Predicting when seam carved images become unrecognizable

Sam Cunningham

April 29, 2008

Acknowledgements

I would like to thank my advisors, Shriram Krishnamurthi and Michael Tarr for all of their help along the way. You not only provided motivation for my progress, but were instrumental in helping target my research, design the experiment, calculate the results, and prepare for my presentation. Additionally, I would like to thank Cesalie Stepney for her help in performing statistical calculations and editing this document. Finally, thank you to all of the participants in my experiment.

Contents

1 Abstract

As the content of the Internet is brought to mobile devices with small screens of varying aspect ratios, a technology is needed to resize images to fit these displays. Seam carving is one such technology. It uses the content of an image to remove thin strips of pixels, called *seams*, to both shrink images and change their aspect ratios. One difficulty with the technique, however, is that as seam carving progresses, distortions accumulate that eventually make the image unrecognizable. I proposed an experiment to compare seam carving to other techniques for resizing images, by determining the size at which an image ceased to be easily recognizable. I then proposed an algorithm to automatically predict this size, and used the results from the experiment to evaluate that algorithm. The data from this study suggested that by automatically detecting when an image becomes unrecognizable, it is possible to shrink the width of an image using seam carving until this occurs, and then continue using uniform scaling. This method produced a superior result over seam carving alone.

2 Introduction

Seam carving is a method for resizing images, with consideration of their content, to adapt them to displays of varying sizes and proportions. In the process, small distortions may result from changes to aspect ratio or information that is removed in order to fit the image within the desired dimensions. The method iterates over the image, shrinking it by either one row or column at a time, so that these small distortions accumulate and eventually make the image unrecognizable. I sought to determine when this change occurs, and whether it is possible to automatically judge when a seam carved image becomes unrecognizable, either across all humans or on a per-human basis. This information could potentially allow algorithms to combine seam carving with less disruptive techniques such as seam carving in order to resize and re-proportion images without losing recognizability.

Applications of seam carving are numerous, but let us consider one example. As the content of the Internet begins to appear on mobile devices, such as the Apple iPhone (Apple Computer, Cupertino, CA), web browsers must adapt images designed for desktop monitors to much smaller displays. Doing so poses two challenges: first, mobile displays have much smaller dimensions than their desktop counterparts, and second, the aspect ratios of a display and of those images may differ. An intelligent solution will not only shrink the image, but also fill the entire display rather than leave some portions blank, and will not arbitrarily crop the image, which can potentially lose important information. Seam carving offers a solution to both these problems by considering the content of an image while adapting it to an arbitrary size and width.

Seam carving achieves this goal by shrinking the size of an image by one row or column at a time. It then iterates until it reaches the desired dimensions. At each step, seam carving calculates the energy of every individual pixel using an *energy function*, usually the image gradient. Next, it calculates the minimum weight seam in the image and removes it. A *seam* is a chain of pixels, one from each row (to shrink the width) or column (to shrink the height), in which each pixel either neighbors two other links or is at the edge of the image. The *weight* of a seam is the sum of the energies at each pixel it contains. The authors of the seam carving technique, Avidan and Shamir (2007), suggested that by removing seams, seam carving dramatically increases the average energy across the image. This is very different from uniform scaling, where shrinking the image has little effect on the average energy.

Distortions resulting from seam carving are inevitable, both because the algorithm alters the aspect ratio of the image and because it removes information. Fortunately, small distortions are not always significant. Most images on the web serve only to convey a particular idea, which these small distortions do not interrupt. However, large disruptions will occur if images are seamed too far, causing subjects in the images, such as people and faces, to become misshapen and not easily recognizable. The degree of this distortion varies greatly between images, and is highly dependent upon their content. For instance, images that are highly condensed become more distorted than other types of images (Avidan and Shamir, 2007). These types of images may be more successfully resized by uniform scaling. One way to distinguish between these is for the author of a web page to insert a tag into the HTML code that specifies the resizing method to perform. Another is by detecting this automatically.

This experiment examined whether there exists a certain point at which a seam carved image become unrecognizable. Additionally, it related this point to the average energy across the image. The results showed that such a correlation exists, and subsequently, that the average energy is useful in predicting this point. I proposed both an experiment to determine when an image becomes unrecognizable and a prediction tool which uses the average energy across the image to predict the point at which this change occurs.

3 Methods

3.1 Experimental setup

Participants.

The participants in this study were all undergraduate and graduate-level students at Brown University, ages 20 to 30, and 22.7 years old on average. Twelve of the students were female (60%), and eight male (40%). This age demographic represents some of the primary users of the Internet, and is therefore very applicable to many target applications of seam carving. Students all were given informed consent per Brown IRB guidelines and were compensated for their participation.

Stimuli.

Internet-based applications also motivated the selection of test images, which were taken from the New York Times website (`http://www.nytimes.com`). Fifty images were selected from many different sections of the site and included images of people, portraits, landscapes, and some cartoons. The dimensions of these images varied, though most were approximately 600 pixels wide and 300-400 pixels tall. The smallest image was 190 pixels wide by 240 pixels tall, and the largest was 575 pixels wide by 500 pixels tall.

Figure 1: Screenshot of the experimental window

Design.

Because seam carving is too computationally intensive to be performed in real time, images were carved in advance and the results were compiled into a high-quality movie. Participants used a slider to display a particular frame from the movie; positions towards the left selected wider images while positions towards the right selected thinner images (Figure 1). Uniform scaling was performed in real time. All images were displayed on a web page using the Apple QuickTime plug-in (Apple Computer, Cupertino, CA).

Each of the fifty images was displayed three times using techniques of seam carving, scaling and a combination of both. The order of the images and of the techniques was randomized; however, the combination technique always came after seam carving. The starting position of the slider was also randomized uniformly across the second and third quartiles of the available widths. This was intended to reduce bias by the subjects that may be caused by viewing the entire image first. Since applications of the technique may not begin with the original image, this was considered a more accurate testing method. Students were then instructed, “Adjust the movie until the image is as small as possible while still being easily recognizable. With a quick glance, you should be able to get the gist of the photo without being distracted by misshapen elements.” I called the width at which this occurs the *fail width* of an image.

I decided that only the width of the image would be seamed, in order to change the aspect ratio and therefore necessitate some distortion to the image. In seam carving, this distortion results from the removal of seams, while in uniform scaling, objects appear elongated. The decision also focuses the problem at hand on changing the aspect ratio to match that of a display, rather than simply reducing the size of the image.

3.2 Seam carving

The seam carving implementation used for this experiment was designed to be similar to that described by Avidan and Shamir (2007). The energy function was a measure of the gradient at each individual pixel (Efford, 2000, 164). The image was converted to the perceptually uniform LAB color space (Durrett, 1987, 96-9). Then, the vertical and horizontal gradients were determined using Prewitt kernels (Efford, 2000, 164-5) over a 3×3 pixel neighborhood for each color channel and summed together. The absolute values of each gradient were added together to create the final energy value.

A seam was considered to be a chain of individual pixels, one from each row of the image. Linking pixels were restricted to those within a one-pixel radius of each other. The minimum seam across the image was computed using dynamic programming, building seams from one side of the image to the other. The minimum seam was then removed. In the case that multiple seams shared a minimum weight, one of these options was selected at random. This avoided bias towards seams from a particular section of the image. To remove subsequent seams, additional iterations were performed and both the gradient and seams were recomputed each time.

This seam carving implementation did not employ any smoothing technique after seaming the image. Avidan and Shamir (2007) suggested that seam carving could be used in combination with a Poisson reconstruction to reduce the sharp contrast observed where seams had been removed. For the purposes of this experiment, I determined that smoothing would not have a significant effect on the recognizability of the images.

4 Results

4.1 Comparing techniques: seam carving vs. uniform scaling

The results of this experiment provided a fail width, the smallest width at which the image was easily recognizable, for each image with each of the three techniques: seam carving, uniform scaling, and a combination of both. The average width of the original images was 576.7 pixels, and the average fail width for seam carving was 311.4 pixels (a 46.0% reduction), $sd = 151.0$, 201.5 pixels for uniform scaling (a 65.1% reduction), $sd = 98.7$, and 204.9 pixels for the combination technique (a 64.5% reduction), $sd = 102.8$ (Table 1). The deviation in these results was remarked to be relatively large. This suggests that images tend to become unrecognizable very gradually, and that an algorithm attempting to predict this point may be able to return many acceptable values.

Based on the results, uniform scaling of an image had a clear advantage over simple seam carving, evidenced by the fact that images remained easily recognizable at smaller widths when calculated with scaling than with seam carving. To determine the statistical significance of the results, a matched-pair within subjects t-test was computed between the data from seam carving and from uniform scaling, $t(19) = -10.58, p < .001$. However, in the test using the combined method, subjects were able to continue reducing the widths of already maximally seamed images by using uniform scaling, producing a significant advantage over seam carving alone $t(19) = -8.05, p < .001$. This final result was not significantly different from uniform scaling alone, $t(19) = 0.53, p = .603$.

To determine the cause of the deviation in the results, the standard deviation on a per-subject basis was computed for seam carving. Participants' answers for each image were computed as a percentage of the average fail width. For instance, if an image had an average fail width of 400 pixels but a particular participant determined that the fail width was 350 pixels, the answer is recorded as $350/400 = 87.5\%$. The average standard deviation in these percentages was 23.0% for seam carving. This deviation is far less than the 39.9% observed across the group, indicating that subjects had individual preferences or definitions of the recognizability of an image that were consistently either smaller or larger than the average for the general population. This suggested that it may be possible to tailor a prediction of the fail width to individual users for optimal performance.

4.2 Predicting the usefulness of seam carving

The goal of this experiment was to determine whether the average energy across an image could be used to determine when an image undergoing seam carving was no longer recognizable by human observers. This study considered the average energy across the image, the change in the average energy during seam carving, and the percentage change as three possible measures of the images through which this point could be determined. Evidence from the seam carved images showed that as an image's width shrunk, the average energy across the image increased continuously. Therefore, given a threshold value, it was possible to stop shrinking the size of the images at the last width for which the average energy (or change in average energy, percent change, etc.) was less than or equal to the threshold. This was the *predicted fail width*.

I considered the accuracy of a particular threshold the number of images for which the fail width was greater than or equal to the average result of the human experimental data minus the standard deviation for that image. I calculated the residual by averaging the square of the difference between each predicted fail width and the corresponding empirical fail width, and then taking the square root. I decided that the optimum threshold value would be the value for which the average residual was minimized and the accuracy was approximately 98%. This accuracy would indicate that only one image out of fifty was too small to be easily recognizable within the range of standard deviation.

The most successful threshold I obtained to match the entire set of images and empirical data was when the percent change in image width was used as a measure. This value was 43.67%, with an accuracy of 94.0% and average residual of 79.31 pixels. This average residual was significantly less than the standard deviation of the experimental results, and represented a viable result for real-world applications. The next-best threshold was when change in average energy over the image was used as a measure, with the threshold set at 10.22, providing an accuracy of 96.0% and average residual of 83.08 pixels. Thresholds for several other measures appear in Table 2, and the graph of their residuals appears in Figure 2.

I also validated the prediction method using a jackknife procedure (Agency, 2008). Rather than using the results from all images and all participants, I withheld one image, subject, or both at a time and then optimized the threshold. Finally, I averaged together the thresholds, accuracy and residuals obtained. This technique ensured that the thresholds would generalize better to images and participants besides those included in this experiment. The statistics obtained confirm that the percent change in image width and change in average energy over the image provide very good predictors of the fail width.

5 Discussion

Automated seam carving is vulnerable to removing too much of an image or distorting subjects such that they become unrecognizable. Results from this experiment showed that distortion from seam carving prevented images from being easily recognizable far sooner than uniformly scaled images. It was therefore useful to know when this distortion began to occur, and to modify the algorithm by switching to uniform scaling. Doing so enabled the images to remain recognizable even at much smaller widths.

Furthermore, this experiment showed that either the percent reduction in image width or the change in average energy over a seam carved image was a useful predictor of when an image started to become unrecognizable, which I called the *fail width*. I have proposed an algorithm which performs seam carving on the image until the change in average energy across the image from the original exceeds a threshold value. Data from this experiment showed that this algorithm can provide a prediction that is well within the range of acceptable values.

6 Limitations

This experiment focused on seam carving using a specific algorithm to carve away columns from a limited set of images. This makes the results very useful for the applications described, but potentially less useful in others. The choice of algorithm, for instance, could be tailored to certain tasks by recognizing particular shapes in an image and creating a mask to seam carve around those objects. In this case, the specific results from this experiment may not be able to describe the new algorithm.

The images used in this experiment were also very similar in size, shape, and composition, which may have led to some bias in the results. While the results are valuable when used with many images found on the Internet, it may be practical to repeat the experiment for larger, full-resolution images. Despite these limitations with the actual results, the techniques used in this experiment are highly flexible. Few changes would be needed to use this experiment with other algorithms or images.

7 Additional Research

7.1 Alternative energy functions

One convenient property of seam carving is that the energy function is flexible and can be modified to improve the way images are seam carved, or to combine multiple elements. Avidan and Shamir (2007), for instance, tested a saliency measure, Harris-corners measure, and eye gaze measurement as replacements for the standard gradient measure. They also suggested that particular objects can be painted, manually or automatically, with higher or lower energies in order to preserve them or remove them from a scene.

In order to evaluate whether one energy function is "better" or "worse" than another, the methods of this experiment can be performed with seam carved images that use different energy functions. Instead of comparing seam carving to uniform scaling, the experiment can compare one type of seam carving to another. One function is judged "better" than another if the average fail width is lower. Such a test could also reveal whether a particular energy function works better for some types of images and worse for others.

Given time constraints, my own evaluations are based on a modification to this technique. I compared each energy function against a baseline — the seam carved images used in the experiment (LAB gradient energy measure), at the average empirical fail width. These baseline images were compared side-by-side against images that were seam carved to the same widths using the new energy function. I then judged whether each test image was more or less easily recognizable than the baseline. The energy function received a score: the number of images that were better than the baseline minus the number that were worse. A positive number (maximum 50) indicated that the new method was better than the baseline, and a negative number (minimum -50) indicated that it was worse. This test was designed such that the subjective judgment calls would be unlikely to vary between participants, and it would therefore provide meaningful information even with a single observer.

I performed this test with two different energy functions: smoothing the image prior to calculating the LAB gradient, and image saliency. The smoothing technique used a 5×5 pixel Gaussian filter, $\sigma = 5$ (Efford, 2000, 155-6), to transform images before calculating the gradient. The method earned a score of -2, suggesting that the results were slightly less recognizable than the baseline. The image saliency energy function achieved a score of -16, indicating that it performed worse than the baseline. Selections from these results are shown in Figure 3.

7.2 Color space and the gradient

The choice to calculate the image gradient using a LAB color space was motivated by the desire to reflect the differences between pixels as humans perceive them. This incorporated color information into the calculations as well as other factors. However, gradients are often calculated by first transforming the image into grayscale. To compare seam carving with these two different gradients, I repeated the simple experiment above in order to compare the two gradients to each other.

The grayscale gradient earned a score of 8, indicating that it performed somewhat better than the LAB gradient. This suggested that luminosity differences between different subjects in the images were more significant than color differences. A selection of these results are shown in Figure 4.

7.3 Springs and jumps

Seam carving stipulates that a seam consists of one pixel from each column or row of an image and that those pixels neighbor one another. This has a computational advantage in that it limits the total number of seams in an image and thus reduces the number of computations that need to be performed. However, seams may achieve smaller total energies if they are allowed to "spring" to pixels that are further away, or "jump" if the difference between two neighboring pixels exceeds a certain threshold.

I implemented springs in seam carving by attaching a weight to the spring length. Springing to a non-neighboring cell increased the energy proportionally with the square of the spring length, thus favoring smaller springs. In my implementation, springs were given a weight of 0.056, meaning that a spring of length x would have a weight of $0.056 \cdot (x^2)$ plus the energy at that pixel. I observed that when seam carving images with springs, the seams removed were often clustered in one particular region of the image, more so than without springs. This caused lots of disruption in that part of the image. Images tended to be equivalent to or worse than the baseline seam carved images.

Jumping allows seams to jump to a different part of the image if the energy at all of the three neighboring pixels exceeds a certain threshold. In my own experiments, I found that 0.5 was the optimum threshold. Even at this level, however, shifts of large parts of the image occurred at each jump. For instance, if a seam ran down the left side of an image until it hit an object and then jumped to the right side, all of the pixels in between would shift against each other. This created large disruptions and made all of the images worse in comparison to the regular seam carving. Examples of both springs and jumps are shown in Figure 5.

7.4 Face detection

One suggestion that Avidan and Shamir (2007) make is that seam carving can be used in conjunction with a mask. A person could, for instance, manually paint over the important objects in an image, and only the seams around that image would be removed. The same could be done with some automated methods, such as face detection. I used the OpenCV face detection implementation in order to create a mask over faces, and then performed seam carving normally (Library, 2008).

Face detection earned a score of -1, suggesting that there was little difference between this seam carving implementation and the baseline. However, viewing the images revealed many cases in which face detection did provide a significant gain, shown in Figure 6. In many cases, the face detection algorithm detected not only faces, but also returned many false positives, and it was in these images that the seam carving became worse. Additionally, it would be useful for the detector to return an outline of the face, rather than a circle. These flaws in the algorithm prevent seams from passing through some unimportant areas where they otherwise would. If improvements to the detection algorithm can be made, it will become more similar to manually painting over the subjects and may yet provide an advantage in seam carving.

8 Conclusions and Future Work

The two experiments described here showed how seam carved images, calculated using various techniques, can be compared against one another and against other resizing techniques. The first experiment also helps us determine when images cease to be easily recognizable. This information is useful in switching between methods in order to improve the overall result.

This work suggested that the seam carving technique could be improved upon in various ways. First, stopping seam carving at the fail width and switching to uniform scaling provided a significant advantage over seam carving alone. Also, seam carving images with a grayscale gradient offered a better result than a LAB color space gradient. Other modifications to the energy function may garner additional improvements.

One way to improve the prediction of the fail width may be to incorporate multiple measures of the image into the calculation. For example, rather than using the average energy alone, this measure may be more useful when considered in conjunction with the minimum energy. Also, the prediction may be improved by tailoring the results to particular users.

It is also curious that the combination of seam carving and uniform scaling provided a less satisfactory result than uniform scaling alone. Perhaps the instructions led subjects to seam the images too much. It would seem that the combination results should be at least as good as the uniform scaling results, never worse, because in the worst case seam carving can be stopped at the very first frame and uniform scaling performed from there. So, rather than asking subjects to seam carve the images as far as possible while the image remains, "easily recognizable," perhaps they should be instructed to seam carve them as far as possible while it does not distort any portion of the image necessary for its comprehension.

Using this reasoning, if a seam is removed from one of the principle subjects, seam carving should stop immediately rather than continuing. If this does provide better results, then combining seam carving with automatic detection of the subjects should provide substantial gains. Face detection, for example, could be used to find the people which are the primary focus of the image, and prevent any seams from passing through them. If the detection algorithm can be perfected, the benefit to seam carving should be significant.

With improvements in automatically seam carving images, the technique will become more and more useful in applications, including the resizing of images to fit displays of different aspect ratios. Experimental methods such as those described here will continue to serve as important measures of how successful one algorithm is compared with another.

References

U.S. Environmental Protection Agency. Jackknife — statistical primer, 2008. URL `http://www.epa.gov/bioindicators/statprimer/jackknife.html`. [Online; accessed 22-April-2008].

Shai Avidan and Ariel Shamir. Seam carving for content-aware image resizing. *ACM Transactions on Graphics*, 26(3), 2007. URL `http://www.faculty.idc.ac.il/arik/imret.pdf`. [Online; accessed 22-April-2008].

H. John Durrett, editor. *Color and the Computer*. Academic Press, Inc., Orlando, FL, 1987.

Nick Efford. *Digital Image Processing*. Pearson Education, Ltd., Harlow, England, 2000.

Open Source Computer Vision Library. Face detection using opencv, 2008. URL `http://opencvlibrary.sourceforge.net/FaceDetection`. [Online; accessed 22-April-2008].

A Appendix

	Seam Carving			Uniform Scaling			Combination		
Image	**Average**	**Std Dev**	**Std Error**	**Average**	**Std Dev**	**Std Error**	**Average**	**Std Dev**	**Std Error**
1	401.3	123.35	27.58	214.6	99.17	22.17	218.7	109.29	24.44
2	302.0	137.80	30.81	223.0	82.41	18.43	210.6	83.66	18.71
3	220.6	122.67	27.43	198.9	108.89	24.35	153.5	64.36	14.39
4	169.8	57.82	12.93	153.2	85.01	19.01	125.1	52.63	11.77
5	198.2	104.15	23.29	224.6	97.82	21.87	173.5	81.54	18.23
6	200.7	77.91	17.42	170.3	65.68	14.69	156.6	74.69	16.70
7	112.2	45.07	10.08	73.2	36.49	8.16	77.3	38.18	8.54
8	412.4	127.39	28.49	242.0	88.74	19.84	253.3	92.60	20.71
9	299.0	145.65	32.57	182.0	79.82	17.85	210.7	90.74	20.29
10	315.4	124.60	27.86	210.3	95.29	21.31	230.4	110.45	24.70
11	366.6	138.20	30.90	199.8	72.08	16.12	238.3	95.10	21.26
12	224.3	117.86	26.35	162.0	57.89	12.95	149.4	62.37	13.95
13	381.0	157.99	35.33	191.2	91.41	20.44	216.8	93.18	20.84
14	356.4	129.52	28.96	242.6	108.67	24.30	199.2	82.20	18.38
15	287.1	136.56	30.54	237.6	116.77	26.11	198.4	74.34	16.62
16	283.0	122.30	27.35	228.5	90.39	20.21	219.0	101.89	22.78
17	231.6	115.82	25.90	199.4	93.26	20.85	167.7	69.13	15.46
18	329.1	126.85	28.36	208.8	117.05	26.17	224.9	95.53	21.36
19	258.2	83.57	18.69	237.3	98.18	21.95	180.4	64.65	14.46
20	357.4	149.48	33.43	204.3	113.73	25.43	240.8	119.44	26.71
21	444.7	136.16	30.45	182.1	83.75	18.73	228.6	87.98	19.67
22	331.6	128.04	28.63	225.2	94.04	21.03	244.8	124.52	27.84
23	391.9	125.19	27.99	180.6	77.41	17.31	237.5	108.34	24.23
24	347.0	149.44	33.42	240.6	110.73	24.76	206.3	80.32	17.96
25	440.4	95.88	21.44	182.2	98.58	22.04	195.9	101.55	22.71
26	293.7	169.22	37.84	201.0	94.38	21.10	239.3	134.93	30.17
27	188.2	105.98	23.70	191.3	116.84	26.13	134.2	83.92	18.76
28	324.9	114.01	25.49	246.8	100.09	22.38	203.6	63.23	14.14
29	280.7	104.66	23.40	201.1	121.54	27.18	167.1	53.12	11.88
30	196.1	71.30	15.94	204.8	110.83	24.78	176.7	71.75	16.04
31	297.0	141.03	31.54	209.2	89.63	20.04	250.6	146.11	32.67
32	302.1	128.44	28.72	162.4	86.04	19.24	176.6	71.30	15.94
33	294.1	116.06	25.95	242.5	129.64	28.99	209.0	78.14	17.47
34	368.2	174.61	39.04	184.6	85.88	19.20	203.2	84.92	18.99
35	506.0	161.86	36.19	205.8	105.28	23.54	279.7	127.20	28.44
36	336.7	177.48	39.69	180.0	84.22	18.83	248.2	144.95	32.41
37	418.3	107.99	24.15	223.7	110.56	24.72	248.0	126.06	28.19
38	203.9	104.00	23.26	173.7	94.17	21.06	172.3	89.66	20.05
39	286.1	87.01	19.46	216.0	101.59	22.72	196.2	65.87	14.73
40	351.2	190.42	42.58	178.0	70.48	15.76	217.5	124.32	27.80
41	490.5	138.91	31.06	191.7	75.59	16.90	242.0	135.39	30.27
42	285.3	120.27	26.89	209.3	101.09	22.61	198.9	127.17	28.44
43	255.5	125.30	28.02	205.8	92.52	20.69	183.4	78.22	17.49
44	327.6	113.68	25.42	151.1	88.65	19.82	186.5	99.88	22.33
45	330.8	153.07	34.23	158.9	88.36	19.76	211.4	96.27	21.53
46	251.3	137.20	30.68	194.0	96.43	21.56	182.8	103.02	23.04
47	422.9	160.66	35.92	230.4	120.01	26.83	264.4	115.25	25.77
48	396.4	120.26	26.89	220.9	109.78	24.55	282.5	112.84	25.23
49	337.1	117.89	26.36	254.6	86.87	19.43	246.2	98.89	22.11
50	163.1	94.05	21.03	223.3	89.38	19.99	168.7	100.51	22.47
Averages	**311.4**	**151.0**	**4.77**	**201.5**	**98.7**	**3.12**	**204.9**	**102.8**	**3.25**

Table 1: Results by image

	Energy Over Image			Energy Over Seam			Ratio of Seam Average Energy	Percent Reduction
	Average	Change	Percent Change	Average	Change	Percent Change	to Image Average Energy	in Image Width
All Items								
Threshold	31.20	10.22	44.35%	12.24	8.54	148.57%	31.94%	42.67%
Accuracy	98.0%	96.0%	98.0%	98.0%	98.0%	98.0%	96.0%	94.0%
Residual	196.71	83.93	90.74	152.15	105.86	150.09	191.06	80.11
Image Jackknife								
Threshold	27.92	7.78	24.91%	10.87	8.15	110.78%	26.58%	41.52%
Accuracy	98.0%	100.0%	98.0%	98.0%	98.0%	100.0%	100.0%	97.9%
Residual	213.05	106.76	125.65	167.94	109.33	166.32	228.76	81.22
Subject Jackknife								
Threshold	31.24	10.13	43.33%	12.25	8.53	148.87%	31.85%	42.17%
Accuracy	97.4%	96.5%	96.6%	97.8%	96.5%	96.7%	96.5%	96.8%
Residual	196.51	84.76	91.68	152.22	106.18	150.40	191.91	81.63
Image and Subject Jackknife								
Threshold	27.90	7.79	24.92%	10.79	8.07	110.88%	26.52%	40.93%
Accuracy	99.5%	99.7%	98.2%	98.8%	99.3%	99.4%	100.0%	99.3%
Residual	213.16	107.21	125.99	169.07	110.72	166.23	229.17	83.88

Table 2: Optimized prediction thresholds by image measure, with and without jackknifing

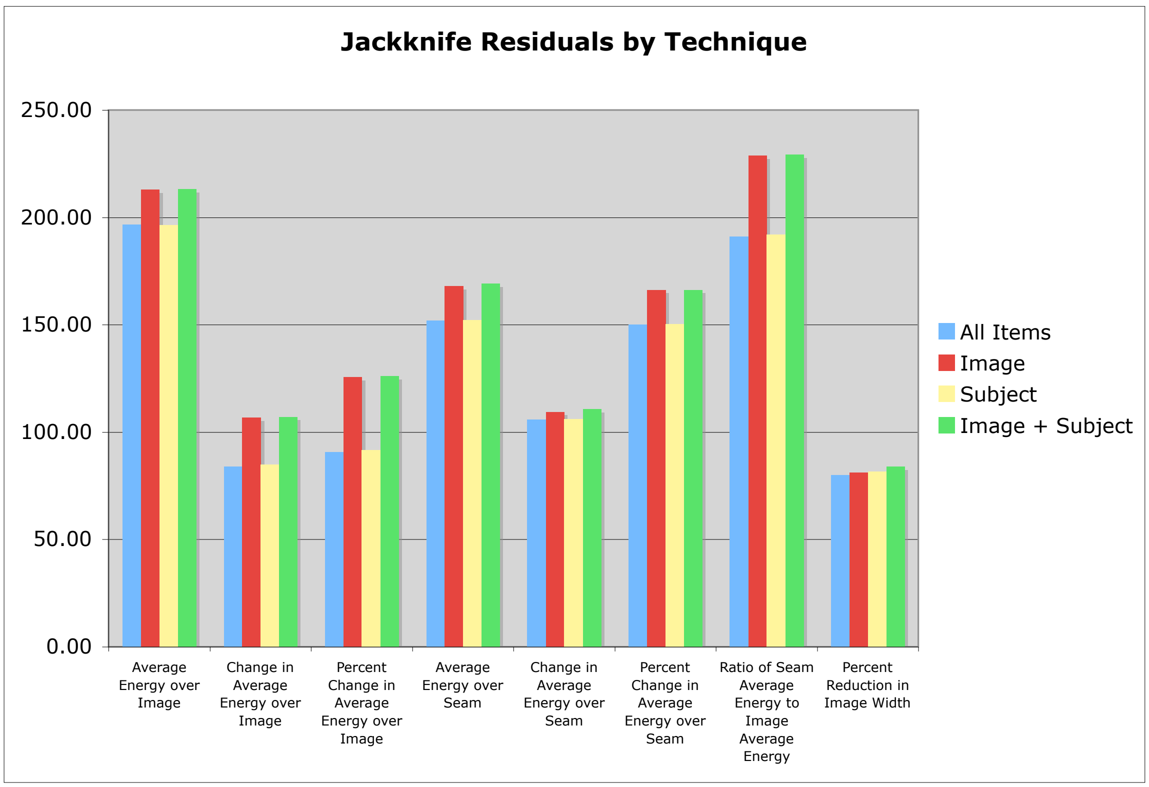

Figure 2: Jackknife Residuals by Technique

Figure 3: Alternative energy functions and their effects on seam carving

Original images

Grayscale gradients

LAB Gradients

Seam carved images, grayscale gradient

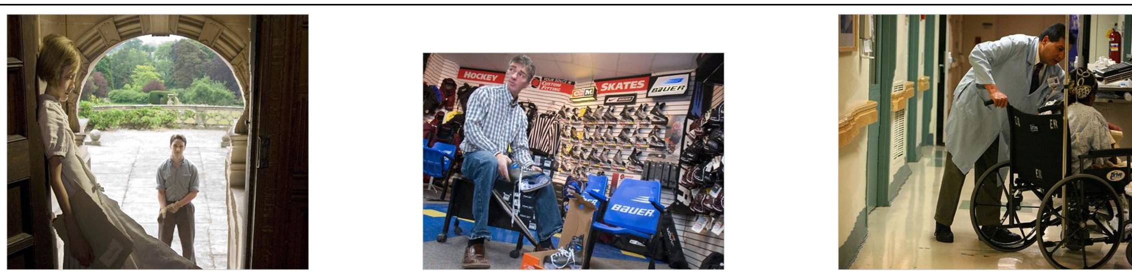

Seam carved images, LAB gradient

Figure 4: Effects of color space on seam carving

Original images

Seam carving with springs, seams shown in red

Seam carving with springs

Seam carving with jumps, seams shown in red

Seam carving with jumps

Figure 5: Seam carving with springs and jumps

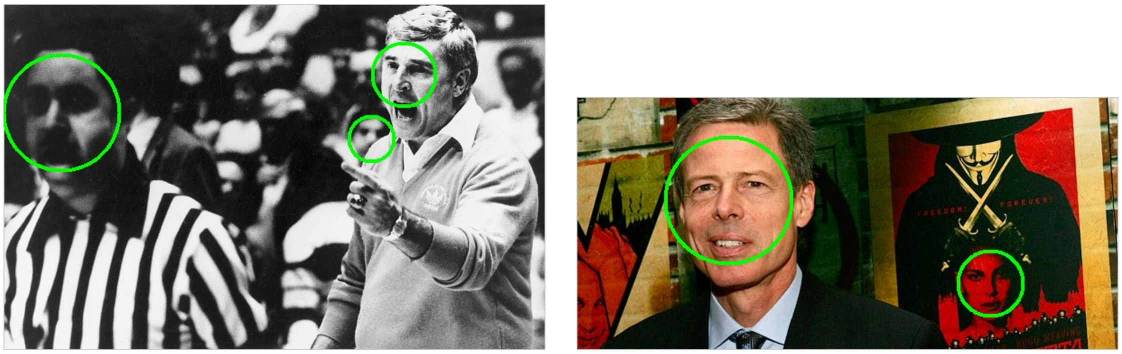

Original images, with detected faces circled

New gradients

Standard seam carving

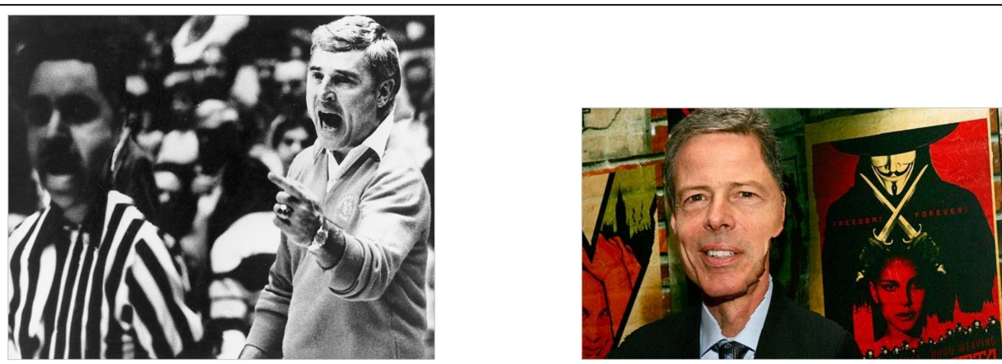

Seam carving with face detection

Figure 6: Modifications to seam carving with face detection