

Image Segmentation

How Photographers Can See Like Machines

Image Segmentation

To better study and look at images, segment the images (something I learned from machine learning):

How do machines see images?

To start off, I want to state that humans are far superior to machines in almost every way. Let us not get fooled by the 'golden calf' of AI (artificial intelligence):

The purpose is for us to leverage machines to make better humans.

Not,

Trying to use humans to make better machines.

To re-iterate:

Let us leverage machine learning to better improve ourselves.

So I don't think the purpose is for humans to see like machines (remember Arnold S as the Terminator, when he sees the world in black-and-red). What I propose is that we study how machines see images, as a way to guide how we can better see as humans.

Or in other words:

Let us study machine learning, to better understand how we (as humans) think and see!

Simplify



Fig. 6. Fully convolutional networks improve performance on PASCAL. The left column shows the output of our most accurate net, FCN-8s. The second shows the output of the previous best method by Hariharan *et al.* [14]. Notice the fine structures recovered (first row), ability to separate closely interacting objects (second row), and robustness to occluders (third row). The fifth and sixth rows show failure cases: the net sees lifejackets in a boat as people and confuses human hair with a dog.

The interesting thing about how machines see images:

They simplify it.


When I study my own black and white photos, I believe my best photos are my simple photos. Even if you study great artists like Picasso and Matisse, they always strived for elegant images-- simple and elegant. Think of Matisse's cutouts -- he was able to abstract the human body into the simplest shape, form, and color.


Figure-to-ground



Figure (the subject) and ground (background): you want maximal contrast between the two elements.

One thing we can use: "Gaussian Blur" in Photoshop, to abstract the images, to just see the position of our visual elements, and to simplify the contrast:





Does the photo work well as a small thumbnail?



To better understand whether your compositions are strong or not-- look at them as small thumbnails. If your photos work as small thumbnails, it probably means your composition is strong!










Bounding Boxes


There is already artificial intelligence that can automatically put 'bounding boxes' (little red boxes) around subjects. This is how self-driving cars see pedestrians (not to run them over):





Even more insane-- machines can automatically color, segment, and classify different subjects in a scene (real time):





FCIS								
	Mask R-CNN							
		backbone	AP	AP ₅₀	AP ₇₅	AP _S	AP _M	AP _L
		MNC [10]	ResNet-101-C4	24.6	44.3	24.8	4.7	25.9
FCIS [26] +OHEM		ResNet-101-C5-dilated	29.2	49.5	-	7.1	31.3	50.0
FCIS+++ [26] +OHEM	ResNet-101-C5-dilated	33.6	54.5	-	-	-	-	
Mask R-CNN	ResNet-101-C4	33.1	54.9	34.8	12.1	35.6	51.1	
Mask R-CNN	ResNet-101-FPN	35.7	58.0	37.8	15.5	38.1	52.4	
Mask R-CNN	ResNeXt-101-FPN	37.1	60.0	39.4	16.9	39.9	53.5	

Instance segmentation mask AP on COCO test-dev. MNC and FCIS are the winners of the COCO 2015 and 2016 segmentation challenges, respectively. Mask R-CNN outperforms the more complex FCIS+++.




A simple way we (humans) can look at photographs:

Let us put little red boxes around different visual elements in photos.

The Cherry on Top

For example, I like to look at the 'cherry on top' for subjects in the photograph-- small details which I find interesting. For example, in this photograph, this little boy in the background is a detail I love:







But this is the difficult thing-- most people (when they see the picture really small, perhaps on a phone), they won't see the detail of the boy in the background:



This is why I like printing out photos big-- so you can actually see small details!

And to be honest, 99% of your viewers won't notice these 'cherry on top' details in the photo. But it doesn't matter-- as long as *YOU* see the detail and like it-- that is enough!

Color Segmentation



We can also apply gaussian blur to color photos, to better understand the color palette of our photos:



I used the color picker in Photoshop (hotkey "I" for "I[eye] dropper"). And painted these colors:



If I make the background grey, you can see the colors:



So you can see the basic colors: rogue red, pale yellow, and green in the bottom-left.


Keep the colors simple




For color photography, I generally think the simpler, the better.

For example this picture of this woman, the primary color is pink (with Gaussian Blur applied):







The final with the background subtracted:



Which details do you look at in a photo?




If you use the concept of a 'bounding box' -- you can use it as a tool to identify things you find interesting in a scene-- like the woman smoking a cigarette, and two Red Bulls in her left hand.



Feature Map





When you are shooting photos, try to abstract your vision -- simplify it.

For example, this is why I recommend shooting high-contrast black and white mode, because it helps you see better when you're taking photos.

For example on a RICOH GR II, I recommend using the 'high contrast black and white mode' in JPEG, and when you're composing your photos, keep the photos simple.

Silhouette



This is how silhouettes work -- we create a 'feature map' (the gradients between black, grey, and white).

Edges

Also another way to look at photos; try to simplify the edges, to see the textures in the photo:





Silhouette of Faces




Another idea:

As humans, we intuitively know the proportions of a human face, and also the silhouette of faces.

To see better, let us apply maximal contrast to the photo:



Then, a cyan cutout of the woman:





Conclusion: The Future of Photography



I'm currently doing more studying on Augmented Reality (AR), especially the 'AR CORE' platform from Google, as well as their "Tensor Flow" platform. The basic idea I have is this:

One day, while you're taking photos on your phone, you will get compositional grids, and 'bounding boxes' to help assist your compositions.

I already have this functionality on "intelligent Auto" mode on my Lumix G9-- when I am composing a scene, the LCD screen will automatically put a yellow box around my subject in the frame.

Augmented Reality for Photographers

I think the future of photography is this:

Photographers will leverage artificial intelligence (AI) and machine learning to take better photos (improved composition, or awareness of subjects).

The future is already here

For further reading, I recommend:

1. Research "AR CORE" from Google
2. Research 'computational photography'
3. Check out the Google AI Experiments
4. Download some "AR" (augmented reality) apps on your phone or iPad and play around with it
5. Think about how photographers, individuals, and artists can leverage virtual reality, augmented reality, and machine learning to make better art!

ERIC