unlike the commonly used forward energy seam carving of [15], as shown in Sect. 5.1. This allows a clearer understanding of our energy and allows direct comparison to results generated by other optimization methods. It also produces demonstrably better results on many images, e.g. see Fig. 1.

In summary, our key contributions are: (1) The definition of retargeting as a binary graph labelling problem. (2) An efficient optimization scheme using seam carving operations, given energy terms from a well defined general family. (3) An improved seam carving algorithm.

We next describe related work. In Sect. 3 we then define the visibility map, from which we derive the general form of our improved seam carving algorithm (Sect. 4). In Sect. 5 we compare this to related methods. In Sections 6 and 7 we describe various energy and optimization options. In Sect. 8 we show results, and present our improved seam carving algorithm. Finally, in Sect. 9 we conclude and discuss future directions.

## 2 Related Work

Scaling and cropping have been long used in image editing and retargeting. Automatic methods have been used to guide these simple processes [16, 17]. However, these operations have fundamental limitations. Scaling keeps uninteresting parts of the image, while distorting structured objects such as faces and man-made objects when the scaling is non-uniform. A good crop maintains only the interesting parts of the image, but not all may fit within the desired output image size.

Cropping has been generalized to more flexible pixel removal methods. These methods may remove areas of uninteresting content while being able to rearrange the image to better show all of the interesting parts.

Seam carving [2, 15] fits into this category. Using simple low level energy terms, this algorithm iteratively removes pixels. Retargeted images at a range of sizes can be quickly generated. Its simplicity, speed and effectiveness has led it to be used as component of many retargeting methods including [4, 6–8, 11, 16, 20]. Most use seam carving as a complete algorithm, with the exceptions of our previous work [11] where we extended seam carving to better protect objects during retargeting, and [7] which redefines seams for video retargeting to achieve improved results.

Other methods also operate by pixel removal. These include shift map image editing [12], which optimizes a mapping from pixels in the output image to pixels in the input image. For retargeting they add a label ordering constraint which maintains the ordering of pixels in the input image in the output image. In this case, the result can equivalently be generated by removing pixels.

Shift map image editing without this constraint, and other algorithms which generate outputs in terms of input pixels, e.g. [1, 5, 13, 14], also owe much of their effectiveness to pixel removal. However, allowing pixel re-arrangement and duplication gives greater flexibility, which must be appropriately constrained.

These methods have a number of drawbacks. When approximating scaling through downsampling, these methods suffer the same problem of causing non-uniform scaling of structured objects. Also, they may lead to discontinuities in lines and curves in the image, which can be very visually disturbing.

These issues have motivated other paradigms for retargeting. Non-linear warping/interpolation is used in [9, 10, 18, 19] among others to determine the output image. Pixel estimation is used in [3, 17] to minimize a patch-based bidirectional image similarity. The patch match algorithm [3] achieves interactive speeds, and allows very effective user interaction to be used to preserve lines and structured regions. However, these methods can be complex to implement, and usually require the optimization to be re-run from the beginning for each target size.

However, despite the drawbacks of seam carving, it is still popular in practice due to its simplicity, speed and effectiveness in a wide range of images. This motivates our aim to better understand and improve seam carving, which we do through the framework of the visibility map we introduce in the next section.

### 3 Visibility Map

In image editing and retargeting methods such as seam carving, the output image is created from the input image by simply removing pixels and squashing them together. This operation can be naturally defined through a binary graph labelling problem as follows. For each pixel  $(r, c)$  in the input image, there is one node that can take on one of two labels  $X_{r,c} \in \{0, 1\}$ . If the node label is 1 the corresponding pixel is visible in the output image; if 0, it is non-visible. In other

words the output image is generated by only showing the pixels whose nodes are labelled with a 1. We refer to this graph as the visibility map over the image.

Example visibility maps for cropping, downsampling, seam carving and our improved seam carving of an image are shown in Fig. 1. Due to the simple relationship between the input image and output image via the visibility map, this representation provides an intuitive framework through which to view pixel-removal methods.

Regarding notation: throughout the paper we will denote row indices with  $r$  and column indices with  $c$ , in an image  $R \times C$  in size.

## 4 Seam Carving Operations to Optimize a Visibility Map

The visibility map allows retargeting to be formulated as a binary labelling problem. However, the structure of the problem does not allow for a simple solution by standard binary labelling methods. For example, a fixed number of pixels must be removed from each row and column to maintain a rectangular output image. This constraint would take the form of higher order cliques in the graph which may make its solution intractable.

Instead, we determine a seam carving based approach for optimization. An important property of algorithms based on seam carving is their computational efficiency. At each iteration, the optimization is a dynamic program. We want to retain this efficiency, while extending the method to optimize a well-defined graph labelling problem. In this section, we show how this can be achieved.

### 4.1 Energy

We consider a general energy over a visibility map  $\mathbf{X}$  for retargeting that allows for efficient optimization by seam carving operations, as described in the following section. Throughout, we assume that vertical seams are being removed, without loss of generality. This energy takes the form

$$E(\mathbf{X}) = \sum_{r,c} \psi_{r,c}^U(\mathbf{X}) + \sum_{r,c_1 < c_r} \psi_{r,c_1,c_r}^H(\mathbf{X}) + \sum_{r > 1, c_u, c_d} \psi_{r,c_u,c_d}^V(\mathbf{X}) . \quad (1)$$


**Unary Terms.**  $\forall r, c$

$$\psi_{r,c}^U(X_{r,c}) = E_{r,c}^U[X_{r,c} \neq 0] , \quad (2)$$

where  $[\cdot]$  is the indicator function.

**Horizontal Contact Terms.** These are potential functions over higher order cliques defined  $\forall r, c_1 < c_r$  as

$$\psi_{r,c_1,c_r}^H(X_{r,c_1}, \dots, X_{r,c_r}) = \begin{cases} E_{r,c_1,c_r}^H, & X_{r,\{c_1,c_r\}} = 1, X_{r,\{c_1+1,\dots,c_r-1\}} = 0 \\ 0, & \text{otherwise} \end{cases} . \quad (3)$$



**Fig. 2.** Contact term potentials over higher order cliques are turned on when certain pixels come into contact. Nodes are shaded to show example configurations which turn these potentials on

**Vertical Contact Terms.** These are potential functions over higher order cliques defined  $\forall r > 1, c_u, c_d$  as

$$\psi_{r,c_u,c_d}^V(X_{r-1,1}, \dots, X_{r-1,c_u}, X_{r,1}, \dots, X_{r,c_d}) = \begin{cases} E_{r,c_u,c_d}^V, & X_{r-1,c_u} = 1, X_{r,c_d} = 1, \sum_{c=1}^{c_u-1} X_{r-1,c} = \sum_{c=1}^{c_d-1} X_{r,c} \\ 0, & \text{otherwise} \end{cases} \quad (4)$$

Note that the specific terms we use for  $E_*^*$  are defined later in Sect. 6.

The contact terms are so named because of their special sparsity properties. These potentials are only non-zero for configurations where certain pixels are brought into contact, e.g. for the horizontal terms only when nodes  $(r, c_l)$  and  $(r, c_r)$  are labelled 1 and all nodes in between are labelled 0. This is illustrated in Fig. 2. Although there are a huge number of these potentials, only a small number are “turned on” for each image configuration. This is the main reason why seam carving operations can be applied to minimize this energy.

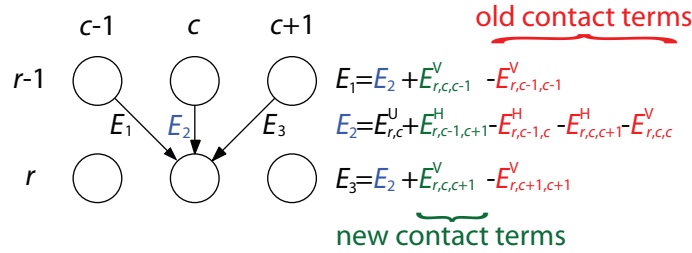
We also place locality constraints on the non-zero values these potential functions can take. We enforce that the term  $E_{r,c_l,c_r}^H$  may not be a function of the properties of any pixels other than those on row  $r$ , and  $E_{r,c_u,c_d}^V$  of any pixels other than those on rows  $r-1$  and  $r$ .

## 4.2 Optimization

In this section we show how the form of energy described in the previous section may be optimized by seam carving operations.

Let us first recap the seam carving method. Seam carving greedily removes one seam per iteration, where a seam is defined as an 8-connected path across the image with one pixel per row. Dynamic programming is used to efficiently optimize for the seam with lowest energy, with order  $O(RC + R)$ . This process of optimizing for a seam to remove we refer to as a *seam carving operation*. We define these operations explicitly in order to distinguish them from seam carving *algorithms*, which also define the energy terms to be used.

In terms of the visibility map, seam carving can be understood as follows. From an initial all-ones labelling, each seam ‘removed’ encodes a label switch



**Fig. 3.** Energy terms for seam carving operations. Green terms relate to new pixel contact, red terms to old pixel contact

of nodes with label 1 to label 0. This process is iterated until the target size is acquired.

We now explain why our energy can be optimized using seam carving operations. During the forward pass of dynamic programming, as each pixel in the seam is chosen, the seam pixel in the row above is already known, conditioned on the current pixel being contained in the optimal seam. With this information, it is clearly possible to determine the correct non-zero energy terms as described in Sect. 4.1, which are subject to locality constraints.

It is also possible to determine which pixels are newly brought into contact, as used in [15]. Hence it is known which of the potential functions over higher order cliques are turned on and off, as these depend only on pixel contact.

It is therefore clear that, at each iteration, seam carving operations can find the seam to remove which results in the minimum energy visibility map. The energy terms defined in Sect. 4.1 are used in dynamic programming over the current image as shown in Fig. 3. The (red) old contact terms are paid negatively at the shown locations as an equivalent but simpler alternative to paying these terms positively everywhere *except* at the locations shown. This is because this distortion would remain for seams elsewhere in the image. Note that these old contact terms are only paid if they were previously paid as (green) new contact, i.e. if the pixels referenced were not neighbours in the input image.

## 5 Relationship to Other Methods

### 5.1 Forward Energy Seam Carving

Our new algorithm results in the generalized energy terms for the seam carving operations shown in Fig. 3. The terms of forward energy seam carving [15] are similar, but with a key difference: they pay only the new contact terms, and not the terms related to old contact. Seam carving can thus be thought of as “forgetting” the original image and only taking into account distortion introduced at that iteration.

This means that forward energy seam carving does not optimize for an energy defined over a visibility map, and therefore not for an energy defined simply




Fig. 4. Seam carving forward energy is dependent on the seam removal order

between the input and output images. This can also be seen from the fact that the energy is dependent on the order of seam removal, as illustrated in Fig. 4. Consider the four neighbouring pixels shown and their horizontal contact terms. Removing pixels b then c brings into contact a and c with cost  $E_{ac}$  in the first seam and then a and d with cost  $E_{ad}$  in the second seam. Removing pixels c then b incurs costs  $E_{bd}$  and then  $E_{ad}$ , which are different in general.

This makes the seam carving forward energy harder to understand. An energy defined between the input and output allows better energy modelling and also comparison to other methods which produce visibility maps.

### 5.2 Shift Map Image Editing

The visibility map framework also has close connections with the shift map framework described by Pritch et al. [12]. The shift map is a multi-label mapping over the output image, where the label describes the shift between the pixel in each position and its original position in the input image.

The shift map is related to the visibility map. When a label ordering constraint is enforced, a shift map result can be represented by a visibility map. This relationship is illustrated in Fig. 5. If the shift map is given by  $M_{r,c}$ , then this relationship can be written formally as

$$X_{r,c} = \begin{cases} 1, & \exists (u,v) \text{ such that } ((u,v) + M_{u,v}) = (r,c) \\ 0, & \text{otherwise} \end{cases} \quad (5)$$




Fig. 5. The shift map is closely related to the visibility map. The corresponding entries for pixel (2,3) are highlighted

Comparing the two representations, while the shift map offers a clear description of the energy terms as shown in [12] and is not limited to maintaining pixel ordering, it yields a multi-label problem. The visibility map is an alternative which poses a binary labelling problem with the pixel ordering being implicitly enforced. In the context of image retargeting, or other problems where a pixel ordering may be desirable, the use of the visibility map therefore yields a simpler formulation than the equivalent shift map.

## 6 Improved Seam Carving – Energy

We have shown in Sect. 4 that seam carving operations can be used to optimize for a well defined energy over a visibility map. This gives us the key advantage of greater intuition into the behaviour of the algorithm, given a defined energy. This intuition can be used in designing a good energy for the problem. With this in mind, we now consider a number of different options for the energy terms.

For the contact energy terms, we consider the general form

$$\begin{aligned} E_{r,c_l,c_r}^H &= D_{r,c_l,c_r}^H + S_{r,c_l,c_r}^H \\ E_{r,c_u,c_d}^V &= D_{r,c_u,c_d}^V + S_{r,c_u,c_d}^V . \end{aligned} \quad (6)$$

We now describe options for these terms.

### 6.1 Distortion Terms

These terms measure the distortion created in the output image. The following notation is used:  $D_{r,c_l,c_r}^H$  is the horizontal distortion term, and  $D_{r,c_u,c_d}^V$  the vertical distortion term,  $\mathbf{I}$  is the input image, with magnitude  $|\mathbf{I}|$ . All terms have an order term  $n_D$ . In our experiments, we consider  $n_D \in \{1, 2\}$ .

#### Magnitude Distance.

$$\begin{aligned} D_{r,c_l,c_r}^H &= ||I|_{r,c_l} - |I|_{r,c_r}|^{n_D} \\ D_{r,c_u,c_d}^V &= ||I|_{r-1,c_u} - |I|_{r,c_d}|^{n_D} . \end{aligned} \quad (7)$$

For  $n_D = 1$ , this is the contact energy used in forward energy seam carving [15].

#### RGB Distance.

$$\begin{aligned} D_{r,c_l,c_r}^H &= \sum_{x \in \{R,G,B\}} |I_{r,c_l}^x - I_{r,c_r}^x|^{n_D} \\ D_{r,c_u,c_d}^V &= \sum_{x \in \{R,G,B\}} |I_{r-1,c_u}^x - I_{r,c_d}^x|^{n_D} . \end{aligned} \quad (8)$$

This energy is similar to the above, but makes use of differences in RGB colour rather than intensity magnitude.



### Relative RGB Distance.

$$\begin{aligned}
 D_{r,c_1,c_r}^H &= \sum_{x \in \{R,G,B\}} |I_{r,c_1}^x - I_{r,c_r-1}^x|^{n_D} + |I_{r,c_r}^x - I_{r,c_1+1}^x|^{n_D} \\
 D_{r,c_u,c_d}^V &= \sum_{x \in \{R,G,B\}} |I_{r-1,c_u}^x - I_{r-1,c_d}^x|^{n_D} + |I_{r,c_d}^x - I_{r,c_u}^x|^{n_D} .
 \end{aligned} \tag{9}$$

For  $n_D = 2$ , this is part of the contact energy used in shift map image editing [12].

## 6.2 Seam Terms

We can regularize the spatial distribution of the seams and thus provide an explicit regularization against seam ‘clumping’. Such clumping can occur when the distortion cost of removing a single clump of seams is lower than the cost of removing seams spread throughout the image, resulting in a visually disturbing seam of high distortion in the output image.

We consider two different possibilities. The following notation is used:  $S_{r,c_1,c_r}^H$  is the horizontal shift control term, and  $S_{r,c_u,c_d}^V$  the vertical shift control term.

### Repeat Cost for Intermediate Seams.

$$\begin{aligned}
 S_{r,c_1,c_r}^H &= (c_r - c_1 - 1)D_{r,c_1,c_r}^H \\
 S_{r,c_u,c_d}^V &= (c_u - c_d - 1)D_{r,c_u,c_d}^V .
 \end{aligned} \tag{10}$$

This seam term ensures that some cost is paid for each seam that has been removed, with that cost given by the energy of the currently visible seam.

### Average Unary Cost for Intermediate Seams.

$$\begin{aligned}
 S_{r,c_1,c_r}^H &= \frac{(c_r - 1) - (c_1 + 1)}{(c_r - 1) - (c_1 + 1) + 1} \sum_{c=c_1+1}^{c_r-1} E_{r,c}^U \\
 S_{r,c_u,c_d}^V &= (c_u - c_d - 1)D_{r,c_u,c_d}^V .
 \end{aligned} \tag{11}$$

This seam term leads to an approximation of the forward energy seam carving algorithm [15]. Consider removing a pixel with an already removed neighbour. Additional distortion energy and seam term energy ( $\Delta D + \Delta S$ ) is paid. The additional seam term energy, using this measure, approximates the previous distortion energy by the average unary term of the previously removed pixels ( $\Delta S \approx D_{\text{old}}$ ). The remaining energy is the new distortion energy ( $\Delta D + \Delta S = D_{\text{new}} - D_{\text{old}} + \Delta S \approx D_{\text{new}}$ ), which is what the forward energy seam carving algorithm optimizes for.

Note that a similar approximation for the vertical term is hard to define, so the repeating cost is again used.

### 6.3 Unary Energy Terms

By inspection of the form of the contact energies, it can be seen that they alone do not well model the problem of retargeting. Consider a retarget in which an area of interesting texture is completely removed, leaving only a homogeneous background. The contact energy terms will be low due to the low visible distortion, but clearly this is not the best way to retarget the image. In the terminology of [17], the contact terms provide a measure of *coherence* but not of *completeness*.

A unary term may be used to model this loss. We consider using the following as a simple saliency-based unary term, with a variable order given by  $n_U$ :

$$E_{r,c}^U = \left( \left| \left( \frac{\partial}{\partial x} \mathbf{I} \right)_{r,c} \right| + \left| \left( \frac{\partial}{\partial y} \mathbf{I} \right)_{r,c} \right| \right)^{n_U}. \quad (12)$$

For  $n_U = 1$ , this is the unary energy used in backward energy seam carving [2]. In our experiments, we consider  $n_U \in \{1, 2\}$  and also consider not using the unary term.


## 7 Improved Seam Carving – Optimization

### 7.1 Refinement

A key advantage of having a well defined energy over a visibility map is that we can compare different optimization techniques and combine them to achieve a lower energy. We consider an optimization step based on the observation that we can not only remove seams, but also put them back in. We refer to this as visibility map refinement. Using refinement steps may allow a lower energy to be reached by allowing greater flexibility to explore the solution space.

At each refinement step, we run our improved seam carving algorithm in the visibility map for pixels labelled 0 (non-visible) in the visibility map, and relabel the pixels in the optimal seam to 1 (visible) as illustrated in Fig. 6. We run such a refinement step at each iteration after the seam removal step. We then run another removal step to maintain the current image size. If the energy is decreased, we keep this new labelling proposal. We repeat this until the overall energy is no longer decreased.

Note that by relabelling seams in the visibility map, the property of 8-connected seam removal in the current image is not preserved. However, we did not observe this as causing any lack of pixel consistency in our results.



**Fig. 6.** Refinement by seam carving in the non-visible pixels of the visibility map




Fig. 7. Simple example of linear blending with  $w = 0.25$

## 7.2 Blending

It is also possible to relax the visible/non-visible interpretation of the visibility map. We consider instead an interpretation of a label of 0 as indicating a low weight  $w$  in a blending operation. If pixel  $(r, c)$  is labelled 0, we use a linear interpolation to blend it into its horizontally neighbouring pixels  $\forall x \in \{R, G, B\}$

$$\begin{aligned} I_{r,c-1}^x &= wI_{r,c}^x + (1-w)I_{r,c-1}^x \\ I_{r,c+1}^x &= wI_{r,c}^x + (1-w)I_{r,c+1}^x . \end{aligned} \quad (13)$$

This is illustrated in Fig. 7.

Note that this simple blending operation can be taken into account before calculating the energy terms, and therefore can be optimized for directly. In our experiments where we use blending, we use  $w = 0.25$ .

## 8 Results

We have described a framework for visibility map optimization for retargeting using seam carving operations, and given a range of energy terms and optimization options. We collected a set of 100 images of different kinds of scenes from flickr.com and ran all 288 combinations of these options on them.<sup>6</sup>

We give a selection of the many results we generated in Fig. 8 to demonstrate our findings, and show a larger selection in the supplementary material<sup>7</sup>.

To compare the results, we tried to rank them by average bidirectional similarity [16, 17]. However, we found a poor alignment with human judgement.

**Overall Trends.** We found that the use of distortion energy terms alone, without unary or seam terms, gave poor results. As can be seen from the results in the supplementary material, such an energy favours the creation of few high energy seams over distribution of the error over the image. This was found to be much improved by the use of a unary term. For good results, we found the use of seam terms to be necessary.

With regard to the different optimization options, we found that use of blending reduced the energy in 99.6% of our results, and refinement in 68.2%. Blending

<sup>6</sup> All code is available at [www.vision.ee.ethz.ch/~mansfiea/improvingsc/](http://www.vision.ee.ethz.ch/~mansfiea/improvingsc/) under the GNU General Public License

<sup>7</sup> Also available at [www.vision.ee.ethz.ch/~mansfiea/improvingsc/](http://www.vision.ee.ethz.ch/~mansfiea/improvingsc/)

gave an average energy reduction of 19.3% but refinement an average energy increase of 2.3%. This is possible because the use of refinement only guarantees the energy is the same or lower at each iteration. This result shows that this greediness in some cases leads to an increased energy of the end result. However, for both of these techniques, we found that in practice their effect on the visual appearance of the images was limited.

**Improved Seam Carving.** From our results we chose the following parameters. We use the forward energy seam carving distortion energy (7) with  $n_D = 1$ , unary terms with  $n_U = 1$  and the seam carving-approximating seam term (11). Refinement and blending had little visual effect on our results, so neither are used.

**Representative Results.** A small sample of results from our improved seam carving are shown, with the seam carving and scaling results, in Fig. 8. In these images, cropping would clip interesting areas out. Scaling shows the whole scene, but may include uninteresting areas at the cost of distortion due to non-uniform scaling (see (b)).

The results of seam carving show all the interesting areas, but may include line discontinuities and other distortion (see (a) and (c)). Our improved seam carving distributes seams more evenly in these areas and reduces these artefacts, while maintaining good performance in images where seam carving does well (see (b)). In non-structured images such as landscape seams, all methods perform well (see (d)).

## 9 Conclusions

In this work we introduced the visibility map, which can be used to define retargeting as a binary graph labelling problem. We described a general energy that can be efficiently optimized using seam carving operations. From tests on a large training database, we presented a new, improved seam carving algorithm.

Many works build upon seam carving as a complete algorithm. Our improved seam carving thus may also be used to improve these methods.

However, the improvements we are able to show are relatively minor and do not overcome the major problems of seam carving. Indeed, these problems clearly cannot be solved by simple low level pixel removal methods. Seam carving and related pixel removal methods fundamentally cannot retarget smooth curves to smooth curves. Methods making use of low level information do not know which areas of an image are structured such that non-uniform scaling would be distorting. The limits of such methods seem now to have been reached.

Alternatives do exist, as described in Sect. 2. These have problems of their own, typically in optimization. Nevertheless, it is clear that work using more general image synthesis frameworks, additional images (e.g. video, image databases, stereo cameras), intelligent use of user input and automated feature detection

(e.g. lines, vanishing points, artificial structure) will strongly shape future methods in retargeting. Combining these sophisticated methods with the success of existing simple methods is also a promising direction.

**Acknowledgements** We would like to thank the following users of Flickr for allowing us to use their work under the Creative Commons License: telmo32 for Fig. 8(a), Tambako the Jaguar for (b), Amir K. for (c) and Michal Osmenda for (d).

## References

1. Agarwala, A., Dontcheva, M., Agrawala, M., Drucker, S., Colburn, A., Curless, B., Salesin, D., Cohen, M.: Interactive digital photomontage. In: SIGGRAPH
2. Avidan, S., Shamir, A.: Seam carving for content-aware image resizing. In: SIGGRAPH (2007)
3. Barnes, C., Shechtman, E., Finkelstein, A., Goldman, D.B.: PatchMatch: A randomized correspondence algorithm for structural image editing. In: SIGGRAPH (2009)
4. Chen, B., Sen, P.: Video carving. In: Eurographics Short Papers (2008)
5. Cho, T.S., Butman, M., Avidan, S., Freeman, W.: The patch transform and its applications to image editing. In: CVPR (2008)
6. Dong, W., Zhou, N., Paul, J.C., Zhang, X.: Optimized image resizing using seam carving and scaling. ACM SIGGRAPH (2009)
7. Grundmann, M., Kwatra, V., Han, M., Essa, I.: Discontinuous seam-carving for video retargeting. In: CVPR (2010)
8. Han, D., Wu, X., Sonka, M.: Optimal multiple surfaces searching for video/image resizing - a graph-theoretic approach. In: ICCV (2009)
9. Kim, J.S., Kim, J.H., Kim, C.S.: Adaptive image and video retargeting technique based on fourier analysis. In: CVPR (2009)
10. Krähenbühl, P., Lang, M., Hornung, A., Gross, M.: A system for retargeting of streaming video. In: SIGGRAPH (2009)
11. Mansfield, A., Gehler, P., Van Gool, L., Rother, C.: Scene carving: Scene consistent image retargeting. In: ECCV (2010)
12. Pritch, Y., Kav-Venaki, E., Peleg, S.: Shift-map image editing. In: ICCV (2009)
13. Rother, C., Bordeaux, L., Hamadi, Y., Blake, A.: AutoCollage. In: SIGGRAPH
14. Rother, C., Kumar, S., Kolmogorov, V., Blake, A.: Digital tapestry. In: CVPR (2005)
15. Rubinstein, M., Shamir, A., Avidan, S.: Improved seam carving for video retargeting. In: SIGGRAPH (2008)
16. Rubinstein, M., Shamir, A., Avidan, S.: Multi-operator media retargeting. SIGGRAPH (2009)
17. Simakov, D., Caspi, Y., Shechtman, E., Irani, M.: Summarizing visual data using bidirectional similarity. In: CVPR (2008)
18. Wang, Y.S., Tai, C.L., Sorkine, O., Lee, T.Y.: Optimized scale-and-stretch for image resizing. In: SIGGRAPH Asia (2008)
19. Wolf, L., Guttman, M., Cohen-Or, D.: Non-homogeneous content-driven video-retargeting. In: ICCV (2007)
20. Zhang, X., Hua, G., Zhang, L., Shum, H.Y.: Interest seam image. In: CVPR (2010)

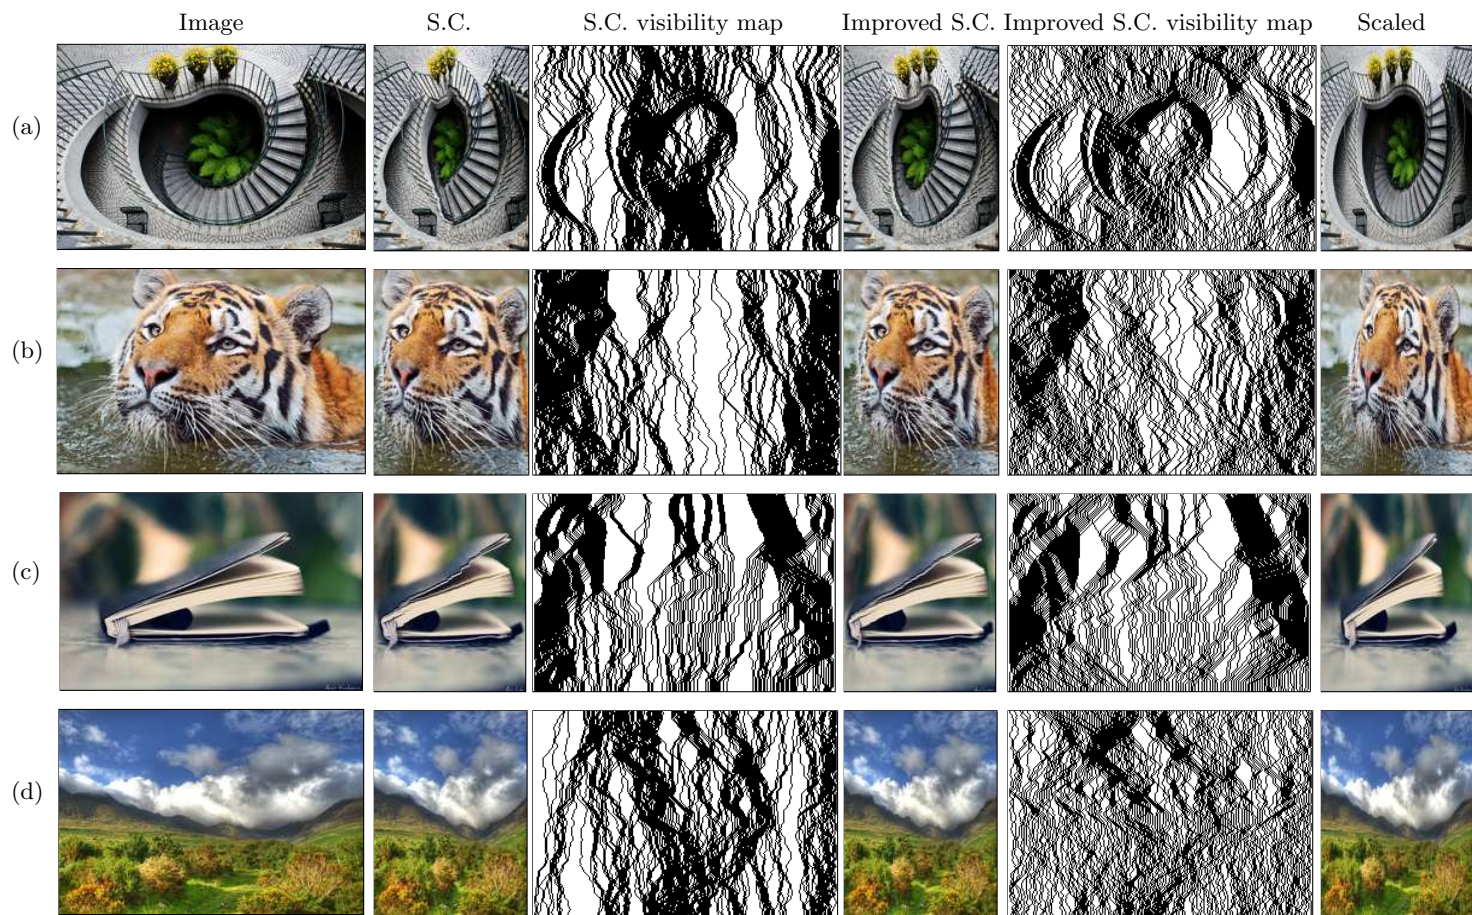


Fig. 8. Results comparing seam carving (S.C.), improved S.C. and scaling for halving image width