Volumetric Seam Carving

MS Thesis Defense

Dachao Sun, Computer Science

November 21, 2016

School of Computing, Clemson University

Committee:

- Dr. Donald House (chair)
- Dr. Joshua Levine (co-chair, **advisor**)
- Dr. Brian Dean

CLEMSON School of COMPUTING

“The seam carving removes the best ‘column’ to resize an image in 2D $I(x, y)$. In Vis., we work on 3D images from medical imaging and all kinds of scientific simulations. It will be handy to have a tool like this, as we will instead need to find the best ‘plane’ to carve with.”

April 2015

Introduction

Introduction

Why?

Track 1: **Volume data reduction**

Higher-resolution volumes, storage/rendering challenges; calls for reduced-size volumes.

Will seam carving do a good job?

Track 2: **Topological analysis (data *understanding*)**

What happens to *isosurfaces* when volume is "seam carved" — does the "content-aware" seam carving also preserve topological features?

Introduction

Why?

Track 1: **Volume data reduction**

Higher-resolution volumes, storage/rendering challenges; calls for reduced-size volumes.

~~Will~~ How would seam carving ~~do a good job~~ perform?

Track 2: **Topological analysis (data *understanding*)**

What happens to *isosurfaces* when volume is "seam carved" — ~~does the "content-aware"~~ How does seam carving ~~also preserve~~ affect topological features?

Volumetric image

Volumetric image / 3D image refers to the digitalized, three-dimensional image in the form of a **three-dimensional array of (usually) scalars** in computer memory.

By acqusition: Computed Tomography (CT), Magnetic Resonance Imaging (MRI), Ultrasound Imaging, scientific simulations.

Analysis goal: (1) **prepare for visualization**; (2) geometric alignment; (3) automatic structure extraction; (4) **data understanding**: segmentation, labeling, feature detection.

Introduction

Volumetric image: visualization

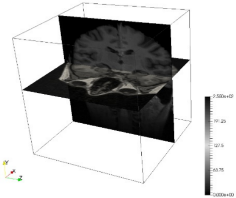

Slicing

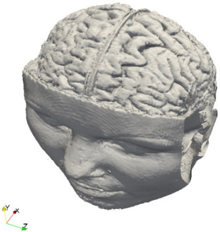

Isosurfacing

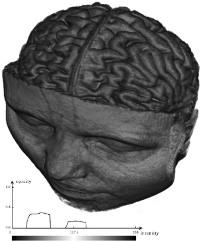

volume rendering

Seam Carving

Seam Carving

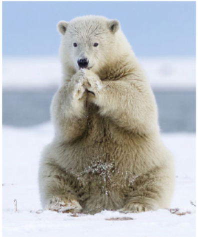

Crop

Scale

Content-aware

Seam Carving

Seam Carving

Seam Carving

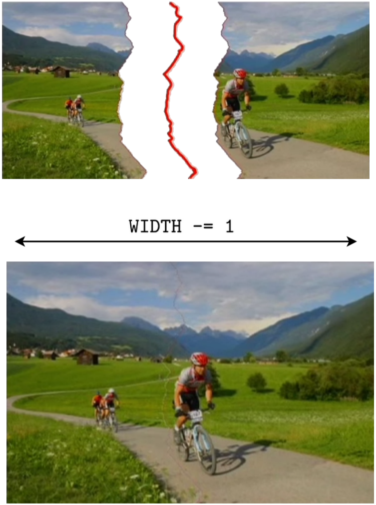

Ariel Shamir “The Vision of Image”, TEDx talk

Seam Carving

2D image:

$$\bar{\mathbf{I}} = \Big\{ I(x,y) \;\big|\; x \in [0,w), y \in [0,h) \Big\}$$

Energy function (L_1-norm of gradient):

$$e_1\big(x,y\big) \;=\; \left|\frac{\partial}{\partial x} I(x,y)\right| + \left|\frac{\partial}{\partial y} I(x,y)\right|$$

Vertical seam / x-seam:

$$\mathbf{s^x} \;=\; \{s^x_j\} \;=\; \Big\{\Big(\; x(j),\; j \;\Big)\Big\}$$

$$s.t.\;\; j \in \mathbb{Z},\; 1 \leqslant j \leqslant h,\; \text{and}$$

$$\forall j,\; \big|x(j) - x(j-1)\big| \leqslant 1$$

Seam Carving

Optimal x-seam:

$$\mathbf{s}^{\mathbf{x}*} = \min_{\mathbf{s}^{\mathbf{x}}} \sum_{j=1}^{h} e_1\left(s_j^x\right)$$

Solved using **dynamic programming**

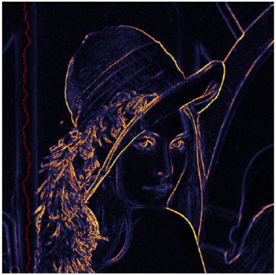

12 (12)	**15** (15)	**28** (28)	**25** (25)
		232 (217)	

= 217 + min{15,28,25}

Seam Carving

Proposal: Shai Avidan and Airel Shamir, *Seam carving for content-aware image retargeting* ACM Transactions on Graphics, 26(3), 2007.

It is:

- an image retargeting approach
- a discrete approach
- content-aware
- **improved and extended to retarget videos (spatial-temporal volumes)**

Seam Carving

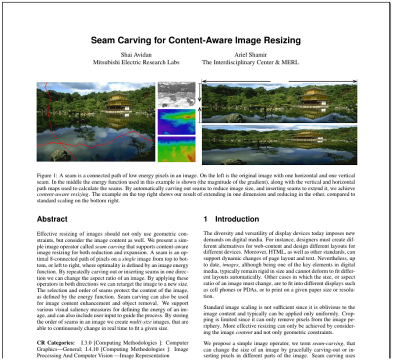

Seam Carving for Content-Aware Image Resizing

Shai Avidan, Mitsubishi Electric Research Labs
Ariel Shamir, The Interdisciplinary Center & MERL

Figure 1: A seam is a connected path of low energy pixels in an image. On the left is the original image with one horizontal and one vertical seam. In the middle the energy function used in this example is shown (the magnitude of the gradient), along with the vertical and horizontal path maps used to calculate the seams. By automatically carving out seams to reduce image size, and inserting seams to extend it, we achieve *content-aware resizing*. The example on the top right shows our result of extending in one dimension and reducing in the other, compared to standard scaling on the bottom right.

Abstract

Effective resizing of images should not only use geometric constraints, but consider the image content as well. We present a simple image operator called *seam carving* that supports content-aware image resizing for both reduction and expansion. A seam is an optimal 8-connected path of pixels on a *single* image from top to bottom, or left to right, where optimality is defined by an image energy function. By repeatedly carving out or inserting seams in one direction we can change the aspect ratio of an image. By applying these operators in both directions we can retarget the image to a new size. The selection and order of seams protect the content of the image, as defined by the energy function. Seam carving can also be used for image content enhancement and object removal. We support various visual saliency measures for defining the energy of an image, and can also include user input to guide the process. By storing the order of seams in an image we create *multi-size* images, that are able to continuously change in real time to fit a given size.

CR Categories: I.3.0 [Computing Methodologies]: Computer Graphics—General; I.4.10 [Computing Methodologies]: Image Processing And Computer Vision —Image Representation

1 Introduction

The diversity and versatility of display devices today imposes new demands on digital media. For instance, designers must create different alternatives for web-content and design different layouts for different devices. Moreover, HTML, as well as other standards, can support dynamic changes of page layout and text. Nevertheless, up to date, *images*, although being one of the key elements in digital media, typically remain rigid in size and cannot deform to fit different layouts automatically. Other cases in which the size, or aspect ratio of an image must change, are to fit into different displays such as cell phones or PDAs, or to print on a given paper size or resolution.

Standard image scaling is not sufficient since it is oblivious to the image content and typically can be applied only uniformly. Cropping is limited since it can only remove pixels from the image periphery. More effective resizing can only be achieved by considering the image *content* and not only geometric constraints.

We propose a simple image operator, we term *seam-carving*, that can change the size of an image by gracefully carving-out or inserting pixels in different parts of the image. Seam carving uses

2007

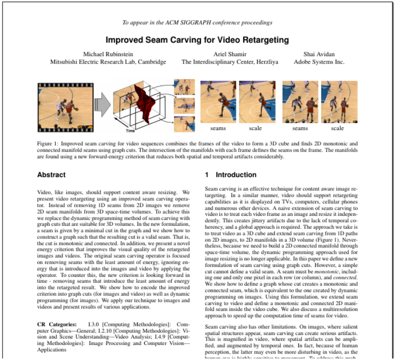

To appear in the ACM SIGGRAPH conference proceedings

Improved Seam Carving for Video Retargeting

Michael Rubinstein, Mitsubishi Electric Research Lab, Cambridge
Ariel Shamir, The Interdisciplinary Center, Herzliya
Shai Avidan, Adobe Systems Inc.

Figure 1: Improved seam carving for video sequences combines the frames of the video to form a 3D cube and finds 2D monotonic and connected manifold seams using graph cuts. The intersection of the manifolds with each frame defines the seams on the frame. The manifolds are found using a new forward-energy criterion that reduces both spatial and temporal artifacts considerably.

Abstract

Video, like images, should support content aware resizing. We present video retargeting using an improved seam carving operator. Instead of removing 1D seams from 2D images we remove 2D seam manifolds from 3D space-time volumes. To achieve this we replace the dynamic programming method of seam carving with graph cuts that are suitable for 3D volumes. In the new formulation, a seam is given by a minimal cut in the graph and we show how to construct a graph such that the resulting cut is a valid seam. That is, the cut is monotonic and connected. In addition, we present a novel energy criterion that improves the visual quality of the retargeted images and videos. The original seam carving operator is focused on removing seams with the least amount of energy, ignoring energy that is introduced into the images and video by applying the operator. To counter this, the new criterion is looking forward in time - removing seams that introduce the least amount of energy into the retargeted result. We show how to encode the improved criterion into graph cuts (for images and video) as well as dynamic programming (for images). We apply our technique to images and videos and present results of various applications.

CR Categories: I.3.0 [Computing Methodologies]: Computer Graphics—General; I.2.10 [Computing Methodologies]: Vision and Scene Understanding—Video Analysis; I.4.9 [Computing Methodologies]: Image Processing and Computer Vision—Applications

1 Introduction

Seam carving is an effective technique for content aware image retargeting. In a similar manner, video should support retargeting capabilities as it is displayed on TVs, computers, cellular phones and numerous other devices. A naive extension of seam carving to video is to treat each video frame as an image and resize it independently. This creates jittery artifacts due to the lack of temporal coherency, and a global approach is required. The approach we take is to treat video as a 3D cube and extend seam carving from 1D paths on 2D images, to 2D manifolds in a 3D volume (Figure 1). Nevertheless, because we need to build a 2D connected manifold through space-time volume, the dynamic programming approach used for image resizing is no longer applicable. In this paper we define a new formulation of seam carving using graph cuts. However, a simple cut cannot define a valid seam. A seam must be *monotonic*, including one and only one pixel in each row (or column), and *connected*. We show how to define a graph whose cut creates a monotonic and connected seam, which is equivalent to the one created by dynamic programming on images. Using this formulation, we extend seam carving to video and define a monotonic and connected 2D manifold seam inside the video cube. We also discuss a multiresolution approach to speed up the computation time of seams for video.

Seam carving also has other limitations. On images, where salient spatial structures appear, seam carving can create serious artifacts. This is magnified in video, where spatial artifacts can be amplified, and augmented by temporal ones. In fact, because of human perception, the latter may even be more disturbing in video, as the

2008

Related Work

Related work on seam carving:

geometry, depth camera, volume

Ellen Dekkers and Leif Kobbelt. Geometry seam carving. *Computer-Aided Design*, 46:120–128, 2014.

Jianbing Shen, Dapeng Wang, and Xuelong Li. Depth-aware image seam carving. *IEEE transactions on cybernetics*, 43(5):1453–1461, 2013.

Qichao Wang, Yubo Tao, and Hai Lin. Surface carving-based automatic volume data reduction. *The Visual Computer*, 31(11):1459–1470, 2015.

Volumetric Seam Carving

Volumetric Seam Carving: Pipeline

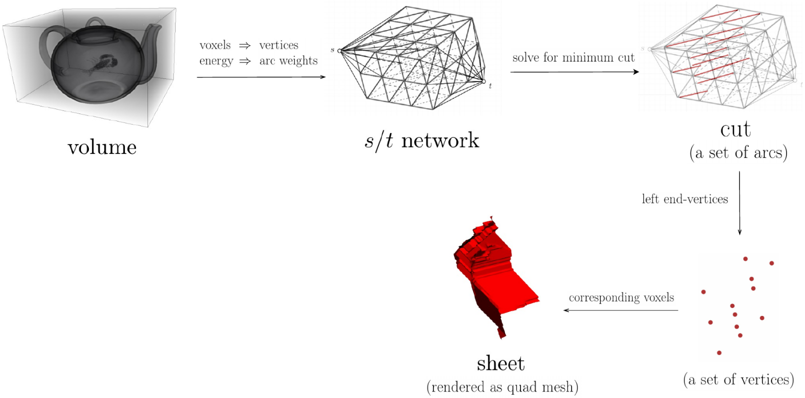

Volumetric Seam Carving: Removing a “Sheet”

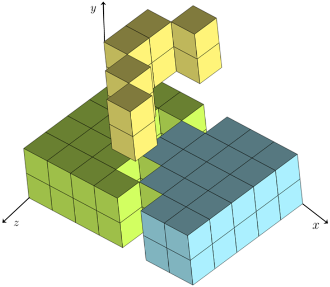

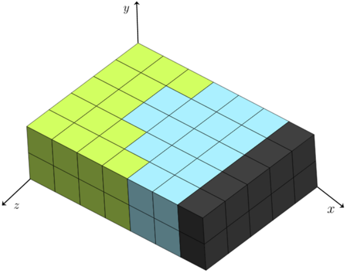

Volumetric Seam Carving: Graph Cut

Graph cut in an s/t network

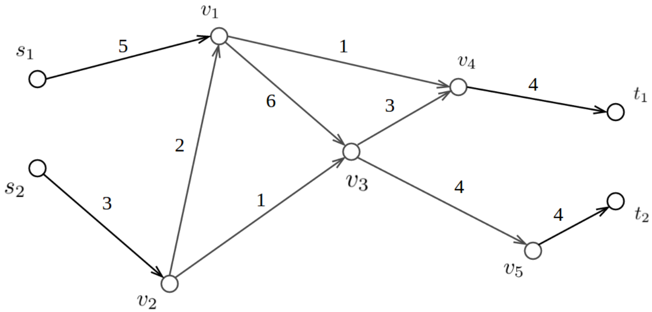

Volumetric Seam Carving: Graph Cut

Graph cut in an s/t network

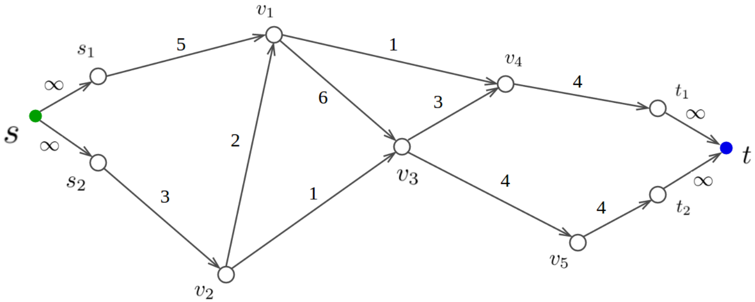

Graph cut in an s/t network

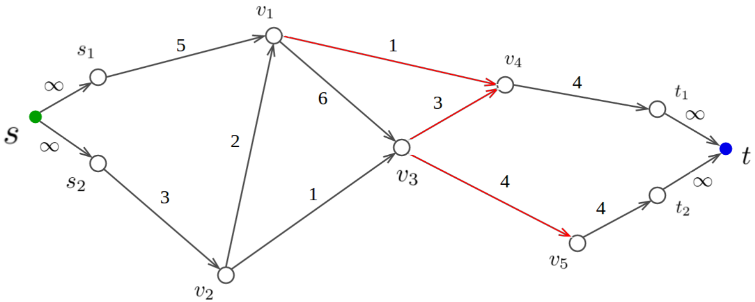

Direct outcome: a **partitioning of vertices into two disjoint subsets**.

Volumetric Seam Carving: Definitions

Image:

$$\bar{\mathbf{I}} = \Big\{ I(x,y,z) \;\Big|\; x \in [0,w), y \in [0,h), z \in [0,d) \Big\}$$

Energy function:

$$e_1(x,y,z) \;=\; \left|\frac{\partial}{\partial x} I(x,y,z)\right| + \left|\frac{\partial}{\partial y} I(x,y,z)\right| + \left|\frac{\partial}{\partial z} I(x,y,z)\right|$$

x-*sheet*:

$$\mathbf{s^x} \;=\; \{s^x_{jk}\} \;=\; \Big\{ \Big(\; x(j,k), \;\; j, \;\; k \; \Big) \Big\}$$

$s.t.\;\; j,\; k \in \mathbb{Z},\; 1 \leqslant j \leqslant h,\; 1 \leqslant k \leqslant d,$ and

$\mathbf{s^x}$ satisfies both *monoticity* and *connectivity* [RSA08]

Optimal x-*sheet*:

$$\mathbf{s^{x*}} \;=\; \min_{\mathbf{s^x}} \; \sum_{j=1}^{h} \sum_{k=1}^{d} e_1\Big(s^x_{jk}\Big)$$

Volumetric Seam Carving: Workflow

(1) Compute energy function field (volume)

Using the L_1-norm of gradient definition of energy, compute at each voxel position the energy value (a scalar). Takes $O(n)$ for both time and memory space where $n = w \cdot h \cdot d$.

(2) Construct graph (from voxels of the volume)

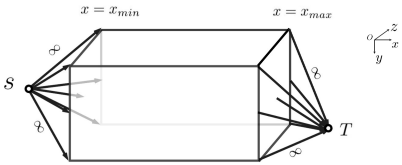

Volumetric Seam Carving: Workflow

Adding non-terminal arcs

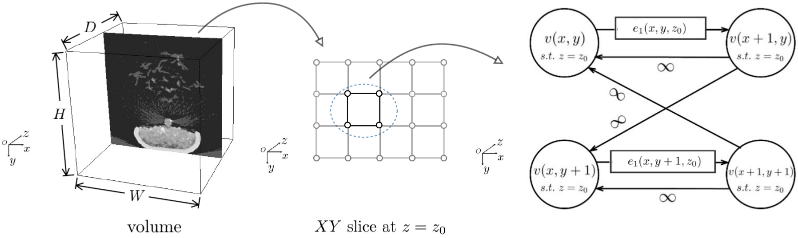

adding one node/arc $\approx$ constant time; $m \approx n$

totaling $O(n)$ for graph construction (time and memory)

Volumetric Seam Carving: Workflow

(3) Solve for the minimum cut

Yuri Boykov and Vladimir Kolmogorov. An experimental comparison of min-cut/max-flow algorithms for energy minimization in vision. *IEEE transactions on pattern analysis and machine intelligence*, 26(9):1124–1137, 2004.

- Algorithm with worst-case complexity $O\left(mn^2\ |C|\right)$
- "Significantly outperforms" other standard algorithms in practice
- **C++ library written and shared by the author**

Optimization:

Forward vs. backward energy

Volumetric Seam Carving: Forward Energy

Issue: joining two sides after seam removal

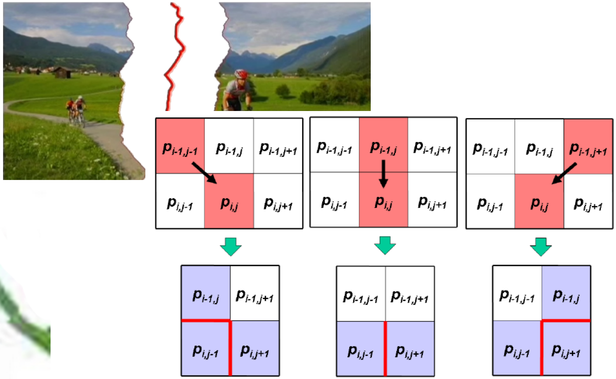

Volumetric Seam Carving: Forward Energy

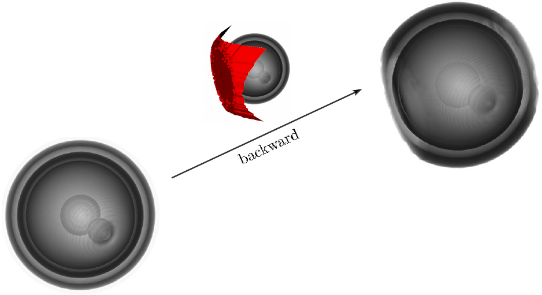

Volumetric Seam Carving: Forward Energy

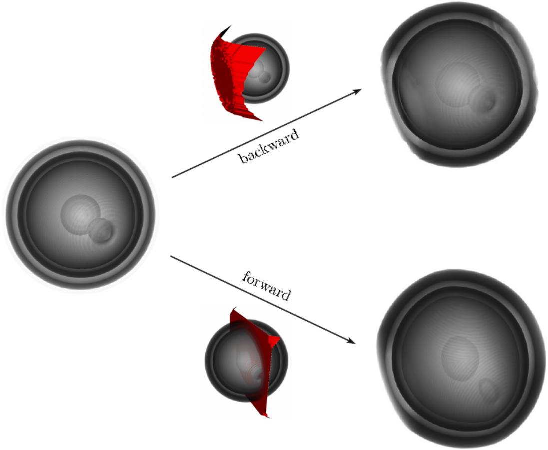

Volumetric Seam Carving: Forward Energy

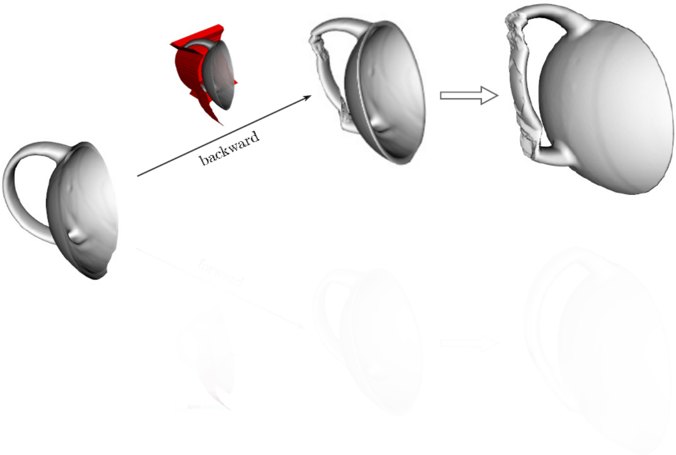

Volumetric Seam Carving: Forward Energy

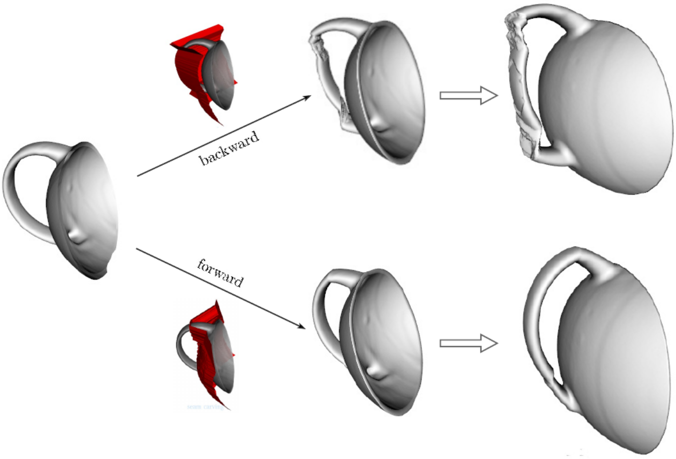

Volumetric Seam Carving: Isosurface Protection

Protect certain isosurface(s) as the volume "sheet carved"

Many of the *isosurface*s contain interesting and important features of the dataset. When downsizing a volume, we'd like to persist those features by finding a way to let the operator protect certain isosurface(s) of interest.

Extracting an isosurface $I^{-1}(c)$ from 3D grid data:

Marching Cubes, 1987 by Lorensen and Cline

Volumetric Seam Carving: Isosurface Protection

2D grid: Marching Squares

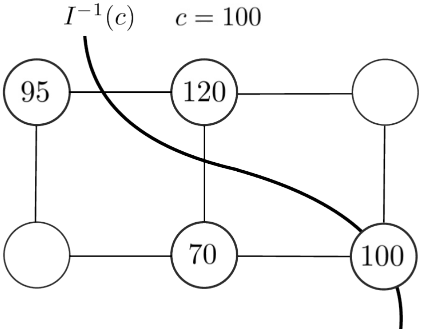

Volumetric Seam Carving: Isosurface Protection

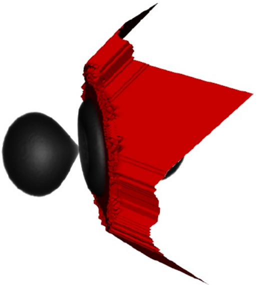

Volumetric Seam Carving: Isosurface Protection

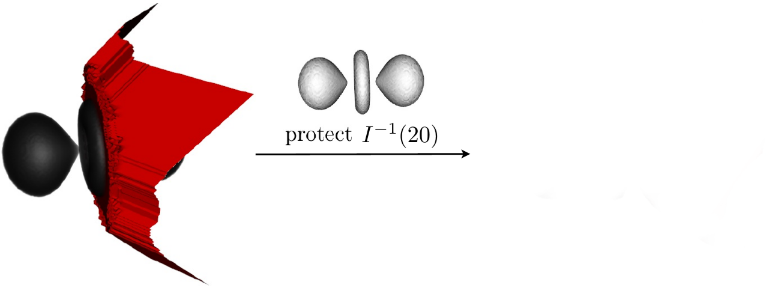

Volumetric Seam Carving: Isosurface Protection

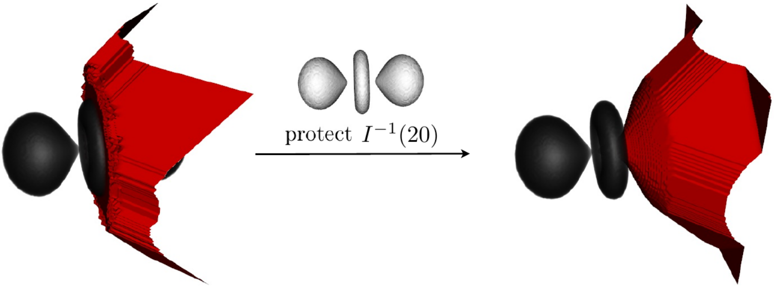

Volumetric Seam Carving: Isosurface Protection

Two-line algorithm:

PENALIZE-WEIGHT (c, E)

For each arc e in E do:

If $\Big(c - I\big(e.left\big)\Big) \cdot \Big(c - I\big(e.right\big)\Big) \leq 0$ do:

Set weight of arc e to ∞

Encoding Opacity Transfer Function

Volumetric Seam Carving: Encoding Opacity

- Main technique of visualization: *volume rendering*.

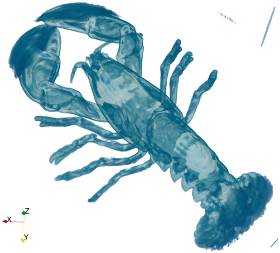

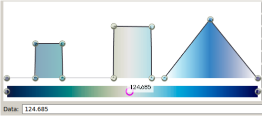

- Spatial perception, interactive control of visual representation.

- The center of volume rendering is *transfer function* design, especially the **opacity** transfer function.

Volumetric Seam Carving: Encoding Opacity

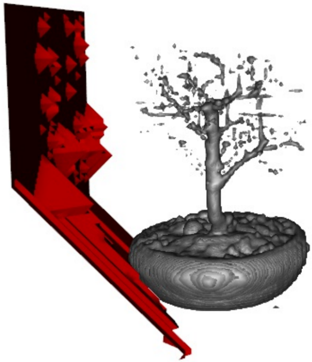

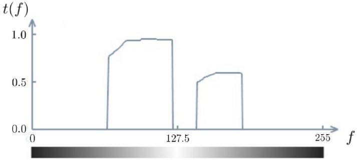

Actual data ← **mismatch** → viewed data

Volumetric Seam Carving: Encoding Opacity

To encode opacity transfer function into volumetric seam carving,

Notation	Range	Actual range	Meaning
t	$\mathbb{R} \to \mathbb{R}$	$[0, 255] \to [0, 1]$	opacity transfer function
I	$\mathbb{R}^3 \to \mathbb{R}$	$\mathbb{R}^3 \to [0, 255]$	scalar field (function) of the image
$t \circ I$	$\mathbb{R}^3 \to \mathbb{R}$	$\mathbb{R}^3 \to [0, 1]$	t compose I
$\big(t \circ I\big) \cdot I$	$\mathbb{R}^3 \to \mathbb{R}$	$\mathbb{R}^3 \to [0, 255]$	scalar field, weighted values
e	$\mathbb{R}^3 \to \mathbb{R}$	—	energy function

What we see is ~~I~~ $\big(t \circ I\big) \cdot I$

Options: (1) $e\big(t \circ I\big)$; (2) $e\big((t \circ I) \cdot I\big)$

Experimental Results

Experimental Results

Comprehensive test: 16 datasets

Tier	Name	Width	Height	Depth	N	Description
1	nucleon	41	41	41	68,921	Simulation of the two-body distribution probability of a nucleon in the atomic nucleus ^{16}O if a second nucleon is known to be positioned at $r' = (2f_m, 0, 0)$.
	silicium	64	34	34	73,984	Simulation of a silicium grid.
	fuel	64	36	36	82,944	Simulation of fuel injection into a combustion chamber. The higher the density value, the less presence of air.
	neghip	64	64	64	262,144	Simulation of the spatial probability distribution of the electrons in a high potential protein molecule.
2	hydro-genAtom	128	104	104	1,384,448	Simulation of the spatial probability distribution of the electron in an hydrogen atom, residing in a strong magnetic field.
	heart	131	104	104	2,307,303	MRI scan of human heart.
	statueLeg	166	221	93	3,411,798	CT scan of a leg of a bronze statue.
	lobster	253	253	56	3,584,504	CT scan of a lobster contained in a block of resin.
3	MRbrain	189	199	99	3,723,489	MRI scan of human brain.
	engine	157	216	128	4,340,736	CT scan of two cylinders of an engine block.
	golfball	163	166	171	4,626,918	CT scan of a golf ball.
	Boston-Teapot	241	139	178	5,962,822	CT scan of the SIGGRAPH 1989 teapot with a small version of the AVS lobster inside.
4	bonsai	181	221	256	10,240,256	CT scan of a bonsai tree.
	aneurism	216	226	241	11,464,656	Rotational b-plane x-ray scan of the arteries of the right half of a human head. A contrast agent was injected into the blood and an aneurism is present.
	foot	231	256	256	15,138,816	Rotational b-plane x-ray scan of a human foot. Tissue and bone are present in the dataset.
	skull	256	256	256	16,777,216	Rotational b-plane x-ray scan of phantom of a human skull.

Experimental Results

Task: reduce 1/3 of width (x-dimension) for all volumes.

Two other trival methods for comparison:

- Volumetric seam carving (forward energy, opacity encoded)
- Scaling with bilinear interpolation
- "Best slice": find and remove among the w YZ-slices the one with lowest total energy.

Experimental Results

Criteria of success: **visual quality**

1. **Local** shapes (e.g. tail fins of the lobster)
2. **Global** shape: persist content in the original image.

Volumetric seam carving outperforms

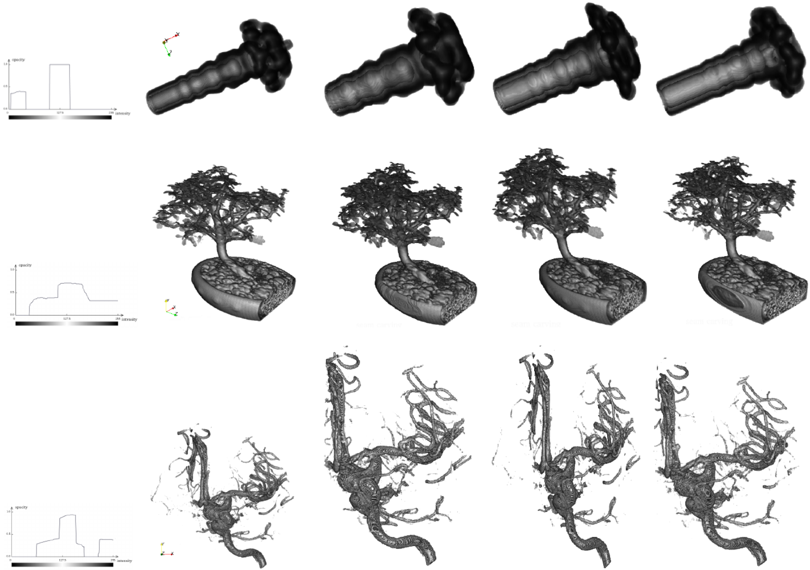

Volumetric seam carving is outperformed

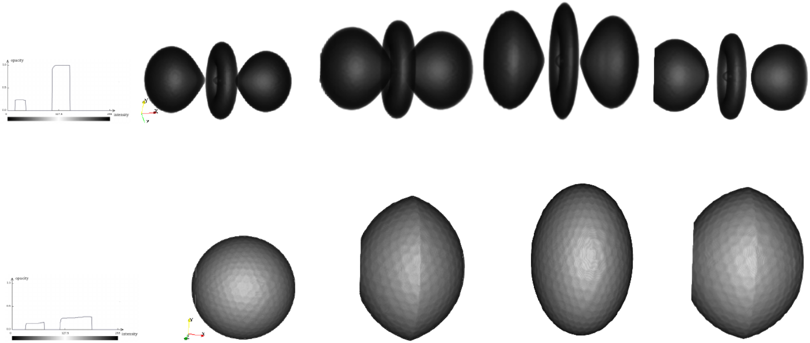

No “overall” advantage among three

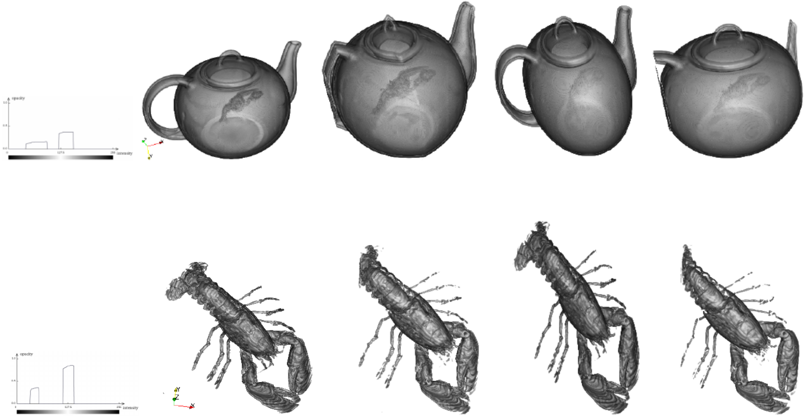

No “overall” advantage among three

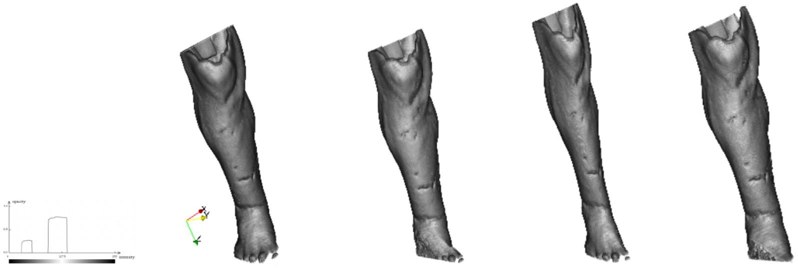

Experimental Results

$$e\big(t \circ I\big) \textbf{ versus } e\big((t \circ I) \cdot I\big) \text{ — “BostonTeapot”}$$

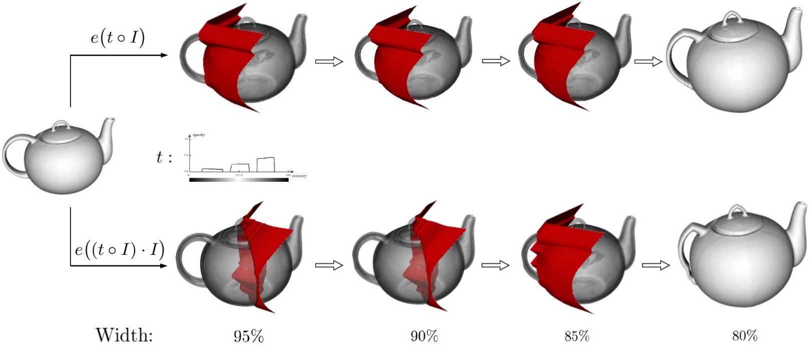

Conclusion

Conclusion: What We Learned

- **Flexibility in encoding various information, in order to enhance performance and match user's need** — most-important advantage of graph cut formulation, volumetric seam carving

- Encoding opacity: $e\Big(t \circ I\Big)$ instead of $e\Big((t \circ I) \cdot I\Big)$

Conclusion: Phases of Volume Reduction

1. **Cropping:** subset selection, can be manual or automatic (e.g. "best slice" approach)

2. **Specification of transfer function(s):** user-driven, central to the task of volume rendering (e.g. *semi-automatic generation of transfer functions*, G. Kindlmann 1998

3. **Sheet removals *exterior* to object:** as long as the sheet wraps but does not go through the object.

4. **Sheet removals *intersect* with object:** forward energy, isosurface protection turned on.

Stopping criteria: when loss of important structures is unavailable. Level of reduction applied should vary across datasets.

Conclusion: Limitation and Future Work

- **Biggest limitation:** lacking a well-defined and application oriented **evaluation metric**.
- Keeping track of topological features using the *Contour Trees*, to observe the change as volume is sheet carved, and finally encode this information into the operator.
- Segmentation along with *distance fields*, which may be useful in surgical planning where a sheet provides a partitioning of a subset of the volume.

Software System (demo)

Summary: Acknowledgement

Acknowledgement

- "Make Simple Tasks Simple!"
- ~~Popular~~ good science
- Dr. Wang, Dr. Duchowski
- Visual computing faculty, the department
- "*Trust recursion*"
- "To tell you the truth, $\cdots$"
- To ~~confirm~~ learn

Wilhelm Rontgen took this radiograph of his wife's left hand, shortly after his discovery of *X*-rays.

— December 22, 1895

Thank you

dachaos@g.clemson.edu

Q & A