Spectral Reconstruction from LAB Data Using Neural Networks

Shankhya Debnath

September 13, 2024

1 Introduction

Spectral reconstruction refers to the process of recovering the spectral reflectance of an object from limited colorimetric data, such as CIE LAB values. This process is crucial in many applications, such as digital imaging, color reproduction, and computer vision, where accurate spectral information is required for tasks like metamerism reduction and accurate color rendering under varying illuminants.

The challenge of spectral reconstruction stems from the fact that human color vision is trichromatic, meaning that the perception of color is reduced to just three numbers (e.g., CIE XYZ or LAB), while the full spectral reflectance is typically a continuous function. This leads to a non-invertible problem since multiple spectral distributions can map to the same colorimetric values. However, under controlled conditions, machine learning techniques such as neural networks can learn an approximate mapping from LAB values to spectral reflectance, enabling spectral reconstruction.

2 Experiments

In this work, I have aimed to reconstruct the spectral reflectance of color samples from their CIE LAB values using a neural network. The Spectral Materials Database was used for extracting spectral data and their corresponding LAB values. The input to the model is a 3-dimensional vector representing the LAB values, and the output is a 39-dimensional vector representing the spectral reflectance over 39 wavelengths, ranging from 360 nm to 780 nm at 10 nm intervals.

Let:

$$\mathbf{L}_{input} = [L, a, b]$$

denote the input LAB values, and

$$\mathbf{S}_{output} = [r_{360}, r_{370}, \ldots, r_{780}]$$

represent the spectral reflectance values over the 39 wavelengths.

3 Data Preprocessing

Before feeding the data into the model, several preprocessing steps are performed. The dataset consists of color samples with corresponding spectral reflectance values and LAB values. To ensure data consistency, rows with missing LAB or spectral values are removed. The spectral data is extracted from the dataset and filtered to include only samples with exactly 39 measurements, corresponding to the 39 wavelengths.

The input LAB values are standardized using a process known as z-score normalization, defined as:

$$z = \frac{x - \mu}{\sigma}$$

where x is a feature (e.g., LAB values), μ is the mean, and σ is the standard deviation of the training set. Standardizing the input ensures that the model can converge faster by providing input features with a similar scale.

4 Neural Network Model

A feed-forward neural network is employed to approximate the mapping from LAB values to spectral reflectance. The network consists of an input layer, three hidden layers, and an output layer. Each hidden layer applies a nonlinear activation function, ReLU (Rectified Linear Unit), which is defined as:

$$f(x) = \max(0, x)$$

The ReLU activation helps the network to capture complex relationships between input and output, as it introduces non-linearity into the model.

The layers are structured as follows: - **Input layer**: Takes a 3-dimensional vector corresponding to the LAB values. - **Hidden layers**: Three fully connected layers with 64, 128, and 256 neurons, respectively. Each hidden layer applies ReLU activation. - **Output layer**: A fully connected layer with 39 neurons representing the predicted spectral reflectance values.

The neural network's operation can be represented mathematically as:

$$\mathbf{h}_1 = \text{ReLU}(\mathbf{W}_1\mathbf{L}_{input} + \mathbf{b}_1)$$

$$\mathbf{h}_2 = \text{ReLU}(\mathbf{W}_2\mathbf{h}_1 + \mathbf{b}_2)$$

$$\mathbf{h}_3 = \text{ReLU}(\mathbf{W}_3\mathbf{h}_2 + \mathbf{b}_3)$$

$$\mathbf{S}_{output} = \mathbf{W}_4\mathbf{h}_3 + \mathbf{b}_4$$

where $\mathbf{W}_i$ and $\mathbf{b}_i$ are the weights and biases for each layer i, and $\mathbf{h}_i$ represents the hidden layer activations.

5 Training the Model

The model is trained using a dataset of LAB and spectral reflectance pairs. The objective is to minimize the mean squared error (MSE) between the predicted and actual spectral reflectance values. The loss function is defined as:

$$\text{MSE} = \frac{1}{n}\sum_{i=1}^{n}(\mathbf{S}_{output,i} - \mathbf{S}_{true,i})^2$$

where $\mathbf{S}_{output,i}$ is the predicted spectral reflectance for the i-th sample, and $\mathbf{S}_{true,i}$ is the true spectral reflectance.

The model is optimized using the Adam optimizer, which is an adaptive learning rate optimization algorithm. Adam computes individual learning rates for each parameter and updates the weights by minimizing the loss function using backpropagation.

The model is trained for 100 epochs with a batch size of 16, and a validation split of 0.1 is used to evaluate the model's performance on unseen data during training. The model's performance is evaluated using several metrics, including:

- **Mean Squared Error (MSE)**
- **Mean Absolute Percentage Error (MAPE)**
- **Coefficient of Determination (R-squared)**

6 Results

After training, the model is evaluated on the test data. The spectral reflectance predicted by the model is compared with the true spectral reflectance. For visualization, we plot the actual and predicted spectral reflectance values for 20 randomly selected samples.

Each plot shows the spectral reflectance values across 39 wavelengths, where the x-axis represents the wavelength (360 nm to 780 nm), and the y-axis represents the reflectance value. The actual spectral reflectance is plotted alongside the predicted values to visually assess the model's performance.

The table below shows the performance of the model on both the training and testing datasets, as measured by three metrics: Mean Squared Error (MSE), Mean Absolute Percentage Error (MAPE), and the Coefficient of Determination (R-squared).

Metric	Training	Testing
MSE	17.2498	21.6648
MAPE	1.0702	0.6152
R-squared	0.9465	0.9282

Table 1: Performance Metrics for Training and Testing Data

7 Observations

The results of the model, as shown in the table, indicate that the neural network performs reasonably well in reconstructing spectral data from LAB values. The following inferences can be drawn from the performance metrics:

- **Mean Squared Error (MSE)**: The training MSE is 17.25, and the testing MSE is 21.66. The relatively low values of MSE suggest that the model is able to approximate the relationship between the LAB inputs and the spectral reflectance outputs effectively. However, the slightly higher MSE on the testing data suggests some degree of overfitting, meaning that the model might generalize slightly less effectively to unseen data.
- **Mean Absolute Percentage Error (MAPE)**: The MAPE for the training set is 1.07%, while for the test set it is 0.62%. These low values indicate that, on average, the model's predictions are very close to the true spectral reflectance values, with errors less than 1.1% across both sets. This shows the model's capability to maintain accuracy in its predictions across various samples.
- **R-squared ($\mathbf{R^2}$)**: The R-squared values for training and testing are 0.9465 and 0.9282, respectively. These values are close to 1, implying that the model explains over 92% of the variance in the spectral data. This suggests a strong fit of the model to the data, confirming that the neural network has captured the essential relationships needed to predict spectral reflectance from LAB values.
- **Generalization**: While the testing MSE is higher than the training MSE, the relatively small difference between the two suggests that the model generalizes fairly well to unseen data. The close values of the training and testing $\mathbf{R^2}$ scores reinforce this inference, indicating that the model has avoided significant overfitting.

Taken together, the three metrics (MSE, MAPE, and $\mathbf{R^2}$) indicate that the model is capable of predicting spectral reflectance with a high degree of accuracy and minimal error, even on test data. The results demonstrate the effectiveness of using a neural network for this spectral reconstruction task.

8 References

1. J. Alstan Jakubiec (2022). Data-driven selection of typical opaque material reflectances for lighting simulation, LEUKOS, DOI: 10.1080/15502724.2022.2100788.

2. J. Alstan Jakubiec (2016). Building a database of opaque materials for lighting simulation. In PLEA 2016–Cities, Buildings, People: Towards Regenerative Environments, Proceedings of the 32nd International Conference on Passive and Low Energy Architecture.