Afronet: Shopping Afrowear from Anywhere

Oluwasanya Awe ~ Robel Mengistu ~ Vikram Sreedhar
oawe@stanford.edu ~ robel@stanford.edu ~ vsreed@stanford.edu

Motivation

The African clothing industry is a $20 Billion market, largely overlooked by American retailers [1]. The industry presently benefits from strong tailwinds and consumers have greater buying power than ever before.

The Status Quo is Frustrating

Figure 1. A typical afrowear retail experience.

Our Solution

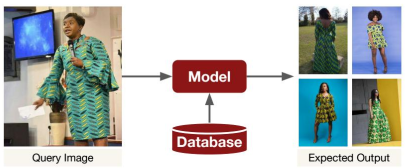

Figure 2. High-level diagram showing our model's input and outputs.

Dataset

We scraped African shopping websites for over 15K images of labelled afrowear with 14 different categories.

Triplet Data Sampling

Semi-hard swatch triplets
38K+ triplets that violate the margin constraint. Theses negatives were generated using Photoshop scripts.

Hard negatives category triplet
60K+ triplets of different categories that violate the triplet constraint.

Figure 3. Sample semi-hard swatch triplets

Model

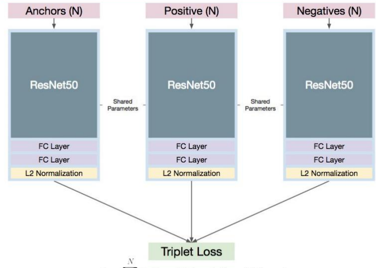

$$L = \sum_{i=1}^{N} [||f_i^a - f_i^p||_2^2 - ||f_i^a - f_i^n||_2^2 + \alpha]_+$$

where f_i^* is the learned embedding of $*$.

Figure 4. Triplet loss architecture trained on the same network with shared parameters.

We used a fine-tuned ResNet50 to extract features of images such that similar-looking images are closer together in the vector space. We apply our trained model on the database of stock images and return the top-K images to the query image [2].
We constrained the output of our network to unit length in order to have a constant margin invariant to embedding length.

Metrics

Color Histogram Correlation

$$d(H_1, H_2) = \frac{\sum_I (H_1(I) - \bar{H}_1)(H_2(I) - \bar{H}_2)}{\sqrt{\sum_I (H_1(I) - \bar{H}_1)^2 \sum_I (H_2(I) - \bar{H}_2)^2}} \qquad \bar{H}_k = \frac{1}{N} \sum_J H_k(J)$$

We extracted a square from the center of each image based on half of the minimum height and width of either of the images being compared. We then computed the correlation between their two color histogram based on the above formula where N is total number of histogram bins [3].

Mean Average Precision

For our accuracy metric, we used Mean Average Precision on rank *n* to evaluate our model [4].

$$MAP(\hat{y}, y) = \frac{1}{N} \sum_{i=1}^{N} \frac{\sum_{j=1}^{i} score(y_j, \hat{y}_j)}{i}$$

Results

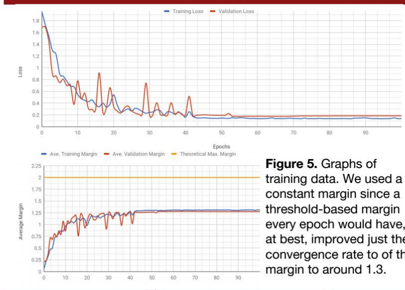

Figure 5. Graphs of training data. We used a constant margin since a threshold-based margin every epoch would have, at best, improved just the convergence rate to of the margin to around 1.3.

Triplets	Model 1	Model 2	Model 3
Train Margin	**1.31**	1.30	1.29
Val. Margin	**1.28**	**1.28**	1.24

Model 1	Train MAP	Test MAP
	0.779	0.5779

Table 1. Table of results from experiments.

Models and Important Hyperparameters

1. Model 1: FC Layers: 3, Final Dim: 512, Margin 2.0
2. Model 2: FC Layers: 3, Final Dim: 1024, Margin 3.0
3. Model 3: FC Layers: 3, Final Dim: 1024, Margin 2.0

Discussion

Expectations: Our model performed slightly better than the baseline as it was able to learn that the background should be ignored.

Impact of Background: We cropped as much of the background as possible, however, the distracting background colors still negatively affected our similarity metrics.

Measuring Style: Cropping to the center of images helped improve our test accuracy for color using histogram correlation between the query and output images.

Margin Size: Setting a higher margin generally led to better results in training. But from margins greater than 2.0, there was no significant impact on the margin after it converged.

Sample Outputs

Figure 7. Sample outputs from baseline and Model 1

Future Work

Learning Margins: Updating margins from backpropagating loss.
Training on More Data: Gathering more non-stock data and extending the dataset to include more categories.
Exploring Similarity Metrics: Categorizing data using attribute semantics.

References

[1] Niyi Aderibigbe. Why the world should invest in African fashion. 2014.
[2] Convolutional features for keras' pretrained neural nets. 2017.
[3] Rosebrock, Adrian. "How-To: 3 Ways to Compare Histograms Using OpenCV and Python." PyImageSearch, 14 July 2014.
[4] Christopher D Manning, Prabhakar Raghavan, Hinrich Schütze, et al. Introduction to information retrieval, volume 1. Cambridge university press Cambridge, 2008.
[5] Icon Images from Noun Project (Gan Khoon Lay, Mateo Zlatar).