



APPLICATIONS OF A CNN FOR AUTOMATIC CLASSIFICATION OF OUTSOLE FEATURES

Miranda Tilton, M.S.

Feb 20, 2020

What is the probability of a coincidental match?



What is the probability of a coincidental match?

Given a particular shoe outsole...

1. Define the comparison population
2. Sample N shoes from the comparison population
3. Count the number of similar shoes S from the comparison population that are similar to the given shoe
4. Estimate the probability of a coincidental match:

$$\hat{p} = \frac{S}{N}$$

Obstacles: Characterizing Comparison Populations

- No 100% complete database of all shoes
 - manufacturer, model, size, tread style, manufacturing molds
- Shoe purchases vs. frequency of wear
- Local populations may differ wildly (Benedict, et al., 2014)



Comparison Population

How to collect data from the comparison population?

1. Build a low profile scanner, place in a high traffic area
2. Scan shoes of those walking past
3. Create a local-area database of relevant scans

Goal:

$$\hat{p} = \frac{S}{N}$$

This is an engineering problem

Comparison Population

Assume a machine exists that can scan shoe outsoles of pedestrians

1. Identify relevant features within the scans



Comparison Population

Assume a machine exists that can scan shoe outsoles of pedestrians

1. Identify relevant features within the scans
2. Define similarity for shoe images



Comparison Population

Assume a machine exists that can scan shoe outsoles of pedestrians

1. Identify relevant features within the scans
2. Define similarity for shoe images
3. Assess the frequency of similar shoes in the sampled data




Relevant Features

Use features other than make/model and size to characterize shoes

- Knockoffs often have very similar tread patterns
- Similar styles have similar tread patterns across brands
- Unknown shoes can still be classified and assessed

Dr. Martens	Eastland	Timberland
		
Work 2295 Rigger	1955 Edition Jett	6" Premium Boot

Relevant Features

Bowtie	Chevron	Circle
		
Line	Polygon	Quadrilateral
		
Star	Text	Triangle
		

Used to separate shoes by make/model in (small) local samples
(Gross, et al., 2013)

IMAGE ANALYSIS AND FEATURE DETECTION

Image Analysis

Goal: Identify geometric tread features in images of shoe outsoles

- Robust to different lighting conditions, rotation, image quality
- Fast processing of new images
- Identify features that are explainable to practitioners




Analogy: Human Classification

Our eyes detect features in an image, and our brains learn to connect features with labels after seeing many examples.




Analogy: Human Classification



CONVOLUTIONAL NEURAL NETWORKS

CNN Architecture



Transfer Learning


- Very deep CNNs can require > 1 million images to optimize performance
- Using a model base trained on different input data can greatly reduce the required number of images
- Our approach
 - Use pretrained convolutional base: VGG16
 - Train a new model head

VGG16

- Pre-trained CNN (Simonyan, et al., 2014)
 - Trained on 1.3 million images from ImageNet (Krizhevsky, et al., 2012)
 - Simple structure

FITTING CoNNOR: CONVOLUTIONAL NEURAL NETWORK FOR OUTSOLE RECOGNITION

Acquire Data



ShoeScrapR
package


over 80,000
images scraped
since April 2018



- Jenny Kim
- Ben Wonderlin
- Mya Fisher
- Holden Jud
- Miranda Tilton
- Charlotte Roiger
- Susan VanderPlas
- Joe Zemmels
- and others

- LabelMe Annotation Tool used as a web interface - creates XML files with labels and coordinates. (Russell, et al., 2008)
- 27,710 regions labeled with one or more geometric objects
- 37,562 labels

Label Data




Model Training

- 256 x 256 pixel images
- Training data (60%):
 - 1x Augmented images (rotation, skew, zoom, crop) to prevent overfitting
 - Class weights used to counteract uneven class sizes
- Validation and test data (20% each)
- Fit using the keras package in R, which provides a high-level API for the tensorflow library




Model Training




Binary Cross-entropy Loss:

$$-y \log(p) - (1-y) \log(1-p)$$


Evaluating the Model



Evaluating the Model




Evaluating the Model



For multi-label images, only incorrect predictions contribute to off-diagonal probabilities

EER_i used as the cutoff

Interpreting the model Class Activation Maps




Blue: Prediction matches image label

Grey: Prediction does not match image label

Heatmaps are scaled by class. Yellow = high activation

Interpreting the model Class Activation Maps




Blue: Prediction matches image label

Grey: Prediction does not match image label

Heatmaps are scaled by class. Yellow = high activation

Interpreting the model Class Activation Maps



Blue: Prediction matches image label

Grey: Prediction does not match image label

Heatmaps are scaled by class. Yellow = high activation

Project Summary

- Geometric shapes provide a convenient feature space for assessing shoe similarity
- Transfer learning allows application of CNNs to much smaller datasets
- CoNNOR performs well
 - Reduction in feature space: $256 \times 256 \times 3 \rightarrow 9$
 - 88% accuracy; many errors attributable to data labeling

References

- [1] I. Benedict, E. Corke, R. Morgan-Smith, et al. "Geographical variation of shoeprint comparison class correspondences". In: *Science and Justice* 54.5 (2014), pp. 335–337.
- [2] S. Gross, D. Jeppesen, and C. Neumann. "The variability and significance of class characteristics in footwear impressions". In: *Journal of Forensic Identification* 63.3 (2013), p. 332.
- [3] A. Krizhevsky, I. Sutskever, and G. E. Hinton. "ImageNet Classification with Deep Convolutional Neural Networks". In: *Advances in Neural Information Processing Systems* 25. Ed. by F. Pereira, C. J. C. Burges, L. Bottou and K. Q. Weinberger. Curran Associates, Inc., 2012, pp. 1097–1105.
- [4] B. C. Russell, A. Torralba, K. P. Murphy, et al. "LabelMe: A Database and Web-Based Tool for Image Annotation". En. In: *International Journal of Computer Vision* 77.1-3 (May. 2008). 02464, pp. 157–173. ISSN: 0920-5691, 1573-1405. DOI: [10.1007/s11263-007-0090-8](https://doi.org/10.1007/s11263-007-0090-8). URL: <http://link.springer.com/10.1007/s11263-007-0090-8>.
- [5] K. Simonyan and A. Zisserman. "Very Deep Convolutional Networks for Large-Scale Image Recognition". En. In: *arxiv.org* (Sep. 2014). URL: <https://arxiv.org/abs/1409.1556>.

Tools

- R Packages and Toolkits:
 - Modeling: `keras`, `tensorflow`
 - Data Wrangling: `magrittr`, `dplyr`, `lubridate`, `stringr`, `tidyr`, `purrr`, `furrr`
 - Image Processing: `jpeg`, `imager`, `magick`
 - Annotation Manipulation: `sf`, `sp`
 - Visualization: `ggplot2`, `viridis`, `ggcorrplot`, `deepviz`, `tidygraph`, `ggraph`, `shiny`
 - XML/Web Scraping: `xml2`, `XML`, `rvest`, `RSelenium`
 - Slides/Documentation: `rmarkdown`, `xaringan`, `knitr`
- Other Software: Docker, Selenium, LabelMe Annotation Tool (w/ Matlab toolbox), gimp image editor

THANK YOU