

Dataset

The CIFAR-10 dataset consists of 60000 32x32 colour images in 10 classes, with 6000 images per class. There are 50000 training images and 10000 test images. The dataset is divided into five training batches and one test batch, each with 10000 images. The test batch contains exactly 1000 randomly-selected images from each class. Here are the classes in the dataset, as well as 10 random images from each:




Figure 1: Examples of CIFAR-10 dataset

In our project, for training set, we use 40000 images randomly mixed with training images and test images in the CIFAR-10 dataset; for test set, we use 10000 of the remaining images adding random noise chosen from $\{-1,0,1\}$, together with 20000 “fake” images to prevent certain form of cheating. The test score is based on the prediction of the “real” images.

Classifiers

Each student implemented one classification method of their choice. As a summary, one student chooses SVM, two students choose the decision tree, and the remaining students all choose neural networks of certain form (MLP, CNN, etc.). The list of classifiers is below:

- ❑ Decision tree (including random forest and other methods based on decision tree)
- ❑ Multilayer Neural Network
- ❑ K-nearest neighbors
- ❑ Kernel classification
- ❑ Support Vector Machine
- ❑ Linear or Logistic regression, with any regularization

Methods

We also consider the size of the training sets used and the training error for students using NN. For the training size, two students use half of the training data for their predictor, and two students use data augmentation to create more training images, while other students use the entire dataset with cross-validation. For training error, three students are in the interpolation regime (i.e. training error ~ 0) while others are in the classical regime.

Results

Below is a plot of the error rates on the test set for the classifiers implemented. The average test error is about 37.5%, with a standard deviation of 22.2%. Note that in the plot, “Actual Error” is the L_{01} loss on the test set, and “Predicted Error” is each student’s error estimate obtained from the training data.

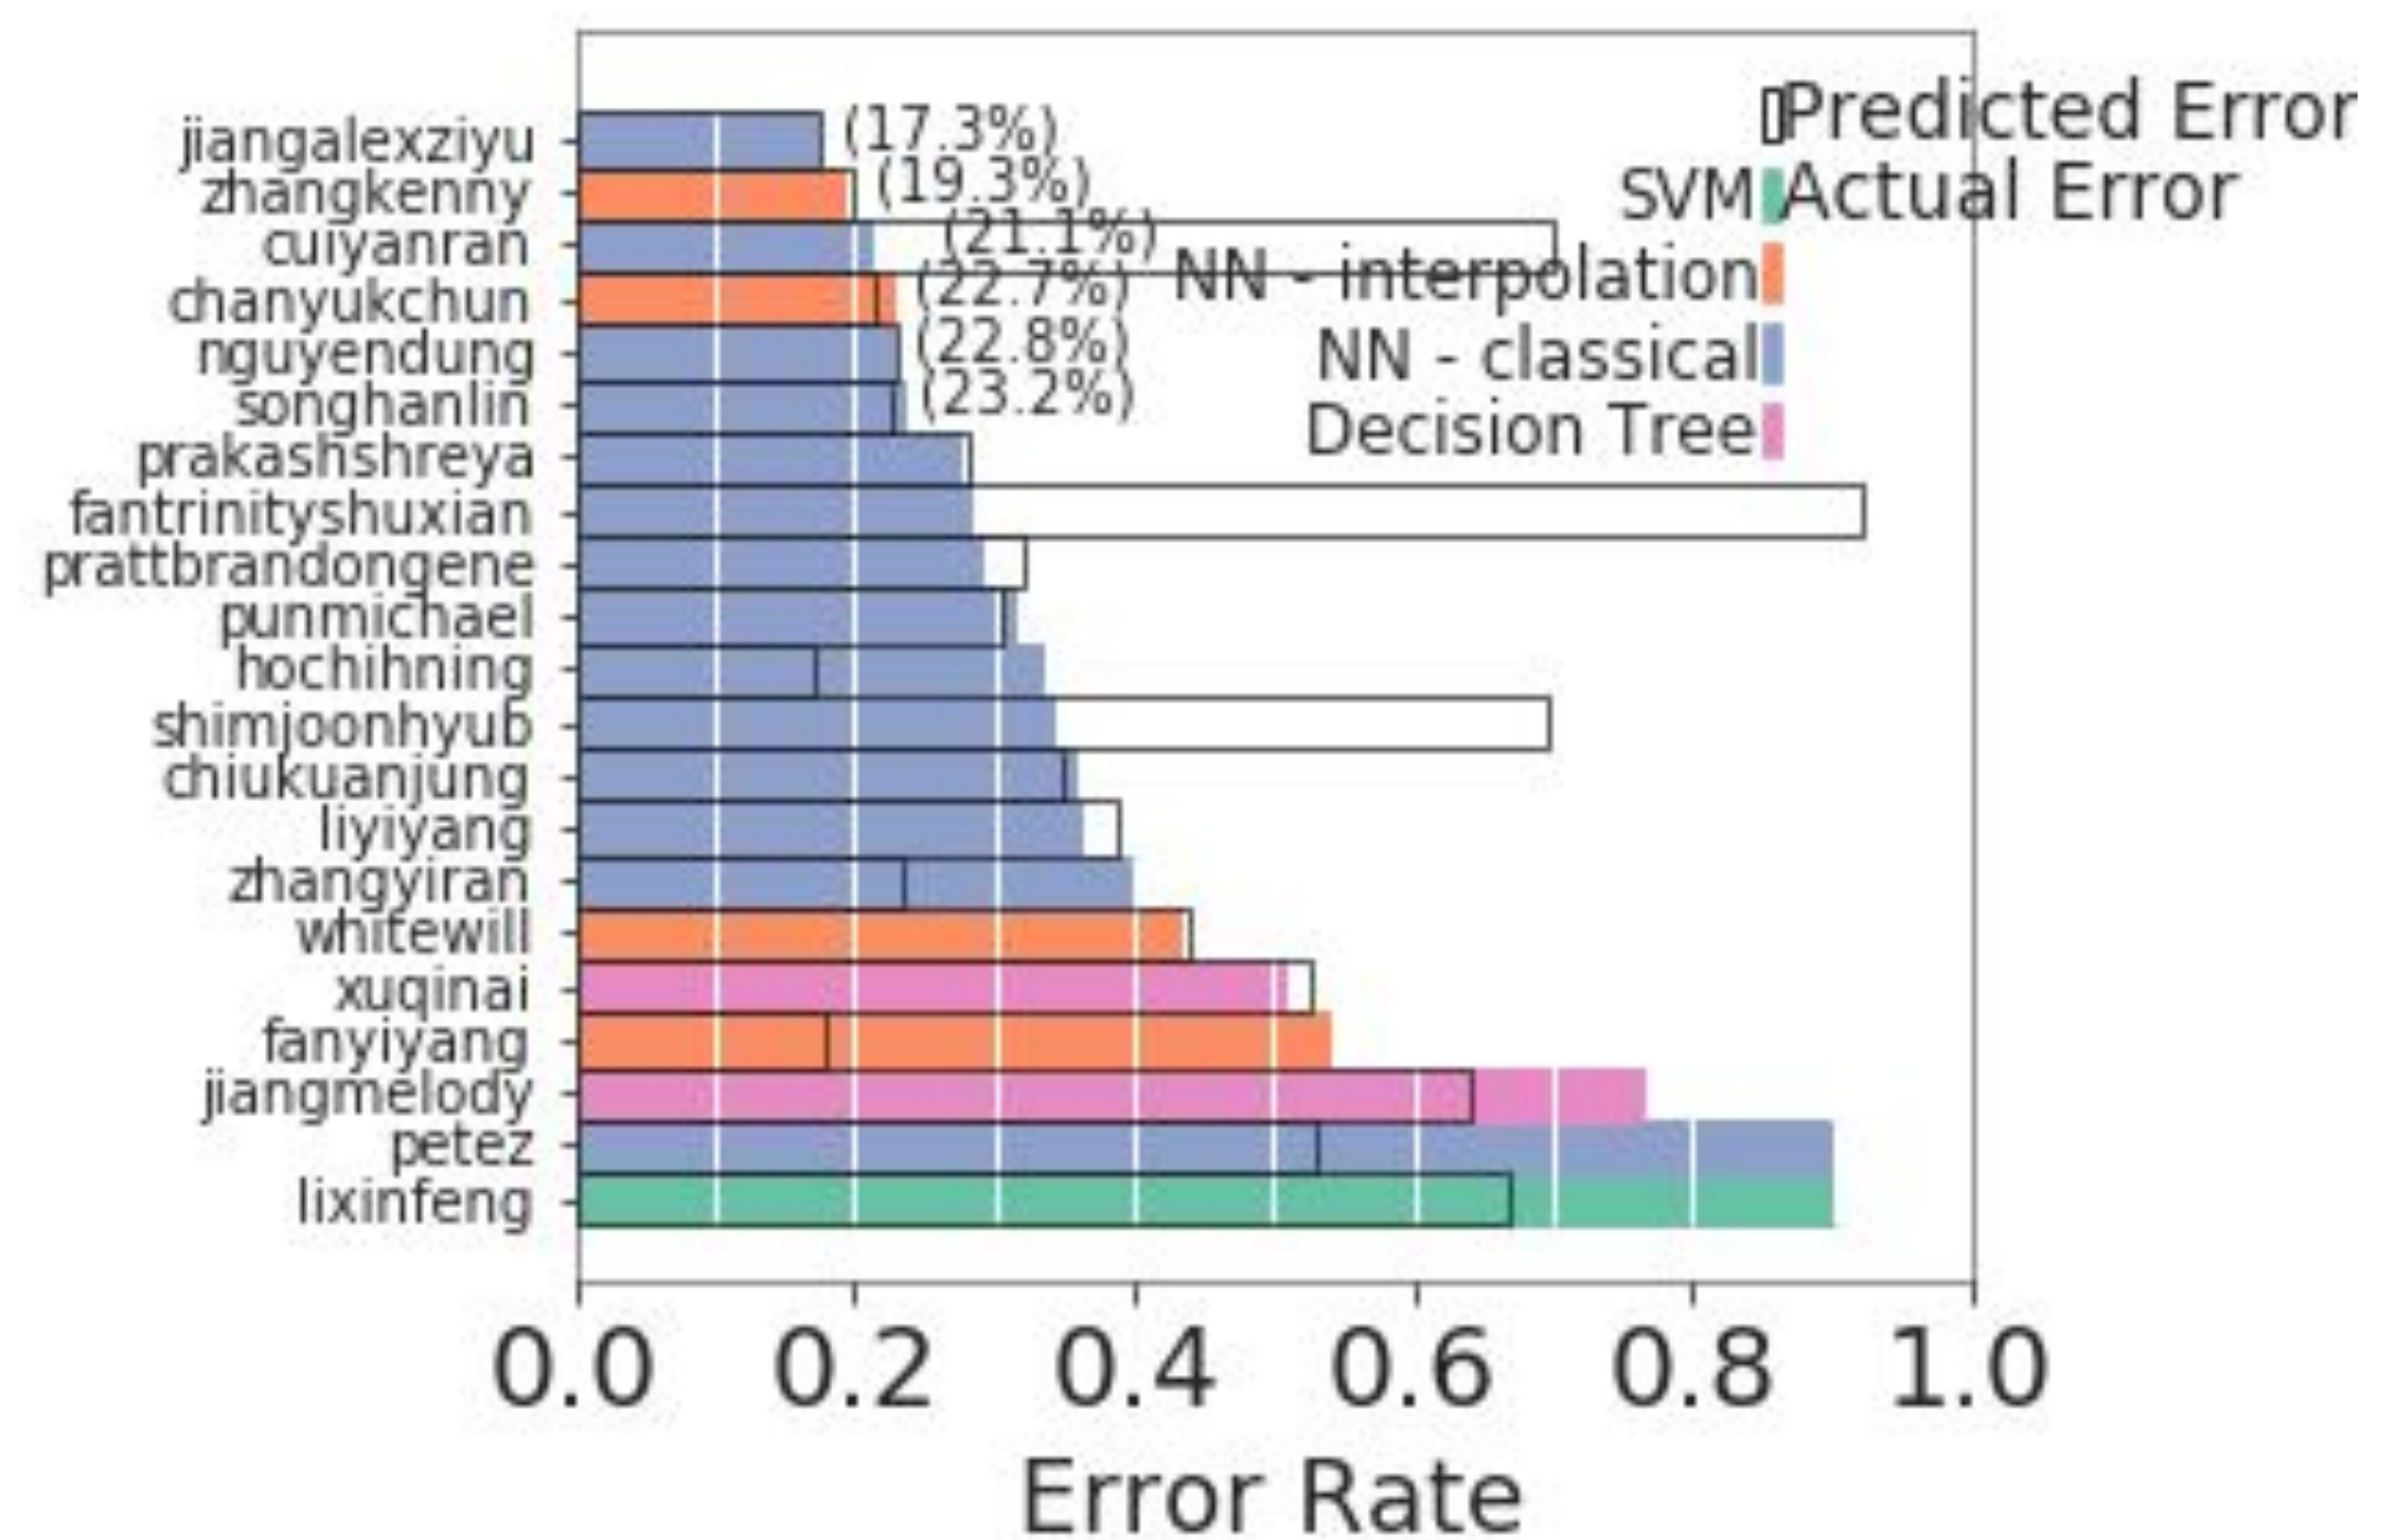


Figure 2: Test results

Awards

- ❑ Most accurate neural network classifier: Alex Ziyu Jiang
- ❑ Most accurate non-NN classifier: Qinai Xu
- ❑ Most “optimistic” estimate: Yiyang Fan
- ❑ Most “pessimistic” estimate: Trinity Shuxian Fan
- ❑ Most “efficient” classifier in terms of training size: Kenny Zhang