Text Line Extraction Using Seam Carving

Christopher Stoll*, Yingcai Xiao, and Zhong-Hui Duan
Department of Computer Science, The University of Akron, Akron, Ohio, USA
* cas2@zips.uakron.edu

Abstract - *This paper describes a novel technique for information extraction based upon seam carving. Modifications needed to adapt the seam carving process to the new problem domain are explained. Then, two output methods, direct area detection and information masking, are described. Finally, potential modifications to the technique, which could make it suitable for use in other domains such as general computer vision or bioinformatics, are discussed.*

Keywords: text line extraction, seam carving

1. Introduction

While researching the possible application of the eigenface technique [1, 2] towards optical character recognition it was discovered that the technique only worked well when the candidate characters, the characters to be recognized, were scaled and centered exactly as the characters which were used for training. Unlike faces, characters do not share fixed points of reference that can be used to scale and rotate them into a standard form; a robust text line extraction method is required. Existing text extraction approaches, such as X-Y Cut, run-length smearing, and whitespace analysis, did not produce output suitable for the proposed new recognition approach. It is proposed that to achieve the desired results of extracting text from scanned documents, a seam carving approach be used.

With an iterative seam carving approach it will be possible to segment the image, split out the lines of text, and split out the characters from the detected lines. The early iterations identify the text blocks, top and bottom of the characters are identified during the horizontal line splitting step, and the left and right of the characters are identified during the vertical character splitting step. Given the precise boundaries of the candidate characters, they can be projected onto a vector with a preset number of dimensions, and that vector can be used for image recognition. If this approach where to be viable it could further eliminate the need for specific deskewing, connected component analysis, and other preprocessing algorithms. The orientation of the relevant seams could be determined either at the line level or at the character level.

This paper will begin by providing a brief explanation of seam carving, then a novel approach to text line extraction based upon seam carving will be described. Modifications to the seam carving technique, which are necessary to apply seam caving to text line recognition, will be discussed. Then, two variations on the approach, direct area detection and information masking, will be described. Direct area detection attempts to directly locate the boundaries of text lines, while information masking provides a mask under which information is located. Finally, future work and possible alternative applications will be explored.

2. Proposed Approach

The seam caving approach to text extraction has a lot in common with whitespace analysis. The main goal of the seam carving approach described in this paper is to use the whitespace in a document to identify where the interesting information lies; rather than looking for rectangular covers which are joined into large seams, the seam carving approach attempts to find whitespace seams directly.

Consider a blank document, or a white image. With a properly constructed seam tracing function, one that attempts to maintain straight lines, all horizontal seams would go straight across the image. If a single letter or word is added to the center of the document, as seams are examined, from the top down or from left to right, they will begin to deviate as they get closer to the letter (Figure 1). The deviation should reach its maximum around the center of the word, then switch direction, and begin to deviate less.

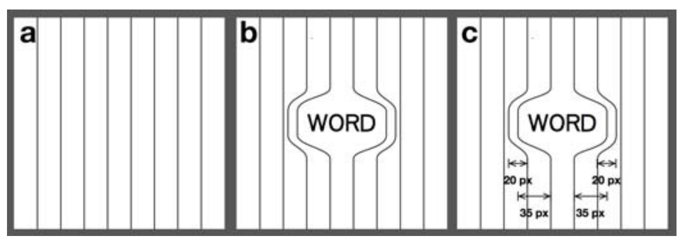

Figure 1: (a) seams in a blank document (b) seams in a document which contains a word (c) seam deviation amounts in a document which contains a word

Seen in this way, just as physical matter in space causes distortions in space-time, information in a document causes distortions in the document space. With the seam carving approach to text segmentation the distortions made by the information are used to identify the boundaries of information within the document. The seam carving approach just needs slight modification to better support this new approach.

2.1. Seam Carving

Seam Carving was originally described by Avidan and Shamir as a technique for content aware image resizing. [3] Their insight was that it could be possible to resize images through the removal of unimportant information from the image rather than though cropping or scaling. The use of seams, which travel a jagged path completely through the image, allow each row or column to remain the same width as the less important pixels are removed.

2.1.1. Preprocessing Functions

The original seam carving paper does not mention explicit preprocessing steps, and they may not be appropriate for image retargeting, but are needed when the approach is applied to text extraction. First, color images are converted to greyscale since color information is less valuable for the purposes of text extraction.

In addition to converting color images to greyscale, different contrast enhancement techniques were experimented with. Otsu binarization was used, but it did not consistently yield the best results for this application. In many situations passing each pixel through a simple cosine function yielded the best results. In common photo editing applications this is known as curve adjustment, here a cosine function was used for the shape of the curve.

2.1.2. Energy Functions

Since the goal of seam carving is to remove pixels which will not be noticed, Avidan and Shamir considered the images' "energies." Each pixel in an image will have an energy value which essentially represents how similar it is to the pixels around it. There are various methods for calculating energies, and since text is normally contrasted against its background almost any method should work. Here we will use the same simple gradient magnitude function used by Avidan and Shamir.

2.1.3. Seam Traversal

Given the results of an energy function, seams must next be calculated for the image. For image retargeting the image can be treated as a graph and seams can be removed iteratively, but Avidan and Shamir chose to take a dynamic programming approach.

For vertical seams the dynamic programming matrix is filled in a top-to-bottom, left-to-right manner; a seam value is calculated for each pixel by adding the minimum of the left three pixels in the current pixel's 8-connected neighborhood (left above, left, left below) to the current pixel's energy value.

For image retargeting, when analyzing vertical seams, the right column of pixels can be checked for minimums once the dynamic programming matrix is filled. The minimum values represent the starting points for seams which can be removed. From these points backtracking is used to remove pixels.

2.1.4. Seam Traversal Modifications

To apply seam carving towards text extraction the seam traversal procedure is modified. First, with text images, seam values create "shadows" which can obscure information following large features (Figure 2b). This effect could be due to the dynamic programming approach's tendency to propagate errors; to check that, a greedy approach was considered.

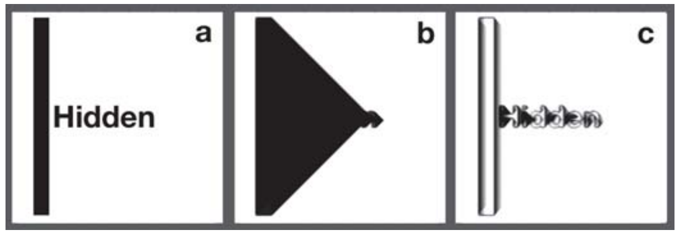

Figure 2: Example of seam "shadows." Original (a), seam values (b), seam values with decrement (c)

The use of a greedy approach, where just the next minimum is taken rather than adding it to the present value, results in a seam value map which is nearly identical to the energy map. The greedy approach does eliminate shadows, however the final results are much worse (Figure 3b). It turns out that the shadows are an important part of this approach. Like a spearhead the shadows guide seams around interesting portions of the image; without the chamfered edges, seams would be more likely to go straight through perpendicular edges rather than flowing around them.

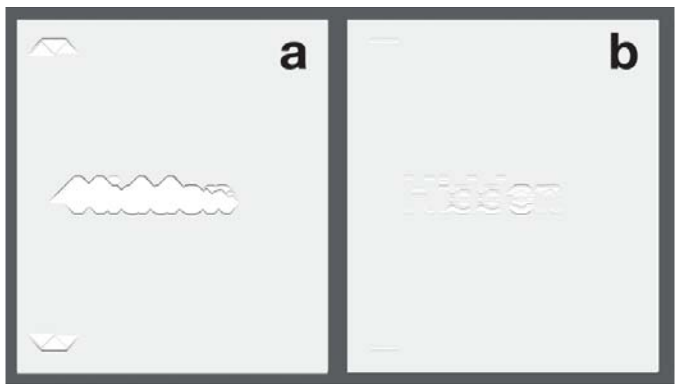

Figure 3: All seams traced (seam traversal darkens image, so white areas have no seam traversal), dynamic programming approach with decrement (a) and greedy approach (b)

Since the greedy approach did not yield the desired results, the dynamic programming approach will be modified to optimize the effect of the shadows. The dynamic programming matrix is filled as described by Avidan and Shamir, except that each pixel's resulting seam value is reduced by some value k if it is greater than zero (Equation 1). This is a key modification, without this decrement seams would never travel between text lines due to how seam values are carried across the image. A k value of 1 was experimentally determined to yield satisfactory results (Figure 2c). Based upon what was learned form taking the greedy approach, it is likely that the optimal k value is image dependent. Further research needs to be done to identify the optimal k value.

$$
\begin{aligned}
M(i, j) &= e(i, j) \\
&+ \min(M(i-1, j-1), M(i-1, j), M(i-1, j+1)) \\
&- k
\end{aligned}
\tag{1}
$$

2.1.5. Text Extraction Steps

Additional, unique steps are required to use seam carving for text line extraction. Rather than examining the last column or row for minimum values, each pixel in the terminal row or column is examined. In fact, finding the seams of minimum value does not generally work with text images. For most images containing text the last row or column will contain a large number of minimum values since the background of most documents is some uniform color. So, backtracking must be performed for each of the pixels in the terminal row or column. Note that since we are not removing any lines, energy values will not be recomputed between seam traversals.

The process of backtracking must also be modified to support the new problem domain. For image retargeting the actual seam path matters little, whereas for text line extraction it is expected to proceed in a straight line since text is normally written in lines. So, in addition to considering the values of the next step when backtracking, the overall deviation from straight is also considered.

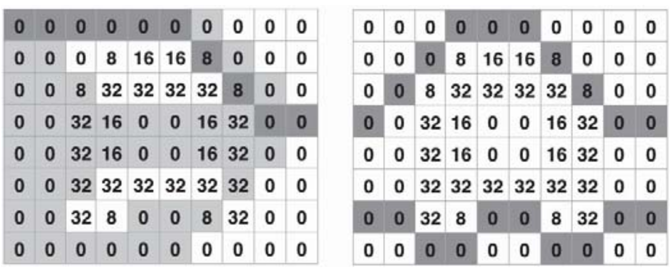

0	0	0	0	0	0	0	0	0	0
0	0	0	8	16	16	8	0	0	0
0	0	8	32	32	32	32	8	0	0
0	0	32	16	0	0	16	32	0	0
0	0	32	16	0	0	16	32	0	0
0	0	32	32	32	32	32	32	0	0
0	0	32	8	0	0	8	32	0	0
0	0	0	0	0	0	0	0	0	0

0	0	0	0	0	0	0	0	0	0
0	0	0	8	16	16	8	0	0	0
0	0	8	32	32	32	32	8	0	0
0	0	32	16	0	0	16	32	0	0
0	0	32	16	0	0	16	32	0	0
0	0	32	32	32	32	32	32	0	0
0	0	32	8	0	0	8	32	0	0
0	0	0	0	0	0	0	0	0	0

Figure 4: Possible (light grey) and probable (dark grey) seam paths are shown for a standard seam carving approach (left). Top and bottom seam paths (dark grey) are shown for backtracking with added deviation constraints (right)

If a seam has deviated from vertical or horizontal then the algorithm will force the seam back in that direction. When there are more than one possible backtracking move with the same value, the algorithm has discretion to choose which path to take. For the purposes of this text extraction approach, the algorithm must take the center path whenever possible, unless there is a net deviation. When there is a deviation the implementation must take the choice that minimizes the deviation. This results in seams which minimize the deviation and minimize the energy of the path.

2.2. Information Extraction

From here there are two distinct approaches to identifying the distortions in the image seams. Direct area detection uses heuristics to identify information as originally proposed, whereas information masking crates an image mask based upon where no seams travel.

2.2.1. Direct Area Detection

The first way to identify distortions in the seams is to check for seams which have the most net deviation from straight, as originally discussed. As a horizontal seam traverses across the image each column's deviation from the starting row is recorded. The sum of these deviations is the net deviation. When a seam's net deviation changes sign it represents the bottom of an area, and the seam prior to it represents the top of an area. Long straight runs at the beginning and end are trimmed off.

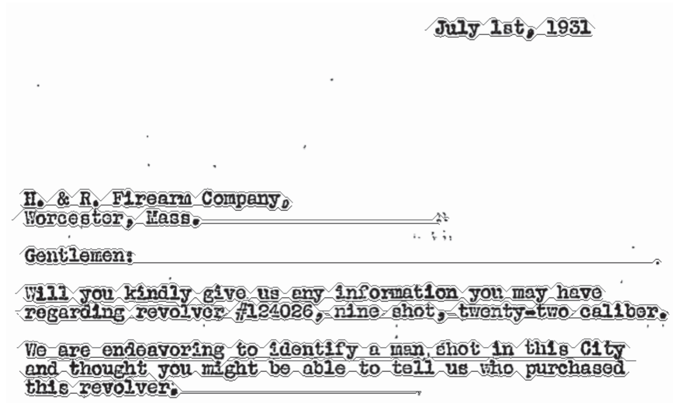

July 1st, 1931

H. & R. Firearm Company,
Worcester, Mass.

Gentlemen:

Will you kindly give us any information you may have regarding revolver #124026, nine shot, twenty-two caliber.

We are endeavoring to identify a man shot in this City and thought you might be able to tell us who purchased this revolver.

Figure 5: Direct Area Detection

In practice, seams defining the beginning and ending boundaries (upper and lower for horizontal seams or left and right for vertical seams) are rarely consecutive. If a beginning seam has not been found (without a corresponding ending seam) then local maximum net deviations are sought. Once a local maximum is found, a local minimum net deviation is sought. This process repeats as the entire image is processed.

2.2.2. Information Masking

Another approach to identify distortions in the seams is to look for pixels where seams overlap. As seams divert around information they will tend to go through the pixels just outside the upper and lower boundary. If a matrix is created to keep track of how many seams pass through each pixel, then the matrix cells nearest to information will have higher values (black areas in Figure 6) and matrix cells where information resides will have zero values (white areas in Figure 6).

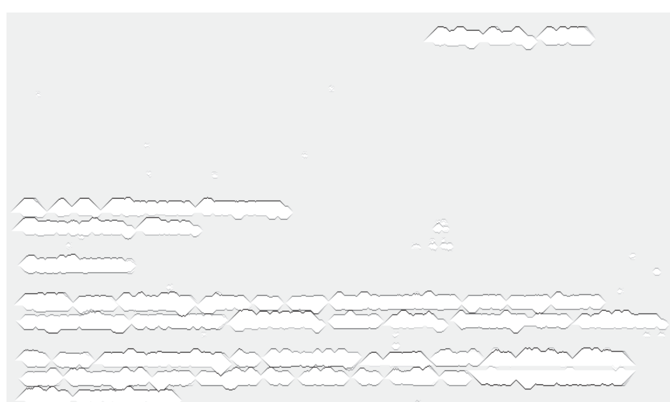

Figure 6: Information Masking

2.2.3. Method Comparison

Direct area detection and information masking both work reasonably well when the document is simple, but information masking automatically works much better on more complicated document layouts. When there are three columns of irregularly laid out words the difference in character size between the columns prevents lines from being properly identified. In the cases where lines are identified correctly, they are actually in different columns and should not be treated as single lines. Another issue with direct area detection is that words near the edge are sometimes not properly detected. The problem appears to be that the leading and trailing lines are trimmed prematurely due to the lack of seam run length. Finally, direct area detection, without further enhancements, begins to fail faster on documents which are skewed; this is especially problematic since handling skew is normally a strength of the seam carving approach.

Information masking is superior to direct area detection when the target text is skewed. However, when text skew reaches a certain point even information masking, under the current implementation, begins to degrade in performance. Currently, direct area detection clearly begins to break down at two degrees of skew, and stops working completely by the time the text is skewed by eight degrees. Information masking, since it does not rely upon heuristics, continues to work at eight degrees of skew, but the seam markers at the beginning and ending of lines becomes stretched.

3. Discussion

In order to be a viable solution the shortcomings mentioned in the above sections should be overcome, otherwise there is less compulsion to use this technique over the established techniques. The problem of skew is minor; it has been shown, through the information masking output, that the raw seam carving technique can identify text that is skewed. It is only the heuristics for direct area detection that need to be improved. And, if all else fails, a dedicated deskewing algorithm can be implemented as a preprocessing step.

The problem of handling more complicated text layouts, however, is more significant. Again, the raw seam carving technique, as shown through the information masking approach, can identify text regardless of its position, but keeping the text logically grouped is, as it always has been, more problematic.

3.1.1. Improved Skew Handling

In order to improve the handling of skewed text the approach could be modified so that it does not attempt to force seams to travel at precisely 90 or 180 degrees. Currently the algorithm encourages seams to end in the same row or column in which they started. A better method would be to identify deviation trends; if the seam is steadily deviating more, then it could be assumed that the information that is displaying it is skewed. Rather than attempting to force seams to travel at exactly 90 degrees or 180 degrees the algorithm could attempt to force the seam to travel along the trend angle.

Since direct area detection performs so poorly at detecting starting and ending seams when the text is skewed, that portion should also be improved. Currently it simply evaluates the magnitude of the net deviation, once it changes sign the top and bottom seams are presumed to have been found, but this is not necessarily true with skewed text. Checking for a sign change in magnitude should account for the overall angle of seams within a certain range.

In all cases the heuristics used should be improved through more thorough data analysis, which in turn requires improved parameterization of the algorithm. The values used for heuristics are the result of manually examining the algorithm's performance over a very small data set. The use of a larger and more varied set of example images should yield better heuristic values.

3.1.2. Handling Complicated Layouts

There are various possible techniques to overcome the problems of locating text within columns and other more complicated layouts. Most of these possible techniques are based upon methods described in other text extraction approaches.

The first option is to run the seam carving algorithm both vertically and horizontally, then using a heuristic to determine which direction is more appropriate to split in first. One possible heuristic is total resistance. The idea is that, as each of the seams are explored, when they are forced to deviate from a straight path they are experiencing resistance. For direct area detection the net deviation for each seam is calculated, so it would be trivial to sum these quantities. The sum net deviations for horizontal and vertical seam traversal should then be compared. The direction with less resistance, or a lower sum net deviation, is the direction which should be explored first. This technique helps when identifying an information mask, but it is still very crude. It tends to favor moving horizontally (for left-to-right or right-to-left texts) and does not help improve direct area detection.

Another approach to higher level segmentation would be to take a kD-tree or X-Y cut type of approach. Instead of initially tracing all the seams, the minimum seams will be found, as is done with typical seam carving. The difference comes in how "minimum seam" is defined. Rather than checking the terminal value of the seam this approach would need to find the seam with the minimum net deviation. Since there are likely to be multiple seams with the same minimum net deviation, this approach would need to find the center seam of the largest grouping of minimum net deviation seams. The identified seam would be used to split the image and the process would be ran again on each of the halves. This was not implemented as a part of this project, but is a goal of future research.

3.2. Extracting Finer Details

Setting aside this approach's current limitations, more of its benefits will be described. Given a non-skewed document which contains simple content with a manhattan layout, it is possible to iteratively run the same seam carving process on identified sub areas. If the first pass of the algorithm identified sentences (Figure 5), then the next pass of the algorithm will identify phrases, words, or letters (Figure 7). The technique is simply used against areas identified in the first step, but in a perpendicular direction. This process can be repeated until the base case or halting condition, which is not yet fully defined, is reached.

Figure 7: Seams identified in previously identified areas of interest. The first line (left) and second line (right) identified in the document from Figure 5

3.3. Asymptotic Analysis

The first step in the seam carving process is to read in an image and create an appropriate data structure. The example program performs brightness, contrast, and energy calculations as distinct steps, but this is done to allow for flexibility during testing, a production program would combine these into a single step. The Difference of Gaussians could be performed in one pass given the appropriate data structure. So, assuming that the image has a width of n and a height of m, then this process will take O(*nm*).

The next step is to fill the seam matrix, which requires iterating over the image again, taking O(*nm*). Finally, all the seams must be traversed. Considering horizontal seam evaluation, each of the n starting pixels is considered as a starting point, and the seam must cross the entire height m of the image, so this process takes another O(*nm*). The runtime performance for a single pass is 3 * O(*nm*), and 5 * O(*nm*) when both horizontal and vertical seams are evaluated. So, this approach performs in O(*nm*) time, where n is image width and m is image height. This conclusion is supported by empirical analysis (Figure 8).

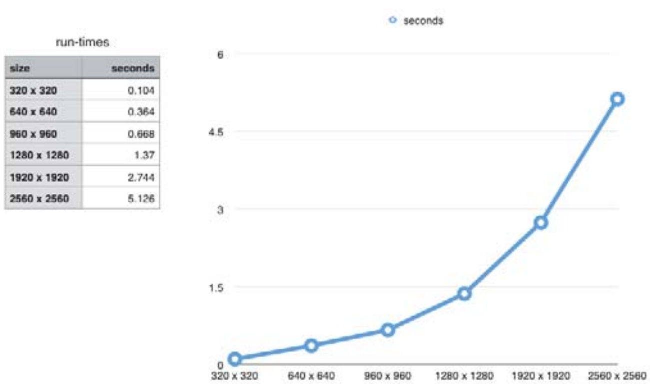

Figure 8: Actual run-time performance data; both directions, average of 20 runs

It should be noted that although vertical and horizontal seam analysis both perform in essentially O(n^2) time, real world performance can be drastically different due to algorithm implementation and hardware limitations. If the image data structure uses row major order, then cutting horizontal seams can take twice as long to accomplish on some hardware due to how memory is cached and moved to the processor.

3.4. Related Work

Seam carving has previously been used in the area of handwriting recognition, however this approach takes a slightly different perspective on the problem. Saabini and El-Sana describe a method for applying seam carving to handwritten text line extraction [11]. When they surveyed the research on text line extraction for handwriting recognition they also found that most of the techniques relied upon some sort of connected component analysis, so they devised a novel technique based upon seam carving.

The algorithm described by Saabini and El-Sana calculates the image's or document's energy in such a way (using a signed distance transformation) that the selected seam paths went through the lines of handwritten text rather than attempting to identify the whitespace. Once the text line's center line is found it expands vertically to identify the full height of the line.

One major advantage of their approach is that small components, such as the dot above an "i" can be included in the row based upon what the algorithm learns about row heights. The new technique described in this paper can be susceptible to leaving out such small components, especially when processing images which have not first been subdivided. In practice, for single lines of text, the technique described here automatically includes small pieces of information due to how the heuristics are set find maximum deviation.

Another benefit of Saabini and El-Sana's technique is that since it is tracking lines of text, it can handle multi-skew and lines that actually touch each other. For the technique described in this paper, even if a moving average was used to determine line skew, it will still have difficulty handling multi-skew. This technique can also cope with lines that touch at least as well as the method described by Saabini and El-Sana.

Asi, Saabini and El-Sana subsequently demonstrated that an enhancement to Saabini and El-Sana's approach could identify the precise boundaries between text lines [19]. Their method's ability to do this is entirely dependent upon it ability to find the text lines' medial path, which in turn relies upon single columns of handwritten information. Whereas their work is likely to perform superiorly in identifying lines of free-form handwritten text, the method described in this paper is likely to excel in typed or printed documents. The method described here is intended to run recursively, so there is some expectations that the information it finds will be subsequently subdivided, whereas their approach is intended to identify information using a single pass.

3.5. Future Work

The existing implementation of seam carving for text line extraction performs as well as other notable text extraction methods. Like them it has trouble with skewed text and more complicated layouts, traditional problems in this area of research. Fortunately, there are some techniques to possibly overcome these shortcomings. The first improvements should be in the area of handling skewed text. The technique should consider some sort of moving average of the seam angles, but more research will need to be done in this area.

Also, experiments should be conducted where different energy maps are used. The use of signed distance transforms (SDT), as Saabini and El-Sana, should be evaluated[11]. This approach will be looking for local maximum (the whitespace), rather than local minimum. Also, the approach of using elliptical convolution kernels which align to the direction of travel for the image energy calculation, as described by Zhang and Tan, should be explored [19]. Zhang and Tan also described a method for limiting the distance which energy should be propagated which should be incorporated into this approach.

Another area of future work is the implementation of a kD-tree or X-Y cut approach to higher level document segmentation. This will help define the structure of the document and improve actual text extraction. Since a dynamic programming approach has been taken and due to the nature of seams within text documents, it should be possible to reuse seam data as the document is subdivided, thus reducing computation time. A check may be required to ensure that the border pixels of the sub-area all contain zero weights. Also, a more precisely defined base condition would need to be defined to limit the number of subdivision attempts.

An alternative approach to higher level document segmentation using seam carving would be to use the idea of “zooming,” or otherwise considering the level of detail (LOD). This idea is based upon how people process information within their view; people can get an overview of what they see or they can understand fine details within it, but those are generally two distinct steps. To simulate this process the candidate image could have its size reduced and contents blurred to decrease the definition of finer details. Then, when seam carving is ran it will not find sentences, for example, but rather it would find paragraphs or columns.

The approach of using LOD could bring multiple benefits. First, since the seam caving approach runs in polynomial time, a reduction in image size would greatly in crease runtime performance. However, the dynamic programming matrix would have to be recreated at each zoom level which could negate any realized gains. If this modification was combined with the kD-tree modification, it would make identification of higher level divisions easier. The idea of “zooming out” also presents opportunities for early recognition or at least setting a context for subsequent recognition steps. It may be possible to run the low resolution areas of interest through an algorithm which quickly detects basic shapes, or one that uses textons [12] to detect textures. If this approach were viable it would have implications outside of the realm of OCR; it could be used in the field of bioinformatics or for general computer vision.

Existing implementations of this approach have been experimentally applied to bioinformatics. There are techniques being developed to identify carcinomas based upon examining differences of protein profiles; presently the differentiation of the protein profiles is done by manual inspection [16]. Enhanced version of the seam carving approach may allow for algorithmic differentiation of protein profiles.

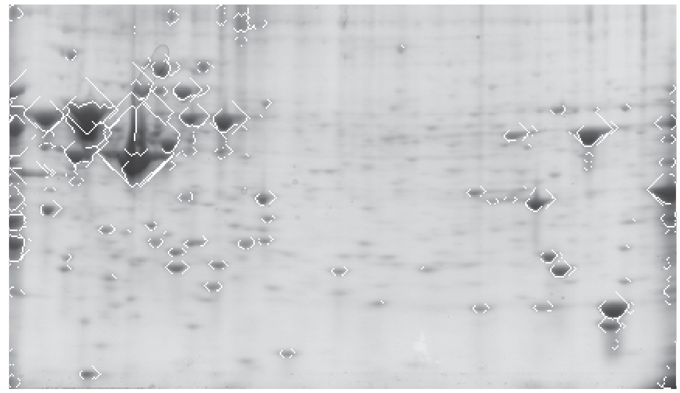

Figure 9: Experimental application toward protein differentiation; white lines around dark areas show boundaries found running this approach both vertically and horizontally

Successful application to bioinformatics may further lead to applications toward general computer vision. Consider an image which contains a stop sign. If the algorithm identifies the sign as being interesting in one of the initial iterations, and if the octagonal shape could be detected, then it may not even be necessary to continue processing the area — if a person sees a stop sign out of the corner of their eye they do not need to read the word “STOP” to understand the meaning of it. For most driving adults, seeing a red octagon is enough to understand that the symbol means stop. Also, if “STOP” is being recognized but some of the letters are illegible, the fact that the word being recognized is within an octagon would induce the algorithm to respond that the word is “STOP.”

3. Conclusions

A promising new approach to text extraction has been described based upon seam carving. The seam carving approach brings many of the same benefits of white space analysis techniques; it operates largely parameter free, it does not make assumptions about document layout and it works regardless of text direction or page orientation. And though it also has many of the same limitations that whitespace analysis, there are some promising methods unique to this approach which can help overcome those shortcomings and allow for it to be used outside of the original problem domain.

Applying the seam carving approach to text line extraction reinforced that the original technique was designed for images and not text documents. When seam values are calculated, the original technique causes “shadows” to be cast which can block out smaller pieces of information. The shadows represent errors which, unlike images, easily get propagated through the mostly blank space of a text document. The shadows were reduced by adding decay to the seam value function.

The approach described by Avidan and Shamir also relies upon the ability to find minimal terminal seam values from which to backtrack. For text documents the terminal value of all seams is likely to be zero. This is due to the amount of whitespace, space without information, in text documents. Images have a higher information density and are not normally padded with headspace. Modifying the technique so that every seam terminus is examined is necessary to apply seam carving to text line extraction.

The original seam carving technique further allowed for indeterminate seam paths (determined by the implementation). For resizing images it is probably beneficial to have some randomness in the seam path, but that works against text line extraction. Text normally flows in straight lines, and it makes sense to consider that when looking to extract text lines. To apply seam carving to text extraction the concept of net deviation was introduced; in addition to minimizing the seam path cost the net deviation of the seam is also minimized.

These minor innovations have enabled the seam carving algorithm to be applied to task of text line extraction; it builds upon existing techniques, yet represents a truly novel approach. The potential benefits of perfecting this approach could have impacts beyond optical character recognition, and is thus worthy of more research.

4. References

1. L. Sirovich and M. Kirby, “Low-dimensional procedure for the characterization of human faces,” Journal of the Optical Society of America A, Volume 4, Issue 3, pp. 519-524, 1987.

2. M. Turk and A. Pentland, “Face recognition using eigenfaces,” Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 586–591, 1991.

3. S. Avidan, and A. Shamir, "Seam carving for content-aware image resizing,” ACM Transactions on graphics (TOG), Volume 26, Number 3, 2007.

4. K. Y. Wong, R. G. Casey, and F. M. Wahl, “Document Analysis System,” IBM Journal of Research and Development, Volume 26, Issue 6, pp. 647-656, 1982.

5. G. Nagy and S. Seth, “Hierarchical representation of optically scanned documents,” International Conference on Pattern Recognition - ICPR, 1984.

6. L. O’Gorman, “The Document Spectrum for Page Layout Analysis,” IEEE Transactions on Pattern Analysis and Machine Intelligence, Volume 15 Issue 11, pp. 1162-1173, 1993.

7. H. S. Baird, “Background Structure In Document Images,” In Advances in Structural and Syntactic Pattern Recognition, pp. 17-34, 1992.

8. T. M. Breuel, “Two Geometric Algorithms for Layout Analysis,” Document Analysis Systems, 2002.

9. T. M. Breuel, “Robust least square baseline finding using a branch and bound algorithm,” Document Recognition and Retrieval, SPIE, pp. 20-27, 2002.

10. K. Kise, A. Sato, and M. Iwata, “Segmentation of Page Images Using the Area Voronoi Diagram,” Computer Vision and Image Understanding, Volume 70, Number 3, pp. 370-382, 1998.

11. R. Saabni and J. El-Sana, “Language-Independent Text Lines Extraction Using Seam Carving,” 2011 International Conference on Document Analysis and Recognition (ICDAR), pp. 563-568, 2011.

12. B. Julesz, “Textons, the Elements of Texture Perception, and their Interactions,” Nature, Volume 290, pp. 91–97, 1981.

13. J. Ha, R. M. Haralick, and I. T. Phillips, “Recursive X-Y Cut using Bounding Boxes of Connected Components,” Third International Conference on Document Analysis and Recognition, pp. 952-955, 1995.

14. L. O’Gorman, “Image and document processing techniques for the RightPages electronic library system,” International Conference on Pattern Recognition, pp. 260-263, 1992.

15. G. Nagy, J. Kanai, M. Krishnamoorthy, M. Thomas, and M. Viswanathan, ‘‘Two Complementary Techniques for Digitized Document Analysis” ACM Conference on Document Processing Systems, 1988.

16. T. Shi, F. Dong, L. S. Liou, Z. H. Duan, A. C. Novick, J. A. DiDonato, ‘’Differential Protein Profiling in Renal-Cell Carcinoma” Molecular Carcinogenesis, Volume 40, pp. 47-61, 2004.

17. Medical Article Records Groundtruth (MARG), National Library of Medicine, http://marg.nlm.nih.gov, 2005.

18. S. Firilli, F. Leuzzi, F. Rotella, F. Esposito, ‘’A Run Length Smoothing-Based Algorithm for Non-Manhattan Document Segmentation,” Italy, 2012.

19. A. Asi, R. Saabni and J. El-Sana, “Text Line Segmentation for Gray Scale Historical Document Images,” International Workshop on Historical Document Imaging and Processing (HIP), 2011.

20. X. Zhang and C. L. Tan, “Text Line Segmentation for Handwritten Documents Using Constrained Seam Carving”