# CS231N: Classifying Artistic Eras using Deep CNNs

Yangyang Yu, Daniel Hsu, Olivier Jin

## Problem Statement

Artistic eras are often seen as fluid boundaries with no clear classification. Therefore, our goal is to objectively define each period using one of three techniques: Constructing a custom deep CNN, utilizing transfer learning, and applying transfer learning while fine-tuning the last convolution layer. With these methods, we demonstrate that deep CNNs are a viable solution to classifying artwork.

## Background

The difficulties in defining an artistic era include:

- The wide range of subject matter (portraits, still life, etc)
- Shared artistic techniques between eras

Previous papers relied on a wide range of methods, some involving CNNs, others focusing directly on lower-level features such as color, brush strokes, and texture patterns.

## Dataset Characterization

We used the Pandora18K dataset:

- 18,038 images over 18 artistic styles
- Ranges from from Byzantine to pop art.
- Non-uniform distribution of images across classes

Most images in our dataset have an aspect ratio of 1:1 with width and height around 500 pixels. Since CNNs are designed for a specific input size, we use Keras to rescale the images to 500x500 pixels. This lets us keep most images relatively unchanged,

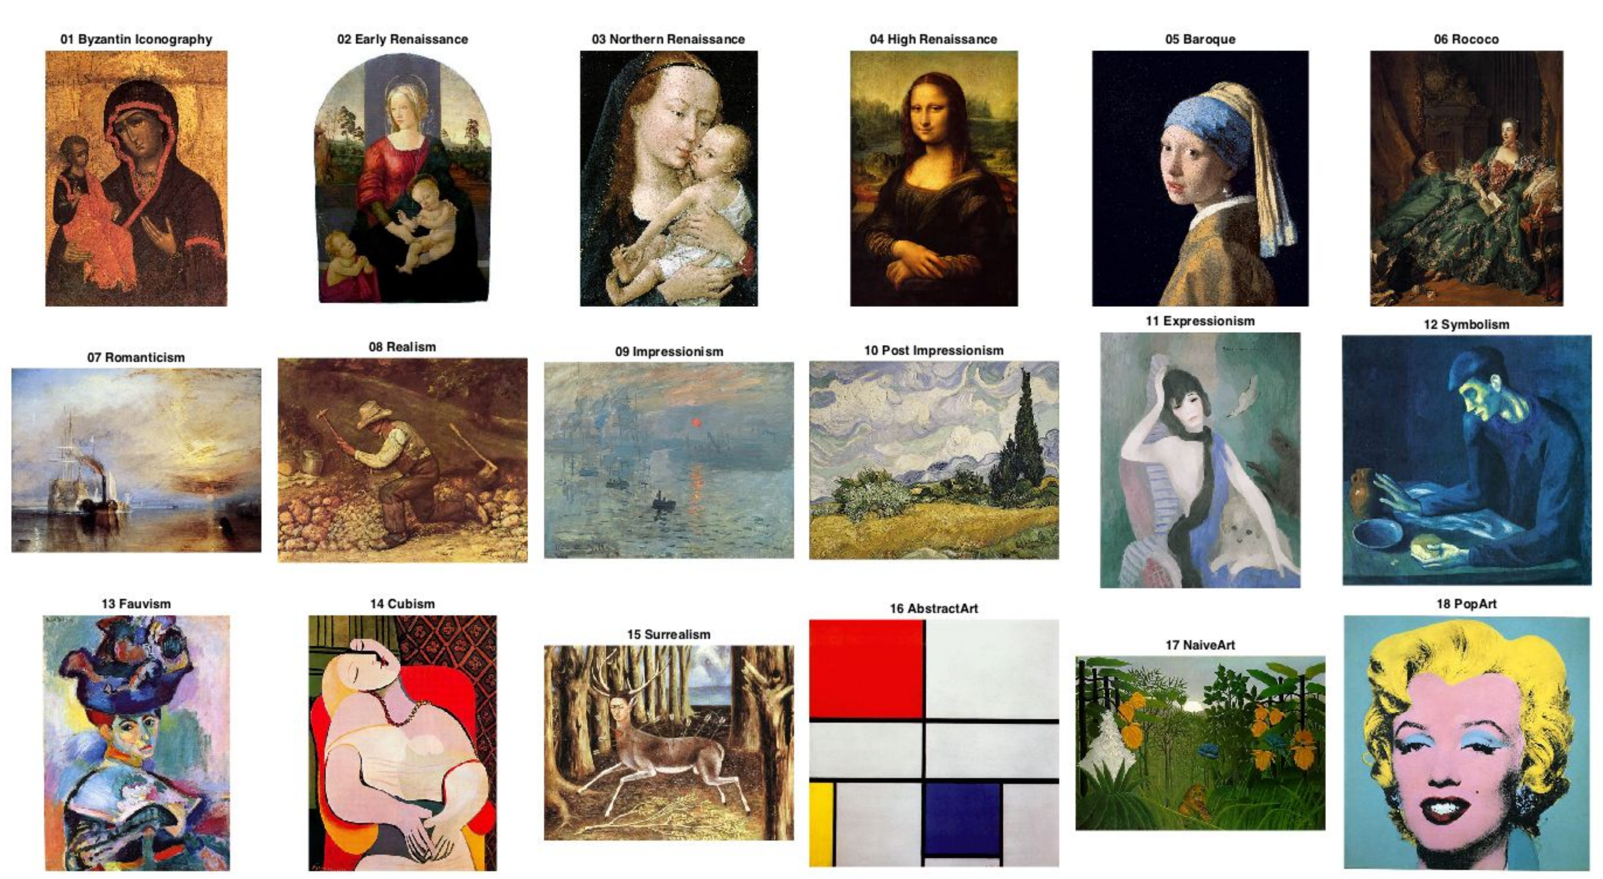

## Custom CNN Architecture

Custom architecture based on both VGGNet and Inception:

- Four 3x3 conv layers
- Max-pooling layer
- An Inception-style layer
  - 1x1, 3x3, 5x5, pool
- Two dense layers
- Softmax loss

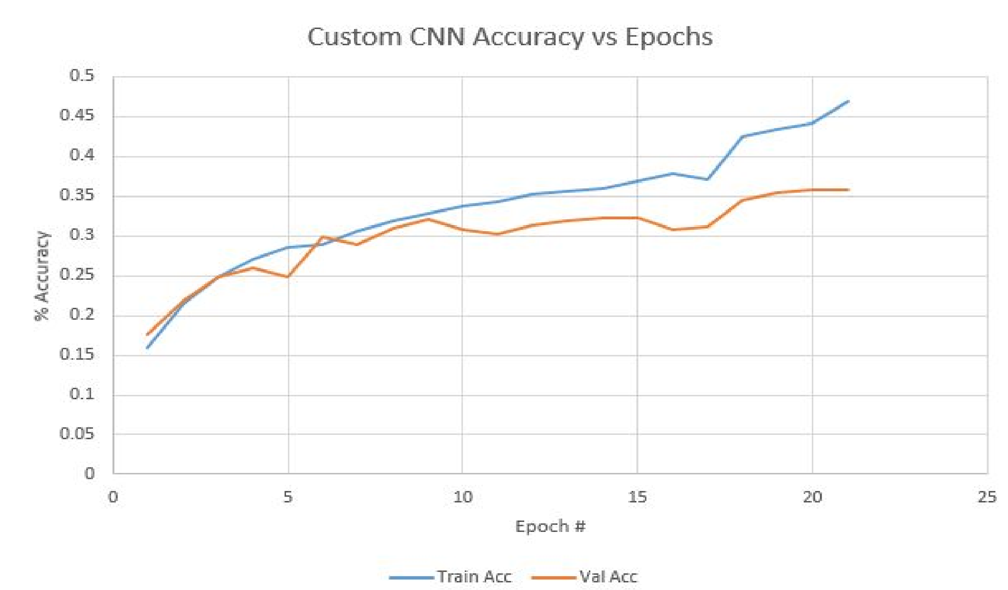


Final val accuracy: 35.76%

## Transfer Learning w/ Inception

Transfer learning used on InceptionV3 pre-trained with ImageNet

- Pre-trained InceptionV3
  - No dense layers
- 70% dropout + dense
- 70% dropout + dense
- Softmax loss
- Store weights to HDF5 file

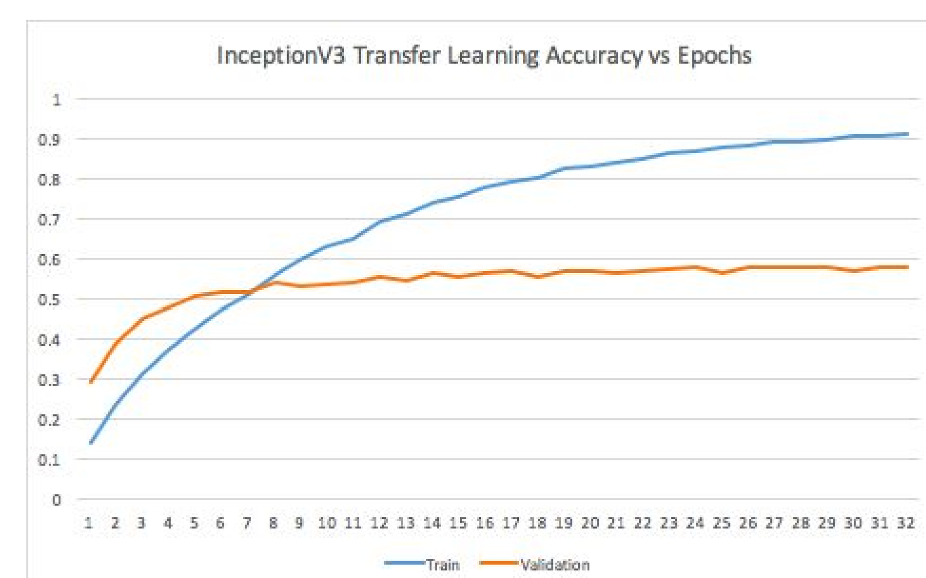


Final val accuracy: 57.84%

## Transfer Learning + Fine Tuning

Model initialized using weights learned from transfer learning

- Pre-trained InceptionV3
- Load weights from HDF5 file
- Freeze all layers
  - Unfreeze last conv layer
  - Unfreeze FC layers

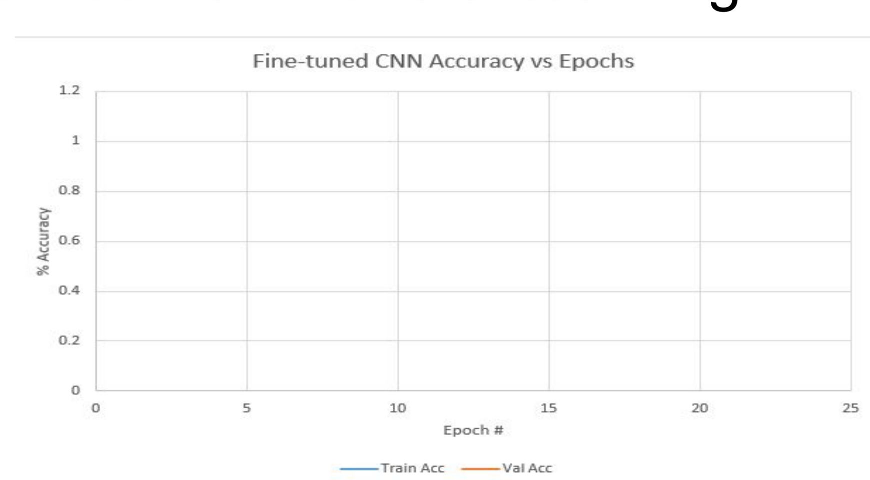


Final val accuracy: TBD%

## Findings & Discussion

We used test accuracy to identify our best technique. This metric demonstrates that transfer learning is the best approach for our task. Our final test accuracy of 56.6% is better than the accuracy found in Florea et al[1]

**Observations:**

- Adjacent eras are harder to differentiate
- Some eras have similar features (Impress. vs Post-Impress.)
- Byzantine art is most unique and easily classified

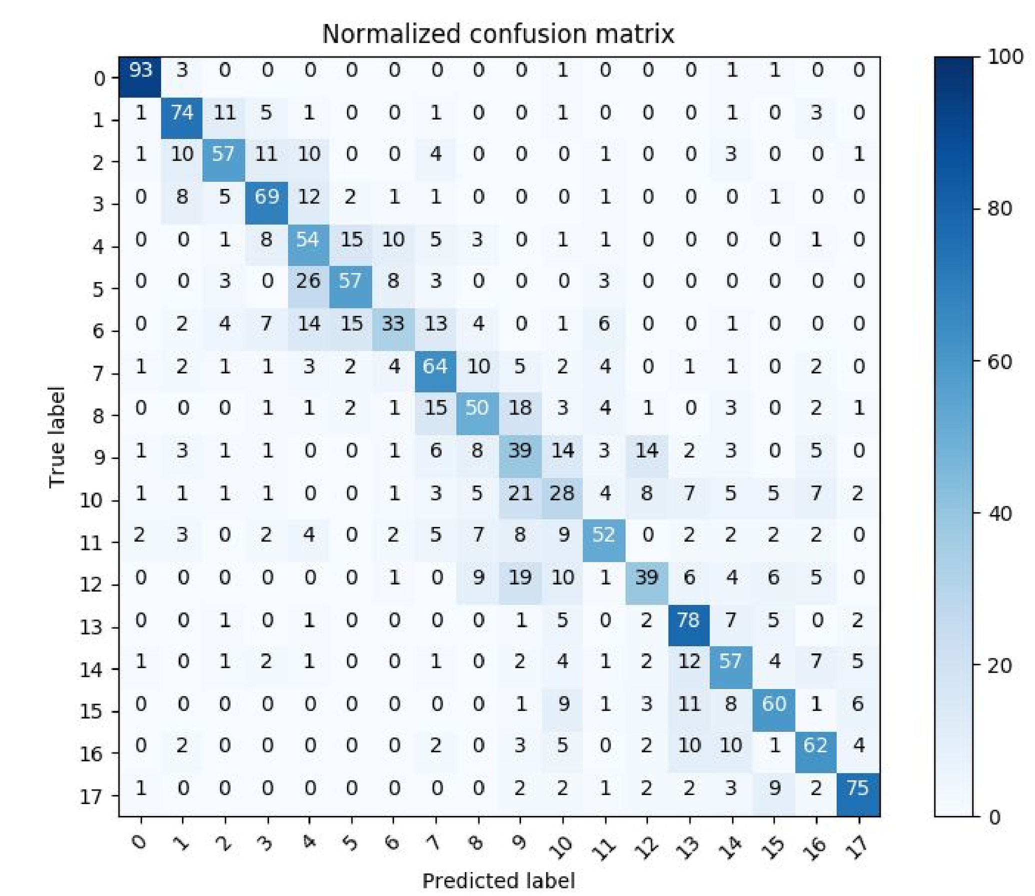


[1] C. Florea, C. Toca, F. Gieske," Artistic Movement Recognition by Boosted Fusion of Color Structure and Topographic Description", WACV 2017, Santa Rosa, USA, pp. 569-577

## Conclusion

We achieved high baselines from all three of our methods. Our custom CNN architecture and the prelearned weights from ImageNet on Google's Inception architecture are both well suited for extracting features for artistic style. We plan on pursuing transfer learning with other networks such as ResNet and extending transfer learning to additional layers.

## Acknowledgements

We would like to acknowledge Florea et al. for their efforts in compiling the Pandora dataset and publishing initial findings.