

# Unsupervised Learning for Representing Visual Concepts

Michael Chang


<http://arxiv.org/pdf/1602.06822v1.pdf>

# Automatic Image Captioning



**“A group of young people playing a game of frisbee.”**


# Scene Segmentation



# So. Much. Data.



# Hours of Manual Labor




# Task-Specific



**“A group of young people playing a game of frisbee.”**


# Computer Vision as Inverse Graphics



Disentangled representation: objects, parts, pose, lighting

Reflects causal structure in the world

# Model





# Learning Binary Decisions


- Hard binary decisions are hard to learn
  - no gradients
- Continuation method for smoothing between soft weighting and hard decision

$$w'_i = \frac{(w_i + \mathcal{N}(0, \sigma^2))^\gamma}{\sum_j w_j^\gamma}$$


# Unsupervised Learning for Representing Visual Concepts





# Unsupervised Learning for Representing Visual Concepts



52



1



17



# Unsupervised Learning for Representing Visual Concepts

