



Segmentation

Segmentation

- Semantic Segmentation: by class
- Instance Segmentation: by object
- Panoptic Segmentation: mixed
 - Instance segmentation of “things”
 - Semantic segmentation of “stuff”



Datasets

- MS COCO
- Driving Datasets
- Simulated Datasets



Applications

- Driving
- MS Kinect
- Graphics



FCN


- Upsampling
 - Linear
 - Upconvolutions
- Offset predictions
 - NaN testing
- Pre-trained classifiers



Improving FCN: Dilation

- Dilated convolutions
 - Pre-trained models still work
- Context module
 - Dilated layers on top of a complete network
 - Finer detail in segmentation

a		b		c
d		e		f
g		h		i




Improving FCN: Deformable Conv.

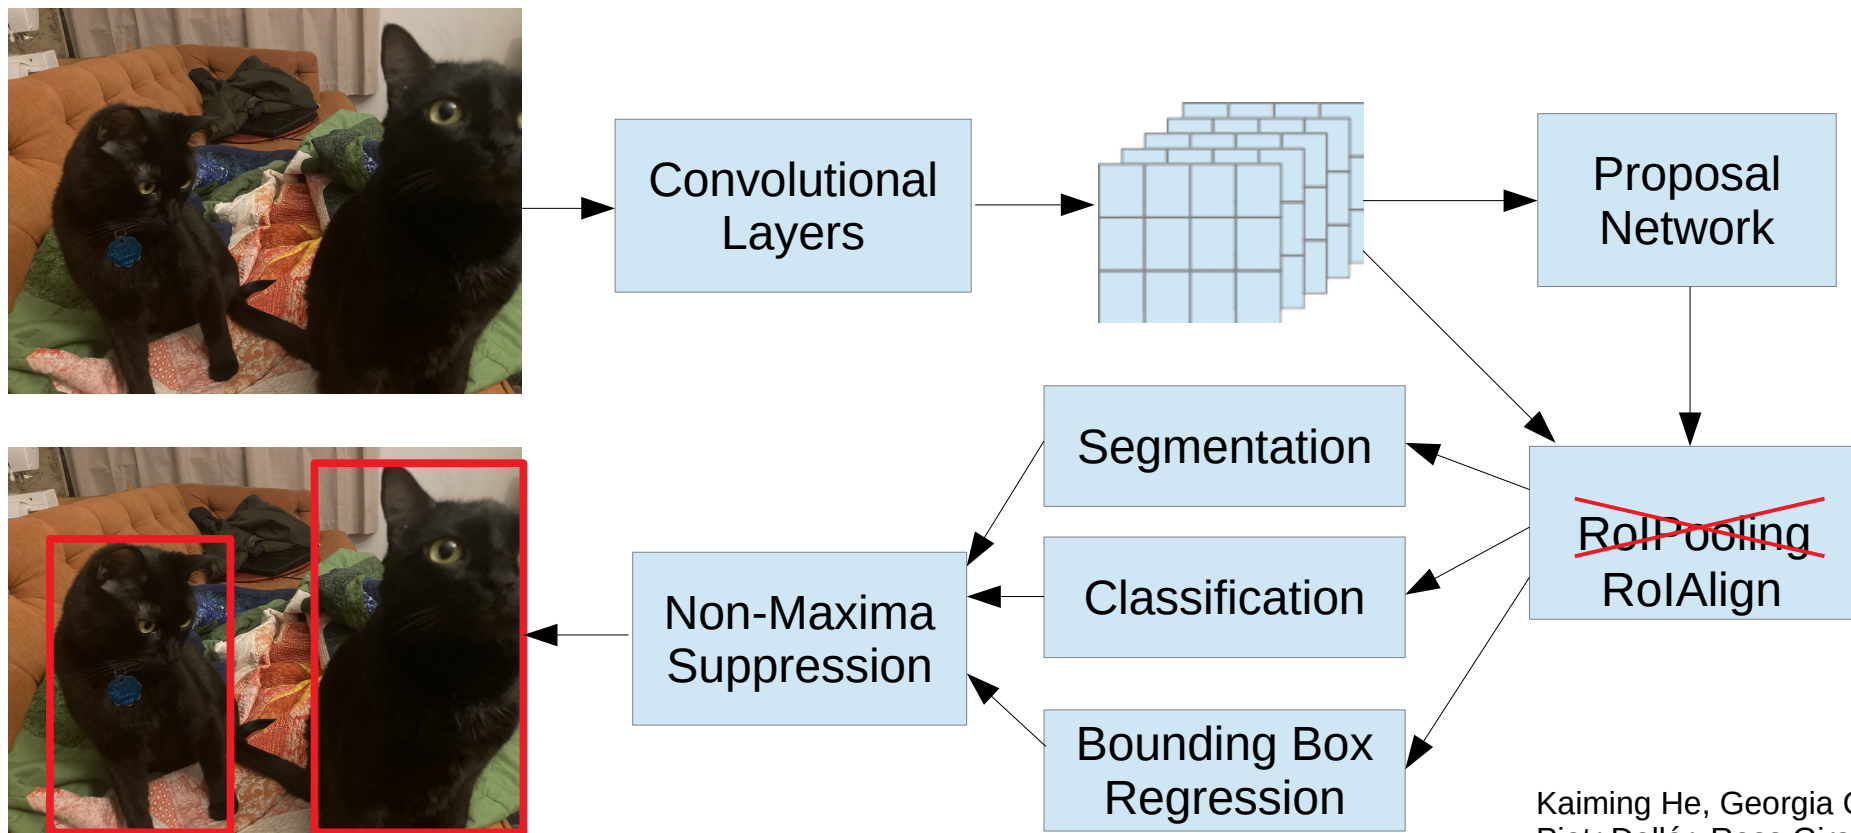
Train network to predict locations of the dilated kernel pixels

a	b	c
d	e	f
g	h	i

	a	b		c
			e	
d			f	
	g			
		h		i



Mask RCNN




Object Representation: Box

Simple

Easy to label

Compact



Occlusion


Too big

Not enough
information

Object Representation: Mask

More detail

Better with occlusion
and partial objects



Hard to label

Still not that much
information

Object Representation: Keypoint

Low dimensional

Can have useful
information



Class specific

Object Representation: 3D model

- Inverse graphics
- Lots of detail
- Too much detail
- Labeling is extremely hard
- (Sometimes) Surprisingly little information for all that detail