

Seam Carving for Enhancing Image Usability on Mobiles

Nicholas Hill
FIT Lab
Computer Science Department
Swansea University
Wales, UK
SA2 8PP
nick@nick-hill.com

Parisa Eslambolchilar
FIT Lab
Computer Science Department
Swansea University
Wales, UK
SA2 8PP
p.eslambolchilar@swansea.ac.uk

ABSTRACT

In this paper we discuss a visual interface for presenting large images on small screen devices. We use the seam carving algorithm as an example technique to improve the readability of images, experimenting with it in various ways. We explain the results of a pilot study where we examined the use of seam carving algorithms on a small screen, making comparisons with simple other image processing methods. This is the first example of a visual interface which uses the seam carving algorithm for mobile devices. We discuss scenarios for the use of such techniques.

Categories and Subject Descriptors

H.1.2 [User/Machine Systems]: Human factors, Human information processing.

General Terms

Experimentation, Human Factors.

Keywords

Human-machine Interaction, Image processing, seam carving, small screen devices

1. INTRODUCTION

Increasingly, people are capturing their lives to present their emotions, plans and actions with their mobile phone camera – to remember or share their joy and experience. Since availability and capacity of flash memory cards for mobile devices have increased dramatically in recent years, many people store their photos on the phone. However, visual interfaces for browsing photos have crucial limitations in small screen devices. These devices have a limited amount of screen space on which to display information. Designing interfaces for mobile computers is problematic as there is a very limited amount of screen resource on which to display information and users' eyes are often needed on the environment rather than the interface (i.e. that they can look where they are going), therefore, output is limited [1,3]. Also, low graphics resolution and few colours in these devices do not help designers to design complicated visual interfaces.

Reducing the size of an image without losing the important information in the image is an ongoing research and it becomes

more challenging on handhelds. One way around these problems would be seam carving that resizes the image while preserving its original energy as much as possible [2]. One advantage of using this technique is that users would not have to use any buttons to view the image, such as zooming in and out of the image, and could instead view the most important parts of the image as an accurate overview. As Shneiderman [8] has stated there are many benefits to mixing overview and context together; the context is the same size as the overview.



Figure 1. A picture with the lowest energy seams through the image (shown in red).

2. SEAM CARVING

Avidan and Shamir's Seam Carving [2] involves a dynamic programming algorithm that runs on an image energy map and allows images to be resized while preserving as much of its original energy as possible. This is depicted in Figure 2. One example of an energy map involves simple gradient magnitude.

In order to reduce the size of an image in a direction by one pixel, a connected series of neighbouring pixels are calculated from one end of the image to the other. This seam is the shortest energy path through the image. The pixels that make up the line are removed and the image is shrunk in the required direction. For example, Figure 1 shows the lowest cost seams in both the horizontal and vertical directions, which are then removed to shrink the image.

Seam carving is repeatedly performed until the desired size is reached. Although the final image is smaller, it retains the largest possible amount of energy and it is still possible to identify features in the image.

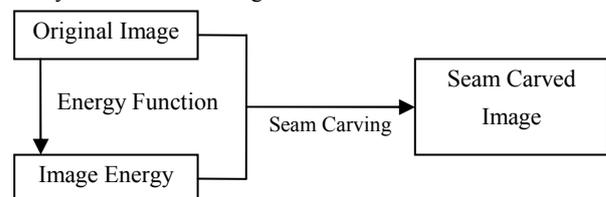


Figure 2. The process of seam carving an image.

3. ALGORITHM EXPERIMENTATION

The seam carving algorithm was implemented as an exercise prior to conducting a pilot study. The idea was to implement seam carving as naively as possible initially and then move on to implementing a dynamic version of the algorithm.

3.1 Straight Line Seam Carving

The simplest implementation of the seam carving algorithm for images involves summing the energies along the columns and / or rows of the image and then removing those columns or rows that have the least total energy. This is depicted in Figure 3. The algorithm does not take very long to implement, and executes almost instantly. This is because each pixel only needs to be considered once.

3.2 Random Seam Carving

The next step involves allowing the lines to be nonparallel and to deviate around high energy regions in an image. Random strategies were used for this step. The algorithms were easy to implement and quick to execute, especially when compared to the run time of the dynamic programming solution.

The simplest random strategy involves performing several completely random ‘walks’ on an image, accumulating energy values as the path is composed. After several random paths have been calculated, the path with the least total energy is selected and the pixels that lie in it are removed. The number of random seams calculated prior to selection can be varied.

An alternative is to calculate some random starting points for seams, and then carve a path intelligently by selecting the least energy value in the three possible pixels on the next row or column.

10	1	7	8	7	10	1	7	8	7
10	12	2	8	8	10	12	2	8	8
7	10	8	1	8	7	10	8	1	8
9	11	3	10	12	9	11	3	10	12
9	1	9	9	10	9	1	9	9	10
45	35	29	36	45					

Figure 3. Seams calculated for simple greyscale images using straight line and dynamic seam carving respectively.

3.3 Dynamic Seam Carving

The final step involved implementing seam carving using dynamic programming. The use of dynamic programming [13] for calculating the shortest paths in this graph takes advantage of the nature of the overlapping subproblems and optimal substructure present. Images carved using this technique retain as much of the image’s original energy as possible. An example seam is shown in Figure 3.

3.4 Discussion

By examining the results in Figure 4, it appears that straight-line seam carving is not well suited to photographic images. It may be well suited to simple monochrome drawings containing large empty spaces, but even then, an alternative seam carving method is still expected to perform better.

Although a significant amount of energy has been preserved using random seam carving, there is clearly some distortion. For example, the buildings look distorted. Increasing the number of random seams per column is likely to improve the quality of a final, seam carved image.

Images created with intelligent random seam carving are almost as good, and in some cases better than dynamic seam-carved images. Unfortunately, drawbacks include the fact that selected random seams can still chop off high energy parts of an image due to the nature of the algorithm

The maximum amount of energy has been preserved in the image carved using dynamic programming, although the buildings on the left part of the carved image look quite

distorted. It is possible that these images could be improved by using a better energy function.

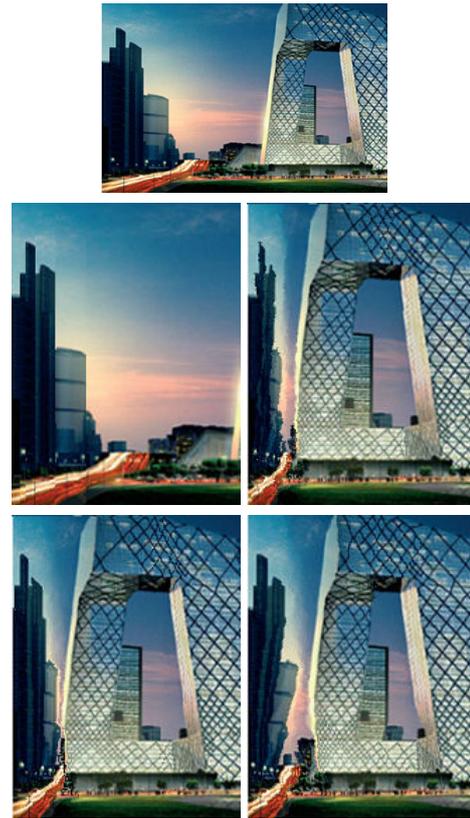


Figure 4. An original image carved using straight lines (top left), random (top right), intelligent random (bottom left) and dynamic (bottom right) seam carving algorithms.

It may be more beneficial to implement intelligent random seam carving on mobile devices rather than dynamic seam carving, because it runs quickly (even with relatively few seams).

4. PILOT STUDY

A pilot study involving the use of prototype software was conducted. It was designed to examine how users felt about viewing large images that had been adapted for the small screens on mobile devices.

4.1 Method

A simple image browsing program was developed in Visual Basic .NET for the pilot study. It was designed to present users with different versions of a selected image.

The main form comprises a blank space of 240 pixels wide and 320 pixels high. This is similar to many small screen devices. On the bottom, there are Previous and Next buttons allowing users to scroll between four different presentations, or *variants* of a selected image.

When a new image is selected, the original, unmodified version of it is displayed in a new window that must be dismissed prior to users interacting further with the application. Users can then use the previous and next buttons to move between the four variants of the image.

The four types of image were named as ‘cropped’, ‘resized’, ‘seam carved’ and ‘smart seam carved’:



Figure 5. The main interface of the software used for the pilot study.

Cropped: The centre of the image relates to the centre of the small screen and any remaining part of the image appears outside of the region is then deleted, or cropped.

Resized: Images are resized and forced to fit within the viewing region in both dimensions, utilising the full extents of the screen.

Seam Carved: In seam carved images, images are resized in both directions using seam carving so that they fit the viewing area whilst losing the least amount of image energy possible.

Smart Seam Carved: An adjustment to simple seam carving where the image is first resized to fit one dimension, and then the remainder of the image is seam carved to fit the other. Smart seam carving was considered because of an inherent problem with naïve seam carving. The results of the algorithm depend on the original size of the image.

Smart seam carving was employed in the pilot study to determine whether or not users prefer a hybrid of naïve resizing and seam carving. That is, if they preferred smart seam carved images over seam carved images. Seam carved and smart seam carved versions of images were prepared with IntuImage [9] and Paint .NET [14].

Six images were prepared for users to look at, some well suited to seam carving and others not so well suited. Each of the original images was larger than the 240x320 resolution small screen display, and consisted of a photo of a beach, park, baby, university, field and map.

The beach photo was the best candidate for seam carving, while the baby image was thought to be the least desirable in seam carved form and unsolicited feedback was expected. The university and park images contained high energy regions throughout.

Cropped images were assumed to be the most desirable when the main object(s) of focus were already in the centre of the image, and least desirable when images contain significant features outside of the viewing region.

Resized images were thought to be the most desirable when the image typically contains few large objects that would scale well, such as a face. They were thought to be the least desirable when the images contain many small features that do not scale well, and can in fact be lost when the image is resized.

Seam carved images were thought to be the most desirable when the original images contain large regions of little energy that can be removed without altering the context of the image

too much. They were thought to be the least desirable for images containing large areas of energy throughout the image.



Figure 6. Cropped (top left), resized (top right), seam carved (bottom left) and smart seam carved (bottom right) versions of the beach image.

It was generally assumed that users would be capable of appreciating the benefits that seam carving brings to resizing large images.

4.2 Procedure

The pilot study involved the use of a questionnaire primarily. A total of six participants were asked to discuss and rate the six images presented by the software. Of the six participants, three were female and three were male. Each of the participants were in their 20s and had owned at least one photographic-capable mobile device. Experience working with images on their mobile devices ranged between a few times a week to a few times every few months. Some users owned digital cameras, which often also contain small screens. Each session typically took twenty minutes to complete.

The pilot study was conducted in two phases, the *picture viewing* phase and the *post viewing* phase. In the *picture viewing* phase, users were asked to look at different versions of each of the images used in this study. Each user was first shown the unmodified image in its original size. Once they had examined the image, it became unavailable and the user was allowed to move forwards and backwards between the four alternative versions of the image.

Users were advised to examine the readability of the image, not necessarily to consider the aesthetic quality, or 'niceness' of the image as a primary measurement to its level of suitability. Users were also instructed on what seam carving is and what it was designed to do.

Users were then asked to answer a few questions about the variants of each image before moving on to the next image. Non-qualitative questions involved users rating the seam carved images over other types, from a scale of 0 to 10.

The *post viewing* phase occurred after the user had finished looking at and answered questions about all of the images, and users were asked to answer some further general questions about seam carving and viewing images on small screens, which

aimed to gather opinions that the user may have built during the initial phase.

4.3 Results and Discussion

As expected, participants strongly indicated a preference to the seam carved versions of the beach image over both the cropped and resized versions, often giving the maximum score of ten. This was because the beach image was the best suited to seam carving. Conversely, ratings for the seam carved baby image over both the cropped and resized versions never peaked above a score of two, indicating every participant's strong disapproval of the use of seam carving for faces. One participant refused to choose between seam carving and smart seam carving because she did not want to look at the images more than once in order to determine a preference. Interestingly, this negative response was the most prevalent amongst each of the females, with each of them making a direct negative comment about the photo. As a conclusion, it seems to be the case that it is absolutely vital to 'protect' faces as much as possible from seams since they consist of regions of low energy and are carved by the seam carving algorithm, resulting in severely distorted facial features. This was already considered by Avidan and Shamir.

Other images caused mixed results. For example, users were generally undecided over whether or not seam carving worked well on the images of the park, university, field or world map. The scores for the map image are surprising, as it was hoped that people would appreciate the benefits of seam carving on schematic drawings, but participants were undecided on this. A possible reason is that participants treated the shape of each country as more important than simply being able to see and know where and what each of the countries were, even if they were somehow made to exist in a small space.

Seam carved photos were preferred over smart seam carved photos only three out of 29 times, meaning that participants had a very strong preference for smart seam carved versions of an image over seam carved versions. Users wished to view images in the correct aspect ratio, even if valuable space on the viewing area is lost. There was an appreciation that cropped images remove important details, such as is the case for the map image. There are many cases where users found the cropped versions acceptable, and this was typically when the images were only slightly larger than the screen size itself or when the main features in the image were in the middle and therefore not cut off in the cropping operation.

Reflecting on these discussions, it seems that it is not the best thing to seam carve every image prior to showing it to the user. Instead, several facts about an image should be determined prior to showing an image to a user using a particular mechanism. For example, potential criteria include the original size and type of the image, the total amount of energy it contains, the locations of the important features in the image and its aspect ratio. If these facts are determined before showing an image, it would be possible to select the best visualisation. For example, in some cases it would be more beneficial to show users a cropped version of an image if most of the energy is concentrated around a particular point and in other cases it would be more beneficial to show users a resized version of the image if most of its energy is evenly spread throughout.

It may be possible to determine these facts automatically with some additional algorithms, and then automatically select the best visualisation and present that variant of the desired image to the user. The benefit to this is that no user interaction would be required.

5. CONCLUSIONS

Seam carving is a useful algorithm for resizing images containing large areas of low energy and small areas of high energy, but there are challenges. Although participants agreed that it can work quite effectively for some photos, there are many cases where seam carving performs poorly in communicating the content of an image to the user. It is important to determine *how* to present an image to a user rather than simply seam carving every one. This determination should be done as automatically as possible. For example, a checking algorithm could run on an image and make a decision about what to do to the image before presenting it.

6. FURTHER WORK

Further work will involve finding out whether or not it is possible to automatically determine the best way to present an image to the user, and in finding out whether or not the users prefer to view these automatically modified versions (which require no interaction) over the originals (which require interaction such as panning or zooming).

7. ACKNOWLEDGEMENTS

The authors gratefully acknowledge the support of FIT Lab and Computer Science Department and students in Swansea University who participated in the user study.

8. REFERENCES

- [1] Brewster, S. A. and Murray, R. (2000). Presenting dynamic information on mobile computers. In *Personal Technologies*, pages 209--212. Springer-Verlag.
- [2] Shai Avidan, Ariel Shamir. Seam Carving for Content-Aware Image Resizing, ACM SIGGRAPH 2007, pages 1-10.
- [3] Matt Jones and Gary Marsden. Mobile Interaction Design, Chapter 10. Wiley (2006) ISBN: 0-470-09089-8.
- [4] Nicholas Hill. Usable Images on Mobile Devices – Project Proposal. December 2007, Swansea University.
- [5] Rafael C. Gonzalez, Richard E. Woods. Digital Image Processing, Third Edition. ISBN 0-13-168728-X.
- [6] Feng Liu and Michael Gleicher. Automatic Image Retargeting with Fisheye-View Warping. ACM UIST 2005, pages 153-162.
- [7] Hao Liu, Xing Xie, Wei-Ying Ma, Hong-Jiang Zhang. Automatic Browsing of Large Pictures on Mobile Devices. ACM 2003.
- [8] Ben Shneiderman. The Eyes Have It: A Task by Data Type Taxonomy for Information Visualizations. HCI laboratory, University of Maryland, IEEE 1996, pages 336-343.
- [9] IntuImage – Free Seam Carving Software version 0.3. <http://www.intuimage.com/>.
- [10] Derek Rowntree. Statistics Without Tears – An Introduction for Non-Mathematicians. Penguin, ISBN 978-0-14-013632-6.
- [11] Lance Good, Benjamin B. Bederson, Mark Stefik, Patrick Baudisch. Automatic Text Reduction for Changing Size Constraints. ACM CHI 2002 Short Talk.
- [12] Huiying Liu, Shuqiang Jiang, Qingming Huang, Changsheng Xu, Wen Gao. Region-Based Visual Attention Analysis with its Application in Image Browsing on Small Displays. ACM 2007.
- [13] Richard Bellman. Dynamic Programming. 1957, Princeton University Press Dover paperback edition (2003), ISBN 0486428095.
- [14] Paint .NET. <http://www.getpaint.net>