



OPEN

# Detection of ChatGPT fake science with the xFakeSci learning algorithm

Ahmed Abdeen Hamed<sup>1,4</sup> & Xindong Wu<sup>2,3,4</sup>

Generative AI tools exemplified by ChatGPT are becoming a new reality. This study is motivated by the premise that “AI generated content may exhibit a distinctive behavior that can be separated from scientific articles”. In this study, we show how articles can be generated using means of prompt engineering for various diseases and conditions. We then show how we tested this premise in two phases and prove its validity. Subsequently, we introduce xFakeSci, a novel learning algorithm, that is capable of distinguishing ChatGPT-generated articles from publications produced by scientists. The algorithm is trained using network models driven from both sources. To mitigate overfitting issues, we incorporated a calibration step that is built upon data-driven heuristics, including proximity and ratios. Specifically, from a total of a 3952 fake articles for three different medical conditions, the algorithm was trained using only 100 articles, but calibrated using folds of 100 articles. As for the classification step, it was performed using 300 articles per condition. The actual label steps took place against an equal mix of 50 generated articles and 50 authentic PubMed abstracts. The testing also spanned publication periods from 2010 to 2024 and encompassed research on three distinct diseases: cancer, depression, and Alzheimer’s. Further, we evaluated the accuracy of the xFakeSci algorithm against some of the classical data mining algorithms (e.g., Support Vector Machines, Regression, and Naive Bayes). The xFakeSci algorithm achieved F1 scores ranging from 80 to 94%, outperforming common data mining algorithms, which scored F1 values between 38 and 52%. We attribute the noticeable difference to the introduction of calibration and a proximity distance heuristic, which underscores this promising performance. Indeed, the prediction of fake science generated by ChatGPT presents a considerable challenge. Nonetheless, the introduction of the xFakeSci algorithm is a significant step on the way to combating fake science.

**Keywords** ChatGPT, Generative AI, Fake publication, Human-generated publications, ML Algorithm, Fake science

With Large Language Models (LLMs) and generative AI tools (e.g., ChatGPT)<sup>1</sup> becoming a new reality, our world finds itself in a state of controversy. On one hand, there exists a camp of optimists who perceive their potential and seek to harness them. On the other hand, there are doubters who remain skeptical, seeking validation and further assessments to discern how this new paradigm will impact our lives. This division provides strong motivation for this study and catalyzes efforts towards providing a tool that assesses the capability of generating fake science by ChatGPT. Undoubtedly, real science, documented in scientific publications, stands as one of the most sacred sources of knowledge due to its invaluable contribution to future discoveries<sup>2–4</sup>. Historically, the widespread of various predatory journals has fueled the rise of fake science<sup>5–8</sup>, with the influence of social media exacerbating its impact<sup>9,10</sup>. Particularly during the global Coronavirus pandemic, the dissemination of misinformation regarding the importance of vaccination has led to its rejection by certain groups, endangering public health<sup>11–13</sup>. Another alarming incident during the pandemic involved the spread of an article reporting fake findings linking vitamin D deficiency to the death of 99 percent of the studied population. Although retracted, the misinformation was amplified by a DailyMail news article, spreading globally<sup>14,15</sup>. It is imperative to safeguard

<sup>1</sup>Complex Adaptive Systems and Computational Intelligence Laboratory, State University of New York at Binghamton, Binghamton, NY 13902, USA. <sup>2</sup>Hefei University of Technology, Key Laboratory of Knowledge Engineering with Big Data (the Ministry of Education of China), Hefei 230009, China. <sup>3</sup>China Electronics Data Industry Group, Shenzhen 518057, China. <sup>4</sup>These authors contributed equally: Ahmed Abdeen Hamed and Xindong Wu. ✉email: ahamed1@binghamton.edu; xwu@hfut.edu.cn

the authenticity of scientific publications from fraud or any influential factors that compromise the integrity of this crucial source of knowledge<sup>16</sup>.

In this article, we demonstrate how the emergence of ChatGPT (alongside many other generative AI tools) has impacted our society today: (1) the launch of many special issues and themes to study, assess, analyze, and test the impact and potential of ChatGPT<sup>17–21</sup>, (2) the adoption of new policies by journals regarding ChatGPT authorship<sup>16,22–26</sup>, (3) the development of ChatGPT plugins and inclusion in professional services such as Expedia and Slack<sup>27</sup>, and (4) the creation of educational tools (e.g., Wolfram) and the potential development of learning and educational support tools (e.g., Medical Licensing Examination<sup>28</sup>).

## Literature background

Due to the significance of fake science and the imperative for plagiarism detection mechanisms, numerous researchers have investigated fake science detection and employed diverse methodologies. Here, we list some of the most relevant works.

Chaka addressed the detection of content generated by tools such as ChatGPT and assessed it in five different languages. The work used a prompt-engineering approach to generate content in languages including Spanish, French, and German. The research concluded that detecting machine-generated content poses a significant challenge, and further investigations are much needed to protect against plagiarism<sup>29</sup>.

Cingillioglu also tackled the problem of identifying AI-generated essays using the Support Vector Machine algorithm (SVM)<sup>30</sup>. The work reported a 100% accuracy for human-created articles; however, their approach failed to state the accuracy of ChatGPT-generated essays<sup>31</sup>.

Elkhatat et al. evaluated 15 paragraphs generated by ChatGPT (versions 3.5 and 4.0) using means of other automatic detection tools such as Copyleaks<sup>32</sup> and CrossPlag<sup>33</sup>. The authors concluded that AI detection tools were able to predict ChatGPT-generated content in isolation without mixing with other content. However, when the generated content was perturbed using human-written responses, all detection tools failed to produce accurate or consistent results<sup>34</sup>.

In their editorial effort, Anderson et al. also presented similar issues related to ChatGPT-generated articles. They further noted the lack of potent algorithms to detect AI-generated content. The authors concluded by mandating the need for publishers to introduce detection tools throughout the lifecycle of publications<sup>35</sup>.

Rashidi et al. presented their tool that determines whether an article is human-created or otherwise. Similar to ours, the study considered publication abstracts, which were collected from top-quality journals in the period of 1980–2023. Using a text-generated AI detector, the tool identified 8% of real publications as machine-generated. The authors concluded their work by affirming the significance of advancing such research directions in the effort of protecting good science<sup>36</sup>.

With this study, we respond to the urgent call to combat fake science by presenting the xFakeSci algorithm. The algorithm is primarily designed to detect ChatGPT-generated content and distinguish it from real PubMed abstracts of published articles. The algorithm operates in two modes: (a) a Single mode, where it is trained from a single source to predict one class, and (b) a Multi mode where the algorithm processes various types of resources to predict the correct label for each class. As a network-driven algorithm, it is trained by a model that utilizes the largest connected components (LCC) as an admissible heuristic. To mitigate overfitting issues, we implemented a data-driven calibration step that ensures the accurate prediction of documents, utilizing a small sample of training articles (100 articles from each source). In the Methods section, we will provide detailed explanations of the following components: (1) the prompt-engineering process of how fake documents were generated in three different diseases using ChatGPT, (2) the phases of testing the premise that content generated from ChatGPT may exhibit certain characteristics that reveal its identity and make it distinguishable from real science, (3) the xFakeSci algorithm, which we designed to predict the class of a given document and determine whether it is real or fake. We outline the computational steps, including constructing the network training models, calibrating using data-driven heuristics, testing the algorithm in multiple contexts (diseases) and data from various publication periods, and finally, (4) benchmarking it against the most common classical data mining algorithms as a verification step.

## Results

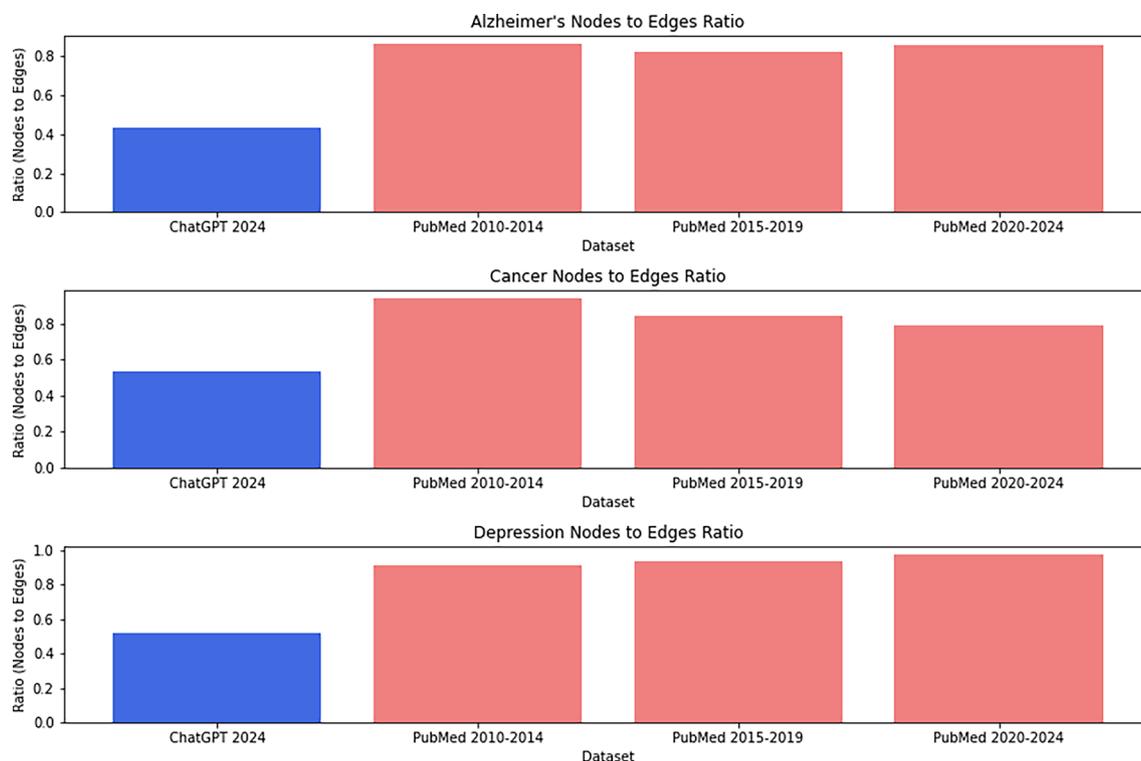
### Outcome of evaluating the premise

As mentioned earlier, we investigated the premise that of whether AI-generated content may exhibit unique characteristics that differ from those observed in scientific articles. We tested this intuition two stages:

#### *Phase I: Analysis of topological properties of network training models*

We constructed two types of network training models: one derived from content generated through prompt-engineering with ChatGPT, and the other from PubMed abstracts. We examined the structural properties of these network models in terms of the number of nodes and edges. These analyses were conducted within the contexts of three diseases: Alzheimer's, cancer, and depression. The node counts computed from ChatGPT training models were 519, 559, and 577, respectively. In contrast, the number of nodes generated from scientific publications varied across different time periods: for the years 2010–2014, it was 742, 755, and 801; for 2015–2019, it was 774, 828, and 755; and for 2020–2024, it was 817, 817, and 790. Regarding edge counts, ChatGPT training models exhibited 1194, 1050, and 1108 edges, whereas publication network models produced 861, 803, and 878 edges for the years 2010–2014; 940, 977, and 826 for 2015–2019; and 958, 1030, and 809 for 2020–2024.

These findings, presented in Table 1, suggest that ChatGPT-generated datasets generally have fewer nodes compared to scientific articles. However, our analysis also revealed that ChatGPT network models tend to have a higher number of edges relative to publication datasets. This observation is visually depicted in Figure 1, which



**Figure 1.** Nodes to edges ratios for different datasets ChatGPT vs scientific articles.

Dataset	Alzheimer's		Cancer		Depression	
	Nodes	Edges	Nodes	Edges	Nodes	Edges
ChatGPT 2024	519	1194	559	1050	577	1108
PubMed 2010–2014	742	861	755	803	802	878
PubMed 2015–2019	774	940	828	977	775	826
PubMed 2020–2024	817	958	817	1030	790	809

**Table 1.** Phase I premise testing: summary of nodes and edges for training models from different sources: PubMed vs ChatGPT datasets.

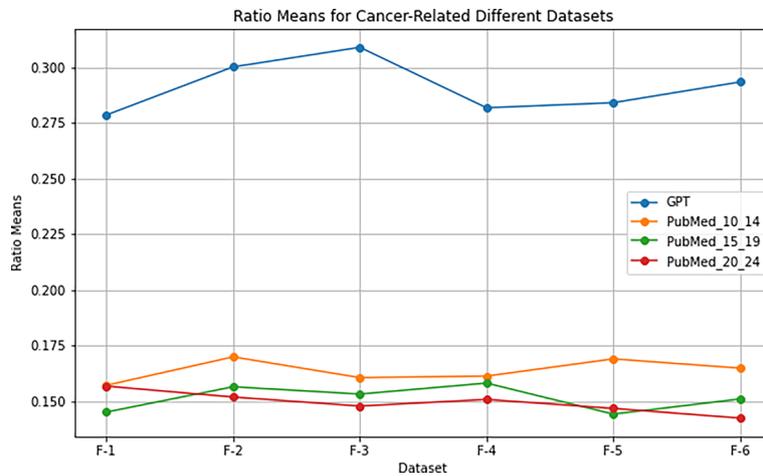
highlights the strikingly lower node-to-edge ratios of ChatGPT models compared to network models derived from scientific articles.

#### Phase II: Further testing the distinctive behavior in ChatGPT-generated documents

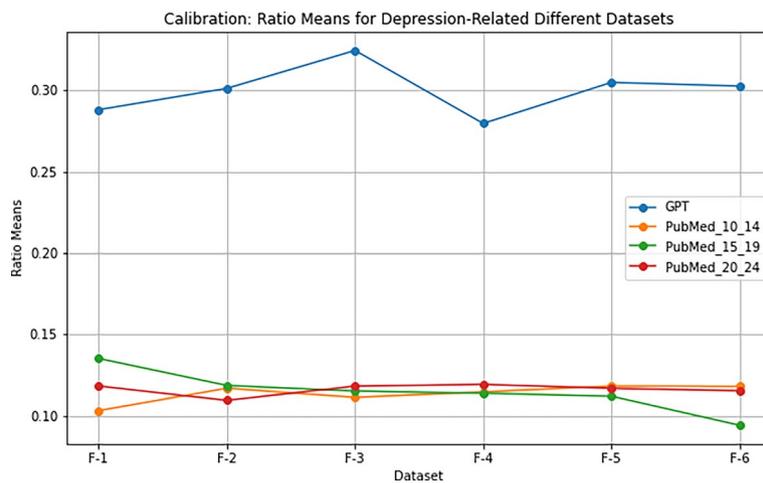
To further investigate the premise, we conducted a test to analyze the mean ratios of contributing bigrams extracted from k-Folds against the document word count. This analysis aimed to establish a baseline for assessing the contribution of bigrams to the overall content structure. The results revealed a consistent pattern across all three disease datasets. Specifically, ChatGPT-generated datasets exhibited significantly higher ratios than their scientific publication counterparts in each of the k-Folds used. For instance, in the Alzheimer's disease dataset, ChatGPT scores were (0.27, 0.30, 0.30, 0.28, 0.28, 0.29), while scientific publications from 2010–2014 scored (0.16, 0.17, 0.16, 0.16, 0.17, 0.16), for 2015–2019 (0.15, 0.16, 0.15, 0.16, 0.14, 0.15), and for 2020–2024 (0.15, 0.15, 0.14, 0.15, 0.14, 0.14). These findings are consistent across the other two diseases, as evident in Table 2. Figures 2 and 3 clearly demonstrate that the k-Folds ratios calculated from ChatGPT-generated data are significantly higher than those derived from scientific publications across different years and scopes. They further illustrate a similar pattern for the cancer and depression datasets. This evidence reinforces the notion that ChatGPT-generated content may exhibit distinct characteristics compared to scientific articles.

#### Outcome of label prediction of multi-mode classification experiments

To establish confidence in our method and ensure consistent performance of the xFakeSci algorithm, we conducted two types of experiments in the subject area of three different diseases. Additionally, we performed experiments to evaluate whether the year of publication plays a role in class prediction. This section presents the outcomes of experiments utilizing ChatGPT-generated text obtained algorithmically using ChatGPT



**Figure 2.** Comparison of calibrating ratio means for the cancer disease.



**Figure 3.** Comparison of calibrating ratio means for the depression disease.

	Alzheimer's				Cancer				Depression			
	GPT	Pub 10-14	15-19	20-24	GPT	10-14	15-19	20-24	GPT	10-14	15-19	20-24
F-1 Means	0.27	0.15	0.14	0.15	0.27	0.14	0.15	0.17	0.28	0.10	0.13	0.11
F-2 Means	0.30	0.16	0.15	0.15	0.29	0.14	0.15	0.17	0.30	0.11	0.11	0.10
F-3 Means	0.30	0.16	0.15	0.14	0.26	0.15	0.16	0.15	0.32	0.11	0.11	0.11
F-4 Means	0.28	0.16	0.15	0.15	0.25	0.14	0.15	0.15	0.27	0.11	0.11	0.11
F-5 Means	0.28	0.16	0.14	0.14	0.25	0.13	0.16	0.15	0.30	0.11	0.11	0.11
F-6 Means	0.29	0.16	0.15	0.14	0.28	0.14	0.16	0.17	0.30	0.11	0.09	0.11

**Table 2.** Phase II premise testing: summary of ratio means for different diseases and datasets.

prompt-engineering, as outlined in Algorithm 1, and scientific publications retrieved from the PubMed web portal<sup>37</sup> related to the Alzheimer's, cancer, and depression diseases.

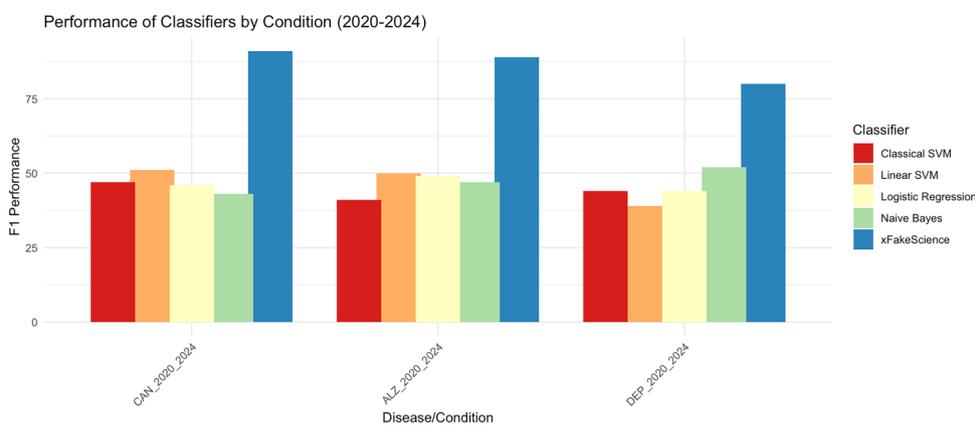
Here, we present the results of multi-mode experiments, where xFakeSci was trained using a combination of ChatGPT and PubMed abstracts and evaluated on a dataset of unseen documents from all three diseases. Specifically, we trained xFakeSci using an equal-sized dataset of ChatGPT-generated and PubMed abstracts. Then, we calibrated the algorithm using the exact number of k-Folds for each disease. For the PubMed dataset, we used abstracts of articles published between 2020 and 2024.

For each disease, we tested xFakeSci on 100 articles, comprising 50 PubMed abstracts and 50 ChatGPT-generated documents. Table 3 summarizes the performance of xFakeSci in this mode, capturing the number of true positives (TP), true negatives (TN), false positives (FP), and false negatives (FN) from both the publication and ChatGPT test items. Using the F1 measure, we note that xFakeSci scored 80%, 91%, and 89% for depression, cancer, and Alzheimer's, respectively. Figure 4 provides a comprehensive analysis of the classification results.

Table 3 demonstrates that xFakeSci detected all 50 PubMed publications for each disease (TP=50). Additionally, we observed that the algorithm identified the ChatGPT-generated documents to varying extents (TN=25, 41, 38) for depression, cancer, and Alzheimer's, respectively. It remains concerning that ChatGPT is classified as PubMed with (FP=25), indicating that 50% of the test documents are misclassified as real publications. Further research is needed to investigate and improve performance.

### Outcome of publication period as a factor for the multi-mode classification experiments

The purpose of this section is to test whether the publication period is a factor in making predictions and assigning labels to a dataset of documents with mixed classes (ChatGPT vs PubMed articles). Here, we use the F1 metric as a measure to present our results and explain the associated performance of our algorithm (captured



**Figure 4.** Multi-Mode Experiments: F1 classification scores for (Cancer, Alzheimer's, and Depression) for publications gathered in period (2020-2024).

Classifier used	Disease dataset tested	TP	FP	FN	TN	F1 Score (%)
xFakeSci	Depression	50	25	0	25	80.00
xFakeSci	Cancer	50	9	0	41	91.74
xFakeSci	Alzheimer's	50	12	0	38	89.29

**Table 3.** Multi-mode experiments: xFakeSci performance evaluation for recent publications (2019- - 2024).

Publication Period (2014–2019)	Depression (%)	Cancer (%)	Alzheimer's (%)
xFakeSci	80	92	90
Naive Bayes	50	43	44
Linear SVM	40	51	48
Classical SVM	44	47	40
Logistic regression	45	47	46
Publication Period (2010–2014)	Depression (%)	Cancer (%)	Alzheimer's (%)
xFakeSci	80	94	89
Naive Bayes	51	46	45
Linear SVM	38	52	48
Classical SVM	42	50	41
Logistic regression	42	50	48

**Table 4.** Multi-mode experiments: F1 classification scores by older periods.

The xFakeSci Algorithms	The SOTA Algorithms
[FP]: gpt_misclassified_as_pubmed 1 0.158730158730	Classifier: Multinomial Naive Bayes
[FP]: gpt_misclassified_as_pubmed 2 0.172413793103	TP: 9, FP: 41, FN: 25, TN: 25
[FP]: gpt_misclassified_as_pubmed 3 0.171052631579	F1 Score: 0.43103448275862066
[FP]: gpt_misclassified_as_pubmed 4 0.163043478261	Classifier: Linear SVM
[FP] [GPT ORIGIN: ] X PUBMED PREDICTED => ratio:	TP: 22, FP: 28, FN: 23, TN: 27
[FP]: gpt_misclassified_as_pubmed 6 0.324561403509	F1 Score: 0.5142857142857142
[FP] [GPT ORIGIN: ] X PUBMED PREDICTED => ratio:	Classifier: Support Vector Machine (SVM)
[FP]: gpt_misclassified_as_pubmed 8 0.231707317073	TP: 20, FP: 30, FN: 25, TN: 25
[FP] [GPT ORIGIN: ] X PUBMED PREDICTED => ratio:	F1 Score: 0.47619047619047616
-----COUNTS-----	Classifier: Logistic Regression
PUBMED CLASSIFIED (TP): 50	TP: 17, FP: 33, FN: 25, TN: 25
CHATGPT CLASSIFIED (TN): 41	F1 Score: 0.46296296296296297
[PUBMED ORIGINAL: ] X GPT PREDICTED : (FN) 0	
[GPT ORIGINAL: ] X PUBMED PREDICTED : (FP) 9	
-----	
F1 Score: 0.9174311926605504	

**Figure 5.** Showing (TP, TN, FP, FN) Metrics and F1 Scores for xFakeSci vs classical data mining algorithms using cancer dataset from 2020–2024.

in Table 4). For each disease dataset extracted from three periods (2020–2024; 2015–2019; 2010–2014), we computed the F1 score.

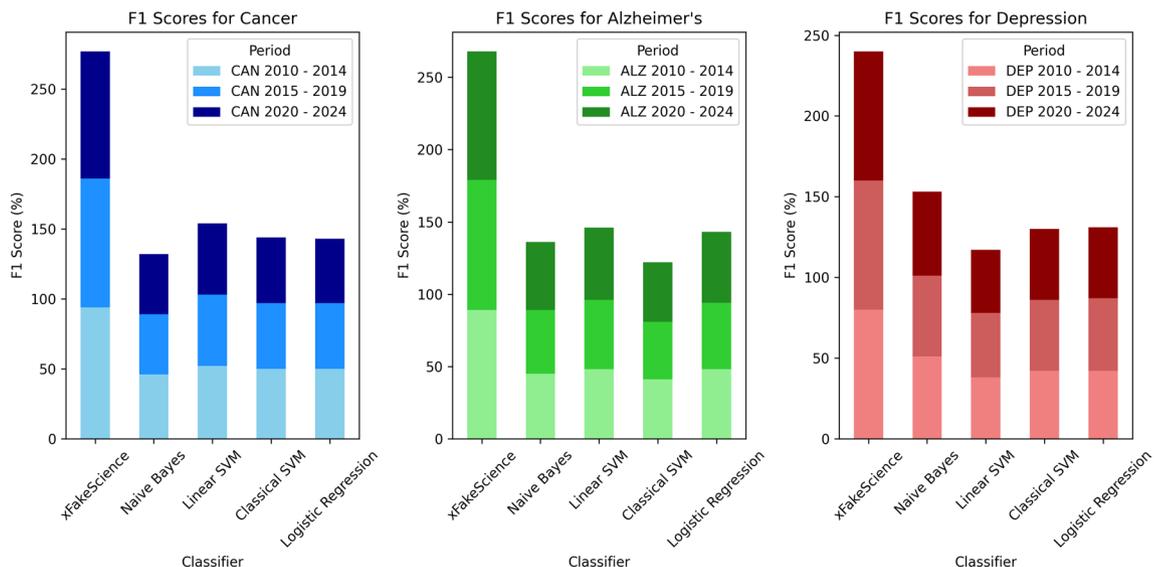
For the cancer disease dataset, the F1 scores recorded were 91%, 92%, and 94% for the three different periods. These scores show a consistent improvement in predicting the labels with older publication datasets, from 2010 to the present.

For the depression disease dataset, the F1 score remained constant over time at 80%, indicating that the pattern did not show deterioration over time.

For the Alzheimer’s disease dataset, the scores showed a slight improvement from 89% (in 2020–2024) to 90% in (2015–2019), but it dropped back to 89% for the (2010–2014) disease dataset. While the pattern of prediction improvements did not hold as we analyzed older publications, the score did not degrade below the F1 score of 89% for (2020–2024).

### Outcome of xFakeSci performance analysis against other data mining algorithms

We compared the performance of xFakeSci against some of the most common state-of-the-art algorithms. Specifically, we conducted various performance evaluation experiments against the following algorithms: (1) Naive Bayes, (2) Support Vector Machine (SVM), (3) Linear Support Vector Machine (SVM), and (4) Logistic Regression. Some of these algorithms are listed among the Top-10 Data Mining Algorithms<sup>38</sup>. To establish fairness, we trained each of the algorithms using the exact training dataset used against xFakeSci in multi-mode (where we train and test with mixed datasets). For training, we used the first 100 PubMed abstracts and the first 100 documents of the ChatGPT-generated dataset. As for testing, we used a combination of 50 PubMed abstracts followed by 50 ChatGPT-generated documents (Fig. 5).



**Figure 6.** Accuracy of xFakeSci vs classical data mining algorithms showing publication year as a prediction factor.

Classifier / publication (2020–2024)	Depression (%)	Cancer (%)	Alzheimer's (%)
xFakeSci	80	91	89
Naive Bayes	52	43	47
Linear SVM	39	51	50
Classical SVM	44	47	41
Logistic regression	44	46	49

**Table 5.** Multi-mode experiments: F1 Classification scores from most recent publications.

Each algorithm was used as a blackbox and received input from two sources (training and test), and in turn, produced a detailed analysis in the form of (TP, TN, FP, FN). Using these metrics, the F1 score was computed accordingly. Figure 6 visually depicts the F1 scores observed over time for the 5 algorithms (including xFakeSci) in three different diseases, presented by sub-figures, one for each disease. Each sub-figure shows stacked bars for the three periods of publications, and each bar represents the F1 score resulting from a given algorithm. Table 5 captures the performance analysis of the xFakeSci algorithm against other classical data mining algorithms (Naive Bayes, Linear SVM, Classical SVM, and Logistic Regression). The performance is presented using the F1 metric for publications spanning between 2020 and 2024. The table shows that xFakeSci scores ranged from 80% to 91%, while those of the data mining algorithms fluctuated between 43% and 52%.

Moving to the period of 2014–2019, xFakeSci demonstrated F1 scores ranging from 80 to 92%, compared to other data mining algorithms which exhibited F1 scores in the range of 43–51%. Lastly, during the period of 2010–2014, the F1 scores achieved by xFakeSci fluctuated between 80 and 94%, whereas the F1 scores of the other data mining algorithms were recorded between 38% and 52% as shown in Table 4. Figure 5 shows a screenshot providing evidence of achieving 91% accuracy measured by the F1 metric, while the F1 scores of the other data mining algorithms fluctuated from 43 to 51%. All three sub-figures show a consistent pattern where xFakeSci clearly outperforms all the other four algorithms. The F1 score is calculated using Eq. 1, and the analysis was done using the scikit-learn library<sup>39</sup>.

$$F1 = \frac{2 * TP}{2 * TP + FP + FN} \quad (1)$$

## Methods

### Data collection

We compiled two distinct types of datasets for the study: (1) **Literature dataset:** To establish a baseline for comparison and train the xFakeSci algorithm, we utilized the PubMed archive to retrieve scientific articles. We employed three search queries: (a) “Alzheimer’s disease and co-morbidities,” which we generated 1196 JSON records, (b) “cancer and co-morbidities,” generated 1243 JSON records, and (c) “depression” which generated 1513 JSON records. To assess the influence of publication year on fake science detection, we conducted these searches at five-year intervals, resulting in three distinct datasets for each disease (2010–2024). (2) **ChatGPT-generated dataset:** We obtained this dataset by programmatically prompting the ChatGPT API (version: 3.5, model: “gpt-3.5-turbo-16k”) to generate simulated articles. The prompt-engineering process comprises two primary components:

- **Prompt Engineering ChatGPT for Simulated Article Generation:** We employed ChatGPT, a generative AI tool, for generating simulated articles. We implemented a prompt-engineering technique to guide the process of text generation in three distinct diseases: Alzheimer’s, cancer, and depression.
- **Predicting Fake Science:** We devised a network-centric learning algorithm, called xFakeSci, which we trained on text documents of published scientific articles and ChatGPT-generated documents. The algorithmic steps involved in this process are explained in the subsequent sections.

### ChatGPT prompt engineering for article generation

We utilized the ChatGPT API (version: 3.5, model: “gpt-3.5-turbo-16k”) to engineer prompts for generating simulated articles in the subject areas of real published articles (Alzheimer’s, cancer, and depression). These prompts were parameterized using information from the real articles (search keywords used for article retrieval and the average number of words) to make them comparable to real abstracts. They included three key elements: (a) the role of the prompt: “a biomedical researcher,” (b) the request example: “Generate a list of 20 simulated PubMed-style abstracts,” (c) topic example: “the Alzheimer’s disease,” and (d) specifications: each article must contain ID, Title, and Abstract fields. We also instructed prompts to generate a valid JSON response with these specifications. The number of words helped offset any bias and made the fake articles comparable to the precise level of detail required (the 200–250 words range is a common requirement by many prominent biomedical informatics journals). Table 6 captures the search queries and the number of fake articles generated in the JSON format.

The prompt-engineering process is computationally described in Algorithm 1. Although this process was done programmatically, due to the timeout limit, we executed it to produce 20 simulated articles at a time. This prompt-engineering approach enabled us to generate a large corpus of simulated articles that closely resembled

Search query	Count (JSON Records)
Alzheimer's disease and co-morbidities	1196
Cancer and co-morbidities	1243
Depression	1513
Total number of generated articles	3952

**Table 6.** Result counts of prompt engineering.

real scientific publications in terms of structure, content, and overall style. This dataset played a crucial role in training the xFakeSci algorithm, enabling it to accurately distinguish between real scientific articles and machine-generated ones.

**Require:** diseases = ["Alzheimer's", "Cancer", "Depression"]

**Require:** *article\_number* = 20.

**Require:** *abstract\_length\_range* = (200-250) of words in each article.

**for** For each disease *d* in diseases **do**

[System Role: ] You are a biomedical researcher specialized in studying *d* disease.

[Request Topic:] disease.

[Request Content:] Generate a list of *article\_number* simulated PubMed-style abstract.

[Request Specs:] Documents must contain: "GPT-ID, Title, and Abstract"; and abstract must be between 200-250 words".

[Request Details:] Provide disease and co-morbidities detailed information.

[Response Format:] A valid JSON format returned as an array of valid JSON records.

**end for**

**Return** an array list of ChatGPT-generated article in JSON format.

**Algorithm 1.** The computational process of ChatGPT prompt engineering for article generation

### Prediction of fake articles using xFakeSci

The xFakeSci algorithm is a network-driven label prediction algorithm designed to distinguish between real scientific articles and machine-generated ones. This entails that the algorithm has two main tasks: training the model and testing it to detect the label of entirely new documents that have never been seen before. In this section, we introduce the computational steps that describe the prediction process, starting with (1) the construction of network training models, (2) the calibration of the algorithm, and (3) the label prediction for each of the ChatGPT-generated articles.

#### Model derivation and network construction

Since our training model is network driven from text, we used the Term Frequency-Inverse Document Frequency (TF-IDF)<sup>40-45</sup> to extract word features as building blocks of the training models. The TF-IDF algorithm can be configured to generate two consecutive words (known as bigrams) that may prove significant across an entire dataset<sup>46,47</sup>. Equation 2 shows the mathematical representation of the TF-IDF, where  $tf(t, d)$  is the frequency of bigram  $t$  in document  $d$ , and  $idf(t, D)$  is the inverse document frequency of bigram  $t$  in the document set  $D$ .

The term frequency of a term  $t$  is calculated as the ratio between the number of occurrences of the term divided by the total number of terms in a document  $d$  as show in Eq. 3. The inverse document frequency ( $idf(t, D)$ ) is calculated as in Eq. 4 where  $N$  is the total number of documents in the collection and  $df(t, D)$  is the number of documents containing the bigram  $t$ .

$$TF-IDF(t, d, D) = tf-idf(t, d, D) = tf(t, d) \times idf(t, D) \quad (2)$$

$$TF(t, d) = \frac{\text{Number of occurrences of term } t \text{ in document } d}{\text{Total number of terms in document } d} \quad (3)$$

$$IDF(t, D) = \log \left( \frac{N}{1 + df(t, D)} \right) \quad (4)$$

To construct a training model, we extracted bigrams to form a network model as follows: the individual words of a bigram served as nodes, and edges represented the relationship between the bigrams. To illustrate the utility of bigrams in constructing the training model, let's consider a scenario concerning the "depression disease": bigrams such as "mental health," "health condition," and "condition worsen" form connections based on the common words they share, enabling a network that can be analyzed for various purposes. Using this mechanism, we constructed two distinct training models: one from the abstracts of published literature (labeled as the "PUBMED" class), and another from ChatGPT-generated text (labeled as the "GPT" class). To ensure fairness and prevent biases, both models were constructed from the same number of documents (100 abstracts and 100

ChatGPT-generated). Both datasets were processed using identical series of steps, including removing stopwords and sentence tokenization.

Algorithm 2 outlines the steps involved in building the network model from bigrams. We applied this algorithm twice, once to create a publication training model and another to create a ChatGPT model. We recorded the corresponding statistics (numbers of nodes and edges) for each model in Table 1. The initial observations revealed a consistent pattern where models constructed from ChatGPT-generated text exhibited the lowest number of nodes, yet they also maintained the highest number of edges. The resulting models showed disconnected components and fragmented communities, requiring pruning. This need was satisfied by applying the Largest Connected Components (LCC) algorithm<sup>48</sup>, which ensured that the resulting networks maintained high connectivity. The LCC presents an admissible pruning heuristic due to the presence of high-degree nodes that promote network stability and robustness<sup>49–52</sup>.

---

**Require:**

$\mathcal{D}$ : [document dataset (PubMed or ChatGPT)]

$\mathcal{T}$ : training graph

$\mathcal{B}$ : empty list of bigrams

**Ensure:**

$\mathcal{G}$ : fully populated graph

```

1: for each document  $d$  in  $\mathcal{D}$  dataset (either from ChatGPT or PubMed) do
2:    $B \leftarrow$  Compute Term Frequency - Inverse Term Frequency  $\text{tf-idf}(d)$ 
3:   for each bigram  $b$  in  $\mathcal{B}$  encoded as source and target  $b(s,t)$  do
4:     Form an edge  $e(s,t)$  for the unigram constituents of the bigram  $b$ .
5:     if  $e$  does not exist in the training graph:  $\mathcal{T}$  then
6:       Add  $e$  to the training model  $\mathcal{T}$ .
7:     end if
8:   Reset  $B \leftarrow []$  to populate with next document
9:   end for
10: end for
11: return  $\mathcal{T}$  a fully populated graph from dataset bigrams

```

---

**Algorithm 2.** Network construction from bigrams computed from text documents (PubMed, and ChatGPT)

### Evaluating the premise of ChatGPT's distinctive behavior

As mentioned earlier, the training models are constructed from the first 100 articles from each dataset. To test the premise of how ChatGPT may exhibit distinctive behavior, we divided the remaining articles into  $k$ -Folds, each containing 100 articles. The main idea of such a test was to measure the impact of each fold on the corresponding training model, specifically how the bigrams extracted from each of the folds altered the Largest Connected Components (LCCs) of their respective data types.

For the Alzheimer's disease dataset, we constructed three training models: the impact of bigrams was determined by calculating the mean ratio between the number of bigrams contributing to the LCC and the total number of words in each article within a fold. This process is captured by measuring the average contribution rate of the bigrams of a given fold. Algorithm 3 provides the pseudocode for this step.

---

**Require:** Training model  $M$ ,

**Require:** One fold of 100 articles from  $k$ -folds,

**Require:** Computed ratios of each fold  $R \leftarrow []$

▷ Initialize as an empty list

**Ensure:** Fold mean

```

1: for each document  $d$  in  $D$  do
2:    $\text{doc\_ratio} \leftarrow 0.0f$  ▷ Initialize using a value of 0.0f
3:    $\text{bigram\_count} \leftarrow 0$  ▷ Initialize using a value of 0
4:    $\text{doc\_wc} \leftarrow \text{compute\_word\_count}(d)$ 
5:    $B \leftarrow$  Compute bigrams using TF-IDF
6:   for each bigram  $b$  in  $B$  do
7:     if  $b$  contributes an edge to the corresponding training model  $M$  then
8:        $\text{bigram\_count} += 1$  ▷ Increment the count by 1
9:     end if
10:   end for
11:    $\text{doc\_ratio} \leftarrow \frac{\text{bigram\_count}}{\text{doc\_wc}}$ 
12:    $R.\text{append}(\text{doc\_ratio})$  ▷ Keep track of the ratios
13: end for
14:  $\text{fold\_mean} \leftarrow$  compute the mean of the ratios  $\text{mean}(R)$ 
15: return  $\text{fold\_mean}$ 

```

---

**Algorithm 3.** Compute the mean of ratios for a corresponding fold

We summarized the analysis of each dataset and disease in Table 2. The initial observations indicated that the ChatGPT ratios fluctuated between 27 and 30% for the Alzheimer's disease, 27 and 29% for cancer, and 28 and 32% for depression. In contrast, the ratios derived from scientific articles ranged between 14 and 16% for the

	Alzheimer's				Cancer				Depression			
	GPT	Pub 10-14	15-19	20-24	GPT	10-14	15-19	20-24	GPT	10-14	15-19	20-24
Lower	0.27	0.15	0.14	0.14	0.25	0.13	0.15	0.15	0.28	0.10	0.09	0.10
Upper	0.30	0.16	0.15	0.15	0.29	0.15	0.16	0.17	0.32	0.11	0.13	0.11

**Table 7.** Model calibration: summary of lower and upper bound ranges for different diseases and datasets by year periods: 2010–2014, 2015–2019, 2020–2024.

Alzheimer's disease, 14 and 17% for cancer, and 9 and 11% for depression. The full analysis of these results will be discussed in the “Results” section. The analysis demonstrated that the ratios of ChatGPT-generated documents were significantly different from those computed from scientific articles. This distinction serves as proof that the premise is indeed true. Furthermore, this knowledge provides lower and upper bounds for each disease, offering more guidance to the algorithm to predict the label while avoiding issues of overfitting. Algorithm 4 demonstrates how the ratios are computed. For brevity, we only demonstrate for the depression disease. Table 7 presents the corresponding lower and upper bounds, which are necessary during the calibration phase.

---

```

1: procedure COMPUTERANGES( $R, R'$ )
2:   Input: List of ChatGPT k-Fold ratios  $R: [r_0, r_1, r_2, \dots, r_k]$ 
3:   Input: List of PubMed k-Fold ratios  $R': [r'_0, r'_1, r'_2, \dots, r'_k]$ 
4:   Output: Dictionary of key/value pairs where the key is the name of the dataset, and the value is the lower/upper bound
      range.
5:    $gpt\_ranges \leftarrow []$ 
6:   for each ratio in  $R$  do
7:      $gpt\_min \leftarrow \min(R)$ 
8:      $gpt\_max \leftarrow \max(R)$ 
9:      $gpt\_lower\_upper \leftarrow \text{range}(gpt\_min, gpt\_max)$ 
10:     $gpt\_ranges.append(gpt\_lower\_upper)$ 
11:  end for
12:   $pubmed\_ranges \leftarrow []$ 
13:  for each ratio in  $R'$  do
14:     $pubmed\_min \leftarrow \min(R')$ 
15:     $pubmed\_max \leftarrow \max(R')$ 
16:     $pubmed\_lower\_upper \leftarrow \text{range}(pubmed\_min, pubmed\_max)$ 
17:     $pubmed\_ranges.append(pubmed\_lower\_upper)$ 
18:  end for
19:  return [ $gpt\_ranges, pubmed\_ranges$ ]
20: end procedure

```

---

**Algorithm 4.** Computing the calibrating ratio ranges from k\_Folds

*Label prediction of articles: real vs. fake*

Testing the premise in the above section demonstrated the fundamental differences in content behavior between fake ChatGPT articles and real publications. In this section, we present the learning algorithm, which is the main contribution of this paper. During the Coronavirus global pandemic, our previous work addressed the challenge of detecting fake news and science as an emerging infodemic<sup>15</sup>. However, this work was limited by the lack of comprehensive machine-generated datasets that could adequately assess performance in the presence of fake data. Now, with the advent of ChatGPT and generative AI technologies, we can generate diverse datasets using prompt-engineering algorithms as demonstrated above. Additionally, the previous work did not make use of data-driven insights, which we incorporate as a calibration step. This is an intermediate phase that takes place after the training phase and before the label prediction phase. Due to these factors, the previous work was limited to a single-mode label prediction using a single type of dataset. Therefore, it was necessary to split the dataset into training and test sets. The following is a complete comparison of the previous work and the current features presented by xFakeSci, such as content type, configuration parameters, classification mode, calibration, and classification, as presented in Table 8.

As shown in Table 8, the xFakeSci algorithm is particularly designed to address multi-mode classification. Therefore, it is expected to train the algorithm using two or more independent types of data. Consequently, the algorithm also expects a hybrid test set of mixed types and will produce more accurate labels for each type. However, such modes suffer from what is known as the “overfitting” issue<sup>53–56</sup>. The introduction of the calibration step (by calculating the lower/upper bound ratios captured in Table 7) was to guide the decision of the final label prediction and avoid such an issue. The table demonstrates a clear separation of lower/upper bound ratios. Therefore, we further utilize such a mechanism by incorporating a calibrating step to further guide the classification process without having to train the algorithm with too many samples. The algorithmic steps for the calibration process are explained in Algorithm 4. Though the ranges provide an extra net for predicting the label, it is also possible that some document instances may fall outside the specified ranges of the datasets,

Feature	NeoNet algorithm	xFakeSci algorithm
Content type	Publication/news	publications/ChatGPT fake articles
Calibration	Not-supported	data-driven calibration
Configuration Parameters	Min-support and min-confidence	lower-bound/higher-bound ratio means
Distance	No supported	uses proximity distance for out-of-range instances
Classification	Single-mode classification (news against news)	single/multi-mode classification (fake against real)
Classification/Year	Not-supported	comparative classification by the year of publication
Multi-Mode Performance	Not-supported	F1 score 80%-94%

**Table 8.** Comparison between NeoNet and xFakeSci algorithms.

which could result in not predicting a label correctly. Therefore, we introduced a proximity heuristic that favors the shortest distance to the ranges driven from the individual datasets (real or ChatGPT-generated) and assigns a label accordingly. Eqs. 5 and 6 demonstrate how the distance is calculated.

$$\text{distance} = \min(|\text{point} - \text{range}_{\text{end}}|, |\text{point} - \text{range}_{\text{start}}|) \quad (5)$$

$$\text{proximity} = \min(\text{distance}_d, \text{distance}_{d'}) \quad (6)$$

Algorithm 5 illustrates the computational steps for multi-mode execution, demonstrating the complexity involved, including the proximity distance. To use the algorithm in detecting fake science, it must be trained using two different types of data: (1) a real publication dataset and (2) ChatGPT-generated articles. The algorithm also expects the ratio means of each data source, which are computed using the calibration algorithm.

---

```

1: Input: [ChatGPTModel  $G$ , PubMedModel  $P$ ]
2: Input: [ChatGPT – ratios  $\text{range}_1$ , PubMed – ratios  $\text{range}_2$ ]
3: Input: a dataset  $D$  of mixed fake and real documents
4: for each document  $d$  in  $D$  dataset do
5:    $\text{ratio}_1 \leftarrow \text{compute\_model\_contribution}(d, G)$ 
6:    $\text{ratio}_2 \leftarrow \text{compute\_model\_contribution}(d, P)$ 
7:   if  $\text{ratio}_1$  in  $\text{range}_1$  then
8:     assign GPT as a label to document  $d$ 
9:   else if  $\text{ratio}_2$  in  $\text{range}_2$  then
10:    assign PubMed as a label to document  $d$ 
11:   else
12:      $\text{dist}_{\text{range}_1} \leftarrow \text{compute\_distance\_to\_range}(\text{ratio}_1, \text{range}_1)$ 
13:      $\text{dist}_{\text{range}_2} \leftarrow \text{compute\_distance\_to\_range}(\text{ratio}_2, \text{range}_2)$ 
14:     if  $\text{dist}_{\text{range}_1} < \text{dist}_{\text{range}_2}$  then
15:       assign GPT as a label to article  $d$ 
16:     else if  $\text{dist}_{\text{range}_2} < \text{dist}_{\text{range}_1}$  then
17:       assign PubMed as a label to article  $d$ 
18:     end if
19:   end if
20:   capture stats  $\leftarrow$  score prediction
21: end for
22: return stats

```

---

**Algorithm 5.** Multi-Mode execution of the xFakeBib algorithm

## Discussion

In a world where generative AI has become widespread, various studies aimed to investigate the potential issues of using ChatGPT to generate fake science. The literature review showed a desperate need to advance the algorithmic approaches to discern real publications from fake ones, especially, when they are mixed. Our study aimed to address such issues incrementally. Specifically, we first tested the intuition of whether the content generated by ChatGPT may exhibit unique characteristics that distinguish it from real science. We explored this task using prompt engineering, where we created engineered datasets on the subjects of Alzheimer’s, cancer, and depression diseases. In this work, we contributed a prompt-engineering algorithm on how to generate simulated content to evaluate this premise. Working with plain text (using publication abstracts or generated from ChatGPT), using the TF-IDF algorithm is a common approach to generate bigrams that can be used to construct more complex models.

Our initial observation of networks generated from ChatGPT content is that they are highly connected and contain fewer nodes compared to networks constructed from real publication text. Additionally, when we calculated the ratios of the number of bigrams against the total number of words of documents on k-Folds,

we found that the ratios of ChatGPT content are much higher than scientific abstracts. These two indications supported our intuition that ChatGPT documents exhibit distinguishable behavior than PubMed abstracts. One interpretation of this observation could be due to the inherent design of the ChatGPT engine. As observed, ChatGPT is optimized to generate highly convincing content by predicting the next correlated terms statistically using a Large Language Model. On the other hand, scientists prioritize accurate documentation of hypotheses testing, scientific experiments, and careful explanation of observations. Describing science in terms of highly correlated words is not a goal of scientists. Clearly, the difference in goals may contribute to less connectivity in scientific publications.

Further, we introduced the xFakeSci algorithm, a learning algorithm that predicts a label for a given article. In the Methods section, we showed that it is designed to operate in two modes: (1) Single-mode: where only one type of articles from the same source is used for training and a new set of documents from the same pool is used for predicting the label of an article; and (2) Multi-mode: where the algorithm was trained from two sources and a hybrid train model (of real and generated datasets) was constructed to make the predictions. The single-model is trivial; therefore, we focused our experiments on demonstrating the multi-mode. We performed several experiments to do the following: (1) to test and measure how the xFakeSci algorithm predicts labels of ChatGPT generated documents for a given disease when mixed with scientific abstracts, (2) to evaluate whether the algorithm performs consistently using various datasets of different diseases not only one disease, (3) to test whether the year of the publication plays a role in predicting ChatGPT generated documents when mixed with publications from various periods (2020–2024, 2015–2019, 2010–2014), and (4) to benchmark the algorithm against a baseline of some of the most common data mining algorithms. Our results for each experiment used the TP, TN, FP, FN metrics and F1 scores.

When testing whether the year of publication plays a role in label prediction, we observed F1 scores of 91%, 92%, and 94% for cancer-related publications across different periods. This suggests a pattern of better detection of ChatGPT articles when mixed with older publications. However, identifying newer publications proved more challenging. For the Alzheimer's disease, while no improvement was observed, degradation was also absent. As mentioned earlier, the Alzheimer's datasets were the smallest among all datasets, limiting the calibration process due to fewer k-Folds compared to other diseases. In the case of depression, the algorithm exhibited consistent performance with an F1 score of 80% across all periods. It's plausible that mental health data acquisition posed limitations, potentially constraining resources from this specific area. Testing this hypothesis involves measuring document similarity between PubMed and ChatGPT sources using lexical and semantic analysis.

Upon benchmarking xFakeSci against classical data mining algorithms, we observed an interesting pattern: xFakeSci correctly predicted all the scientific publications in all the experiments we performed. However, other algorithms misclassified publications as ChatGPT and vice versa (true positives, false positives, false negatives, and true negatives). xFakeSci, however, needed improvement in predicting true negatives (ChatGPT documents), as many ChatGPT documents were labeled as true positives (real publications). In all the experiments, the F1 scores of xFakeSci ranged between 80 and 94%. In contrast, the other data mining algorithms showed much lower performance, with F1 scores ranging between 32 and 52%. We attribute the high performance of xFakeSci to the calibration process, which was guided by ratios and proximity distances. Although the training model remained lightweight, both heuristics provided more guidance for predicting fake articles. This novel calibration method benefits from an abundance of data, without suffering from overfitting issues like other common classification algorithms. Clearly, the xFakeSci algorithm does not suffer from such a deficiency in identifying real articles when mixed with ChatGPT-generated content

While xFakeSci is designed to distinguish fake science from real, it can be applied to various types of text data, including clinical notes, clinical trial summaries, and interventions. With the widespread adoption of generative AI tools such as ChatGPT and Google Bard, ethical concerns may arise, such as clinicians using ChatGPT to generate clinical notes, potentially resulting in erroneous entries with serious consequences. In such cases, our algorithm may serve as a forensic tool to identify potentially fake portions of these reports.

While we have highlighted the potential for harm posed by ChatGPT and similar tools, it is also important to recognize their positive generative capabilities. For instance, ChatGPT played a crucial role in providing our algorithm with simulated data, which was essential for our work during the global pandemic in detecting fake news and publications<sup>15</sup>. Moreover, ChatGPT can generate code snippets as building blocks for various basic tasks, including data visualization, across diverse programming languages. We are currently exploring this capability to construct workflows for life sciences applications. Additionally, the ChatGPT engine can effectively convert semi-structured content into popular formats like JSON, XML, and others. While these capabilities are undoubtedly useful, they necessitate the development of ethical standards to ensure responsible use of such tools.

Another intriguing potential use is that, when creatively engineered, ChatGPT could function as a valuable teaching assistant for academics and school teachers. It could potentially generate various ways to present questions while maintaining the integrity of the original content. Furthermore, ChatGPT could revolutionize scientific writing by providing support in addressing grammatical errors, typography, and paraphrasing, particularly for those whose native language is not English<sup>57</sup>.

## Conclusions and future directions

When we asked a high school student about their knowledge of ChatGPT, they responded, “Do you mean that tool that does my homework for you?” Indeed, ChatGPT is an incredibly sophisticated tool with a wide range of impressive capabilities. Since the rise of ChatGPT, many new research topics have opened a new generative door, and many long-standing questions are now being investigated. However, the most significant concern associated with ChatGPT and other generative AI tools is that they could pose a threat to the future of science. If younger

generations utilize ChatGPT to plagiarize, it could undermine the integrity of research and learning, potentially having a negative impact on the development of future pioneers.

While learning algorithms, such as xFakeSci, can assist in identifying fake science, there is an ethical obligation to use generative AI tools responsibly and regulate their usage<sup>16</sup>. It is worth noting that certain countries, such as Italy, have taken the extreme step of banning ChatGPT. While the authors believe such measures may be drastic, addressing ethical concerns is a new frontier that must be tackled. As ChatGPT itself states, “It is up to individuals and organizations to use technology like mine in ways that promote positive outcomes and minimize any potential negative impacts.” As advised by Anderson et al., it is also the responsibility of publishers and those involved in the production of science to play a proactive role in promoting good science. This includes raising awareness of the importance of implementing advanced fake science detection algorithms, including ours, and activating the use of technologies to distinguish fake research and fabricated findings<sup>35</sup>.

Looking ahead, there are several avenues for future research based on our current work: (1) conducting a preprocessing step (e.g., clustering) to group more closely related publications together (e.g., breast cancer, prostate cancer, and others), or separate diseases from co-morbidities. The use of knowledge graphs may be a powerful tool to use in continuing to investigate this research direction; (2) further experimentation in training and calibrating the xFakeSci algorithm by utilizing heuristics learned from preprocessing steps and the discoveries of clusters; and (3) testing the algorithm on more than two data sources (clinical reports, publications, and ChatGPT-generated documents).

### Code and Data availability

Both code and the dataset required for executing the algorithm are available in the xFakeSci GitHub repository available at “[<https://github.com/drahmedabdeenamed/xFakeSci>]”. Due to the file size limitation that can be uploaded on GitHub, some of the data files could not be uploaded, and must be requested from the authors if needed.

Received: 11 July 2023; Accepted: 3 July 2024

Published online: 14 July 2024

### References

1. Chatgpt. Online: <https://chat.openai.com> (2023). Accessed 15 Aug 2023.
2. Synnestevedt, M. B., Chen, C. & Holmes, J. H. Citespace ii: visualization and knowledge discovery in bibliographic databases. In *AMIA annual symposium proceedings*, vol. 2005, 724 (American Medical Informatics Association, 2005).
3. Holzinger, A. et al. On graph entropy measures for knowledge discovery from publication network data. In *Availability, Reliability, and Security in Information Systems and HCI: IFIP WG 8.4, 8.9, TC 5 International Cross-Domain Conference, CD-ARES 2013, Regensburg, Germany, September 2-6, 2013. Proceedings 8*, 354–362 (Springer, 2013).
4. Usai, A., Pironti, M., Mital, M. & Aouina Mejri, C. Knowledge discovery out of text data: a systematic review via text mining. *J. Knowl. Manag.* **22**, 1471–1488 (2018).
5. Thaler, A. D. & Shiffman, D. Fish tales: Combating fake science in popular media. *Ocean Coastal Manag.* **115**, 88–91 (2015).
6. Hopf, H., Krief, A., Mehta, G. & Matlin, S. A. Fake science and the knowledge crisis: ignorance can be fatal. *Royal Soc. Open Sci.* **6**, 190161 (2019).
7. Ho, S. S., Goh, T. J. & Leung, Y. W. Let’s nab fake science news: Predicting scientists’ support for interventions using the influence of presumed media influence model. *Journalism* **23**, 910–928 (2022).
8. Frederickson, R. M. & Herzog, R. W. Addressing the big business of fake science. *Molecular Therapy* **30**, 2390 (2022).
9. Rocha, Y. M. et al. The impact of fake news on social media and its influence on health during the covid-19 pandemic: A systematic review. *J. Public Health* **31**, 1–10 (2021).
10. Walter, N., Brooks, J. J., Saucier, C. J. & Suresh, S. Evaluating the impact of attempts to correct health misinformation on social media: A meta-analysis. *Health Commun.* **36**, 1776–1784 (2021).
11. Loomba, S., de Figueiredo, A., Piatek, S. J., de Graaf, K. & Larson, H. J. Measuring the impact of covid-19 vaccine misinformation on vaccination intent in the uk and usa. *Nat. Human Behav.* **5**, 337–348 (2021).
12. Lewandowsky, S., Ecker, U. K., Seifert, C. M., Schwarz, N. & Cook, J. Misinformation and its correction: Continued influence and successful debiasing. *Psychol. Sci. Public Interest* **13**, 106–131 (2012).
13. Myers, M. & Pineda, D. Misinformation about vaccines. *Vaccines for biodefense and emerging and neglected diseases* 255–270 (2009).
14. Matthews, S. & Spencer, B. Government orders review into vitamin d’s role in covid-19. Online: <https://www.dailymail.co.uk/news/article-8432321/Government-orders-review-vitamin-D-role-Covid-19.html> (2020). Accessed on 13 Apr 2024.
15. Abdeen, M. A., Hamed, A. A. & Wu, X. Fighting the covid-19 infodemic in news articles and false publications: The neonet text classifier, a supervised machine learning algorithm. *Appl. Sci.* **11**, 7265 (2021).
16. Hamed, A. A., Zachara-Szymanska, M. & Wu, X. Safeguarding authenticity for mitigating the harms of generative ai: Issues, research agenda, and policies for detection, fact-checking, and ethical ai. *iScience* **27**, 108782. <https://doi.org/10.1016/j.isci.2024.108782> (2024).
17. Eysenbach, G. et al. The role of chatgpt, generative language models, and artificial intelligence in medical education: A conversation with chatgpt and a call for papers. *JMIR Med. Edu.* **9**, e46885 (2023).
18. IEEE special issue on education in the world of ChatGPT and other generative AI. Online: <https://iee-edusociety.org/iee-special-issue-education-world-chatgpt-and-other-generative-ai> (2023). Accessed 13 Apr 2024.
19. Financial innovation. Online: <https://jfin-swufe.springeropen.com/special-issue---chatgpt-and-generative-ai-in-finance> (2023). Accessed 13 Apr 2024.
20. Special issue “language generation with pretrained models”. Online: [https://www.mdpi.com/journal/languages/special\\_issues/K1Z08ODH6V](https://www.mdpi.com/journal/languages/special_issues/K1Z08ODH6V) (Year). Accessed 13 Apr 2023.
21. Call for papers for the special focus issue on ChatGPT and large language models (LLMs) in biomedicine and health. <https://academic.oup.com/jamia/pages/call-for-papers-for-special-focus-issue> (Year). Accessed 4 July 2023.
22. Leung, T. I., de Azevedo Cardoso, T., Mavragani, A. & Eysenbach, G. Best practices for using ai tools as an author, peer reviewer, or editor. *J. Med. Internet Res.* **25**, e51584. <https://doi.org/10.2196/51584> (2023).
23. Null, N. The PNAS journals outline their policies for ChatGPT and generative AI. *PNAS Updates* <https://doi.org/10.1073/pnas-updates.2023-02-21> (2023).
24. Brainard, J. As scientists explore ai-written text, journals hammer out policies. *Science* **379**, 740–741 (2023).

25. Fuster, V. *et al.* Jacc journals' pathway forward with ai tools: The future is now. *JACC: Adv.* **2**, 100296. <https://doi.org/10.1016/j.jacadv.2023.100296> (2023).
26. Flanagan, A., Bibbins-Domingo, K., Berkwitz, M. & Christiansen, S. L. Nonhuman “authors” and implications for the integrity of scientific publication and medical knowledge. *Jama* **329**, 637–639 (2023).
27. Chatgpt plugins. Online: <https://openai.com/blog/chatgpt-plugins> (2023). Accessed 13 Apr 2023.
28. Gilson, A. *et al.* How does chatgpt perform on the united states medical licensing examination? the implications of large language models for medical education and knowledge assessment. *JMIR Med. Edu.* **9**, e45312 (2023).
29. Chaka, C. Detecting ai content in responses generated by chatgpt, youchat, and chatsonic: The case of five ai content detection tools. *J. Appl. Learn. Teac.* <https://doi.org/10.37074/jalt.2023.6.2.12> (2023).
30. Vapnik, V. N. An overview of statistical learning theory. *IEEE Trans. Neural Netw.* **10**, 988–999 (1999).
31. Cingilloglu, I. Detecting ai-generated essays: the chatgpt challenge. *Int. J. Inf. Learn. Technol.* **40**, 259–268 (2023).
32. Copyleaks: AI & machine learning powered plagiarism checker. Online: <https://copyleaks.com/>. Accessed 13 Apr 2024.
33. Crossplag: Online plagiarism checker. Online: <https://crossplag.com/>. Accessed 13 Apr 2024.
34. Elkhatat, A. M., Elsaid, K. & Almeer, S. Evaluating the efficacy of ai content detection tools in differentiating between human and ai-generated text. *Int. J. Edu. Integrity* **19**, 17 (2023).
35. Anderson, N. *et al.* Ai did not write this manuscript, or did it? can we trick the ai text detector into generated texts? the potential future of chatgpt and ai in sports & exercise medicine manuscript generation. *BMJ Open Sport Exercise Med.* <https://doi.org/10.1136/bmjsem-2023-001568> (2023).
36. Rashidi, H. H., Fennell, B. D., Albahra, S., Hu, B. & Gorbett, T. The chatgpt conundrum: Human-generated scientific manuscripts misidentified as ai creations by ai text detection tool. *J. Pathol. Inf.* **14**, 100342 (2023).
37. NLM, N. L. o. M. National center of biotechnology information. Online: <https://pubmed.ncbi.nlm.nih.gov/>. Accessed on 25 Jan 2024.
38. Wu, X. *et al.* Top 10 algorithms in data mining. *Knowl. Inf. Syst.* **14**, 1–37 (2008).
39. Pedregosa, F. *et al.* Scikit-learn: Machine learning in Python. *J. Machine Learn. Res.* **12**, 2825–2830 (2011).
40. Aizawa, A. An information-theoretic perspective of tf-idf measures. *Inf. Process. Manag.* **39**, 45–65 (2003).
41. Qaiser, S. & Ali, R. Text mining: use of tf-idf to examine the relevance of words to documents. *Int. J. Comput. Appl.* **181**, 25–29 (2018).
42. Ramos, J. *et al.* Using tf-idf to determine word relevance in document queries. In *Proceedings of the first instructional conference on machine learning*, vol. 242, 1, 29–48 (Citeseer, 2003).
43. Trstenjak, B., Mikac, S. & Donko, D. Knn with tf-idf based framework for text categorization. *Proc. Eng.* **69**, 1356–1364 (2014).
44. Wu, H. C., Luk, R. W. P., Wong, K. F. & Kwok, K. L. Interpreting tf-idf term weights as making relevance decisions. *ACM Trans. Inf. Sys. (TOIS)* **26**, 1–37 (2008).
45. Zhang, W., Yoshida, T. & Tang, X. A comparative study of tf<sup>α</sup> idf, lsi and multi-words for text classification. *Expert Syst. Appl.* **38**, 2758–2765 (2011).
46. Tan, C.-M., Wang, Y.-F. & Lee, C.-D. The use of bigrams to enhance text categorization. *Inf. Process. Manag.* **38**, 529–546 (2002).
47. Hirst, G. & Feiguina, O. Bigrams of syntactic labels for authorship discrimination of short texts. *Literary Linguistic Comp.* **22**, 405–417 (2007).
48. Dorogovtsev, S. N., Mendes, J. F. F. & Samukhin, A. N. Giant strongly connected component of directed networks. *Phys. Rev. E* **64**, 025101 (2001).
49. Kitsak, M. *et al.* Stability of a giant connected component in a complex network. *Phys. Rev. E* **97**, 012309 (2018).
50. Beygelzimer, A., Grinstein, G., Linsker, R. & Rish, I. Improving network robustness by edge modification. *Phys. A Stat. Mechan. Appl.* <https://doi.org/10.1016/j.physa.2005.03.040> (2005).
51. Zhang, G., Duan, H. & Zhou, J. Network stability, connectivity and innovation output. *Technol. Forecast. Soc. Change* <https://doi.org/10.1016/j.techfore.2016.09.004> (2017).
52. Bellingeri, M. *et al.* Link and node removal in real social networks: A review. *Front. Phys.* <https://doi.org/10.3389/fphy.2020.00228> (2020).
53. Genkin, A., Lewis, D. D. & Madigan, D. Large-scale bayesian logistic regression for text categorization. *Technometrics* **49**, 291–304 (2007).
54. Feng, X. *et al.* Overfitting reduction of text classification based on adablm. *Entropy* **19**, 330 (2017).
55. Deng, X., Li, Y., Weng, J. & Zhang, J. Feature selection for text classification: A review. *Multimed. Tools Appl.* **78**, 3797–3816. <https://doi.org/10.1007/s11042-018-6083-5> (2019).
56. Khurana, A. & Verma, O. P. Optimal feature selection for imbalanced text classification. *IEEE Trans. Artif. Intell.* **4**, 135–147. <https://doi.org/10.1109/TAL.2022.3144651> (2023).
57. Conroy, G. How chatgpt and other ai tools could disrupt scientific publishing. *Nature* **622**, 234–236 (2023).

## Acknowledgements

This research is supported by the European Union's Horizon 2020 research and innovation programme under grant agreement Sano No 857533 which is carried out within the International Research Agendas programme of the Foundation for Polish Science, co-financed by the European Union under the European Regional Development Fund, and created as part of the Ministry of Science and Higher Education's initiative to support the activities of Excellence Centers established in Poland under the Horizon 2020 program based on the agreement No “MEiN/2023/DIR/3796”, and the National Natural Science Foundation of China (NSFC) under grant 62120106008. The authors also acknowledge Laila Hamed for her valuable perspective on ChatGPT.

## Author contributions

A.H. conceived the idea(s), A.A. and X.W. designed the experiment(s), A.A. and X.W. analyzed the results. Both authors written and reviewed the manuscript.

## Competing interests

The authors declare no competing interests.

## Additional information

**Correspondence** and requests for materials should be addressed to A.A.H. or X.W.

**Reprints and permissions information** is available at [www.nature.com/reprints](http://www.nature.com/reprints).

**Publisher's note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2024