electronics

Article

Explainable Prediction of Crowdfunding Success Using Hierarchical Attention Network

SeungHun Lee , Muneeb A. Khan and Hyun-chul Kim *

Department of Software, Sangmyung University, Cheonan 31066, Republic of Korea; mr.leesh90@gmail.com (S.L.); muneebkhan046@gmail.com (M.A.K.)
* Correspondence: hyunchulk@gmail.com

Abstract

Crowdfunding has emerged as an alternative funding source among entrepreneurs, businesses, and industries. In recent years, research on machine learning-based project classification models has been conducted with the aim of predicting the success of crowdfunding campaigns, both for entrepreneurs and investors. However, most of the research has focused on classification approaches using non-content information such as project metadata, creators' behavior, and social history, but there have been few attempts to use text content data per se, particularly in order to provide explanations and evidence for how the prediction decisions were made. To address this point, we propose to use a deep learning-based approach called Hierarchical Attention Network (HAN) to predict the success of crowdfunding campaigns and provide explanation and justification of the prediction decisions using attention weights. We collect publicly available data of crowdfunding campaigns and build our success prediction model with an accuracy of 86.38% and 87.29%, using an *Updates* section and backers' comments in a *Comments* section, respectively. We also explore the feasibility of early success prediction during the funding period (up to 2 months), with as much as 80.99% accuracy in 1 to 2 months. Finally, we examine word and sentence attention weight scores to clarify key factors in predicting crowdfunding success.

Keywords: crowdfunding; success prediction; deep learning approach; hierarchical attention network; explainable AI

1. Introduction

In recent years, crowdfunding has emerged as an alternative to traditional forms of financing for entrepreneurs, businesses, and industries. This method of funding involves collecting sums of money from a large number of individuals rather than relying on traditional venture capital investment. The global crowdfunding industry has seen remarkable growth. As of 2026, The market has an estimated value of $18.54 billion, projected to reach $57.69 billion by 2035 at an average annual growth rate of 11.6% [1].

Crowdfunding can be broadly classified into four types [2]: reward-based, equity-based, debt-based, and donation-based. In reward-based crowdfunding (e.g., Kickstarter.com), supporters receive a reward for their contribution, such as a product or service related to the project. Equity-based crowdfunding allows supporters to receive a share of the company in return for their investment, while debt-based crowdfunding involves supporters loaning money to the project and receiving regular interest payments. In donation-based crowdfunding, supporters make donations to a project without receiving anything in return. One of the most prominent examples of crowdfunding success is the

Academic Editors: Mehdi Ghayoumi and Kambiz Ghazinour

Received: 24 December 2025
Revised: 23 January 2026
Accepted: 27 January 2026
Published: 28 January 2026

Pebble Watch campaign on Kickstarter, which raised more than $10 million [3]. This campaign set the precedent for what could be achieved through crowdfunding and has since been followed by numerous successful campaigns across various industries, including technology, arts, and social causes. The average success rate of crowdfunding projects on Kickstarter, calculated as the percentage of projects that successfully achieve their funding goals, is 40%, despite the significant growth in the number of projects and funding on crowdfunding platforms [4]. This indicates that the outcome of a crowdfunding campaign is not always certain and is influenced by various factors, such as marketing strategy and market demand. Hence, understanding the factors contributing to the success of crowdfunding and predicting the likelihood of project success is one of the most prominent research challenges in the field of crowdfunding.

Previous research in this area has focused mainly on using non-content information such as project metadata or statistical information (e.g., goal money amount, geographical location, the number of updates, etc.) and campaign creators' behavior, social, or history information, as well as classifying successful projects using traditional machine learning algorithms including logistic regression [2], Hidden Markov Models and Support Vector Machines [5], and decision trees [6]. Our previous work was the first to demonstrate the feasibility of accurately predicting project success using only text content, without utilizing any project metadata or statistically-derived numerical features [7].

Building on previous research (1), this study expands the dataset size and category range by refining the sampling criteria. The dataset comprises 5384 projects from both the technology and art categories, including only those with at least one update and one comment, whereas previous studies analyzed a randomly selected sample of 2000 projects from the technology category alone. (2) We systematically optimize the hyperparameters of the Hierarchical Attention Network (HAN), experimentally determining the optimal number of sentences and words per sentence for input. To further ensure balanced evaluation across multiple data splits, cross-validation was used. This rigorous and scientific approach enhances both the robustness and generalizability of the results. (3) Finally, by leveraging attention weights, this work identifies the most informative words and sentences for accurate prediction, providing insights into the reasoning behind prediction decisions. Contributions of this paper are listed as follows:

1. In this study, we collected publicly available crowdfunding campaign data from one of the most popular crowdfunding sites, Kickstarter.com. We trained and evaluated success prediction models on two different major categories of crowdfunding campaigns—Technology and Art, using text content data only.
2. We propose adopting HAN [8] to predict crowdfunding success by effectively modeling the multi-level structure (word and sentence) of campaign texts. Our model achieves state-of-the-art level accuracies of 86.38% and 87.29% using only textual data from the *Updates* and *Comments* sections, respectively.
3. We also explore the feasibility of early prediction of campaign success using text content data only, showing our model achieves 59.55% of prediction accuracy on the very first day of campaign launch. As more contents arrive in *Updates* and *Comments* sections later on, our model achieves 70.96% to 80.99% (with *Updates* section text only) and 65.99% to 74.49% (with backers' comments in *Comments* section text only), within one to two months.
4. This paper presents an empirical analysis of attention weight scores at both word and sentence levels generated by HAN, to examine which words and sentences in *Updates* and *Comments* sections affected the prediction results (i.e., the success of the campaigns) most. For instance, we found that in *Updates* section of successful projects, sentences where the creator explicitly requests backers to share project information

via social media platforms like Facebook and Twitter are prominent. In *Comments* section of failed projects, sentences expressing dissatisfaction due to the creator's lack of response to investors' messages stand out. We believe that our findings provide valuable insights for both researchers and practitioners, such as project creators and backers, helping them make informed decisions and increase the success rate of their campaigns and investments.

The organization of this paper is as follows: In Section 2, a review of related work is provided. The data set and the methodology are described in Section 3. The experimental setup of our model, results from training the model, and visualizations of word/sentence attention scores are covered in Section 4. The paper concludes in Section 5.

2. Related Work

Crowdfunding has recently attracted significant interest from academia and industry and is used for various purposes, including funding for artistic activities, philanthropic endeavors, and startup funding. As a result, researchers have studied various aspects of crowdfunding, including motivation, success factors, fraud detection, and project recommendations. Mollick [2] examined the characteristics that influence the success of crowdfunding projects and discovered a correlation between social networks, project quality, and the success of crowdfunding campaigns. He found that platform design aspects such as user experience, usability, and information accessibility play a vital role in influencing the behavior of both entrepreneurs and investors, and ultimately influence the success of a crowdfunding campaign. Greenberg and Mollick [9] investigated the impact of the composition of the founding team on the success of new ventures. They compared the performance of solo-founded enterprises with that of team-founded ventures and found that solo-founded enterprises have a higher likelihood of survival and generate comparable or higher revenue than team-founded ventures, particularly those founded by two-person teams.

The success of crowdfunding campaigns is largely influenced by online social information and relationships [10], such as the sharing of project information through social networks such as Twitter [5,11], the presence of information on Facebook and the number of Facebook friends [2], as well as offline relationships with friends, relatives, and family members [12]. Project videos play a very important role in the success of crowdfunding [2,13,14]. Ma and Palacios [13] found that by coding and analyzing crowdfunding campaign videos, entrepreneurs can improve fundraising results by reducing video pitch distractions, appearing earlier, and minimizing public speaking anxiety. Furthermore, contextual information depicted regarding the "workspace" (e.g., "office," "toys", "teamwork", and "collaboration," etc.) was determined to be advantageous for the success of the crowdfunding campaign [14]. Other factors that greatly impact the success of a campaign include providing regular and high-quality updates on progress [15], the creator's backing history [16], the connection between the creator and the backers [17], and the creator's personal characteristics [18,19]. Yuan et al. [4] introduced a hierarchical framework to identify diverse persuasion tactics, including the creator's credibility, emotional appeals, and logical reasoning, within project descriptions. Geiger and Moore [20] identified the number of backers as a primary mediating mechanism, demonstrating that campaign characteristics such as text volume, videos, and positive tone fundamentally enhance funding outcomes by attracting a larger crowd. Zhang and Lau [21] combined multimodal features of text, audio, and video with explicit cues such as backers count and creator history to improve crowdfunding success prediction using CNN, RNN, and BERT.

One of the primary reasons projects fail on Kickstarter is the lack of enough potential investors [22]. To address this, An et al. [22] and Rakesh et al. [17] created automatic recommendation systems that match projects with potential investors based on various factors, such as project characteristics, personal details, location, and network connections. Gerber et al. [23] found that social connectivity, which strengthens commitment to an idea through creator's opinions and feedback, and belonging to a community that shares interests and values among backers, has become an important factor in attracting investors' interest.

Much research has also been conducted to analyze and detect fraudulent crowdfunding campaigns using linguistic features. Siering et al. [24] have demonstrated the feasibility of detecting fraudulent crowdfunding projects using linguistic characteristics, with studies showing the correlation between the use of linguistic and content-based cues and the accuracy of detecting fraudulent campaigns. Cumming et al. [25] examined the probability of detecting fraudulent campaigns, and found that the likelihood of detecting fraudulent campaigns is also significantly correlated with factors such as the details of the campaign description, the background of the creators, the available social information and the characteristics of the campaign. Lee et al. [26] found that linguistic characteristics such as the number of words and sentences, personal pronouns, and spatial-temporal words were useful in detecting fraudulent projects.

Cheng et al. [27] studied predicting the success of crowdfunding projects using multimodal deep learning techniques and multiple data modalities (i.e., images, videos, and textual descriptions) from the project description. Yu et al. [28] collected a dataset of crowdfunding projects and associated characteristics, including project information, backer information, and project updates, then trained a deep neural network model (i.e., Multi-Layer Perceptron (MLP)) to predict the success of the projects. Shi et al. [29] proposed a novel deep learning method based on audio analytics that can extract audio features to predict the fundraising results of projects, with 83.8% F1 score. In our recent study, we used a sequence-to-sequence (seq2seq) deep neural network model with sentence-level attention and HAN to classify successful projects with 89% accuracy [7].

Our approach includes the following novelties: (i) we utilize raw text content data from *Campaign*, *Updates*, and *Comments* sections to train our model and categorize successful projects, whereas prior research primarily examined non-content information such as project metadata or statistical information (e.g., goal money amount, geographical location, the number of updates, etc.), and campaign creators' behavior, social, or history information to classify successful or fraudulent crowdfunding projects. (ii) We employ the HAN, which has demonstrated superior performance in Natural Language Processing (NLP) tasks, instead of traditional machine learning algorithms. (iii) We use the attention weight scores of sentences and words to show and explain which words and sentences were most informative and helpful for accurate classification. In contrast, previous studies primarily focused on classification using deep learning methods without providing a detailed explanation of why the classification was made, lacking a way to clarify the reasons behind it.

3. Methodology

The overall workflow of the proposed framework for predicting and interpreting crowdfunding success is illustrated in Figure 1.

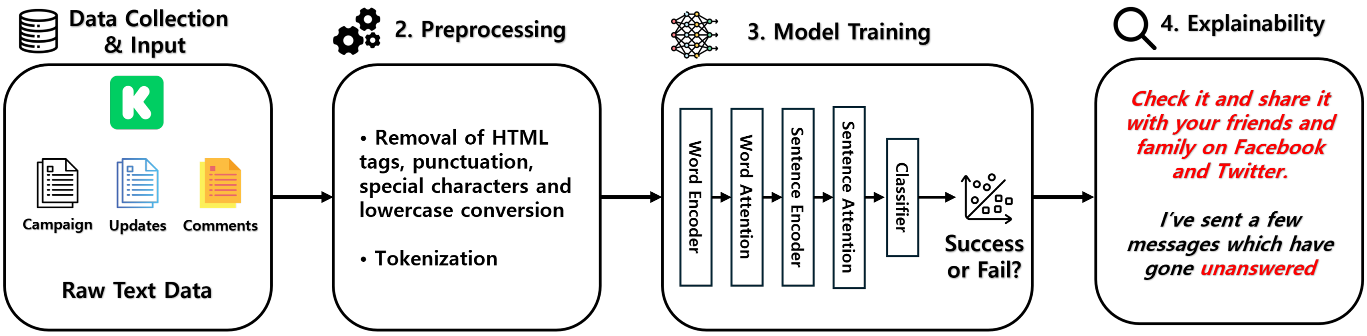

Figure 1. Overall workflow of the proposed crowdfunding success prediction framework.

3.1. Data Set

In our study, we collected 344,544 campaigns from Kickstarter.com, one of the most popular and publicly available crowdfunding sites, between April 2009 and October 2018. Table 1 presents the distribution of campaigns across 15 main categories defined by Kickstarter, with the percentage of each outcome status (Successful, Failed, Canceled, Suspended) shown in parentheses. Each category exhibits distinct characteristics such as funding goals, objectives, and reward types [2,30,31]. For example, the technology category has a relatively low success rate (22.2%) with a strong emphasis on engineering and mathematical objectives, while the Art category shows a higher success rate (42.9%) with more artistic, cultural, and emotional emphasis. Therefore, we chose the Technology and Art categories for comparison and examined how the contrasting characteristics of these two categories influence success.

Table 1. The total number of campaigns and the percentage of campaigns in each category (2009–2017). The percentage in parentheses represents the proportion of each status within that category (e.g., Successful(%) = number of successful projects/total # of campaigns in that category).

Category	Successful (%)	Failed (%)	Canceled (%)	Suspended (%)	Total # of Campaigns
Film & Video	23,298 (40.0)	30,374 (52.2)	4445 (7.6)	93 (0.2)	58,210
Publishing	12,976 (34.1)	22,342 (58.7)	2687 (7.1)	55 (0.1)	38,060
Music	17,757 (46.2)	17,996 (46.9)	2562 (6.7)	91 (0.2)	38,406
Games	14,556 (39.7)	16,187 (44.2)	5730 (15.6)	183 (0.5)	36,656
Technology	6836 (22.2)	19,670 (63.8)	3931 (12.8)	370 (1.2)	30,807
Design	11,804 (40.4)	14,658 (50.2)	2541 (8.7)	182 (0.6)	29,185
Art	11,441 (42.9)	13,383 (50.2)	1782 (6.7)	69 (0.3)	26,675
Fashion	6330 (32.9)	10,860 (56.5)	1911 (9.9)	111 (0.6)	19,212
Food	4817 (27.0)	11,526 (64.6)	1399 (7.8)	105 (0.6)	17,847
Comics	6349 (55.4)	4367 (38.1)	718 (6.3)	23 (0.2)	11,457
Photography	3549 (34.6)	5894 (57.5)	755 (7.4)	46 (0.4)	10,244
Theater	5873 (57.6)	3854 (37.8)	445 (4.4)	17 (0.2)	10,189
Crafts	2358 (27.6)	5549 (64.9)	581 (6.8)	60 (0.7)	8548
Journalism	1114 (21.9)	3352 (66.0)	558 (11.0)	57 (1.1)	5081
Dance	2463 (62.1)	1314 (33.1)	175 (4.4)	15 (0.4)	3967
Total	131,521 (38.2)	181,326 (52.6)	30,220 (8.8)	1477 (0.4)	344,544

The crowdfunding campaign consists of three key sections: the *Campaign* section, where the project creators present their ideas with videos and images; the *Updates* section, where they keep the backers informed of progress; and the *Comments* section, which serves as a platform for communication between the project creators and the backers. Therefore, we collect all available text content from all sections. Our goal is to classify successful projects by analyzing their content and context based on raw text, thus raw text content data is required. Therefore, we selected campaigns that had at least one text entry in both the *Updates* and *Comments* sections for analysis. Among the technology campaigns, there were 6055 successful and 3688 failed campaigns with at least one update and comment, while in the art category, there were 6898 successful and 1346 failed campaigns. To balance the data set, we randomly undersampled based on the number of failed art campaigns (i.e., 1346), resulting in 2692 campaigns selected from each category of technology and art.

3.2. Hierarchical Attention Network

We propose using HAN to tackle the challenge of accurately predicting the success of crowdfunding campaigns [8]. HAN is designed to understand the hierarchical structure of a document through a combination of Recurrent Neural Networks (RNNs) and an attention mechanism. This architecture consists of four stages: word encoding, word attention, sentence encoding, and sentence attention. In the first stage, a word vector is generated using a Bidirectional Gated Recurrent Unit (Bi-GRU) [32] capable of understanding both the forward and backward contexts. In the second stage, the most relevant information is identified through word focus weights, which are then used to compute the final word vector by combining the word vector and its attention weight. In the third and fourth stages, sentence focus weights and the sentence vectors are calculated using a process similar to that of the word encoding phase. By leveraging the strengths of RNNs and the attention mechanism, HAN can effectively capture the hierarchical structure of a document and extract key information for accurate prediction.

3.2.1. Gated Recurrent Unit (GRU)

The GRU [33] is a type of RNN designed to address the vanishing gradient problem. This issue arises when a model struggles to retain information across long input sequences, making it difficult to capture long-term dependencies. GRU mitigates this problem through gating mechanisms that regulate the flow of information, allowing the model to retain essential information while discarding irrelevant data.

The GRU contains two gates: the update gate z_t and the reset gate r_t, which work together to control the flow of information from the previous hidden state h_{t-1} and the current input x_t to the current hidden state h_t. These gates enable the GRU to selectively preserve or discard information, thereby improving its ability to learn long-term dependencies and make accurate predictions. The update gate z_t determines how much past information from the previous hidden state h_{t-1} is retained versus how much new information from the current input x_t is incorporated. The update gate z_t is calculated as shown in Equation (1):

$$z_t = \sigma(W_z x_t + U_z h_{t-1} + b_z) \tag{1}$$

where σ is the sigmoid function, W_z is the weight matrix for the input in the update gate, U_z is the weight matrix for the previous hidden state in the update gate, and b_z is the bias vector. The reset gate r_t controls how much influence the previous hidden state h_{t-1} has on the candidate hidden state. When r_t is close to zero, the model effectively ignores the previous hidden state. The reset gate r_t is computed according to Equation (2):

$$r_t = \sigma(W_r x_t + U_r h_{t-1} + b_r) \tag{2}$$

where W_r is the weight matrix for the input in the reset gate, U_r is the weight matrix for the previous hidden state in the reset gate, and b_r is the bias vector. The candidate hidden state $\widetilde{h_t}$ is then determined as expressed in Equation (3):

$$\widetilde{h_t} = tanh(W_h x_t + r_t \odot (U_h h_{t-1}) + b_h) \tag{3}$$

where $\odot$ is the element-wise product. Finally, the current hidden state h_t is updated as shown in Equation (4):

$$h_t = (1 - z_t) \odot h_{t-1} + z_t \odot \widetilde{h_t} \tag{4}$$

We utilize a Bi-GRU for sequence modeling in our neural network. Unlike conventional RNNs, which process sequences only in a single direction, Bi-GRU captures both forward and backward contexts of the sequence. This is achieved using two GRU units: one encodes the sequence in the forward direction, and the other encodes it in reverse. The two GRU outputs are then concatenated to construct the final representation, effectively integrating bidirectional information. This final hidden representation serves as input to subsequent layers for downstream tasks such as text classification or machine translation.

3.2.2. Word Encoder

HAN is based on document-level classification. Given a document consisting of L sentences s_i, where the s_ith sentence has T_i words represented as w_{it}, with $t \in [1, T_i]$. We go on to detail a step-by-step procedure for building a document-level vector through the utilization of a hierarchical structure by progressively combining word vectors.

For a given sentence containing words w_{it}, where $t \in [1, T_i]$, the words are first transformed into vectors using the FastText [34] embedding matrix W_e. The transformation is defined as shown in Equation (5):

$$x_{it} = W_e w_{it}, t \in [1, T_i] \tag{5}$$

We utilized a Bi-GRU to obtain annotations for words. This process summarizes information from both directions for each word and incorporates contextual information into the annotations. The Bi-GRU consists of a forward GRU $\overrightarrow{f}$ and a backward GRU $\overleftarrow{f}$. The forward GRU reads the sentence s_i starting from w_{i1} to w_{iT_i}, while the backward GRU reads it from w_{iT_i} to w_{i1}. The forward and backward hidden states are calculated as shown in Equation (6) and Equation (7), respectively:

$$\overrightarrow{h_{it}} = \overrightarrow{\mathrm{GRU}}(x_{it}), t \in [1, T_i] \tag{6}$$

$$\overleftarrow{h_{it}} = \overleftarrow{\mathrm{GRU}}(x_{it}), t \in [T_i, 1] \tag{7}$$

We get an annotation for a given word w_{it} by concatenating forward hidden state $\overrightarrow{h_{it}}$ and backward hidden state $\overleftarrow{h_{it}}$ into one representation, $h_{it} = [\overrightarrow{h_{it}}, \overleftarrow{h_{it}}]$. This representation summarizes the information of the entire sentence with w_{it} as the center.

3.2.3. Word Attention

In our approach, we recognize that certain words carry more weight in determining the meaning of a sentence. To identify these keywords, we employ an attention mechanism to focus on the most impactful words in the sentence. These keywords are used to generate a sentence vector that represents the meaning of the sentence. To obtain the hidden representation u_{it} of the word annotation h_{it}, we apply a MLP. This calculation is shown in Equation (8):

$$u_{it} = tanh(W_w h_{it} + b_w) \tag{8}$$

The importance of each word is determined by calculating the similarity between u_{it} and the context vector at the word level u_w. We obtain a normalized importance weight α_{it} for each word using a softmax function, as expressed in Equation (9):

$$\alpha_{it} = \frac{\exp(u_{it}^\top u_w)}{\sum_t \exp(u_{it}^\top u_w)} \tag{9}$$

The sentence vector s_i is then calculated as a weighted sum of the word annotations based on the importance weights, as shown in Equation (10):

$$s_i = \sum_t \alpha_{it} h_{it} \tag{10}$$

3.2.4. Sentence Encoder

To derive the document vector from the sentence vector s_i, we utilize a Bi-GRU to encode the sentences. The forward and backward hidden states for each sentence are calculated as shown in Equation (11) and Equation (12), respectively:

$$\overrightarrow{h_i} = \overrightarrow{\text{GRU}}(s_i), i \in [1, L] \tag{11}$$

$$\overleftarrow{h_i} = \overleftarrow{\text{GRU}}(s_i), i \in [L, 1] \tag{12}$$

The annotation of each sentence i is produced by concatenating the forward hidden state $\overrightarrow{h_i}$ and the backward hidden state $\overleftarrow{h_i}$, effectively encapsulating the information within the sentence. The sentence vector $h_i = [\overrightarrow{h_i}, \overleftarrow{h_i}]$ is then used to summarize the information of both neighboring sentences, providing a comprehensive representation of the sentence i in the context of the surrounding sentences.

3.2.5. Sentence Attention

We have integrated the attention mechanism to identify and emphasize sentences that play a pivotal role in the accurate classification of a document. The mechanism employs a sentence-level context vector u_s, which assesses the significance of each sentence in the document. The hidden representation u_i and the normalized importance weight α_i are calculated as shown in Equation (13) and Equation (14), respectively:

$$u_i = tanh(W_s h_i + b_s) \tag{13}$$

$$\alpha_i = \frac{\exp(u_i^\top u_s)}{\sum_i \exp(u_i^\top u_s)} \tag{14}$$

The document vector v is then generated as a weighted sum of sentence annotations, as expressed in Equation (15):

$$v = \sum_i \alpha_i h_i \tag{15}$$

The result of this evaluation is then used to generate a document vector v that summarizes all the information from the sentences and serves as a high-level representation for the document classification task. This document vector v is used to compute the final classification probabilities. Specifically, we pass the document vector through a softmax layer to obtain the probability distribution over the possible classes, as shown in Equation (16):

$$p = \text{softmax}(W_c v + b_c) \tag{16}$$

The training loss is defined as the negative log likelihood of the true labels, as defined in Equation (17):

$$L = -\sum_{d} \log p_{dj} \tag{17}$$

3.3. Performance Metrics

We utilize four evaluation metrics to assess the accuracy of our success campaign detection model: accuracy, AUC, precision, and recall.

- (Overall) accuracy: the ratio of the projects correctly classified as successful or failed out of the total number of projects contained in our data set. We use this measure to assess the accuracy of the classifier across the entire dataset, as shown in Equation (18).
- AUC: AUC is the "Area Under the Receiver Operating Characteristic (ROC) curve". ROC is a probability curve that plots the True Positive Rate (TPR) against the False Positive Rate (FPR) at various classification thresholds. AUC measures the overall performance of the classifier by calculating the area under the ROC curve, with a value ranging from 0 to 1. An AUC score of 1 indicates a perfect classifier, meaning it correctly distinguishes all positive and negative instances across all thresholds.

$$Accuracy = \frac{TP + TN}{TP + TN + FP + FN} \tag{18}$$

The next two metrics aim to evaluate the effectiveness of classification performance, particularly in identifying successful campaigns.

- Precision: the ratio of True Positives to the sum of True Positives and False Positives, indicating the proportion of predicted successful campaigns that were actually successful. True Positives are the number of correctly classified successful campaigns, False Positives are failed projects falsely labeled as success. This calculation is expressed in Equation (19):
- Recall: the ratio of True Positives to the sum of True Positives and False Negatives, indicating the proportion of actual successful campaigns that were correctly identified. False Negatives represent successful campaigns that were incorrectly classified as failed, leading to missed positive cases. The recall is determined as shown in Equation (20):

$$Precision = \frac{TP}{TP + FP} \tag{19}$$

$$Recall = \frac{TP}{TP + FN} \tag{20}$$

4. Results

4.1. Experimental Settings

To process our documents for analysis, we first applied several preprocessing steps to clean the raw text data crawled from Kickstarter's HTML pages. This included removing HTML tags, eliminating excessive whitespace, converting all text to lowercase, and removing special characters. After preprocessing, we use the Natural Language Toolkit (NLTK) [35] to divide them into sentences and then break each sentence into its constituent tokens. NLTK offers a comprehensive suite of tools for tokenization, stemming, and part-of-speech tagging, as well as the tools to build and implement models for text categorization, semantic reasoning, and information extraction. We retain only the 50,000 most frequently used words to construct the vocabulary, replacing all other words with a designated UNK token.

Figure 2a shows the distribution of the number of sentences and the number of words in a sentence in each category section. Most sections contain fewer than 50 sentences, with the maximum number of sentences (excluding outliers) being approximately 100. Typically, the technology category contains more sentences than art in all sections. In particular, we observed that the *Comments* section in the technology category had more sentences, suggesting a more active conversation between backers and creators. Figure 2b shows a similar pattern for the distribution of the number of words in a sentence, with an average of around 20 words (and a maximum of almost 50). When implementing the HAN algorithm, it is necessary to set a maximum sentence length for embedding and a maximum number of words per sentence. To determine the optimal parameters, experiments were conducted with sentence lengths ranging from 10 to 100 (i.e., 10, 20, 30, 50, 100) and word counts per sentence ranging from 10 to 50 (i.e., 10, 20, 30, 40, 50) to determine their impact on the campaign success prediction. Based on these experiments, we selected a sentence length of 100 and a word count of 30 as the optimal parameters. Detailed experimental results and the rationale for these parameter choices are described in Appendix A and the following section.

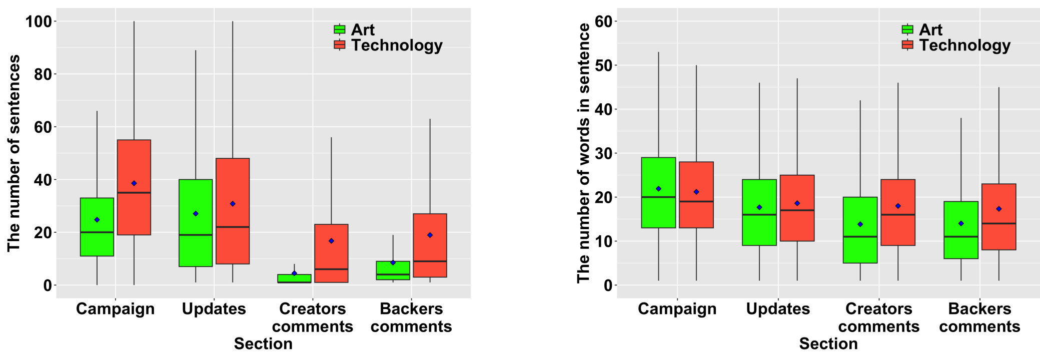

(**a**) (**b**)

Figure 2. A box plot visualizing the distribution of sentence and word counts per sentence across different category sections. (**a**) The number of sentences; (**b**) The number of words in a sentence. The horizontal line of the box represents the median, and the blue circle indicates the mean. The first quartile (25th percentile) and third quartile (75th percentile) are, respectively, represented by the bottom and top of the box. The whiskers of the box plot extend from either end of the box to the minimum and maximum values, with outliers excluded.

We used a 300-dimensional pre-trained FastText model (crawl-300d-2M) [34] to obtain word embeddings for our text data. In our experiment, the dimension of the GRU was set at 150 and then doubled to 300 through the combination of forward and backward GRUs. The context vectors for words and sentences also had a dimension of 300. The training process for our models used 20 epochs with early stopping and a batch size of 32. We used the RMSprop optimization algorithm with a learning rate of 0.001. We employed Keras's "ReduceLROnPlateau" [36] and "ModelCheckpoint" [37] callback APIs to regulate the learning rate and maintain the optimal model. The learning rate was reduced by a factor of 0.2 if no improvement was detected after three consecutive epochs, continuing until the lower limit was reached. The initial learning rate was 1×10^{-3}, with a lower bound of 1×10^{-5}. The "ModelCheckpoint" function saved the best model based on the validation loss at each epoch.

To ensure the robustness of our results and prevent data leakage, we employed 10-fold cross-validation. The dataset was randomly partitioned into ten disjoint folds, and each iteration used a training, validation, and test split ratio of 8:1:1. This process was repeated ten times so that every project in our dataset was evaluated as a test set exactly once. L2 regularization and a dropout ratio of 0.2 were applied. All experiments were implemented using Keras 3.12, TensorFlow 2.20, and Python 3.10. Model training and evaluation were performed on a system equipped with an NVIDIA Quadro RTX 6000 GPU (24GB VRAM).

4.2. Classification Performance Evaluation

Figures A1–A3 illustrate that the accuracy of project success prediction was influenced by both the number of sentences and the number of words per sentence. Analysis of backers' comments in the *Comments* section for the technology category (as shown in Figure A2) revealed that the highest prediction accuracy was achieved with 30 words per sentence, with accuracy improvements from 70.25% to 87.29% as the number of sentences increased from 10 to 100. As a result, 30 words per sentence and 100 sentences were identified as the optimal parameters for achieving the highest accuracy.

As shown in Table 2, the backers' comments in *Comments* section for the technology category yielded the highest results with a precision of 91.15%, recall of 82.66%, an overall accuracy of 87.29%, and an AUC score of 0.941. Furthermore, the "technology + art" and art categories showed the highest accuracy in the *Updates* section, with 86.38% and 80.85% successful projects classification, respectively. This result highlights the importance of engaging in dialogue and taking into account the feedback of backers who have invested in the project as it progresses, as well as consistently providing updates on the project's progress. The technology category outperformed the art category in all sections (excluding the *Campaign* section), which is likely due to the larger number of sentences in the technology (as shown in Figure 2) providing more information to train the classification model. Further analysis of the *Campaign* section's average image usage revealed that technology projects utilized an average of 16.84 images per campaign, while art projects employed 7.25 images per campaign. The crucial information in images, particularly in technology projects such as product photographs, technical diagrams, and prototype demonstrations, helps explain the relatively lower prediction accuracy in the *Campaign* section. The *Updates* section classified the successful projects with high accuracy, particularly with precision scores of 86.68% and 82.36% for technology and art, respectively. However, creators' comments in *Comments* section for the art had a limited impact on the prediction of success, with an accuracy of only 59.56%. As shown in Figure 2, the average number of sentences per creator's comments in art was 4.46, which was relatively low compared to the other sections of both categories. This suggested that increasing the amount of textual content, particularly the number of sentences, may be necessary to build an effective prediction model.

Table 2. Classification performance of each section by category.

Category	Content	Precision	Recall	Accuracy	AUC
Technology + Art	*Campaign*	60.23	59.41	59.55	0.640
	Updates	**84.31**	**89.47**	**86.38**	**0.925**
	Creators' comments in *Comments*	66.90	79.85	69.47	0.762
	Backers' comments in *Comments*	84.03	82.08	83.09	0.908
Technology	*Campaign*	65.70	58.14	61.63	0.688
	Updates	86.68	**84.14**	85.43	0.920
	Creators' comments in *Comments*	78.58	79.38	78.59	0.854
	Backers' comments in *Comments*	**91.15**	82.66	**87.29**	**0.941**
Art	*Campaign*	61.69	69.13	62.38	0.679
	Updates	**82.36**	**80.26**	**80.85**	**0.898**
	Creators' comments in *Comments*	56.79	79.54	59.56	0.643
	Backers' comments in *Comments*	76.91	78.36	77.26	0.836

Table 3 displays a comparison of various classification methods that have been used to predict the success of crowdfunding campaigns based on different features. The methods used different types of data such as project and creator behavior information, social network traits like Twitter and Facebook, as well as videos, images, audio, and text data available in the *Campaign*, *Updates*, and *Comments* sections. Yu et al. [28] achieved the highest accuracy of 93.2% using project and creator information as features with the MLP method. The multimodal method employed by Shi et al. [29] achieved the highest F1 score of 0.838 by utilizing project and creator information along with audio features from *Campaign* section. In contrast, our method based on HAN performed well, achieving an accuracy of 87.29% and an AUC score of 0.941, using only raw text content from each *Campaign*, *Updates*, or *Comments* as features. These results suggested that raw text content data alone can effectively predict crowdfunding success. Furthermore, they indicated that our model could be further improved by incorporating additional features beyond textual content, such as non-content information (e.g., project metadata, goal amount, geographical location, number of updates) and campaign creators' behavioral, social, or historical information.

Table 3. Performance comparison of classification methods for predicting campaign success using different features.

Paper	Method	Features	Accuracy	F1
Chung and Lee (2015) [38]	AdaboostM1	Project & Creator info. + Twitter	84.2%	–
Shi et al. (2021) [29]	DNN	Project & Creator info. + Audio from *Campaign*	–	0.838
Cheng et al. (2019) [27]	Multimodal (CNN, BoW)	Project info. + Text & Image from *Campaign*	–	0.753
Lai, Lo and Hwang (2017) [39]	XGBoost	Project info. + Text from *Updates* and *Comments*	92.4%	–
Kaminski and Hopp (2020) [40]	Logistic Regression	Video & Speech & Text from *Campaign*	72%	0.720
Yu et al. (2018) [28]	MLP	Project & Creator info.	93.2%	–
Yuan et al. (2023) [4]	PSM-PEM	Proejct info. + Text from *Campaign*	86.6%	0.799
Zhang and Lau (2024) [21]	CNN, RNN, BERT	Project & Creator info. + Video & Audio & Text from *Campaign*	82.2%	–
Ours	**HAN**	**Raw text from *Campaign* OR *Updates* OR *Comments***	**87.29%**	**0.867**

Kickstarter follows an "all-or-nothing" policy [41], meaning that if a project does not reach its target funding within the set time frame, which can range from one to two months, the project is considered a failure and the creators receive nothing. Crowdfunding success prediction plays a crucial role in the decision-making process for creators and backers, as it provides valuable information about a project's potential. For creators, an accurate success prediction can help them assess their chances of securing funding and adjust their campaign strategies to increase their chances of success. For backers, predicting the success of a project enables them to save time by avoiding investments in projects with low funding prospects and identifying projects with higher quality and greater potential.

As shown in Table 2, on the very first day of the project start, using the campaign description available in the technology and art categories, the project was successfully classified with an accuracy of 59.55%. Figure 3 illustrates the average classification performance as time passes from the start of the project. As shown in Figure 3g–i, in the case of *Updates* and *Comments* sections, due to the lack of text being uploaded in the immediate aftermath of the project start, the classification was not effective. But over time, the accuracies of classifying successful projects using *Updates* section and backers' comments in *Comments* section for the "technology + art" categories have increased significantly from 70.96% to 80.99% and 65.99% to 74.49% in one to two months, surpassing the classification accuracy of the *Campaign* section. This result indicated that the *Updates* section and backers' comments in the *Comments* section played the most important role and effectively classified even when relying solely on raw text content. Figure 3a–f demonstrate a consistent increase in precision across all sections over time. Notably, the technology category demonstrated superior precision compared to the art and "technology + art" categories throughout all sections, with improvements of up to 12.43% and 15.63% over the "technology + art" and art categories, respectively. In contrast, the recall for both creators' and backers' comments in *Comments* section did not improve during the initial 10 days, indicating that the model was not effectively learning in the early phase. However, as time progressed and additional data was incorporated, recall began to gradually increase after one month, indicating an improvement in the model's generalization capabilities. This pattern reflected the model's evolving ability to discern complex patterns as it processed more diverse data, ultimately achieving a better balance between precision and recall.

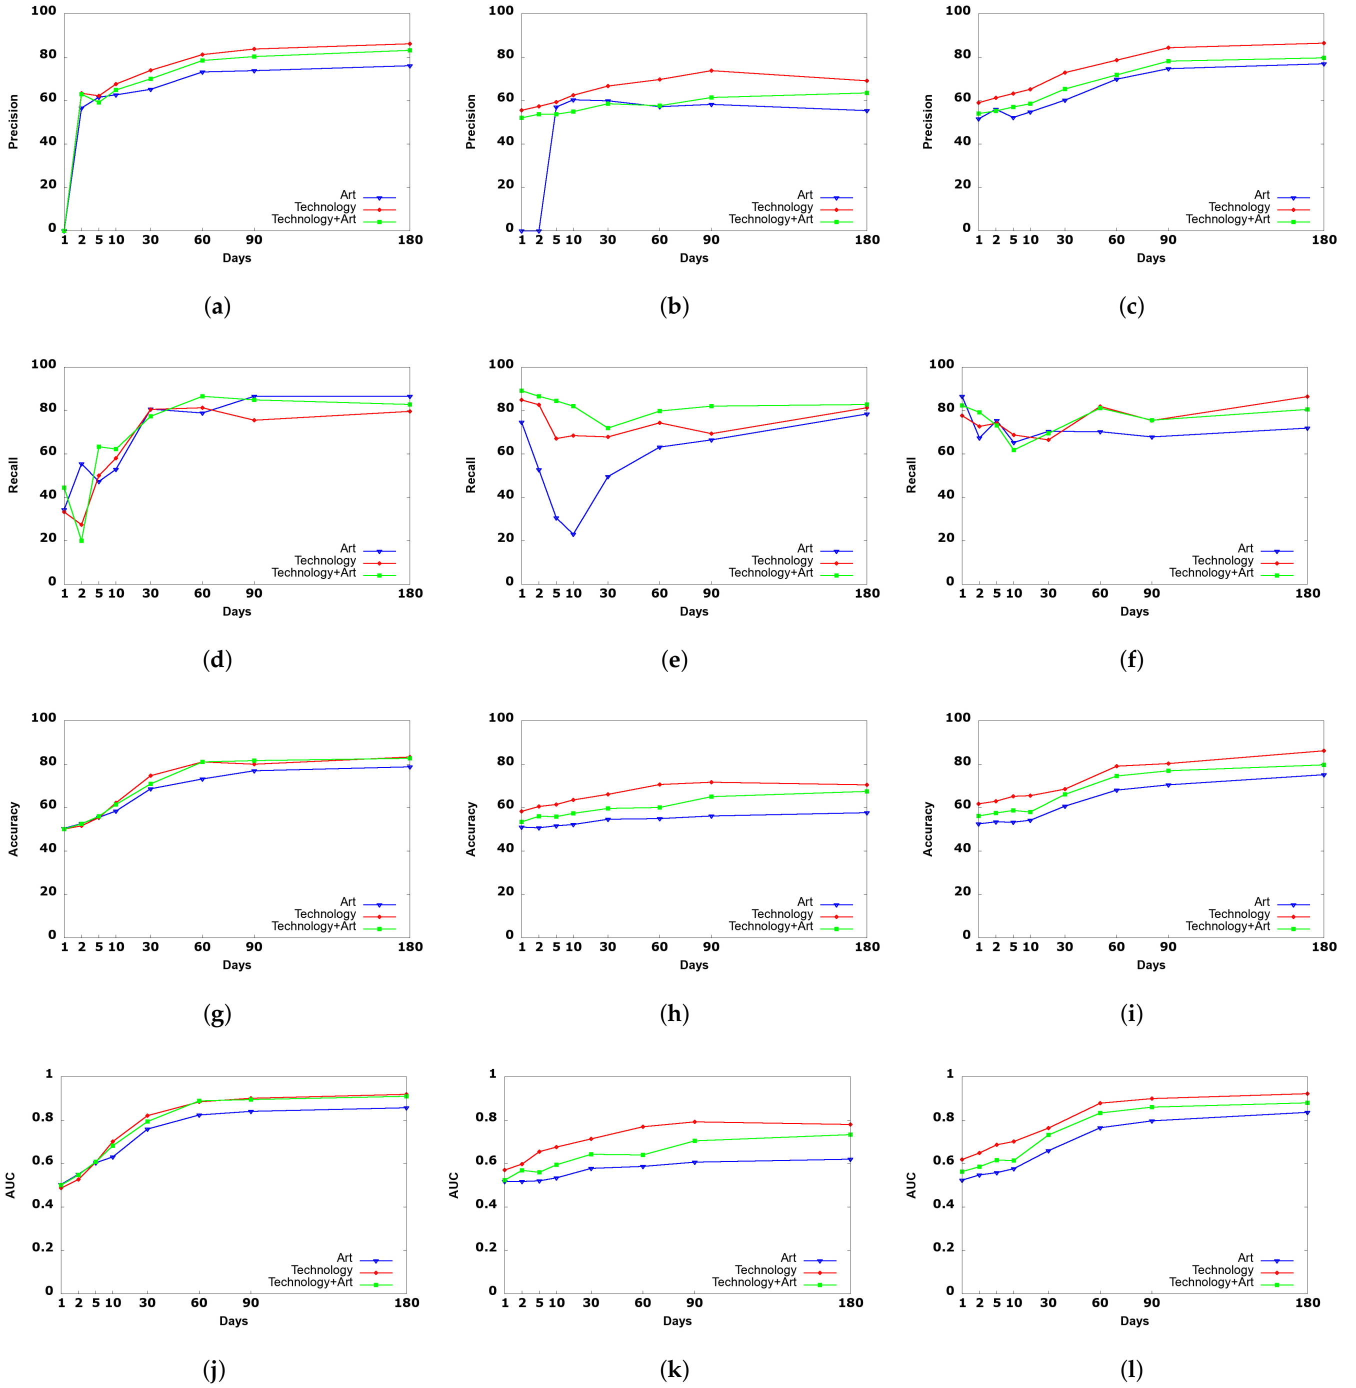

Figure 3. Estimated average classification performance vs. elapsed time (days). (**a**–**c**) Precision for updates, creator comments, and backers comments; (**d**–**f**) recall; (**g**–**i**) (overall) accuracy; (**j**–**l**) AUC.

From a temporal perspective, previous research has emphasized the predictive power of early campaign dynamics. Etter et al. [5] demonstrated that integrating financial trajectories, such as cumulative pledged money and the number of backers, can achieve over 76% accuracy within just 4 h of a campaign's launch. Similarly, Chung and Lee [38] developed a comprehensive model by combining static project metadata and social media signals with these temporal funding patterns, achieving 83.6% accuracy within the initial 10% of the project duration. However, while these approaches offer practical value for early outcome monitoring, it is important to note that financial trajectories are inherently linked to the success criterion itself. Consequently, predicting success based on these trends can be viewed as tracking the progress of an outcome rather than identifying the underlying drivers of success.

In contrast, our results showed that while text-only prediction exhibited a 'cold start' characteristic, accuracy improved as the *Updates* and *Comments* sections accumulated more content over one to two months, reaching up to 80.99%. Our methodology demonstrated that text-only data alone could maintain predictive power without relying on auxiliary variables such as financial trajectories, social capital, or project metadata. Furthermore, as demonstrated in Section 4.3, the analysis of attention weights enabled us to pinpoint the specific sentences and words most influential in the prediction process.

4.3. Explainability Analysis

The HAN model utilizes an attention mechanism that allows it to focus on specific parts of the input data, making it easier to understand the reasoning behind its predictions and decisions. By using context-dependent words/sentences attention weights, we uncover which words and sentences were most useful in our successful project prediction model. Even the same words or sentences can have low or high levels of attention weights depending on preceding or succeeding words/sentences and order or content, because attention weights are context dependent [8]. As shown in Table 2, the *Updates* section and backers' comments in the *Comments* section yielded the best classification results, leading us to focus on these two sections for a more detailed analysis. In the following examples, sentences with particularly high word/sentence attention weights are marked in red. The sentences and words highlighted in red represent the highest attention weights at the document and sentence levels, respectively. Specifically, red sentences indicate the most influential sentences within the document, while red words denote the most influential words within each sentence.

4.3.1. Updates Section

Xu et al. [15] found that *Updates* section is important to predict successful projects, particularly in promoting the project to backers and encouraging the spread of project information and rumors through social networks. In addition, maintaining updates on the project's progress and providing information on new products are also crucial. The creators wrote an update asking to share the information with friends and family through Facebook and Twitter, which had the highest sentence weight score among several sentences in our model. An example would be:

> *"Our team is well-versed in knowing how to mitigate privacy and security issues. Read and share the article about MirroCool in KnowTechie We are pretty excited about this article! Check it and share it with your friends and family on Facebook and Twitter. Here is the link of the news article <url>"*

Phrases such as "milestone reached", "update", "halfway", and other terms associated with progress updates, as identified by Xu et al. [15], were found to be effective indicators for forecasting success. An example would be:

> *"I'm excited to start delivering your rewards in the new repair truck soon! Best, Pete Major milestone reached"*

A detailed and reliable description of the product, particularly in the technology category, which enhances backers' trust by explaining the product's functionality in an understandable and factual way. An example would be:

> *"The de-limer works by electrolysis. This is where a small DC current is passed though a liquid.The de-limer consists of two anodes and one cathode. The middle part of the device is the cathode and the top and bottom parts are the anodes. The anodes must be in contact with the metal of the device to be cleaned. If there is a layer of limescale between the de-limer and your kettle then it will not work."*

The failed projects contain only expressions of gratitude and lack information on project progress and products. This aligns with Xu et al. [15], who found that while simple appreciation messages have little impact on success, providing details about new rewards and project progress is more influential. The high attention score assigned to the appreciation messages further underscores this issue, suggesting that failed projects tend to prioritize thanking backers over delivering substantial updates or incentives. An example would be:

> *"Thank you for all the support! It's been 1 week now since we've launched our Kickstarter and it's been a very exciting ride. We wanted to say thank you so much for your pledges, they really mean a lot to Court and the team. We are going to continue pushing forward and doing all we can to spread the message of recovery to those who need it. If you have any questions for us you can ask them in the FAQ section and we'll answer them as quickly as we can. Thanks again for your support and the support you give to all those who are struggling today with addiction."*

4.3.2. Backers' Comments in Comments Section

Tables 4 and 5 provide examples of sentences and words that play a crucial role in predicting successful and failed projects based on backers' comments. The key content in backers' comments primarily includes feedback on product delivery, reviews of the product, and discussions related to its features. Additionally, some comments criticize the creator's behavior and management of the project, often revealing issues with communication and trust.

Table 4. Examples of sentences and words with high attention weights in success project backers' comments.

"after battery replacement worked two weeks well"
"beautiful artwork and craftmanship"
"let's get this thing funded."
"have they all been shipped now"
"a cool 20 physical reward like maybe a print would have been great"
"Mar 27 2018 or 8 weeks "Add ons Today, we FINISHED shipping all the add ons (apart from maker kits and additional batteries)"
"I am sure that this project has a great potential"

Table 5. Examples of sentences and words with high attention weights in failure project backers' comments.

"I've sent a few messages which have gone unanswered."
"he have tried to use the trust of the backers in order for him to obtain an external investors funding"
"I have backed almost 500 campaigns, most have not made me angry. You need to learn some customer service. I will be withdrawing my support for your campaign—even though it is most likely the pledges will be refunded anyway when you do not meet your goal."
"I haven't received my order"
"good luck with the campaign"

5. Conclusions

In this paper, we collected publicly available data from Kickstarter, one of the most prominent crowdfunding sites, to evaluate our success prediction model on technology and art projects. We employ the HAN model to predict the success or failure of crowdfunding campaigns. By leveraging only text content data from updates and backers' comments, our model achieves high accuracy of 86.38% and 87.29%, respectively. Our results also demonstrate the feasibility of early project success detection with a maximum accuracy of 80.99% within 1 to 2 months. Through the examination of word/sentence attention weight scores, our model provides explanations and justifications for the classification. For example, (1) in the *Updates* section, creators explicitly requested backers to share project information on social media platforms like Facebook and Twitter, and (2) in the *Comments* section, concerns about the creators' lack of response to investors' messages were observed. This study contributes to the field of crowdfunding by providing a novel and effective solution to help entrepreneurs make informed decisions and increase their campaigns' success rates.

This study has important implications for predicting the success of crowdfunding campaigns; however, several limitations remain. This study focused on only two categories and utilized a restricted dataset of 6898 campaigns that necessarily included at least one Update and one Comment. Consequently, the model faces challenges in predicting success for projects lacking such textual interactions, and the temporal coverage does not reflect the most recent market trends. To address these limitations, we aim to advance this work in future studies by incorporating a broader range of categories and a larger volume of projects. As Kickstarter encompasses various diverse categories beyond Technology and Art, analyzing text content across all these domains will enable a comprehensive understanding of the factors contributing to campaign success. This approach will allow us to identify category-specific trends and analyze how text content impacts success in each context, accounting for the varying patterns revealed in previous research [2,42]. Moreover, we plan to extend this framework by applying transformer-based or large language models, such as BERT [43], GPT [44], PaLM [45], to more recent datasets extending up to 2025. These models are particularly adept at capturing long-range dependencies [43,46], and fine-tuning them for the specific downstream task of crowdfunding prediction is expected to further enhance performance and generalizability. Lastly, while our study achieved a high accuracy of 87.29% using only raw text content data from the backers' comments in *Comments* section, integrating additional data types such as multimedia data (e.g., video [40], image [27], audio [29]), project metadata [2,6], statistical information [5,47], and campaign creators' behavior, social, or history information [15,16,39], which have been identified as advantageous in previous research, could potentially enhance our model.

Author Contributions: Conceptualization, S.L., and H.-c.K.; methodology, S.L., and H.-c.K.; software, S.L., and H.-c.K.; validation, S.L., and H.-c.K.; investigation, S.L., and H.-c.K.; resources, H.-c.K.; data curation, S.L.; writing—original draft preparation, S.L., and H.-c.K.; writing—review and editing,

S.L., M.A.K., and H.-c.K.; visualization, S.L., and H.-c.K.; supervision, H.-c.K.; project administration, H.-c.K. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by a 2023 Research Grant from Sangmyung University, South Korea.

Data Availability Statement: The data will be made available by the authors upon request.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

HAN	Hierarchical Attention Network
MLP	Multi-Layer Perceptron
seq2seq	sequence-to-sequence
RNNs	Recurrent Neural Networks
Bi-GRU	Bidirectional Gated Recurrent Unit
GRU	Gated Recurrent Unit
ROC	Area Under the Receiver Operating Characteristic
TPR	True Positive Rate
FPR	False Positive Rate
NLTK	Natural Language Toolkit
NLP	Natural Language Processing

Appendix A. Performance Comparison According to the Number of Sentences and the Number of Words in a Sentence in Each Section by Category

Appendix A.1. Technology + Art

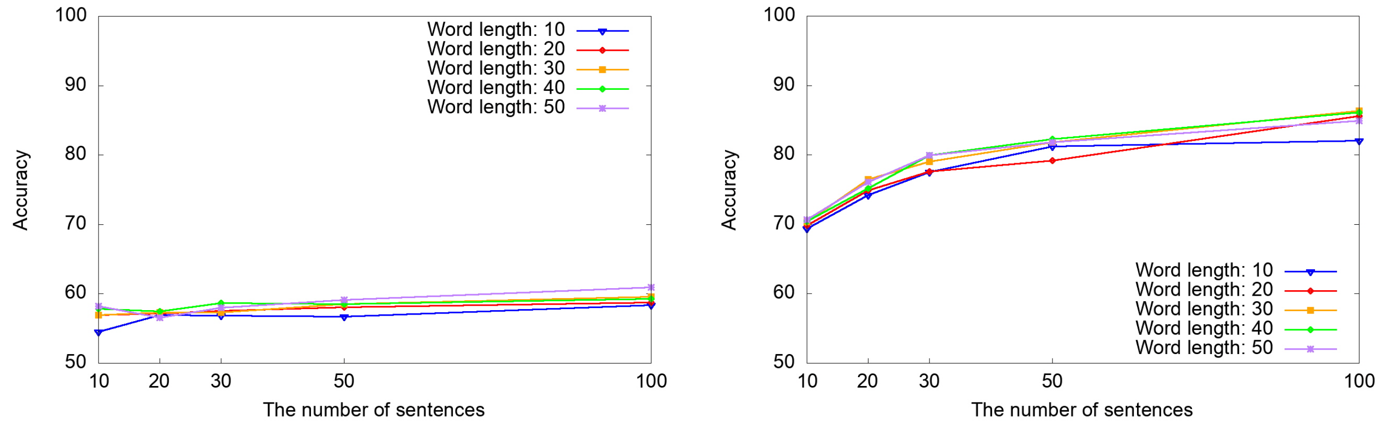

(**a**) *Campaign* (**b**) *Updates*

Figure A1. *Cont.*

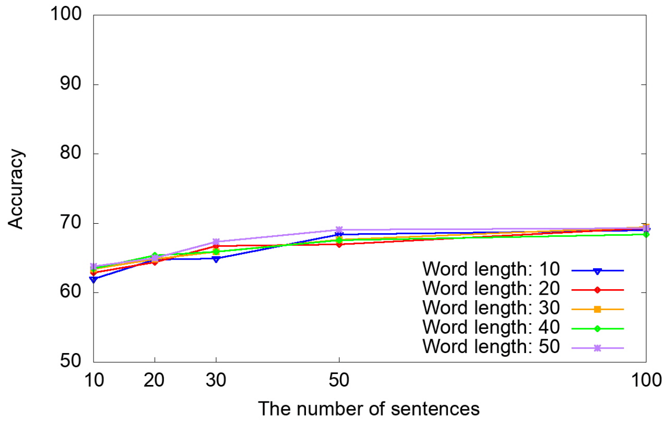

(**c**) *Creators comments*

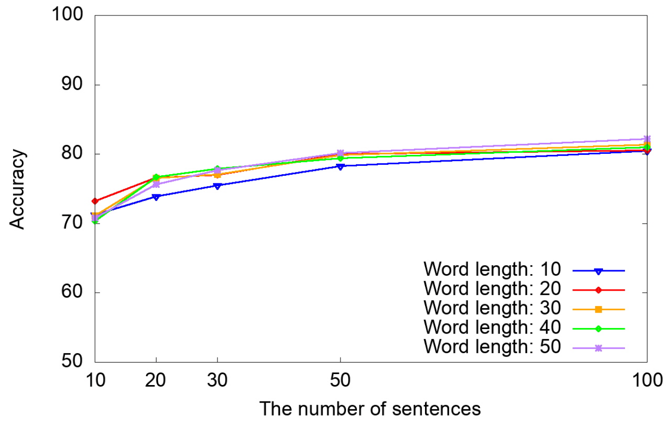

(**d**) *Backers comments*

Figure A1. Comparison of performance according to the number of sentences and the number of words in a sentence in each section of combined Technology and Art.

Appendix A.2. Technology

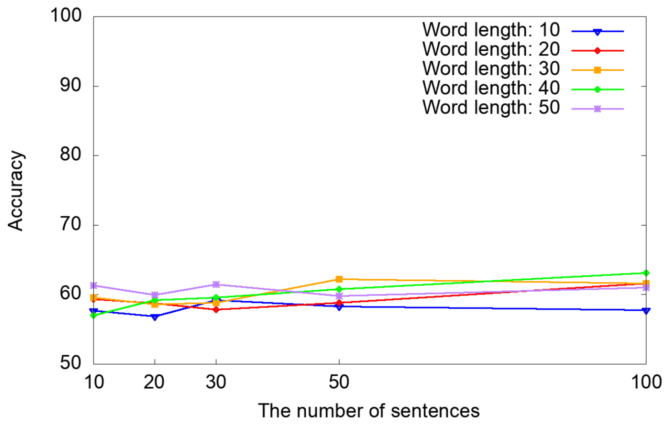

(**a**) *Campaign*

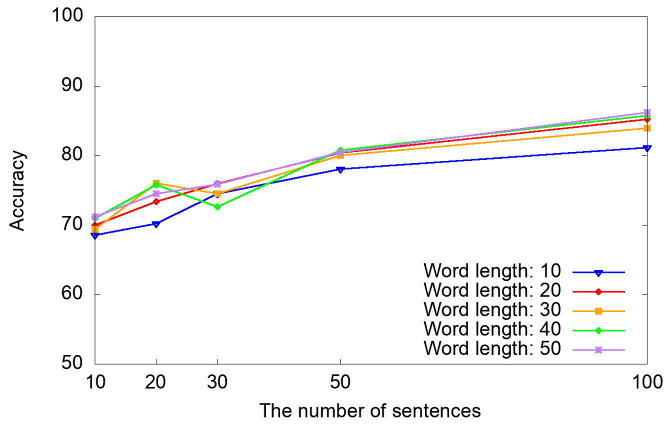

(**b**) *Updates*

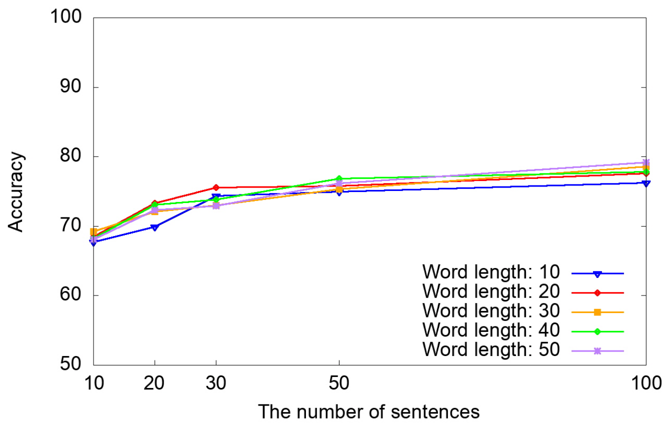

(**c**) *Creators comments*

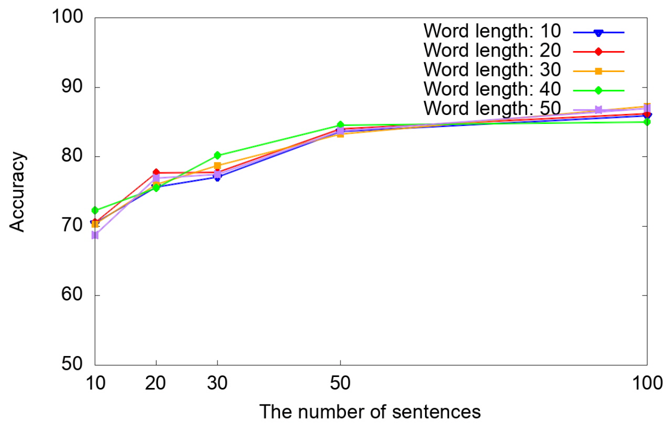

(**d**) *Backers comments*

Figure A2. Comparison of performance according to the number of sentences and the number of words in a sentence in each section of Technology.

Appendix A.3. Art

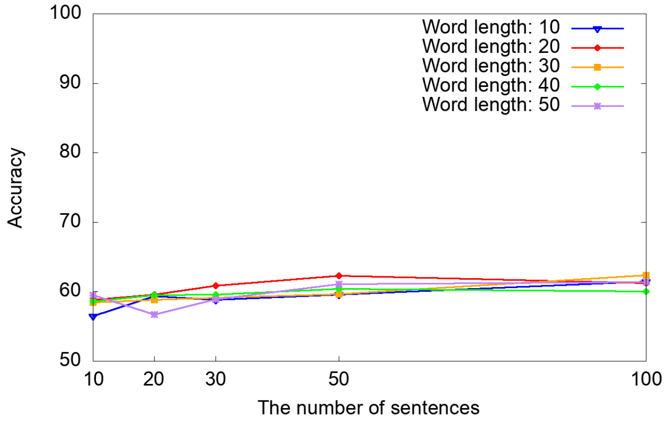

(**a**) *Campaign*

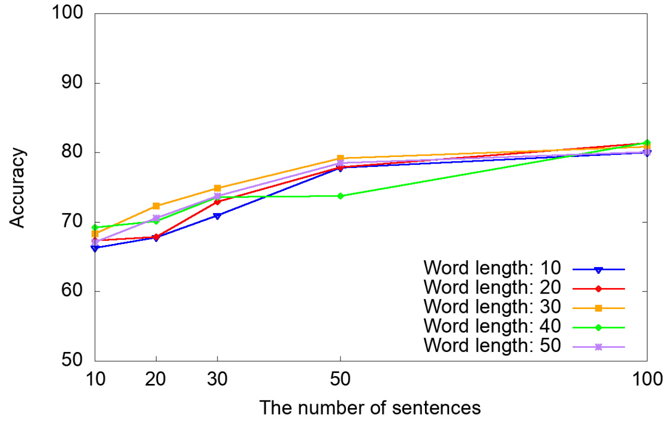

(**b**) *Updates*

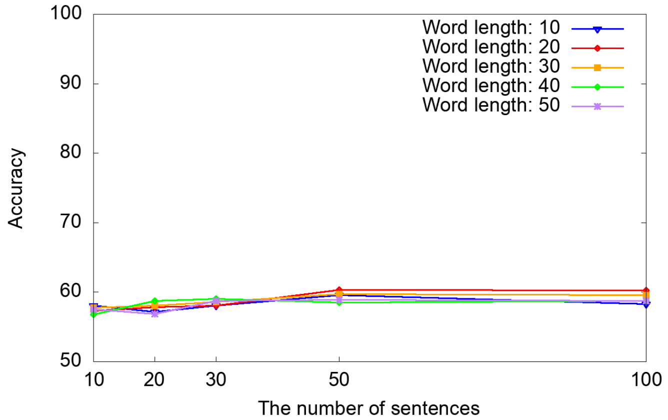

(**c**) *Creators comments*

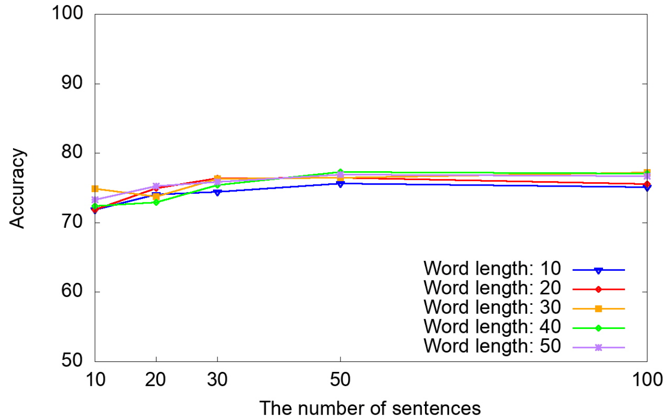

(**d**) *Backers comments*

Figure A3. Comparison of performance according to the number of sentences and the number of words in a sentence in each section of Art.

References

1. Crowdfunding Market Size & Share [2026–2035]. Available online: https://www.businessresearchinsights.com/market-reports/crowdfunding-market-100103 (accessed on 19 January 2026).
2. Mollick, E. The dynamics of crowdfunding: An exploratory study. *J. Bus. Ventur.* **2014**, *29*, 1–16. https://doi.org/10.1016/j.jbusvent.2013.06.005.
3. Pebble Time—Awesome Smartwatch, No Compromises by Pebble Technology—Kickstarter. Available online: https://www.kickstarter.com/projects/getpebble/pebble-time-awesome-smartwatch-no-compromises (accessed on 19 December 2025).
4. Yuan, H.; Qiao, D.; Wei, Q. A Persuasion-Based Framework for Predicting Crowdfunding Success. 2023. Available online: https://ssrn.com/abstract=4658668 (accessed on 21 January 2026).
5. Etter, V.; Grossglauser, M.; Thiran, P. Launch Hard or Go Home!: Predicting the Success of Kickstarter Campaigns. In Proceedings of the ACM Conference on Online Social Networks (COSN), Boston, MA, USA, 7–8 October 2013; pp. 177–182. https://doi.org/10.1145/2512938.2512957.
6. Greenberg, M.D.; Pardo, B.; Hariharan, K.; Gerber, E. Crowdfunding support tools: Predicting success & failure. In Proceedings of the Conference on Human Factors in Computing Systems-Extended Abstracts (CHI-EA), Paris, France, 27 April–2 May 2013; pp. 1815–1820. https://doi.org/10.1145/2468356.2468682.
7. Lee, S.; Lee, K.; Kim, H.-C. Content-based Success Prediction of Crowdfunding Campaigns: A Deep Learning Approach. In Proceedings of the ACM Conference on Computer Supported Cooperative Work and Social Computing (CSCW), Jersey City, NJ, USA, 3–7 November 2018; pp. 3–7. https://doi.org/10.1145/3272973.3274053.

8. Yang, Z.; Yang, D.; Dyer, C.; He, X.; Smola, A.; Hovy, E. Hierarchical attention networks for document classification. In Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL HLT), San Diego, CA, USA, 12–17 June 2016; pp. 1480–1489. https://doi.org/10.18653/v1/N16-1174.
9. Greenberg, J.; Mollick, E. Sole Survivors: Solo Ventures Versus Founding Teams. 2018. Available online: https://ssrn.com/abstract=3107898 (accessed on 18 December 2025).
10. Zhang, Y.; Tan, C.D.; Sun, J.; Yang, Z. Why do people patronize donation-based crowdfunding platforms? An activity perspective of critical success factors. *Comput. Hum. Behav.* **2020**, *112*, 106470. https://doi.org/10.1016/j.chb.2020.106470.
11. Aleksina, A.; Akulenka, S.; Lubloy, A. Success factors of crowdfunding campaigns in medical research: Perceptions and reality. *Drug Discov. Today* **2019**, *24*, 1413–1420. https://doi.org/10.1016/j.drudis.2019.05.012.
12. Agrawal, A.K.; Catalini, C.; Goldfarb, A. *The Geography of Crowdfunding*; Technical Report; National Bureau of Economic Research: Cambridge, MA, USA, 2011. https://doi.org/10.3386/w16820.
13. Ma, Z.; Palacios, S. Image-mining: Exploring the impact of video content on the success of crowdfunding. *J. Mark. Anal.* **2021**, *9*, 265–285. https://doi.org/10.1057/s41270-021-00133-8.
14. Li, Y.; Xiao, N.; Wu, S. The devil is in the details: The effect of nonverbal cues on crowdfunding success. *Inf. Manag.* **2021**, *58*, 103528. https://doi.org/10.1016/j.im.2021.103528.
15. Xu, A.; Yang, X.; Rao, H.; Fu, W.-T.; Huang, S.-W.; Bailey, B.P. Show Me the Money!: An Analysis of Project Updates during Crowdfunding Campaigns. In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI), Toronto, ON, Canada, 26 April–1 May 2014; pp. 591–600. https://doi.org/10.1145/2556288.2557045.
16. Zvilichovsky, D.; Inbar, Y.; Barzilay, O. Playing Both Sides of the Market: Success and Reciprocity on Crowdfunding Platforms. 2015. Available online: https://ssrn.com/abstract=2304101 (accessed on 18 December 2025).
17. Rakesh, V.; Choo, J.; Reddy, C.K. Project recommendation using heterogeneous traits in crowdfunding. In Proceedings of the 9th International AAAI Conference on Web and Social Media (ICWSM), Oxford, UK, 26–29 May 2015; pp. 26–29. https://doi.org/10.1609/icwsm.v9i1.14624.
18. Thies, F.; Wessel, M.; Rudolph, J.; Benlian, A. Personality matters: How signaling personality traits can influence the adoption and diffusion of crowdfunding campaigns. In Proceedings of the 24th European Conference on Information Systems (ECIS), Istanbul, Turkey, 12–15 June 2016.
19. Moreno-Moreno, A.; Sanchis-Pedregosa, C.; Berenguer, E. Success factors in peer-to-business (P2B) crowdlending: A predictive approach. *IEEE Access* **2019**, *7*, 148586–148593. https://doi.org/10.1109/ACCESS.2019.2946858.
20. Geiger, M.; Moore, K. Attracting the crowd in online fundraising: A meta-analysis connecting campaign characteristics to funding outcomes. *Comput. Hum. Behav.* **2022**, *128*, 107061. https://doi.org/10.1016/j.chb.2021.107061.
21. Zhang, Z.; Lau, R.Y. Exploiting multimodal features and deep learning for predicting crowdfunding successes. In Proceedings of the 2024 IEEE International Conference on Omni-layer Intelligent Systems (COINS), London, UK, 29–31 July 2024; pp. 1–6. https://doi.org/10.1109/COINS61597.2024.10622132.
22. An, J.; Quercia, D.; Crowcroft, J. Recommending Investors for Crowdfunding Projects. In Proceedings of the 23rd International Conference on World Wide Web (WWW), Seoul, Republic of Korea, 7–11 April 2014; pp. 261–270. https://doi.org/10.1145/2566486.2568005.
23. Gerber, E.M.; Hui, J.S.; Kuo, P.Y. Crowdfunding: Why people are motivated to post and fund projects on crowdfunding platforms. In Proceedings of the ACM SIGCHI Conference on Computer-Supported Cooperative Work & Social Computing (CSCW), Seattle, WA, USA, 11–15 February 2012; p. 10.
24. Siering, M.; Koch, J.-A.; Deokar, A.V. Detecting fraudulent behavior on crowdfunding platforms: The role of linguistic and content-based cues in static and dynamic contexts. *J. Manag. Inf. Syst.* **2016**, *33*, 421–455. https://doi.org/10.1080/07421222.2016.1205930.
25. Cumming, D.; Hornuf, L.; Korami, M.; Schweizer, D. Disentangling Crowdfunding from Fraudfunding. *J. Bus. Ethics* **2021**, *182*, 1103–1128. https://doi.org/10.1007/s10551-021-04942-w.
26. Lee, S.; Shafqat, W.; Kim, H.-C. Backers Beware: Characteristics and Detection of Fraudulent Crowdfunding Campaigns. *Sensors* **2022**, *22*, 7677. https://doi.org/10.3390/s22197677.
27. Cheng, C.; Tan, F.; Hou, X.; Wei, Z. Success prediction on crowdfunding with multimodal deep learning. In Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI), Macao, China, 10–16 August 2019; pp. 10–16. https://doi.org/10.24963/ijcai.2019/299.
28. Yu, P.F.; Huang, F.M.; Yang, C.; Liu, Y.H.; Li, Z.Y.; Tsai, C.H. Prediction of Crowdfunding Project Success with Deep Learning. In Proceedings of the IEEE 15th International Conference on e-Business Engineering (ICEBE), Xi'an, China, 12–14 October 2018; pp. 12–14. https://doi.org/10.1109/ICEBE.2018.00012.
29. Shi, J.; Yang, K.; Xu, W.; Wang, M. Leveraging deep learning with audio analytics to predict the success of crowdfunding projects. *J. Supercomput.* **2021**, *77*, 7833–7853. https://doi.org/10.1007/s11227-020-03595-2.
30. Mollick, E. Delivery Rates on Kickstarter. 2015. Available online: https://ssrn.com/abstract=2699251 (accessed on 18 December 2025).

31. Dey, S.; Duff, B.; Karahalios, K.; Fu, W.T. The art and science of persuasion: Not all crowdfunding campaign videos are the same. In Proceedings of the ACM Conference on Computer Supported Cooperative Work Companion (CSCW), Portland, OR, USA, 25 February–1 March 2017; pp. 755–769. https://doi.org/10.1145/2998181.2998229.
32. Schuster, M.; Paliwal, K.K. Bidirectional recurrent neural networks. *IEEE Trans. Signal Process.* **1997**, *45*, 2673–2681. https://doi.org/10.1109/78.650093.
33. Cho, K.; Merrienboer, B.v.; Bahdanau, D.; Bengio, Y. On the properties of neural machine translation: Encoder-decoder approaches. *arXiv* **2014**, arXiv:1409.1259.
34. Bojanowski, P.; Grave, E.; Joulin, A.; Mikolov, T. Enriching Word Vectors with Subword Information. *Trans. Assoc. Comput. Linguist.* **2017**, *5*, 135–146. https://doi.org/10.1162/tacl_a_00051.
35. NLTK Natural Language Toolkit. Available online: https://www.nltk.org/ (accessed on 19 December 2025)
36. ReduceLROnPlateau. Available online: https://keras.io/api/callbacks/reduce_lr_on_plateau/ (accessed on 19 December 2025)
37. ModelCheckpoint. Available online: https://keras.io/api/callbacks/model_checkpoint/ (accessed on 19 December 2025)
38. Chung, J.; Lee, K. A long-term study of a crowdfunding platform: Predicting project success and fundraising amount. In Proceedings of the 26th ACM Conference on Hypertext & Social Media, Guzelyurt, Cyprus, 1–4 September 2015; pp. 1–4. https://doi.org/10.1145/2700171.2791045.
39. Lai, C.-Y.; Lo, P.-C.; Hwang, S.-Y. Incorporating comment text into success prediction of crowdfunding campaigns. In Proceedings of the 21st Pacific-Asia Conference on Information Systems (PACIS), Langkawi, Malaysia, 16–20 July 2017.
40. Kaminski, J.C.; Hopp, C. Predicting outcomes in crowdfunding campaigns with textual, visual, and linguistic signals. *Small Bus. Econ.* **2020**, *55*, 627–649. https://doi.org/10.1007/s11187-019-00218-w.
41. Cumming, D.J.; Leboeuf, G.; Schwienbacher, A. Crowdfunding Models: Keep-It-All vs. All-or-Nothing. 2019. Available online: https://ssrn.com/abstract=2447567 (accessed on 18 December 2025).
42. Horisch, J. Crowdfunding for environmental ventures: An empirical analysis of the influence of environmental orientation on the success of crowdfunding initiatives. *J. Clean. Prod.* **2015**, *107*, 636–645. https://doi.org/10.1016/j.jclepro.2015.05.046.
43. Devlin, J.; Chang, M.-W.; Lee, K.; Toutanova, K. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv* **2018**, arXiv:1810.04805.
44. OpenAI. GPT-4 Technical Report. *arXiv* **2023**, arXiv:2303.08774.
45. Google. PaLM 2 Technical Report. *arXiv* **2023**, arXiv:2305.10403.
46. Brown, T.B.; Mann, B.; Ryder, N. Language models are few-shot learners. *arXiv* **2020**, arXiv:2005.14165.
47. Li, Y.; Rakesh, V.; Reddy, C.K. Project Success Prediction in Crowdfunding Environments. In Proceedings of the Ninth ACM International Conference on Web Search and Data Mining (WSDM), San Francisco, CA, USA, 22–25 February 2016; pp. 247–256. https://doi.org/10.1145/2835776.2835791.