Depth-Adaptive Superpixels

Jan Möller Robert Posch

Technische Universität München
Department of Informatics
Computer Vision Group

October 5, 2015

Outline

1. Introduction
2. Superpixel Algorithms
3. CUDA Implementation
4. Conclusion and Future Work

Outline

1 Introduction

2 Superpixel Algorithms

3 CUDA Implementation

4 Conclusion and Future Work

Introduction

- Subdivide image into meaningful regions (superpixels)
- Performance criteria:
  - Spatial proximity
  - Color similarity
  - Adherence to edges
  - Structural resemblance
  - Run time
- Approach by modified k-means algorithm

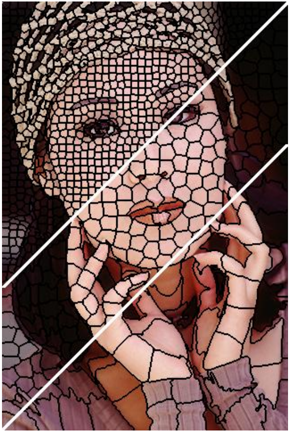

[Achanta et al., 2012]

Outline

1 Introduction

2 Superpixel Algorithms

3 CUDA Implementation

4 Conclusion and Future Work

Simple Linear Iterative Clustering (SLIC)

- Generate a cluster seed for each superpixel
- Cluster seeds have position and color
- For each pixel compute distance to closest cluster
- Distance function has terms for
  - spatial distance
  - color distance
- Re-compute cluster-position and -color by averaging over affiliated pixels
- Iterate with new values
- Speed-up by only considering small window around cluster

[Achanta et al., 2012]

Simple Linear Iterative Clustering (SLIC)

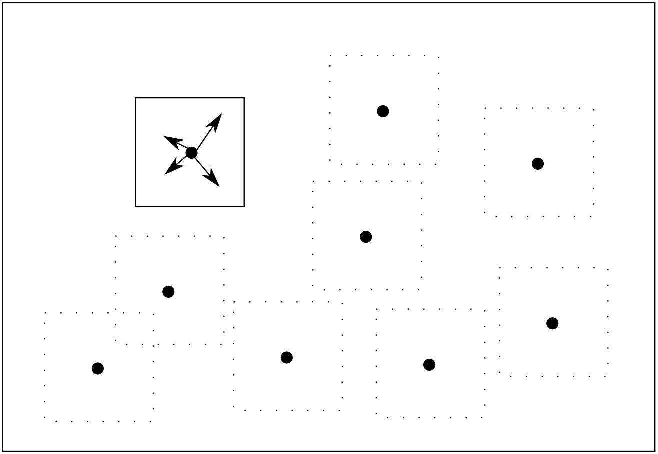

Search window during distance computation

Simple Linear Iterative Clustering (SLIC)

- Pro:
  - Produces similar-sized superpixels
  - Good color adherence
  - Faster than similar algorithms
- Con:
  - Needs post-processing step to make sure pixels are connected with cluster center
  - No structural resemblance

[Achanta et al., 2012]

Depth-Adaptive Superpixels

- Incorporates 3D information into algorithm
- Computes density image from depth gradients
  - High density in distant and perspectively deformed areas
  - Cluster seeds are distributed with respect to density
- Spawns more cluster seeds in high-density regions
- Computes normal vectors from depth gradients
- Adds normal-term to distance function
- Update step similar to SLIC

[Weikersdorfer et al., 2012, Weikersdorfer, 2013]

Overview

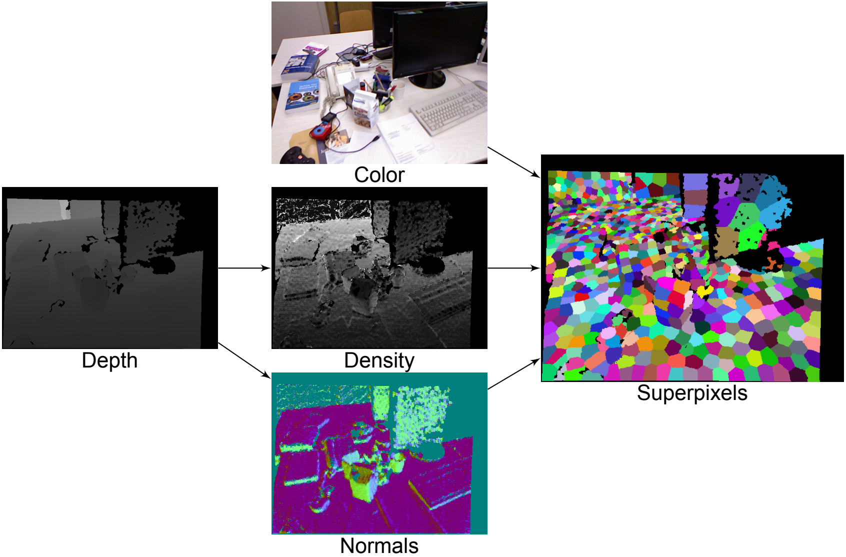

Depth-Adaptive Superpixels

- Pro:
  - Superpixel shape respects 3D objects
  - Doesn't completely rely on color
- Con:
  - Non-trivial cluster generation
  - Speed deficit

[Weikersdorfer et al., 2012, Weikersdorfer, 2013]

Outline

1 Introduction

2 Superpixel Algorithms

3 CUDA Implementation

4 Conclusion and Future Work

GPU Implementation

- Density computation is well posed for parallelization
- Algorithm used for seed generation inherently sequential
  - Uses Floyd-Steinberg dithering
  - By computing on CPU additional memory transfers necessary
  - Can be hidden by parallel execution
- Distance computation challenging
  - Start a kernel for each cluster in parallel
  - Dynamically calculate search window size
  - Race conditions with overlapping windows

Reminder: SLIC

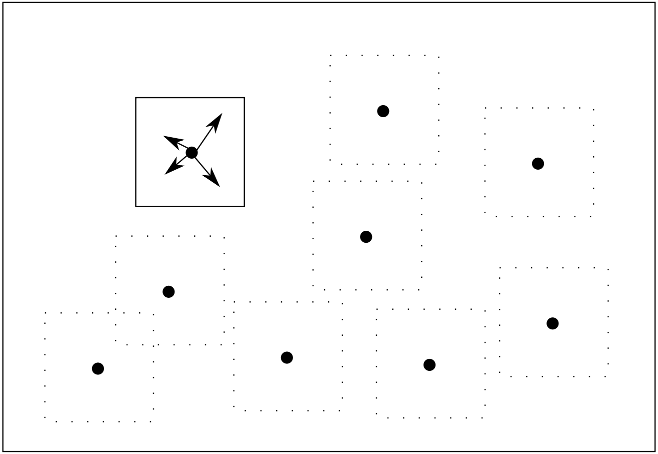

Search window during distance computation

Race conditions

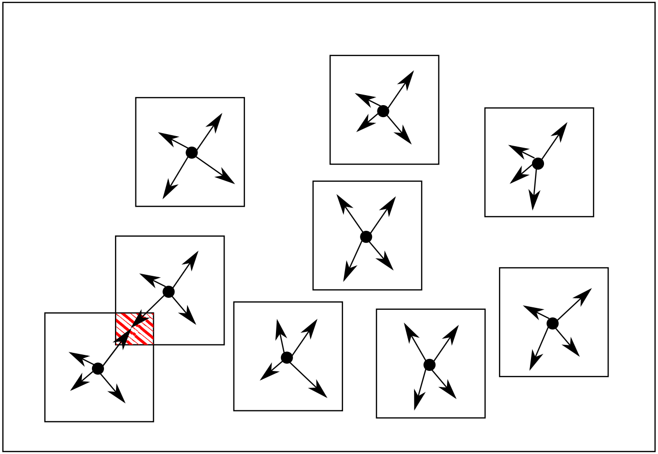

Overlapping windows cause race conditions

Use of atomicMin

- Use atomicMin to avoid race conditions
  - Cluster index and distance need to be set at once
  - No mutexes in CUDA
- atomicMin only defined for integer data types
- Align distance and index such that they can be represented by a single 64-bit integer
- First four bytes determine the outcome of minimum operation
  - Store distance into the first four bytes
  - Store cluster index into last four bytes

Live Demo

Outline

1 Introduction

2 Superpixel Algorithms

3 CUDA Implementation

4 Conclusion and Future Work

Performance Comparision

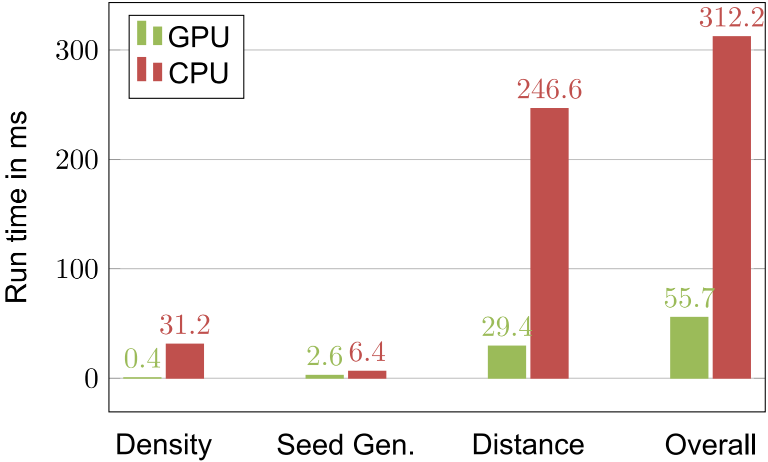

Conclusion

- Atomics don't necessarily have a big performance impact
- Simple algorithms may not easily transfer to massively parallel architectures
- Organization of spatial data can be challenging if it doesn't follow a trivial alignment

Future Work

- Ensure pixels are connected to affiliated cluster center
- Cluster seed generation on GPU
- Further optimization

Bibliography I

[Achanta et al., 2012] Achanta, R., Shaji, A., Smith, K., Lucchi, A., Fua, P., and Süsstrunk, S. (2012).
Slic superpixels compared to state-of-the-art superpixel methods.
IEEE Transactions on Pattern Analysis and Machine Intelligence.

[Weikersdorfer, 2013] Weikersdorfer, D. (2013).
Efficiency by Sparsity: Depth-Adaptive Superpixels and Event-based SLAM.
Dissertation, Technische Universität München.

[Weikersdorfer et al., 2012] Weikersdorfer, D., Gossow, D., and Beetz, M. (2012).
Depth-adaptive superpixels.
IEEE International Conference on Pattern Recognition.