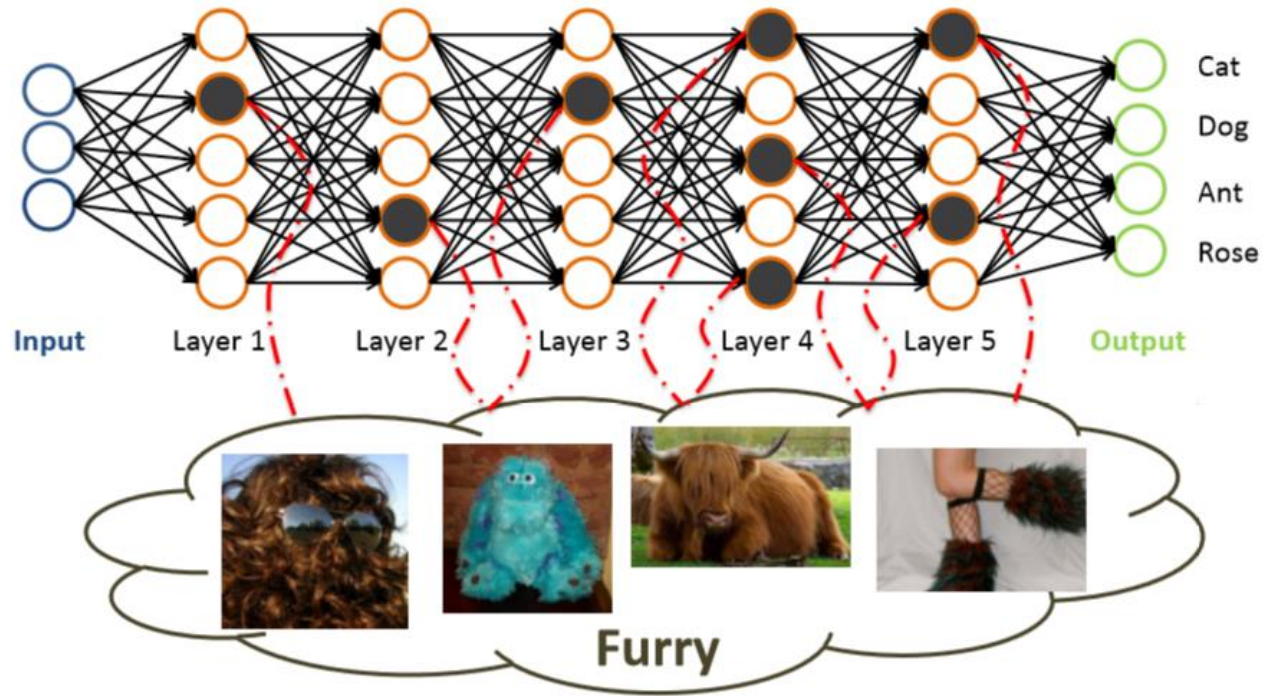



# On the Relationship between Visual Attributes and Convolutional Networks



Victor Escorcia, Juan Carlos Niebles, Bernard Ghanem

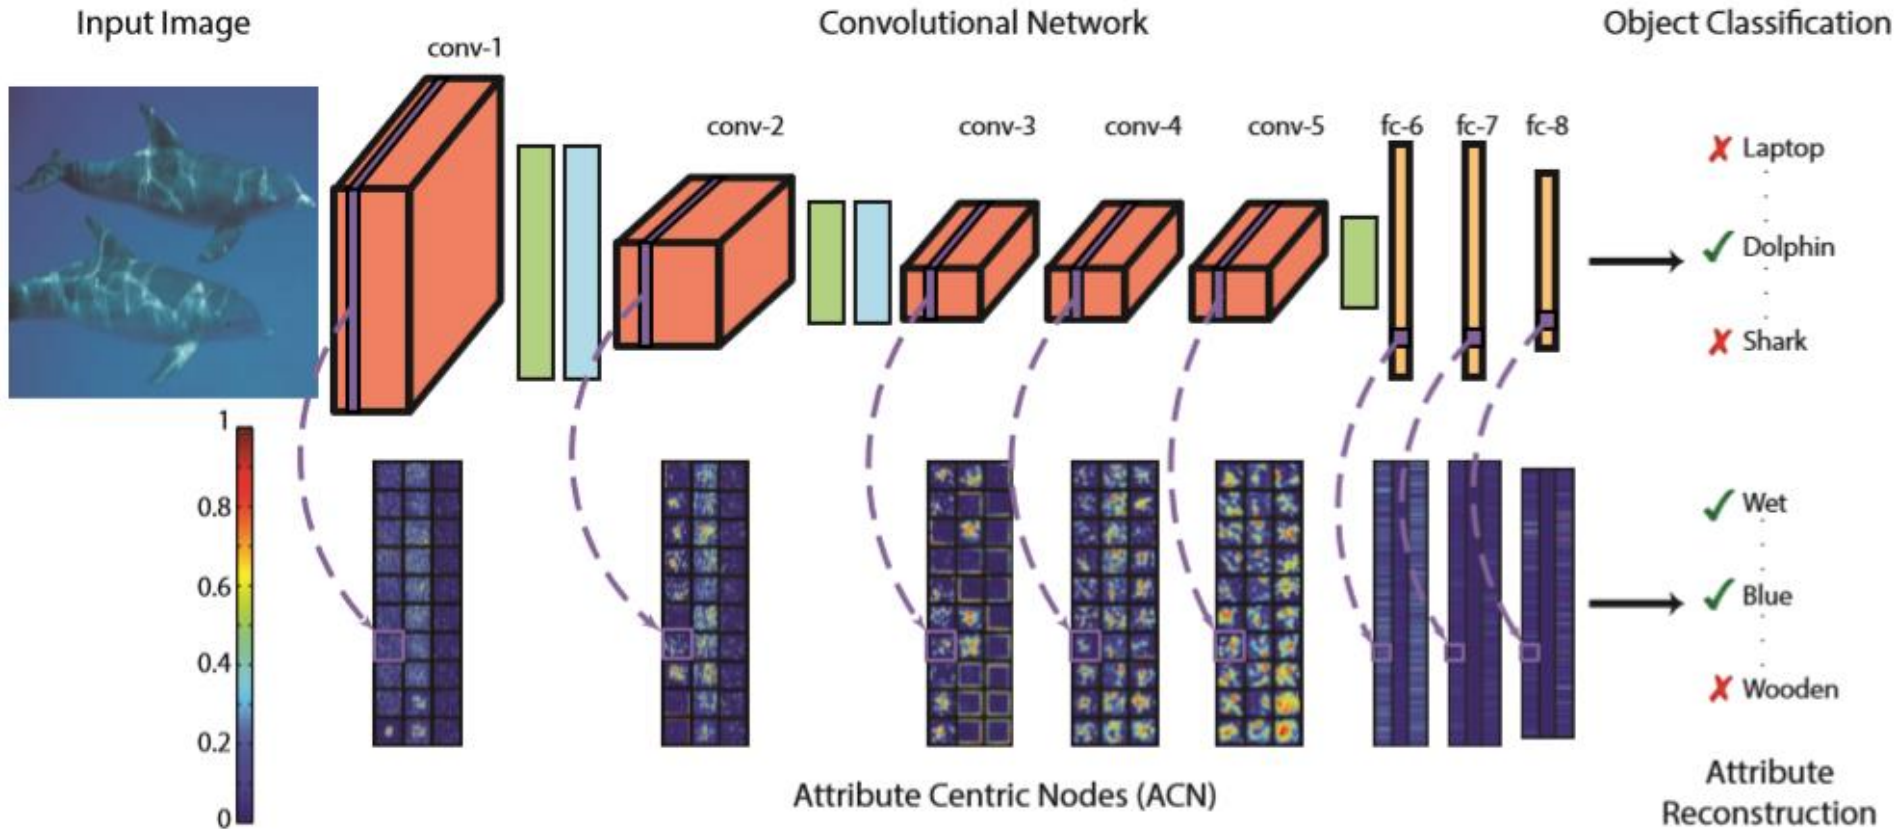
King Abdullah University of Science and Technology (KAUST), Saudi Arabia  
Universidad del Norte, Colombia

# Overview




We empirically show that a sparse number of nodes in a conv-net, trained to recognize objects (e.g. 'dolphin'), inherently encode information about semantic visual attributes (e.g. 'wet'). We call these activation locations Attribute Centric Nodes (ACNs).

# Main Findings



1. ACNs exist. They are sparse and powerful representations to recover attributes.
2. ACNs are unevenly distributed throughout the network and are attribute specific. ACNs of co-occurring attributes tend to be similar.
3. ACNs are important for object classification.


# 1. ACNs Existence



We found that only a small percentage of ACNs is required to produce near perfect reconstruction. At  $\mu = 200$  (1.2% of the conv-net nodes), the AUC-ROC score stabilizes, which is evidence that ACNs are truly sparse in the conv-net.

$\mu$  is a hyper-parameter to trade-off sparsity and reconstruction. Check details in the paper.

# 1. ACNs Existence



Attribute Group	[15]	Our
Color (8 attr)	87.5%	91.6%
Texture (7 attr)	77.5%	86.7%
Pattern (2 attr)	63.4%	86.7%
Shape (3 attr)	83.6%	88.7%
Overall (20 attr)	80.8%	<b>89.0%</b>

ACNs are powerful features for attribute prediction. They record an improvement around 10% against multiple hand-crafted features with non-linear kernels [15].


# 1. ACNs Existence

Attribute Group	Our	SVM FC-6	SVM FC-7	SVM FC-8
Color (11 attr)	90.5%	88.1%	89.6%	87.5%
Texture (8 attr)	87.7%	84.3%	83.9%	85.5%
Pattern (2 attr)	86.7%	87.6%	86.9%	85.9%
Shape (4 attr)	88.8%	90.4%	90.9%	89.0%
Overall (25 attr)	<b>89.0%</b>	87.2%	87.7%	86.9%

In general, attribute prediction based on ACNs is better than attribute prediction of one robust classifier trained over the activations of a single fully-connected layer.

It suggests that some attributes take advantage of activations from lower hidden layers.


## 2. Distribution of ACNs



We study the importance of the ACN with respect to their localization in the conv-net, for each attribute. The average normalized contribution of each layer is reported in parentheses.




## 2. Distribution of ACNs



ACN localization is attribute dependent. For example, color attributes tend to be represented well by activations from the convolutional layers.

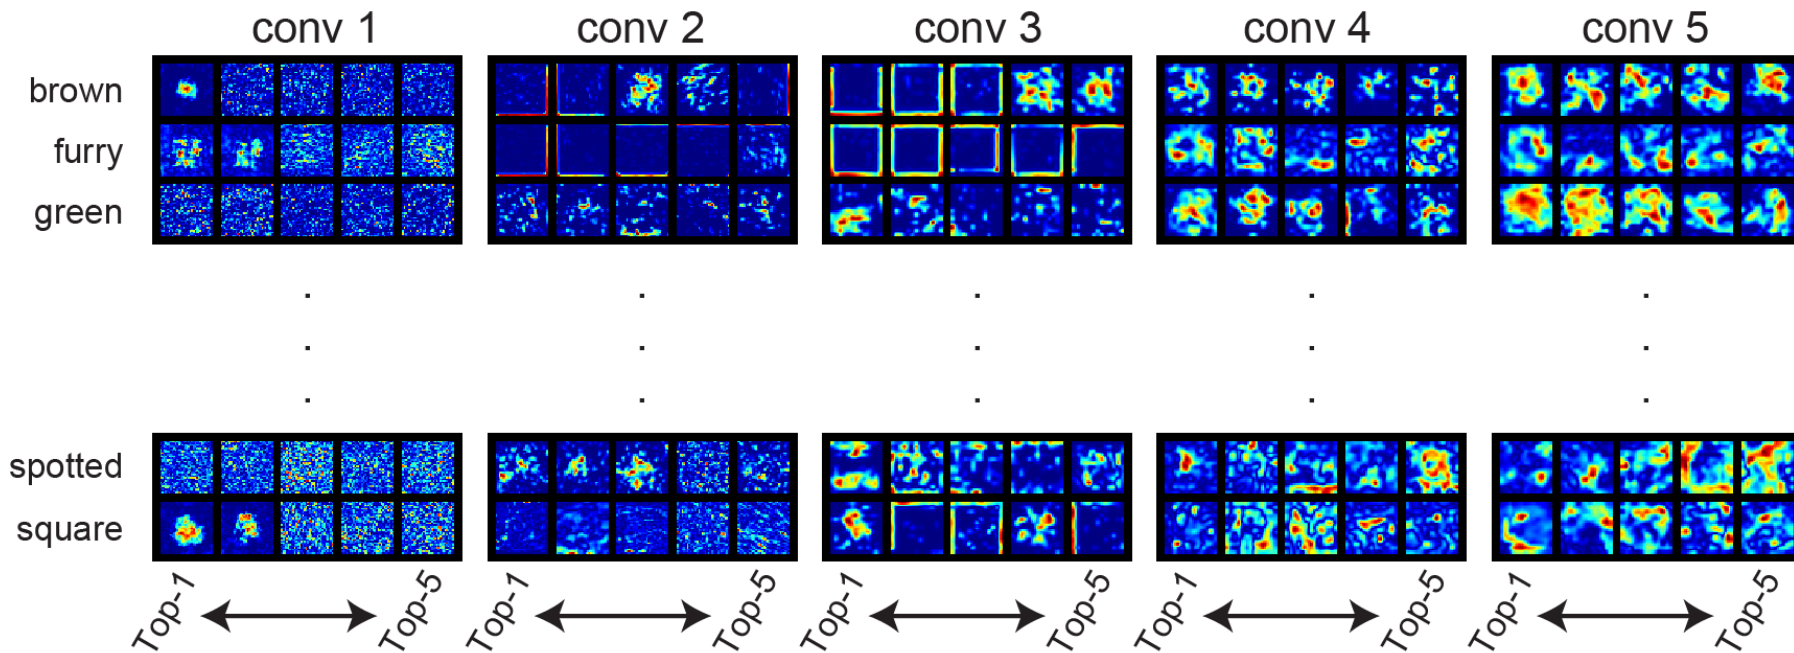


## 2. Distribution of ACNs



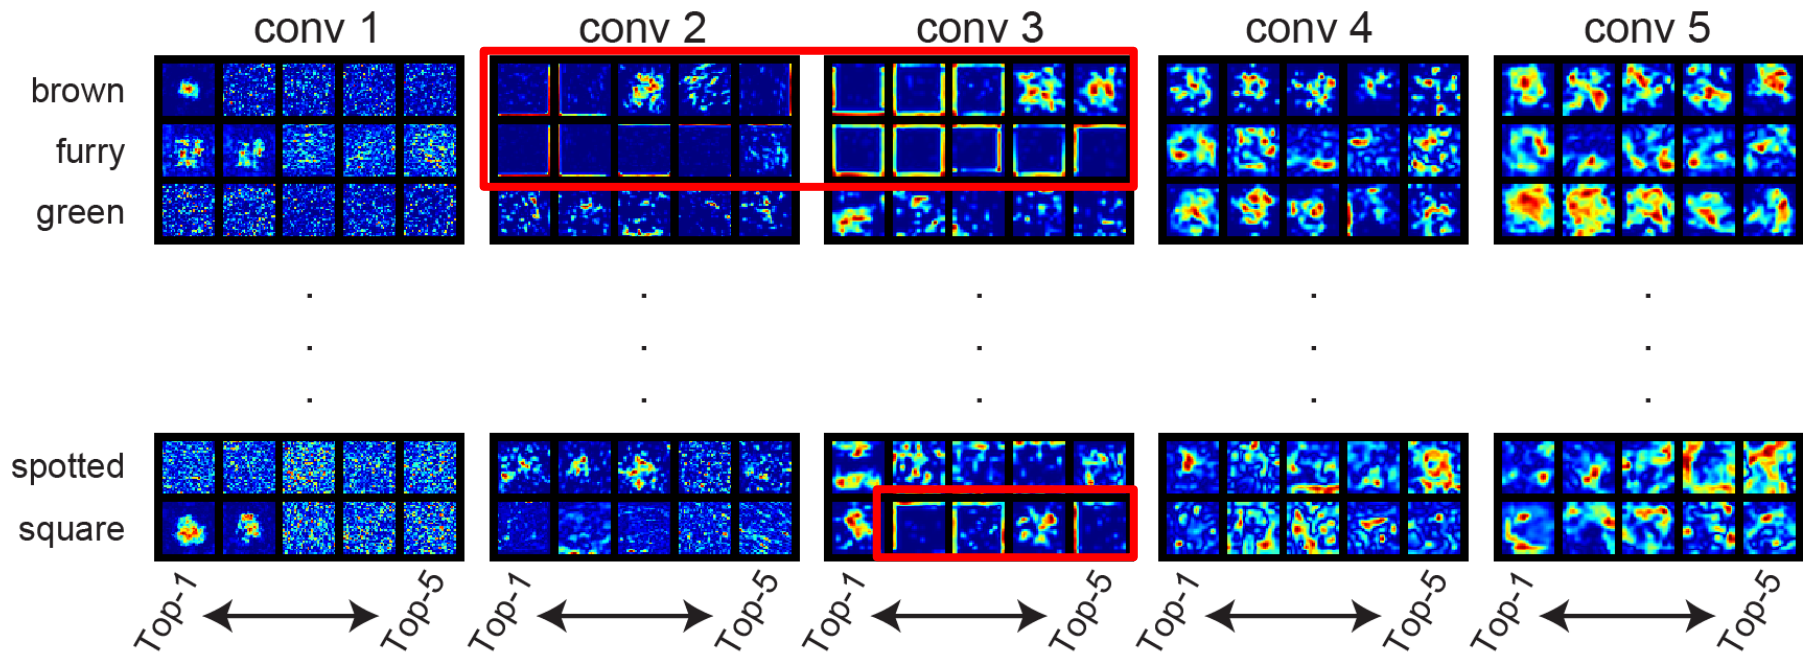
In general, the most *important* activations are localized on higher layers. This is experimental support for the common practice of using features from the fully-connected layers

## 2. Distribution of ACNs



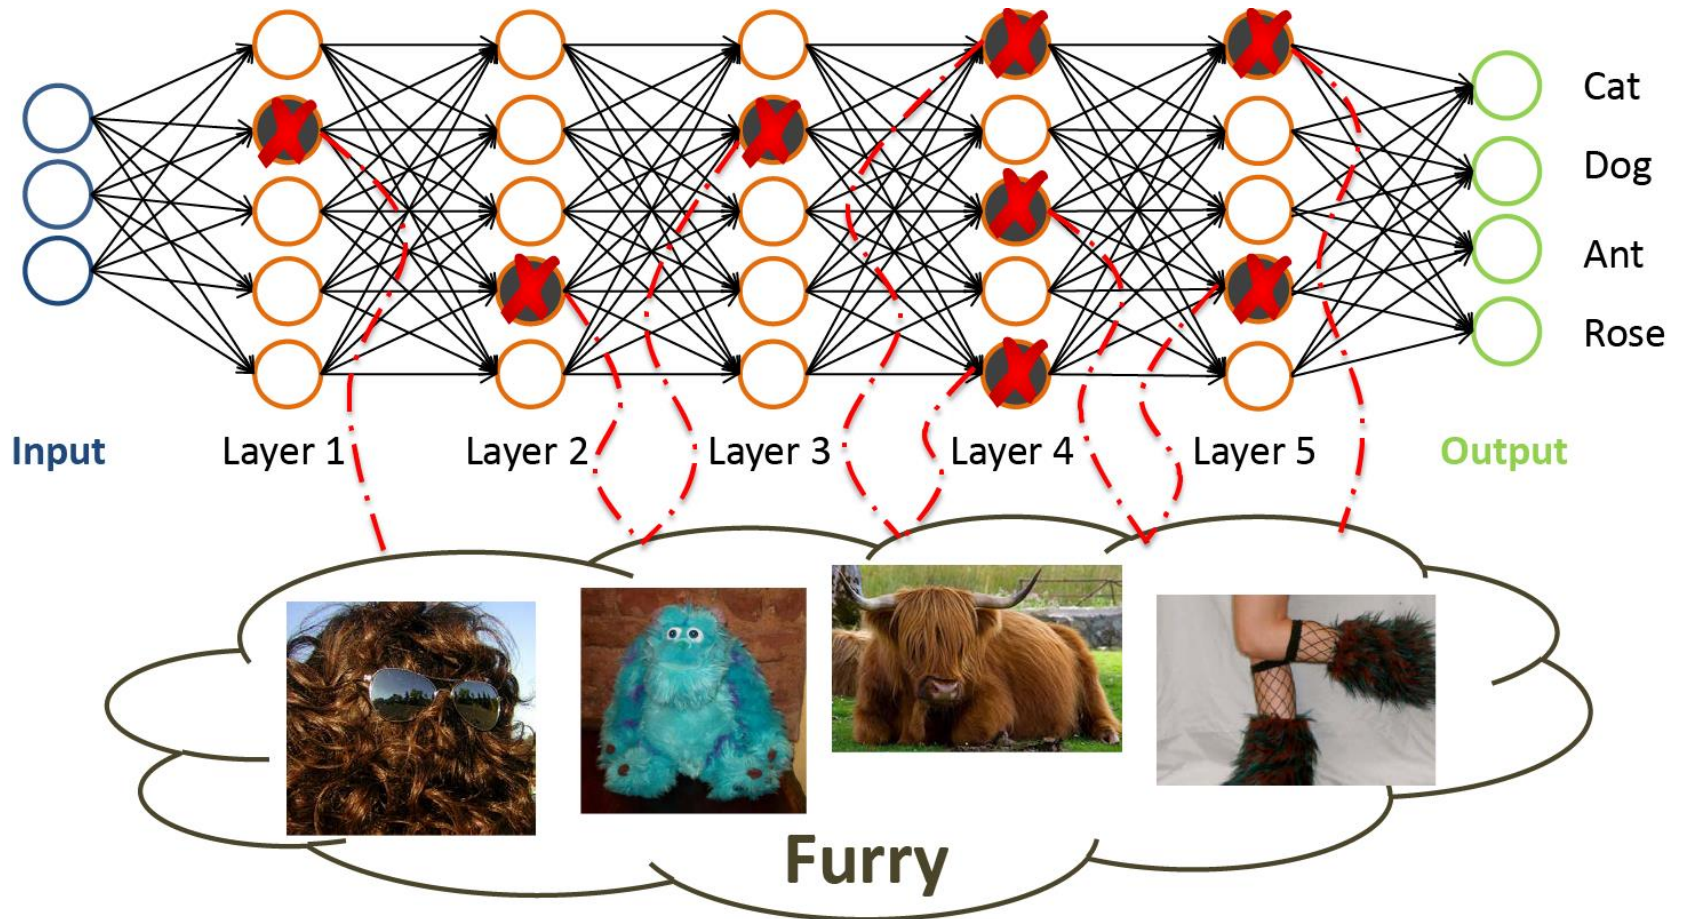
The visual pattern of ACNs in convolutional layers is diverse. It tends to be spatially structure-less in lower convolutional layers while it becomes more centralized and spatially contiguous in higher layers.

## 2. Distribution of ACNs



Interestingly, some layers (e.g. conv2 and conv3) of some attributes contain ACNs in local patterns around image borders suggesting some form of learned context.


# 3. Impact of ACNs on Object Recognition



We perform an ablation study to measure the impact of ACNs on object classification.

We compare ACN ablation to random node ablation.

# 3. Impact of ACNs on Object Recognition




We ablated ACNs associated with all 25 absolute attributes. Both ablation methods show a steep drop-off as  $\mu$  increases. However, the difference in accuracy between the two is significant.

This suggests that ACNs encode important information used by the conv-net for recognition.

The small standard deviation of random ablation is scaled up to be visually perceivable.























# 3. Impact of ACNs on Object Recognition



We also remove ACNs for each attribute separately. Similarly, we observe that ACN ablation produces a more drastic drop-off than random nodes ablation.

The small standard deviation of random ablation is scaled up to be visually perceivable.

# 3. Impact of ACNs on Object Recognition

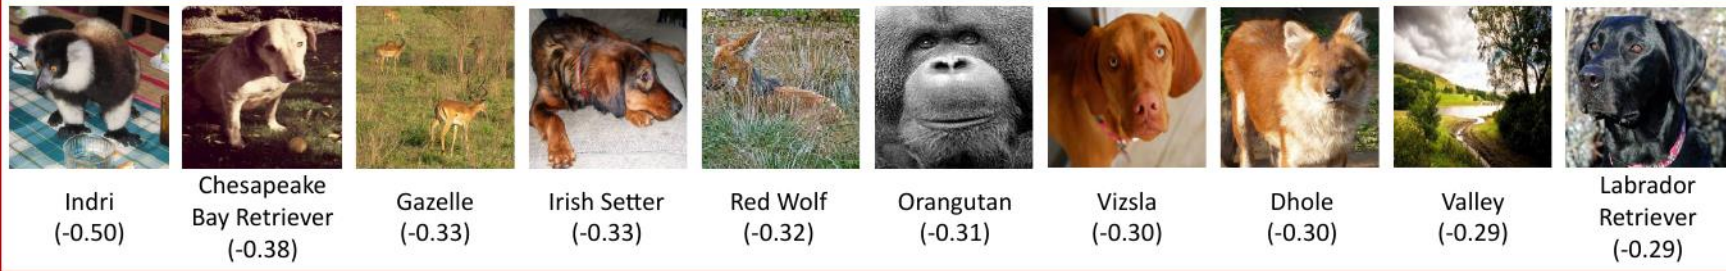
Green										
Red	Butternut Squash (-0.50)	Pomegranate (-0.42)	Acorn (-0.36)	Spaghetti Squash (-0.36)	Banana (-0.31)	Acorn Squash (-0.31)	Orange (-0.29)	Guacamole (-0.26)	Strawberry (-0.26)	Beer Glass (-0.24)
Orange										
Yellow	Wing (0.06)	American Chameleon (0.05)	Sweatshirt (0.05)	Sports Car (0.05)	Analog Clock (0.05)	B&G Garden Spider (0.05)	Space Heater (0.04)	Maillot (0.04)	Academic Gown (0.03)	Washbasin (0.03)

We also found semantic relationships between ablated ACNs of specific attributes and **the most** (and **least**) affected object categories. The drop-off in mean average precision is reported in parentheses.

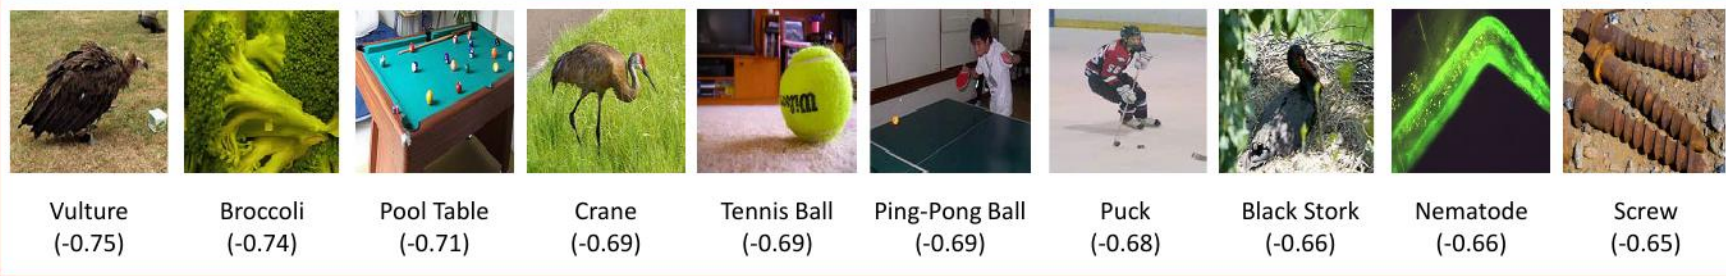


# 3. Impact of ACNs on Object Recognition

Black  
Brown  
Furry




Green  
Brown  
Wooden





# 3. Impact of ACNs on Object Recognition

Gray  
Metallic  
Shiny



Furry  
Gray  
White




# Acknowledgements



BG: Competitive  
Research Funding

VE: Visiting Student  
Research Program  
(VSRP)



[caffe.berkeleyvision.org](http://caffe.berkeleyvision.org)



JCN: Microsoft  
Research Faculty  
Fellowship



Barranquilla - Colombia

