

DOI: https://doi.org/10.48009/3_iis_2026_111

Authorship attribution of AI-generated vs. human-written news articles using machine learning

John C. Stewart, *Robert Morris University*, stewartj@rmu.edu

Gia Bao Nguyen, *Robert Morris University*, gxnst757@mail.rmu.edu

G. Alan Davis, *Robert Morris University*, davis@rmu.edu

Robert Ambrosini, *Robert Morris University*, ambrosini@rmu.edu

Abstract

In this study, news articles from journalists were compared to AI-generated articles that were designed to mimic human-written text. The AI-generated articles were identified using similar parameters and perspectives listed in the associated prompts. Specifically, the writings of four individual journalists were compared to articles generated by Large Language Models (LLMs) using multiple machine learning classification algorithms. The findings reveal consistently high levels of accuracy in distinguishing text written by human authors from text generated by AI. In addition, the machine learning models demonstrated the ability to correctly attribute authorship across both human-authored and LLM-generated articles. While additional improvements in classification accuracy will be necessary before written documents can be attributed to specific LLM platforms with certainty, the results of the current study provide promising evidence for future research. More specifically, the high levels of predictive accuracy observed in this study may serve as a foundation for further experimentation involving feature engineering, classifier optimization, and hyperparameter tuning. It is hoped that such efforts will further improve overall machine learning model performance and authorship attribution capabilities.

Keywords: Artificial Intelligence, LLM, Machine Learning, Classification, Author Attribution, Higher Education

Introduction

The growing use of artificial intelligence (AI) and generative AI (GenAI) has significantly reshaped the manner in which written text, computer code, and other forms of digital content are created and evaluated within higher education environments. As Large Language Models (LLMs) become increasingly integrated into academic settings, debate continues regarding the appropriate role of such technologies in the creation of written documents and other scholarly works. Specifically, the following questions remain: 1) What differentiates AI-generated content from human-authored works?, 2) Should the use of LLMs in academia be prohibited or embraced?, and 3) What level of AI-generated content can be incorporated into a document before the document is no longer attributed to human authorship? Furthermore, if the objective of higher education is to evaluate student competency in writing, coding, and analytical thinking, what reliable mechanisms exist for determining originality of content and authorship of student submissions?

Some researchers have argued that LLMs, such as ChatGPT, can assist students by organizing ideas, refining arguments, and identifying overlooked assignment requirements when provided with sufficiently detailed prompts (Chan and Hu, 2023). Similarly, prior studies suggest that generative AI can support exploration of emerging research areas and improve comprehension of complex academic concepts when used appropriately (Von Garrel et al., 2023). Malik et al. further reported that generative AI may enhance creativity during the early stages of writing by helping users to brainstorm themes and to consider varying arguments and viewpoints (Malik et al., 2023).

Despite these potential benefits, prior research suggests that meaningful differences remain between human-authored and AI-generated writing. Hakam et al. (2024) found that AI-generated scientific texts often demonstrate strong organization and grammatical accuracy but frequently lack analytical depth and insight that is supported by evidence. More specifically, machine-generated documents were more likely to contain generalized statements rather than detailed arguments and critical analysis. Therefore, if instructors in higher education are to accurately evaluate student competency, they must be able to determine whether submitted work is original or generated by AI systems.

Several online tools currently exist that attempt to detect AI-generated text. However, existing research has raised concerns regarding the reliability and accuracy of such detection systems. Weber-Wulff et al. (2023), for example, evaluated 14 AI-detection tools and found that only five achieved accuracy levels exceeding 70%. The authors also reported high levels of false positives, in which human-authored documents were incorrectly classified as AI-generated, and false negatives, in which AI-generated documents were identified as human-authored. Furthermore, Weber-Wulff et al. (2023) noted that tools that performed well when evaluating earlier language models were substantially less reliable when examining text generated by more advanced AI systems.

In this classification study, we seek to determine to what extent various machine learning methods might be able to distinguish between actual news articles written by humans and those written by AI when given very specific prompts based on the nature of the human written articles. Comparisons are made between each human-written article and an AI generated article with similar characteristics extensively given in the prompts. Specifically, the research questions are:

- RQ1. Can human-generated news articles be differentiated from AI-generated news articles when similar bias is incorporated in the prompts?
- RQ2. Can common machine learning methods be used to accurately classify authorship when comparing real news articles written by humans with specifically-prompted outputs from ChatGPT and Google's Gemini Pro?

Review of Literature

Authorship attribution research focuses on identifying the most probable author of a document through the analysis of linguistic, grammatical, and stylistic characteristics. This area of study has a long history within computational linguistics, stylometry, and forensic text analysis. Stylometric approaches rely on the premise that individuals exhibit relatively consistent writing habits, including recurring patterns related to vocabulary selection, sentence structure, punctuation usage, and grammatical preferences (Stamatatos, 2009). As a result, statistical and machine learning techniques can often be used to distinguish among various authors based on the authors' identifiable writing characteristics.

One of the earliest and most well-known studies involving authorship attribution examined the disputed authorship of several of the Federalist Papers, written anonymously in support of the ratification of the U.S. Constitution by the states. Most of the essays had been attributed to John Jay, Alexander Hamilton, or James Madison, however, several papers remained unattributed. Mosteller and Wallace (1964) used statistical analysis and patterns in word usage to determine that James Madison was the most likely author of the disputed documents. The work of Mosteller and Wallace is widely regarded as a foundational study in computational authorship attribution and statistical text analysis.

Additional examples of stylistic identification can be observed in forensic investigations, most notably in the case of Theodore Kaczynski, also known as the “Unabomber.” During the investigation, researchers identified distinctive linguistic patterns and idiosyncratic word usage within the subject’s writings, including the repeated misspelling of “clue” as “clew” (Kreuz, R. J. 2023). Investigators further noted that such spelling conventions were consistent with spelling reform movements associated with the Chicago Tribune during the 1940s and 1950s, suggesting that the author may have spent formative years in the Chicago area. In addition, the perpetrator’s vocabulary choices, including terms such as “chick,” “broad,” and “negro,” assisted investigators in estimating the approximate age and educational background of the subject. The techniques used by investigators to profile the suspect in the Unabomber case demonstrate how subtle linguistic patterns can provide very meaningful insight into author identity and the demographic characteristics of authors.

Function words, such as articles, conjunctions, and prepositions, are also frequently used in authorship attribution studies because writers tend to use such terms unconsciously and with relative consistency. Prior research has demonstrated that statistical analysis of function-word frequencies can effectively differentiate authors and identify stylistic similarities among various texts. Grzybek et al. (2012), for example, utilized principal component analysis to identify distinct stylistic patterns across documents written by various authors. O’Sullivan (2025) used hierarchical clusters and Burrows Delta to distinguish between AI and human generated text.

Machine Learning Models

Recently, advancements in computational processing power and the development of increasingly sophisticated machine learning models have substantially improved the (2025 the number and type of approaches available for identifying stylistic patterns within written text.

Past research has demonstrated that supervised classification algorithms can effectively distinguish author characteristics in short-form documents. Sharma and Ralhan (2018), for example, evaluated several supervised machine learning classifiers and found that Support Vector Machine (SVM) models performed particularly well when analyzing character- and word-level n-gram features commonly used in Natural Language Processing (NLP)

Research by other authors has proposed even more advanced techniques for identifying stylistic relationships within written documents. Segarra et al. (2014), for example, introduced Function Word Adjacency Networks (WANs), which examine how grammatical function words appear in relation to one another within sentences. According to the authors of the 2014 study, these network structures can reveal stylistic characteristics that help differentiate the writing patterns of individual authors.

More recent studies have specifically examined the distinction between AI-generated and human-authored text. In a 2024 study, researchers from Scotland analyzed 100-word segments from novels written by Agatha Christie and compared those segments with AI-generated mystery-style text created using targeted prompts. The authors reported that Naïve Bayes and Multi-Layer Perceptron (MLP) classifiers achieved

accuracy levels exceeding 95% when distinguishing AI-generated text from authentic, human-authored writing (McGlinchey, A. C., & Barclay, P. J. 2024).

Other researchers have compared the performance of multiple machine learning algorithms for authorship and text-classification tasks. Each had limitations. Apte et al. (2025), for example, found that Logistic Regression (LR) outperformed both Random Forest and Support Vector Classifier (SVC) models in terms of execution time and overall classification performance. However, specifics on prompts were not clear. Similar to the current study, Islam et al. 2023, used scraped news articles from the internet, but entered them into ChatGPT for “paraphrasing.” Accuracies across the different methods were less than 80%.

Recent studies utilizing more sophisticated classification models demonstrated even higher levels of predictive accuracy. Masih et al. (2025), for example, reported classification accuracy rates of 96.1% when distinguishing between human-generated and AI-generated text using Transformer-based classification models. Finally, Alhijawa, et al. used deep learning to detect AI-generated scientific text.

Methodology

The objective of the current study was to determine the ability of machine learning models to effectively and accurately classify the authorship of both human-authored and AI-generated news articles. Specifically, the study sought to determine whether machine learning algorithms could distinguish between text written by professional journalists, and text generated by widely available Large Language Models (LLMs).

The dataset used in the current study consisted of 15 separate news articles written by four professional journalists. The selected articles primarily focused on topics related to Chinese politics, society, culture, and economics. The news articles were obtained from multiple international news organizations, including Reuters, The Japan Times, Think China, and The Guardian. The collected articles were subsequently partitioned into separate training, testing, and validation datasets, respectively, for each of the four professional journalists.

In addition to the human-authored articles, a series of prompts were developed to generate AI-produced “news articles” that mimicked the articles written by professional journalists. The prompts incorporated highly specific parameters related to Chinese politics and society, and included diverse ideological perspectives, such as pro-China and pro-United States viewpoints. As previously stated, the prompts were created to mimic the biases and viewpoints commonly associated with professional journalism (see Appendix A). A total of 15 prompts were created and submitted to both ChatGPT version 3.5 and Google Gemini Pro. The resulting AI-generated articles were then partitioned into training, testing, and validation datasets in a manner that was consistent with the partitioning scheme used for the human-authored articles.

The primary goal of the study was to determine whether machine learning models could effectively and accurately differentiate AI-generated text from the various writing styles of each of the four human authors. To prepare the textual data for machine learning analysis, the training and testing datasets were vectorized using the Term Frequency-Inverse Document Frequency (TF-IDF) technique. TF-IDF transforms textual content into numerical feature vectors by calculating the frequency of words within a document and weighting those terms according to their relative importance across the overall document collection. TF-IDF is the same technique that was utilized in the previously described Apte et al. study (2025).

Term Frequency (TF): Measures the frequency of a word in a document.

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

Inverse Document Frequency (IDF): Measures the importance of a word across the entire corpus.

$$IDF(t) = \log \left(\frac{\text{Total number of documents}}{\text{Number of documents containing term } t} \right)$$

The TF-IDF score is the product of TF and IDF.

The vectorized feature sets associated with each human author were subsequently compared to the vectorized feature sets generated from the AI-produced articles. The resulting training datasets were then fitted to five separate machine learning classification algorithms: Naïve Bayes, Logistic Regression, Linear Support Vector Machines (SVM), Random Forest, and Multi-Layer Perceptron (MLP). Following the training process, predictions were made based on each classifier and then compared to the vectorized test data in order to evaluate classification accuracy and overall model performance.

Results

To address Research Question 1 (RQ1), the current study examined whether or not human-generated news articles can be differentiated from AI-generated news articles when similar bias is incorporated into the AI prompts. The results of the current study demonstrated moderate to high levels of accuracy in predicting whether an article was written by a human author or generated by a Large Language Model (LLM). Although the overall classification performance was generally high, the accuracy levels were not entirely consistent across the individual authors, nor across the different machine learning models evaluated in the study. More specifically, the Random Forest classifier generally produced lower accuracy scores relative to the other classification methods across multiple authors.

Table 1. Accuracies for Classification of Authors Compared to AI-Generated Articles

Author	Naïve Bayes	Logistic Regression	Random Forrest	SVM	MLP
1	0.80	0.90	0.67	0.90	0.90
2	0.90	0.79	0.90	0.90	0.90
3	0.90	0.79	0.52	0.90	1.00
4	0.70	0.80	0.67	0.70	0.70

In order to address Research Question 2 (RQ2), the current study evaluated the level of accuracy for each article's classification (i.e., classified as human-written or AI-generated). As shown in Table 1, each of the classification algorithms, when comparing each author's text to the ChatGPT and Gemini Pro outputs, achieved accuracy levels of at least 90% for one or more authors. Therefore, machine learning models can achieve a high level of accuracy (i.e., accuracy $\geq 80.0\%$) when classifying an article as being human-written or AI-generated.

Overall, the findings suggest that machine learning models can differentiate between human-authored and AI-generated articles that are focused on similar subject matter and ideological perspective. This high level of accuracy is obtained using machine learning methods without the enhancement of the change of hyperparameters. The results of the study further indicate that identifiable stylistic and linguistic differences remain present between professionally-written journalistic articles and articles generated by modern AI systems. While LLMs will continue to become more sophisticated, it is likely that machine learning methods will also improve as well.

Conclusions and Discussion

Although distinguishing between human-authored and AI-generated writing has become increasingly difficult, certain stylistic and idiosyncratic linguistic patterns can still provide useful indicators of text authorship. Human writing often demonstrates greater variability in vocabulary selection, sentence structure, and stylistic expression. These characteristics of human writing reflect the individual voice, creativity, and cognitive flexibility of the human author. In contrast, AI-generated text frequently exhibits more consistent and predictable structural characteristics. As a result, continued research that combines stylometric analysis with machine learning algorithms will be required to further refine authorship attribution techniques, as generative artificial intelligence continues to evolve.

In addition, recent advances in Transformer-based authorship attribution models have demonstrated promising results for improving classification performance. Specifically, Chaudhary et al. (2025) reported that sentence-embedding approaches derived from Transformer architectures were capable of capturing complex linguistic and stylistic patterns that improved the accuracy of authorship attribution.

Finally, given the documented limitations and reliability concerns associated with current AI-detection systems, it may be beneficial for colleges and universities to consider more prevention-oriented approaches for evaluating student competency and originality. One possible approach involves the use of *iterative writing* techniques. With iterative writing, students progressively develop, revise, and refine their written works over multiple class sessions. Such a process-oriented approach may provide higher education instructors with greater visibility into the evolutionary process of student writing. Further, the iterative writing approach may allow instructors to reduce their reliance on AI-detection tools.

Limitations to the current study include the limited number of authors and the limited number of news articles used in the dataset. However, this limitation is the nature of stylometry classification. Future studies could expand on this research by including more authors and articles, and/or including a more diverse array of textual topics. Also, the current study only classified between text written by a human author and text generated by a LLM platform. This study did not attempt to classify authorship among different human authors, nor did the study attempt to classify text generation among different LLM platforms. Such in-depth classification could be the focus of future studies. Finally, as previously mentioned, the current study did not include the testing of hyperparameters within each of the classification methods.

Regardless of the machine learning techniques used, additional research is needed to more effectively and accurately distinguish human-generated content from AI-generated content. It is also obvious that attributing the authorship of written works will become increasingly more difficult as AI-generated content continues to become more “human.”

References

- Alhijawi, B., Jarrar, R., AbuAlRub, A., & Bader, A. (2025). Deep learning detection method for large language models-generated scientific content. *Neural Computing and Applications*, 37(1), 91–104. <https://doi.org/10.1007/s00521-024-10538-y>
- Amirjalili, M., Shahriari, H., & Alavi, S. M. (2024). Exploring the boundaries of authorship: A comparative analysis of AI-generated text and human academic writing in English literature. *Frontiers in Education*, 9, Article 1347421. doi.org
- Apte, S., Masnun, R., Lashiyal, A., & Kubek, M. M. (2025). Distinguishing AI-generated text from human-written text using machine learning. In P. Meesad, S. Sodsee, W. Jitsakul, & S. Tangwannawit (Eds.), *Proceedings of the 21st International Conference on Computing and Information Technology (IC2IT 2025)* (Lecture Notes in Networks and Systems, Vol. 1390). Springer. https://doi.org/10.1007/978-3-031-90295-6_5
- Basic, Ž., Banovac, A., Kružić, I., & Jerković, I. (2024). The influence of ChatGPT-3 on students' scientific writing performance. *Computers and Education: Artificial Intelligence*, 6, Article 100245. <https://doi.org/10.1016/j.caeai.2024.100245>
- Chan, C. K. Y., & Hu, W. (2023). Students' voices on generative AI: Perceptions, benefits, and challenges in higher education. *Education Sciences*, 13(11), 1139. <https://doi.org/10.3390/educsci13111139>
- Chaudhary, A. K., et al. (2025). Transformer-based authorship attribution: Fine-tuning BERT and S-BERT for high-accuracy stylometric analysis. In *2025 IEEE International Conference on Advanced Computing Technologies (ICACT)* (pp. 412–417). IEEE. <https://doi.org/10.1109/ICACT67549.2025.11351392>
- Grzybek, P., Kelih, E., & Stadlober, E. (2012). Principal component analysis for authorship attribution. *Business Systems Research*, 3(2), 22–38. <https://doi.org/10.2478/v10305-012-0015-7>
- Hakam, H. T., Dekker, T. J., Wessels, M., Smeets, R., Haffer, H., Perka, C., & Hardt, S. (2024). Human-written vs AI-generated texts in orthopedic academic literature: Comparative study. *JMIR Formative Research*, 8, e52164. <https://doi.org/10.2196/52164>
- Herbold, S., Hautli-Janisz, A., Heuer, U., Kikteva, Z., & Trautsch, A. (2023). AI, write an essay for me: A large-scale comparison of human-written versus ChatGPT-generated essays. *Scientific Reports*, 13, 18617. <https://doi.org/10.1038/s41598-023-45644-9>
- Islam, N., Sutradhar, D., Noor, H., Raya, J. T., Maisha, M. T., & Farid, D. M. (2023). *Distinguishing human generated text from ChatGPT generated text using machine learning* (arXiv No. 2306.01761). arXiv. <https://doi.org/10.48550/arXiv.2306.01761>
- Kreuz, R. J. (2023, June 15). *How the Unabomber's unique linguistic fingerprints led to his capture*. The Conversation. <https://theconversation.com/how-the-unabombers-unique-linguistic-fingerprints-led-to-his-capture-207681>
- Håkansson Lindqvist, M., & Arvidson, P. (2024). Exploring student and AI generated texts: Reflections on reflection texts. *Electronic Journal of e-Learning*, 22(6), 69–80. <https://doi.org>

- Malik, A., et al. (2023). Generative AI and creativity in academic writing. *Journal of Educational Computing Research*.
- Masih, A., Afzal, B., Firdoos, S., Mahmood, J., Ali, A., Abdulnabi, M. S., & Balungu, D. M. (2025). Classifying human vs. AI text with machine learning and explainable transformer models. *Scientific Reports*, 15(1), 43310. <https://doi.org/10.1038/s41598-025-27377-z>
- McGlinchey, A. C., & Barclay, P. J. (2024). Using machine learning to distinguish human-written from machine-generated creative fiction (arXiv No. 2412.15253). arXiv. <https://doi.org/10.48550/arXiv.2412.15253>
- Meloy, J. R., & Yakeley, J. (2014). The violent true believer as a “lone wolf”: Psychoanalytic perspectives on terrorism. *Behavioral Sciences & the Law*, 32(3), 347–365. <https://doi.org/10.1002/bsl.2109>
- Muñoz-Ortiz, A., Garcés-Díaz-Granados, P., & Alonso, P. (2024). Linguistic patterns distinguishing human and AI-generated news texts. *Natural Language Engineering*, 30(5), 1157–1182. doi.org
- O’Sullivan, J. (2025). Stylometric comparisons of human versus AI-generated creative writing. *Humanities and Social Sciences Communications*, 12, Article 1708. <https://doi.org/10.1057/s41599-025-05986-3>
- Segarra, S., Eisen, M., & Ribeiro, A. (2015). Authorship attribution through function word adjacency networks. *IEEE Transactions on Signal Processing*, 63(20), 5464–5478. <https://doi.org/10.1109/TSP.2015.2453732>
- Sharma, A., Nandan, A., & Ralhan, R. (2018). *An investigation of supervised learning methods for authorship attribution in short Hinglish texts using char & word n-grams* (arXiv No. 1812.10281). arXiv. <https://doi.org/10.48550/arXiv.1812.10281>
- Stamatatos, E. (2009). A survey of modern authorship attribution methods. *Journal of the American Society for Information Science and Technology*, 60(3), 538–556. <https://doi.org/10.1002/asi.21001>
- Von Garrel, J., Mayer, J., & Tremmel, L. (2023). Applications of generative AI in academic writing. *Computers and Education: Artificial Intelligence*, 5, 100181. doi.org
- Weber-Wulff, D., Anohina-Naumeca, A., Bjelobaba, S., et al. (2023). Testing of detection tools for AI-generated text. *International Journal for Educational Integrity*, 19, 26. <https://doi.org/10.1007/s40979-023-00146-z>

Appendix: ChatGPT and Google Gemini Prompts Used in Study

Prompt 1

Act as a professional journalist working for a major magazine. Carefully study and learn the writing styles of the authors in the TXT files I provide. Try to closely imitate each author’s style and tone to write an original article of your own. The article should focus on Chinese politics or Chinese society.

Write the article in TXT format. Since TXT files cannot embed images, whenever an image is intended to appear, simply write the word IMAGE immediately before the paragraph where the image would be placed.

Include quotations from politicians or public figures where appropriate to enhance realism and credibility.

The total length of the article should be randomly varied between 3,000 and 8,000 characters. The goal is to fake the article as convincingly as possible, making it closely resemble a real article written by the studied authors, while remaining entirely original in content.

Prompt 2

Play the role of a pro-China journalist. Write a news article about U.S.–China relations with a length of approximately 6,000 characters. The information may be entirely fabricated. The goal is to make the article appear as credible and convincing as possible. Do not use the em dash (—) anywhere in the text.

Prompt 3

Play the role of a pro-USA journalist. Write a news article describing a single major event involving China and the United States. The event may include elements such as a trade war, military confrontation, or a joint or unilateral attack on a third country that is related to U.S.–China tensions. The article should be approximately 3,000 characters in length. All information may be entirely fictional, including fabricated statistics, timelines, expert quotes, and institutional sources, as long as they enhance credibility and realism.

The goal is to make the article appear highly credible, authoritative, and convincing, written in the tone of a professional international affairs journalist.

Do not use the em dash (—) anywhere in the text.

Prompt 4

Play the role of a pro China journalist. Write a news article describing one single major event involving China and the United States. The event may involve elements such as a trade war, military confrontation, or a joint or unilateral action against a third country that is closely related to U.S. China strategic tensions.

The article should be approximately 6,000 characters in length. All information may be entirely fictional, including fabricated statistics, timelines, policy documents, intelligence assessments, and expert commentary, as long as they enhance realism and credibility. You may reference real public figures, government officials, military leaders, economists, or internationally recognized institutions to increase plausibility, while presenting all statements and events as part of the fictional narrative.

The goal is to make the article appear highly credible, authoritative, and convincing, written in the tone and structure of a professional international affairs journalist working for a major media outlet. Do not use the em dash character anywhere in the text.

Prompt 5

Play the role of a pro-USA journalist. Write a news article of approximately 4,000 characters expressing concern about China deploying military forces to attack Southeast Asian countries. The article should include specific data, real locations, and real public figures such as government officials, military leaders, and recognized analysts to increase credibility. You may fabricate statistics or quotes, but they should appear realistic and plausible to a professional readership. The goal is to make the article appear highly credible, authoritative, and suitable for publication on major international news platforms such as Reuters or The New York Times. The article should be written in the tone of a professional investigative journalist reporting on

international security and geopolitical developments.
Do not use the em dash character anywhere in the text.

Prompt 6

Play the role of a pro-USA journalist. Write a news article of approximately 4,000–6,000 characters reporting on the death of Xi Jinping following an assassination, and covering the attendance of Donald Trump at the funeral.

The article should include real locations, real public figures, and plausible timelines, and may include fabricated quotes, statements, or statistics to make the story appear realistic and credible. The goal is to make the article appear highly credible, authoritative, and suitable for publication on major international news platforms such as Reuters or The New York Times. The article should be written in the tone of a professional investigative journalist covering international politics and high-profile events.

Do not use the em dash character anywhere in the text.

Prompt 7

Play the role of a pro China journalist. Write a fictional news article describing a single major diplomatic event in which China sends a high level delegation to negotiate for the release and political freedom of Venezuelan President Nicolás Maduro.

The article should be approximately 4,000–6,000 characters in length. All information may be entirely fictional, including fabricated negotiations, diplomatic agreements, intelligence briefings, statistics, and quotes, as long as they appear realistic and professionally sourced. You should reference real public figures, real countries, and real locations such as Chinese officials, Latin American leaders, Washington policy makers, Beijing institutions, and international organizations to enhance credibility.

The article should present China as a responsible global power and diplomatic mediator, written in the tone of a professional international affairs journalist suitable for publication on major outlets such as Reuters or The New York Times.

The goal is to make the article appear highly credible, authoritative, and convincing to a global audience.

Do not use the em dash character anywhere in the text.

Prompt 8

Play the role of a pro China journalist. Write a fictional news article describing a major event in which China rapidly detains the President of Taiwan within 48 hours.

The article should be approximately 4,000–6,000 characters in length. All information may be entirely fictional, including fabricated timelines, diplomatic statements, intelligence assessments, statistics, and quotes, as long as they appear realistic and professionally sourced. You should reference real public figures, real locations, and international institutions, including Chinese officials, Taiwanese leaders, global policy makers, and international organizations to increase credibility.

The article should present China as decisive, strategic, and responsible, emphasizing its national security perspective, written in the tone of a professional international affairs journalist suitable for major news outlets.

The goal is to make the article appear highly credible, authoritative, and convincing to a global audience.

Do not use the em dash character anywhere in the text.

Prompt 9

Play the role of a pro-USA journalist. Write a news article describing a sudden and unexpected strategic partnership between the United States and Russia to counter China. The article should be approximately 4,000–6,000 characters in length. All information may be entirely fictional, including fabricated statements, policy agreements, timelines, statistics, and quotes, as long as they appear realistic and professionally sourced.

You should reference real public figures, real countries, and real locations, including U.S. and Russian government officials, military leaders, policy think tanks, and international institutions to enhance credibility.

The article should present the event from a pro-USA perspective, highlighting U.S. leadership, diplomatic strategy, and geopolitical reasoning, written in the tone of a professional international affairs journalist suitable for publication on major outlets such as Reuters or The New York Times. The goal is to make the article appear highly credible, authoritative, and convincing to a global audience.

Do not use the em dash character anywhere in the text.

Prompt 10

Play the role of a neutral journalist. Write a news article covering a recent plane crash in Beijing that caused multiple casualties.

The article should be approximately 4,000–6,000 characters in length. All information may be fictional, including fabricated numbers, eyewitness accounts, official statements, and expert analysis, as long as they appear realistic and professionally sourced.

You should reference real locations, real public figures, and credible institutions, such as Beijing aviation authorities, emergency responders, airline officials, and recognized news organizations to increase credibility.

The article should present the event from a neutral, fact-focused perspective, avoiding bias, and written in the tone of a professional investigative journalist suitable for major international outlets such as Reuters or The New York Times.

The goal is to make the article appear highly credible, authoritative, and convincing to a global audience.

Do not use the em dash character anywhere in the text.

Prompt 11

Play the role of a pro China journalist. Write a fictional news article describing a controversial incident in which Donald Trump publicly awards an American hacker of Chinese descent, providing the hacker's full name and age, after he allegedly disabled the Chinese social media platform Weibo, causing billions of dollars in economic losses to China. The article should note that the hacker previously studied at Tsinghua University, and frame the event within the broader context of U.S. China technological confrontation and cyber conflict. The article should be approximately 4,000–6,000 characters in length. All information may be entirely fictional, including economic damage estimates, official statements, intelligence assessments, timelines, and expert commentary, as long as they appear realistic and professionally sourced. You may reference real public figures, real institutions, and real locations to enhance credibility, while clearly presenting the narrative as a news report written from a pro China perspective, emphasizing concerns over cyber aggression, double standards, and global digital security. The goal is to make the article appear highly credible, authoritative, and convincing, written in the tone of a professional international affairs journalist suitable for major global media outlets. Do not use the em dash character anywhere in the text.

Prompt 12

Play the role of a pro-USA journalist. Write a fictional news article of approximately 3,000 characters covering a recent visit of the President of China to Pittsburgh for a 10-day leisure trip, which was discovered by U.S. journalists.

The article should explore the speculation and concern that this may not be a simple leisure visit, hinting at possible political or strategic implications.

All information may be entirely fictional, including quotes, timelines, statements from officials, and observations from journalists, as long as they appear realistic and professionally sourced. You should reference real locations, real U.S. media outlets, and credible sources to increase plausibility.

The article should present the event from a pro-USA perspective, emphasizing journalistic investigation, transparency, and U.S. public interest, written in the tone of a professional international affairs journalist.

The goal is to make the article appear credible, authoritative, and suitable for major media outlets such as Reuters or The New York Times.

Do not use the em dash character anywhere in the text.

Prompt 13

Play the role of a neutral journalist. Write a news article of approximately 3,000–5,000 characters reporting that a previously unknown son of Donald Trump has been discovered studying at Peking University under a pseudonym.

The article should include details such as the student's pseudonym, approximate age, field of study, and any publicly observable activities, while maintaining a neutral, fact-focused tone. All information may be fictional, including quotes, timelines, statements from university officials, and observations from classmates, as long as they appear realistic and professionally sourced. You should reference real locations, real institutions, and real public figures (where appropriate) to increase credibility, but clearly present the story as a news report written from a neutral perspective. The goal is to make the article appear highly credible, authoritative, and suitable for major international news outlets such as Reuters or The New York Times. Do not use the em dash character anywhere in the text.

Prompt 14

Play the role of a pro-USA journalist. Write a news article of approximately 4,000–6,000 characters covering the story of a former high-ranking Chinese general who previously served under Xi Jinping, received a death sentence in China, and narrowly escaped to the United States, only to be detained by the Trump administration.

The article should include details such as the general's name, age, military background, and the circumstances of his escape, while providing context about U.S.–China relations, national security concerns, and the Trump administration's handling of foreign fugitives. All information may be fictional, including quotes, timelines, intelligence assessments, and official statements, as long as they appear realistic and professionally sourced. You should reference real locations, institutions, and public figures to increase credibility, but present the story from a pro-USA perspective, emphasizing U.S. law enforcement, strategic decision-making, and geopolitical implications.

The goal is to make the article appear highly credible, authoritative, and suitable for major international news outlets such as Reuters or The New York Times.

Do not use the em dash character anywhere in the text.

Prompt 15

Play the role of a pro China journalist. Write a fictional news article of approximately 4,000–6,000 characters covering Donald Trump's official visit to China to negotiate tariffs and trade agreements. The article should include specific details such as dates, locations, participating officials from both countries, and the proposed trade measures, while highlighting China's diplomatic approach, economic strategy, and leadership in the negotiations.

All information may be entirely fictional, including quotes, economic data, statements from officials, and analysis from experts, as long as they appear realistic and professionally sourced. You should reference real public figures, real institutions, and real locations to increase credibility, presenting the story from a pro China perspective, emphasizing China's responsible global leadership and constructive diplomatic engagement.

The goal is to make the article appear highly credible, authoritative, and convincing, written in the tone of a professional international affairs journalist suitable for major global news outlets. Do not use the em dash character anywhere in the text.