GPU Accelerated Time-of-Flight Super-Resolution for Image-Guided Surgery

Jens Wetzl¹, Oliver Taubmann¹, Sven Haase¹, Thomas Köhler^{1,2}, Martin Kraus^{1,2}, Joachim Hornegger^{1,2}

BVM Heidelberg, 04.03.2013

¹ Pattern Recognition Lab, FAU Erlangen-Nuremberg
² Erlangen Graduate School in Advanced Optical Technologies (SAOT)







Outline

Introduction

Image-Guided Surgery Range Imaging

Methods

Super-Resolution Nonlinear Optimization

Experiments and Results

Summary and Outlook



Introduction

Image-Guided Surgery Range Imaging







Image-Guided Surgery

- Scan surface in operating room, register with preoperative CT/MR data
- Show tracked instruments aligned with CT/MR data
- Example application: Image-guided abdominal surgery



Image source: transplantsurg.wustl.edu



Range Imaging

- Surface information can be acquired with a Time-of-Flight camera
- Pros:
 - (+) Dense range data
 - (+) Real-time frame rates
- Cons:
 - (-) Low spatial resolution
 - (-) Poor signal-to-noise ratio



Range Imaging

• Surface information can be acquired with a Time-of-Flight camera

• Pros:

- (+) Dense range data
- (+) Real-time frame rates
- Cons:
 - (-) Low spatial resolution
 - (-) Poor signal-to-noise ratio



Super-Resolution Nonlinear Optimization







Super-Resolution

- Multi-frame super-resolution
 - Given: Sequence of preregistered low resolution images
 - Goal: Combine information in a single high resolution image



Idealized Sub-Pixel Motion Gain

• Utilizing **sub-pixel motion** improves effective resolution



Super-Resolution

- Multi-frame super-resolution
 - Given: Sequence of preregistered low resolution images
 - Goal: Combine information in a single high resolution image



Idealized Sub-Pixel Motion Gain

• Utilizing **sub-pixel motion** improves effective resolution



See: Pickup LC. Machine learning in multiframe image super-resolution. PhD Thesis, University of Oxford; 2008.

- Low-resolution (LR) images generated from ideal high-resolution (HR) image
- Affine motion, a Gaussian point spread function (PSF) and decimation





- Low-resolution (LR) images generated from ideal high-resolution (HR) image
- Affine motion, a Gaussian point spread function (PSF) and decimation





- Low-resolution (LR) images generated from ideal high-resolution (HR) image
- Affine motion, a Gaussian point spread function (PSF) and decimation





- Low-resolution (LR) images generated from ideal high-resolution (HR) image
- Affine motion, a Gaussian point spread function (PSF) and decimation





- Low-resolution (LR) images generated from ideal high-resolution (HR) image
- Affine motion, a Gaussian point spread function (PSF) and decimation





- Low-resolution (LR) images generated from ideal high-resolution (HR) image
- Affine motion, a Gaussian point spread function (PSF) and decimation





• Based on a forward model

- Low-resolution (LR) images generated from ideal high-resolution (HR) image
- Affine motion, a Gaussian point spread function (PSF) and decimation



Inversion of that model combined with a smoothness prior

- Formulated as a nonlinear optimization problem
- Yields a maximum a posteriori (MAP) estimate



Objective Function

$$\boldsymbol{x}^* = \underset{\boldsymbol{x}}{\operatorname{argmin}} \left(\|\boldsymbol{W}\boldsymbol{x} - \boldsymbol{y}\|_2^2 + \lambda \cdot \|\mathbf{h}_{\delta}(\boldsymbol{D}\boldsymbol{x})\|_1 \right)$$
(1)

- **W** System matrix, models image generation
- **x** HR image, linearized

- Wx y Residuum to model prediction
 - λ Prior strength, adjusts smoothness
 - $h_{\delta}(\cdot)$ Pseudo-Huber loss function, applied element-wise
 - **Dx** Discrete Laplacian of the ideal image



System Overview



Computation of system matrix, evaluation of objective function and nonlinear optimization all done on the GPU using NVIDIA CUDA.



Sparse System Matrix – Challenges

• $\boldsymbol{W} \in \mathbb{R}^{m \times n}$ is prohibitively large

$$m \ = \ \#_{ ext{LR pixels}} \cdot \#_{ ext{LR images}}$$

 $n \ = \ \#_{ ext{HR pixels}}$

- ightarrow Typical size e.g. (200² \cdot 10) imes 400² = 64 \cdot 10⁹ elements
- \rightarrow Set point spread function to zero beyond 3 $\cdot\,\sigma$
- ightarrow becomes sparse and can be stored in (GPU) memory
- Computing MULTIPLY and TRANSPOSE-MULTIPLY efficiently

$$r(\boldsymbol{x}) = \|\boldsymbol{W}\boldsymbol{x} - \boldsymbol{y}\|^{2} \qquad ($$

$$\nabla r(\boldsymbol{x}) = 2\boldsymbol{W}^{T}(\boldsymbol{W}\boldsymbol{x} - \boldsymbol{y}) \qquad ($$

 \rightarrow Store both CRS* and CCS* representations of W

* Compressed {Row, Column} Storage

(2)

(3)



Sparse System Matrix – Challenges

• $\boldsymbol{W} \in \mathbb{R}^{m \times n}$ is prohibitively large

$$m = \#_{ ext{LR pixels}} \cdot \#_{ ext{LR images}}$$

 $n = \#_{ ext{HR pixels}}$

- ightarrow Typical size e.g. $(200^2 \cdot 10) \times 400^2 = 64 \cdot 10^9$ elements
- \rightarrow Set point spread function to zero beyond 3 $\cdot\,\sigma$
- \rightarrow W becomes sparse and can be stored in (GPU) memory
- Computing MULTIPLY and TRANSPOSE-MULTIPLY efficiently

$$r(\boldsymbol{x}) = \|\boldsymbol{W}\boldsymbol{x} - \boldsymbol{y}\|^{2}$$
(4)

$$\nabla r(\boldsymbol{x}) = 2\boldsymbol{W}^{T}(\boldsymbol{W}\boldsymbol{x} - \boldsymbol{y})$$
(5)

 \rightarrow Store both CRS* and CCS* representations of W

* Compressed {Row, Column} Storage

(2)

(3)



Nonlinear Optimization

- Objective function gradient **nonlinear** in HR image pixels
 - \rightarrow Iterative descent method for optimization
- We chose L-BFGS (limited-memory Broyden-Fletcher-Goldfarb-Shanno)
 - Quasi-Newton variant, high convergence rate
 - Stores low-rank approximation of inverse Hessian
 - \rightarrow Low memory footprint, well suited for GPU implementation
- Implementation of a **general-purpose library** for unconstrained nonlinear optimization on the GPU



Nonlinear Optimization

- Objective function gradient **nonlinear** in HR image pixels
 - \rightarrow Iterative descent method for optimization
- We chose L-BFGS (limited-memory Broyden-Fletcher-Goldfarb-Shanno)
 - Quasi-Newton variant, high convergence rate
 - Stores low-rank approximation of inverse Hessian
 - \rightarrow Low memory footprint, well suited for GPU implementation
- Implementation of a **general-purpose library** for unconstrained nonlinear optimization on the GPU



Experiments and Results







Experiments



Photo of a porcine liver we measured with a PMD CamCube 3.0



Qualitative Results





HR frame (10 input frames, $2 \times$ upsampling factor)



Qualitative Results





Mesh from LR frame

Mesh from HR frame (10 input frames, $2 \times$ upsampling factor)



Quantitative Results

	Raw	SR
RMSE	8.11 mm	6.29 mm
Median	4.92 mm	3.16 mm
Std. dev.	\pm 0.34 mm	\pm 0.16 mm

Error statistics for registered meshes (Parameters: 10 input frames, 2× upsampling factor)

- RMSE: Root-mean-square error
- Median: Median of absolute distances
- Std. dev.: Standard deviation of absolute distances

 \Rightarrow RMSE improved by more than 20%



Quantitative Results



(GPU: NVIDIA GTX 580, input frame size: 200 × 200)

Measured runtimes for varying

- sequence lengths (with $2 \times$ upsampling factor)
- upsampling factors (with 10 input frames)









• Super-resolution framework (including optimizer) on the GPU

- Comprehensive image generation model
- MAP estimate (\rightarrow smoothness prior)
- Interactive frame rates
- What's left to do?
 - Integrate robust affine registration
 - Combine registration and super-resolution
- Our code is available under an open source license (CC-BY 3.0)
 - \Rightarrow Feel free to download it from our website and experiment!

www5.cs.fau.de/map-superresolution (coming soon)



• Super-resolution framework (including optimizer) on the GPU

- Comprehensive image generation model
- MAP estimate (\rightarrow smoothness prior)
- Interactive frame rates
- What's left to do?
 - Integrate robust affine registration
 - Combine registration and super-resolution
- Our code is available under an open source license (CC-BY 3.0)
 - \Rightarrow Feel free to download it from our website and experiment!

www5.cs.fau.de/map-superresolution (coming soon)



• Super-resolution framework (including optimizer) on the GPU

- Comprehensive image generation model
- MAP estimate (\rightarrow smoothness prior)
- Interactive frame rates
- What's left to do?
 - Integrate robust affine registration
 - Combine registration and super-resolution
- Our code is available under an open source license (CC-BY 3.0)
 - \Rightarrow Feel free to download it from our website and experiment!

www5.cs.fau.de/map-superresolution (coming soon)



Acknowledgements

We gratefully acknowledge funding of the **Erlangen Graduate School in Advanced Optical Technologies** (SAOT) by the **German Science Foundation** (DFG) in the framework of the excellence initiative, as well as the support by the DFG under Grant No. HO 1791/7-1. This research was funded/supported by the **Graduate School of Information Science in Health** (GSISH) and the TUM Graduate School.





Thanks for listening!



Qualitative Results (Addenda)



Preprocessed LR frame



HR frame (10 raw frames, $2 \times$ upsampling factor)