




Visualizing Nodes in a Deep Network

Stella Yu¹, Karl Zipser²

¹ICSI, Computer Science Department, U.C. Berkeley, (stellayu@berkeley.edu)


²Helen Wills Neuroscience Institute, U.C. Berkeley, CA



original



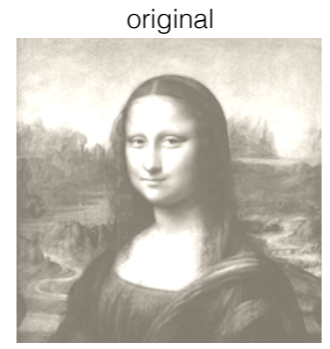
woman



promontory



woman and promontory



original

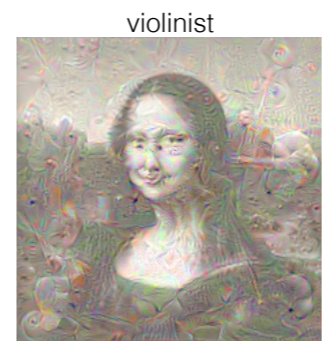
Modifying an image based on specificity of individual person category nodes yields distinct individual characteristics. Modifications can be featural as well as contextual.



woman



man




violinist



redhead



economist



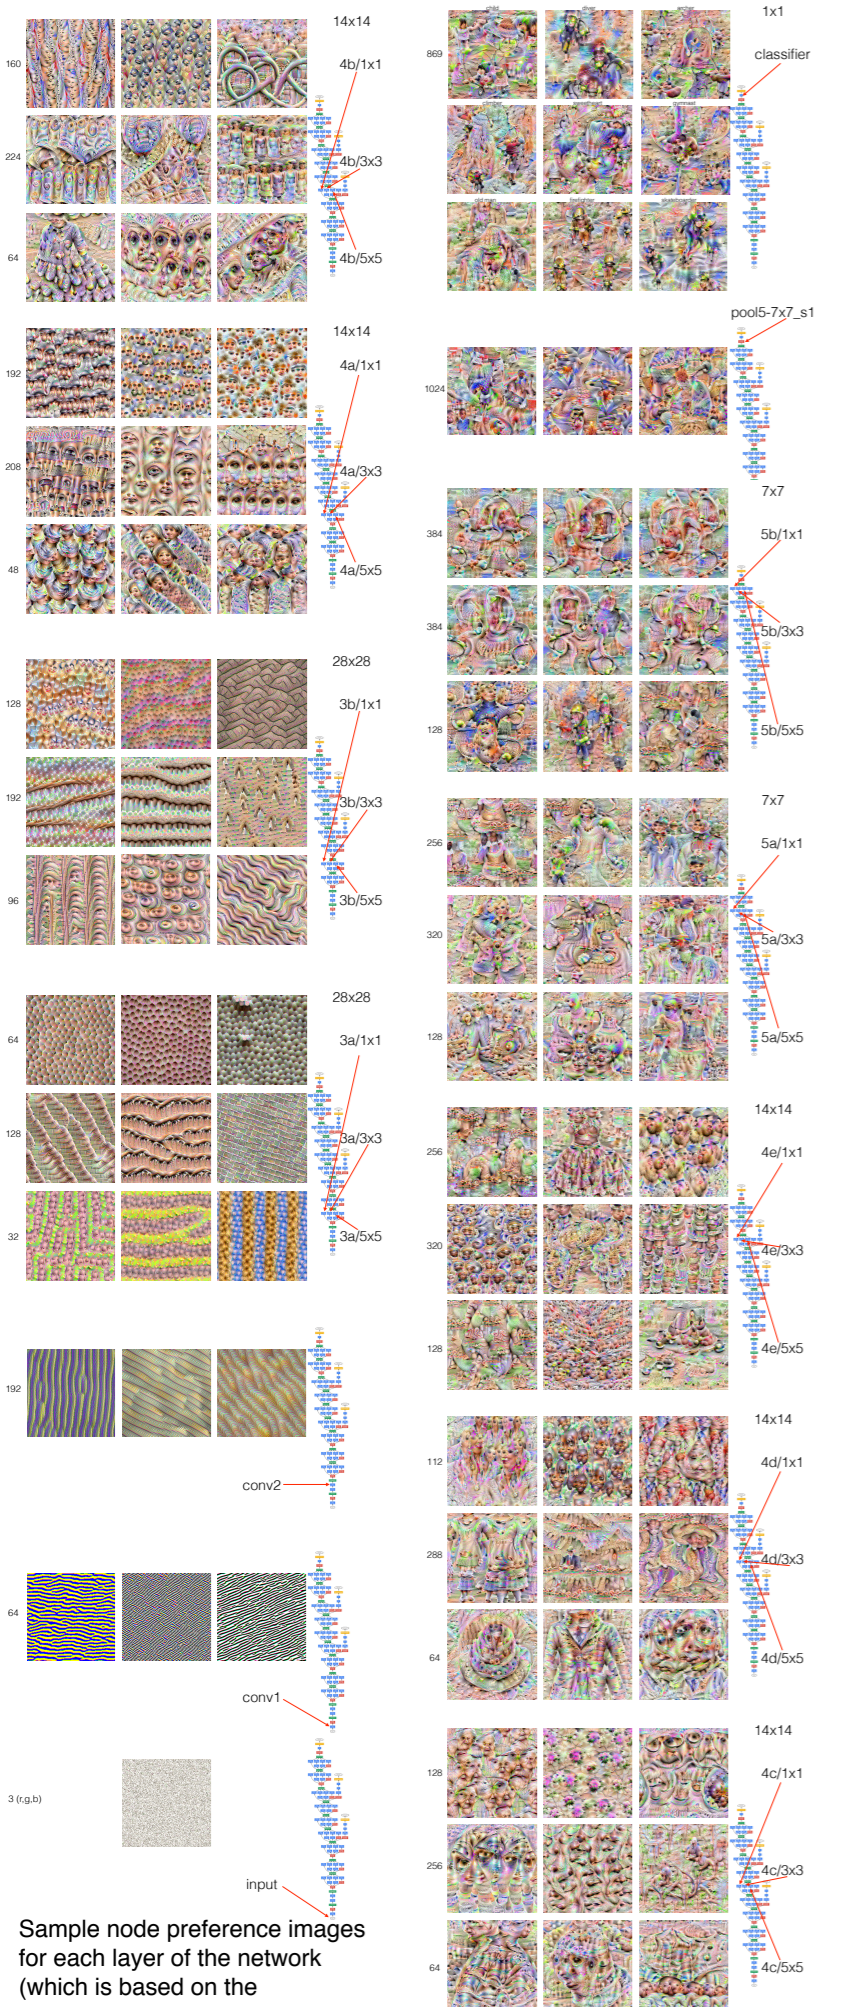
Dalai Lama




"beauty"




"adonis"



Sample node preference images for each layer of the network (which is based on the GoogleLeNet architecture).



How does RF specificity depend on RF position for convolutional nodes? For nodes with small RFs, there is no position-dependence for the vast majority of RF positions, but for nodes with large RFs, position can have a dramatic effect on RF specificity at most positions.



Tiling nodes of a single type reveals large scale structure preferences

