Predicting the Popularity of Images Through Machine Learning

Jessica Liao, Arjun Krishna, Dr. Klaus Mueller

Stony Brook University

Introduction

Background

People develop opinions about the aesthetic appeal of an image at first sight. Being able to predict these preferences will enable an improved understanding of overall human predilections. In order to accomplish this, the user study Infographic Aesthetics: Designing for the First Impression, by Harrison, Reinecke and Chang, was used. The study conducted an online survey, having participants score an image's appeal on a 9-point Likert Scale after being exposed to it for 500 milliseconds. The researchers ran 30 trials per subject, making a total of 38340 different samples.

Objective

Design a neural network that is able to predict the appeal of an image to a person, based on his or her gender, age, and education level. A Likert Scale, a rating from 1 to 9, will be used as a form of measurement for the appeal of an image.

Strongly Dislike	Dislike	Moderately Dislike	Mildly Dislike	Neutral	Mildly Like	Moderately Like	Like	Strongly Like
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)

Study Results

- Females prefer more colorful and less complex infographics
- Males generally prefer fewer and less saturated colors and are relatively unaffected by different levels of complexity
- Preference for simple infographics slightly increases with age and education level
- A participant's gender, age, and education level explains 34% of the variance of participants' ratings of appeal

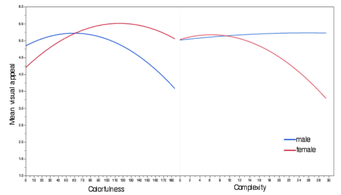

Method

Dataset

The study had individual programmers filter through a set of 2,000 images to eliminate those that were considered controversial, previously seen, lacking in data visualization, and non-english. As a result, a total of 330 images were able to be used for the study and our neural network. Each image was then resized into 64px x 64px to make training the neural network smoother and quicker.

The dataset was divided in order to train the network. To do this, a filter of the participants' gender, age and education level was applied to the overall dataset to produce a training set. The largest subgroup was chosen as the training set: 8370 samples from high school females, between the ages of 10 - 19.

Neural Network

A neural network is a form of machine learning. It is composed of multiple layers of computational units called neurons; each neuron multiplies an initial number by a value known as a weight. The sum of all the products in the layer is taken and is then adjusted by a bias. The result of this computation is then passed into an activation function, such as Sigmoid and ReLU. The result of the function is an estimation of the output; however, because it is an estimation, there is error. Using backpropagation, the network is able to adjust the values of the weights and bias and therefore make the estimations more accurate.

PyTorch was used in the making of this particular neural network. Its library nn utilizes a fully-connected ReLU network with one hidden layer, trained to predict y from x by minimizing squared Euclidean distance.

The network was first trained with a subset of the images (training set) before moving on to the remaining images (validation set). It then predicted the appeal scores of these images based on the participants' gender, age, and education level.

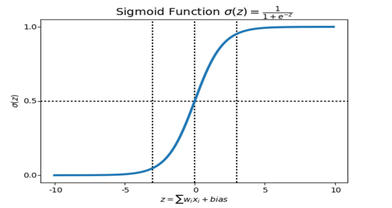

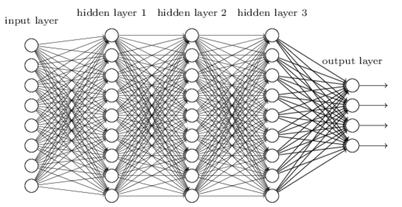

Results and Conclusion

Final Training Results

- After training, the network was able to predict 29.01% of the scores accurately. This is nearly three times better than chance and proofs the success of the developed scheme.
- The overall errors in the training results were able to slightly decrease; however, there were peaks where the error increased.

Future Work

- In order to achieve a better proportion of accurate predictions, the neural network must be able to analyze the images and recognizing trends between the images, participants and scores.
- New samples are also needed to create a more accurate neural network.