

University of Illinois at Urbana-Champaign
Dept. of Electrical and Computer Engineering

ECE 101: Computing Technologies and the Internet of Things

Computer Vision

© 2022 Steven S. Lumetta and Romit Roy Choudhury. All rights reserved.

1

Sense, Compute, Actuate, Communicate

Consider monitoring a street intersection.

What is being sensed?

“Visual” information, which could be in various frequencies, using both passive and active sensors.



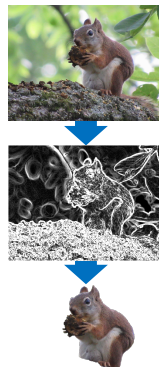
2

Sense, Compute, Actuate, Communicate

What about computation?

It might require **multiple steps**:

1. **combine sensor data** over time
2. **identify and classify** types of objects/entities
3. **understand the physics** of the things identified



3

Computation: Systems Need to Learn and Predict Behavior

Think about young animals learning to hunt—they need to study the motion of their prey.

If we want systems **to act “intelligently,”**
◦ the systems **need to learn the differences**
◦ between humans, dogs, squirrels, cars, trucks, and so forth.



4


Puzzles Play with Our Brain's Visual Processing

Let's have some fun.
Humans like puzzles.

What do you see in the image on the right?

Two faces?
Or a vase?

Most people can see both if they try.




5

Puzzle Missing Clues Found in the Real World

Why can we see both?

The **visual clues** our brains use are **missing or undermined**:

- consistent **coloring** (forced here)
- **depth** perception (no depth in image)
- **shading** (removed here)
- **structure** (none visible)
- **motion** (faces move like faces; vases don't move)



6

A More Realistic Image has Clues

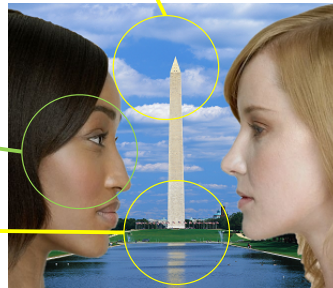
Do you see a vase now?

(You can, but you really have to use your imagination!)

Color, shading, and structure match our idea of a vase.

Shading and structure not right for a vase shape.

Really colorful vase?



7

Tools and Techniques Used by Our Brains

What goes on in our brain to enable us to tell the difference?

Some of everything we talked about:

- **edge detection** by light/color,
- **segmentation** into areas of common color,
- use of depth perception and shading to **estimate 3D structure**, and
- **motion tracking** to estimate body dynamics.

8

Can We Be Fooled by Statues?

Body dynamics? Yes...

A realistically colored and shaped

- **dog statue can trick us**
- until we decide that
- it should have moved but did not.

Look carefully and you may notice

- the “fur” doesn’t look like fur;
- your brain is inferring material properties from lighting.



9

Modern Sensors Make these Tasks Easier

As with speech and language,

- **accurate interpretation** of vision
- **requires lots of information.**

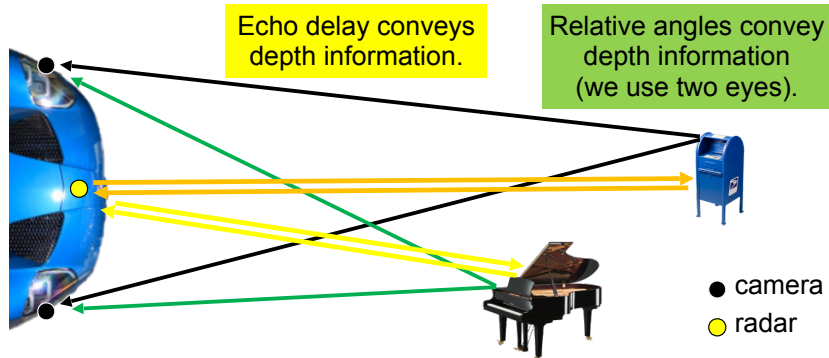
Modern technology does make the task easier.

Rather than depending

- on purely passive sensors (as do humans),
- we **make use of active sensors**, such as lidar and radar.

10

Depth Can be Gauged Passively or Actively



11

Depth Perception Adds to Active Sensor Information

Why use depth perception?

Active sensors rely on specific frequencies,

- which **may not reflect** from all materials,
- **and may require** correct **shape**.

For example,

- some radar (NOT mmWave) reflects only from metals;
- stealth aircraft use surfaces angled to avoid reflection.

12

Feature Registration Enables Fusion of Sensor Data

Modern systems use both

- **active** sensors generate “point clouds” (with depth labels), **and**
- **passive sensors** to generate images.

Specific points on objects

- must be **identified and equated**
- (called **feature registration**)
- in order to combine the sensor information.

Example: which pixels/points correspond to the left front leg of the mailbox?



13


Feature Identification Enables Recognition of Individuals

Feature extraction / registration

- also **used for recognizing** things
- and **instances of things**.

Example: given a human face,

- extract a set of points:
- nostrils, eyebrows, lips, and so forth.
- Relative positions of features
- allow matching of the individual face.



14

Must Understand Motion to Integrate Data Over Time

All **sensors** may **suffer from occlusion**:

- an object is partially or completely
- **behind something else**.

One “**instant**” of data

- allows us to **build**
- a **partial 3D map** of the surroundings.

To combine information over time, we **must**

- identify objects and **associate** them
- **with types of movement** (mailbox: doesn't move!).



15

Edge Detection: Find Boundaries Between Segments


Let's look more closely at one of the first steps: **edge detection**.

In humans,

- the first level of neurons in the visual cortex
- performs similar processing,
- **finding the boundaries between different parts of an image**.

For example, in the image to the right,

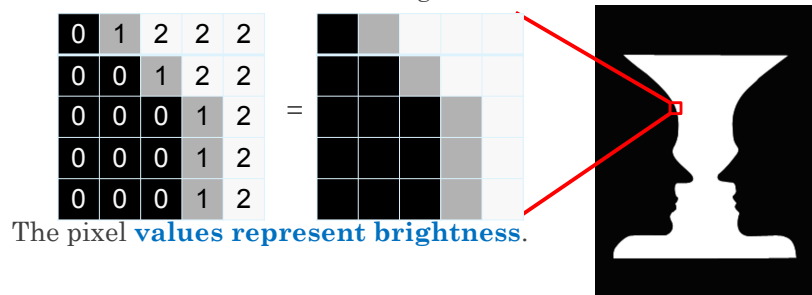
- the yellow region is the “edge”
- between the vase and the rest of the image.



16

Monochrome Pixel Values Represent Image Brightness

Let's first zoom in to a small image section.



17

Filter Captures Notion of a Vertical Edge

Let's use this **3x3 filter** of values.*

- For each pixel,
- we center the 3x3 filter
 - over the pixel,
 - multiply each overlapping pixel
 - with a filter value, and
 - sum all of the results.

-1	0	1
-2	0	2
-1	0	1

vertical edge
filter

The sum indicates whether a vertical line appears at that pixel.

*The filter was designed by I. Sobel and G. Feldman in 1968.

18

Example of Computing One Pixel's Vertical Edge Value

Choose a pixel, align, ...



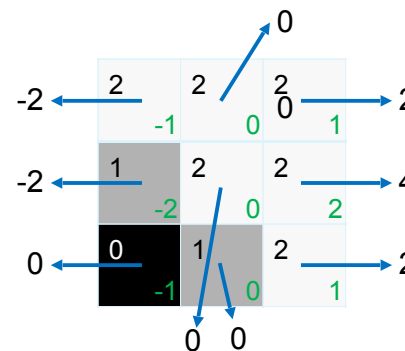
19

Continuing the Computation of One Pixel's Value

... multiply ...

... and sum ...

Answer is 4!



20

The Computed Values Highlight Vertical Lines

Do the same for every pixel (called **convolution** with the filter).



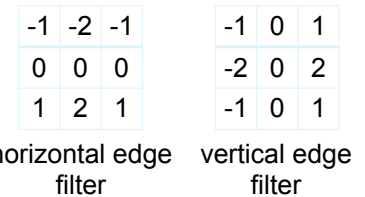
(I duplicated the boundary cells for computation.)

With Measure for Horizontal Lines, Find Line Direction

A second 3x3 filter detects horizontal lines.

By combining the two results,

- we **obtain line strength**
- and line **direction** at each pixel.



horizontal edge filter vertical edge filter

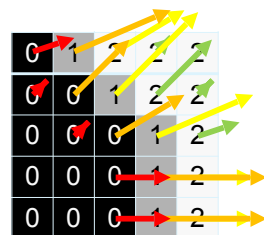
Line Directions (Perpendicular) Identify Segment Bounds

The vectors drawn to the right illustrate the result of the filters

- (colors are just for clarity).

- By **finding continuous sequences**
- of adjacent vectors
 - with similar direction**,
 - we find curves that split the image.

The **curves bound segments** of the image, **which correspond to objects**.



Example: Magnitude of Edge Vectors of ECEB



Example: Magnitude of Edge Vectors of CSL



25

Boundaries Define Segments Corresponding to Objects

Using the boundaries,

- we **split** an **image into segments**.
- In our simple image, there are only two.

Segments

- can be matched against object models
- to find appropriate labels.

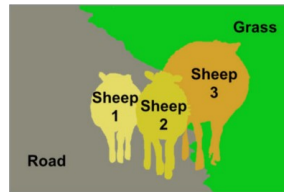



26

ML Filters (Convolutions!) Can Identify Regions

Or we can use ML to train convolutional filters

- that **mark regions according to object type**,
- then look for the regions within the segmented image.



27

Fewer Object Types Means Easier Identification

As with speech,

- a **smaller “vocabulary” of objects**
- **makes identification easier**.

For example, one might omit

- models of lions and elephants
- from many systems.

And almost no systems need to be able

- to spot a rampaging Tyrannosaurus Rex,
- despite the fact that most humans could
- (or so they think).




28

Some Objects Take Time to Identify Correctly

“At first it seemed a little speck,
And then it seemed a mist;
It moved and moved, and took at last,
A certain shape, I wist.”

—Samuel Taylor Coleridge,
The Rime of the Ancient Mariner

Sometimes, it **takes time and observations**
◦ of the motion of an object
◦ in order **to identify** it.



29

ML and Traditional Models May Mix

Towards this end,
◦ when correct identification matters,
◦ the ML-based part of the system
◦ is likely to be backed
◦ by a more traditional set of object models.

As with our brains,
◦ these **models** can be **used in both directions**:
◦ feed **shapes and movement in** to the model
◦ **to classify an object, or,**
◦ **given** an object's **classification, predict**
the kind of **motions** that it can make.

30

“Visual” Data Also Fused with Other Types

Finally, “**visual**” data are likely to be
integrated with other types of information.

Let's say that my car is following
another vehicle, for example.

**How do we know when the
vehicle in front slows down?**

1. **Brake lights**
2. **WiFi notification** from vehicle (in future)
3. “Visual” **speed/depth change** detection



31

Most Sensors Must be Treated as Unreliable

1. Brake lights
2. WiFi notification from vehicle (in future)
3. “Visual” speed/depth change detection

The **first two may fail**:

- slowing may occur without braking, and
- brake lights can go out or be obscured.
- WiFi notifications may not arrive, or
may not be sent from older cars.



32

Trust What You See?

1. Brake lights
2. Notification from vehicle (in future)
3. “Visual” speed/depth change detection

Only change detection is reliable

- (and even there, we need to have multiple sensors—ours can also fail),
- but it is **also delayed**—takes time to notice.



33

Safety-Critical Systems Require Safety Envelopes

In practice, systems require a “**safety envelope**”:

- make sure that the car
- has **enough time to react safely**
- **to changes** in the environment.



34


Safe Decision Making Relies on Layers of Models and Data

Information on slowdown

- must be fed into **models for motion planning**,
- which in turn rely on accurate **models of vehicle dynamics**,
- which rely on sensing the **vehicle's load**.

Average car weight is 4,000 pounds,
with a cargo load of about 1,000 pounds.

That's a 25% change in mass,
assuming humans obey rules.



35

Vision is a Powerful Tool

Where will modern computer vision take us?

To understand where we might go,
let's look to the past...

36

Powerful Vision has Appealed for Millenia

Many human mythologies

- include beings associated with
- **enormous powers of vision,**

Sometimes through use of many eyes, as with Argos.

Others through extremely powerful sight, as with Heimdall.



37

What if You Could See Everything, All the Time?

What would you do

- if you had **eyes everywhere,**
- **all the time?**

Obviously, that's too much information to process oneself.

But now imagine that it is filtered, so you **see only what is interesting.**

What will you do?




38

Parenting? Let us Lend an Eye!

PlayCam's kid-tracking cameras

- **ensure** that your **child** is **safe** from harm and
- **notifies you** instantly **of** any **unusual activity** or behavior!

You can finally relax!



39

Birdwatching? Do it at Home!

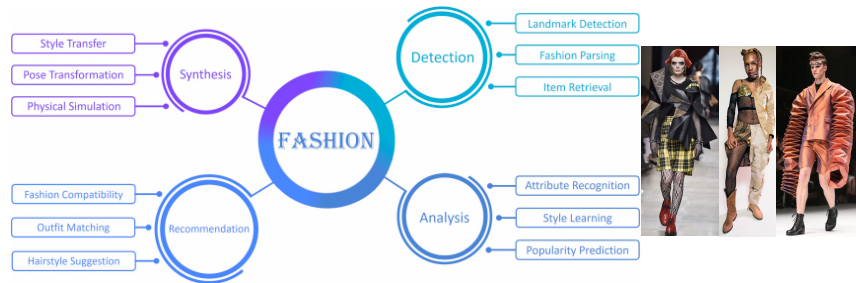
Bird Watch is a virtual reality birdwatching simulator which places participants inside a virtual forest populated by ten common North American birds.

Some predictions about birdwatching!



40

Style? Fashion? We've Got You Covered!

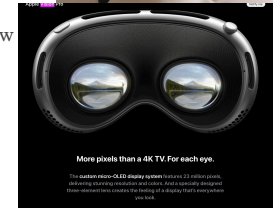


41

Virtual Vision: This is Your Life!

The LiDAR Scanner and TrueDepth camera work together to create a fused 3D map of your surroundings, enabling Apple Vision Pro to render digital content accurately in your space.

EyeSight reveals your eyes and lets those nearby know when you're using apps or fully immersed in an experience. When someone approaches, Apple Vision Pro simultaneously lets you see the person and reveals your eyes to them.



42

The Last Phases: Actuation and Communication

Let's close the loop by returning to the cycle: sense, compute, actuate, communicate.

Once a vision-capable system has interpreted what it sees, **it can respond**.

Each system will be different:

- **drive** the vehicle **safely**,
- **authenticate humans** and **report unusual activity** or
- **deliver messages** when human returns home.



43

Robots Also Need Vision!

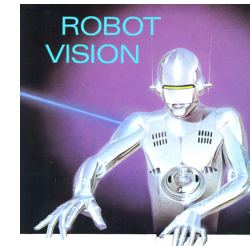
Of course, robots will also need vision.

Robots that look and act like humans are not many years away.

Processing visual signals

- will be an important part in enabling such robots
- to act and communicate with humans.

Whether they look like or unlike us seems to be a matter of choice.



Berthold Klaus Paul Horn

44

Humanoid Robots: State of the Art

American Society of Mechanical Engineers' 2020 lineup of humanoid robots and virtual people (below)...



45

Terminology You Should Know from These Slides

- edge detection
- segmentation
- segments (of an image)
- depth perception
- motion tracking
- point cloud vs. image
- feature registration
- occlusion
- filter
- convolution
- safety envelope

46

Concepts You Should Know from These Slides

- human use of visual clues to identify objects
- why depth perception is useful together with active sensors
- use of feature extraction to recognize individuals
- why understanding movement is useful to “intelligence”
- integration of information over time and across sensors
- role of safety envelope in ensuring safety

47