



Twitter Sentiment Detection

Researcher: Loyiso Jiya

Co-researcher: Daiyaan Sataar

Supervisor: M Ghaziasgar

Co-supervisor: R Dodds

→ Recap

→ Preliminary Testing

→ Testing and
Optimization

→ Results

→ Demo

Overview

- Recap
- Preliminary Testing
- Testing and Optimization
- Results
- Demonstration

→ Recap

→ Preliminary Testing

→ Testing and Optimization

→ Results

→ Demo

Recap: Background

- **What is sentiment detection?**
 - Attempts to determine the feeling conveyed by any text
 - Can be positive, neutral, negative, or range
 - Text can be documents, forum posts or posts on social media
 - AKA sentiment analysis, opinion mining, opinion extraction
- **Why sentiment detection?**
 - Companies need feedback
 - conducted market research, focus groups, etc.
 - Inefficiency, time-consuming and costly
 - reviews after service is issued (i.e. MTN customer service)
 - reviews when given are often bias (mostly negative)

→ Recap

→ Preliminary Testing

→ Testing and Optimization

→ Results

→ Demo

Recap: Progress

- Work Done:
 - Literature review
 - Data pre-processing
 - Implementation: Initial Training and testing
 - Exhaustive testing and optimization
- Scope:
 - Implement ResNet and XGBoost for performance comparison in sentiment detection.

→ Recap

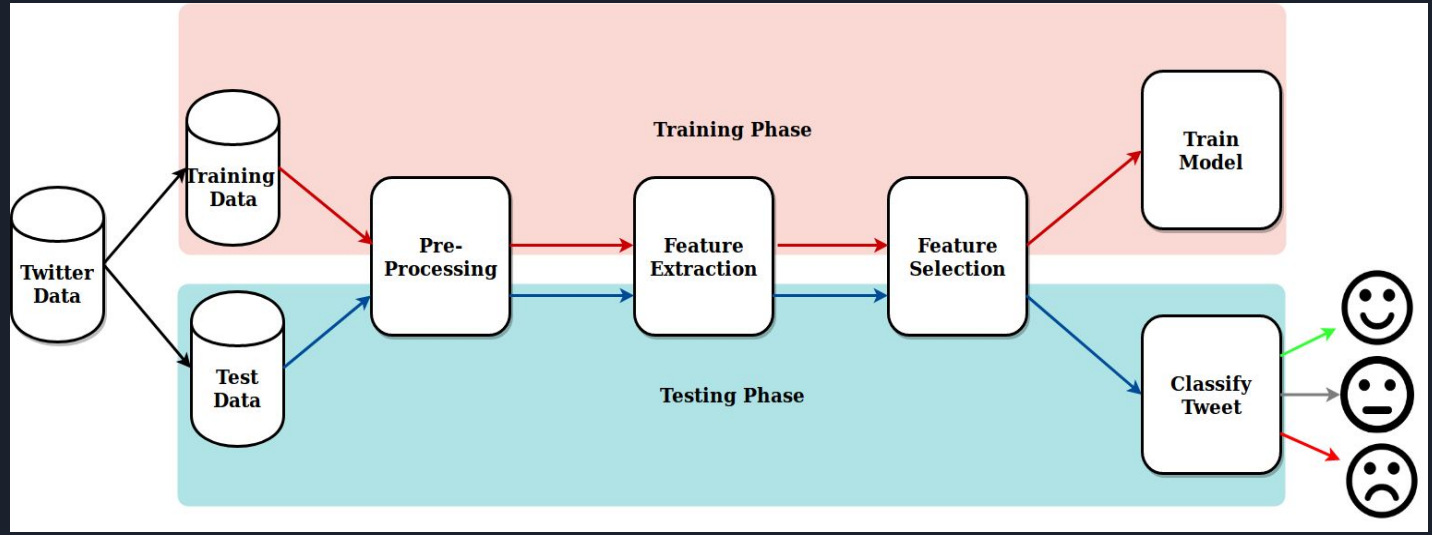
→ Preliminary Testing

→ Testing and Optimization

→ Results

→ Demo

Recap: Proposed Architecture





Recap: Implementation

→ Recap

→ Preliminary Testing



→ Testing and Optimization



→ Results



→ Demo

→ Recap

→ Preliminary Testing

→ Testing and Optimization

→ Results

→ Demo

Preliminary Testing

- Testing Environment
 - Windows/Linux Operating System
 - 16GB RAM
 - NVIDIA GeForce GTX 1060 6GB memory
 - NVIDIA CUDA drivers
 - Jupyter notebook with python3 kernel
- Data sets
 - Stanford Twitter Corpus
 - 250 000 samples used
 - 80% training and 20% for Testing

→ Recap

→ Preliminary Testing

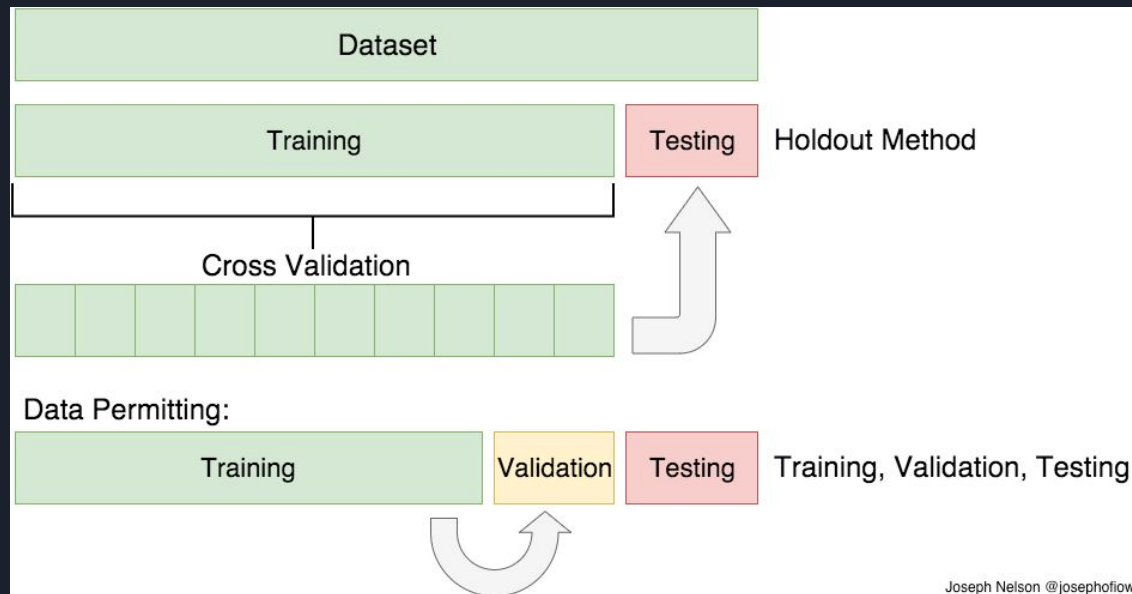
→ Testing and Optimization

→ Results

→ Demo

Preliminary Testing

- Testing For machine learning



→ Recap

→ Preliminary Testing


→ Testing and Optimization

→ Results

→ Demo

Testing and Optimization

- Cross validation:



→ Recap

→ Preliminary Testing

→ Testing and Optimization

→ Results

→ Demo

Testing and Optimization

- Grid search:

Parameters	Search Space
max_depth	[10, 20, 30, 40, 50, 60, 70, 80, 90, 100]
n_estimators	[20, 40, 60, 80, 100, 120, 140, 160, 180, 200]
learning_rate	[0.1, 0.01, 0.001, 0.03, 0.04]

→ Recap

→ Preliminary Testing

→ Testing and Optimization

→ Results

→ Demo

Results

TABLE I: Baseline Tests for Desired Models

Prediction	F1-score (%)		
	XGBoost	ResNet(Adam Optimizer)	ResNet-10(Gradient Descent Optimizer)
negative	45	67	0
positive	70	72	67
overall	70	72	67

TABLE II: Improved results Desired Models

Prediction	F1-score (%)		
	XGBoost	ResNet(Adam Optimizer)	ResNet-10(Gradient Descent Optimizer)
negative	78	71	74
positive	73	80	76
overall	70	75	75

→ Recap

→ Preliminary Testing

→ Testing and
Optimization

→ Results

→ Demo

Demo



Thank you.

Any Questions?