


July 14<sup>th</sup> 2010  
Daniel Kennedy

## Large Displacement Optical Flow



Paper by Thomas Brox, Christopher Bregler and Jitendra Malik

Presented at CVPR 2009

# Outline

Overview of the method

Region matching

Calculating the flow

Results

# Overview



Comparison of the two main ways of estimating optical flow:

## Lucas-Kanade Optical Flow

Obtain very dense, accurate flow fields  
However downsampling can smooth too much and cause a loss of information. This favors the movements of larger objects over smaller ones

## Descriptor matching

We can track the salient parts of the image  
However, structure is lost and we're at the mercy of outliers.  
A sparse flow field is created and interpolating the results leads to inaccurate results overall



# Overview

## Their method:

Use a directed variational optical flow:

Do a hierarchical segmentation on the image

After segmenting the image, calculate the flow taking into account the information received from the segmentation

## The benefits of such an approach:

Regions in the segmented image are likely to coincide with objects

We obtain a dense field instead of interpolating a sparse field

The whole image is covered

Use of multiple matching hypotheses takes care of outliers  
(more details later)

# Region Matching

## Compute the regions:

- Uses a boundary detector:

  - Takes into account texture

  - Creates a hierarchy in the boundaries detected

  - Boundaries are thus more likely to correspond to objects or parts of objects

## Region descriptors:

- To each region they fit two descriptors:

  - One that takes into account orientation (16 orientation histograms)

  - The other to take into account color (mean RGB color of the 16 sub-parts)

  - Correspondences between regions is thus done by taking the Euclidean distances of both descriptors

  - Correspondences are then thresholded according to given parameters to exclude false matches

# Region Matching

## Refinement of the descriptor matching:

Despite thresholding the descriptors, outliers still remain

This is due to the small size of these regions which makes the region hard to describe

Need to create a confidence measure for potential correspondences

They take into account the actual distance between two patches as well as the deformation

Note: instead of computing the confidence measure for each pair, the algorithm chooses the 10 nearest neighbors from the previous step

# Flow Calculation

The bulk of the work has been done:


The descriptor matching and its refinement is the main focus of the paper

The flow is calculated by minimizing equation (5)


This equation takes into account the data constraints and correspondence calculations shown in the previous section

# Results


The results are an improvement over previous techniques:



Interpolated region correspondences



Coarse-to-fine Method from [4]




Proposed method



Two frames compared



# Results



The process has trouble with repeating backgrounds

Some outliers still remain

Some very fast moving objects are still overlooked