


# CPSC 425: Computer Vision



## Lecture 18: Scale Invariant Features (SIFT)

# Menu for Today (October 19, 2018)

## Topics:

- Scale Invariant Feature Transform (SIFT)
- SIFT Detector
- SIFT Descriptor
- Analysis of stability

## Readings:


- **Today's** Lecture: Forsyth & Ponce (2nd ed.) 5.4  
“Distinctive Image Features for Scale-Invariant Keypoints
- **Next** Lecture: Forsyth & Ponce (2nd ed.) 10.4.2, 10.1, 10.2

## Reminders:

- **Assignment 3:** Texture Synthesis is **out**, due on **October 29th**

# Today's “**fun**” Example: Dazzle Camouflage

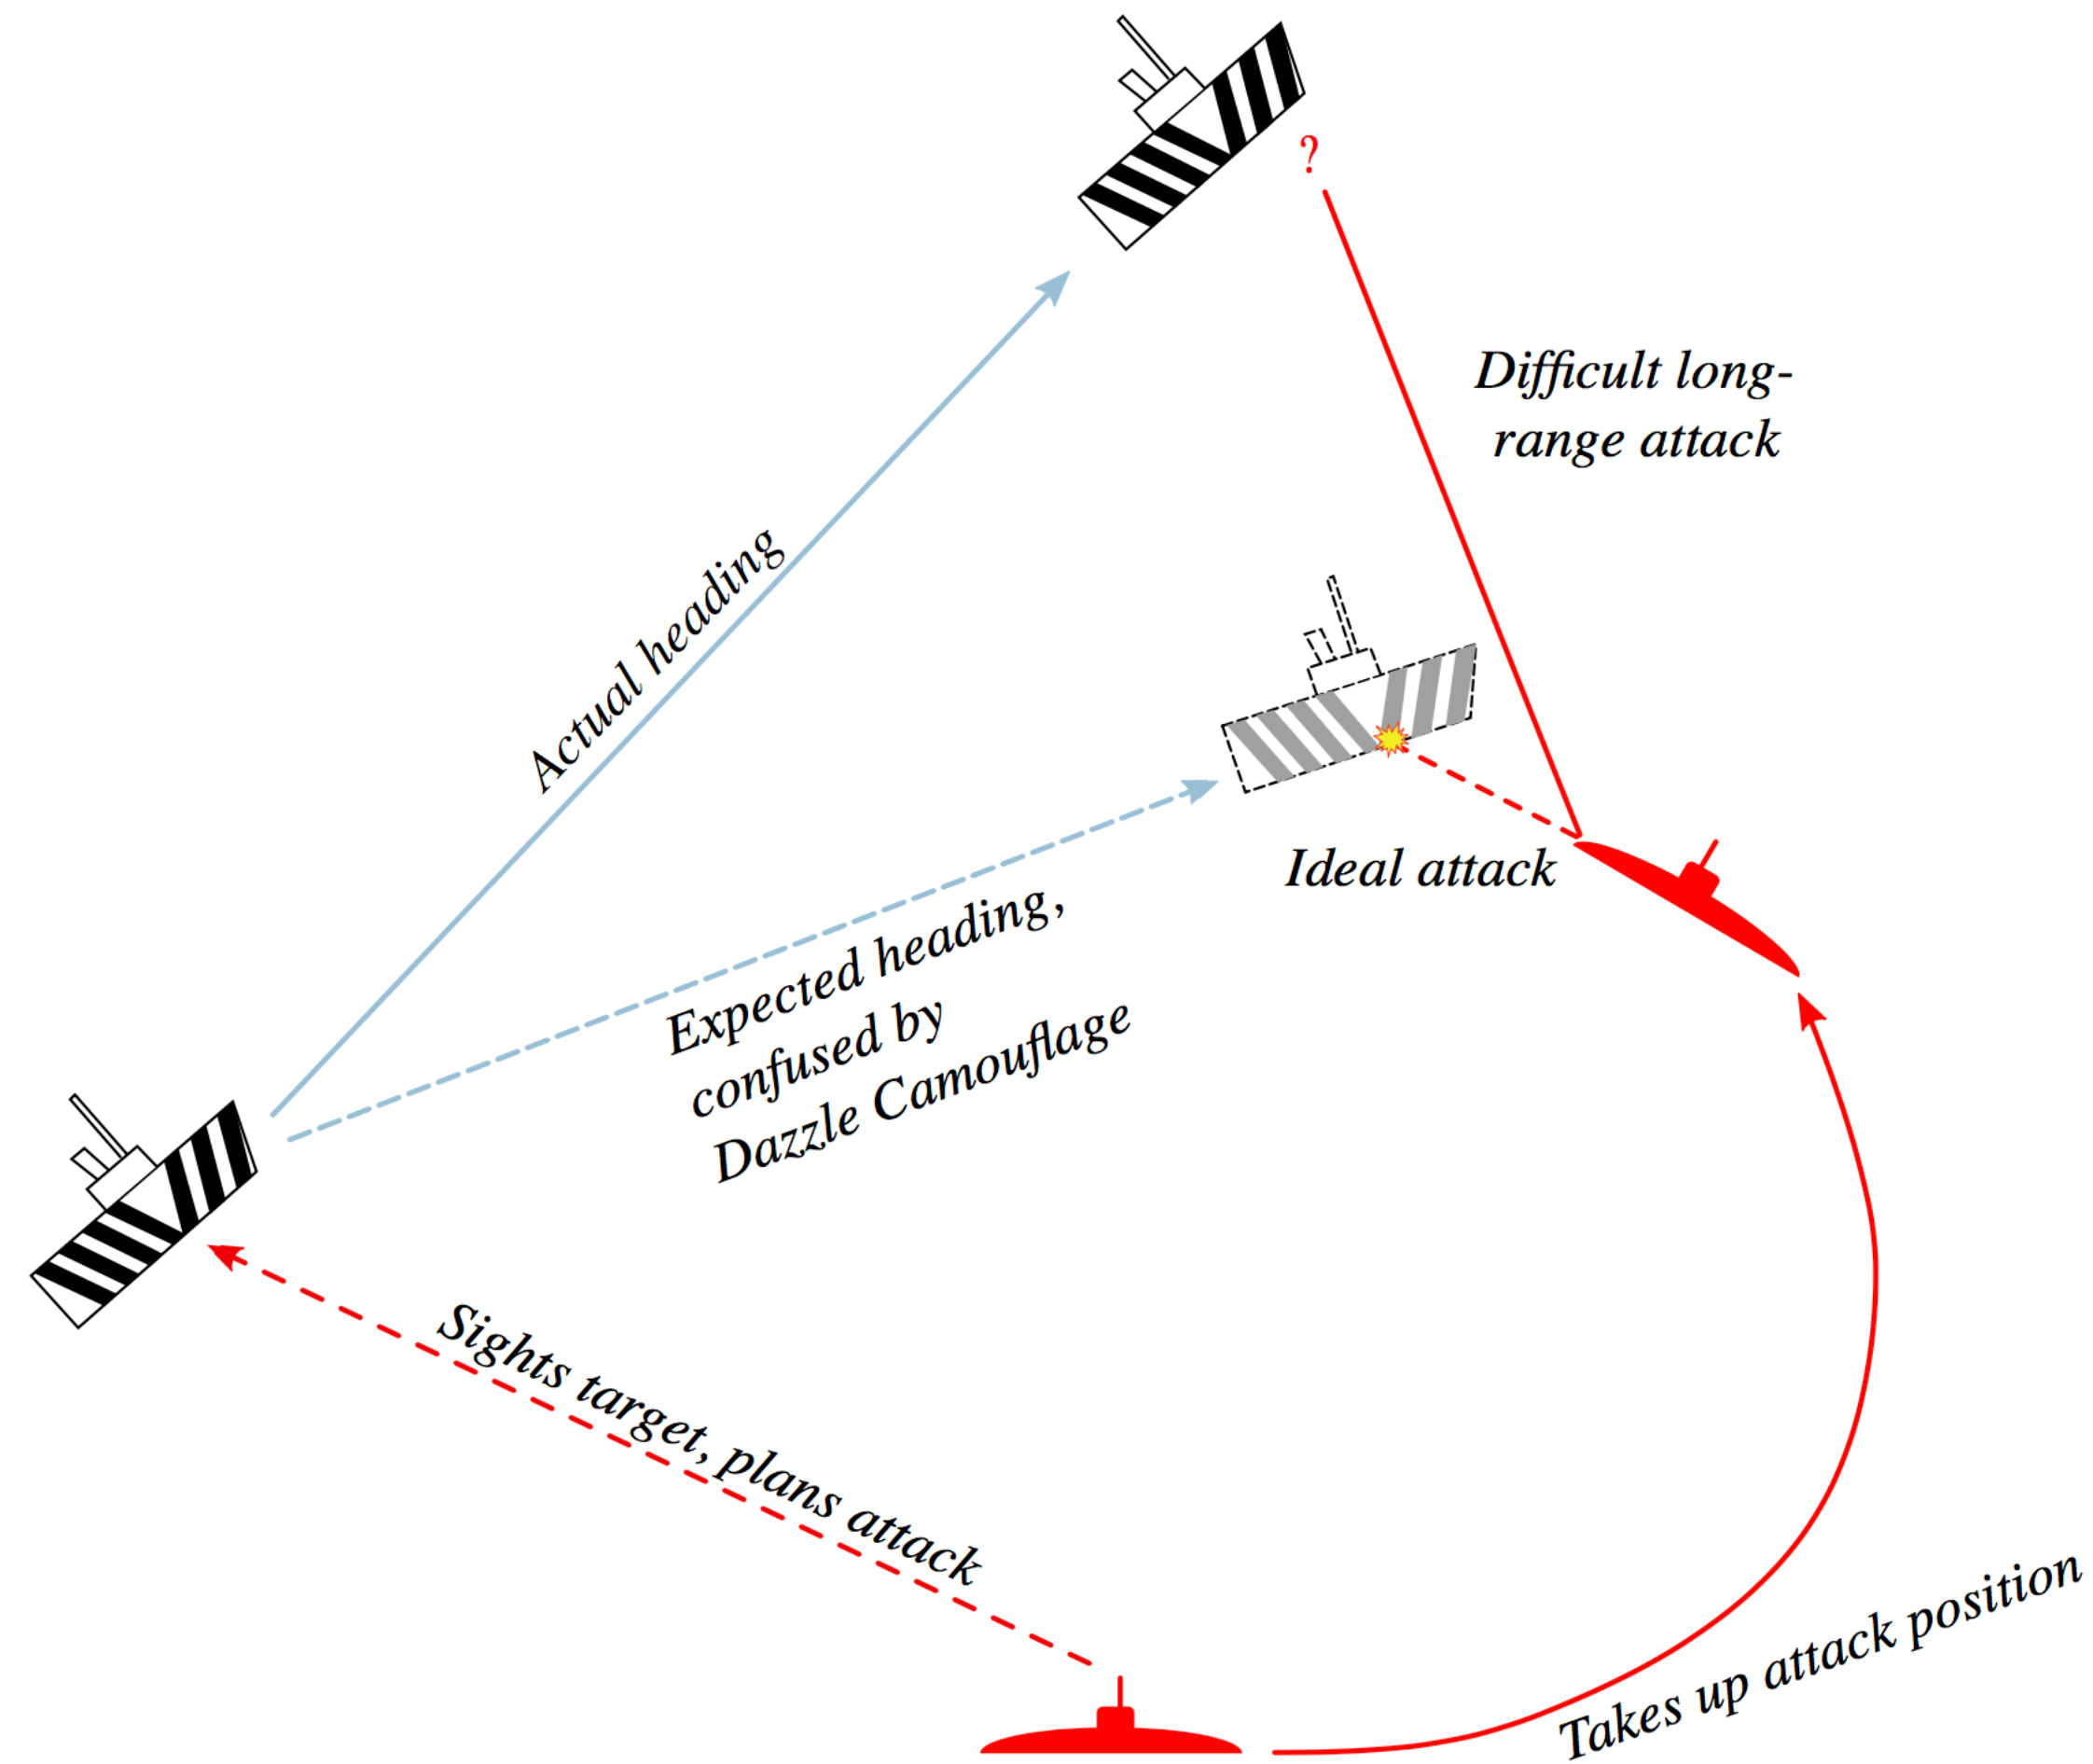
A type of ship camouflage that uses strongly contrasted colours and shapes to make it difficult to estimate the ship's speed and heading





# Today's "fun" Example: Dazzle Camouflage

A type of ship camouflage that uses strongly contrasted colours and shapes to make it difficult to estimate the ship's speed and heading







# Lecture 16: Re-cap

- Human **colour perception**
  - colour matching experiments
  - additive and subtractive matching
  - principle of trichromacy
- **RGB** and **CIE XYZ** are linear colour spaces
- **Uniform colour space**: differences in coordinates are a good guide to differences in perceived colour
- **HSV** colour space: more intuitive description of colour for human interpretation
- (Human) **colour constancy**: perception of intrinsic surface colour under different colours of lighting



# Back to **Good Local Features**



Where are the good features, and  
how do we match them?




# Photometric Transformations





# Geometric Transformations




objects will appear at different scales,  
translation and rotation

Lets assume for the moment we can figure out where the good features (patches) are ... how do we **match** them?

# Intensity Image

Just use the pixel values of the patch




Perfectly fine if geometry and appearance is unchanged  
(a.k.a. template matching)

What are the problems?



# Intensity Image

Just use the pixel values of the patch




Perfectly fine if geometry and appearance is unchanged  
(a.k.a. template matching)

What are the problems?

How can you be less sensitive to absolute intensity values?

# Image Gradients / Edges

Use pixel differences




Feature is invariant to absolute intensity values

What are the problems?

# Image Gradients / Edges

Use pixel differences



Feature is invariant to absolute intensity values

What are the problems?

How can you be less sensitive to deformations?



# Where does **SIFT** fit in?

Representation	Result is...	Approach	Technique
intensity	dense (2D)	template matching	(normalized) correlation, SSD
edge	relatively sparse (1D)	derivatives	$\nabla^2 G$ , Canny
“corner”	sparse (0D)	locally distinct features	Harris, <b>SIFT</b>

# Object **Recognition** with Invariant Features

**Task:** Identify objects or scenes and determine their pose and model parameters


## **Applications:**

- Industrial automation and inspection
- Mobile robots, toys, user interfaces
- Location recognition
- Digital camera panoramas
- 3D scene modeling, augmented reality



# David Lowe's Invariant Local Features

Image content is transformed into local feature coordinates that are invariant to translation, rotation, scale, and other imaging parameters



SIFT Features



# Advantages of Invariant Local Features



**Locality:** features are local, so robust to occlusion and clutter (no prior segmentation)

**Distinctiveness:** individual features can be matched to a large database of objects

**Quantity:** many features can be generated for even small objects

**Efficiency:** close to real-time performance

# Scale Invariant Feature Transform (**SIFT**)



SIFT describes both a **detector** and **descriptor**

1. Multi-scale extrema detection


2. Keypoint localization

3. Orientation assignment


4. Keypoint descriptor



# 1. Multi-scale Extrema Detection



# 1. Multi-scale Extrema Detection



Gaussian




Laplacian




# 1. Multi-scale Extrema Detection

Detect maxima and minima of Difference of Gaussian in scale space



# 1. Multi-scale Extrema Detection — Sampling Frequency

More points are found as sampling frequency increases, but accuracy of matching decreases after 3 scales/octave



## 2. Keypoint Localization

- After keypoints are detected, we remove those that have **low contrast** or are **poorly localized** along an edge

## 2. Keypoint Localization

— After keypoints are detected, we remove those that have **low contrast** or are **poorly localized** along an edge

How do we decide whether a keypoint is poorly localized, say along an edge, vs. well-localized?



## 2. Keypoint Localization

— After keypoints are detected, we remove those that have **low contrast** or are **poorly localized** along an edge

How do we decide whether a keypoint is poorly localized, say along an edge, vs. well-localized?

$$C = \begin{bmatrix} \sum_{p \in P} I_x I_x & \sum_{p \in P} I_x I_y \\ \sum_{p \in P} I_y I_x & \sum_{p \in P} I_y I_y \end{bmatrix}$$

## 2. Keypoint Localization

— After keypoints are detected, we remove those that have **low contrast** or are **poorly localized** along an edge

How do we decide whether a keypoint is poorly localized, say along an edge, vs. well-localized?

— Lowe suggests computing the ratio of the eigenvalues of  $\mathbf{C}$  (recall Harris corners) and checking if it is greater than a threshold

— Aside: The ratio can be computed efficiently in fewer than 20 floating point operations, using a trick involving the trace and determinant of  $\mathbf{C}$  - no need to explicitly compute the eigenvalues

## 2. Keypoint Localization

### Example:





- (a)  $233 \times 189$  image
- (b) 832 DOG extrema
- (c) 729 left after peak value threshold
- (d) 536 left after testing ratio of principal curvatures



### 3. Orientation Assignment

- Create **histogram** of local gradient directions computed at selected scale
- Assign **canonical orientation** at peak of smoothed histogram
- Each key specifies stable 2D coordinates ( $x$  ,  $y$  , scale, orientation)



# 4. Keypoint Description


We have seen how to assign a location, scale, and orientation to each key point  
— **keypoint detection**

— The next step is to compute a **keypoint descriptor**: should be robust to local shape distortions, changes in illumination or 3D viewpoint

— Keypoint detection is not the same as keypoint description, e.g. some applications skip keypoint detection and extract SIFT descriptors on a regularly spaced grid

# 4. SIFT Descriptor

- Thresholded image gradients are sampled over  $16 \times 16$  array of locations in scale space (weighted by a Gaussian with sigma half the size of the window)
- Create array of orientation histograms
- 8 orientations  $\times$  4  $\times$  4 histogram array





# 4. SIFT Descriptor

How many dimensions are there in a SIFT descriptor?

(**Hint:** This diagram shows a 2 x 2 histogram array but the actual descriptor uses a 4 x 4 histogram array)

