# Human Activity Classification

Aristos Athens, Navjot Singh, Zachary Blum
aristos@stanford.edu, navjot@stanford.edu, zblum25@stanford.edu

## Executive Summary

Activity recognition is an important task in several healthcare and consumer product applications. With the recent proliferation of human sensory data, we wanted to see how well a machine learning model can classify human activities using heart rate data and IMU readings from a user's hand, chest, and ankle. Since hand IMU and heart rate data is more ubiquitous and accessible than chest and ankle data through the use of smart watches, we also decided to compare our logistic regression, SVM, decision tree, and neural net models between the full dataset and a limited dataset with only hand IMU and heart rate data. With data from the full dataset, all models were all able to perform with high accuracy on all recorded activities. Using the limited dataset, the models performed only slightly worse than the full dataset.

## Data and Features

- PAMAP2 Dataset from the UCI machine learning repository [1].
- Data features include 9-axis IMU data streams for sensors on each of hand, chest, and ankle and subject heart rate.
- 1.9 million data points of 52 features each, spread over nine subjects.
- 18 different activity IDs, including sitting, walking, running, folding laundry, and cycling.
- For comparison, the "limited" dataset contains only hand IMU and heart rate data.

## References

[1] "PAMAP2 Physical Activity Monitoring Dataset." (2012, August). Retrieved from http://archive.ics.uci.edu/ml/datasets/pamap2+physical+activity+monitoring
[2] Mark Schmidt, Nicolas Le Roux, Francis Bach. Minimizing Finite Sums with the Stochastic Average Gradient. Mathematical Programming B, Springer, 2017
[3] "Lecture 19: Decision Trees." (2017, Nov. 7). Retrieved from https://web.stanford.edu/class/stats202/content/lec19.pdf

## Models

**Preprocessing:**
Both logistic regression and SVM involved a preprocessing step of subtracting the mean and dividing by the standard deviation of the training set from each data point.

$$z = \frac{(x-u)}{s} \qquad u = \textit{training mean} \qquad s = \textit{training standard deviation}$$

### Logistic Regression

L2 regularization was employed along with Stochastic Average Gradient Descent [2] as a solver because of its scalability in speed for large datasets.

### Support Vector Machine

Radial Basis Function Kernel was used because of its ability to generate non-linear boundaries. The rbf kernel also performed better than the linear and polynomial kernels.

### Decision Trees

**Gini Loss:** $\sum_{m=1}^{|T|} q_m \sum_{k=1}^{K} p_{mk}(1-p_{mk})$ where $p_{mk}$ is the proportion of examples in class k present in region $R_m$, and $q_m$ is the proportion of examples in $R_m$ from tree T with |T| different $R_m$ regions [3]

***Boosting***: AdaBoost--ensembles 500 weak decision trees.
***Random Forest***: Ensembles 100 trees, where only the square root of the total number of features are considered at each split.

### Deep Learning

**MultiLayer Perceptron**: Neural network architecture consisting of two hidden layers of size (n, 512) and (512, 512), input layer with size (n, 1), output one-hot encoded to size (k, 1). We use ReLU activation and categorical cross entropy loss.

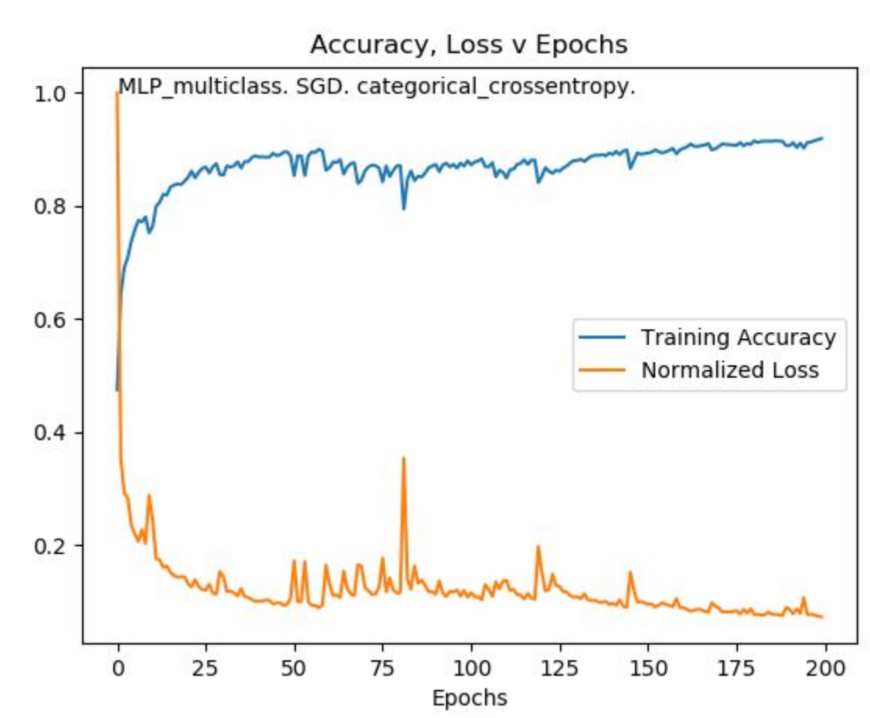


## Results

5-fold cross-validation was conducted on different combinations of data features. The best performing features (full three-IMU + heart rate, and hand-IMU + heart rate) are shown below.

| Model | Dataset | Training Accuracy (%) | Test Accuracy (%) |
|---|---|---|---|
| Logistic Regression | Hand IMU + HR | 64.26 | 64.25 |
| | Three IMUs + HR | 82.06 | 82.06 |
| SVM | Hand IMU + HR | 94.04 | 91.94 |
| | Three IMUs + HR | 99.83 | 99.09 |
| Decision Trees | Hand IMU + HR | 99.65 | 87.32 |
| | Three IMUs + HR | 99.77 | 91.93 |
| Boosted Decision Trees | Hand IMU + HR | 99.93 | 93.43 |
| | Three IMUs + HR | 100.00 | 98.71 |
| Random Forest | Hand IMU + HR | 100.00 | 93.32 |
| | Three IMUs + HR | 100.00 | 97.82 |
| Neural Net | Hand IMU + HR | 84.78 | 81.4 |
| | Three IMUs + HR | 93.33 | 95.63 |

## Discussion & Future Work

- The dataset provides input features that likely would not be present in real-world applications, like chest and ankle IMUs. We found we could get relatively good performance using just hand IMU and heart rate, the type of data one might get from a smart watch.
- Logistic regression unsurprisingly performed the worst as it is a linear classifier.
- As expected, ensembling (random forest and boosting) improved test accuracy over the original decision trees.
- The neural net consistently provided high accuracies at the cost of long train times and relatively slow classification. In the future we would try using RNNs to classify more complex tasks that depend on sequential lower level actions.
- In the future we would like to test these models using real IMU's. In particular, we would want to see if a low-compute embedded device could perform classifications with NN's or SVM's in real time, in addition to computationally cheaper decision trees.