

# Homework 3

## Text classification using LSTM and Logistic Regression on AG News Dataset

**Name:** Zinia Sultana Joti

**NJIT ID:** 31730019

## Introduction

In this report, we will perform a text classification task using classical machine learning and deep neural network algorithm on a publicly available data: AG News dataset. The goal is to

- classify news articles into predefined categories.
- Apply two different algorithms:
  - classical ml model - Logistic Regression (LR) with TF-IDF features
  - deep learning model- Long Short-Term Memory (LSTM).
- Compare their performance.
- Understand differences in their modeling approaches.

## Dataset

The AG's corpus contains 496,835 categorized news articles from more than 2000 news sources. The AG's news topic classification dataset is used as a text classification benchmark. It contains 4 largest classes from the original corpus, each containing 30,000 training samples and 1,900 test samples. The total number of training samples is 120,000 and testing samples is 7,600. The file contains a list of classes with corresponding labels. Dataset was split into training, validation and test sets.

The classes in the dataset are World, Sports, Business, Sci/Tech.

```
[zj98@login02 Code]$ python prepare_data.py
{'text': "Wall St. Bears Claw Back Into the Black (Reuters) Reuters - Short-sellers,
Wall Street's dwindling\band of ultra-cynics, are seeing green again.", 'label': 2
}
Train size: 120000
Test size: 7600
Saved train/val/test in data folder.
Train: 108000
Validation: 12000
Test: 7600
Label distribution:
Business: 30000 samples
Sci/Tech: 30000 samples
Sports: 30000 samples
World: 30000 samples
```

Figure 1: Dataset information

## Preprocessing

### Logistic regression

Logistic regression is a traditional machine learning algorithm for text classification. It uses feature extraction techniques.

Bag-of-words TFIDF: BoG is constructed by using most frequent words from training subset. Here, the count of each word is used as the features, is also defined as ‘term-frequency’. The inverse document frequency (IDF) is computed by taking the logarithm of the ratio of total number of documents and number of documents with the word in the training subset. The features are normalized by the largest feature value.

The following preprocessing steps are applied:

1. TF-IDF vectorization
2. Lowercasing
3. Stop words removal (using default “English” stop words)
4. Feature limit to 20,000 terms

### LSTM

LSTM is designed to capture sequential dependencies in text. It processes text as ordered sequences rather than independent features.

The pipeline for LSTM focuses on preparing text sequences:

1. Tokenization: converting text into words/subwords
2. Vocabulary building based on training data
3. Padding sequences to fixed length for batch processing
4. No stop word removal required to preserve sentence structure and context for sequence modeling.

## Models

### Logistic regression

Model: logistic regression

Input: TF-IDF features

Activation function: Sigmoid for binary, Softmax for multi-class classification

The model learns a linear decision boundary that separates classes based on importance of input features.

## LSTM

Long Short-Term Memory is a version of recurrent neural network which addresses the vanishing gradient problem faced in RNN. It uses gates activation layers) to forget irrelevant data and remember past knowledge that the network has seen so far. A LSTM consists of three different gates for different purposes.

1. Forget Gate
2. Input Gate
3. Output Gate

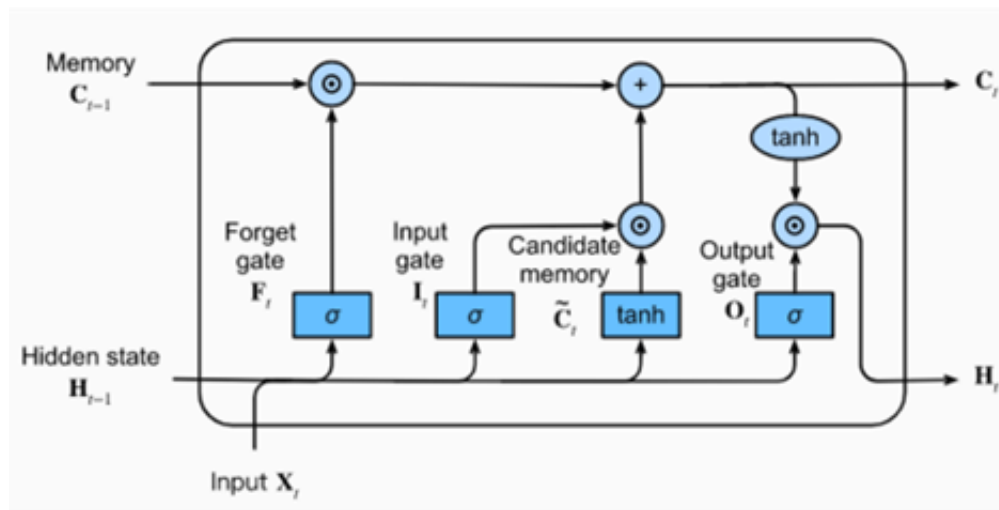


Figure 2: LSTM architecture

## Training setup

Batch size: 32

Embedding dimensions: 100

Hidden dimensions: 128

Learning rate: 0.001

Epochs: 10

Loss function: Cross Entropy Loss

Optimizer: Adam

Maximum sequence length: 100

Minimum word frequency: 2

Vocab size: 52162

Number of classes: 4

## Evaluation Metrics

There are few metrics to evaluate a model's performance. We will use accuracy, precision, recall, F1 score. These metrics are derived from the confusion matrix.

### Confusion Matrix

	Actually positive	Actually negative
Predicted positive	True positives (TP)	False positives (FP)
Predicted negative	False Negatives (FN)	True Negatives (TN)

$$Accuracy = \frac{TP + TN}{TP + FP + FN + TN}$$

$$Precision, P = \frac{TP}{TP + FP}$$

$$Recall, R = \frac{TP}{TP + FN}$$

$$F1 = \frac{2PR}{P + R}$$

## Results

Both Logistic Regression (LR) and LSTM models achieved strong performance on the AG News classification task, with test accuracy more than 90%. However, LR performed slightly better than LSTM across all the evaluation metrics.

Table 1 below illustrates the evaluation metrics where logistic regression consistently shows better performance than LSTM.

Table 1: Evaluation metrics for logistic regression and LSTM model

Model	Loss	Accuracy	Precision	Recall	F1 score
Logistic Regression	0.285912354	0.915263158	0.915083334	0.915263158	0.915071952
LSTM	0.353681161	0.908289474	0.908631597	0.908289474	0.908389392

The performance differences are quite small, Figure 3 demonstrates the variation in values of the metrics.

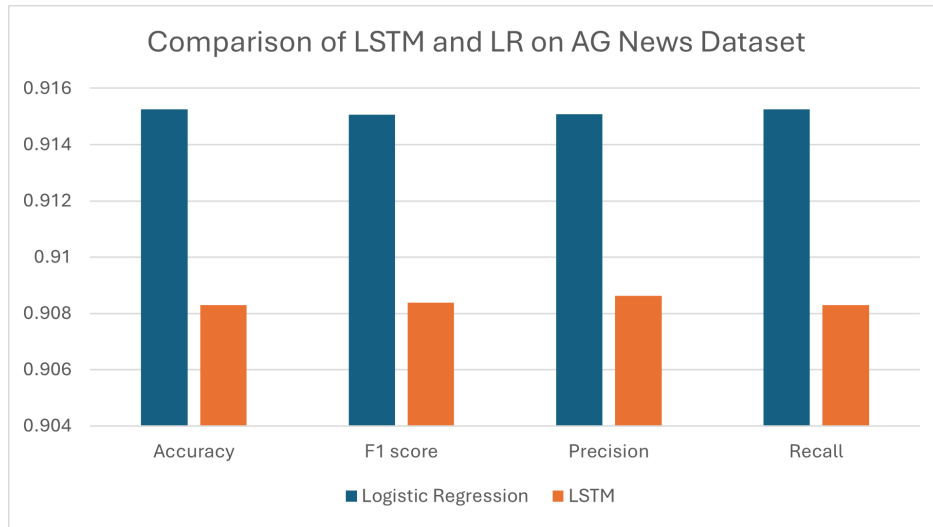


Figure 3: Final comparison of all metrics between Logistic regression and LSTM

### Logistic regression output:

Logistic Regression shows strong and stable performance on both validation and test sets, indicating good generalization.

Table 2: Performance of Logistic Regression on validation set and test set

Dataset	Accuracy	Precision	Recall	F1 score
Validation set	0.9215	0.921276	0.9215	0.921301
Test set	0.915263158	0.915083334	0.915263158	0.915071952

### LSTM training curves:

Performance of LSTM training improved rapidly after first few epochs.

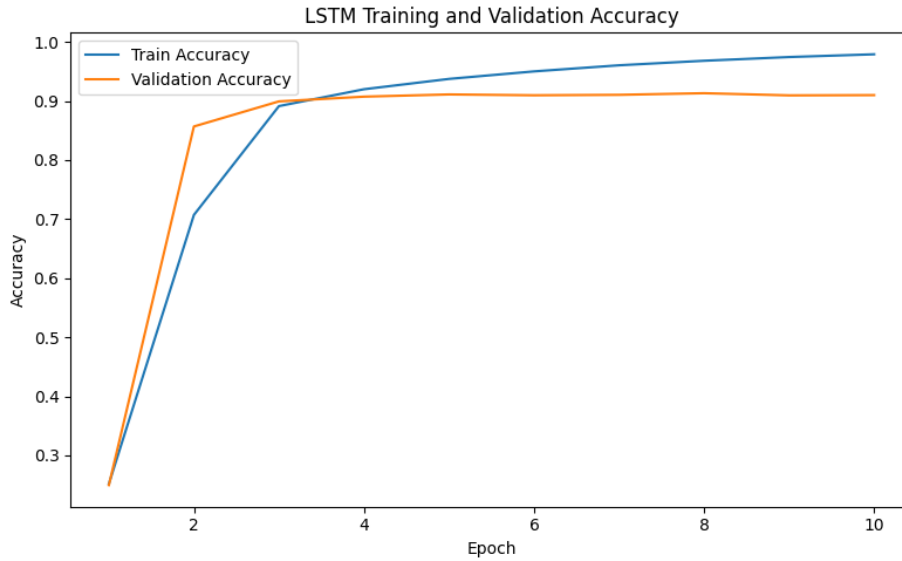


Figure 4: Training and Validation Accuracy of LSTM

Both training and validation accuracy increased significantly, showing fast learning. The validation accuracy stabilizes around 91%, while the training accuracy continues to increase. This indicates mild overfitting on training data.




Figure 5: Training and Validation F1 Score of LSTM

Figure 5 shows a similar trend for F1 score. Validation performance stabilizing after around epoch 3–5.




Figure 6: Training and Validation loss of LSTM

The training loss decreases steadily across epochs, indicating models are learning effectively. The validation loss initially decreases but after about 4–5 epochs, it begins to increase. The model starts to overfit in later epochs, as a result it performs worse on unseen data.

## Discussion

Logistic Regression combined with TF-IDF features, treats text as bag-of-words representation and based on word frequency it identifies the importance of words. In the AG News structured dataset, articles are categorized by topic-specific keywords, this approach is highly effective and computationally efficient.

On the contrary, LSTM is a sequence-based deep learning model that captures the word order and context in a sentence. LSTM can learn complex linguistic patterns; however, it requires more data and longer training to perform effectively. The number of epochs for this project was only 10 to reduce computational costs. Additionally, LSTM models can improve by tuning hyperparameters such as embedding size, hidden dimensions and sequence length. Moreover, LR ran significantly faster (only 15s), making it more practical choice for this task. By contrast, LSTM required 31 minutes as it involves iterative training and higher computational costs.

Overall, the results demonstrate that traditional machine learning model (in this case, Logistic regression), can remain highly competitive for text classification tasks. While LSTM provides more expressive frameworks.

## Conclusion

In brief, both models showed remarkable accuracy on the AG News dataset, while Logistic Regression outperformed LSTM in all metrics. We can imply that simple models with TF-IDF features can be highly effective on structural datasets and on text classification tasks. On the other hand, LSTM can be used where contextual data is required, however, it may require fine tuning its hyperparameters and large dataset.

## Appendix: Code

### 1. prepare\_data.py

```
1 from datasets import load_dataset
2 from collections import Counter
3 from sklearn.model_selection import train_test_split
4 import pandas as pd
5 import os
6
7 os.system("cls")
8 # load dataset
9 dataset = load_dataset("ag_news")
10
11 train_data = dataset["train"]
12 test_data = dataset["test"]
13
14 print(train_data[0])
15 print("Train size:", len(train_data))
16 print("Test size:", len(test_data))
17
18 train_df = pd.DataFrame(train_data)
19 test_df = pd.DataFrame(test_data)
20
21 # split train data into train and validation sets
22 train_df, val_df = train_test_split(
23     train_df, test_size=0.1,
24     random_state=42,
25     stratify=train_df["label"]
26 )
27
28 os.makedirs("data", exist_ok=True)
29
30 train_df.to_csv("data/train.csv", index=False)
31 val_df.to_csv("data/val.csv", index=False)
32 test_df.to_csv("data/test.csv", index=False)
33
34 print("Saved train/val/test in data folder.")
```

```

35 print("Train:", len(train_df))
36 print("Validation:", len(val_df))
37 print("Test:", len(test_df))
38
39 label_counter = Counter(train_data["label"])
40 # print("Label counts:", label_counter)
41 label_names = {
42     0: "World",
43     1: "Sports",
44     2: "Business",
45     3: "Sci/Tech"
46 }
47 print("Label distribution:")
48 for label, count in label_counter.items():
49     print(f"{label_names[label]}: {count} samples")

```

Listing 1: prepare\_data.py

## 2. data\_helper.py

```

1 import re
2 from collections import Counter
3 import pandas as pd
4 import torch
5 from torch.utils.data import Dataset, DataLoader
6
7 def load_data(train_path, val_path, test_path):
8     train_df = pd.read_csv(train_path)
9     val_df = pd.read_csv(val_path)
10    test_df = pd.read_csv(test_path)
11
12    return train_df, val_df, test_df
13
14 def clean_text(text):
15     text = str(text).lower().strip()
16     text = re.sub(r'<.*?>', '', text) # Remove HTML tags
17     text = re.sub(r'[^a-zA-Z0-9\s]', '', text) # Remove non-
18     alphanumeric characters (except spaces)
19     return text

```

```

19
20 def tokenize(text):
21     return clean_text(text).split()
22
23 def build_vocab(texts, min_freq=2):
24     counter = Counter()
25     for text in texts:
26         tokens = tokenize(text)
27         counter.update(tokens)
28
29     vocab = {
30         "<PAD>": 0,
31         "<UNK>": 1
32     }
33
34     for token, freq in counter.items():
35         if freq >= min_freq:
36             vocab[token] = len(vocab)
37
38     return vocab
39
40 def text_to_sequence(text, vocab):
41     tokens = tokenize(text)
42     return [vocab.get(token, vocab["<UNK>"]) for token in tokens]
43
44 def pad_sequence(seq, max_len, pad_value=0):
45     if len(seq) < max_len:
46         seq = seq + [pad_value] * (max_len - len(seq))
47     else:
48         seq = seq[:max_len]
49     return seq
50
51 class AGNewsDataset(Dataset):
52     def __init__(self, dataframe, vocab, max_len=100):
53         self.texts = dataframe["text"].tolist()
54         self.labels = dataframe["label"].tolist()
55         self.vocab = vocab
56         self.max_len = max_len

```

```

57
58     def __len__(self):
59         return len(self.texts)
60
61     def __getitem__(self, idx):
62         text = self.texts[idx]
63         label = self.labels[idx]
64
65         seq = text_to_sequence(text, self.vocab)
66         seq = pad_sequence(seq, self.max_len)
67
68         return torch.tensor(seq), torch.tensor(label)
69
70 def create_dataloader(train_df, val_df, test_df, vocab, batch_size
71 =32, max_len=100):
72     train_dataset = AGNewsDataset(train_df, vocab, max_len)
73     val_dataset = AGNewsDataset(val_df, vocab, max_len)
74     test_dataset = AGNewsDataset(test_df, vocab, max_len)
75
76     train_loader = DataLoader(train_dataset, batch_size=batch_size,
77                               shuffle=True)
78     val_loader = DataLoader(val_dataset, batch_size=batch_size,
79                             shuffle=False)
80     test_loader = DataLoader(test_dataset, batch_size=batch_size,
81                              shuffle=False)
82
83     return train_loader, val_loader, test_loader

```

Listing 2: data\_helper.py

### 3. lstm.py

```

1 import torch
2 import torch.nn as nn
3
4 class LSTMClassifier(nn.Module):
5     def __init__(self, vocab_size, embed_dim, hidden_dim, output_dim
6                 , pad_idx, dropout=0.5):
7         super().__init__()

```

```

7
8     self.embedding = nn.Embedding(
9         num_embeddings=vocab_size,
10        embedding_dim=embed_dim,
11        padding_idx=pad_idx
12    )
13
14    self.lstm = nn.LSTM(
15        input_size=embed_dim,
16        hidden_size=hidden_dim,
17        batch_first=True
18    )
19
20    self.fc = nn.Linear(hidden_dim, output_dim)
21    self.dropout = nn.Dropout(dropout)
22
23    def forward(self, text):
24        x = self.embedding(text)
25        _, (hidden, _) = self.lstm(x)
26        hidden = hidden[-1]
27        hidden = self.dropout(hidden)
28        output = self.fc(hidden)
29    return output

```

Listing 3: lstm.py

#### 4. train\_lstm.py

```

1 import argparse
2 import torch
3 import torch.nn as nn
4 from sklearn.metrics import accuracy_score,
5     precision_recall_fscore_support
6 import pandas as pd
7 import os
8
9 from data_helper import load_data, build_vocab, create_dataloader
10 from lstm import LSTMClassifier

```

```

11 def train_one_epoch(model, dataloader, optimizer, criterion, device)
    :
12     model.train()
13     total_loss = 0
14     all_preds = []
15     all_labels = []
16
17     for texts, labels in dataloader:
18         texts, labels = texts.to(device), labels.to(device)
19
20         optimizer.zero_grad()
21         outputs = model(texts)
22         loss = criterion(outputs, labels)
23         loss.backward()
24         optimizer.step()
25
26         total_loss += loss.item()
27         preds = torch.argmax(outputs, dim=1).cpu().numpy()
28         all_preds.extend(preds)
29         all_labels.extend(labels.cpu().numpy())
30
31     avg_loss = total_loss / len(dataloader)
32     acc = accuracy_score(all_labels, all_preds)
33     precision, recall, f1, _ = precision_recall_fscore_support(
34         all_labels, all_preds, average="weighted", zero_division=0
35     )
36
37     return avg_loss, acc, precision, recall, f1
38
39 def evaluate(model, dataloader, criterion, device):
40     model.eval()
41     total_loss = 0
42     all_preds = []
43     all_labels = []
44
45     with torch.no_grad():
46         for texts, labels in dataloader:
47             texts, labels = texts.to(device), labels.to(device)

```

```

48         outputs = model(texts)
49         loss = criterion(outputs, labels)
50         total_loss += loss.item()
51
52         preds = torch.argmax(outputs, dim=1).cpu().numpy()
53         all_preds.extend(preds)
54         all_labels.extend(labels.cpu().numpy())
55     avg_loss = total_loss / len(dataloader)
56     acc = accuracy_score(all_labels, all_preds)
57     precision, recall, f1, _ = precision_recall_fscore_support(
58         all_labels, all_preds, average="weighted", zero_division=0
59     )
60     return avg_loss, acc, precision, recall, f1
61
62 def main():
63     parser = argparse.ArgumentParser()
64     parser.add_argument("--batch_size", type=int, default=32)
65     parser.add_argument("--embed_dim", type=int, default=100)
66     parser.add_argument("--hidden_dim", type=int, default=128)
67     parser.add_argument("--epochs", type=int, default=10)
68     parser.add_argument("--lr", type=float, default=0.001)
69     parser.add_argument("--max_len", type=int, default=100)
70     parser.add_argument("--min_freq", type=int, default=2)
71     args = parser.parse_args()
72
73     device = torch.device("cuda" if torch.cuda.is_available() else "
74         cpu")
75     print(f"Using device: {device}")
76
77     train_df, val_df, test_df = load_data(
78         "data/train.csv",
79         "data/val.csv",
80         "data/test.csv"
81     )
82     vocab = build_vocab(train_df["text"], min_freq=args.min_freq)
83     pad_idx = vocab["<PAD>"]
84     vocab_size = len(vocab)
85     num_classes = len(train_df["label"].unique())

```

```

85
86     print("Train size:", len(train_df))
87     print("Validation size:", len(val_df))
88     print("Test size:", len(test_df))
89     print("Vocab size:", vocab_size)
90     print("Number of classes:", num_classes)
91
92     train_loader, val_loader, test_loader = create_dataloader(
93         train_df, val_df, test_df,
94         vocab=vocab,
95         batch_size=args.batch_size,
96         max_len=args.max_len
97     )
98
99     model = LSTMClassifier(
100         vocab_size=vocab_size,
101         embed_dim=args.embed_dim,
102         hidden_dim=args.hidden_dim,
103         output_dim=num_classes,
104         pad_idx=pad_idx,
105         dropout=0.5
106     )
107     model = model.to(device)
108
109     criterion = nn.CrossEntropyLoss()
110     optimizer = torch.optim.Adam(model.parameters(), lr=args.lr)
111
112     best_val_f1 = 0.0
113     best_model_path = f"best_lstm.pt"
114
115     train_losses = []
116     val_losses = []
117     train_accs = []
118     val_accs = []
119     train_precs = []
120     val_precs = []
121     train_f1s = []
122     val_f1s = []

```

```

123
124 # Training loop
125 for epoch in range( args.epochs):
126     train_loss, train_acc, train_prec, train_rec, train_f1 =
127         train_one_epoch(
128             model, train_loader, optimizer, criterion, device
129         )
130     train_losses.append(train_loss)
131     train_accs.append(train_acc)
132     train_precs.append(train_prec)
133     train_f1s.append(train_f1)
134
135     val_loss, val_acc, val_prec, val_rec, val_f1 = evaluate(
136         model, val_loader, criterion, device
137     )
138     val_losses.append(val_loss)
139     val_accs.append(val_acc)
140     val_precs.append(val_prec)
141     val_f1s.append(val_f1)
142
143     print(f"Epoch {epoch+1}/{args.epochs}")
144     print(f"Train Loss: {train_loss:.4f} | Acc: {train_acc:.4f}
145           | Prec: {train_prec:.4f} | Rec: {train_rec:.4f} | F1: {
146           train_f1:.4f}")
147     print(f"Val Loss: {val_loss:.4f} | Acc: {val_acc:.4f} | Prec
148           : {val_prec:.4f} | Rec: {val_rec:.4f} | F1: {val_f1:.4f}"
149         )
150
151     if val_f1>best_val_f1:
152         best_val_f1 = val_f1
153         torch.save(model.state_dict(), best_model_path)
154         print(f"New best model saved with F1: {best_val_f1:.4f}"
155             )
156
157 history_df = pd.DataFrame({
158     "epoch": list(range(1, args.epochs + 1)),
159     "train_loss": train_losses,
160     "val_loss": val_losses,

```

```

155     "train_acc": train_accs,
156     "val_acc": val_accs,
157     "train_prec": train_precs,
158     "val_prec": val_precs,
159     "train_f1": train_f1s,
160     "val_f1": val_f1s
161 })
162 history_df.to_csv(f"lstm_training_history.csv", index=False)
163
164 # Load best model and evaluate on test set
165 model.load_state_dict(torch.load(best_model_path))
166 test_loss, test_acc, test_prec, test_rec, test_f1 = evaluate(
167     model, test_loader, criterion, device
168 )
169
170 os.makedirs("results", exist_ok=True)
171
172 summary_df = pd.DataFrame([
173     "model": "lstm",
174     "test_loss": test_loss,
175     "test_accuracy": test_acc,
176     "test_precision": test_prec,
177     "test_recall": test_rec,
178     "test_f1": test_f1
179 ])
180
181 summary_df.to_csv(f"results/lstm_test_results.csv", index=False)
182
183 print("\nFinal Test Results:")
184 print(f"Test Loss: {test_loss:.4f} | Acc: {test_acc:.4f} | Prec:
185     {test_prec:.4f} | Rec: {test_rec:.4f} | F1: {test_f1:.4f}")
186
187 if __name__ == "__main__":
188     main()

```

Listing 4: train\_lstm.py

## 5. train\_logreg.py

```
1 import pandas as pd
2 from sklearn.feature_extraction.text import TfidfVectorizer
3 from sklearn.linear_model import LogisticRegression
4 from sklearn.metrics import accuracy_score,
5     precision_recall_fscore_support
6 import os
7 from sklearn.metrics import log_loss
8
9
10 def load_data():
11     train_df = pd.read_csv("data/train.csv")
12     val_df = pd.read_csv("data/val.csv")
13     test_df = pd.read_csv("data/test.csv")
14     return train_df, val_df, test_df
15
16
17 def evaluate_model(model, X, y):
18     preds = model.predict(X)
19
20     acc = accuracy_score(y, preds)
21     precision, recall, f1, _ = precision_recall_fscore_support(
22         y, preds, average="weighted", zero_division=0
23     )
24
25     return acc, precision, recall, f1
26
27
28 def main():
29     os.makedirs("results", exist_ok=True)
30
31     # Load data
32     train_df, val_df, test_df = load_data()
33
34     X_train = train_df["text"]
35     y_train = train_df["label"]
```

```

36
37 X_val = val_df["text"]
38 y_val = val_df["label"]
39
40 X_test = test_df["text"]
41 y_test = test_df["label"]
42
43 # TF-IDF vectorization
44 vectorizer = TfidfVectorizer(
45     lowercase=True,
46     stop_words="english",
47     max_features=20000
48 )
49
50 X_train_tfidf = vectorizer.fit_transform(X_train)
51 X_val_tfidf = vectorizer.transform(X_val)
52 X_test_tfidf = vectorizer.transform(X_test)
53
54 # Logistic Regression model
55 model = LogisticRegression(
56     max_iter=1000,
57     random_state=42
58 )
59
60 model.fit(X_train_tfidf, y_train)
61
62 # Validation metrics
63 val_acc, val_prec, val_rec, val_f1 = evaluate_model(model,
64     X_val_tfidf, y_val)
65
66 print("Validation Results")
67 print(f"Accuracy : {val_acc:.4f}")
68 print(f"Precision: {val_prec:.4f}")
69 print(f"Recall    : {val_rec:.4f}")
70 print(f"F1-score  : {val_f1:.4f}")
71
72 # Test metrics

```

```

72     test_acc, test_prec, test_rec, test_f1 = evaluate_model(model,
73         X_test_tfidf, y_test)
74
75     print("\nTest Results")
76     print(f"Accuracy : {test_acc:.4f}")
77     print(f"Precision: {test_prec:.4f}")
78     print(f"Recall    : {test_rec:.4f}")
79     print(f"F1-score  : {test_f1:.4f}")
80
81     probs = model.predict_proba(X_test_tfidf)
82     test_loss = log_loss(y_test, probs)
83
84     # Save summary
85     summary_df = pd.DataFrame([
86         "model": "logistic_regression",
87         "val_accuracy": val_acc,
88         "val_precision": val_prec,
89         "val_recall": val_rec,
90         "val_f1": val_f1,
91         "test_accuracy": test_acc,
92         "test_precision": test_prec,
93         "test_recall": test_rec,
94         "test_f1": test_f1,
95         "test_loss": test_loss
96     ])
97
98     summary_df.to_csv("results/logreg_test_results.csv", index=False)
99
100     print("\nSaved results to results/logreg_test_results.csv")
101
102 if __name__ == "__main__":
103     main()

```

Listing 5: train\_logreg.py

## 6. visual.py

```

1 import pandas as pd
2 import matplotlib.pyplot as plt
3 import os
4
5 def plot_lstm_loss(lstm_history):
6     plt.figure(figsize=(10, 5))
7     plt.plot(lstm_history["epoch"], lstm_history["train_loss"],
8             label="LSTM Train Loss")
9     plt.plot(lstm_history["epoch"], lstm_history["val_loss"], label=
10             "LSTM Val Loss")
11     plt.xlabel("Epoch")
12     plt.ylabel("Loss")
13     plt.title("LSTM Training and Validation Loss")
14     plt.legend()
15     plt.savefig(os.path.join("plots", "lstm_loss_plot.png"))
16     plt.show()
17
18 def plot_lstm_accuracy(lstm_history):
19     plt.figure(figsize=(10, 5))
20     plt.plot(lstm_history["epoch"], lstm_history["train_acc"], label=
21             "LSTM Train Accuracy")
22     plt.plot(lstm_history["epoch"], lstm_history["val_acc"], label="
23             LSTM Val Accuracy")
24     plt.xlabel("Epoch")
25     plt.ylabel("Accuracy")
26     plt.title("LSTM Training and Validation Accuracy")
27     plt.legend()
28     plt.savefig(os.path.join("plots", "lstm_accuracy_plot.png"))
29     plt.show()
30
31 def plot_lstm_f1(lstm_history):
32     plt.figure(figsize=(10, 5))
33     plt.plot(lstm_history["epoch"], lstm_history["train_f1"], label=
34             "LSTM Train F1")
35     plt.plot(lstm_history["epoch"], lstm_history["val_f1"], label="
36             LSTM Val F1")
37     plt.xlabel("Epoch")
38     plt.ylabel("F1 Score")

```

```

33     plt.title("LSTM Training and Validation F1 Score")
34     plt.legend()
35     plt.savefig(os.path.join("plots", "lstm_f1_plot.png"))
36     plt.show()
37
38 def plot_test_comparison(lstm_test, logreg_test):
39     metrics = ["test_accuracy", "test_precision", "test_recall", "
40               test_f1"]
41     metric_names = ["Accuracy", "Precision", "Recall", "F1 Score"]
42
43     lstm_values = [lstm_test[m].iloc[0] for m in metrics]
44     logreg_values = [logreg_test[m].iloc[0] for m in metrics]
45
46     x = range(len(metrics))
47     width = 0.35
48
49     plt.figure(figsize=(10, 5))
50     plt.bar([i - width/2 for i in x], lstm_values, width=width,
51            label="LSTM")
52     plt.bar([i + width/2 for i in x], logreg_values, width=width,
53            label="Logistic Regression")
54     plt.xlabel("Metrics")
55     plt.ylabel("Score")
56     plt.title("Final Test Results Comparison")
57     plt.xticks(list(x), metric_names)
58     plt.legend()
59     plt.savefig(os.path.join("plots", "test_comparison.png"))
60     plt.show()
61
62 def main():
63     os.makedirs("plots", exist_ok=True)
64
65     lstm_history = pd.read_csv("lstm_training_history.csv")
66     logreg_history = pd.read_csv("results/logreg_training_history.
67                                  csv")
68
69     lstm_test = pd.read_csv("results/lstm_test_results.csv")
70     logreg_test = pd.read_csv("results/logreg_test_results.csv")

```

```
67
68     plot_lstm_loss(lstm_history)
69     plot_lstm_accuracy(lstm_history)
70     plot_lstm_f1(lstm_history)
71     plot_test_comparison(lstm_test, logreg_test)
72
73 if __name__ == "__main__":
74     main()
```

Listing 6: visual.py