

Computer Vision Lesson Plan

Overview—*Computer Vision*

Summary

Computers today are being used to accomplish tasks that require using one or more of the five senses. Vision - seeing objects and identifying them - is something that humans can do relatively easily but computers cannot. One way we help teach computers to “see” is to identify where edges are in a picture. From there, the computer can find shapes and compare against a known database of shapes that a human has entered.

Timeline

What	Time	Where
Introduction	10	Computer Vision Slides.pdf
Star Edge Detection	10	Star Worksheet
Cat Edge Detection	10	Cat Worksheet Edge Detection – Cat Slides
Image Recognition	10	Kangaroo Worksheet
What’s It All About	10	

Materials

- ✓ Computer Vision PowerPoint slides.
- ✓ Each child will need pencil or pen and the following worksheets:
 - Worksheet 1: Edge Detection (Star)
 - Worksheet 2: Edge Detection (Cat)
 - Worksheet 3: Kangaroo worksheet

Introduction- Whole Class

Computer Vision is an exciting field of computer science that is still in its infancy. The act of seeing objects, and identifying objects, is something humans can do very well and without much thought. The same task, however, is a nontrivial problem for computers.

Have you ever wondered how your cell phone camera can detect faces in a picture automatically, and focus on them? Or how Google's self-driving car "sees" the road? Computers don't have eyes, but with clever thinking, humans have taught computers to use camera images to sense their surroundings.

One way computers see what is around them is using edge detection to identify shapes and then objects. This is what will be explored in today's activity.

Discussion

Begin the class by getting students to think about what it means to see objects. How do you identify a spoon from a fork? When you see a person with a tennis racket, do you think they are hitting a lemon or a tennis ball? Why?

Humans are very good at identifying objects through sight. We use context – pieces of the surrounding environment – to fully understand and accurately describe what we are looking at.

Once students reach a consensus that seeing is harder than it seems, talk about what problems require vision to solve. Some examples include:

- Recognizing faces in a photo
- Recognizing license plates on a toll road
- Counting objects

What are some activities humans do that require sight?

- Driving – we have to look for other cars, people, and bicycles
- Sports – hand/eye coordination
- Walking – we can see rocks, inclines, and stairs
- Reading books (the paper variety, not on your Kindle/Nook)

After discussing what you can do *with* eyesight, ask students how hard it would be to do the same activities if they were blind.

Use the accompanying *Computer Vision* PowerPoint slides to show the class real world examples of computer vision in action. Notes are attached to each slide explaining what the computer would see, as well as additional discussion questions (to see the notes in pdf, press View->Tools->Comment). For convenience, notes are repeated below.

- Slide 2: What is Computer Vision?
 - The boxes in the image correspond to objects that have been identified in the image.
- Slide 3: Scene Matching
 - This is an example of reconstructing an image using computer vision.
 - The building is cropped out of the original image and a program takes a large collection of images, picks out the ones that are likely from the same location, and constructs an image of what the scene might look like behind the building.
- Slide 4: Google Self Driving Car
 - This is an image of how Google's self-driving car sees the world.
 - A model of the world is created using a laser mounted on top of the car. This information is combined with stored inch-precision maps of the area the car is operating in.
 - The boxes represent objects that the car has identified, such as other cars and pedestrians.
 - Limitations: currently cannot differentiate similar objects such as a rock and crumpled up paper.
 - Discussion point: could this be safer than human driven cars? So far, it has not caused a single accident in over 700,000 miles driven.
- Slide 5: OCR Recognition
 - Discussion point: which questions are easier for a human or computer to answer?
 - What is the product of two 10-digit numbers? (computer)
 - Is this painting beautiful? (human)
 - What is the topic of this article? (human)
 - How many words are in this article? (computer)
- Slide 6: Edge Detection
 - This is an easy task for a computer.
 - Compare neighboring pixels in image. If the difference is big enough, there is probably an edge there.
 - Other checks can also be applied. For example, enforce a minimum length of an edge. A single bright pixel surrounded by dark ones shouldn't count as an edge.
- Slide 7: Segmentation
 - Find regions of the image mostly surrounded by edges
 - Each region likely corresponds to an object in the image
- Slide 8: Matching
 - Each image with a known object is generally pre-processed to extract basic information about the object:
 - Shape of the object
 - Color of the object
 - Discussion: Is it always enough to just compare shapes and color?

Worksheets #1 – Edge Detection

If we are trying to program a computer to identify objects in an image, it can be helpful to find the edges in an image. From these edges, we can pick out larger features such as shapes.

Computers represent images as grids of pixels. Ask the students if they know/remember what a pixel (picture element) is. Each pixel has a single color, and we can use numbers to represent the color of a pixel. One of the simplest representations of an image is grayscale, where each pixel can only be a shade of gray. A typical range of numbers is 0 to 255, where 0 is black, 255 is white, and any number between is a shade of gray.

Given a grayscale image, one simple way to find edges is to look at two neighboring pixels and take the difference between their values. If it's big, this means the colors are very different, so it's an edge.

Distribute the star edge detection worksheets to the students. Give them about 5-10 minutes to complete the activity. **DO NOT TELL THEM WHAT SHAPE IS IN THE IMAGE.** That is part of the fun. Let them try to identify what shapes they can see after they detect the edges.

This technique is similar to what computers do. Before you detect the edges, it is very difficult to see the image. But after, you have a clear idea of the dominant shape in the image.

Next distribute the cat edge detection worksheet. Tell the students this one is a bit harder, and they may need to guess what the image is. Give the students 5-10 minutes to complete this activity.

Wrap-up discussion.

How many students figure out what the image is? The power point slides show:

- the original cat picture
- the original picture with the color values overlaid
- the color values without the original picture (corresponds to the worksheet)

Additional Resource

The CS Field Guide has a cool link for edge detection. The Field Guide can be found here: <http://csfieldguide.org.nz/ComputerVision.html>, and the link to the edge detection application can be found here: http://inspirit.github.io/jsfeat/sample_canny_edge.html (you will need a webcam).

Worksheets #2 – Shape Comparison

Kangaroo Worksheet

Another technique computers use to identify objects is to compare them to known shapes.

Distribute the kangaroo worksheet to the students. Give them about 5-10 minutes to complete the activity. Discuss what shapes they used to recognize the kangaroo. Did they use a lot or few shapes?

What are some reasons why a computer might want to be detailed? (*More accurately identify objects on the first try*)

What are some reasons why a computer might want to be general? (*Faster to calculate – answer quick*)

What's It All About Discussion – Whole Class

You can regroup as a class if you have time and discuss links to other fields of computer science:

- Vision is something that is hard for computers. Other hard concepts include the other senses (sounds – talking/listening, for example). If a computer can see and hear well, does it make the computer intelligent? (discussed further in the Artificial Intelligence activity)
- Computers don't "see" colors and shapes. They see numbers. Do you know what numbers the computer reads? (Binary, relates to Binary Numbers and Image Representation activities)

Computer Vision is used in applications that we can see all around our daily lives. Have you ever driven on a toll road without stopping at a toll booth? Toll roads often use high-speed cameras that take a picture of your license plate, "read" the numbers and letters, and pull the address the license plate is registered with in order to process an invoice for the driver.

Social networks like Facebook use computer vision to automatically recognize your face and the faces of your friends in photos uploaded to their website. Tagging pictures is boring, so if a computer can do it for you, you have more time to do other things.

It's important to remember that humans are good at doing hard things, and computers are good at doing easy things... but they do them very quickly. Seeing objects is something that requires a great deal of contextual information that computers have trouble doing. Computer scientists are working every day to come up with better algorithms for recognizing objects in pictures, so that computers can "see" better!

Extension:

Show this interesting video on computer vision: <https://www.youtube.com/watch?v=eQLcDmfmgB0>