Type of Approaches

Different approaches tackle detection differently. They can roughly be categorized into three main types:

- Find **interest points**, followed by Hough voting
- **Sliding windows**: "slide" a box around image and classify each image crop inside a box (contains object or not?)
- Generate **region (object) proposals**, and classify each region ← **We have looked at R-CNN a little bit. Today we'll focus on how to group pixels in super pixels, and regions. Once you have a region, you just attack it with a Neural Network.**

Segmentation

How Can We Get Semantics Quickly?

- Each image has about 1 to 4 **million pixels**.
- That's a lot.

How Can We Get Semantics Quickly?

- Each image has about 1 to 4 **million pixels**.
- That's a lot.
- A real robotics system needs to recognize things really fast if it wants to react to the world around it in real-time

How Can We Get Semantics Quickly?

- Each image has about 1 to 4 **million pixels**.
- That's a lot.
- A real robotics system needs to recognize things really fast if it wants to react to the world around it in real-time
- So the main question really is...

How Can We Get Semantics Quickly?

- Each image has about 1 to 4 **million pixels**.
- That's a lot.
- A real robotics system needs to recognize things really fast if it wants to react to the world around it in real-time
- So the main question really is...

How can we find this cookie before it's gone?

Remember HOG?

- HOG divides an image into cells. There is 64-times less cells than pixels. Two orders of magnitude less. Yet still, there is a lot of them.

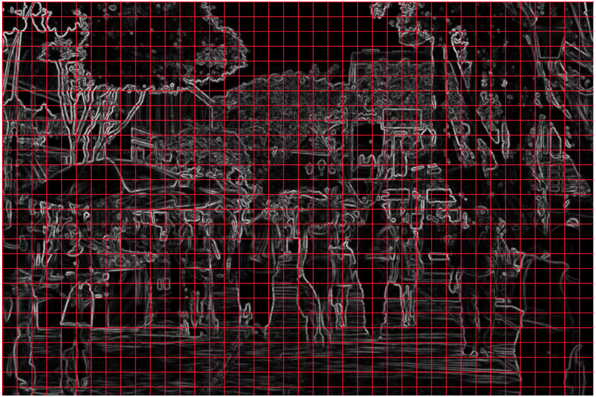

Remember HOG?

- HOG divides an image into cells. There is 64-times less cells than pixels. Two orders of magnitude less. Yet still, there is a lot of them.
- And let's not forget we still need to run the detector, not only tones of locations, but also tones of scales.

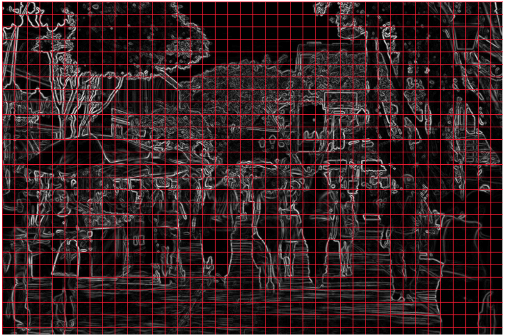

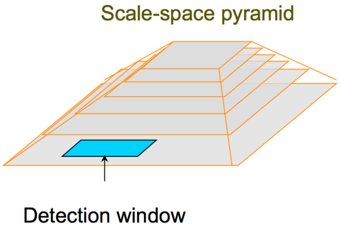

Superpixels

- One of the ideas (which I like quite a lot) is to merge similar pixels into (way less) "superpixels".
- **First question** to ask: How does that help us?
- **Second question**: What properties should superpixels (or any regions) satisfy in order to be useful?

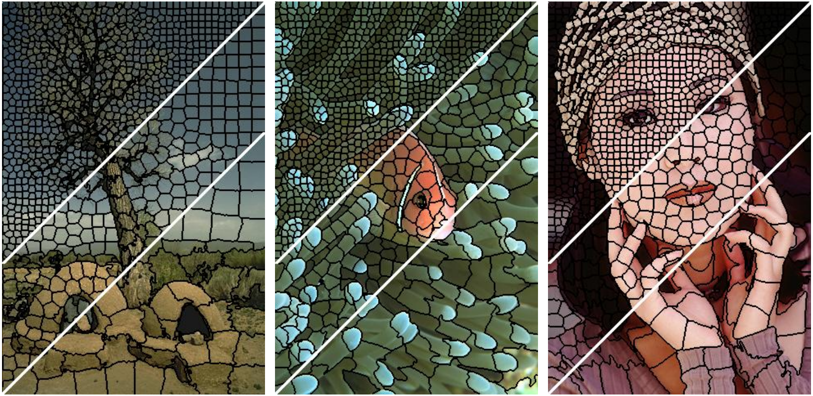

Why Are Superpixels Useful?

- Example 1: How can we find all *road* pixels in this image?

Why Are Superpixels Useful?

- Compute superpixels. A 4million pixel image is converted into only 500 superpixels. And now?

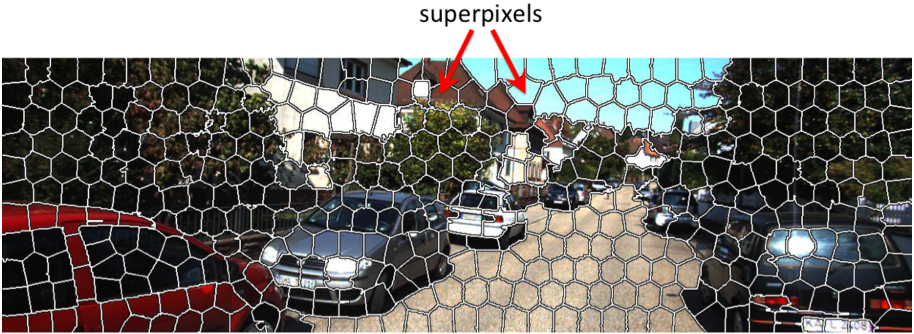

[Superpixels computed by Jian Yao, PhD student at UofT. Thanks Jian!]

Why Are Superpixels Useful?

- Possible idea: Compute features on each superpixel and train a classifier for road/non-road. Use this classifier at test time

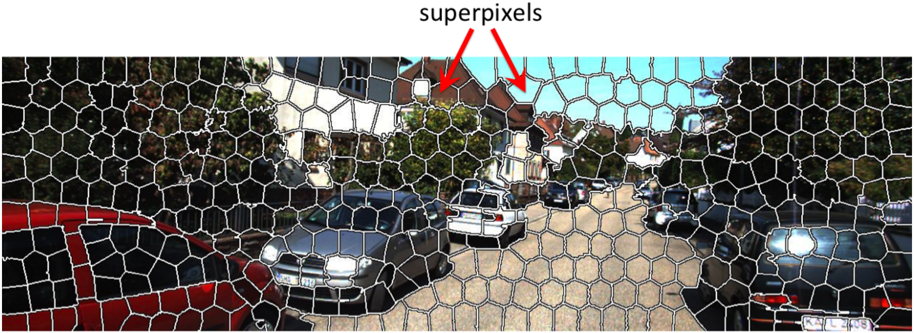

[Superpixels computed by Jian Yao, PhD student at UofT. Thanks Jian!]

Why Are Superpixels Useful?

- Compute features for each superpixel. (make sure to normalize them, e.g, to norm or max value 1; classifiers will work better)
- Different superpixels have different number of pixels. How can I compute my feature vector (has to have the same dimension for each superpixel)?

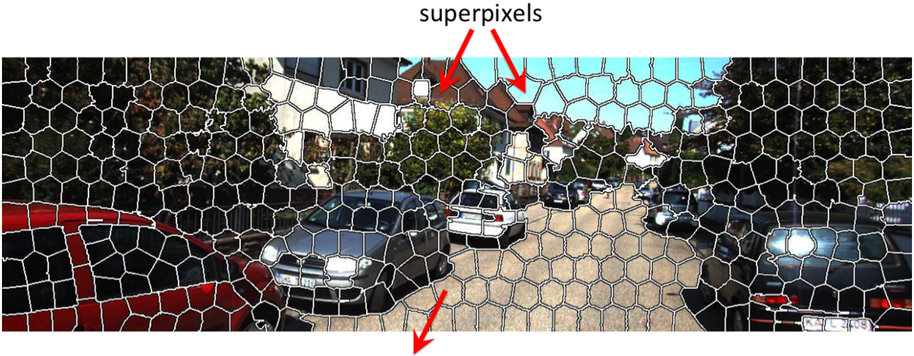

$$\mathbf{f} = [\text{color}, \text{gradients}, x, y, \ldots]^T$$

How?

Why Are Superpixels Useful?

- Average or better: **histogram** (typically the more dimensions that the feature vector has, the better the classifiers work). Histograms allow you to inflate a feature to multiple dimensions. And now?

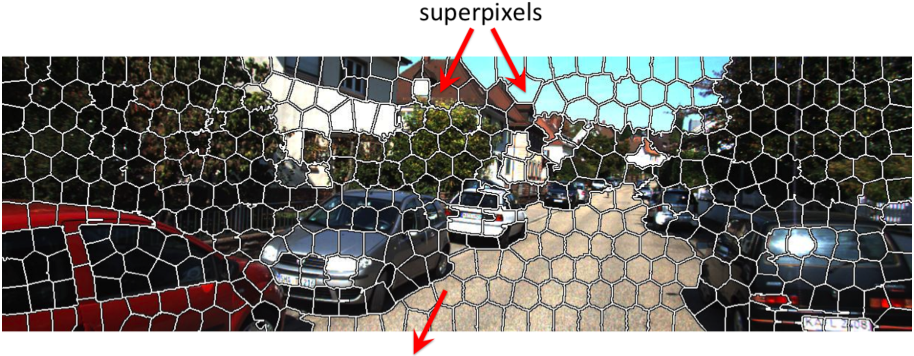

$$\mathbf{f} = [\text{color}, \text{gradients}, x, y, \ldots]^T$$

Can be: 1) average across pixels in superpixel, 2) a histogram, 3) a histogram of visual words

Why Are Superpixels Useful?

- Collect (randomly sample) a bunch of positive features (from regions annotated as *road*) and negative features (from regions of non-road).

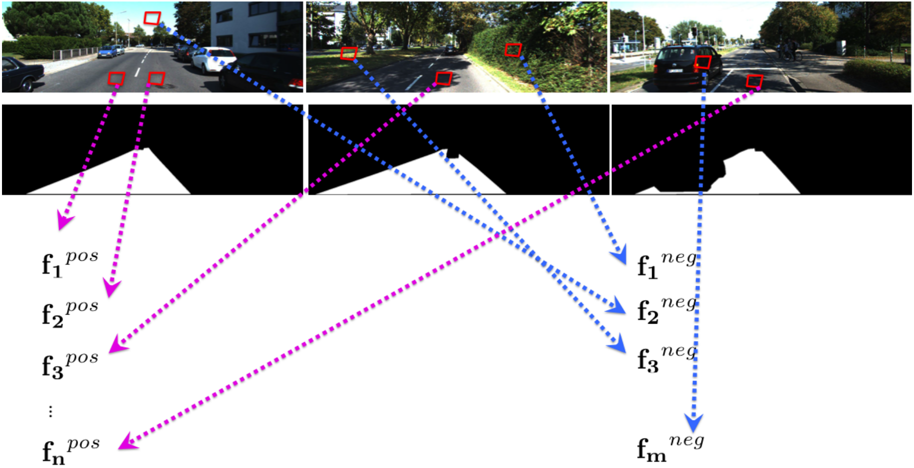

Why Are Superpixels Useful?

- Train your favorite classifier. And at test time?

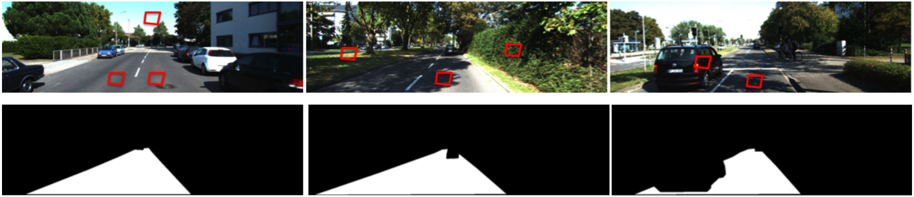

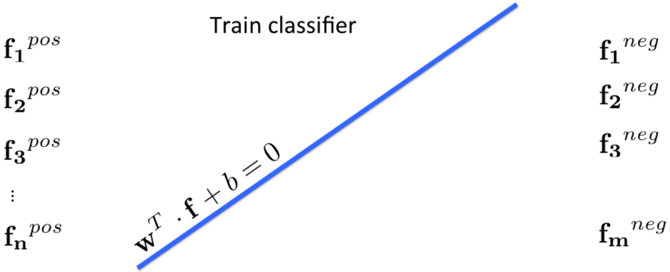

Why Are Superpixels Useful?

- Some problems can emerge, e.g., shadows, or grayish things on buildings that can confuse the classifier

Shadows look like anything else (car, dark tree, etc). Prediction will be bad.
Can we do something about it?

- We can make use of **depth** (e.g, road is typically not 1m above our eye level). Which brings us to:
- Example 2: We can use superpixels also for de-noising stereo results

(a) left image (b) right image

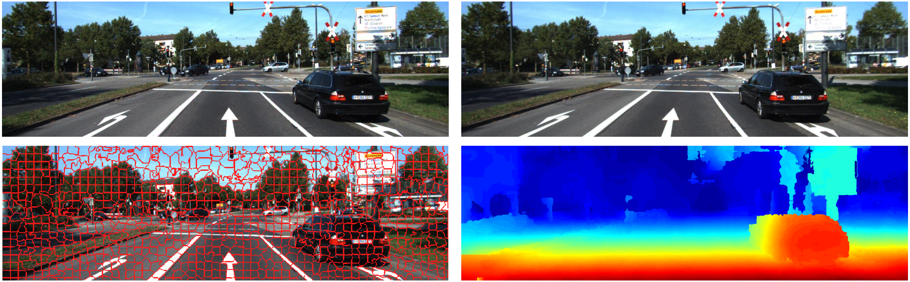

(c) superpixels (d) estimated disparity

Figure: `http://ttic.uchicago.edu/~dmcallester/SPS/index.html`

Why Are Superpixels Useful?

- (back to) Example 1: Now just simply augment superpixel features with 3D features.
- And hope for the best.

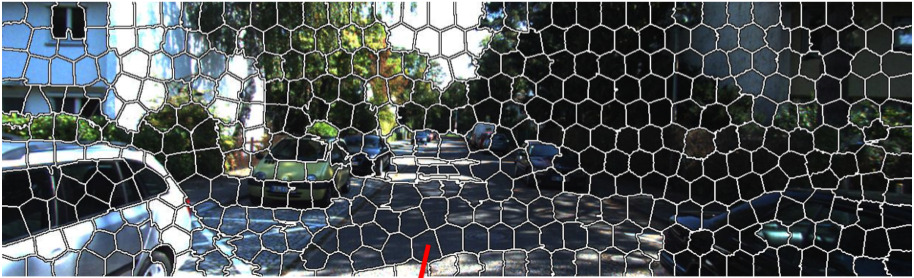

$$\mathbf{f} = [\text{color}, \text{gradients}, x, y, Y, \text{3d features (gradients in 3D?)}, \ldots]$$

Add depth informed features

Why Are Superpixels Useful?

- (back to) Example 1: Now just simply augment superpixel features with 3D features.
- And hope for the best.

Road Estimation Evaluation

UM_ROAD

	Method	Setting	Code	MaxF	AP	PRE	REC	FPR	FNR	Runtime	Environment	Compare
1	DDN			**93.65 %**	88.55 %	**94.28 %**	93.03 %	**2.57 %**	6.97 %	2 s	GPU @ 2.5 Ghz (Python + C/C++)	
R. Mohan: Deep Deconvolutional Networks for Scene Parsing. 2014.												
2	CNN1			91.73 %	**92.08 %**	91.10 %	92.36 %	4.11 %	7.64 %	2 s	1 core @ 2.5 Ghz (C/C++)	
Anonymous submission												
3	CNN		code	91.22 %	91.35 %	91.22 %	91.23 %	4.00 %	8.77 %	2 s	1 core @ 2.5 Ghz (C/C++)	
4	HybridCRF			90.99 %	85.26 %	90.65 %	91.33 %	4.29 %	8.67 %	2 s	1 core @ 2.5 Ghz (C/C++)	
Anonymous submission												
5	NNP			90.50 %	87.95 %	91.43 %	89.59 %	3.83 %	10.41 %	5 s	4 cores @ 2.5 Ghz (Matlab)	
Anonymous submission												
6	HIM			90.07 %	79.98 %	90.79 %	89.35 %	4.13 %	10.65 %	7 s	>8 cores @ 2.5 Ghz (Python + C/C++)	
D. Munoz, J. Bagnell and M. Hebert: Stacked Hierarchical Labeling. European Conference on Computer Vision (ECCV) 2010.												
7	ANM			89.76 %	86.50 %	90.59 %	88.94 %	4.21 %	11.06 %	1s	1 core @ 2.5 Gh (Matlab)	
Anonymous submission												
8	FusedCRF			89.55 %	80.00 %	84.87 %	**94.78 %**	7.70 %	**5.22 %**	2 s	1 core @ 2.5 Ghz (C/C++)	
L. Xiao, B. Dai, D. Liu, T. Hu and T. Wu: CRF based Road Detection with Multi-Sensor Fusion. Intelligent Vehicles Symposium (IV) 2015.												
9	FCN_LC			89.36 %	78.80 %	89.35 %	89.37 %	4.85 %	10.63 %	0.03 s	GPU @ 2.5 Ghz (Python + C/C++)	

Andrew Berneshawi's last year's CSC420 project

Why Are Superpixels Useful?

- Example 3: You can also train superpixel classifiers to predict **surface labels** from a **monocular image**

Figure: http://web.engr.illinois.edu/~dhoiem/projects/context/

Why Are Superpixels Useful?

- Example 4: You can also train superpixel classifiers to predict a variety of **semantic classes**

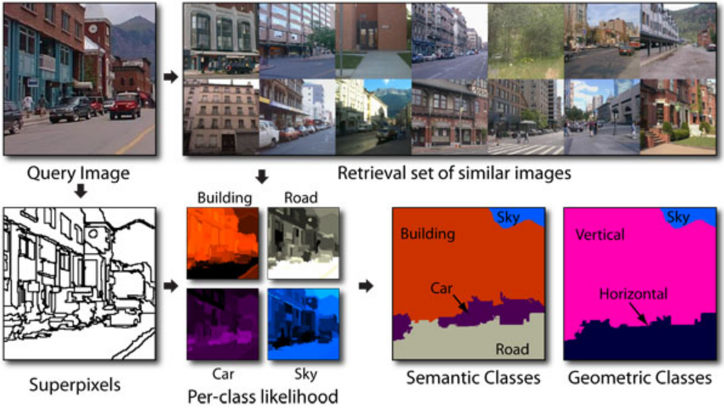

Figure: `https://www.youtube.com/watch?v=UZ_NTYCFlHk`

[Results from: J. Tighe and S. Lazebnik, SuperParsing: Scalable Nonparametric Image Parsing

Why Are Superpixels Useful?

- Remember the RGB-D image in Assignment 3?

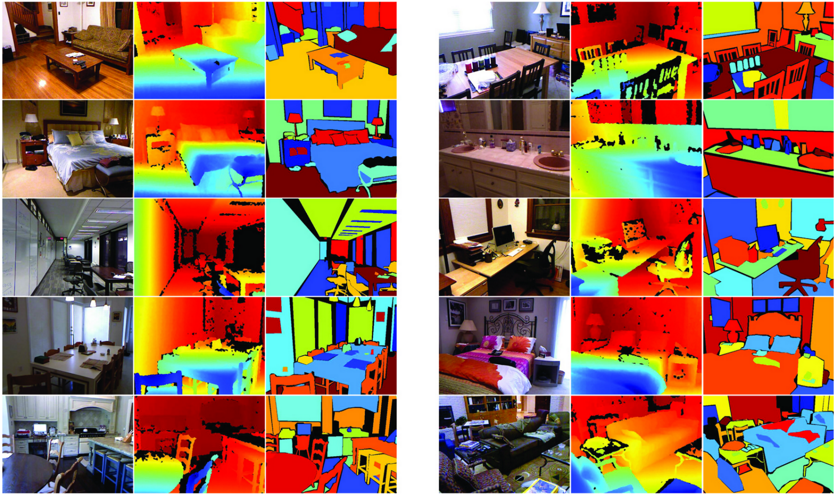

Figure: http://cs.nyu.edu/~silberman/datasets/nyu_depth_v2.html

Why Are Superpixels Useful?

- Example 5: You can also train superpixel classifiers with RGB-D features to predict a variety of **semantic classes**

(a) Color Image (b) Depth (c) Contours (d) Semantic Segmentation

Figure: http://www.cs.berkeley.edu/~sgupta/pdf/GuptaArbelaezMalikCVPR13.pdf

Why Are Superpixels Useful?

- Example 6: And of course also to find our cookie.

[Superpixels computed by Jian Yao, PhD student at UofT. Thanks Jian!]

Why Are Superpixels Useful?

- Example 7: We can use them for tracking. And many more things.

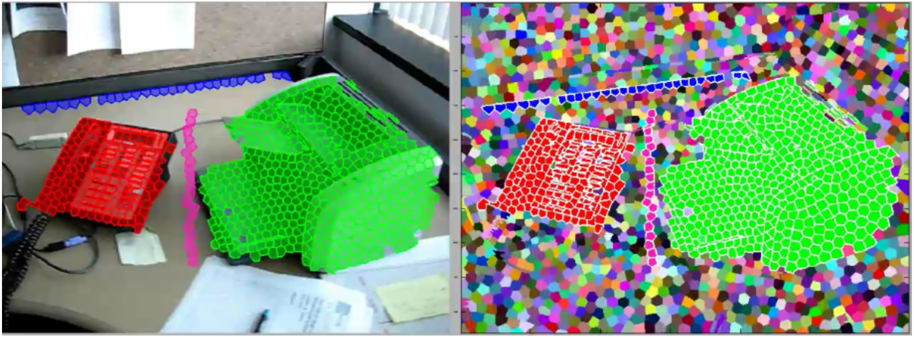

Figure:
`http://groups.csail.mit.edu/vision/sli/projects.php?name=temporal_superpixels`

Segmentation

- Question 2: What properties should our segmentation have?

- Should be fast to compute
- Should not merge different objects (undersegmentation)

Types of segmentations

Oversegmentation

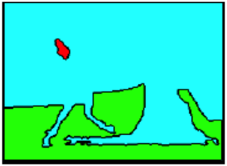

Undersegmentation

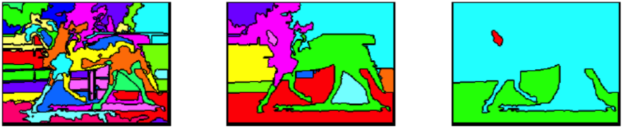

Multiple Segmentations

Slide by D. Hoiem

Segmentation Algorithms

There are a lot of them out there

- Felzenswalb and Huttenlocher's Graph-based Segmentation
- SLIC
- SEEDS
- Shi and Malik's Normalized Cuts
- Malik's group: Probability of boundary (gPb)
- Grundmann et al: Hierarchical Graph-based Video Segmentation
- Chang et al., Temporal Superpixels

For most the code is available!

Segmentation Algorithms

There are a lot of them out there

- Felzenswalb and Huttenlocher's Graph-based Segmentation $\rightarrow$ Let's quickly look at this one
- SLIC
- SEEDS
- Shi and Malik's Normalized Cuts
- Malik's group: Probability of boundary (gPb)
- Grundmann et al: Hierarchical Graph-based Video Segmentation
- Chang et al., Temporal Superpixels

For most the code is available!

Felzenswalb and Huttenlocher's Graph-based Segmentation:
`http://cs.brown.edu/~pff/segment/`

Felzenswalb & Huttenlocher

- Take an image.

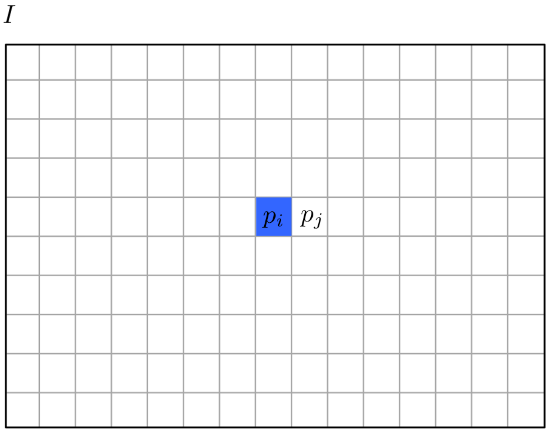

Felzenswalb & Huttenlocher

- Compute a weight between a pair of neighboring pixels. A weight reflects **dissimilarity**.

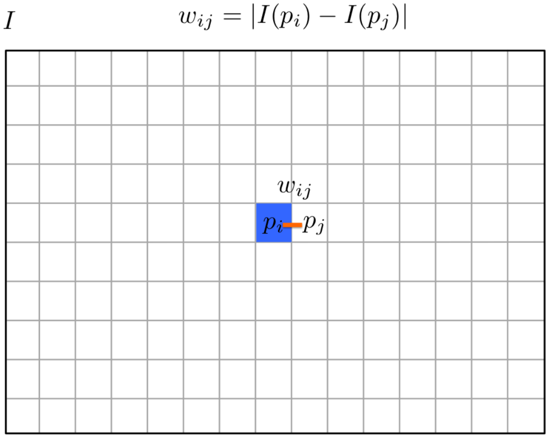

Felzenswalb & Huttenlocher

- Compute weights between a pixel and all 8 of its neighbors.

I

$$w_{ij} = |I(p_i) - I(p_j)|$$

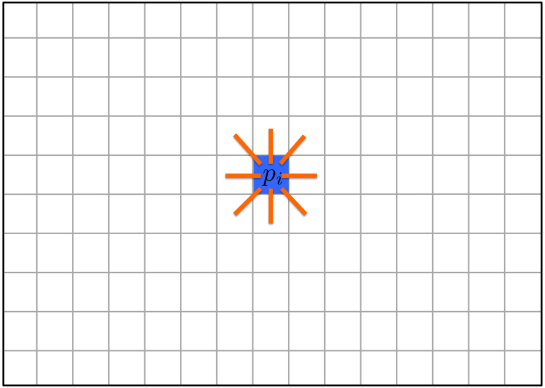

Felzenswalb & Huttenlocher

- And do that for **all pixels** in the image.

I $\qquad w_{ij} = |I(p_i) - I(p_j)|$

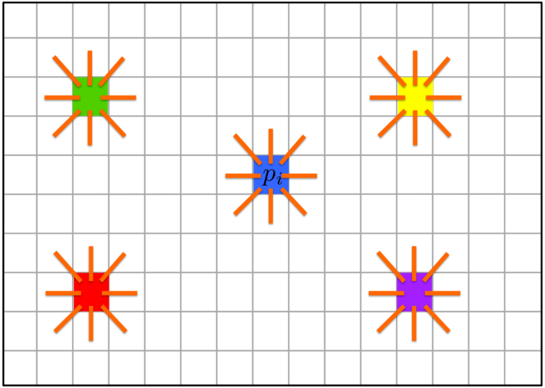

Felzenswalb & Huttenlocher

- Now sort all the weights in the image by non-decreasing values.

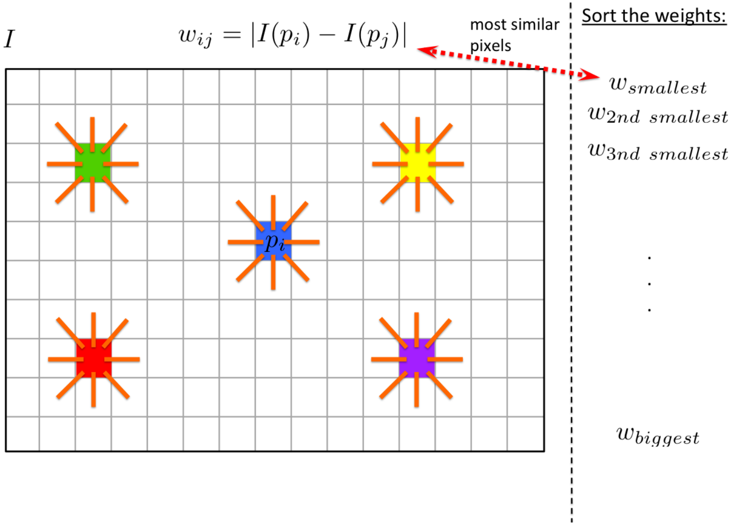

Felzenswalb & Huttenlocher

- Pick the weight from the top of the lest (most similar two neighboring pixels)

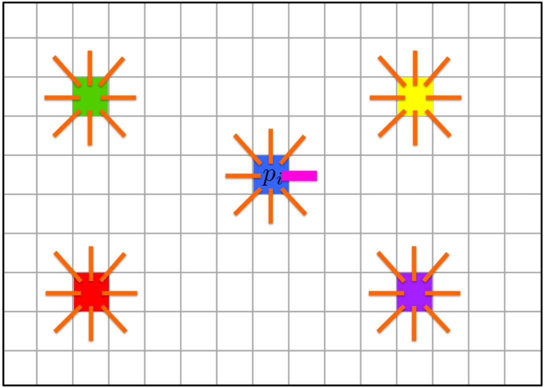

Sort the weights:

$w_{smallest}$

$w_{2nd\ smallest}$

$w_{3nd\ smallest}$

.
.
.

$w_{biggest}$

Felzenswalb & Huttenlocher

- If the weights is smaller than the internal dissimilarities (plus a threshold), **merge** the pixels

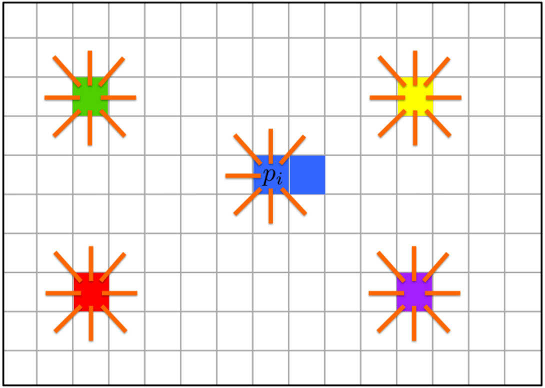

<u>Sort the weights:</u>

$w_{smallest}$

$w_{2nd\ smallest}$

$w_{3nd\ smallest}$

$\vdots$

$w_{biggest}$

Felzenswalb & Huttenlocher

- Take the second weight in the list and do the same.

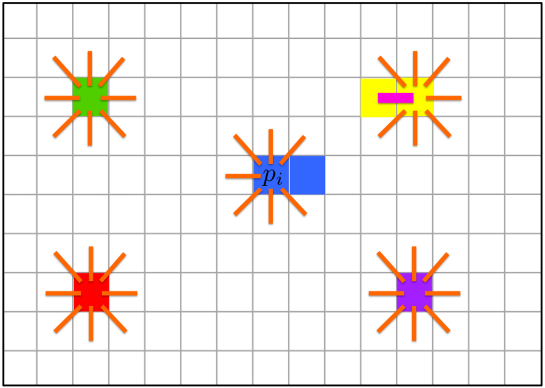

Sort the weights:

$w_{smallest}$

$w_{2nd\ smallest}$

$w_{3nd\ smallest}$

.
.
.

$w_{biggest}$

Felzenswalb & Huttenlocher

- And the third one. Repeat until the end of list, or until the weight is higher than the dissimilarity between the pair of components.

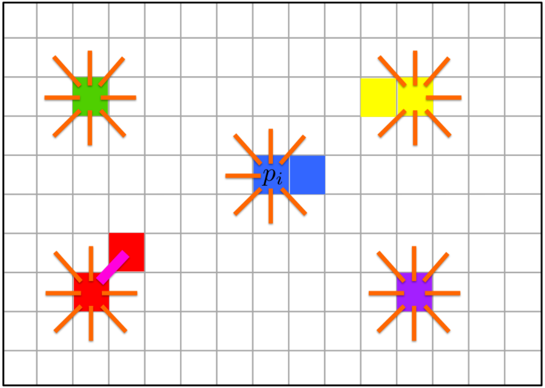

<u>Sort the weights:</u>

$w_{smallest}$

$w_{2nd\ smallest}$

$w_{3nd\ smallest}$

$\vdots$

$w_{biggest}$

Felzenswalb & Huttenlocher

- Result. The algorithm runs real-time.

Segmentation Algorithms

There are a lot of them out there

- Felzenswalb and Huttenlocher's Graph-based Segmentation
- SLIC $\rightarrow$ Let's quickly look at this one too
- SEEDS
- Shi and Malik's Normalized Cuts
- Malik's group: Probability of boundary (gPb)
- Grundmann et al: Hierarchical Graph-based Video Segmentation
- Chang et al., Temporal Superpixels

For most the code is available!

SLIC: `http://ivrg.epfl.ch/research/superpixels`

SLIC Superpixels

- Ask the user for the number of superpixels.
- Initialize the superpixel **centers** in a regular grid.

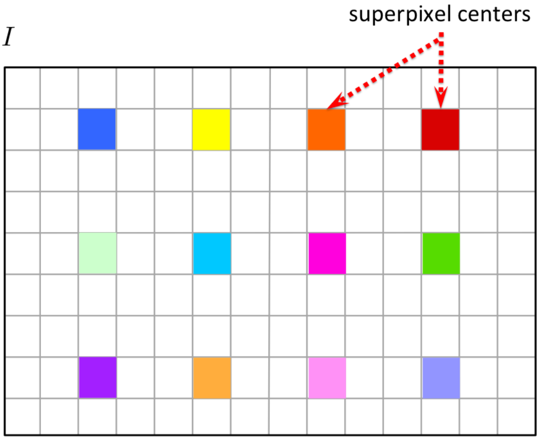

SLIC Superpixels

- Allow each center to shift in a 3×3 window to the location with the smallest gradient.

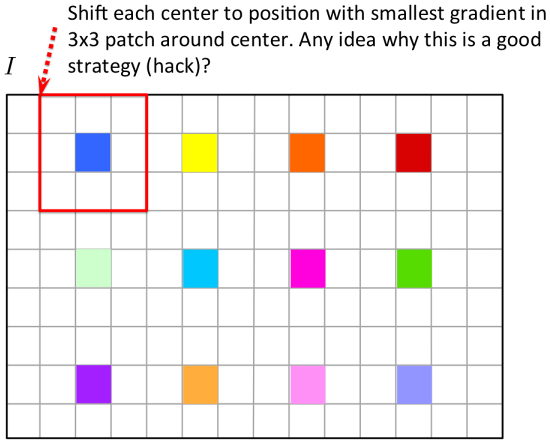

- Compute feature vectors for each pixel.

Compute a feature vector in each pixel (and superpixel center). Feature vector: [color, x, y]

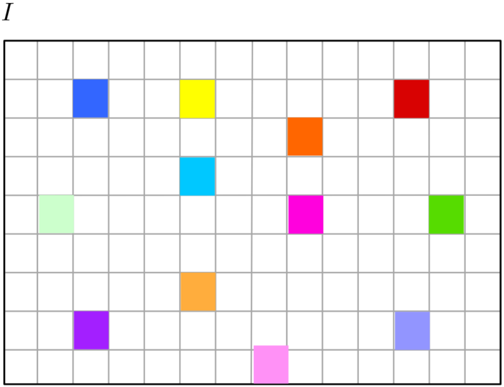

SLIC Superpixels

- Assign each pixel to the closest superpixel center.

Assign all other pixels to the closest center, where ``closest''
is looking at a distance between the feature vectors

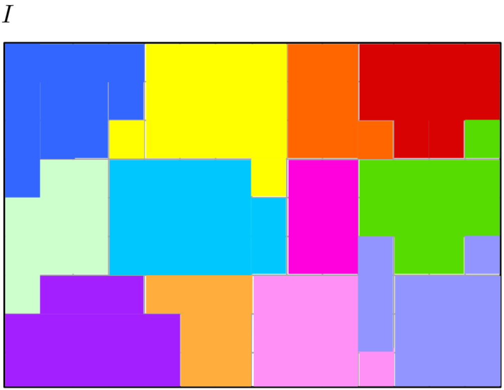

- Re-compute the centers (just average the features across pixels in each superpixel).

Re-compute the superpixel centers (just compute the average feature vector).
Also compute how much each center moved from previous iteration.

I

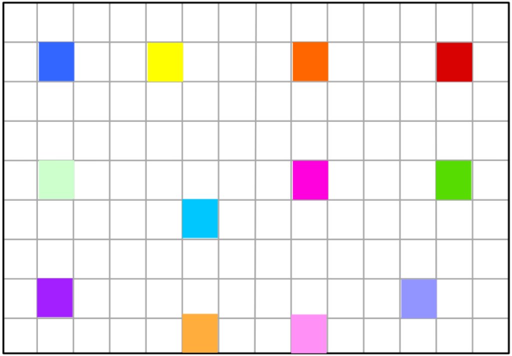

SLIC Superpixels

- Assign again. Iterate until the centers become stable (don't move from the previous iteration).

Assign again, and iterate. Stop when none of the centers moves more than a threshold.

I

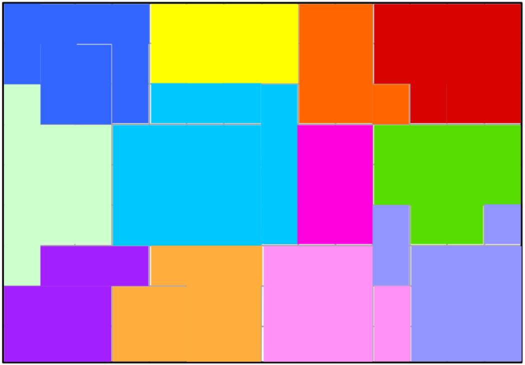

Results

- Results

Results

- Results

Figure: https://www.youtube.com/watch?v=nUnAQUbeymQ

So... Singularity?

- So, you've seen some of the best of computer vision. What do you think, are we approaching singularity?

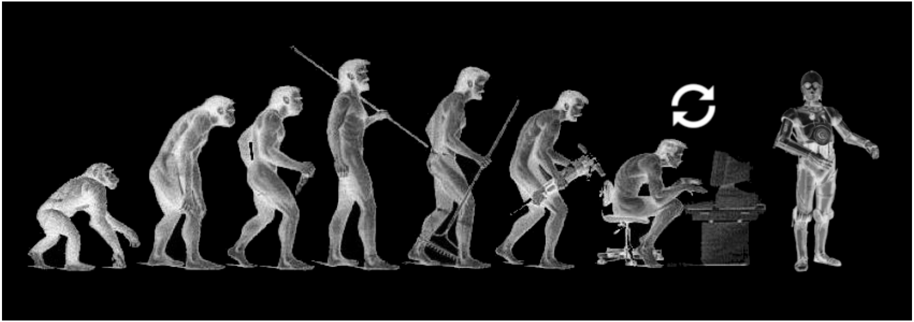

So... Singularity?

- Not just yet...

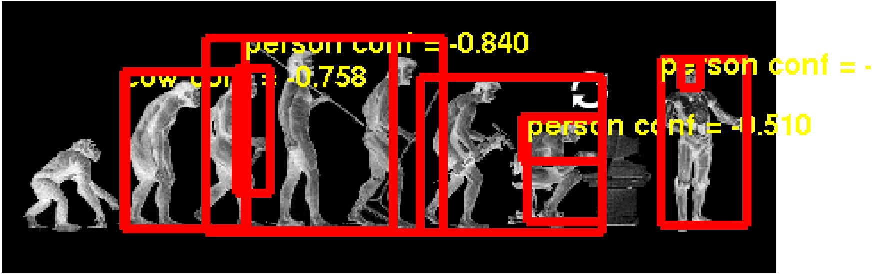

Figure: Results of RCNN

That's It For CSC420... But There Is Much More of Computer Vision For Those Interested!