

Dense Visual Odometry

Christoph Ihrke, Josef Brandl, Duy Nguyen

Technische Universität München
Department of Informatics
Computer Vision Group

October 6, 2015

Outline

- 1 Introduction
- 2 Algorithm
- 3 Implementation
- 4 Performance
- 5 Conclusion

Outline

- 1 Introduction
- 2 Algorithm
- 3 Implementation
- 4 Performance
- 5 Conclusion



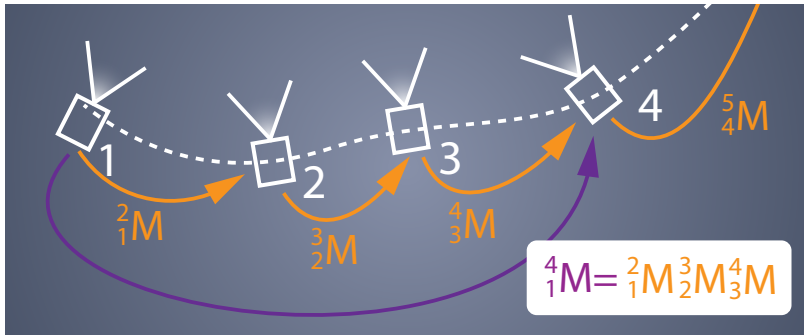
Introduction

- Dense Visual Odometry
 - Visual Odometry: estimate position from images.
 - Dense: every pixel \rightarrow feature
- Motivation
 - often required, e.g: for volumetric reconstruction
 - current CPU implementation can be parallelized.


Outline

- 1 Introduction
- 2 Algorithm**
- 3 Implementation
- 4 Performance
- 5 Conclusion

Overview

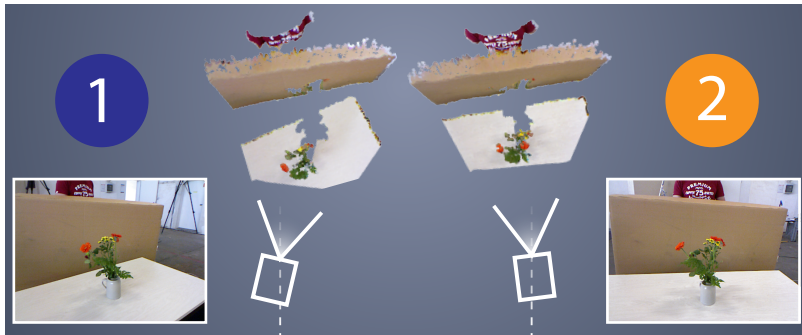


Pyramid approach




- Downsample images.
- Process coarser to finer levels.

Consecutive frame alignment



$$E(\xi) = \sum_{\mathbf{x} \in \Omega} (I_{\text{prev}}(\mathbf{x}) - I_{\text{cur}}(\omega(\mathbf{x}, D_{\text{prev}}(\mathbf{x}), \xi)))^2 \quad (1)$$

Consecutive frame alignment



- Non-linear least squares problem: $\min E(\xi) \rightarrow \xi$
- Algorithms:
 - Gradient descent
 - Gauss-Newton
 - Levenberg-Marquardt
- Robust weights:
 - eliminate outliers
 - Huber, T-Dist

Outline

- 1 Introduction
- 2 Algorithm
- 3 Implementation**
- 4 Performance
- 5 Conclusion



GPU implementation

- Custom kernels:
 - Downsampling
 - Residual & Jacobian:
 - Numeric Jacobian: approximate derivatives
 - Analytic Jacobian: closed form
- cuBLAS for large matrix operations:
 - $A = J^T J$ ($6 \times n \cdot n \times 6$)
 - $b = J^T r$ ($6 \times n \cdot n \times 1$)
- Eigen for some small matrix operations

Lessons learned

- Custom matrix struct with operations on GPU
 - Avoid calling Eigen (in CPU)
 - No memory copying back and forth
- Matrices stored in column major
 - Consistent memory storage: cuBLAS, Eigen and our own
- Use existing library such as cuBLAS for common problem
- **cublasSgemm()** is faster than **cublasSgemv()**

Outline


- 1 Introduction
- 2 Algorithm
- 3 Implementation
- 4 Performance**
- 5 Conclusion

Performance

■ CPU vs. GPU:

	CPU	GPU
Absolute Trajectory Error (ATE)	~0.07	~0.06
Avg. time per frame (ms)	~150	~20

■ Freiburg 1 dataset:





Outline

- 1 Introduction
- 2 Algorithm
- 3 Implementation
- 4 Performance
- 5 Conclusion**

Conclusion

- Problem was easily parallelizable
- Noticable speed-up with GPU

DEMO