

عبدالرحمن الحمادي

# Abdulrahman Alhammadi - Final Project 2\_version\_ Thesis..docx

 Midocean University

## Document Details

### Submission ID

trn:oid::2945:337164319

### Submission Date

Dec 13, 2025, 7:15 PM GMT+2

### Download Date

Dec 13, 2025, 7:54 PM GMT+2

### File Name

Abdulrahman Alhammadi - Final Project 2\_version\_ Thesis..docx

### File Size

1.9 MB

22 Pages

2,868 Words

18,955 Characters

# 100% detected as AI

The percentage indicates the combined amount of likely AI-generated text as well as likely AI-generated text that was also likely AI-paraphrased.

**Caution: Review required.**

It is essential to understand the limitations of AI detection before making decisions about a student's work. We encourage you to learn more about Turnitin's AI detection capabilities before using the tool.

## Detection Groups

-  **45 AI-generated only 100%**  
Likely AI-generated text from a large-language model.
-  **0 AI-generated text that was AI-paraphrased 0%**  
Likely AI-generated text that was likely revised using an AI-paraphrase tool or word spinner.

### Disclaimer

Our AI writing assessment is designed to help educators identify text that might be prepared by a generative AI tool. Our AI writing assessment may not always be accurate (i.e., our AI models may produce either false positive results or false negative results), so it should not be used as the sole basis for adverse actions against a student. It takes further scrutiny and human judgment in conjunction with an organization's application of its specific academic policies to determine whether any academic misconduct has occurred.

## Frequently Asked Questions

### How should I interpret Turnitin's AI writing percentage and false positives?

The percentage shown in the AI writing report is the amount of qualifying text within the submission that Turnitin's AI writing detection model determines was either likely AI-generated text from a large-language model or likely AI-generated text that was likely revised using an AI paraphrase tool or word spinner.

False positives (incorrectly flagging human-written text as AI-generated) are a possibility in AI models.

AI detection scores under 20%, which we do not surface in new reports, have a higher likelihood of false positives. To reduce the likelihood of misinterpretation, no score or highlights are attributed and are indicated with an asterisk in the report (\*%).

The AI writing percentage should not be the sole basis to determine whether misconduct has occurred. The reviewer/instructor should use the percentage as a means to start a formative conversation with their student and/or use it to examine the submitted assignment in accordance with their school's policies.

### What does 'qualifying text' mean?

Our model only processes qualifying text in the form of long-form writing. Long-form writing means individual sentences contained in paragraphs that make up a longer piece of written work, such as an essay, a dissertation, or an article, etc. Qualifying text that has been determined to be likely AI-generated will be highlighted in cyan in the submission, and likely AI-generated and then likely AI-paraphrased will be highlighted purple.

Non-qualifying text, such as bullet points, annotated bibliographies, etc., will not be processed and can create disparity between the submission highlights and the percentage shown.





# **TruthLens: Detecting Fake News Using NLP**

Prepared By

Name: [Abdulrahman Alhammadi] IDNo.: [251002075]

Submitted in fulfillment of the requirements for the degree  
Master of Science in Artificial Intelligence

At

College of Informatics

Midocean University

UAE

Supervisor Name: Dr. Khaled Eskaf

College of Informatics, Midocean University, Comoros

[2025]

# Acknowledgements

I would like to express my sincere gratitude to Dr. Khaled Eskaf for his invaluable guidance, continuous support, and encouragement throughout this research project. His expertise, insightful feedback, and unwavering dedication have significantly contributed to the quality and direction of this thesis.

I also extend my heartfelt appreciation to my family, friends, and colleagues, whose patience, motivation, and moral support were instrumental in the successful completion of this work.

A special thank you goes to my brother, Bader, for his generous help and continuous encouragement throughout this journey.

# Abstract

In the current social networking landscape, web news inconsistencies and the pervasive spread of fake news constitute an urgent global threat, influencing consumer purchases and undermining societal peace. Disinformation can easily manipulate populations and governments, demanding immediate and effective countermeasures. This study responds to the critical need for a more rigorous and empirical investigation into inconsistencies within reported reviews and the broader challenge of Fake News Detection (FND). We conduct an in-depth review, specifically examining the most fundamental and practical application of machine learning (ML) and deep learning (DL) algorithms for detecting fake data. The review systematically classifies techniques used across various datasets, meticulously detailing their performance indicators and simulation media. This comprehensive analysis delivers essential findings and highlights significant challenges, thereby guiding both academic and professional efforts to strengthen the reliability of automated FND models.

## List of Publications

## Table of Contents

|                               |                              |
|-------------------------------|------------------------------|
| List of Abbreviations .....   | v                            |
| List of Figures .....         | Error! Bookmark not defined. |
| List of Tables .....          | ix                           |
| Chapter 1: Introduction.....  | 1                            |
| 1.1 Problem Definition .....  | 1                            |
| 1.2 Thesis Objectives .....   | 1                            |
| 1.3 Thesis Contribution ..... | 2                            |
| 1.4 Thesis Organization.....  | 3                            |

|                                                             |           |
|-------------------------------------------------------------|-----------|
| <b>Chapter 2: Overview.....</b>                             | <b>4</b>  |
| 2.1 Introduction .....                                      | 4         |
| 2.2 Classical Machine Learning Methods .....                | 4         |
| 2.3 Deep Learning Approaches .....                          | 4         |
| 2.4 Hybrid and Optimized Models .....                       | 5         |
| 2.5 Use of Real-world Datasets .....                        | 5         |
| 2.6 Quantitative Results Summary .....                      | 6         |
| 2.7 Research Gaps and Opportunities .....                   | 6         |
| 2.8 Conclusion .....                                        | 7         |
| <b>Chapter 3: Related Work .....</b>                        | <b>8</b>  |
| 3.1 Introduction .....                                      | 8         |
| 3.2 History .....                                           | 8         |
| <b>Chapter 4: Proposed Approach .....</b>                   | <b>9</b>  |
| 4.1 Data Collection and Description .....                   | 9         |
| 4.2 Data Preprocessing .....                                | 9         |
| 4.3 Feature Engineering .....                               | 10        |
| 4.4 Feature Selection .....                                 | 10        |
| 4.5 Model Selection and Optimization .....                  | 11        |
| 4.6 Model Evaluation .....                                  | 12        |
| 4.7 Visualization and Insights .....                        | 12        |
| <b>Chapter 5: Experimental Results and Discussion .....</b> | <b>13</b> |
| 5.1 Experimental Setup .....                                | 13        |
| 5.2 Baseline Model Performance .....                        | 14        |
| 5.3 Feature Selection and Model Refinement .....            | 14        |
| 5.4 Visualization of Delay Factors .....                    | 15        |
| 5.5 Discussion.....                                         | 18        |
| <b>Chapter 6: Conclusion and Future Work .....</b>          | <b>19</b> |
| 6.1 Conclusion .....                                        | 19        |
| 6.2 Future Work .....                                       | 20        |
| <b>References .....</b>                                     | <b>22</b> |

## List of Abbreviations

| Abbreviation | Full Form                                               | Explanation                                                                                    |
|--------------|---------------------------------------------------------|------------------------------------------------------------------------------------------------|
| Kaggle       | -                                                       | A platform for data science competitions, datasets, and collaboration.                         |
| TF-IDF       | Term Frequency–Inverse Document Frequency               | A statistical measure to evaluate the importance of a word in a document relative to a corpus. |
| Word2Vec     | -                                                       | A technique to embed words into vector space using neural networks. Developed by Google.       |
| BERT         | Bidirectional Representations Encoder from Transformers | A transformer-based model by Google for NLP tasks, considers context from both directions.     |
| UI           | User Interface                                          | The point of interaction between the user and a digital product/system.                        |
| FND          | Fake News Detection                                     | Fake News Detection                                                                            |
| ML           | Machine learning                                        | Machine learning                                                                               |

|     |                         |                         |
|-----|-------------------------|-------------------------|
| DL  | Deep Learning           | Deep Learning           |
| NER | Name Entity Recognition | Name Entity Recognition |

|                                                                           |           |
|---------------------------------------------------------------------------|-----------|
| <b>Figure 1: Thesis Structure and Research Flowchart .....</b>            | <b>3</b>  |
| <b>Figure 2: Feature Engineering Workflow.....</b>                        | <b>10</b> |
| <b>Figure 3: Random Forest Classifier .....</b>                           | <b>11</b> |
| <b>Figure 4: Impact of Outlier Removal on Dataset .....</b>               | <b>13</b> |
| <b>Figure 5: Feature Selection Comparison Best vs RFE.....</b>            | <b>15</b> |
| <b>Figure 6: Correlation Matrix for Delay and Weather Variables .....</b> | <b>15</b> |
| <b>Figure 7: Weather Delay vs. Total Arrival Delay .....</b>              | <b>16</b> |
| <b>Figure 8: Weather Count vs. Delay Rate .....</b>                       | <b>17</b> |
| <b>Figure 9: Weather Delay Over Time .....</b>                            | <b>17</b> |
| <b>Figure 10: Monthly Weather Delay Trends .....</b>                      | <b>18</b> |

# Chapter 1

## Introduction

### 1.1 Problem Definition

In today's rapidly evolving digital landscape, the internet has become an essential component of daily life, with billions of users engaging across a wide range of social networking platforms. The unprecedented velocity of information dissemination allows ideas, conversations, and events to spread globally within seconds. A minor incident occurring in a remote village in Japan, for instance, can instantly reach audiences worldwide through audio, video, and real-time updates—immersing individuals in the event as if they were physically present. Although this constant stream of information offers tremendous benefits, it simultaneously raises a critical concern: how can we verify the authenticity of the news we consume, and how can we distinguish factual reporting from fabricated content?

This challenge underscores the urgent need for advanced research and development focused on devising robust methodologies, frameworks, and intelligent systems capable of mitigating the spread of fake news. The proliferation of misinformation has far-reaching consequences that extend beyond individual users. Fake news has emerged as a powerful tool capable of shaping public opinion, manipulating public sentiment, and disrupting social harmony. When strategically deployed, misinformation can fuel societal tensions, instigate polarization, and erode public trust in media and governmental institutions.

Politically, fake news poses a significant threat to national stability and decision-making processes. Misleading narratives and coordinated disinformation campaigns have the potential to influence electoral outcomes, undermine the legitimacy of governments, distort policy debates, and weaken confidence in democratic systems. Several case studies around the world have demonstrated how disinformation can be weaponized to target political actors, interfere in national elections, and amplify ideological divisions.

Economically, the spread of fabricated content can generate substantial financial losses and market instability. False information about companies, sectors, or

economic indicators can cause stock market volatility, damage the reputation of major firms, disrupt supply chains, or be exploited for market manipulation and speculative gains. In an interconnected global economy, a single piece of misleading financial news can trigger cascading effects across multiple industries.

On a societal level, fake news amplifies misinformation, reinforces stereotypes, and encourages the escalation of hate speech and social discord. With many individuals relying heavily on social media as their primary news source, the risk of consuming unverified or deliberately manipulated information has increased significantly—especially with the rise of sophisticated AI-generated content such as deepfakes, which further blurs the line between truth and deception.

Even the sports sector has not been immune to the spread of fake news. Fabricated stories about player transfers, match results, contract disputes, or internal club conflicts can inflame fan reactions, create unnecessary pressure on athletes, and destabilize club management. These rumours often spread rapidly due to the emotional intensity associated with sports communities and can influence public perception and organizational decision-making.

Given the magnitude of these challenges, this scientific paper aims to contribute to the growing body of research on fake news detection by presenting a comprehensive study that incorporates Deep Learning (DL) and Machine Learning (ML) approaches. Our work includes an extensive review of existing literature to identify current methodologies, highlight their limitations, and clarify the research gap that this study intends to address. By comparing our findings with previous studies, we analyze the differences in datasets, feature extraction techniques, model architectures, and evaluation metrics, providing insights into why certain approaches outperform others in specific contexts.

Furthermore, this research leverages state-of-the-art deep learning models to detect and classify fake news with improved accuracy and robustness. Through experimentation, performance evaluation, and comparative analysis, the proposed study demonstrates the effectiveness of these models and offers recommendations for future advancements in the field.

Ultimately, this paper seeks not only to examine the widespread impact of fake news on global stability—politically, economically, socially, and even in sports—but also

to contribute practical and scalable solutions grounded in cutting-edge artificial intelligence technologies.

## 1.2 Thesis Objectives

The primary objective of this study is to investigate and evaluate machine learning and deep learning approaches for the detection of fake news within online and social networking environments. To achieve this goal, the study is guided by the following specific objectives:

- To conduct a systematic review of existing machine learning and deep learning techniques applied to Fake News Detection (FND).
- To analyze and compare the performance of different ML and DL models across publicly available datasets.
- To examine the effectiveness of various feature extraction and preprocessing techniques in identifying inconsistencies within news content.
- To classify fake news detection approaches based on their underlying algorithms, data sources, and evaluation metrics.
- To identify key limitations, challenges, and research gaps in current automated fake news detection systems.

These objectives aim to support the development of more robust, accurate, and reliable fake news detection models for real-world applications.

## 1.3 Thesis Contribution

- NA

## 1.4 Thesis Organization




Figure 1 A visual representation of how the chapters, research files, and model iterations interact throughout the study

# Chapter 2 Overview

## 2.1 Introduction to Flight Delay Forecasting

## 2.2 Classical Machine Learning Methods

## 2.3 Deep Learning Approaches

## 2.4 Hybrid and Optimized Models

## 2.5 Use of Real-world Datasets

## 2.6 Quantitative Results Summary

| Study                          | Model         | Key Findings                                            |
|--------------------------------|---------------|---------------------------------------------------------|
| Yu et al. (2019)               | Decision Tree | 90% accuracy, PR = 0.439                                |
| Zhang et al. (2022)            | LSTM          | RMSE improved by 15% over baseline                      |
| Wu et al. (2022)               | XGBoost       | Best delay predictor among evaluated models             |
| Guo et al. (2021)              | SGDAN         | Effective 3-hour-ahead prediction                       |
| Khan et al. (2022)             | CNN-LSTM      | Accuracy exceeded 94%                                   |
| Li et al. (2022)               | DeepONet      | Achieved >93% accuracy on high-dimensional data         |
| Hu & Zhang (2023)              | CNN-BiLSTM    | Enhanced MAE and RMSE through hybrid deep learning      |
| Pineda-Jaramillo et al. (2024) | LDA + xAI     | Identified temperature and wind speed as top predictors |

## 2.7 Research Gaps and Opportunities

## 2.8 Conclusion

## Chapter 3 Related Work

### 3.1 Introduction

### 3.2 History

## Chapter 4 Proposed Approach

- Data Preprocessing Module

The data preprocessing module plays a vital role in enhancing the performance and reliability of the fake news detection system. Given that the dataset is collected from publicly available sources on the Kaggle website, the raw data contains noise, inconsistencies, and heterogeneous formats that must be addressed before model training. This module is designed to transform raw textual content into a clean,

structured, and informative representation suitable for machine learning and deep learning models.

- Text Processing:

The text processing pipeline applies advanced normalization and cleaning techniques to improve data quality and semantic representation. These techniques include the removal and normalization of special characters, interpretation of emojis to retain sentiment signals often present in social media and online news content, and context-aware lemmatization to unify word forms while preserving meaning. In addition, named entity recognition (NER) is employed to identify and extract key entities such as persons, organizations, and locations, which are particularly important in fake news detection as misleading articles often manipulate or fabricate such entities.

This preprocessing strategy reduces textual noise, preserves contextual meaning, and enhances the discriminative power of the extracted features, thereby improving the overall effectiveness of the fake news classification models.

## Chapter 5

# Experimental Results and Discussion

This chapter presents the results of the developed flight delay prediction model and reflects on its performance. Using real-world data from airline operations, the study assesses the model's ability to classify delays under various configurations. The goal is not only to report accurate metrics but also to interpret their implications in an operational context.

## 5.1 Experimental Setup

All experiments were conducted using Python, utilizing widely adopted libraries such as Scikit-learn, Pandas, NumPy, and Matplotlib. The original dataset comprised over 170,000 records; however, after removing outliers and handling missing values, the final working dataset contained approximately 51,781 records.

With categorical data transformed using one-hot encoding, the total feature space expanded to 424 columns — a necessary step to allow the model to differentiate between operational entities such as airlines and airports. The data was then split into 80% for training and 20% for testing, ensuring an unbiased estimate of model performance.




Figure 2 Demonstrates how preprocessing affects data volume

## 5.2 Baseline Model Performance

The first model configuration used all available features. A Random Forest Classifier, tuned with GridSearchCV, was trained using the following parameters:

- `n_estimators`: 100
- `max_depth`: None.
- `min_samples_split`: 5

The resulting performance metrics on the test set were:

- Test Accuracy: 93.19%
- Precision (On-Time Class): 93%
- Recall (On-Time Class): 99%
- Precision (Delayed Class): 96%
- Recall (Delayed Class): 73%
- F1 Score (Delayed Class): 83%

While overall accuracy was solid, the lower recall for delayed flights (73%) suggests that the model sometimes missed actual delays, a significant concern in real-world operations where underestimating delays can lead to passenger dissatisfaction and logistical disruptions.

### 5.3 Feature Selection and Model Refinement

To improve performance and reduce model complexity, feature selection techniques were applied:

- **SelectKBest (Mutual Information):** Identified operational features like `carrier_ct`, `nas_delay`, and `late_aircraft_ct` as highly predictive.
- **Recursive Feature Elimination (RFE):** Further refined the input set by selecting a combination of delay counts and specific airport codes (e.g., `airport_ADK`, `airport_MRY`) based on their contribution to model accuracy.

When the Random Forest model was retrained using only the RFE-selected features, the improvements were striking:

- **Test Accuracy:** 98.41%
- **Precision (all classes):**  $\geq 98\%$
- **Recall (All Classes):**  $\geq 95\%$
- **F1 Score (Delayed Class):** 0.96

This performance boost underscores the value of careful feature selection. Not only did the model generalize better, but it also became more efficient, requiring fewer inputs to reach higher predictive accuracy. It also performed more consistently across both classes, improving reliability for operational use.



Figure 3 Shows the efficiency of different selection strategies

## 5.4 Visualization of Delay Factors

To explore the reasons behind the model's predictions and offer insights for stakeholders, several visual analyses were conducted:

### Correlation Heatmap:




Figure 4 It illustrates the basis for discussing key predictive features

This visualization revealed strong correlations between weather\_ct, weather\_delay, and arr\_delay, confirming that weather remains one of the most impactful drivers of delays.

- **Scatter Plot (Weather Delay vs. Total Delay):**






Figure 5 It illustrates the basis for discussing key predictive features

A clear upward trend was observed, indicating that increases in weather-related delay counts corresponded with rises in total delays. This pattern was consistent across different airports and periods.

- **Time Series Plot:**



*Figure 6 This figure builds on the scatter plot by showing the frequency impact*



*Figure 7 Demonstrates seasonal delay trends*

Seasonal variations were highlighted, with delay rates peaking in winter months, a trend consistent with known weather patterns in aviation.

These visual tools provided interpretability and helped validate the model's internal logic. For operations teams, such insights can inform planning decisions, especially in weather-sensitive routes or peak travel seasons.




Figure 8 Concludes the visual evidence for seasonal variation

## 5.5 Discussion

The experimental results validate the core hypothesis of this thesis: AI, particularly Random Forests enhanced by thoughtful feature selection, is highly effective for predicting flight delays. Several key takeaways emerged:

- Simplifying the model improves performance: By focusing on the most informative features, the model not only gained accuracy but also became easier to interpret and deploy.
- Weather and carrier delays are central predictors: These features consistently ranked highest in both model importance and visual correlation, indicating their substantial operational impact.
- Visualizations support both analysis and communication: Tools like scatter plots and time series charts help translate technical findings into actionable insights for airline decision-makers.

However, it is essential to note that this model was evaluated using historical data. For real-time deployment, several additional components would be necessary, including integration with live data feeds, more detailed airport-specific attributes, and potentially even passenger-level data for personalized delay forecasting.

# Chapter 6

## Conclusion and Future Work

### 6.1 Conclusion

This study investigated the use of Logistic Regression with TF-IDF for detecting fake news. The results show that the model can effectively distinguish between fake and real news, capturing important patterns in the text.

Logistic Regression performed well as a baseline, even without complex deep learning models, demonstrating that simple, interpretable methods can be reliable for fake news detection. The preprocessing and feature extraction steps, especially TF-IDF, were key to improving the model's performance.

Overall, the study confirms that classical machine learning techniques are a practical and efficient option for automated fake news detection. Future work can build on these results by exploring more advanced models or additional features to further improve accuracy.

### 6.2 Future Work

Based on the findings of this study, several recommendations are proposed:

#### Scientific Recommendations:

- Future studies should explore advanced machine learning and deep learning models, such as transformer-based architectures, to capture deeper contextual and semantic patterns in news content.
- Incorporating additional linguistic and metadata features, including source credibility and temporal patterns, could further enhance model performance.

#### Practical Recommendations:

- Media organizations and online platforms can adopt lightweight and interpretable models like Logistic Regression with TF-IDF as a baseline solution for real-time fake news detection.
- Emphasis should be placed on model transparency and explainability to maintain user trust and support informed decision-making.

#### Research Suggestions:

- Expand research to include multilingual datasets and cross-platform news sources to improve generalizability.
- Investigate hybrid approaches combining classical and deep learning techniques for better accuracy and efficiency.

These recommendations aim to guide both academic research and practical applications in improving automated fake news detection systems.

## References

- Alnabhan, M., & Branco, P. (2024). Fake News Detection Using Deep Learning: A Systematic Literature Review. Ottawa, ON, Canada. Reference
- Mishra, A., & Sadia, H. (2023). A Comprehensive Analysis of Fake News Detection Models: A Systematic Literature Review and Current Challenges. Lucknow, India. Reference
- Yaswanthraj, S., & Nadar, P., & Praveshika, P., & Preethika, KG., Sanchana, NM. (2023). Fake News Detection Using (NLP). India. Reference
- Ahmed, A., & Aljarboun, A., & Donepudi, P., & Choi, M. (2023). Detecting Fake News using Machine Learning: A Systematic Literature Review:). China & USA. Reference