

---

---

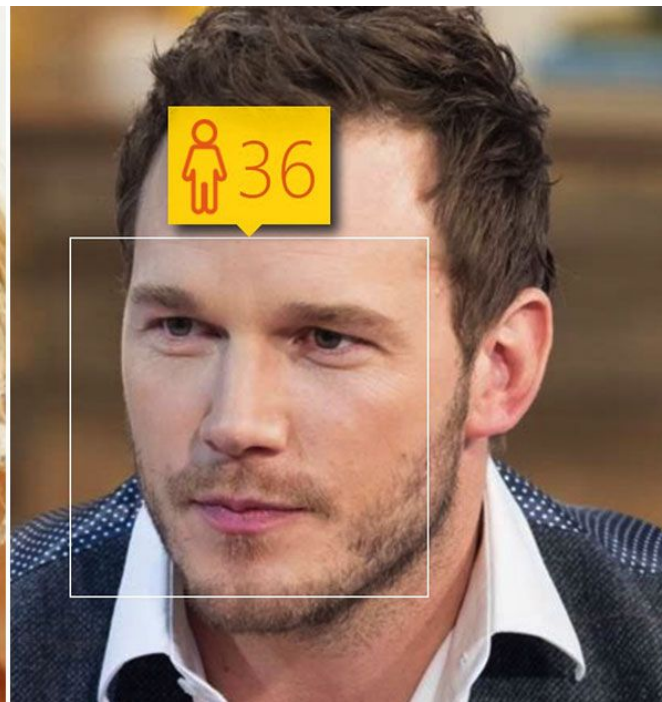
# Building an Age Guesser using Transfer Learning

— Arnav Garg, Shail Mirpuri,  
Darren Tsang —

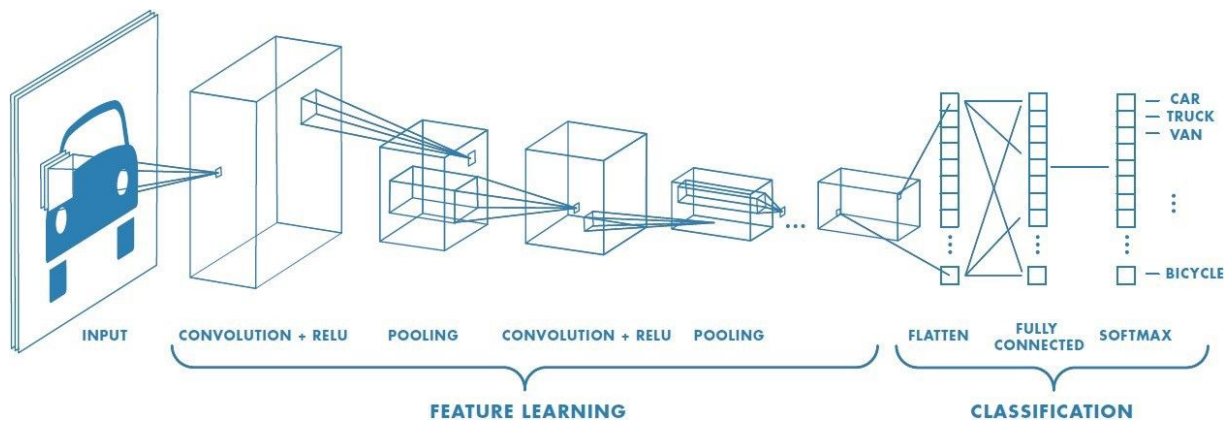
---

---

# Task



# Convolutional Neural Nets (CNNs)




1 <sub>x1</sub>	1 <sub>x0</sub>	1 <sub>x1</sub>	0	0
0 <sub>x0</sub>	1 <sub>x1</sub>	1 <sub>x0</sub>	1	0
0 <sub>x1</sub>	0 <sub>x0</sub>	1 <sub>x1</sub>	1	1
0	0	1	1	0
0	1	1	0	0

Image

4		

Convolved Feature

# Gradient Descent



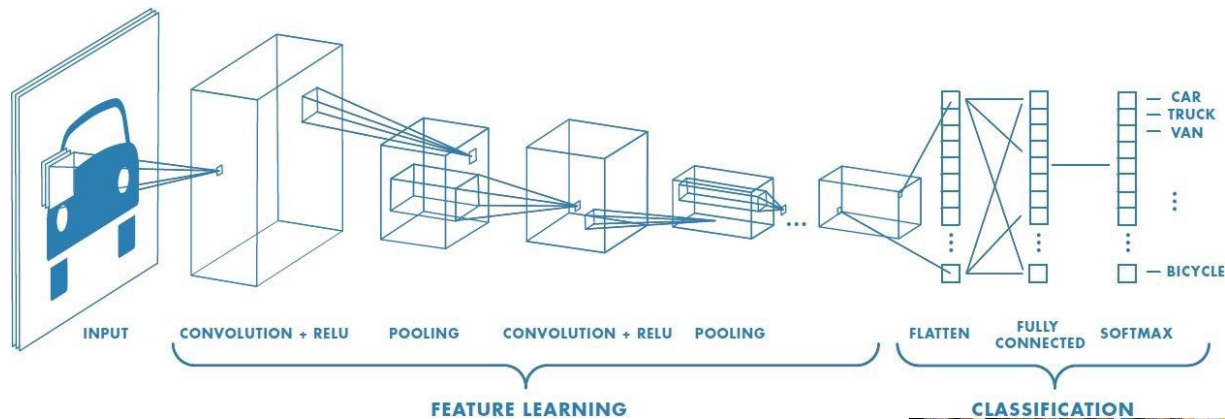
Andrew Ng

Repeat until convergence {

$$\theta_j \leftarrow \theta_j - \alpha \frac{\partial}{\partial \theta_j} J(\theta)$$

}


# Transfer Learning



# VGG-19

Consists of:

- 3 x 3 filters with a stride of 1 and zero-padding
- Max Pooling Layer with a stride of 2
- 19 layers



# The IMDB-WIKI Dataset

- 500,000+ images of celebrities from IMDB and Wikipedia
  - Collected by researchers and publicly available online
- Discarded images where the age did not make sense
  - Negative ages
  - Ages over 100
- Created subsets from cleaned dataset for training, validation, and testing
  - Training used to train model
  - Validation used during training process to ensure overfitting does not happen
  - Testing used at the very end to evaluate how well our model does on new, unseen data
- <https://data.vision.ee.ethz.ch/cvl/rrothe/imdb-wiki/>

# Hyperparameters

- Learning Rate: 0.01
- Batch Size: 10
- Trainable: Last Layer
- Optimizer: Adam
- Random Sample 50,000 images to train our dataset (about 10% of total dataset)



# Results

Good Predictions




Actual Age  
Predicted Age




# Results

## Bad Predictions

45  
27



61  
29




Actual Age  
Predicted Age

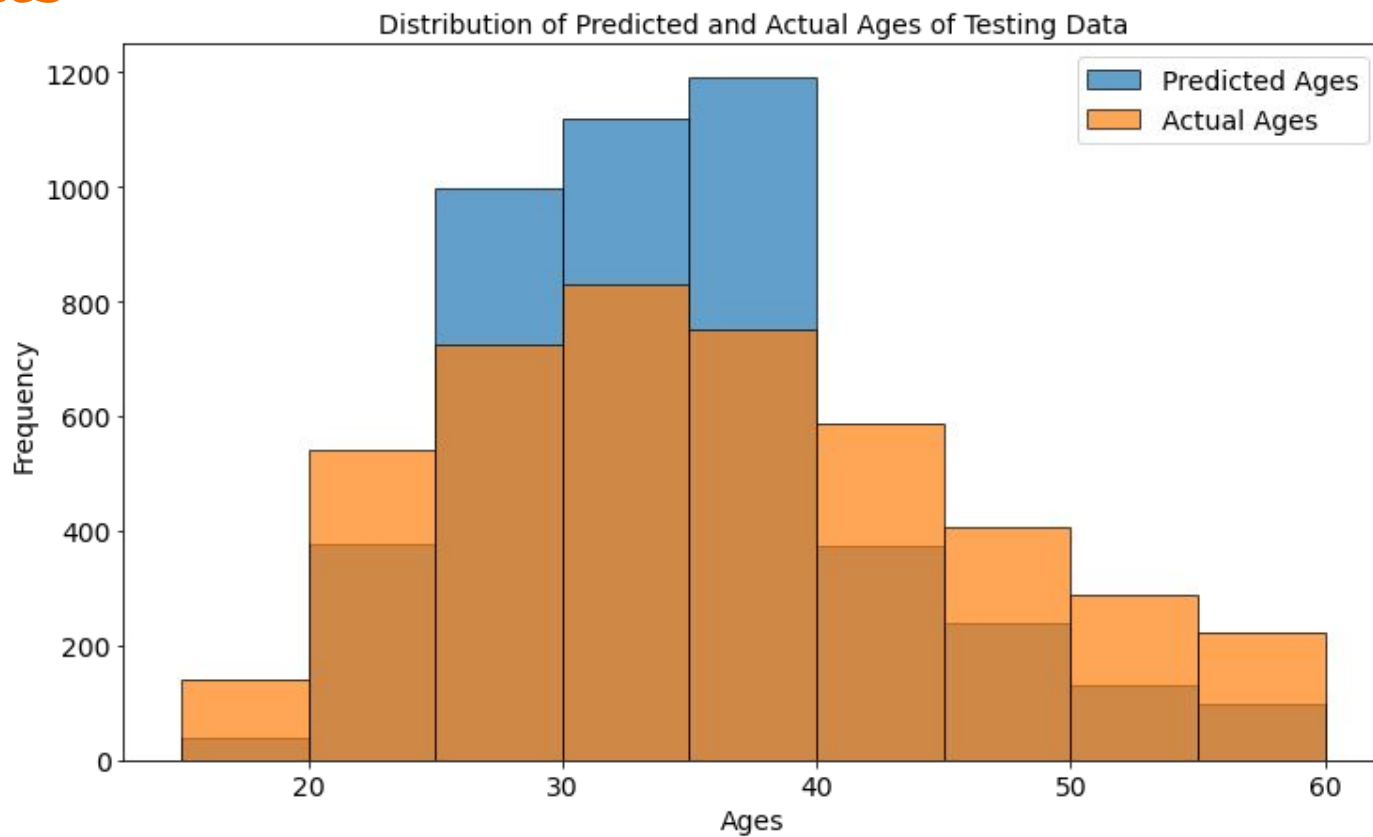
16  
39



38  
21

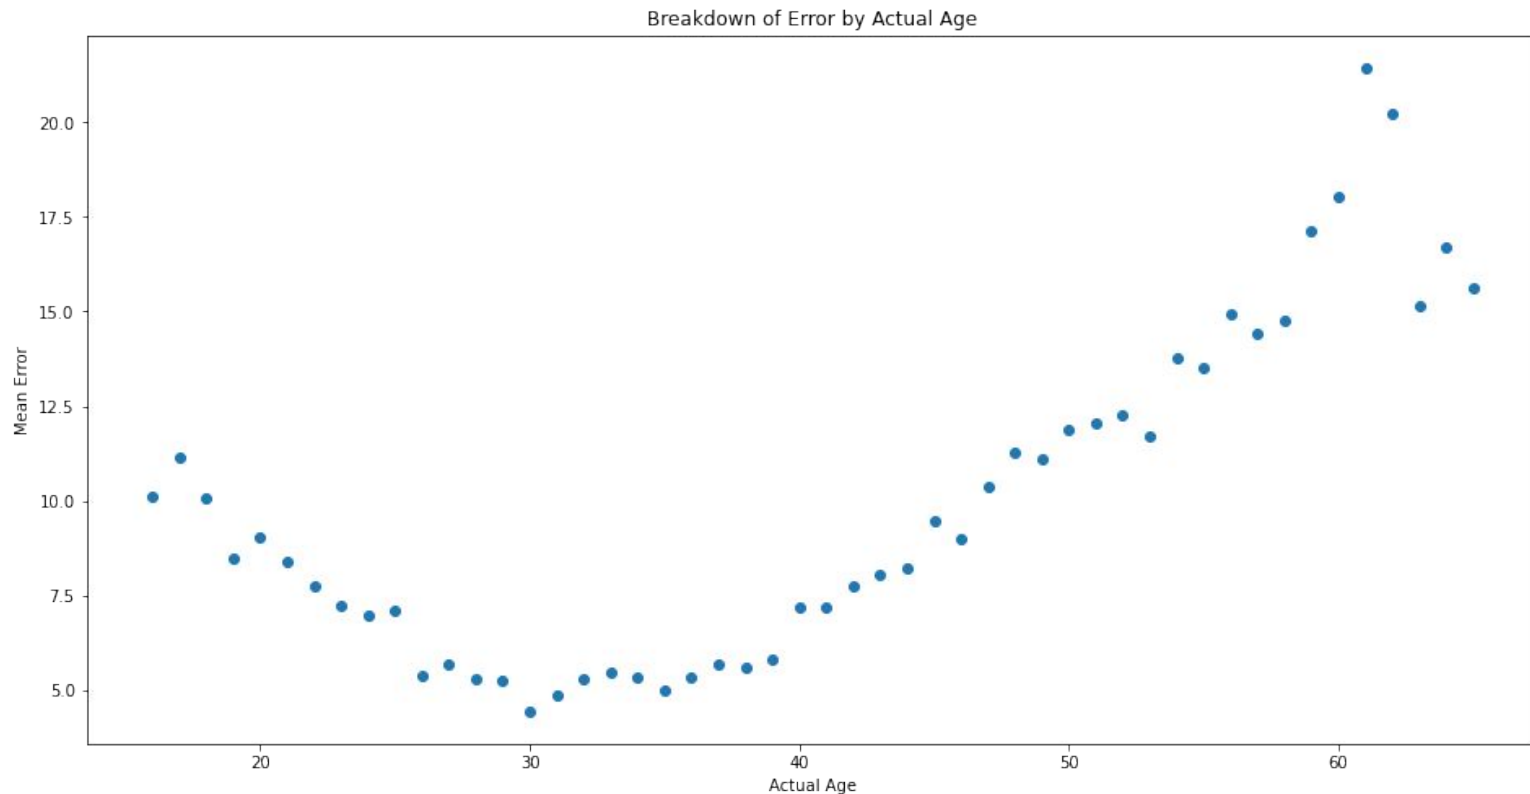


# Results

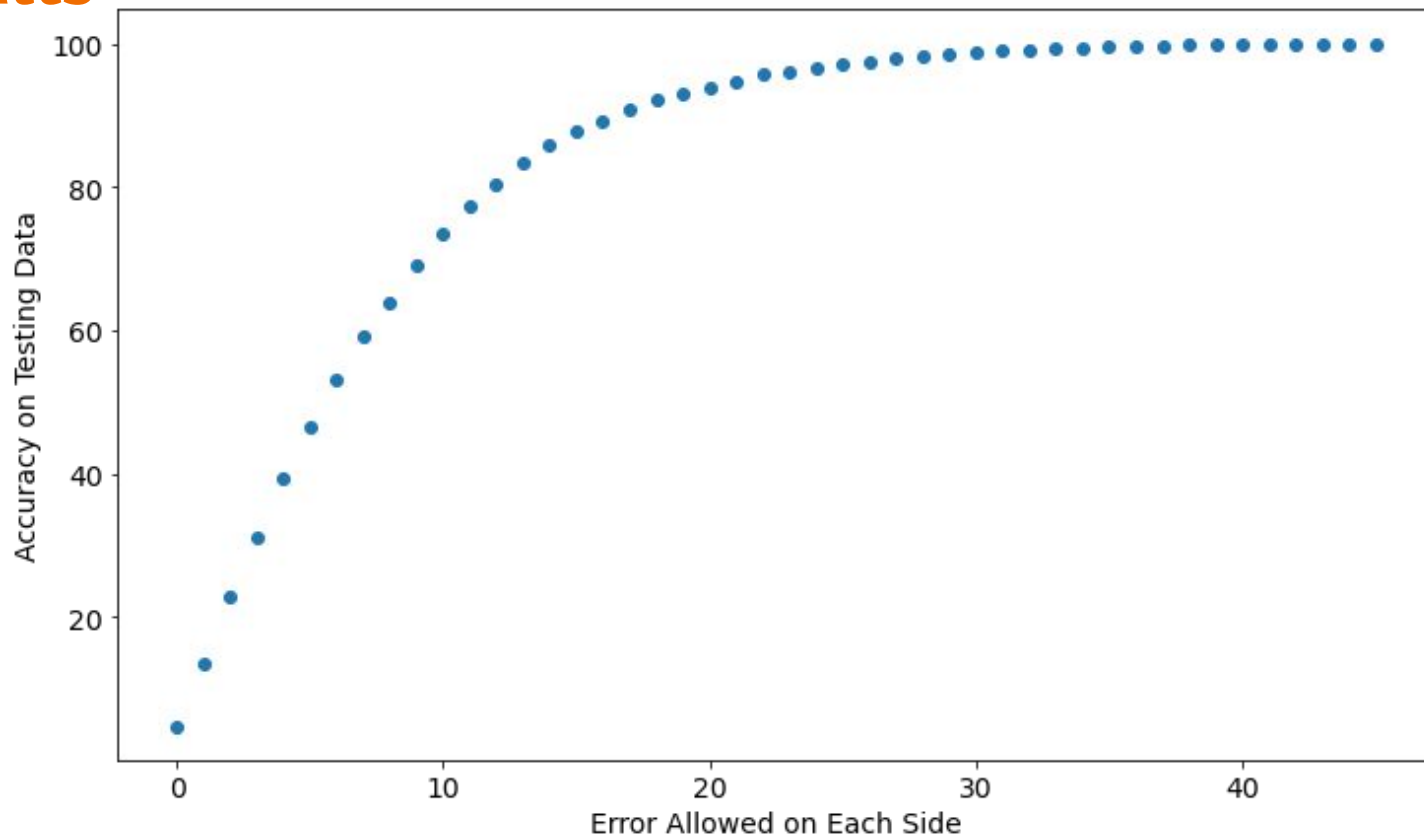


# Results

- **Mean Absolute Error: 7.83** years
- **Standard Deviation of Absolute Error: 6.84**



# Results



# Challenges

- Real age vs apparent age
- Massive dataset, low computing power
- Age skew in dataset



# Future Work

- Make it work on uncropped images
- Multiple people in images
- Train for longer periods of time on full dataset
- Build user application to expand dataset
- Use loss function that better accounts for close guesses
- Look into newer, more advanced network architectures (ResNet, Inception)