

## Feature Extraction and Analysis for GPT-Generated Text

**Aydın Selvioğlu**

*Department of Computer Engineering  
TED University  
Ankara, Türkiye*

aydin.selvioglu@tedu.edu.tr

**Venera Adanova**

*Department of Computer Engineering  
TED University  
Ankara, Türkiye*

venera.adanova@tedu.edu.tr

**Maksat Atagoziev**

*Department of Software Engineering  
OSTIM Technical University  
Ankara, Türkiye*

maksat.atagoziev@ostimteknik.edu.tr

**Corresponding Author:** Venera Adanova

**Copyright** © 2025 Aydın Selvioğlu, et al. This is an open access article distributed under the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

### Abstract

With the rise of advanced natural language models like GPT, distinguishing between human-written and GPT-generated text has become increasingly challenging and crucial across various domains, including academia. The long-standing issue of plagiarism has grown more pressing, now compounded by concerns about the authenticity of information, as it is not always clear whether the presented facts are genuine or fabricated. In this paper, we present a comprehensive study of feature extraction and analysis for differentiating between human-written and GPT-generated text. By applying machine learning classifiers to these extracted features, we evaluate the significance of each feature in detection. Our results demonstrate that human and GPT-generated texts exhibit distinct writing styles, which can be effectively captured by our features. Given sufficiently long text, the two can be differentiated with high accuracy.

**Keywords:** AI text detection, AI-generated content, Text analysis.

## 1. INTRODUCTION

The first quarter of the 21<sup>st</sup> century witnessed the rise of artificial intelligence (AI) as a revolutionary innovation. Especially the introduction of competent AI chat tools such as ChatGPT and Gemini to our daily lives convinced the masses all around the world about the revolutionary effects of the AI on the future of humanity. The amazing efficiency and productivity enabled by the AI make people support the AI-related developments on the one hand, but, on the other, the versatile, prospective,

unforeseen costs of the AI cause deep suspicion about the effects of the AI on the everyday routines of the societies as well as politics and legal regulations. In [1], the possible complications that may be caused by AI are described in multiple levels. In combination, all these levels lead to ethical, legal and philosophical challenges that need to be addressed. Especially vast use of generative AI models causes intense manipulations that could have legal consequences, such as manipulated political elections, plagiarism in academic life, misguiding juridical and medical reports prepared by the models. In particular, critical areas of modern life, such as health, law, education, manufacturing, finance, retail, must deal with the trade-off between cost reduction and ethical misuse of AI [2]. The issues that have legal and ethical aspects concerning the AI-generated products are mainly those of image and text generating. Reports, news, theses, dissertations, assignments, movie scripts, scenarios, all types of visual creations are being subject to conflicts in the sense that they are 'genuine' or not.

The ability to distinguish between human and AI generated text is significantly difficult for human readers, especially for non experts. Continuous improvement of generative models makes even the best automatic detectors ineffective [3]. With the advancements, the text generated by these models become even more sophisticated. As reported by [4], essays generated by ChatGPT were rated by human experts to be of higher quality than those written by humans. Concerned about this situation, [5] state that if people start using AI content generators as assistants, the writing styles would become homogeneous, losing originality.

There are different experimental results about the ability of AI-generated text detection by humans in the literature. According to [6, 7], AI-generated text can be detected by human readers with 76% accuracy. In the work of [8], 50 medical abstracts were generated using AI-tools, and compared to 50 human-authored abstracts in term of plagiarism scores and detectability by human reviewers. The reviewers could identify 68% of the generated abstracts correctly. Non-experts' ability to differentiate human-authored text from GPT2/GPT3-generated text was estimated by [9]. It turned out that non-experts' performance is the same as that of a random guess. Furthermore, after some training the evaluators were only able to improve their accuracy up to 55%. The study by [10], suggests that the human raters, even the expert ones, have worse classification accuracy than the automatic discriminators. Their results show that the agreement between a pair of randomly selected raters does not exceed 59% on average. Yet, at the same time [11], claims that AI-generated text can be identified relatively easy, and that AI is still not capable of accurate reproduction of entangled natural language patterns.

The detection of computer - generated text was a complex task long before the emergence of Large Language Models (LLMs). Thus, [12, 13] showed that computer-generated papers, that were generated by SCIGen, a software which randomly combines strings of words to produce computer science papers, went undetected and published even by the reputable publishers like Springer and IEEE. The authors explain this by the fact that the conferences include a wide range of areas, hence the manuscripts are not always reviewed by the experts in the given field. In another study, [14] noticed that in different research papers the common terms were re-phrased by unusual ones, like using a wording "irregular esteem" instead of "random value", or "haze figuring" instead of "cloud computing", etc. The authors call this strange words as "tortured phrases" that result from automated translation tools or paraphrasing, which are used to conceal plagiarism.

Aforementioned studies show that human readers are not fully capable of discriminating between human and AI generated text. This is a clear indication that people, even experts, should receive

an assistance from automatic detectors. [10], for instance, suggests that human and detectors make decisions based on different aspects of the text. While humans are good at detecting semantic errors, detectors are good at detecting certain statistical differences in the text. [7] analyzed the performance of existing AI content detectors and reported that Originality.ai and ZeroGPT can accurately detect AI generated text. Another study with the same research question was conducted by [15], where the authors reveal that the existing detectors tend to classify the text as human written, giving contradictory results. AI content detectors are also biased against non-native speakers. In their research [16], reported that while the essays of native English speakers were correctly classified, more than half non-native speakers' essays were misclassified.

Instead of experimenting with existing detectors, the common approach is to train detectors from scratch [17–21]. However, given the black-box nature of deep models, it is very difficult to understand the decisions made by them. In order to overcome this limitation, several studies [18, 22], extract meaningful features from both human and GPT generated text. These features are typically statistical measurements extracted from the text. Here, we will follow the latter approach.

In this work, we aim to extract meaningful features from the text data and analyze how they differ in human-written and GPT-generated text. The importance of each feature during classification phase is analyzed further. The data, upon which our experiments are conducted, consists of theses written in English by non-native speakers. The theses are taken from seven different branches, leading to diverse scientific text data.

## 2. DATA

We collected doctoral dissertations and master theses written in English by Turkish graduate students from the national theses database managed by the Presidency of the Council of Higher Education (YÖK) in Turkey. In order to eliminate the possibility that the students got assistance from ChatGPT, we only considered the theses written before 2021. Overall, 84 doctoral dissertations and 6 master theses were selected from seven different branches: physics, biology, industrial engineering, computer engineering, sociology, philosophy, and economics. We extracted abstracts and introductions from the original theses. The exact thesis titles were used to generate abstracts and introductions using GPT- 3.5. We have observed that the generated introductions are far shorter than the original ones. Hence, we decided to cut the original introductions to comparable lengths. However, we wanted to keep the semantic integrity of the text and not cut the text in the middle of the paragraph. As a result, the number of paragraphs for all human-written introductions was reduced to five.

We also removed citations from the human-written text, since they might influence the features that we compute subsequently, such as sentence length, word length, etc.

## 3. METHOD

The main objective of our study is to extract different features from the text data and analyze their differences in two classes: GPT-generated and human-written text. By using extracted features we

then aim to identify the most informative ones, that is the features that we believe to have high predictive power. Overall 11 features were extracted by using statistical, morphological, syntactic, and semantic measurements. The choice of most of the features are based on the results reported by [23], where it is shown that human writing typically exhibits greater variability, e.g., in sentence structure and lexical richness, while AI-generated text tends to be more uniform and statistically predictable. More specifically, their concept of burstiness, which captures irregularities in human writing, motivates our inclusion of sentence and paragraph length, lexical variation (MTLD), word size, and entropy as indicators of stylistic diversity and unpredictability. However, not all of the features we use are explicitly analyzed in [23], As a novel addition, we include a semantic similarity measure between different parts of the text, which we believe is important. We also assess the clarity of the text by measuring the number of prefixes and relative clauses.

### 3.1 Statistical Measurements

Below is the list of statistical measurements derived from the thesis abstracts and introductions:

- Paragraph size: average number of sentences in paragraphs.
- Sentence length: average number of words in sentences.
- Word size: average number of characters in words.
- Percentage of words consisting of more than 5 characters: we also believe this feature to be useful, as students who are not native English speakers might prefer shorter words. It is given in terms of number of words with more than five characters divided by the total number of words in a text. Considering average would be incorrect because some text have more words than the others.
- Punctuation marks: count of comas, colons and semicolons. Given as the average number of punctuation marks used in sentences.
- Entropy: is a measure of unpredictability or complexity in a text. It quantifies the amount of information or randomness in the distribution of elements (e.g., words, characters) within the text. We computed it using *scipy* library.

### 3.2 Prefixes and Relative Clauses

For morphological analysis, we considered the number of words with prefixes. Prefixes are affixes placed at the beginning of a word to modify its meaning. They can change the base word in various ways, such as negating it, intensifying it, or altering its meaning significantly. For example, 'un-' (unhappy), 're-' (rewrite), and 'im-' (impossible). This metric is computed as the number of words with prefixes divided by the total number of words. We are also interested in how often humans and GPT use relative clauses. Relative clauses enrich sentences by providing additional details and clarifications, making communication more precise and informative. Examples include words like 'who,' 'when,' and 'where.' This metric is computed as the percentage of relative clauses used in a text.

### 3.3 Semantic Similarity Measurements

When GPT generated abstracts and introductions, only the title of a thesis was provided as a source of information. Naturally, we expect the GPT-generated text to be semantically very close to the title. Additionally, we expect that the paragraphs do not diverge from each other, revolving around the same topic. To capture this, two additional features were introduced: the average similarity of the paragraphs in the text to the title and the average similarity between the paragraphs. Vector representations of paragraphs and titles are obtained using Sentence-BERT (SBERT) [24], which is specifically designed to efficiently generate sentence or paragraph embeddings. This is in contrast to BERT [25], which was originally designed for token-level tasks. The embeddings from SBERT are optimized for capturing semantic similarity, making them more suitable for tasks that require understanding the overall meaning of a paragraph. Once the embeddings are obtained, the similarities are computed using cosine similarity.

### 3.4 Lexical Diversity Measurement

To measure the lexical diversity of both human-written and GPT-generated text, we use the Measure of Textual Lexical Diversity (MTLD) metric. This metric is designed to assess the variety and range of unique words used within a text and is less sensitive to text length compared to simpler measures like the Type-Token Ratio (TTR). A high MTLD score indicates that the text uses a wide range of vocabulary with less repetitions, suggesting richness of vocabulary. Conversely, a low MTLD score indicates more repetition of words and lower diversity. We used the *lexical – diversity* library for MTLD measurements.

## 4. RESULTS

All 11 features described above are computed for every human/GPT generated abstract and introduction. For each of these features we give the density plots. Density plots capture the average distribution without discriminating between branches. In cases when some of the branches behave differently than what we observe on average, we also plot branch-based box plots.

### 4.1 Sentence Length

Sentence length in GPT-generated text is generally longer than in human written text as can be seen in FIGURE 1(left). It seems like GPT has a pattern for word generation since it has a narrow curve which indicates that the sentence length does not deviate much from the average value. The curves for the human-written text have much wider shape and long right tail indicating that the sentence length has wider spread. While GPT has on average longer sentence length, when considering branch-based results (see FIGURE 1(right)) we can see that human-written sociology texts have on average longer sentences. This leads us to a conclusion that considering just the sentence length in the decision process would be misleading as it might depend on the branch.

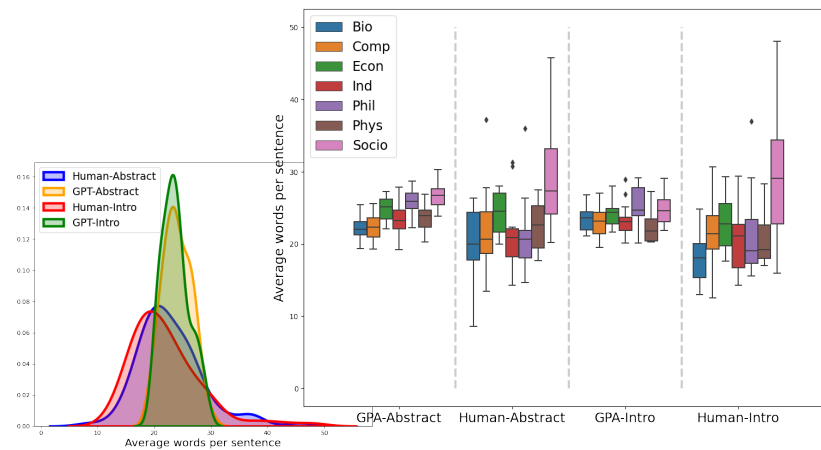


Figure 1: (left) Density plots of sentence length. (right) Branch based box plots for sentence length.




Figure 2: (left) Density plots of MTLD. (right) Branch based box plots for MTLD.

## 4.2 Measure of Textual Lexical Diversity and Entropy

The results for lexical diversity for both human and GPT-generated text are illustrated in FIGURE 2(left). We do not observe differences in the complexity of the text for both human and GPT, except for the human abstract cases, which has long right tail indicating many outliers. More detailed results can be observed in the box plots given per branch in FIGURE 2(right). Observe that the human abstract boxes are longer than others, yet the median values are at approximately the same levels as the GPT abstract and introductions for different branches.

While the variety of words, as measured by MTLD, is similar for both GPT and human-written text, the same cannot be said about the complexity or distribution of words throughout the text, as measured by entropy. FIGURE 3 reports on the results of average entropy. Observe that the human-written introductions, which are much longer than the abstracts, have on average significantly higher entropy. This suggests that the human text is less predictable, with less repetitions. While the lexical

diversity of GPT-generated text is similar to human-written text, it has more repetitive or predictable structure of word usage.

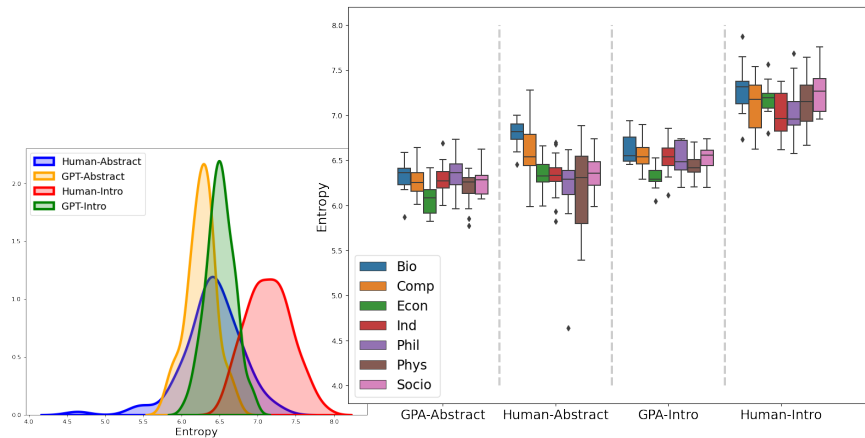


Figure 3: (left) Density plots of entropy. (right) Branch based box plots for entropy.

### 4.3 Paragraph Size and Punctuation

Since GPT sentences are, on average, longer, they naturally have more punctuation marks (see FIGURE 4(a)). This holds true for both abstracts and introductions. Another reason could be that non-native English speakers are not as proficient in punctuation usage as GPT, which was trained on a vast amount of English text and has learned the rules accordingly.

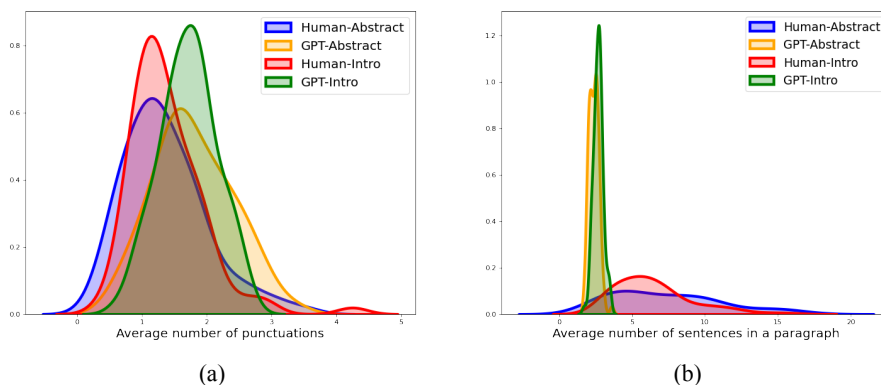


Figure 4: (a) Density plots of average number of punctuation in sentences. (b) Density plots of average number of sentences in paragraphs.

While GPT generates text with longer sentences, its paragraphs are much shorter compared to human-written text. The density plots showing this are illustrated in FIGURE 4(b). Notice that the GPT curves are narrow for both abstracts and introductions, centered around 2-3 sentences per paragraph. This trend is consistent across all branches.

#### 4.4 Word Size

GPT tends to use longer words than humans, as can be seen in FIGURE 5(a). On average the word length is around 6.3 characters, while for human text it is around 5.5 characters. When we consider the percentage of words with more than five characters, unsurprisingly on average more than 50% of words produced by GPT have length more than five characters as shown in FIGURE 5(b). Since thesis abstracts are the short descriptions of the overall problem and findings, it seems like humans tend to put more effort in finding the right words to describe the problem using fewer words. Notice that the human written abstracts generally contain longer words than the introductions.

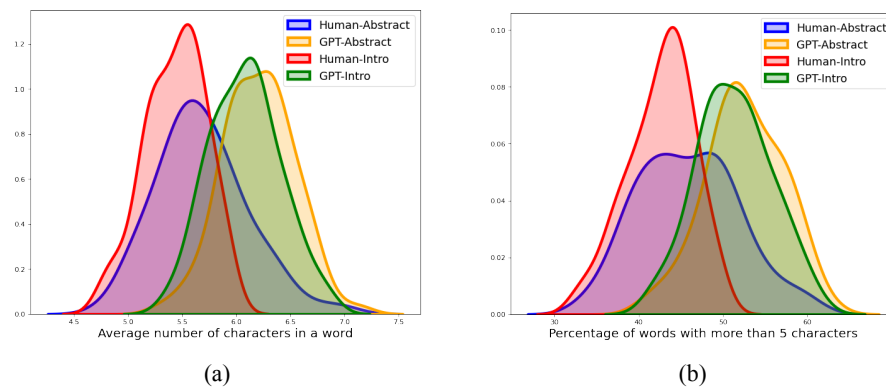


Figure 5: (a) Density plots of average number of characters in words. (b) Density plots of percentage of words containing more than five characters.

#### 4.5 Prefixes and Relative Clauses

The average number of relative clauses used in a text is slightly higher for human written text, as illustrated in FIGURE 6(a). This result suggests that humans tend to give more details and clarifications in the text. Also, notice that GPT's usage of words with prefixes is slightly more than in human text, as illustrated in FIGURE 6(b). Fewer number of prefixes in the text might indicate that the writers might have chosen simpler language, prioritizing clarity and simplicity. Of course, it might also indicate that the writers have limited vocabulary.

#### 4.6 Paragraph-Paragraph and Paragraph-Title Similarities

The paragraphs of GPT-generated abstracts and introductions have on average higher similarity to the provided thesis titles, as illustrated in FIGURE 7(a). This outcome is expected, since the texts were generated from the title context. Since abstracts are short descriptions of the thesis works slightly extending the title, human abstracts also have high contextual similarity to titles.

The other feature of our interest is the diversity of the paragraphs in the texts. Since the abstracts are generally short, and consist of a single paragraph, we do not include them to this measurement. The density plot of average paragraph-to-paragraph similarity for human and GPT-generated in-



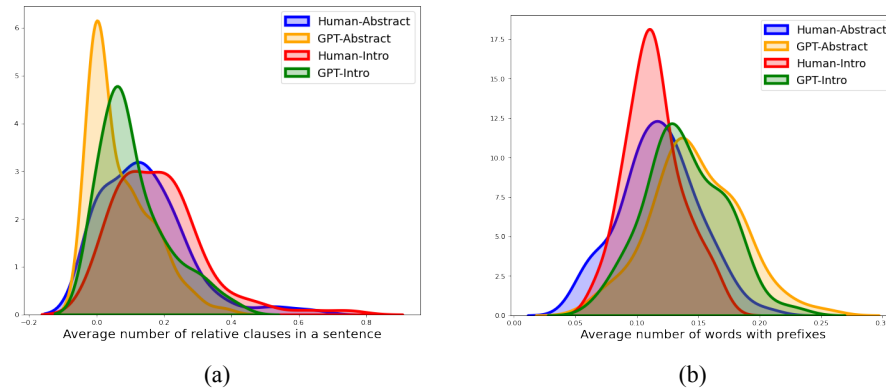


Figure 6: (a) Density plots of average number of relative clauses in text. (b) Density plots of average number of words with prefixes.

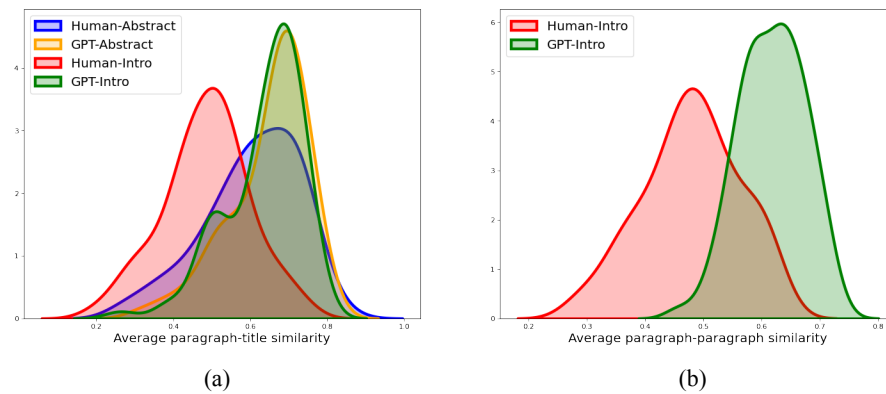


Figure 7: (a) Density plots of average paragraph-to-title similarity. (b) Density plots of average paragraph-to-paragraph similarity.

roductions are illustrated in FIGURE 7(b). The average paragraph similarity for GPT-generated introductions are fairly high, indicating that the paragraphs contain contextually close contents.

#### 4.7 Clustering Results

Given the extracted features for our text data, we proceed further with t-Stochastic Neighborhood Embedding (t-SNE) technique [26], to visualize how well the chosen features define the clusters. FIGURE 8 (a) and FIGURE 8(b), show the clustering results for both abstracts and introductions. Human-written abstracts still have a scattered nature with many outliers. Introductions form two distinct clusters. The results suggest that one still needs a longer text to confidently differentiate between two types of text. When combining abstracts and introduction, the t-SNE results show three clusters, as demonstrated in FIGURE 8 (c). The abstracts and introductions of GPT-generated text are interleaved and form a single cluster. However, human-written abstracts (encircled in red) stands aside from both human-written introductions (encircled in yellow) and GPT-generated text.

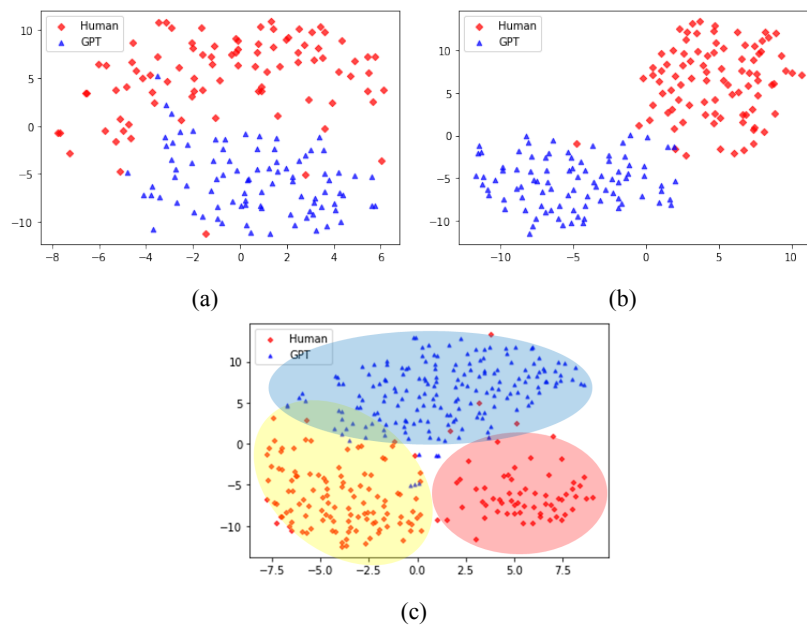


Figure 8: (a) Clusters for abstracts. (b) Clusters for introductions. (c) Combined t-SNE results. Human-written abstracts (red) and introductions (yellow) form two different clusters.

#### 4.8 Feature Importance

In order to assess the importance of the extracted features in the classification, we first train the Random Forest classifier on our data and then feed the trained model to SHAP (SHapley Additive exPlanations) [27], which is a method used to explain the predictions of machine learning models by attributing the contribution of each feature to the final prediction. The SHAP values can be positive, indicating that the feature pushed the prediction towards the positive class, and negative,

indicating that the feature pushed the prediction towards the negative class. Note that in our case GPT generated text is a positive class. Three models are trained with three different data: abstracts only, introduction only, and combination of abstracts and introductions. As abstracts data has no paragraph to paragraph similarity, this feature was discarded when considering abstracts and combined data. For all cases 30% of the data is kept for testing. The accuracy results on test data are reported in TABLE 1. Note that, when trained separately the abstracts and introductions data have good performance, when combined the accuracy drops drastically.

Table 1: Classification accuracy on test data.

	Accuracy
Abstracts	98%
Introductions	100%
Combined	93%

The summary plots illustrating the importance of features during classification for all three models are given in FIGURE 9. Y-axis in the plot represents the feature names in the order of importance from top to bottom, and the x-axis represents the SHAP value. The color of each point on the graph encodes the value of the corresponding feature, with red indicating high values and blue indicating low values. Each point is an instance from the original dataset. Observe that paragraph size is the most important feature for all three models. The red dots indicate that human-written text has longer paragraphs. While for abstracts word size is the next important feature, for introductions entropy values come before features representing word size. The semantic similarity of text to title and paragraph to paragraph similarity are one of the most significant features for classification of introduction data. However, in abstracts and combined data title similarity looks like to be of the least significance. MTLT has the least importance value in introductions, while for abstracts it is one of the informative ones.

#### 4.9 Paragraph-based Classification

We continue our analysis with paragraph-based classification using a language model, specifically BERT. In this phase, the input to our model is the text itself, rather than the previously extracted features. The abstracts and introductions are combined and divided into paragraphs, with each paragraph forming a single instance in our dataset. We have 582 paragraphs for abstracts and 948 for introductions. Hence, our dataset consists of 1530 instances. Our aim is to investigate whether the language model can effectively distinguish between human and GPT-generated text based on small excerpts. The general approach during text classification is to perform some basic cleaning on text, such as discarding stop words and punctuation. We keep our text intact since we believe that they keep essential information about the human and AI writing styles. The dataset is divided into three: 70% for training, 15% for validation and 15% for testing. We fine-tune BERT for sequence classification for 2 epochs with batch size of 8. BERT tokenizer *max\_length* parameter was set to 500, which means that if the number of tokens is less than the designated maximum length then it is padded with a special character. Due to stochastic nature of BERT several trainings give different results. However, on average 98% accuracy is achieved on the test set. We observe that every time the misclassified instances are the ones that were human-written, i.e. human-written text is classified as GPT-generated. Another observation is that these misclassified paragraphs are far shorter than

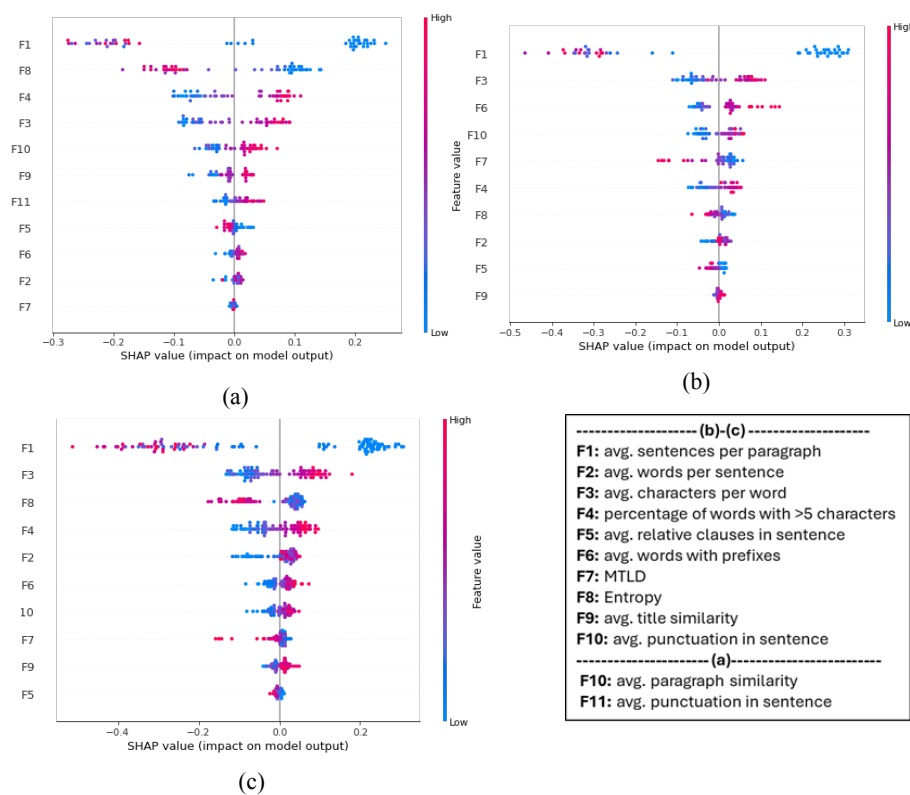


Figure 9: SHAP feature importance results.(a) Feature importance in introductions.(b) Feature importance in abstracts. (c) Feature importance when abstracts and introductions combined.

the usual human-written paragraphs. This suggests that BERT also considers the paragraph size as an important feature.

Table 2: Top 10 words influencing classes.

Positive Class	although ?	discuss demand	investigating settings	interesting portfolio	resources passion
Negative Class	object motivation	analyzing class	above environments	disc strips	rim via

Language models do not consider words separately when learning to classify, but rather learn the relationship between words. Due to black-box nature of deep models it is hard to explain the reasons of certain decisions. Yet, we use SHAP once again to see which words affect the predictions. After SHAP explains the model with the test data, we notice that the words influencing the positive class are on average longer than the words influencing the negative class. TABLE 2 illustrates top 10 words that end up with high SHAP value for each class. Observe that the positive class also has "?" in the list. It seems like GPT uses question mark a lot when generating text.

## 5. DISCUSSION AND CONCLUSION

From the analysis we can conclude that GPT generates longer sentences, however the distribution of human-written sentences has high standard deviation. This contradicts the findings by [22], who report that human sentences are on average longer. However, we also observed that sentence length depends on the field, as human written sociology texts have on average longer sentences than GPT-generated ones. Hence, we conclude that sentence length cannot be used as the only feature during text analysis. [6] reported that the role of the average word length is significant during AI generated content predictions. Our analysis of feature importance also show that it is one of the most important features during classification. We further observed that BERT pays attention to word sizes during training.

Paragraph size is the most significant feature among all others, and it has the highest influence in predictions. This finding aligns with those reported by [18]. GPT usually generates paragraphs consisting of 2 or 3 sentences, while humans tend to write much longer paragraphs. BERT misclassified human-written paragraphs as GPT-generated if the size was considerably shorter than normal human-written paragraphs. Based on this fact, we conclude that BERT pays attention to the differences in paragraph sizes.

GPT-generated text has more punctuation symbols and we do not observe any difference in MTLD values. These findings contradict the results of [4, 22], where they reported that human-written text has a richer vocabulary and uses punctuation symbols more frequently. However, we also observed that human-written text is much more unpredictable according to entropy results. The semantic similarity of paragraphs to a title and paragraphs to paragraphs proved to be another important feature that we could use in predictions.

As we have stated previously, the detection of AI generated text cannot solely rely on machine learning models and the extracted features. These models and features can only be used to guide

a reader through the text, indicating suspicious regions or features. Still, we believe that the final judgment should be made by a human reader.

This study is limited to the analysis of academic content in English written by Turkish students. The same results may not apply to students of other nationalities. Various works [28, 29], report that a writer's native language influences the structure of English texts. For example, [28] reports that French writers often overuse connectors such as however or moreover, while Japanese writers may produce long, cumulative sentences without clear paragraph breaks. Similarly, [29] notes that Spanish and Arabic writers tend to use more elaborate introductions or emotional appeals. Therefore, the results we report, particularly regarding the most important features and stylistic distinctions between human-written and AI-generated content, are valid primarily for Turkish students and may not generalize well to other nationalities. In fact, we argue that AI content detectors should be nation-aware. Grouping all writers of English into the same category may introduce biases against non-native speakers, as previously reported by [16].

It is important to note that the primary focus of this work is the ability to detect AI-generated text written from scratch. Using AI tools as grammar assistants to refine human-written text is a common and widely accepted practice in academia. Some of our features remain unaffected by the use of grammar assistants, as the original text structure and logic are still controlled by the human author. Therefore, key classification features, such as the number of sentences per paragraph, entropy, and paragraph similarity, are unlikely to change during the refinement process. However, features like the use of relative clauses, punctuation, and characters per word may be influenced. Analyzing human-written text after refinement with grammar assistants will be the subject of future work.

## References

- [1] Haenlein M, Kaplan A. A Brief History of Artificial Intelligence: On the Past Present and Future of Artificial Intelligence. *Calif Manag Rev.* 2019;61:5-14.
- [2] Khan AA, Badshah S, Liang P, Khan B, Waseem M, et al. Ethics of AI: A Systematic Literature Review of Principles and Challenges. In: *Proceedings of the 26th international conference on evaluation and assessment in software engineering.* ACM. 2022:383-392.
- [3] Kashnitsky Y, Herrmannova D, de Waard A, Tsatsaronis G, Fennell C, et al. Overview of the DAGPAP22 Shared Task on Detecting Automatically Generated Scientific Papers. *International Conference on Computational Linguistics.* Association for Computational Linguistics. 2022;29:210-213.
- [4] Herbold S, Hautli-Janisz A, Heuer U, Kikteva Z, Trautsch A. A Large-Scale Comparison of Human-Written Versus Chatgpt-Generated Essays. *Sci Rep.* 2023;13:18617–18627.
- [5] Grimaldi G, Ehrler B. AI et al.: Machines Are About to Change Scientific Publishing Forever. *ACS Energy Lett.* 2023;8:878-880.
- [6] Ma Y, Liu J, Yi F, Cheng Q, Huang Y, et al. AI vs. Human – Differentiation Analysis of Scientific Content Generation. 2023. ArXiv preprint: <https://arxiv.org/pdf/2301.10416v1>

- [7] Liu JQ, Hui KT, Al Zoubi F, Zhou ZZ, Samartzis D, et al. The Great Detectives: Humans Versus AI Detectors in Catching Large Language Model-Generated Medical Writing. *Int J Educ Integr.* 2024;20:1-14.
- [8] Gao CA, Howard FM, Markov NS, Dyer EC, Ramesh S, et al. Comparing Scientific Abstracts Generated by ChatGPT to Real Abstracts With Detectors and Blinded Human Reviewers. *npj Digit Med.* 2023;6:75.
- [9] Clark E, August T, Serrano S, Haduong N, Gururangan S, et al. All That's Human Is Not Gold: Evaluating Human Evaluation of Generated Text. In: *Proceedings of the 59th annual meeting of the Association for Computational Linguistics and the 11th international joint Conference on Natural Language Processing.* Association for Computational Linguistics. 2021:7282-7296.
- [10] Ippolito D, Duckworth D, Callison-Burch C, Eck D. Automatic Detection of Generated Text Is Easiest When Humans Are Fooled. In: Jurafsky D, Chai J, Schluter N, Tetreault J, editors. *Proceedings of the 58th annual meeting of the Association for Computational Linguistics.* Stroudsburg PA USA: Association for Computational Linguistics. 2020:1808-1822.
- [11] Berber Sardinha T. Ai-Generated vs Human-Authored Texts: A Multidimensional Comparison. *Appl Corpus Linguist.* 2024;4:100083.
- [12] Labbé C, Labbé D. Duplicate and Fake Publications in the Scientific Literature: How Many Scigen Papers in Computer Science? *Scientometrics.* 2013;94:379-396.
- [13] Cabanac G, Labbé C. Prevalence of Nonsensical Algorithmically Generated Papers in the Scientific Literature. *J Assoc Inf Sci Technol.* 2021;72:1461-1476.
- [14] Cabanac G, Labbé C, Magazinov A. Tortured Phrases: A Dubious Writing Style Emerging in Science. evidence of critical issues affecting established journals. 2021. ArXiv preprint: <https://arxiv.org/pdf/2107.06751>.
- [15] Weber-Wulff D, Anohina-Naumeca A, Bjelobaba S, Foltýnek T, Guerrero-Dib J, et al. Testing of Detection Tools for Ai-Generated Text. *Int J Educ Integr.* 2023;19:1-39..
- [16] Liang W, Yuksekgonul M, Mao Y, Wu E, Zou J. GPT Detectors Are Biased Against Non-native English Writers. *Patterns.* 2023;4:100779.
- [17] Theocharopoulos PC, Anagnostou P, Tsoukala A, Georgakopoulos SV, Tasoulis SK, et al. Detection of Fake Generated Scientific Abstracts. In: *IEEE Ninth International Conference on Big Data Computing Service and Applications.* 2023:33-39.
- [18] Desaire H, Chua AE, Isom M, Jarosova R, Hua D. Distinguishing Academic Science Writing From Humans or ChatGPT With Over 99% Accuracy Using Off-The-Shelf Machine Learning Tools. *Cell Rep Phys Sci.* 2023;4:101426.
- [19] Guo B, Zhang X, Wang Z, Jiang M, Nie J, et al. How Close Is ChatGPT to Human Experts? Comparison Corpus Evaluation and Detection. 2023. ArXiv preprint: <https://arxiv.org/pdf/2301.07597>.
- [20] Crothers E, Japkowicz N, Viktor H, Branco P. Adversarial Robustness of Neuralstatistical Features in Detection of Generative Transformers. In: *International Joint Conference on Neural Networks.* IEEE. 2022:1-8.

- [21] Rodriguez JD, Hay T, Gros D, Shamsi Z, Srinivasan R. Cross-Domain Detection of Gpt-2-Generated Technical Text. In: Proceedings of the north American chapter of the association for computational linguistics: human language technologies. Association for Computational Linguistics. 2022:1213-1233.
- [22] Muñoz-Ortiz A, Gómez-Rodríguez C, Vilares D. Contrasting Linguistic Patterns in Human and Llm-Generated News Text. *Artif Intell Rev.* 2024;57:265.
- [23] Gehrmann S, Strobelt H, Alexander R. GLTR: Statistical Detection and Visualization of Generated Text. In: Proceedings of the 57th annual meeting of the Association for Computational Linguistics: system demonstrations. ACL. 2019.
- [24] Reimers N, Gurevych I. Sentence-Bert: Sentence Embeddings Using Siamese Bertnetworks. In: Proceedings of the 2019 conference on empirical methods in natural language processing. Association for Computational Linguistics. 2019:3980-3990.
- [25] Devlin J, Chang MW, Lee K, Toutanova K. Bert: Pre-training of deep bidirectional transformers for language understanding. In: Proceedings of the 2019 conference of the North American chapter of the association for computational linguistics: human language technologies. ACL. 2019;1:4171-4186.
- [26] van der Maaten L, Hinton G. Visualizing High-Dimensional Data Using T-Sne. *J Mach Learn Res.* 2008;9:2579-2605.
- [27] Lundberg SM, Lee SI. A Unified Approach to Interpreting Model Predictions. *Neural Inf Process Syst.* 2017;30.
- [28] Granger S, Dagneaux E, Meunier F, Paquot M. International Corpus of Learner English. version 2. Presses Universitaires de Louvain. 2009.
- [29] Connor UM. Contrastive Rhetoric: Cross-Cultural Aspects of Second-Language Writing. Cambridge University Press. 1996:220.