




# COMP 141: Probabilistic Robotics for Human-Robot Interaction

Instructor: Jivko Sinapov

[www.cs.tufts.edu/~jsinapov](http://www.cs.tufts.edu/~jsinapov)

# Image Features and Optical Flow





# Reading Assignment

- Chapter 5 of Probabilistic Robotics

# Research Article Presentation

- Sign-up for research article presentation



# Robotics and AI Conferences


- IEEE International Conference on Robotics and Automation (ICRA)
- IEEE International Conference on Intelligent Robots (IROS)
- IEEE International Conference on Development and Learning (ICDL)
- Robotics Science and Systems (RSS)

# Robotics Journals

- IEEE Transactions on Robotics (TRO)
- IEEE Transactions on Autonomous Mental Development (TAMD)
- International Journal of Robotics Research (IJRR)
- Robotics and Autonomous System (RAS)




# Visual Registration and Recognition





# Visual Registration and Recognition



# Registration




# Registration




[[http://vihari.github.io/personal\\_website/images/3dregistration.png](http://vihari.github.io/personal_website/images/3dregistration.png)]



# Registration



# Interest Point Registration






# Interest Point Detection

- Look for image regions that are unusual
  - Leads to unambiguous matches in other images
  - How do we define unusual?

# Suppose we only consider a small window of pixels


- What defines whether a feature is a good or bad candidate?






# Local measure of feature uniqueness


- How does the window change when you shift it?
- Shifting the window in *any direction* causes a *big change*



“flat” region:  
no change in all  
directions




“edge”:  
no change along  
the edge direction



“corner”:  
significant change  
in all directions

Consider shifting the window  $W$  by  $(u,v)$

- how do the pixels in  $W$  change?
- compare each pixel before and after by summing up the squared differences (SSD)
- this defines an SSD “error” of  $E(u,v)$ :



# Harris Detector: Mathematics

Change of intensity for the shift  $[u,v]$ :


$$E(u,v) = \sum_{x,y} w(x,y) \left[ I(x+u, y+v) - I(x,y) \right]^2$$

Window  
function

Shifted  
intensity


Intensity

Window function  $w(x,y) =$




1 in window, 0 outside

or




Gaussian






Eliminate small responses.





Find local maxima of the remaining.






# OpenCV: finding features

`cv::cornerHarris(...)`


[http://docs.opencv.org/2.4/doc/tutorials/features2d/trackingmotion/harris\\_detector/harris\\_detector.html](http://docs.opencv.org/2.4/doc/tutorials/features2d/trackingmotion/harris_detector/harris_detector.html)



# OpenCV: finding features


Shi and Tomasi '94: `cv::goodFeaturesToTrack()`

[http://docs.opencv.org/2.4/modules/imgproc/doc/feature\\_detection.html](http://docs.opencv.org/2.4/modules/imgproc/doc/feature_detection.html)



# Harris Detector: Some Properties

- Rotation invariance



Ellipse rotates but its shape remains the same

*Corner response  $R$  is invariant to image rotation*





# Harris Detector: Some Properties

- Partial invariance to *affine intensity* change


- ✓ Only derivatives are used => invariance to intensity shift  $I \rightarrow I + b$

- ✓ Intensity scale:  $I \rightarrow a I$



# Harris Detector: Some Properties

- But: non-invariant to *image scale*!





All points will be  
classified as **edges**



**Corner !**

# Achieving Scale Invariance

- How do we choose scale?



# Difference-of-Gaussians




# Gaussian Blur

$$\frac{1}{16}$$

1	2	1
2	4	2
1	2	1

3x3 Gaussian Kernel



Computation of the Output Image



# Gaussian Blur Kernels

1/16

1	2	1
2	4	2
1	2	1


1/273

1	4	7	4	1
4	16	26	16	4
7	26	41	26	7
4	16	26	16	4
1	4	7	4	1

1/1003


0	0	1	2	1	0	0
0	3	13	22	13	3	0
1	13	59	97	59	13	1
2	22	97	159	97	22	2
1	13	59	97	59	13	1
0	3	13	22	13	3	0
0	0	1	2	1	0	0

# Finding Keypoints - Scale, Location



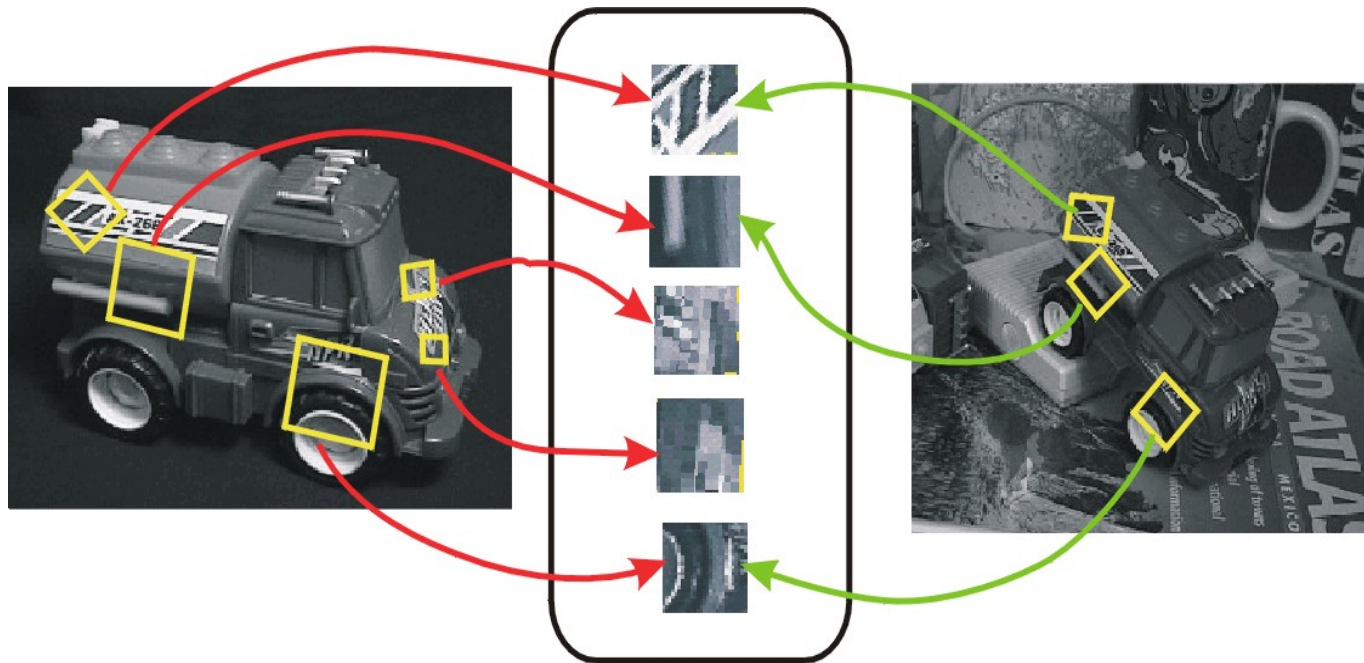


# SIFT descriptor



# Interest Point Descriptors

- Now that we can find interest points, how do we compare them?






# Interest Point Descriptors

- Now that we can find interest points, how do we compare them?
- Answer: compute a numerical feature descriptor describing the orientation, and scale of the interest point

# SIFT descriptor

Basic idea:

- Take 16x16 square window around detected feature
- Compute edge orientation (angle of the gradient - 90 ) for each pixel
- Create histogram of edge orientations




Adapted from slide by David Lowe

# SIFT descriptor


## Full version

- Divide the 16x16 window into a 4x4 grid of cells (2x2 case shown below)
- Compute an orientation histogram for each cell
- $16 \text{ cells} * 8 \text{ orientations} = 128 \text{ dimensional descriptor}$



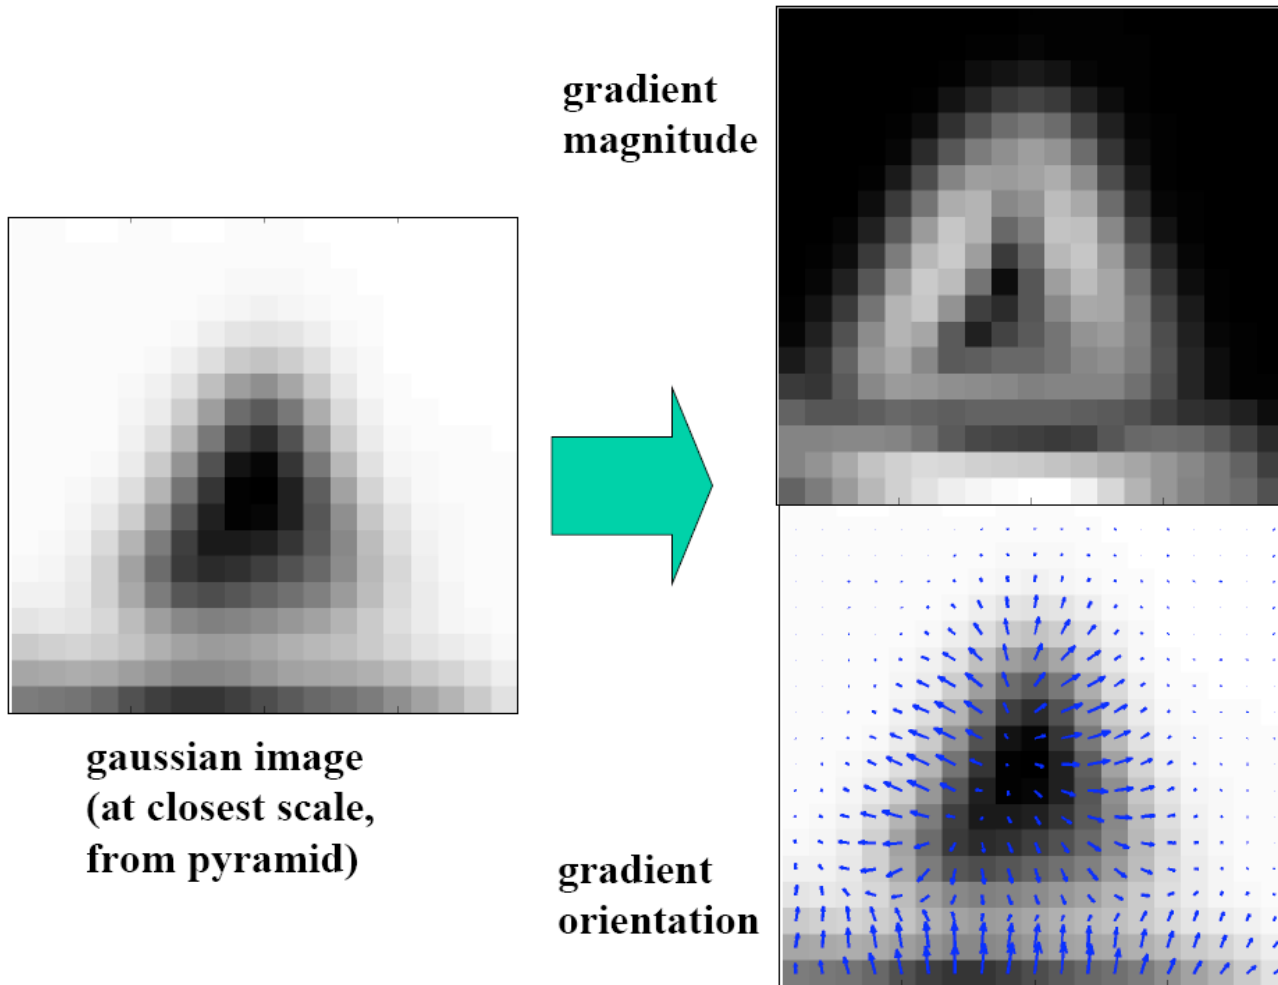
Adapted from slide by David Lowe

# SIFT descriptor




- **Keypoint location = extrema location**
- **Keypoint scale is scale of the DOG image**


# SIFT descriptor



# Computing Angle of Gradient



Angle and magnitude of gradient are computed using 1 and 2-side edge filters:




Patch




# SIFT descriptor

**weighted gradient  
magnitude**




**gradient  
orientation**



**weighted orientation histogram.**

Each bucket contains sum of weighted gradient magnitudes corresponding to angles that fall within that bucket.




**36 buckets**

**10 degree range of angles in each bucket, i.e.**

**$0 \leq \text{ang} < 10$  : bucket 1**

**$10 \leq \text{ang} < 20$  : bucket 2**

**$20 \leq \text{ang} < 30$  : bucket 3 ...**




# The problem with SIFT...


- Slow...
- Copyrighted!
  - ❑ Alternatives: SURF
  - ❑ OpenSURF:
    - <http://opensurf1.googlecode.com/files/OpenSURF.pdf>
  - ❑ Included in OpenCV 2.0+
  - ❑ OpenCV Tutorial:
    - <http://achuwilson.wordpress.com/2011/08/05/object-detection-using-surf-in-opencv-part-1/>

# SURF

- Achieves quicker computation by scaling the filter rather than the image:



**SIFT**



**SURF**

# To summarize...

- Feature detectors:
  - Find interest points in image (e.g., using difference of Gaussians, Harris corner detection, etc.)
- Feature descriptors
  - Each detected feature can be represented by a numerical descriptor encoding orientation, scale, etc.


# Applications



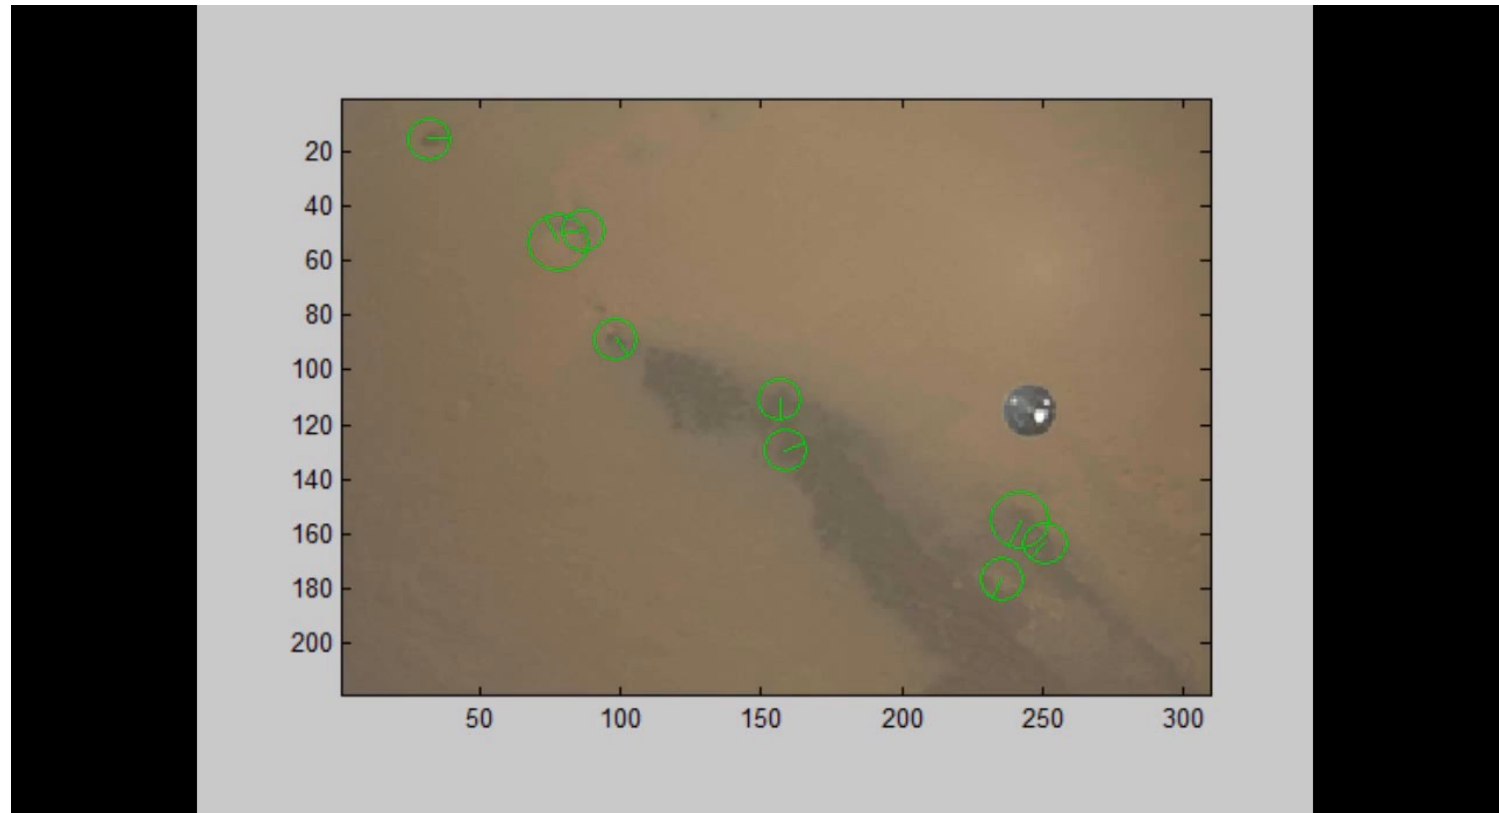
# Object Detection




# Marker-less tracking




# SURF feature tracking during Curiosity Landing



# Visual-odometry



# Image Registration in ROS



[http://wiki.ros.org/find\\_object\\_2d](http://wiki.ros.org/find_object_2d)

# Optical Flow

- Interest key points and feature descriptors are great but suffer from one limitation:
  - They ignore time



# Why Optical Flow?



Not grouped




Proximity




Similarity




Similarity



Common Fate




Common Region




Parallelism



Symmetry



Continuity



Closure

# Why Optical Flow?



Not grouped




Proximity




Similarity



Similarity



Common Fate




Common Region




Parallelism



Symmetry




Continuity




Closure


# Why Optical Flow?



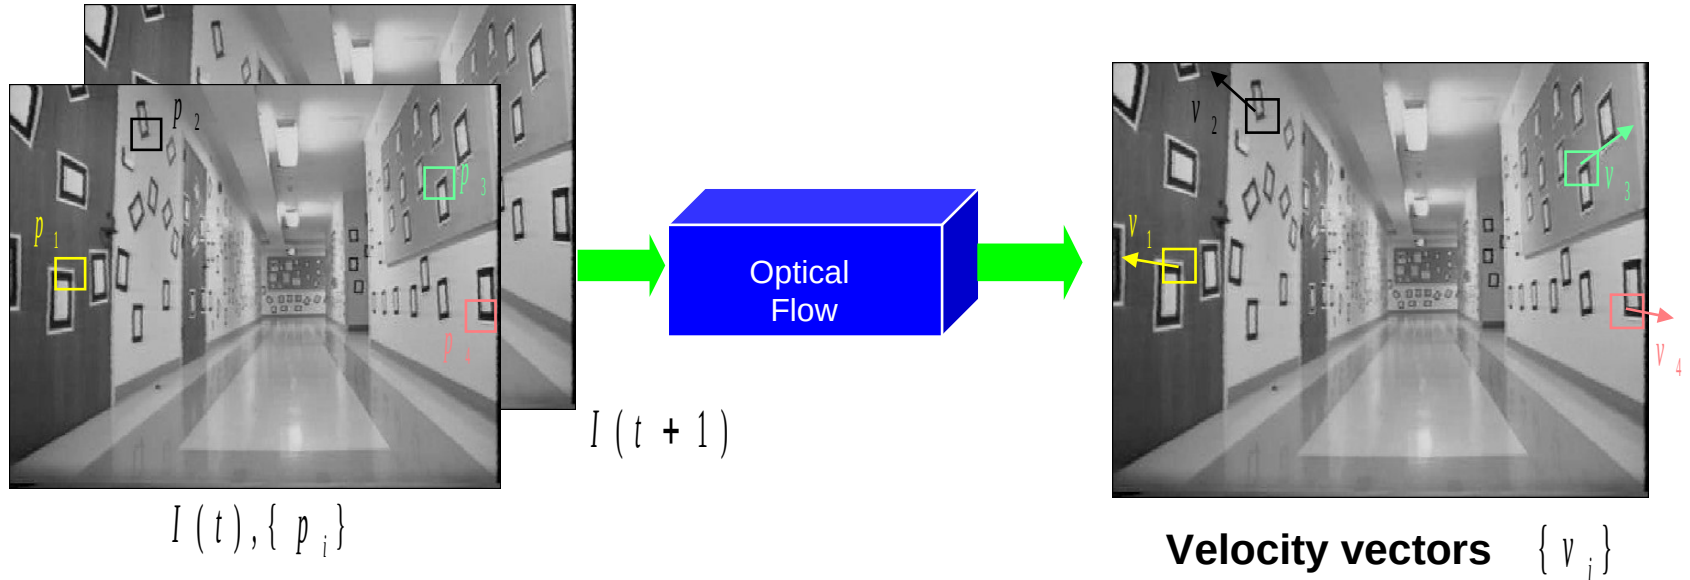
# Why Optical Flow?



# Optical Flow Video



# What is Optical Flow?





“Optical flow is the distribution of apparent velocities of movement of brightness patterns in an image”

- Horn and Schunk, 1981


# Motion Fields

- The motion field is the projection of the 3D scene motion into the image





# Motion Fields and Camera Movement




Length of flow vectors inversely proportional to depth  $Z$  of 3D point

points closer to the camera move more quickly across the image plane

Figure 1.2: Two images taken from a helicopter flying through a canyon and the computed optical flow field.

# Visual-odometry for Drones




# Computing Optical Flow

- Given a set of points in an image, find those same points in another image
- Or, given point  $[u_x, u_y]^T$  in image  $I_1$  find the point  $[u_x + \delta_x, u_y + \delta_y]^T$  in image  $I_2$  that minimizes  $\varepsilon$ :

$$\varepsilon(\delta_x, \delta_y) = \sum_{x=u_x-w_x}^{u_x+w_x} \sum_{y=u_y-w_y}^{u_y+w_y} (I_1(x, y) - I_2(x + \delta_x, y + \delta_y))$$

# Optical Flow Assumptions




## Assumption

Image measurements (e.g. brightness) in a small region remain the same although their location may change.

$$I(x + u, y + v, t + 1) = I(x, y, t)$$

(assumption)


# Spatial Coherence



## Assumption

- \* Neighboring points in the scene typically belong to the same surface and hence typically have similar motions.
- \* Since they also project to nearby points in the image, we expect spatial coherence in image flow.



# Temporal Persistence



Assumption:

The image motion of a surface patch changes gradually over time.

# Dense vs. Sparse Optical Flow





# Code

## ■ MATLAB:

### □ Iterative Pyramidal LK Optical Flow

□ <http://www.mathworks.com/matlabcentral/fileexchange/23142-iterative-pyramidal-lk-optical-flow>

## ■ OpenCV

□ [http://robots.stanford.edu/cs223b05/notes/optical\\_flow\\_demo.cpp](http://robots.stanford.edu/cs223b05/notes/optical_flow_demo.cpp)

# To summarize...

- Feature detectors:
  - Find interest points in image (e.g., using difference of Gaussians, Harris corner detection, etc.)
- Feature descriptors
  - Each detected feature can be represented by a numerical descriptor encoding orientation, scale, etc.

# To summarize...

## ■ Optical Flow

- ❑ Computes how pixels (or features) move from frame to frame
- ❑ Dense optical flow computes a movement vector for each pixel
- ❑ Sparse optical flow computes a movement vector only for a subset of pixels (e.g., the pixels that are interest points)
- ❑ Optical flow can be used to infer movement of objects as well as movement of the camera

# Project Breakout

- How does computer vision relate to your project?
- Will your robot need to process visual data – if so, what are some of the tasks and functions you will need?

THE END

