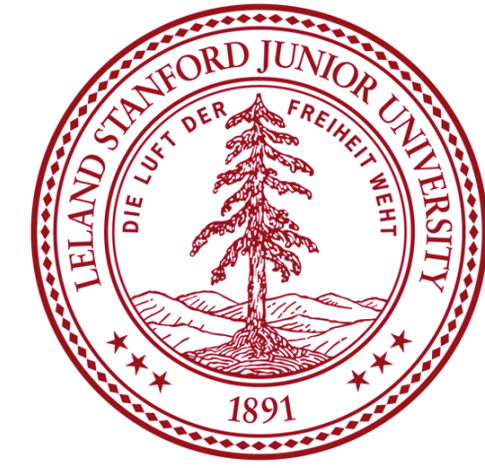


PAINTING GENRE CLASSIFICATION

Presentation link: <https://youtu.be/na6YOaF5jEw>
Soheil Golara
CS230-Fall 2019 Final Project Report



Background

- Using deep learning to determine the genre of a painting
- X: painting image \rightarrow Y: one of 44 genres

Data

- Kaggle's dataset "Painters by numbers"
- 80000 images and their corresponding style, genre, artist, and date
- gray-scale images repeated 3 times to match the RGB dimension
- 4th channel in RGBA images dropped to match the RGB dimension
- Manually downloading more images of rare genres

Models

Transfer Learning in tensorflow

- Input size (300,400,3)
- Adding a shallow (one or two layers) NN to the 'avgpool5' output of a pre-trained VGG19 model

ResNet-50 in Keras

- Input size (128,128,3)
- Trainig all layers

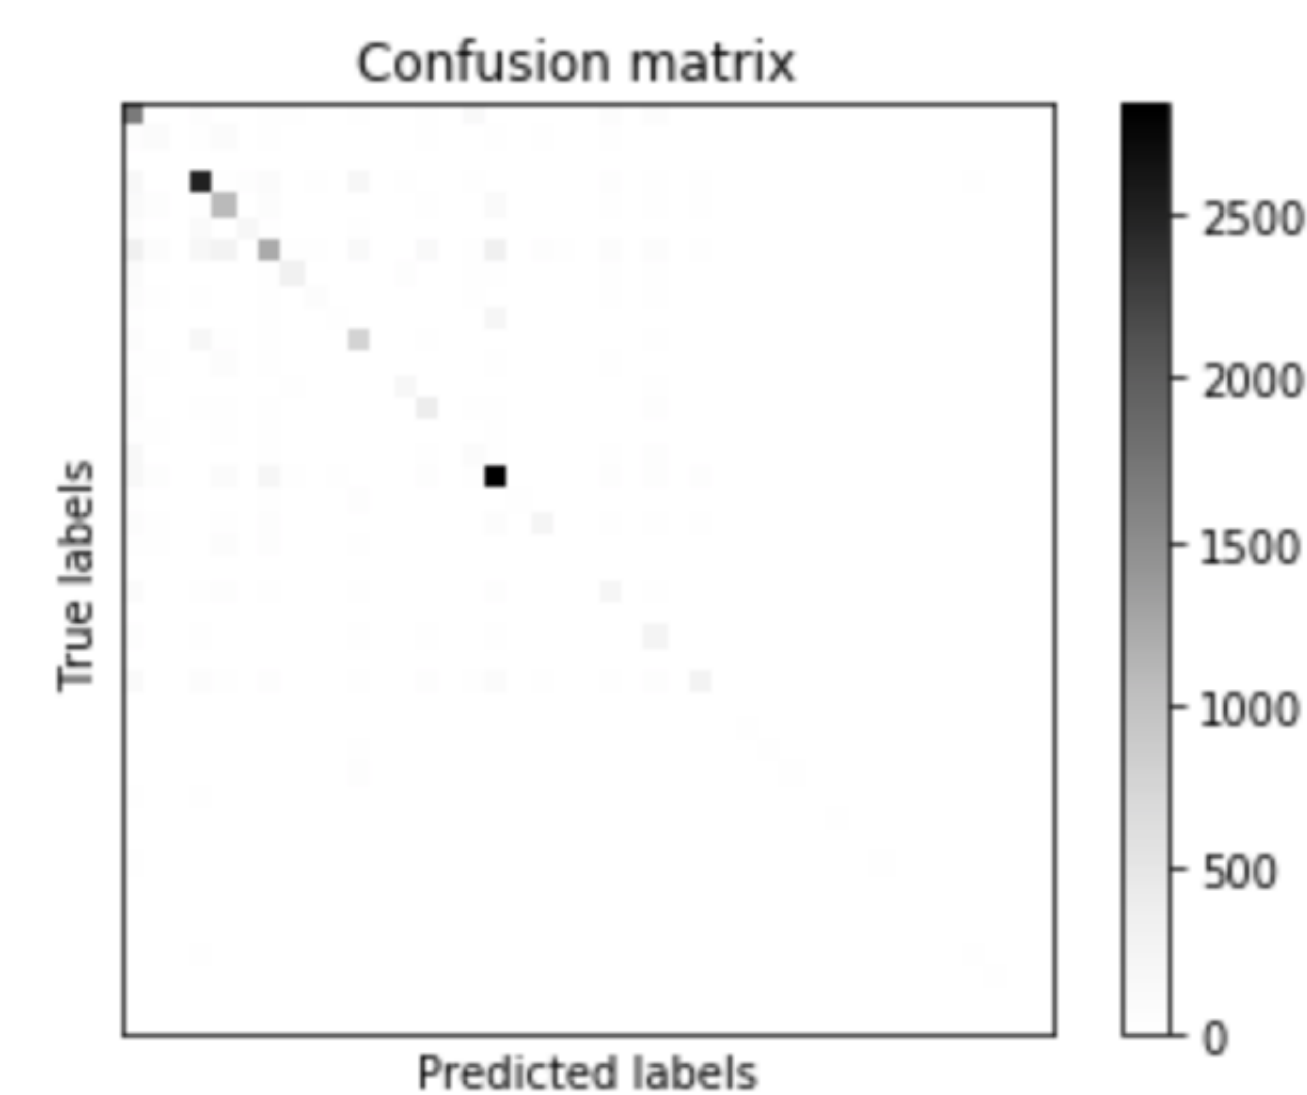
Transfer Learning in Keras using ResNet-50

- Input size (224,224,3)
- Adding a softmax layer

Results and Discussion

Model	Train Acc (%)	Test Acc (%)
VGG-19 (transfer learning)	98 (%)	56 (%)
ResNet-50 (Fully trained)	97 (%)	71 (%)
ResNet-50 (1. Keep top layer, add softmax)	37 (%)	37 (%)
ResNet-50 (2. Remove top layer, add softmax)	69 (%)	61 (%)
ResNet-50 (3. re-train last 24 layers + softmax)	95 (%)	62 (%)

- Models generally have high variance
- Variance problem due to data imbalance
- Confusion matrix show the imbalance in the data



- Rare genres are more likely to get confused with common ones
- Some genres are more susceptible to be confused with each other and penalizing their misclassification more in the loss function could improve the model
- "genre-painting" image misclassified as "figurative":



- "history" image misclassified as "landscape":



Future work

- Data augmentation script to increase number of images of rare genres
- More comprehensive hyperparameters search and trying different ways of regularization
- Using GAN to generate paintings of a particular genre

References

- Sean Chang, Jeffrey Dong, Chen, Patrick James Tanaka Mogan. Determining style of paintings using deep learning and convolutional neural networks.
- Adrian Lecoutre, Benjamin Negrevergne, and Florian Yger. Recognizing art style automatically in painting with deep learning. In Asian conference on machine learning, pages 327-342, 2017.
- Jana Zujovic, Lisa Gandy, Scott Friedman, Bryan Pardo, and Thrasyvoulos N Pappas. Classifying paintings by artistic genre: An analysis of features & classifiers. In 2009 IEEE International Workshop on Multimedia Signal Processing, pages 1-5. IEEE, 2009.
- Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In Advances in neural information processing systems, pages 2672-2680, 2014.
- Gan by example using keras on tensorflow backend. URL <https://towardsdatascience.com/gan-by-example-using-keras-on-tensorflow-backend-1a6d515a60d0>.
- Generating modern arts. URL https://github.com/Anyesh/art_gan.
- Painters by numbers. URL <https://www.kaggle.com/c/painter-by-numbers/>.
- URL <https://github.com/solig/Painting-Genre-Classification.git>.
- Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556, 2014.
- URL <http://www.vlfeat.org/matconvnet/pretrained>.
- URL <https://github.com/anishathalye/neural-style>. 5
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In Proceedings of the IEEE conference on computer vision and pattern recognition, pages 770-778, 2016.
- Ch-14 : Generative adversarial networks (GAN's) with math. URL <https://medium.com/deep-math-machine-learning-ai>.

