CERM: Context-Aware Literature-Based Discovery via Sentiment Analysis

Julio Christian Young ^{b;*} and Uchenna Akujuobi ^a

^aSony AI, Sony Research Inc.

^bGraduate School of Advanced Science and Technology, Tokyo Denki University ORCiD ID: Julio Christian Young https://orcid.org/0000-0002-0483-370X, Uchenna Akujuobi https://orcid.org/0000-0002-7102-7994

Abstract. Driven by the abundance of biomedical publications, we introduce a sentiment analysis task to understand food-health relationship. Prior attempts to incorporate health into recipe recommendation and analysis systems have primarily focused on ingredient nutritional components or utilized basic computational models trained on curated labeled data. Enhanced models that capture the inherent relationship between food ingredients and biomedical concepts can be more beneficial for food-related research, given the wealth of information in biomedical texts. Considering the costly data labeling process, these models should effectively utilize both labeled and unlabeled data. This paper introduces Entity Relationship Sentiment Analysis (ERSA), a new task that captures the sentiment of a text based on an entity pair. ERSA extends the widely studied Aspect Based Sentiment Analysis (ABSA) task. Specifically, our study concentrates on the ERSA task applied to biomedical texts, focusing on (entity-entity) pairs of biomedical and food concepts. ERSA poses a significant challenge compared to traditional sentiment analysis tasks, as sentence sentiment may not align with entity relationship sentiment. Additionally, we propose CERM, a semi-supervised architecture that combines different word embeddings to enhance the encoding of the ERSA task. Experimental results showcase the model's efficiency across diverse learning scenarios.

1 Introduction

We are what we eat - Ludwig Feuerbach. This quote means that the food we eat can have either positive or negative effects on us. Hence, we should aim to consume food that brings us both health and happiness [8]. In recent years, the interest in healthy lifestyles and healthy diets has surged, especially after the pandemic [10, 22]. Ideally, advice regarding an appropriate diet for an individual should be consulted with a nutritionist, considering their expert knowledge for planning a personalized and sustainable diet according to their needs. However, consulting with a nutritionist can often be timeconsuming and require additional funds, making it less accessible. Previous studies have attempted to develop diet recommendation systems by training various artificial intelligence (AI) models on the massive food-related data available [4, 5, 12, 13, 20, 21]. While some studies focus on dietary recommendation systems targeting specific audiences (e.g., dietary planning for people with certain health conditions [5, 21] or age groups [4, 20]), few studies target a broader audience [12, 13].

In [12], researchers developed a dietary recommendation system considering various factors such as demographic information, health conditions, and nutritional needs. The system accurately predicted meal suitability by representing it as a binary classification problem. However, due to its reliance on a neural network for prediction, the system lacked the ability to provide explanations for its decisions. Another study by Kim & Chung [13] combined symbolic AI and neural network approaches to produce a hybrid system that achieves the same goal but can explain its decisions. Due to the additional inference process by the symbolic AI approach, the resulting recommendations can be inherently explained. Although [13] has successfully addressed the explainability issue in the recommendation system, there is still room for improvement to produce a better solution. Previous studies [12, 13] have primarily focused on macronutrients and their impact on human health, neglecting the significance of micronutrients. However, the recommendation process should also consider essential factors like food-disease relationships, nutritional components, and ingredients, as outlined in [17]. Neglecting these components may lead to unintended consequences, such as recommending allergenic meals to users or suggesting the consumption of micronutrients that are unsuitable for a user's medical condition.

A dietary recommendation system that can consider all the information in a meal concerning the user's medical condition can be realized based on two other supporting systems. The first system should be responsible for breaking down all the nutritional information in a meal. Then, based on such information, the second could infer the relationship between each element and the user's medical condition. Research on the first system has been extensively conducted, but research on the second is limited. In addition, research in this field often prioritizes the development of databases [7, 18, 23, 25] that summarize existing relationships rather than adaptive systems that can learn and incorporate new information. Arguably, such a system is much needed given that research related to food science is still developing; thus, knowledge about the effects of compounds on diseases can change based on the latest studies.

With abundant food science-related research articles, learning relationships between food ingredients and biomedical concepts can be structured as a literature-based sentiment analysis task. Specifically, this task involves learning the relationship between two entities, e_1 and e_2 , given a short text (sentence) s where the entities are men-

^{*} Corresponding Author. Email: julio.christian.young@gmail.com.

tioned. This problem is different from a regular sentiment analysis task since the sentiment of s may not necessarily represent the sentiment of e_1 and e_2 . For instance, consider a sentence "*The daily consumption of Ginger is beneficial to the body but ineffective for Diarrhea.*". Although the overall sentiment of the sentence is positive, the relationship between "*Ginger*" and "*Diarrhea*" have a contrasting sentiment. This example demonstrates that the overall sentiment of a sentence may not always indicate the sentiment of the relationship between entities mentioned within it. Furthermore, some sentences involve multiple entity pairs, further complicating the task. Hence, compared to regular sentiment analysis, this task presents a greater challenge. To address this issue, we propose modeling this problem as Entity Relationship Sentiment Analysis (ERSA).

Definition 1 Entity Relationship Sentiment Analysis: Given a sentence s and two entities e_1 and e_2 , where $e_1 \neq e_2$, the goal is to determine the sentiment polarity of the relationship between e_1 and e_2 given s, where the relationship can be classified as either positive, negative, or neutral.

In this study, we also propose CERM - a Context-aware Entity Relationship Prediction Model to address the ERSA problem more effectively. Inspired by promising results in previous studies [1, 2, 3, 9], we combine the abilities of static and contextualized word embeddings models to generate richer representations of the problem inputs. While BERT, as a contextualized word embeddings model, is used to represent the features of the input sentence (s), a static word embeddings model is used to represent the features of each entity $(e_1$ and e_2). The combination of these two models is believed to produce richer representations because static word embeddings are better at capturing general semantic relationships between words (e.g., antonyms and synonyms), while contextualized embeddings are better at capturing more subtle semantic relationships (e.g., negation and sarcasm) [1, 2, 3, 9]. Furthermore, the two models also produce different types of information, and their combination can result in a more comprehensive representation. Furthermore, considering the significant resources required for data labeling, we propose a semisupervised learning (SSL) strategy with the proposed model to leverage unlabeled data during the learning process.

The contributions of this paper include:

- We introduce a new problem called Entity Relationship Sentiment Analysis (ERSA), which enables concept pair relationship inference from a wealth of food science-related research articles through sentiment analysis.
- We propose Context-aware Entity Relationship Prediction Model (CERM) for the ERSA task. Our experiments demonstrate that CERM, our proposed model, outperforms the state-of-the-art semi-supervised text classification methods with application to the ERSA task. We have also showcased the effectiveness of CERM by evaluating its performance on the Aspect-based Sentiment Analysis (ABSA) task.
- We introduce our dataset to facilitate further research in this literature-based sentiment analysis domain.

2 Related Works

2.1 Fine-grained Sentiment Analysis Tasks

Fine-grained sentiment analysis is a task in natural language processing that goes beyond the basic classification of sentiment and seeks to identify more nuanced sentiment levels in text. Entity-sentiment analysis (ESA) [19], aspect-based sentiment analysis (ABSA) [24], and multi-entity sentiment analysis (ME-ABSA) [32] are some examples of fine-grained sentiment analysis tasks. ESA involves identifying the sentiment associated with specific entities mentioned in a text, such as people, organizations, or products. ESA is more complex than regular sentiment analysis due to the presence of multiple entities in a text, and the sentiment towards one entity may contradict the sentiment towards another. For instance, based on a person's tweet, the sentiment towards a specific player may be positive, but the sentiment towards the management may be negative [19].

On the other hand, the ABSA task focuses on assigning multiple sentiments to a given text based on a list of aspects. For example, in [24], researchers define an ABSA task to detect multiple aspects (price, safety, transit, and general) in a comment toward a tourist destination and predict sentiment values for each aspect. As an ABSA task requires a more fine-grained understanding of the language used in the text, it is also considered a more challenging task than a regular sentiment analysis task. In [32], researchers extend the ABSA task to define the ME-ABSA task, where the focus is on identifying the sentiment with specific aspects of multiple entities mentioned in a text. As the ME-ABSA task deals with multiple entities and aspects, it is considered more challenging than ESA and ABSA tasks.

Similar to the previously mentioned tasks, the ERSA task can be classified as a fine-grained sentiment analysis task. While ERSA focuses on predicting the sentiment of the relationship between two named entities, whether explicitly or implicitly expressed in a sentence, ME-ABSA deals with the sentiment of multiple entities and a predefined set of aspects. The critical distinction is that ERSA does not have a predefined set of relationships between entities, making it challenging to map or encode as a ME-ABSA task. In ERSA, the primary objective is to predict the sentiment of the entity relationship itself rather than the sentiment of an entity with respect to a specific aspect, which is the primary goal of ME-ABSA. Therefore, ERSA cannot be directly mapped or encoded as an ME-ABSA task.

2.2 Semi-supervised Learning on Text Classification

Semi-supervised learning is motivated by the time-consuming and expensive process needed for data labeling, especially for data that requires domain-specific knowledge. Furthermore, a much larger pool of unlabeled data is often available than labeled data. There have been several applications of SSL in text classification. In [14], researchers proposed a mechanism for training a machine learning model that utilizes both labeled and unlabeled data. The method involves initially training the model on the labeled data and then using it to generate pseudo-labels for the unlabeled data. The key insight behind this approach is that the model's predictions on the unlabeled data will likely be accurate for at least some examples.

In [31], researchers defined an SSL algorithm by incorporating data augmentation techniques on the unlabeled data. By constraining the model's predictions for the original and augmented versions of the unlabeled data, the model can learn to be consistent. For example, in a traditional sentiment analysis task, if the model predicts a sentiment value of a sentence like "*Scientific research shows that ginger could help ease a sore throat*" to be 'positive,' the model can further enhance its learning by applying a consistency loss that encourages consistent predictions for its augmented versions. The proposed method was also shown to perform well against other SSL algorithms for various text classification tasks. The researchers in [27] further extended the idea by combining consistency regularization and pseudo-labeling techniques. Similarly to [31], FixMatch uses

consistency learning to encourage the model to make similar predictions for slightly perturbed versions of the same input. Instead of using all predictions on the unlabeled data, the FixMatch method only considers pseudo-labels with high confidence during the model training.

In our proposed method, we incorporate consistency regularization on the unlabeled data, following the promising results of the previous two studies. Instead of using BackTranslation (BT) to generate noisy examples, we utilize Easy Data Augmentation (EDA) in our proposed method as it better preserves the original text's meaning and requires lower computational resources. Moreover, considering the input from the problem domain, we also employ cosine embedding loss to generate a better model to learn the relevance of keywords and their corresponding text.

3 ERSA Dataset

The ERSA dataset is a dataset for entity relationship sentiment analysis. This dataset is extracted from the publication text of papers in the PubMed dataset ¹.

From the abstracts and full paper text of Pubmed publications. we extract sentences with one or more mentions of predefined entities. These predefined entities which are obtained from multiple sources²³⁴⁵⁶ can be placed in 6 groups: Genes, Disease, Chemical compounds, Nutrition, and Food ingredients. We then generate entity pairs from entities that appear in the same sentence. Thereby obtaining a dataset that focuses on the relationship between two entities, e_1 and e_2 , in a given sentence, s. Due to the high cost of data labeling, we only select 50,000 entity pairs with their corresponding sentences for labeling. We randomly selected entity pairs with corresponding sentences for labeling to achieve a representative distribution of entity pairs in the dataset. We then invite external data curators to assign the sentiment of selected entity pairs given their corresponding sentences. This labeling is done using the Amazon Mechanical Turk System⁷. Each entity pair and respective sentences are labeled by 3 different people as either positive, negative, or neutral. When there are disagreements, we use the majority label, and if there is no majority consensus, the data point is removed. After cleaning and post-processing the labeled data, we obtained 11, 366 labeled entity pairs. The final labeled dataset has 2,890 positive, 3,191 negative, and 3,011 neutral entity pairs given their corresponding sentences. See Table 1 for data statistics.

Given the labeled data, we randomly divided it into a 70/30 traintest split, where 70% of the data was used for training, and the remaining 30% was reserved for testing. Additionally, we augmented the training set with all available unlabeled data. Table 1 presents descriptive statistics for the dataset. Based on the statistics, it can be seen that the data has a relatively balanced label distribution. Also, it can be observed from the data that each data point contains a unique entity that does not repeat in other data points. This condition can lead to several issues (e.g., limited context and inconsistent quality) in the training process of the static word embeddings model if the learning process is only conducted on labeled data. Given that for the labeled data, each entity is only present in a single sentence;

 Table 1.
 The statistics of ERSA dataset

Property names	Training Set	Testing Set
Labelled data count	9092	2274
- Positive	2890	696
- Negative	3191	785
- Neutral	3011	793
unlabeled data count	50000	-
Word count		
- Average	35.59	35.20
- Min	3	3
- Max	1008	603
# of unique entities	18184	4548
- Chemicals	7823	1971
- Consumables	2450	657
- Diseases	7467	1820
- Nutrients	115	28
- Gene	329	72

there may be insufficient contextual information available to extract relevant features (a limited context issue). Additionally, the lack of diversity in the quality and relevance of the sentences for each entity may significantly impact the model's performance (inconsistent quality). Hence the need for a semi-supervised approach in learning.

4 The Proposed Method

4.1 The Architecture Overview

The proposed method operates based on a combination of static and contextualized word embeddings. Figure 1 illustrates the model architecture of the proposed method.

Given an input of entity pair $\{e_1, e_2\}$ and corresponding sentence, s, CERM extracts features for e_1 and e_2 using static word embeddings layer, T, and for s using contextualized word embeddings layer, B. The entity combination layer D_1 merge the individual embedding $z(e_1)$ and $z(e_2)$ of the entities in the entity pair to generate a rich representation of the entity pair. Similarly, the sentence embedding layer D_2 takes input from B to generate the sentence context embedding. The layers D_1 and D_2 are modeled as MLP layers and are intentionally designed to have the same dimensions to facilitate the learning process on unlabeled data proposed in this research. For a more stable model, we apply the ReLU activation to the output of the linear layers.

In the architecture, T and B can take the form of any pre-trained static word embedding model and contextualized word embeddings, respectively. \oplus is a concatenation operation. For our experiments, we chose to use a FastText model [6] that was trained on both labeled and unlabeled training data for T and BioBERT, a pre-trained BERT model that was trained on a large biomedical text corpus [15], for B. We chose the FastText model for its ability to handle out-of-vocabulary (OOV) words through subword information and BioBERT due to its suitability for the numerous biomedical terms and language used in the dataset. Finally, the classification layer Cconcatenates the learned entity pair embeddings and the sentence embeddings from D_1 and D_2 respectively, and projects it into class probabilities using softmax.

4.2 Task Specific Fine-Tuning for Contextualized Embeddings

Taking inspiration from previous studies in [30], we apply a similar transformation to the input sentence s prior to feature extraction

¹ https://pubmed.ncbi.nlm.nih.gov/download/

² http://ctdbase.org/downloads/#alldiseases

³ https://www.nlm.nih.gov/mesh/meshhome.html

⁴ https://omim.org/

⁵ https://www.ncbi.nlm.nih.gov/geo/info/download.html

⁶ https://pubchem.ncbi.nlm.nih.gov/docs/downloads

⁷ https://www.mturk.com



Figure 1. The proposed CERM model. The left side shows the proposed model's architecture (M) with its underlying entities features component (EC) and sentence features component (SC) for ERSA task. The right side shows the semi-supervised learning approach

to enhance the overall model's learning performance. The transformation is carried out by introducing a new token to the input of the contextualized embedding layer B. This strategy aims to introduce an additional signal that emphasizes the importance of the selected part to the model [30]. Thereby enabling the model to concentrate on critical elements, thus enhancing its performance. The proposed method transforms the input sentence s by including a special token, '[CTX].' Specifically, we add this token before and after each keyword's appearance, e_1 and e_2 , in a sentence s.

Furthermore, fine-tuning of the B model is performed solely on the last few attention layers, as suggested by previous studies [11, 28]. This strategy aims to optimize resource usage and minimize the time required for fine-tuning while preventing overfitting. This is achieved by allowing the B model to preserve its pre-trained knowledge. We apply the gradient accumulation strategy [33] for the parameter update during training. The gradient accumulation strategy smooths out the learning process by reducing the impact of noisy updates and allows the model to generalize on GPUs with limited memory.

4.3 Semi-Supervised Sentiment Analysis Task

In this subsection, we describe how the model can perform semisupervised learning for the ERSA task. We denote the set of input as $X = \{X_l, X_u\}$, where X_l is the set of labeled data samples and X_u is the set of unlabeled samples. The respective set of labeled and unlabeled samples are composed of individual data point x, which represents a tuple (e_1, e_2, s) , where e_1 and e_2 are two entities mentioned in a sentence s. The labeled data X_l has a corresponding set of labels Y such that a labeled data tuple is given as (x, y); $x \in X_l, y \in Y$. We define the classification layer as $f(., \theta)$, with output prediction of $p_{\theta}(y|x)$, which predicts y based on the input x, where θ represents the model parameters. The target is to infer the sentiments of unseen data by learning from labeled and unlabeled data.

4.3.1 Supervised Sentiment learning loss

Given the labeled data samples $x \in X_l$ and the corresponding labels $y^* \in Y$, the prediction probability scores $p_{\theta}(y|x)$ are obtained from

the output of the classification layer $f(., \theta)$. The cross-entropy (CE) loss is computed based on the predicted labels defined below.

$$\mathcal{L}_{CE}(X, Y, \theta) = -\frac{1}{N} \sum_{i=1}^{N} \sum_{j=1}^{L} y_{ij}^* \log p_{\theta}(y_{ij} = 1 | x_i)$$
(1)

where N is the number of labeled samples, L is the number of classes, y_{ij}^* is the ground truth label for the *i*-th sample and *j*-th class (either 0 or 1), and $p_{\theta}(y_{ij} = 1|x_i)$ is the predicted probability of the *i*-th sample belonging to the *j*-th class, given the input x_i and the model parameters θ . The outer sum computes the average over all samples in X. The model will learn a parameter θ that minimizes CE loss.

4.3.2 Unsupervised Consistency Learning

Inspired by a promising result of previous work in [31], our proposed model also uses a data augmentation technique to enable learning from a set of unlabeled data, X_u . We denoted $\epsilon(x)$ as a function representing all augmented versions of x. We minimize the divergence metric between the two distributions as defined in equation 2.

$$\mathcal{L}_{con}(X_u, \theta) = \frac{1}{|X_u|} \sum_{x \in X_u} \mathbb{E}_{\tilde{x} \sim \epsilon(x)} [KL(p_\theta(y|x), p_\theta(y|\tilde{x}))]$$
(2)

where \tilde{x} is a perturbed version of x sampled from the noise distribution ϵ . The Kullback-Leibler (KL) divergence measures the difference between the predicted probability distribution of $p_{\theta}(y|x)$ and the predicted probability distribution of the perturbed version $p_{\theta}(y|\tilde{x})$. We encourage the model to produce similar outputs over the inputs by minimizing the consistency loss. However, unlike previous research that used TF-IDF word replacement and Back-Translation, our model uses Easy Data Augmentation (EDA) [29] as function $\epsilon(x)$. Compared to TF-IDF word replacement, EDA generates more diverse augmented data while computationally less expensive than BT.

4.3.3 Relevancy learning via Similarity Learning

In this subsection, we use the notation $h_{(e_1,e_2)}$ to represent the output features from D_1 corresponding to keywords e_1 and e_2 , and h_s to represent the output features from D_2 corresponding to sentence s. In addition to learning from the unlabeled data, X_u , via consistency loss, the model utilized cosine embedding loss for similarity learning. During the training process, the model learns the relevance between $\{e_1, e_2\}$, and s by minimizing the cosine embedding loss between $h_{(e_1, e_2)}$ and h_s , while also maximizing the loss between $h_{(e_1, e_r)}$ and h_s , where e_r is a randomly selected entity and $e_r \neq e_1 \neq e_2$. Our similarity loss is defined as:

$$\mathcal{L}_{cos}(X_u, \theta) = \frac{1}{|X_u|} \sum_{(e_1, e_2, s) \in X_u} \mathbb{E}_{e_r \sim \Omega(e)}[\phi(e_1, e_2, e_r, s)] \quad (3)$$

where e_r is randomly selected from all keywords that appear in the dataset, $\Omega(e)$, and $\phi(e_1, e_2, e_r, s)$ is a cosine embedding loss defined in Equation 4

$$\phi(e_1, e_2, e_r, s) = 1 - \cos(h_{(e_1, e_2)}, h_s) + \max(0, \cos(h_{(e_1, e_r)}, h_s) - m).$$
(4)

In the Equation 4, cos is a cosine similarity function to measure the similarity or dissimilarity between two input vectors, and m is the margin that controls the degree of separation between the positive and negative pairs in the embedding space. We set the m value to zero in our proposed model. The objective of the ϕ is to ensure that the model encodes the relationship between subject entities and the sentence meaningfully rather than simply memorizing patterns from each feature.

4.3.4 Parameter Learning

To train the network, we jointly optimize $\mathcal{L} = \mathcal{L}_C + \mathcal{L}_{con} + \mathcal{L}_{cos}$, over the model parameters. Loss \mathcal{L}_C is the classification loss as described in section 4.3.1, \mathcal{L}_{con} is the consistency loss as described in section 4.3.2, and \mathcal{L}_{cos} is the similarity loss as described in section 4.3.3.

5 Experiment

To demonstrate the advantage of the proposed method for the ERSA task, we compared its classification performance against various supervised and semi-supervised learning methods. We evaluate the model performance using accuracy and f1 score on the testing set.

For the fully supervised methods, FastText or BERT was used to extract features from the text before feeding them into the classifier. To emphasize the encoding of the ERSA task, a special token '[CTX]' was inserted before and after each occurrence of the entities' keywords in the text before the feature extraction phase. We also evaluate the performance of the proposed model with and without fine-tuning the BERT for the task. The details regarding the feature extraction model used can be found in section 5.1, while the training settings of each classifier can be found in section 5.2.

In this paper, we evaluate the performance of the proposed model on four supervised methods: logistic regression (LR), MLP, LSTM, and Bi-LSTM. To compare the performance of our proposed method against other SSL methods, we employed pre-trained BERT as a classifier and utilized Unsupervised Data Augmentation (UDA) [31], Auxiliary Deep Generative Model (ADGM) [16], and FixMatch [27] as the benchmarks.

5.1 Feature Extractor Settings

Rather than utilizing an existing FastText model, we opted to train our model for the experiment, using a combination of labeled and unlabeled data from the training set and considering the specific language used for the task. There were 17, 955 sentences in total used in the training process. We had previously trained a FastText model on more sentences (57, 955), but we found that the model trained on a smaller dataset achieved better performance, so we decided to train the model using a smaller number. The model is a skip-gram model, trained with negative sampling (n = 5) and sub-word information ($n_gram = [1 - 6]$) to produce an embedding with a size of 200. The initial learning rate and context window are set to 0.1 and 30, respectively.

On the other hand, the BERT model used in the study is a pretrained model trained on a large biomedical text corpus, BioBERT, which is available in [15]. The utilization of BioBERT was conducted through two different settings. Firstly, by utilizing the model solely for feature extraction. And secondly, by jointly fine-tuning the BERT layer for this task. In the later subsection, we will refer to the BioBERT model that acts as a feature extractor as BERT and BERT-F for BioBERT with fine-tuning.

5.2 Learning Parameters for the Classifier

The training process for each classification method involved experimentation with several parameters that could impact the model's performance. To accomplish this, we conducted an exhaustive grid search to find the optimal parameter settings. The best model was selected based on its performance on the test data using the f1-score macro average. PyTorch was used to implement all models except for the logistic regression (LR) model. The scikit-learn implementation was used for the LR. For all models, the learning process utilized Adam optimizer with default parameters. For the baseline models, we implemented a hyper-parameter search and presented the results of the best configuration. For more information about the parameters for each model, please refer to section 5.2.1.

5.2.1 Experiment Reproducibility

Logistic Regression. All parameters were set to default except for the inverse of the regularization coefficient (c) and maximum iteration (max_iter). The max_iter was set to 1000 while c to 0.1. **MLP.** The MLP model was designed with one hidden layer, where the number of neurons in the hidden layer was 200. ReLU activation function was used in the model and was trained using a dropout layer with a dropout rate of 0.5. The model was trained with a batch size of 32 with a learning rate of 5e-05. The maximum number of epochs was set to 50.

LSTM and Bi-LSTM. Both models were designed to have one hidden layer, with the learning rate, batch size, and the number of neurons in the hidden layer set to 1e-05, 64, and 200, respectively. The Bi-LSTM model concatenated the last hidden layer states from both forward and backward phases for classification, whereas the regular LSTM used only the forward phase.

Auxillary Deep Generative Model. For the ADGM implementation, the experiment was only conducted using BERT as the feature extractor. The encoder layer consists of three layers, with [512, 256, 128] hidden neurons. The number of auxiliary variables was set to match the number of latent variables (the third layer). The decoder layer mirrored the encoder layer with respect to the number

	F1-Score (Negative)	F1-Score (Neutral)	F1-Score (Positive)	Macro F1	Accuracy
Logistic Regression	0.67	0.64	0.67	0.66	0.66
MLP	0.68	0.65	0.68	0.67	0.67
LSTM	0.66	0.63	0.64	0.64	0.64
Bi-LSTM	0.66	0.64	0.65	0.65	0.65
ADGM*	0.63	0.59	0.60	0.61	0.61
UDA*	0.68	0.64	0.65	0.66	0.66
FixMatch*	0.65	0.63	0.63	0.64	0.64
CERM*	0.69	0.67	0.68	0.68	0.68
UDA	0.72	0.67	0.71	0.70	0.70
FixMatch	0.72	0.68	0.71	0.70	0.70
CERM	0.72	0.69	0.71	0.71	0.71

 Table 2.
 Comparison of results for the ERSA task. * denotes the semi-supervised methods using the pretrained BERT layer only as a feature extraction layer (i.e., the BERT model is not fine-tuned during training.)

of hidden neurons. During the training process, the maximum epoch was set to 400, and the batch size was set to 256.

BERT Fine-tuning with UDA/ FixMatch. The fine-tuning process was carried out using the parameters obtained from the proposed model to compare the model's performance under the same configuration. Besides that, the lambda value for unsupervised learning was set to one to calculate the final loss, and the threshold value of 0.5 was used for the FixMatch implementation. The epoch was set to 2. The Proposed Model. In the proposed model, the batch size value was set to 64 with a learning rate and epoch of 2e-05. A new token, '[CTX],' was added to the tokenizer and used to encapsulate keywords in the text input for the BERT model. The number of neurons of D_1 and D_2 layers are set to 100. The data augmentation on the unlabeled data was performed using EDA by randomly removing, swapping, inserting, or replacing 20% of the input text.

5.3 Result and Analysis

Table 2 presents the results of all methods in the ERSA task. The MLP model achieved the highest performance in the fully supervised learning settings, followed by the LR model. In such settings, the LSTM and Bi-LSTM models produced lower performance compared to the other two models. This result indicates that the BERT model used as a feature extractor already captures a significant amount of information necessary for the task, so using a more complex classification method only leads to the overfitting of the complete model.

The experiments on semi-supervised learning methods using BERT as just a feature extractor (denoted by '*') also showed similar results. As the ADGM implementation used an extra encoderdecoder and another complex network to capture auxiliary information on top of BERT, the trained model had an inferior performance. Furthermore, training the model using FixMatch or UDA also yields worse model performance, suggesting the inappropriateness of both methods for such settings. While the other two SSL methods failed to improve the classification model performance, CERM showed promising results by achieving a slightly better macro F1 score compared to the best result achieved in fully supervised settings. This result demonstrates how the extra complexity in our proposed model can effectively learn from the unlabeled data. Unlike the ADGM implementation, which suffered decreased performance due to the addition of extra complexity to the model, CERM showed the opposite result. Intuitively, adding an extra network on top of two embeddings that learn different representations can improve performance. This addition allows the model to learn richer representations tailored for

the ERSA task through the combined embeddings.

Lastly, we also compared the performance of other semisupervised learning methods with fine-tuned BERT model with our proposed method. Based on the results, it can be observed that all three semi-supervised learning methods outperformed the supervised methods with fine-tuning. The proposed model achieved the most substantial performance improvement, followed by the UDA and FixMatch methods. Our model's better ability to classify the neutral class was the main reason for its substantial performance, as shown by further analysis of each class F1-score. In this paper, we use 10k unlabeled samples. Adding more unlabeled data for all three SSL strategies resulted in worse model performance. We postulate that for this dataset, due to the limited number of classes, a randomly sampled 10k unlabeled dataset is enough to generalize across the unlabeled data. Thus, adding more unlabeled data may hurt the model's performance, as the model may start to overfit.

5.4 Model Prediction on Special Cases

To demonstrate the model's performance on some non-trivial cases, we show models' predictions on sampled sentences in Table 3.

In Table 3, the CERM column represents the predictions made by our proposed model on the given text, while FM and UDA indicate the predictions made by fine-tuned BERT models using the Fix-Match and UDA SSL strategies, respectively. The first two examples describe multiple entity-to-entity relationships, the next two have a misalignment on the overall sentiment w.r.t the entity-to-entity relationship sentiment, and the last sample expresses mixed sentiments on the same entity-entity relationship. For the last example, a robust model is expected to give a neutral prediction due to the sentiment conflict. Based on these examples, the proposed model accurately predicts entity relationship sentiment in the first four cases. In contrast, models with FM and UDA struggle in the third example, where sentence sentiment doesn't align with entity relationships. UDA model also faces difficulty with multiple entities in the sentence (example 1). However, we also notice the model sometimes fails to predict the correct label for neutral sentiments on sentences expressing mixed sentiment, similar to the other two models (example 5). We postulate that for such cases, incorporating the entity domains/types would enhance the prediction by infusing into the model an additional understanding of the entities.

Table 3. Model Performance on Special Cases.

No.	Sentence	Actual	CERM	UDA	FM
1	while fenugreek raise	+	+	0	+
	testosterone, there is no				
	significant increase for				
	participants that used				
	tongkat ali and tribulus				
2.	consuming high-	-	-	-	-
	potassium foods like				
	honeydew and banana				
	excessively can result in				
	hyperkalemia				
3.	there is still a lack of re-	+	+	0	0
	search on this area but a				
	few suggest a good corre-				
	lation between garlic and				
	sleep quality				
4.	research on highly pro-	-	-	-	-
	cessed foods consumption				
	is good for increase soci-				
	ety awareness on how re-				
	fined sugar can lead to				
	obesity				
5.	Study in B showed a	0	+	+	-
	significant increase in				
	testosterone levels with				
	tongkat ali consumption,				
	while study in A reported				
	a contrasting finding,				
	suggesting a potential				
	negative effect				

5.5 Performance on Another Dataset: ABSA Task

To further validate the effectiveness of our method, we applied it to the task of Targeted Aspect-Based Sentiment Analysis for sentiment prediction, using the Sentihood dataset from a previous study [24]. The dataset consists of 5, 215 annotated sentences that contain one or two mentions of location entities. The names of location entities in the dataset are masked with "location1" and "location2" throughout, so the task does not involve identifying and segmenting named entities. There are 12 aspects included in the dataset, but only the four most frequent aspects were utilized in the experiment. Likewise, to compare our method's sentiment prediction performance with existing methods, we will only consider these four aspects.

However, we conducted two testing scenarios since the provided dataset contains no unlabeled data while our proposed method is trained using a semi-supervised technique. The first scenario only utilized the architecture in the proposed method and conducted fullysupervised learning. Meanwhile, half of the training data was treated as unlabeled in the second scenario, and the semi-supervised learning process was applied. In both scenarios, the best model parameters were determined using the provided validation set, and the final performance of the model was evaluated on the test set using the accuracy metric, consistent with the methodology employed in the previous study.

To encode sentiment prediction in the proposed method, we treat the entity, aspect, and sentence as e_1 , e_2 , and s. As the aspect is often not present in s, to match our problem setup, we concatenated the aspect e_2 at the end of s. We then add the '[CTX]' token before and after each occurrence of e_1 and e_2 in s. The static embedding, T, model used in this task is the pre-trained FastText model on Common Crawl and Wikipedia data⁸. Meanwhile, the contextualized embedding, *B*, model used is the pre-trained BERT on BookCorpus and Wikipedia ⁹. The number of neurons of D_1 and D_2 layers are set to 100. The best training parameters consisted of Adam optimizer with a batch size of 16, a learning rate of 5e-05, and an epoch of 8. To tackle the imbalanced class distribution problem in the dataset, we utilized the weighted loss method to update the weights. Table 4 compares our method and previous baseline methods.

Table 4. Comparison of results for the sentiment prediction in ABSA task.

Methods	Accuracy		
LR-Left-Right**	0.847		
LR-Mask(ngram)**	0.853		
LR-Mask(ngram+POS)**	0.875		
LSTM-Final**	0.820		
LSTM-Location**	0.819		
The proposed method			
- Semi-supervised	0.874		
- Fully-supervised	0.885		

**Retrieved from previous study in [24]

Based on Table 3, it can be observed that our fully-supervised trained model outperforms the best SOTA classification method in previous studies. Furthermore, in the semi-supervised learning setting, the model can achieve a relatively similar accuracy compared to the best classification method in the previous study by only utilizing half of the labeled data. These results demonstrate the benefits of combining pretrained static and contextualized embeddings in the proposed architecture and the advantages of using the proposed loss function, which allows the model to learn relevant information from unlabeled data.

6 Conclusion

In this paper, we propose a new sentiment analysis task called Entity Relationship Sentiment Analysis (ERSA) that focuses on capturing the sentiment between two entities. This task presents a more significant challenge than traditional sentiment analysis tasks, as the sentiment expressed in the sentence may not necessarily reflect the sentiment of the relationship between the entities. Furthermore, we propose an architecture combining two types of word embeddings to encode the ERSA task better. In addition to that, we also propose a method to enable the model to learn in a semi-supervised manner, considering the significant resources required for data labeling. The experimental results demonstrate that our proposed method consistently outperforms other semi-supervised learning methods across various learning scenarios.

In our future work, we aim to enhance the relevancy learning loss by incorporating Multi-class N-pair loss [26]. By doing so, we can emphasize the importance of the entity pairs in the sentence and encourage the model to increase the distance when incorrect words are chosen. By integrating Multi-class N-pair loss, we anticipate improving the model's ability to capture the relationships between entity pairs and ultimately enhance the overall relevancy learning process. In addition, we aim to explore various pre-trained transformerbased models to enhance the generalization of static and contextualized embedding combinations. It would be interesting to evaluate the performance of the proposed method by incorporating a distilled transformer-based model. This evaluation would provide insights into how the learning strategy employed in the proposed method operates on smaller models.

⁸ Available at: https://fasttext.cc/docs/en/crawl-vectors.html

⁹ Available at: https://huggingface.co/bert-base-cased

References

- Israa Alghanmi, Luis Espinosa Anke, and Steven Schockaert, 'Combining BERT with static word embeddings for categorizing social media', in *Proceedings of the Sixth Workshop on Noisy User-generated Text* (W-NUT 2020), pp. 28–33, Online, (November 2020). Association for Computational Linguistics.
- [2] Abdullah I. Alharbi and Mark Lee, 'Multi-task learning using a combination of contextualised and static word embeddings for Arabic sarcasm detection and sentiment analysis', in *Proceedings of the Sixth Arabic Natural Language Processing Workshop*, pp. 318–322, Kyiv, Ukraine (Virtual), (April 2021). Association for Computational Linguistics.
- [3] Abdullah I Alharbi, Phillip Smith, and Mark Lee, 'Enhancing contextualised language models with static character and word embeddings for emotional intensity and sentiment strength detection in arabic tweets', *Procedia Computer Science*, 189, 258–265, (2021).
- [4] Anonnya Banerjee and Nishargo Nigar, 'Nourishment recommendation framework for children using machine learning and matching algorithm', in 2019 International Conference on Computer Communication and Informatics (ICCCI), pp. 1–6, (2019).
- [5] Salliah Shafi Bhat and Gufran Ahmad Ansari, 'Predictions of diabetes and diet recommendation system for diabetic patients using machine learning techniques', in 2021 2nd International Conference for Emerging Technology (INCET), pp. 1–5, (2021).
- [6] Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov, 'Enriching word vectors with subword information', *Transactions of the Association for Computational Linguistics*, 5, 135–146, (2017).
- [7] Shweta Singh Chauhan, Deepak Kumar Sachan, and Ramakrishnan Parthasarathi, 'Focus-db: An online comprehensive database on food additive safety', *Journal of Chemical Information and Modeling*, 61(1), 202–210, (2020).
- [8] Betty M Drees and Brandon Barthel, 'We are what we eat', *Missouri Medicine*, **119**(5), 479, (2022).
- [9] Ashwin Geet d'Sa, Irina Illina, and Dominique Fohr, 'Bert and fasttext embeddings for automatic detection of toxic speech', in 2020 International Multi-Conference on: "Organization of Knowledge and Advanced Technologies" (OCTA), pp. 1–5. IEEE, (2020).
- [10] FoodCircle. Coronavirus Pandemic Triggers Surge in Demand for Healthy Foods. https://www.foodcircle.com/magazine/increasehealthy-foods-health-conscious-consumers.
- [11] Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin De Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly, 'Parameter-efficient transfer learning for nlp', in *International Conference on Machine Learning*, pp. 2790–2799. PMLR, (2019).
- [12] Celestine Iwendi, Suleman Khan, Joseph Henry Anajemba, Ali Kashif Bashir, and Fazal Noor, 'Realizing an efficient iomt-assisted patient diet recommendation system through machine learning model', *IEEE Access*, 8, 28462–28474, (2020).
- [13] Joo-Chang Kim and Kyungyong Chung, 'Knowledge-based hybrid decision model using neural network for nutrition management', *Information Technology and Management*, **21**(1), 29–39, (Mar 2020).
- [14] Dong-Hyun Lee, 'Pseudo-label: The simple and efficient semisupervised learning method for deep neural networks', in Workshop on challenges in representation learning, ICML 2013, (2013).
- [15] Jinhyuk Lee, Wonjin Yoon, Sungdong Kim, Donghyeon Kim, Sunkyu Kim, Chan Ho So, and Jaewoo Kang, 'Biobert: a pre-trained biomedical language representation model for biomedical text mining', *Bioinformatics*, **36**(4), 1234–1240, (2020).
- [16] Lars Maaløe, Casper Kaae Sønderby, Søren Kaae Sønderby, and Ole Winther, 'Auxiliary deep generative models', in *Proceedings of The* 33rd International Conference on Machine Learning, eds., Maria Florina Balcan and Kilian Q. Weinberger, volume 48 of Proceedings of Machine Learning Research, pp. 1445–1453, New York, New York, USA, (20–22 Jun 2016). PMLR.
- [17] Silvia Maggini, Adeline Pierre, and Philip C. Calder, 'Immune function and micronutrient requirements change over the life course', *Nutrients*, 10(10), (2018).
- [18] Kyle McKillop, James Harnly, Pamela Pehrsson, Naomi Fukagawa, and John Finley, 'Fooddata central, usda's updated approach to food composition data systems', *Current Developments in Nutrition*, 5, 596, (2021).
- [19] Margaret Mitchell, Jacqui Aguilar, Theresa Wilson, and Benjamin Van Durme, 'Open domain targeted sentiment', in *Proceedings of the*

2013 Conference on Empirical Methods in Natural Language Processing, pp. 1643–1654, Seattle, Washington, USA, (October 2013). Association for Computational Linguistics.

- [20] Nishargo Nigar and Mohammed nazim uddin, 'An intelligent children healthcare system in the context of internet of things', in 2nd International Conference on Innovative Research in Science, Technology & Management, (06 2018).
- [21] Maiyaporn Phanich, Phathrajarin Pholkul, and Suphakant Phimoltares, 'Food recommendation system using clustering analysis for diabetic patients', in 2010 International Conference on Information Science and Applications, pp. 1–8, (2010).
- [22] M. Poinski. Consumer trends shifting toward health and wellness, ADM finds. https://www.fooddive.com/news/consumer-trendsshifting-toward-health-and-wellness-adm-finds/584388/, sep 3 2020.
- [23] Joseph A Rothwell, Mireia Urpi-Sarda, Maria Boto-Ordoñez, Rafael Llorach, Andreu Farran-Codina, Dinesh Kumar Barupal, Vanessa Neveu, Claudine Manach, Cristina Andres-Lacueva, and Augustin Scalbert, 'Systematic analysis of the polyphenol metabolome using the phenol-explorer database', *Molecular nutrition & food research*, 60(1), 203–211, (2016).
- [24] Marzieh Saeidi, Guillaume Bouchard, Maria Liakata, and Sebastian Riedel, 'SentiHood: Targeted aspect based sentiment analysis dataset for urban neighbourhoods', in *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pp. 1546–1556, Osaka, Japan, (December 2016). The COLING 2016 Organizing Committee.
- [25] Augustin Scalbert, Cristina Andres-Lacueva, Masanori Arita, Paul Kroon, Claudine Manach, Mireia Urpi-Sarda, and David Wishart, 'Databases on food phytochemicals and their health-promoting effects', *Journal of agricultural and food chemistry*, **59**(9), 4331–4348, (2011).
- [26] Kihyuk Sohn, 'Improved deep metric learning with multi-class n-pair loss objective', in Advances in Neural Information Processing Systems, eds., D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett, volume 29. Curran Associates, Inc., (2016).
- [27] Kihyuk Sohn, David Berthelot, Chun-Liang Li, Zizhao Zhang, Nicholas Carlini, Ekin D. Cubuk, Alex Kurakin, Han Zhang, and Colin Raffel, 'Fixmatch: Simplifying semi-supervised learning with consistency and confidence', in *Proceedings of the 34th International Conference* on Neural Information Processing Systems, NIPS'20, Red Hook, NY, USA, (2020). Curran Associates Inc.
- [28] Danilo Vucetic, Mohammadreza Tayaranian, Maryam Ziaeefard, James J. Clark, Brett H. Meyer, and Warren J. Gross, 'Efficient finetuning of bert models on the edge', in 2022 IEEE International Symposium on Circuits and Systems (ISCAS), pp. 1838–1842, (2022).
- [29] Jason Wei and Kai Zou, 'EDA: Easy data augmentation techniques for boosting performance on text classification tasks', in *Proceedings of the* 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP), pp. 6382–6388, Hong Kong, China, (November 2019). Association for Computational Linguistics.
- [30] Taesun Whang, Dongyub Lee, Chanhee Lee, Kisu Yang, Dongsuk Oh, and Heuiseok Lim, 'An effective domain adaptive post-training method for bert in response selection', *Proceedings of the Annual Conference of the International Speech Communication Association, INTERSPEECH*, 2020-October, 1585–1589, (2020).
- [31] Qizhe Xie, Zihang Dai, Eduard Hovy, Thang Luong, and Quoc Le, 'Unsupervised data augmentation for consistency training', in *Advances in Neural Information Processing Systems*, eds., H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, volume 33, pp. 6256–6268. Curran Associates, Inc., (2020).
- [32] Jun Yang, Runqi Yang, Chongjun Wang, and Junyuan Xie, 'Multi-entity aspect-based sentiment analysis with context, entity and aspect memory', *Proceedings of the AAAI Conference on Artificial Intelligence*, **32**(1), (Apr. 2018).
- [33] Liu Zhuang, Lin Wayne, Shi Ya, and Zhao Jun, 'A robustly optimized BERT pre-training approach with post-training', in *Proceedings of the* 20th Chinese National Conference on Computational Linguistics, pp. 1218–1227, Huhhot, China, (August 2021). Chinese Information Processing Society of China.