



Adversarial Capsule Networks for Romanian Satire Detection and Sentiment Analysis

Sebastian-Vasile Echim¹, Răzvan-Alexandru Smădu¹, Andrei-Marius Avram¹,
Dumitru-Clementin Cercel^{1(✉)}, and Florin Pop^{1,2}

¹ Faculty of Automatic Control and Computers, University Politehnica of Bucharest, Bucharest, Romania

sebastian.echim@stud.aero.upb.ro,
{razvan.smadu, andrei.marius.avram}@stud.acs.upb.ro,
{dumitru.cercel, florin.pop}@upb.ro

² National Institute for Research and Development in Informatics - ICI Bucharest, Bucharest, Romania

Abstract. Satire detection and sentiment analysis are intensively explored natural language processing (NLP) tasks that study the identification of the satirical tone from texts and extracting sentiments in relationship with their targets. In languages with fewer research resources, an alternative is to produce artificial examples based on character-level adversarial processes to overcome dataset size limitations. Such samples are proven to act as a regularization method, thus improving the robustness of models. In this work, we improve the well-known NLP models (i.e., Convolutional Neural Networks, Long Short-Term Memory (LSTM), Bidirectional LSTM, Gated Recurrent Units (GRUs), and Bidirectional GRUs) with adversarial training and capsule networks. The fine-tuned models are used for satire detection and sentiment analysis tasks in the Romanian language. The proposed framework outperforms the existing methods for the two tasks, achieving up to 99.08% accuracy, thus confirming the improvements added by the capsule layers and the adversarial training in NLP approaches.

Keywords: Natural Language Processing · Satire Detection · Sentiment Analysis · Capsule Networks · Adversarial Training

1 Introduction

Satirical news is a type of entertainment that employ satire to criticize and ridicule, in a humorous way, the key figures from society, socio-political points, or notable events [27, 38]. Although it does not aim to misinform, it mimics the style of regular news. Therefore, it has a sizeable deceptive potential, driven by the current increase in social media consumption and the higher rates of distrust in official news streams [20].

Furthermore, sentiment analysis is regarded as a successful task in determining the opinions and feelings of people, especially in online shops where customer feedback analysis can lead to better customer service [37]. Limited resources in languages such as Romanian make it challenging to develop large-scale machine learning systems since the largest datasets present up to tens of thousands of examples [27]. Therefore, various

techniques should be proposed and investigated to address these challenges on such datasets.

Adversarial training is an effective defense strategy to increase the robustness and generalization of the models intrinsically. Introduced by Szegedy et al. [33] and analyzed by Goodfellow et al. [8], adversarial examples are augmented data points generated by applying a small perturbation to the input samples. It was initially employed in computer vision, where input images were altered with a small perturbation [8, 18, 36]. More recently, adversarial training gained popularity in NLP. The text input is a discrete signal; therefore, the perturbation is applied to the word embeddings in a continuous space [22]. The application of adversarial training in our experiments is motivated by the potential to improve the robustness and generalization of models with limited training resources.

This paper aims to introduce robust high-performing networks employing adversarial training and capsule layers [28] for satire detection in a Romanian corpus of news articles [27] and sentiment analysis for a Romanian dataset [34]. Our experiments include training models suitable for NLP tasks as follows: Convolutional Neural Networks (CNNs) [12], Gated Recurrent Units (GRUs) [3], Bidirectional GRUs (BiGRUs), CNN-BiGRU, Long Short-Term Memory (LSTM) [10], Bidirectional LSTM (BiLSTM), and CNN-BiLSTM. Starting from Zhao et al. [41], we compare the networks against their adversarial capsule flavors. Next, the best-performing network is subjected to an in-depth analysis concerning the impact on the performance of the capsule model and the training with adversarial examples. Thus, we test the effect of capsule hyperparameters varying the number of primary and condensed capsules [41]. Also, we assess the performance of our model employing Romanian GPT-2 (RoGPT-2) [24] for data augmentation up to 10,000 text continuation examples. Finally, we discuss several misclassified test inputs for the sentiment analysis task.

The main contributions in this work are as follows: (i) we thoroughly experiment with various configurations to assess the performances of the investigated approaches, namely adversarial augmentations and capsule layers; (ii) we show that the best-performing model uses BiGRU with capsule networks, while the most improvements were seen when incorporating RoGPT-2-based augmentations; (iii) we investigate the effects of analyzed components through t-SNE plots [17] and ablation studies; and (iv) we achieve state-of-the-art results on the two Romanian datasets.

2 Related Work

2.1 Capsule Networks in NLP

Firstly presented by Sabour et al. [28], the capsule neural networks are machine learning systems that model hierarchical relationships regarding object properties (such as pose, size, or texture) in an attempt to resemble the biological structure of neurons. Among other limitations, capsule networks are addressing the max pooling problem of the CNNs, which allows for translation invariance, making them vulnerable to adversarial attacks [15]. While it has been demonstrated that capsule networks are successful in image classification [28], there is also a general preference for exploring their potential in NLP tasks, especially in text classification. Several works [11, 42] took the lead in this topic, showing that using different approaches, such as static and dynamic routing, the capsule models provided competitive results on popular benchmarks.

Several studies were performed in topic classification and sentiment analysis using capsule networks. Srivastava et al. [30] addressed the identification of aggression and other activities, such as hate speech and trolling, using a model based on the dynamic routing algorithm [42] involving LSTM as a feature extractor, two capsule layers (namely, a primary capsule layer and a convolutional capsule layer), and finally, the focal loss [16] to handle the class imbalance. The resulting model outperformed several robust baseline algorithms in terms of accuracy; however, a more complex data preprocessing was expected to improve the results further.

For the sentiment analysis task, Zhang et al. [40] proposed CapsuleDAR, a capsule model successfully combined with the domain adaptation technique via correlation alignment [32] and semantic rules. The model architecture consisted of a base and a rule network. The base network employed a capsule network for sentiment prediction, consisting of several layers: embedding, convolutional, capsule, and classification. The rule network involved a rule capsule layer before the classification layer. Extensive experiments were conducted on review datasets from four product domains, which showed that the model achieved state-of-the-art results. Additionally, their ablation study showed that the accuracy decreased sharply when the capsule layers were removed.

Su et al. [31] tackled limitations of Bidirectional Encoder Representations from Transformers (BERT) [4] and XLNet [39], such as local context awareness constraints, by incorporating capsule networks. Their model considered an XLNet layer with 12 Transformer-XL blocks on top of which the capsule layer extracted space- and hierarchy-related features from the text sequence. Experiments illustrated that capsule layers provided improved results compared with XLNet, BERT, and other classical feature-based approaches.

Moreover, Saha et al. [29] introduced a speech act classifier for microblog text posts based on capsule layers on top of BERT. The model took advantage of the joint optimization features of the BERT embeddings and the capsule layers to learn cumulative features related to speech acts. The proposed model outperformed the baseline models and showed the ability to understand subtle differences among tweets.

2.2 Romanian NLP Tasks

In recent years, several datasets have emerged aiming to improve the performance of the learning algorithms on Romanian NLP tasks. Apart from the two datasets used in this work, researchers have also introduced the Romanian Named Entity Corpus (RONEC) [6] for named entity recognition¹, the Moldavian and Romanian Dialectal Corpus (MOROCO) [2] for dialect and topic classification, the Legal Named Entity Recognition corpus (LegalNERo) [26] for legal named entity recognition, and the Romanian Semantic Textual Similarity dataset (RoSTS)² for finding the semantic similarity between two sentences.

Lately, the language model space for Romanian was also improved with the introduction of Romanian BERT (BERT-ro) [5], RoGPT-2, ALR-BERT [23], and

¹ A new version of RONEC is available at <https://github.com/dumitrescustefan/ronec>.

² <https://github.com/dumitrescustefan/RO-STS>.

DistilMulti-BERT [1]. In addition, all the results for these systems have been centralized in the Romanian Language Leaderboard (LiRo) [7], a leaderboard similar to the General Language Understanding Evaluation (GLUE) benchmark [35] that tracks over ten Romanian NLP tasks.

3 Datasets

In this work, we rely on two of the most recent Romanian language text datasets: a corpus of news articles, henceforth called SaRoCo [27], and one composed of positive and negative reviews crawled from a Romanian website, henceforth called LaRoSeDa [34].

3.1 Satirical News

SaRoCo is one of the most comprehensive public corpora for satirical news detection, eclipsed only by an English corpus [38] with 185,029 news articles and a German one [20] with 329,862 news articles. SaRoCo includes 55,608 samples, of which 27,628 are satirical and 27,980 are non-satirical (or regular). Each sample consists of a title, a body, and a label. On average, an entire news article has 515.24 tokens for the body and 24.97 tokens for the title. The average number of sentences and words per sentence are 17 and 305, respectively. The labeling process is automated, as the news source only publishes satirical or regular content.

3.2 Product Reviews

LaRoSeDa is one of the largest corpora for sentiment analysis in the Romanian language. It was created based on the observation that the freely available Romanian language datasets were significantly reduced in size. This dataset totals 15,000 online store product reviews, either positive or negative, for which the ratings were also collected for labeling purposes. Thus, assuming that the ratings might reflect the polarity of the text, each review rated with one or two stars was considered negative. In contrast, the four or five-star labels were considered positive. The labeling process resulted in 7,500 positive reviews (235,474 words) and 7,500 negative reviews (304,813 words). The average number of sentences and words per review is 4 and 36, respectively.

4 Methodology

The generic adversarial capsule network we employ is presented in Fig. 1. It consists of a sub-module that can represent any widely-used NLP model, followed by capsule layers. Concretely, we use primary capsules and capsule flattening layers to facilitate the projection into condensed capsules passed as input for a routing mechanism to obtain the class probabilities. To increase robustness, we feed regular and adversarial samples into the model. In what follows, we detail the employed components.

Word Embeddings. Each word is associated with a fixed-length numerical vector, allowing us to express semantic and syntactic relations, such as context, synonymy, and antonymy. Depending on the model, the embedding representation has various sizes.




Fig. 1. Our generic adversarial capsule architecture, where E_d denotes the embedding size, N_s is the number of sentences, N_w is the number of words per sentence, N_{pc} is the number of primary capsules, N_{cc} is the number of condensed capsules, and N_{cls} is the number of classes to which the routing algorithm will converge.

To use a continuous representation of the input data, we employ two different types of embeddings: BERT- and non-BERT-based. On the RoBERT model [19], we rely on embeddings delivered by the model with a dimension $E_d = 768$, whereas, for the non-BERT models, we abide by Onose et al. [25] in terms of distributed word representations and choose Contemporary Romanian Language (CoRoLa) [21] with an embedding dimension $E_d = 300$, Nordic Language Processing Laboratory (NLPL) [14], having the size $E_d = 100$, and Common Crawl (CC) [9] with $E_d = 300$.

Adversarial Examples. To increase the robustness of our networks, we create adversarial examples by replacing characters in words. Using the letters of the Romanian alphabet, we randomly substitute one character per word, depending on the sentence size: one replacement for less than five words per sentence, two replacements for 5 to 20 words per sentence, and three replacements for more than 20 words per sentence.

Primary Capsule Layer. This layer transforms the feature maps obtained by passing the input through the sub-module into groups of neurons to represent each element in the current layer, enabling the ability to preserve more information. By using 1×1 filters, we determine the capsule \mathbf{p}_i from the projection p_{ij} of the feature maps [41]:

$$\mathbf{p}_i = \text{squash}(p_{i1} \oplus p_{i2} \oplus \dots \oplus p_{id}) \in \mathbb{R}^d \tag{1}$$

where d is the primary capsule dimension, \oplus is the concatenation operator, and $\text{squash}(\cdot)$ adds non-linearity in the model:

$$\text{squash}(\mathbf{x}) = \frac{\|\mathbf{x}\|^2}{1 + \|\mathbf{x}\|^2} \frac{\mathbf{x}}{\|\mathbf{x}\|} \tag{2}$$

Compression Layer. Because it requires extensive computational resources in the routing process (i.e., the fully connected part of the capsule framework), we need to reduce the number of primary capsules. We follow the approach proposed by Zhao et al. [41], which uses capsule compression to determine the input of the routing layer \mathbf{u}_j . Each condensed capsule $\hat{\mathbf{u}}_j$ represents a weighted sum over all the primary capsules:

$$\hat{\mathbf{u}}_j = \sum_i b_i \mathbf{p}_i \in \mathbb{R}^d \tag{3}$$

Routing Layer. It conveys the transition layer between the condensed capsules to the representation layer. It is denoted by a routing method to overcome the loss of information determined by a usual pooling method. In our capsule framework, we choose Dynamic Routing with three iterations [28].

Representation Layer. In the binary classification tasks, the last slice of our generic architecture is represented by the probability of a text input being satirical or regular for SaRoCo and positive or negative sentiment for LaRoSeDa.

5 Experimental Setup

5.1 Model Parameters

Firstly, we use CoRoLa, CC featuring 300-dimensional, and NLPL with 100-dimensional state space vectors for reconstruction at the embeddings level. We choose n -gram kernels with three sizes (i.e., 3, 4, and 5) and 300 filters each for the CNN sub-module. Also, for the Capsule layers, we use $N_{pc} = 8$ primary capsules and $N_{cc} = 128$ condensed capsules, which we fully connect through Dynamic Routing and obtain N_t lists with N_{cls} elements. For each element in the list, the argument of the maximum value represents the predicted label, where “1” is a satirical text or a positive review, whereas “0” is a non-satirical text or a negative review. Secondly, for the GRU and LSTM sub-modules, we employ one layer and a hidden state dimension of 300 for both unidirectional and bidirectional versions. Finally, for the RoBERT model, we choose the base version of the Transformer with vector dimensions of 768, followed by a fully connected layer with the size of 64, tanh activation function, and a fully connected layer with N_{cls} output neurons.

5.2 Training Parameters

The number of texts chosen from SaRoCo is $N_t = 30,000$ (15,000 satirical and 15,000 non-satirical) with a maximum $N_s = 5$ sentences per document and $N_w = 60$ words per sentence. For LaRoSeDa, we use 6,810 positive and 6,810 negative reviews for training, with $N_s = 3$ sentences per document and $N_w = 60$ words per sentence. The optimizer is Adam [13], and the loss function is binary cross-entropy. We set the learning rate to $5e - 5$ with linear decay and train for 20 epochs. The batch size is 32, and the train/validation/test split is 70%/20%/10%.

6 Results

This section presents the performance analysis of our models from quantitative and qualitative perspectives, as well as a comparison with previous works for the chosen datasets.

Initial Results. Table 1 shows our results on the SaRoCo and LaRoSeDa datasets. The experiments with varying embeddings other than RoBERT (i.e., CC, CoRoLa, and NLPL) show that NLPL determines better performance overall. This was unexpected

because CoRoLa covers over one billion Romanian tokens, while CC and NLPL contain considerably fewer tokens. For the SaRoCo dataset, the best model on the CC embeddings uses the BiGRU sub-module, achieving a 95.80% test accuracy. For the CoRoLa corpus, the GRU and BiGRU sub-modules perform equally, resulting in a 95.77% test accuracy. Also, the best NLPL embedding model considers the BiGRU sub-module, scoring a 96.15% test accuracy. On the LaRoSeDa dataset, we find the best model obtaining a 96.06% test accuracy based on GRU with NLPL embeddings. Moreover, training on the RoBERT embeddings brings the highest performance when combined with the BiGRU sub-module, achieving a test accuracy of 98.32% on SaRoCo and 98.60% on LaRoSeDa.

The score differences between our results on the two datasets are less than 0.5%. Therefore, a performance difference is expected due to the more considerable proportion of data for SaRoCo. Thus, there is no concrete insight into whether the satire detection task is more complex than the sentiment analysis one, especially in the binary classification setup. Still, since the training set size for LaRoSeDa is considerably smaller than that of the SaRoCo one, the slight performance difference shows polarization support on sentiment analysis.

We further assess the feature representation quality for each sub-module using the two-dimensional t-SNE visualizations upon the best-performing training results. Figure 2 shows different clustering representations in most cases. For the SaRoCo dataset, the best delimitation is observed on the BiGRU sub-module, which is validated by the best performance achieved for the NLPL embeddings as shown in Table 1. A similar effect applies to the BiGRU sub-module trained and evaluated on LaRoSeDa. Considering these results, the next set of experiments is performed based on the higher performance achieved with and without BERT embeddings, namely, the BiGRU sub-module with RoBERT and NLPL embeddings, respectively.




Fig. 2. t-SNE plots for each sub-module from the best-performing adversarial capsule network. The first row depicts the evaluation on SaRoCo, where blue indicates negative sentiment and orange represents positive one. The second row is for LaRoSeDa, where blue is for the non-satirical text, and orange is for the satirical one. The higher density on SaRoCo is because of a larger test dataset.

Table 1. Accuracy (Acc) of the generic adversarial capsule network with different word embeddings and sub-modules.

Embeddings	Sub-module	SaRoCo		LaRoSeDa	
		Valid. Acc(%)	Test Acc(%)	Valid. Acc(%)	Test Acc(%)
CC (300)	CNN	95.57	95.34	95.52	95.19
	GRU	95.92	95.70	95.29	95.33
	BiGRU	96.02	95.80	95.19	95.53
	CNN-BiGRU	95.90	95.60	95.16	94.39
	LSTM	95.70	95.54	95.09	94.53
	BiLSTM	95.67	95.47	95.19	94.46
	CNN-BiLSTM	95.57	95.00	95.09	95.06
CoRoLa (300)	CNN	95.49	95.60	95.19	95.26
	GRU	95.97	95.77	95.39	95.59
	BiGRU	95.99	95.77	95.46	95.60
	CNN-BiGRU	95.82	95.67	95.42	95.19
	LSTM	95.85	95.70	95.39	94.86
	BiLSTM	95.90	95.70	95.56	95.26
	CNN-BiLSTM	95.65	95.50	95.52	94.73
NLPL (100)	CNN	95.79	95.80	95.29	95.86
	GRU	96.04	95.80	95.92	96.06
	BiGRU	96.10	96.15	95.79	95.83
	CNN-BiGRU	95.60	95.80	95.32	95.19
	LSTM	95.74	95.64	95.52	95.79
	BiLSTM	95.44	95.70	95.29	94.99
	CNN-BiLSTM	95.45	95.57	95.22	95.39
RoBERT (768)	CNN	98.17	98.09	98.50	98.56
	GRU	98.07	98.17	98.39	98.49
	BiGRU	98.27	98.32	98.54	98.60
	CNN-BiGRU	98.07	98.24	98.42	98.45
	LSTM	98.07	98.24	98.36	98.39
	BiLSTM	98.10	98.17	98.46	98.49
	CNN-BiLSTM	98.04	98.24	98.46	98.52
BERT-ro [27]		82.41	73.00	-	-
Char-CNN [27]		73.42	69.66	-	-
HISK+BOWE-BERT+SOMs [34]		-	-	-	90.90

Comparison to Existing Methods. The results of Rogoz et al. [27] on the SaRoCo dataset show a more than 25% gain for our models compared to the BERT-ro approach, while our models outperform the character-level CNN by more than 29%. Human performance is a notable figure in deciding whether a selection of 200 news articles extracted from the dataset is satirical. Rogoz et al. [27] explored the idea, involving

Table 2. Accuracy for various capsule hyperparameters.

Dataset	N_{pc}			N_{cc}		
	2	8	32	32	128	256
SaRoCo	96.07	96.13	96.17	95.95	96.02	96.00
LaRoSeDa	95.23	95.52	95.50	95.01	95.46	95.12

ten human annotators and indicated that the human performance is at 87.35% accuracy. Our approach surpasses this result by more than 11%. In addition, the results shown by Tache et al. [34] on the LaRoSeDa dataset prove the competitive performance of our proposed approach. Thus, our results are 7–8% higher than their best model, HISK+BOWE-BERT+SOMs, which comprises histogram intersection string kernels, bag-of-words with BERT embeddings, and self-organizing maps.

Capsule Hyperparameter Variation. Fig. 1 depicts the hyperparameters of the capsule layers of our generic network, represented by N_{pc} (i.e., the number of primary capsules) and N_{cc} (i.e., the number of condensed capsules). We test the impact of these hyperparameters on the BiGRU sub-module with NLPL embeddings. We present the average for three runs per experiment. The chosen values for the hyperparameters are $N_{pc} = \{2, 8, 32\}$ and $N_{cc} = \{32, 128, 256\}$ (see Table 2).

During experiments, we observed that large values for N_{pc} considerably impact the training time. This is mainly due to the operations over high-dimensional matrices in the *squash*(\cdot) function from the iterative Dynamic Routing algorithm (see Eq. 2). Results from Table 2 support the intuition that a larger N_{pc} would bring better results. The model trained on SaRoCo with $N_{pc} = 32$ achieves the highest accuracy of 96.17%; nevertheless, the difference between choosing 8 and 32 is minimal. For SaRoCo and LaRoSeDa, the best overall performance is achieved in a setting with $N_{cc} = 128$, attaining accuracy scores of 96.02% and 95.46%, respectively. Based on both sets of results, we note that, for better performance, a hyperparameter search should be extended to the capsule hyperparameters.

Ablation Study. Motivated by the noteworthy closeness in performance between the BiGRU-based models with NLPL and RoBERT embeddings, respectively, we perform an ablation study, slicing the generic model into four categories: baselines (i.e., NLPL-BiGRU and RoBERT-BiGRU), adversarial (Adv), Capsule, and Adv+Capsule. The best results on the test datasets are brought by the most complex models in terms of training and architecture, with a 96.02% test accuracy for SaRoCo and a 95.82% test accuracy for LaRoSeDa using the NLPL embeddings, as well as a 98.30% test accuracy for SaRoCo and a 98.61% test accuracy for LaRoSeDa using the RoBERT embeddings (see Table 3).

Regarding model complexity, we determine that except for the adversarial training on a baseline BiGRU model, the performance improves when capsule layers are added on top of it, irrespective of including the perturbed data in training. The increase in performance on the SaRoCo dataset with our model is by 0.45% for the NLPL embeddings and by 0.10% for the RoBERT embeddings. We observe a decrease of 2.73% when the most undersized model (i.e., NLPL-BiGRU) is compared with the most complex one (i.e., RoBERT-BiGRU+Adv+Capsule). For the LaRoSeDa dataset, we gain 1.18%




Fig. 3. t-SNE plots on embedding space for each model from the ablation study.

Table 3. Ablation study.

Model	SaRoCo		LaRoSeDa	
	Valid. Acc(%)	Test Acc(%)	Valid. Acc(%)	Test Acc(%)
NLPL-BiGRU	94.80	95.57	92.73	94.64
+Adv	95.17	95.50	93.17	95.30
+Capsule	95.57	95.67	93.61	95.67
+Adv+Capsule	95.90	96.02	95.61	95.82
RoBERT-BiGRU	98.23	98.20	98.68	98.16
+Adv	98.47	98.00	98.83	97.94
+Capsule	98.33	98.27	98.68	98.46
+Adv+Capsule	98.45	98.30	98.75	98.61

using the NLPL embeddings and 0.45% with the RoBERT embeddings, respectively. Also, the test accuracy difference between the most complex and the most undersized models is 3.97%, determining that the network conveys more value for the sentiment analysis task.

The two-dimensional t-SNE embeddings depicted in Fig. 3 show the contrast between the capsule- and non-capsule-based models. The embeddings obtained with the BiGRU alone feature a specific chained distribution, with clusters defined by halving the sequence. The RoBERT embeddings convey a similar partition. In contrast, the capsule networks will mostly feature well-separated embedding clusters. No significant embedding change occurs when adversarial training is included.

Table 4. Results for RoBERT-BiGRU augmented with RoGPT-2 data in terms of precision (P), recall (R), and accuracy (Acc).

Dataset	Decoder Method	No. of Aug.	P(%)	R(%)	Acc(%)
SaRoCo	Greedy	1,000	98.15	98.18	98.16
		2,500	98.21	98.09	98.15
		5,000	98.36	98.20	98.31
		10,000	99.06	99.08	99.08
	Beam-search-2	1,000	98.24	98.08	98.23
		2,500	98.37	98.29	98.34
		5,000	98.19	98.08	98.17
		10,000	98.58	98.65	98.68
LaRoSeDa	Greedy	1,000	98.39	98.31	98.36
		2,500	98.82	98.52	98.70
		5,000	98.85	98.77	98.87
		10,000	98.94	98.87	98.94
	Beam-search-2	1,000	98.44	98.40	98.43
		2,500	98.72	98.49	98.64
		5,000	98.90	98.80	98.87
		10,000	98.82	98.70	98.77

Data Augmentation. Next, we incorporate the RoGPT-2 text continuation examples on a set of samples using two strategies for the decoder (i.e., greedy and beam-search-2). We perform experiments with the RoBERT-BiGRU model and show that the generative effort increases the overall performance for both tasks (see Table 4). In most cases, the RoBERT embeddings bring increased performance on the LaRoSeDa dataset as a consequence of the polarized effect of the product reviews, being strongly positive or negative. This polarization impact also applies to the models trained on augmented data. Data augmentation using the greedy decoder method achieves the best performance on SaRoCo, with a 99.08% test accuracy, employing 10,000 expanded texts, compared with the best accuracy of 98.68% obtained with beam-search-2. Furthermore, on LaRoSeDa, we determine similar performance on the greedy search algorithm with the best accuracy of 98.94% for 10,000 augmented texts. However, for the second dataset, more generated data will not necessarily determine the best performance as in the beam-search-2 scenario, using 10,000 augmented texts slightly underperforms in contrast with 5,000 examples.

Discussions. RoBERT-BiGRU, augmented with RoGPT-2 samples, correctly classifies 1,344 out of 1,362 examples from the LaRoSeDa test dataset. Due to spatial constraints, Table 5 depicts only the shortest eight misclassified texts out of 18, for which ground truth, predicted, and human annotated labels are shown. Two human annotators concluded from these examples that three indecisions and five classifications contradict the expected ones. The uncertain results and the negative misclassifications are expected to have been 3-out-of-5 stars ratings, which were assumed negative when the dataset was

Table 5. Examples from LaRoSeDa predicted with RoBERT-BiGRU. Ground truth (GT), Predicted (Pred) and Human labels are shown. P stands for Positive, N for Negative, and I for Indecisive.

Romanian text	English translation	GT	Pred	Human
o boxa ok din punct de vedere calitate pret daca este cumparata de unde trebuie. aici nu apare nici numele complet al boxei iar descrierea este saraca, plus pretul cu mult peste cat o gasesti in alte magazine	a good speaker in terms of quality and price if it is bought from the right place. the speaker's full name does not appear here and the description is poor, plus the price is much higher than what you can find in other stores	P	N	I
bun doar pentru incarcare (nu face conexiune, nu incarca rapid modelul nexus x). nu pare sa fie universal. nu realizeaza conexiune. voi mai incerca cu diverse cabluri micro usb si revin daca reusesc sa conectez telefonul la calculator	good only for charging (doesn't connect, doesn't fast charge the nexus x model). it doesn't seem to be universal. it doesn't connect. I will try with various micro usb cables and return if I can connect the phone to the computer	P	N	N
imi place. o bratară feminină care își face bine treaba. se sincronizează foarte bine cu android - samsung. bateria are autonomie zile cu funcția pulse ox activată, fără această funcție scrie că ar avea zile, dar nu am încercat	I like it. a feminine bracelet that does its job well. it synchronizes very well with android - samsung. the battery has an autonomy of days with the pulse ox function activated, without this function, it says it would have days, but I have not tried it	N	P	P
aproape multumit. am cumparat acest produs in urma cu o luna si pana acum doua zile am fost foarte multumit de el. bateria asigura o autonomie de - zile, finisajele sunt ok	almost satisfied. I bought this product a month ago and until two days ago I was very satisfied with it. the battery ensures the autonomy of - days, the finishes are ok	N	P	P
bun. folie calitativ buna dar nepotrivita pentru ecrane curbate. raman - milimetri dezlipiti pe margine. personal as recomanda folie de plastic pentru ecrane curbate dupa experienta asta	good. good quality foil but not suitable for curved screens. it remains - millimetres unglued on the edge. I would personally recommend a plastic film for curved screens after this experience	N	P	I
multumita! este foarte buna sunet clar! doar ca are probleme la conectarea cu bluetooth, il gaseste greu sau face nazuri km a conectare trebuie sa caut de multi ori bluetooth-ul. in rest e ok	pleased! it is a very good clear sound! it's just that it has problems connecting with bluetooth, it finds it hard or it's difficult to connect, I have to look for bluetooth many times. the rest is ok	N	P	P
recomand. claritate, sunet bun si un microfon super, fara fire, doar o cutiuta miniona de incarcare! pretul este mult sub cel de la apple. multumit de produs	I recommend it. clarity, good sound and a great microphone, no wires, just a tiny charging box! the price is much lower than that of apple. happy about the product	N	P	P
decent. il folosesc cu un samsung si nici pe departe nu are incarcare fast charge. daca nu te grabesti si ai rabdare sa astepti, merge	decent. I use it with a Samsung, which doesn't even have a fast charge. It will work if you are not in a hurry and have the patience to wait	N	P	I

created. Furthermore, we observe strongly positive texts such as “I like it. A feminine bracelet that does its job well”, “I was very satisfied with it”, “happy about the product”, “I recommend it”, and “pleased! it is a very good clear sound!” have negative ground

truth in the dataset. However, these are positive examples for the model and human annotators. Thus, we determine noise in the LaRoSeDa dataset, which is expected for datasets gathered from online sources, as the origin of the noise can be introduced by the page user or by automated data extractors.

7 Conclusions

Satire detection and sentiment analysis are important NLP tasks for which literature provides an ample palette of models and applications. Despite the more polarization expected on the product review task in contrast with the increased passivity of satirical texts, our models properly encapsulate the meaning represented by relevant features. In the syntactic and semantic context of our tasks, there is a slight difference in performance for the CC, CoRoLa, and NLPL embeddings, whereas fine-tuning the pre-trained RoBERT model brings up to 3% performance improvement. We showed in many experiments that our parameterized capsule framework can be adapted to specific problems. Moreover, we can improve the capsule network by employing data augmentation using generative models such as RoGPT-2, achieving a maximum gain of 0.6%. Based on our results, the potential of such an architecture is of increased significance, thus enabling further work in this direction.

Acknowledgements. The research has been funded by the University Politehnica of Bucharest through the PubArt program.

References

1. Avram, A.M., et al.: Distilling the knowledge of Romanian BERTs using multiple teachers. In: Proceedings of the Thirteenth Language Resources and Evaluation Conference, pp. 374–384 (2022)
2. Butnaru, A., Ionescu, R.T.: MOROCO: the Moldavian and Romanian dialectal corpus. In: Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics, pp. 688–698 (2019)
3. Cho, K., van Merriënboer, B., Bahdanau, D., Bengio, Y.: On the properties of neural machine translation: encoder-decoder approaches, pp. 103–111 (2014). <https://doi.org/10.3115/v1/W14-4012>
4. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: BERT: Pre-training of deep bidirectional transformers for language understanding. arXiv preprint [arXiv:1810.04805](https://arxiv.org/abs/1810.04805) (2018)
5. Dumitrescu, S., Avram, A.M., Pyysalo, S.: The birth of Romanian BERT. In: Findings of the Association for Computational Linguistics: EMNLP 2020, pp. 4324–4328 (2020)
6. Dumitrescu, Ș.D., Avram, A.M.: Introducing RONEC-the Romanian named entity corpus. In: Proceedings of the 12th Language Resources and Evaluation Conference, pp. 4436–4443 (2020)
7. Dumitrescu, S.D., et al.: LiRo: Benchmark and leaderboard for Romanian language tasks. In: Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 1) (2021)
8. Goodfellow, I.J., Shlens, J., Szegedy, C.: Explaining and harnessing adversarial examples. arXiv preprint [arXiv:1412.6572](https://arxiv.org/abs/1412.6572) (2014)
9. Graves, A., Schmidhuber, J.: Framewise phoneme classification with bidirectional LSTM networks. In: Proceedings. 2005 IEEE International Joint Conference on Neural Networks, vol. 4, pp. 2047–2052 (2005). <https://doi.org/10.1109/IJCNN.2005.1556215>

10. Hochreiter, S., Schmidhuber, J.: Long short-term memory. *Neural Comput.* **9**(8), 1735–1780 (1997)
11. Kim, J., Jang, S., Park, E., Choi, S.: Text classification using capsules. *Neurocomputing* **376**, 214–221 (2020)
12. Kim, Y.: Convolutional neural networks for sentence classification. In: *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 1746–1751 (2014). <https://doi.org/10.3115/v1/d14-1181>
13. Kingma, D.P., Ba, J.: Adam: a method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014)
14. Kutuzov, A., Barnes, J., Velldal, E., Øvreliid, L., Oepen, S.: Large-scale contextualised language modelling for Norwegian. *arXiv preprint arXiv:2104.06546* (2021)
15. Kwabena Patrick, M., Felix Adekoya, A., Abra Mighty, A., Edward, B.Y.: Capsule networks - a survey. *J. King Saud Univ. Comput. Inf. Sci.* **34**(1), 1295–1310 (2022)
16. Lin, T.Y., Goyal, P., Girshick, R., He, K., Dollár, P.: Focal loss for dense object detection. In: *Proceedings of the IEEE International Conference on Computer Vision*, pp. 2980–2988 (2017)
17. Van der Maaten, L., Hinton, G.: Visualizing data using t-SNE. *J. Mach. Learn. Res.* **9**(11), 1–27 (2008)
18. Madry, A., Makelov, A., Schmidt, L., Tsipras, D., Vladu, A.: Towards deep learning models resistant to adversarial attacks. In: *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, 30 April – 3 May 2018, Conference Track Proceedings*. OpenReview.net (2018). <https://openreview.net/forum?id=rJzIBfZAb>
19. Masala, M., Ruseti, S., Dascalu, M.: Robert-a Romanian BERT model. In: *Proceedings of the 28th International Conference on Computational Linguistics*, pp. 6626–6637 (2020)
20. McHardy, R., Adel, H., Klinger, R.: Adversarial training for satire detection: controlling for confounding variables. In: *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pp. 660–665. Association for Computational Linguistics, Minneapolis, Minnesota (2019). <https://doi.org/10.18653/v1/N19-1069>
21. Mititelu, V.B., Tufiş, D., Irimia, E.: The reference corpus of the contemporary Romanian language (CoRoLa). In: *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)* (2018)
22. Miyato, T., Dai, A.M., Goodfellow, I.J.: Adversarial training methods for semi-supervised text classification. In: *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, 24–26 April 2017, Conference Track Proceedings*. OpenReview.net (2017)
23. Nicolae, D.C., Yadav, R.K., Tufiş, D.: A lite Romanian BERT: ALR-BERT. *Computers* **11**(4), 57 (2022)
24. Niculescu, M.A., Ruseti, S., Dascalu, M.: Rogpt2: Romanian gpt2 for text generation. In: *2021 IEEE 33rd International Conference on Tools with Artificial Intelligence (ICTAI)*, pp. 1154–1161. IEEE (2021)
25. Onose, C., Cercel, D.C., Trausan-Matu, S.: SC-UPB at the VarDial 2019 evaluation campaign: Moldavian vs. Romanian cross-dialect topic identification. In: *Proceedings of the Sixth Workshop on NLP for Similar Languages, Varieties and Dialects*, pp. 172–177. Association for Computational Linguistics, Ann Arbor, Michigan (2019). <https://doi.org/10.18653/v1/W19-1418>
26. Păiş, V., Mitrofan, M., Gasan, C.L., Coneschi, V., Ianov, A.: Named entity recognition in the Romanian legal domain. In: *Proceedings of the Natural Legal Language Processing Workshop 2021*, pp. 9–18 (2021)
27. Rogoz, A.C., Gaman, M., Ionescu, R.T.: SaRoCo: detecting satire in a novel Romanian corpus of news articles. In: *Proceedings of the 16th Conference of the European Chapter of the*

- Association for Computational Linguistics: Main Volume. Association for Computational Linguistics, Online (2021). <https://arxiv.org/pdf/2105.06456.pdf>
28. Sabour, S., Frosst, N., Hinton, G.E.: Dynamic routing between capsules. In: *Advances in Neural Information Processing Systems*, vol. 30 (2017)
 29. Saha, T., Ramesh Jayashree, S., Saha, S., Bhattacharyya, P.: Bert-caps: a transformer-based capsule network for tweet act classification. *IEEE Trans. Comput. Soc. Syst.* 7(5), 1168–1179 (2020). <https://doi.org/10.1109/TCSS.2020.3014128>
 30. Srivastava, S., Khurana, P., Tewari, V.: Identifying aggression and toxicity in comments using capsule network. In: *Proceedings of the First Workshop on Trolling, Aggression and Cyberbullying (TRAC-2018)*, pp. 98–105. Association for Computational Linguistics, Santa Fe, New Mexico, USA (2018). <https://aclanthology.org/W18-4412>
 31. Su, J., Yu, S., Luo, D.: Enhancing aspect-based sentiment analysis with capsule network. *IEEE Access* 8, 100551–100561 (2020). <https://doi.org/10.1109/ACCESS.2020.2997675>
 32. Sun, B., Feng, J., Saenko, K.: Correlation alignment for unsupervised domain adaptation. In: Csurka, G. (ed.) *Domain Adaptation in Computer Vision Applications. ACVPR*, pp. 153–171. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-58347-1_8
 33. Szegedy, C., et al.: Intriguing properties of neural networks. In: Bengio, Y., LeCun, Y. (eds.) *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, 14–16 April 2014, Conference Track Proceedings* (2014). <http://arxiv.org/abs/1312.6199>
 34. Tache, A., Gaman, M., Ionescu, R.T.: Clustering word embeddings with self-organizing maps. application on LaRoSeDa - a large Romanian sentiment data set. In: *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pp. 949–956. Association for Computational Linguistics, Online (2021). <https://www.aclweb.org/anthology/2021.eacl-main.81>
 35. Wang, A., Singh, A., Michael, J., Hill, F., Levy, O., Bowman, S.: Glue: a multi-task benchmark and analysis platform for natural language understanding. In: *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pp. 353–355 (2018)
 36. Xiao, C., Li, B., Zhu, J.Y., He, W., Liu, M., Song, D.: Generating adversarial examples with adversarial networks. In: *27th International Joint Conference on Artificial Intelligence, IJCAI 2018*, pp. 3905–3911. International Joint Conferences on Artificial Intelligence (2018)
 37. Xu, G., Meng, Y., Qiu, X., Yu, Z., Wu, X.: Sentiment analysis of comment texts based on BiLSTM. *IEEE Access* 7, 51522–51532 (2019). <https://doi.org/10.1109/ACCESS.2019.2909919>
 38. Yang, F., Mukherjee, A., Dragut, E.: Satirical news detection and analysis using attention mechanism and linguistic features. In: *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pp. 1979–1989. Association for Computational Linguistics, Copenhagen, Denmark (2017). <https://doi.org/10.18653/v1/D17-1211>
 39. Yang, Z., Dai, Z., Yang, Y., Carbonell, J., Salakhutdinov, R.R., Le, Q.V.: XLNET: Generalized autoregressive pretraining for language understanding. In: *Advances in Neural Information Processing Systems*, vol. 32 (2019)
 40. Zhang, B., Xu, X., Yang, M., Chen, X., Ye, Y.: Cross-domain sentiment classification by capsule network with semantic rules. *IEEE Access* 6, 58284–58294 (2018). <https://doi.org/10.1109/ACCESS.2018.2874623>
 41. Zhao, W., Peng, H., Eger, S., Cambria, E., Yang, M.: Towards scalable and reliable capsule networks for challenging NLP applications. In: *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 1549–1559. Association for Computational Linguistics, Florence, Italy (2019). <https://doi.org/10.18653/v1/P19-1150>
 42. Zhao, W., Ye, J., Yang, M., Lei, Z., Zhang, S., Zhao, Z.: Investigating capsule networks with dynamic routing for text classification. In: *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pp. 3110–3119 (2018)