

عبدالرحمن الحمادي

Abdulrahman Alhammadi - Final Project 2_version_ Thesis. 1 3.docx

- Final Thesis Submission after Reduce Similarity
- Thesis_(2)_AI901_Dr_Khaled_Eskaf_Third Trimester 2025 (Moodle PP)
- Midocean University

Document Details

Submission ID

trn:oid::1:3451892038

Submission Date

Dec 31, 2025, 9:35 PM GMT+4

Download Date

Jan 1, 2026, 2:29 AM GMT+4

File Name

xcod....61_111621_sisehT_noisrev_2_tcejorP_laniF_-_idammahIA_namharludbA_عبدالرحمن_الحمادي_05192

File Size

132.3 KB

17 Pages

2,598 Words

14,563 Characters

71% 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

-  **26 AI-generated only 71%**
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 want to thanks Dr. Khaled Eskaf for his guidance and support during the time of this research. His guidance was always helpful and his feedback helped me see all things from a better picture.

I'm also want to be grateful to my wife for her patience and understanding, and to my daughter for bringing happy and reminding me why this work is important. Their love and support provide me the strength to finish this journey. Honestly, I couldn't have done it without them.

A special thank you to my parents for their constant prayers, wisdom, and belief in me. They have always been my foundation, and I truly appreciate everything they have done to help me reach this point. Honestly, I couldn't have done it without them.

Abstract

These days, social media feels like a battlefield of information where truth and lies collide every second. Fake news doesn't just confuse people; it spreads like wildfire, shaping opinions and even influencing what we buy. Beyond personal choices, misinformation can shake economies and destabilize governments. That's why tackling this problem isn't optional it requires fast, decisive action before it spirals out of control.

In this project, we focus on using deep learning (DL) and machine learning (ML) techniques to help detect fake news effectively.

Table of Contents

List of Abbreviations	4
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
 Chapter 2: Overview.....	 4
2.1 Introduction	4
2.2 Classical Machine Learning Methods	4
2.3 Deep Learning Approaches	4
2.4 Quantitative Results Summary	6
2.5 Conclusion	7
 Chapter 3: Proposed Approach	 9
4.1 Data Preprocessing Model.....	9
4.2 Text Processing.....	10
 Chapter 4: Experimental Results and Discussion	 13
5.1 Logistics Regression Model	13

5.2 4.2 Model Evaluating.....,,.....	14
Chapter 5: Conclusion and Future Work	19
6.1 Conclusion	19
6.2 Future Work,.....	20
References	22

List of Abbreviations

- NLP: Natural Language Processing
- Kaggle: A platform for data science competitions, datasets, and collaboration.
- TF-IDF: Term Frequency–Inverse Document Frequency.
- Word2Vec: A technique to embed words into vector space using neural networks. Developed by Google.
- BERT: Bidirectional Encoder Representations from Transformers
- UI: User Interface
- FND: Fake News Detection
- ML: Machine Learning
- DL: Deep Learning
- NER: Name Entity Recognition

List of Figures

Figure 1: Confusion matrix showing the Performance.....	3
Figure 2: Overall performance metrics of the Logistic Regression model...	10
Figure 3: Classification report showing precision, recall, and F1-score.....	11

1 Chapter 1: Introduction

1.1 Problem Definition

Today, the internet is a big part of our daily lives, and billions of people use social media to connect and share. Information now moves faster than ever ideas, conversations, and events can spread across the world in just seconds. For example, something small that happens in a remote village in Japan can quickly reach people everywhere through videos, photos, and live updates, making it feel like you're right there.

While the constant flow of information has many benefits, it also brings an important question: How do we know if the news we read is true? And how can we tell the difference between real reports and fake reports?

These challenges make it clear that we urgently need a good research and smarter systems to stop the global spread of fake news.

Wrong information doesn't just confuse people; it can shake entire communities. Fake news changes how people think, stirs emotions, and slowly eats away at trust. And when it's done on purpose, the damage gets worse. Suddenly, people start questioning the media, doubting the government, and the whole system feels fragile.

Fake news is a serious threat to politics. It doesn't just spread lies; it shakes the stability of countries and messes with how decisions are made. Wrong stories and organized campaigns can flip election results, make people doubt their leaders, and twist political conversations. Over time, trust in democracy starts to crumble. We've seen it happen in many places: fake news used as a weapon to attack politicians, disrupt elections, and deepen divisions among people.

Fake news can hit the economy hard. A single false story about a company or market can send stock prices swinging, damage reputations, and even disrupt supply chains. Some people exploit this chaos to manipulate markets and make quick profits. In today's hyper-connected world, one wrong piece of financial news can ripple across industries and cause serious trouble.

Fake news doesn't just confuse; it tears at the fabric of society. It spreads false information, reinforces harmful stereotypes, and fuels hate speech and division. With so many people getting their news from social media, the chances of running into fake or manipulated content are sky-high. And now, with AI tools like deepfakes, it's even harder to tell what's real and what's not.

Sports isn't untouched by fake news. False reports about transfers, match results, contracts, or internal club issues can stir up fans, put players under pressure, and create chaos for managers. These rumors spread quickly because fans are deeply invested; and that passion can distort opinions and even push sports organizations into making hasty decisions.

These challenges are complex, so I set out to improve how fake news is detected. In this report, I apply deep learning (DL) and machine learning (ML) techniques to tackle the problem. I review existing research to understand current methods, spot gaps, and explain what this thesis aims to solve. I also compare my results with previous studies, analyzing datasets, feature extraction approaches, model designs, and evaluation strategies. This comparison offers insights into why some techniques outperform others in different situations.

This research also explores new deep learning models to improve the accuracy of detecting and classifying fake news. I test these models, evaluate their performance, and compare them with existing approaches. The study highlights how well these models work and offers insights into and recommendations for future research to advance this field

Finally, this paper examines the broader impact of fake news across politics, the economy, society, and even sports. It also proposes practical solutions that can scale, leveraging the latest AI technologies to make a real difference.

1.2 Thesis Objectives

The main goal of this thesis is to evaluate machine learning (ML) and deep learning (DL) techniques for detecting fake news on the internet and social media. To achieve this, I focus on several key steps:

- Reviewing current ML and DL approaches used for fake news detection.

- Testing and comparing different models on publicly available datasets.
- Examining the role of feature extraction and preprocessing in improving detection accuracy.
- Grouping detection methods based on algorithms, data sources, and evaluation strategies.
- Identifying major limitations, challenges, and gaps in existing systems.

Through these steps, I aim to contribute to building more robust, accurate, and reliable fake news detection models for real-world applications.

2 Chapter 2: Overview

2.1 Introduction to FND

Fake news detection means finding news that is wrong or made up. Social media and online platforms grow so fast now, fake news is everywhere and it changes what people think and what choices they make. To find fake news, I use machine learning and deep learning methods to look at the text and decide if it is real or fake.

2.2 Classical Machine Learning Methods

Old machine learning methods like logistic regression, SVM, and random forests still get used a lot for text tasks like finding fake news. These models need feature extraction like TF-IDF or Bag-of-Words to change text into numbers. For example, a fake news system can train a logistic regression model with TF-IDF features from news articles to learn patterns that show fake stuff. Even if deep learning is popular now, these old methods still work good when the dataset is small and when we need results that are easy to understand.

2.3 Deep Learning Approaches

Deep learning methods for finding fake news use neural networks that learn hard patterns in text by themselves, no need to make features by hand. Old methods use things like TF-IDF, but these models use word embeddings like Word2Vec and GloVe, or advanced ones like BERT, to understand meaning better. Models like CNN and RNN (with LSTM and GRU) look at word order and context. Now transformer models like BERT, RoBERTa, and XLNet are most popular because they check both sides of context and give very good accuracy. These models can be trained again on datasets with real and fake news, so they become very strong for understanding language.

2.4 Quantitative Results Summary

Study	Model	Key Findings
Mohammad Anlabhan	RNN	Accuracy - 99%

Study	Model	Key Findings
Paula Branco		
Alok Mishra Halima Sadia	LSTM	Offers strong performance but may lose information depending on feature extraction
Paraveshka Sanchana	Random Forest	Accuracy - 79%
Alim AlAyub Ayman Aljarbuh	Naïve Bayes	Accuracy – 98%

2.5 Conclusion

This thesis looks at using logistic regression with TF-IDF to detect fake news. The results show the model can separate fake news from real news by finding patterns in the text.

Logistic regression worked pretty well as a simple model, even without deep learning. This shows that easy and clear methods can still do a good job for fake news detection. Steps like cleaning the data and using TF-IDF were really important to make the model better.

Overall, this thesis shows old machine learning methods are still useful and practical for fake news detection. In the future, work can try more advanced models or add new features to make accuracy even better.

3 Chapter 3: Proposed Approach

3.1 Data Preprocessing Module

I took the dataset from Kaggle, and the raw data had noise, mistakes, and different formats that needed fixing before training the model. Cleaning the data was really important to make the fake news detection system work better and give good results. This step changed the raw text into clean and organized data that works good for machine learning and deep learning models.

3.2 Text Processing:

Text cleaning has steps to make the data better while keeping the meaning. First, I remove special characters and fix messy stuff. I also check emojis because they show feelings in social media and news. Then I do lemmatization so words look the same but still mean the same. Another step is named entity recognition (NER) to find names of people, places, and organizations since fake news often changes or makes up these names.

This cleaning helps remove noise, keep the context, and make features stronger so the fake news detection models work better.

4 Chapter 4: Experimental Results and Discussion

This part shows the results of the fake news detection model and how well it works. We use real datasets with news articles to test if the model can tell real news from fake news in different settings. The goal is not just to show the numbers but also to explain what they mean for fighting fake news.

4.1 Logistics Regression Model

What is Logistic Regression?

Logistic Regression is a machine learning algorithm for classification.

It is used to answer yes/no questions, such as:

- Fake news or real news?
- Spam or regular email?
- Sick or not sick?

Why is it called regression if it's a classification?

First, it calculates a numerical value, Then, it converts it to a probability between 0 and 1. After that, it decides:

- If the probability is $\geq 0.5 \rightarrow$ Category 1
- If the probability is $< 0.5 \rightarrow$ Category 0

How does it work? (Simply put)

It takes numerical data (like TF-IDF for words), multiplies it by weights, sums the results, and then passes them to a function called Sigmoid. The Sigmoid function (the real heart) converts any number to a value between 0 and 1.

Example:

- Result = 0.9 $\rightarrow$ High probability (Fake News)
- Result = 0.1 $\rightarrow$ Low probability (Real News)

Suppose the news item contains the words:

- “Breaking”
- “Government”

- “Urgent”

The model learns that:

- The word “Urgent” often appears in fake news.
- The word “Government” appears in both.

The model:

- Gives higher weight to the word “Urgent.”
- Summarizes the effect.
- Determines the outcome.

Why is Logistic Regression excellent for NLP?

- Fast
- Simple
- Works well with TF-IDF
- Its results are interpretable (we know the weight of each word).

That's why it's often used in (Fake News Detection)

4.2 Model Evaluating

Logistic regression was used to classify texts using a bag-of-words representation. GridSearchCV was employed to select the best model parameters by testing various solvers, employing 5-fold cross-validation and classification accuracy as performance metrics. After identifying the best model, it was used to predict the test data, its performance was evaluated using standard classification metrics, and finally, the resulting best parameters were presented.

Figure 1 A Confusion matrix showing the performance of the Logistic Regression model in classifying fake and real news.




Figure 2 An Overall performance metrics of the Logistic Regression model, including accuracy, precision, recall, and F1-score for fake news detection.




Figure 3 A Classification report showing precision, recall, and F1-score for fake and real news using the Logistic Regression model.

	precision	recall	f1-score	support
Fake News	0.51	0.53	0.52	2058
Real News	0.52	0.49	0.50	2102
accuracy			0.51	4160
macro avg	0.51	0.51	0.51	4160
weighted avg	0.51	0.51	0.51	4160

5 Chapter 5: Conclusion and Future Work

5.1 Conclusion

This study tested using Logistic Regression with TF-IDF to find fake news. The results show the model can tell fake news from real news by catching important patterns in the text.

Logistic Regression worked pretty good as a simple model, even without using deep learning. This shows that easy and clear methods can still do a good job for detecting fake news. Cleaning the data and applying TF-IDF was really important to make the model perform better.

Overall, this study proves that old machine learning techniques are still useful and practical for fake news detection. In the future, research can try more advanced models or add extra features to improve the accuracy even more.

.

5.2 Future Work

Based on findings of the study I did, there are a several recommendations is proposed:

Scientific Recommendations:

- The future studies should be using the advanced machine learning and deep learning models, for example transformer-based architectures, to find a deeper contextual and semantic patterns in the news articles.
- Incorporating additional linguistic and metadata features, including source credibility and temporal patterns, could further enhance model performance.

Practical Recommendations:

- Social Media companies and online platforms firms can adopt lightweight and interpretable models like Logistic Regression with TF-IDF as a baseline solution for real-time fake news detection.
- We should focus on making the model clear and easy to understand so people can trust it and make better decisions.

Research Suggestions:

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

These recommendations aim to help both researchers and practitioners improve automated fake news detection models.

6 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., & Aljarbounh, A., & Donepudi, P., & Choi, M. (2023). Detecting Fake News using Machine Learning: A Systematic Literature Review:). China & USA. Reference