IV. International Conference on Electrical, Computer and Energy Technologies (ICECET 2024)
25-27 July 2024, Sydney-Australia

THINT: a Deep Learning Framework for Terrorist Threats Identification in Social Media

Marco San Biagio
Digital Technologies Research and Innovation Lab
Engineering Ingegneria Informatica S.p.A.
Palermo, Italy
marco.sanbiagio@eng.it

Valerio Scarfone
Digital Technologies Research and Innovation Lab
Engineering Ingegneria Informatica S.p.A.
Rome, Italy
valerio.scarfone@eng.it

Ernesto La Mattina
Digital Technologies Research and Innovation Lab
Engineering Ingegneria Informatica S.p.A.
Palermo, Italy
ernesto.lamattina@eng.it

Vito Morreale
Digital Technologies Research and Innovation Lab
Engineering Ingegneria Informatica S.p.A.
Palermo, Italy
vito.morreale@eng.it

***Abstract*—This paper presents THINT (THreat INTelligence Framework), a novel framework leveraging deep learning for the Open-Source Intelligence (OSINT) analysis of social media to identify terrorist related threats. Criminal organizations and terrorist groups are continually exploiting social media for radicalization campaigns and various illegal activities. The THINT framework addresses this emerging threat landscape by harnessing advanced deep learning algorithms, enhancing the early detection capability of terrorism-related content. Through meticulous fine-tuning during the training phase, THINT significantly outperforms previous models, offering enhanced identification and classification of threats. Compared to traditional Machine Learning (ML) approaches, such as Support Vector Machines (SVM), our framework not only demonstrates superior accuracy but also reduces classification time. These improvements represent a meaningful innovation in the usage of OSINT investigations to fight against terrorism and organised crime.**

Keywords— threat intelligence, deep learning, countering terrorism, OSINT, social media

I. INTRODUCTION

In the contemporary digital epoch, social media platforms have transcended their initial purpose of fostering social connections, morphing into pivotal components of daily life. As of 2023, platforms such as Facebook, Twitter, Instagram, and TikTok collectively boast over 5.04 billion active users—a figure that represents more than half the global population [1]. This digital renaissance has democratized content creation, enabling users to share a multifaceted array of content, from the mundane to the monumental. The appeal of social media lies in its versatility; platforms serve as repositories for a broad spectrum of user-generated content, including images, videos, and textual posts. This democratization of content creation and dissemination facilitates the exchange of ideas, cultural expressions, and personal experiences, creating a vibrant tapestry of global interaction.

However, the very attributes that make social media platforms incredibly popular, also render them susceptible to misuse by individuals with malicious intent [2]. The anonymity and vast reach provided by these platforms have been exploited for coordinating and promoting activities that range from cyberbullying to acts of terrorism. Notable examples include the use of social media for disseminating propaganda by terrorist organizations [3], coordinating attacks, and broadcasting extremist ideologies. A chilling illustration of this was the live streaming of the Christchurch Mosque shootings in New Zealand in 2019, which underscored the platforms' potential for amplifying acts of violence [4].

Recent scholarly work has significantly advanced our understanding of the multifaceted ways in which social media can be exploited by terrorist organizations and other malicious entities. A pivotal study by [5] delves into the strategies employed by ISIS to leverage social media for spreading its extremist ideology, recruitment, and the incitement of violence. This research underscores the complex challenge of curtailing online radicalization while safeguarding the principles of free expression. Furthermore, [6] offers a comprehensive analysis framework for assessing the role of Internet in violent extremism and terrorism. By proposing six methodological advancements, the author aims to elucidate the intricate dynamics between online platforms and the dissemination of extremist content, urging for a nuanced understanding of the digital ecosystem's influence on terrorism. Additionally, Weimann's investigation [7] into terrorist migration to the dark web reveals a strategic shift of extremist communications towards more obscured sections of the internet. This transition, motivated by increasing surveillance and censorship on mainstream social media, poses new challenges for counter-terrorism efforts, indicating a need for adaptive strategies that can navigate the complexities of the internet's less visible realms.

979-8-3503-9591-4/24/$31.00 ©2024 IEEE

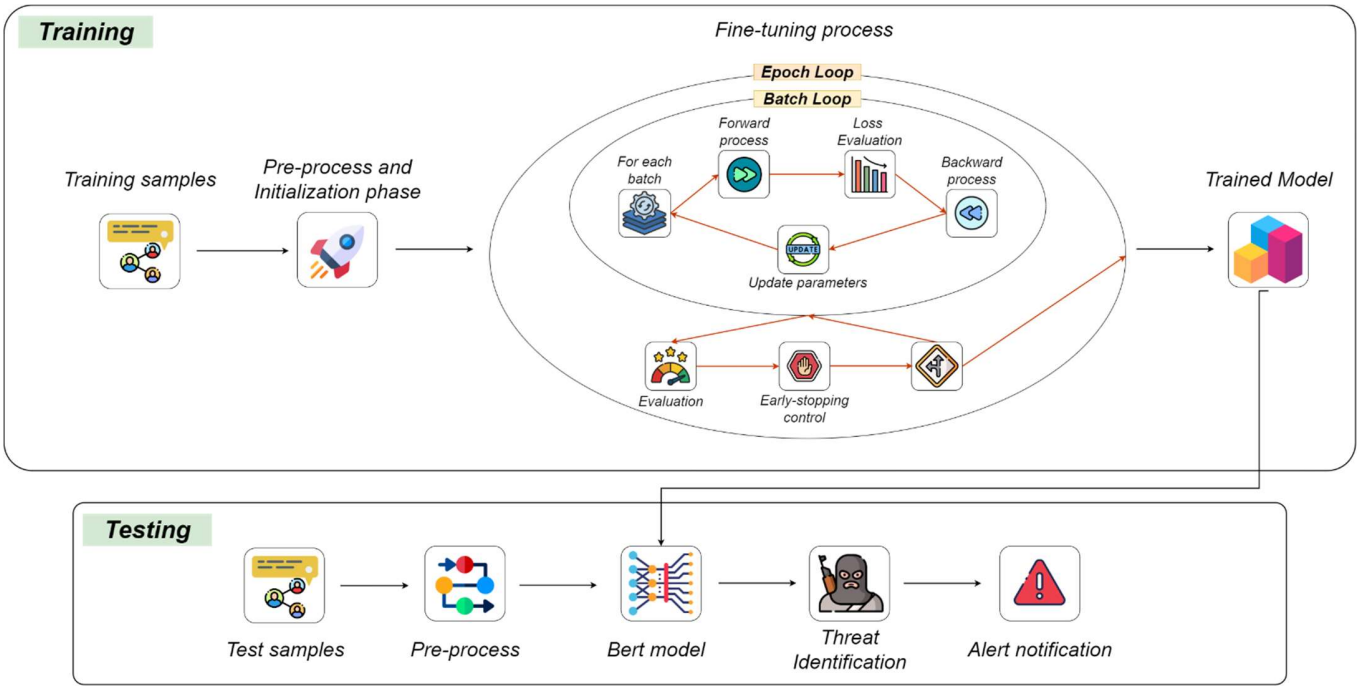

Fig. 1: Pipeline of the proposed method.

These studies collectively highlight the critical importance of ongoing, nuanced research into the misuse of digital platforms by terrorists. They emphasize not only the evolving nature of online radicalization and propaganda but also the imperative for balanced responses that respect human rights while effectively countering extremism.

The voluminous data generated on social media platforms necessitates the development of advanced analytical tools capable of sifting through the noise to identify potential threats. Traditional manual monitoring methods are no longer feasible due to the sheer scale of data, underscoring the need for automated systems that can efficiently process and analyse vast datasets to pinpoint security threats.

In response to the need for more sophisticated security measures, **THINT** (*THreat INTelligence Framework*) introduces a deep learning-based solution for identifying and classifying terrorist-related content on social media platforms. Unlike traditional machine learning methods that require intensive feature engineering and human supervision, THINT utilizes advanced deep neural networks. These networks excel at detecting complex patterns and subtle nuances within data, which enhances both the accuracy and speed of threat detection. Notably, there are few tools in this domain that are specifically designed to scrutinize social media for malicious content, making THINT a critical addition. This approach not only outperforms earlier technologies [8-9] in terms of efficacy but also significantly accelerates the identification process. By harnessing the capabilities of deep learning, THINT represents a significant advancement in the field of open-source intelligence (OSINT), providing law enforcement and intelligence agencies with a robust tool to combat the proliferation of terrorism and extremist ideologies through social media.

II. Literature Review

The use of social media platforms by terrorist organizations for propaganda, recruitment, and communication has been a significant area of research in recent years. Scholars have explored various dimensions of this phenomenon, including detection methodologies, the psychological effects of online radicalization, and the development of counter-narrative strategies. This review will focus on key studies that have contributed to understanding and addressing the challenge of terrorism in the digital domain.

One seminal work in the domain of detecting terrorist content on social media is by Ferrara et al. [10], who investigated the use of social bots by terrorist groups. Their analysis revealed that bots play a significant role in spreading extremist content and propaganda, illustrating the need for advanced detection mechanisms. In [11], the authors delve into the sophisticated use of information warfare by jihadist groups, examining how the Islamic State and similar organizations have utilized social media for recruitment and propaganda. They discuss the implications for Western counter-terrorism strategies, emphasizing the need for a comprehensive approach to counter the virtual caliphate. The study provided in [12] presents an analysis of Twitter networks to identify patterns of support and opposition to ISIS. Using advanced network analysis techniques, the authors offer insights into how ISIS exploits social media for its operations and how these networks can be disrupted. The authors of [13] apply machine learning techniques to predict the spread of extremist content on social media platforms. They focus on identifying users likely to adopt extremist content and the reciprocity of interactions among such users,

providing valuable information for pre-emptive actions against the dissemination of terrorist material. Finally, in [14], the authors explore how YouTube is used for disinformation and crowd manipulation tactics, including those employed by terrorist organizations. The study emphasizes the platform's role in the broader ecosystem of online radicalization and the challenges it poses for detecting and countering extremist content.

The psychological process of radicalization on social media platforms has been another focus area. [15] offers a comprehensive overview of how terrorists exploit social media for psychological warfare, detailing strategies used to radicalize and recruit individuals online. This work underscores the complexity of online radicalization and the multifaceted approach needed to counteract it.

Addressing the proliferation of extremist content, Braddock and Horgan [16] delve into the effectiveness of counter-narratives in combating online radicalization. They argue that well-crafted counter-messages can undermine terrorist propaganda, but they also highlight the challenges in creating impactful narratives.

A key study that addresses the balance between privacy and security in the context of digital surveillance for counter-terrorism purposes is by [17]. Such research critically examines the data collection practices of tech companies and governments, highlighting the potential for abuse and the impact on privacy rights. The paper advocates for a regulatory framework that ensures transparency and accountability in the use of social media data for security purposes. This work ensures operational efficiency of law enforcement while protecting citizens' privacy.

Finally, [18] provides a comprehensive review of literature over the past decade, exploring the dual role of social media in facilitating and combating terrorism, including its use in emergency communications and threat detection.

III. Proposed Method

As detailed in [9], the proposed framework effectively identifies and categorizes potential threats within textual content gathered via OSINT methodologies. The core of the framework lies on the BERT (Bidirectional Encoder Representations from Transformers) algorithm [19], which currently represents the state-of-the-art in the NLP (Natural Language Processing) technologies. BERT marks a significant innovation by using the transformer architecture, which facilitates a dynamic interplay among all elements of the input and output data. This architecture transcends the limitations of previous models that processed text in a linear sequence. The transformative aspect of BERT is its bidirectional processing capability. Unlike earlier approaches that analysed text in a single direction—either from left to right or right to left—BERT evaluates the context of each word by looking at the words that come before and after it, simultaneously. This method allows for a more profound comprehension of the text's context, leveraging the Transformer's ability to attend to all parts of the input data simultaneously. The bidirectional attention mechanism of BERT is adept at discerning subtle and complex relationships within the text, offering a departure from the constraints of unidirectional processing models.

Despite BERT's advanced capabilities in understanding and processing natural language, it is not without its limitations, particularly when addressing highly specialized tasks or niche domains. To harness BERT's full potential for these specific applications, it often necessitates an additional step of fine-tuning on a dataset that is narrowly tailored to the specific task or field in question. This process ensures that the model can adapt its generalized understanding of language to the unique contexts and terminologies of the target domain.

A study by [20] exemplifies this requirement, illustrating how BERT's performance on domain-specific tasks can significantly improve with task-specific fine-tuning. In their research, the authors demonstrate that fine-tuning BERT with data from a targeted domain not only enhances the model's accuracy but also its ability to grasp the subtleties and nuances inherent to specialized fields. This study underscores the importance of fine-tuning in bridging the gap between BERT's general language proficiency and its application to domain-specific challenges.

The proposed approach implements a BERT fine-tuner [21]. It uses PyTorch and the Transformers library on sequence classification tasks. Main parameters used in the fine tuning are: *AdamW* as optimizer, *StepLR* as scheduler and *CrossEntropyLoss* as loss function. *AdamW* and *StepLR* optimize learning dynamics and model convergence, while *CrossEntropyLoss* accurately assesses classification efficacy, together driving superior model performance.

Fig. 1 shows the pipeline of the proposed framework. The procedure is split into two parts:

- **Training:** The first step is the initialization of the essential parameters (i.e., model name, learning rate, batch size, etc.), and loads the pre-trained BERT model and tokenizer. The training method is performed over a specified number of epochs. Within each epoch, the training data is loaded in batches using a custom dataset and data loader. The model is set to training mode, gradients are reset, and forward pass is performed on the input data. Loss is calculated using *CrossEntropyLoss* between the predicted logits and actual labels. Backpropagation is applied to compute gradients, and optimizer updates the model parameters. The process repeats for all batches in the training data (as shown in the figure). After each epoch, the model's performance is evaluated on a separate validation dataset using similar steps as in training but without gradient updates, in which the average validation loss is calculated. The process monitors validation loss. If there's no improvement for a specified number of epochs (patience), training is stopped early to prevent overfitting. Learning rate scheduling (*StepLR*) is applied to adjust the learning rate based on the defined scheduler parameters. Finally, the trained model is saved for future inference. This process iterates until the specified number of epochs or until early stopping criteria are met, resulting in a fine-tuned BERT model capable of classifying sequences based on the provided labels.
- **Testing:** The testing phase, while more straightforward than the preceding steps, follows a critical process. It begins with the preprocessing of

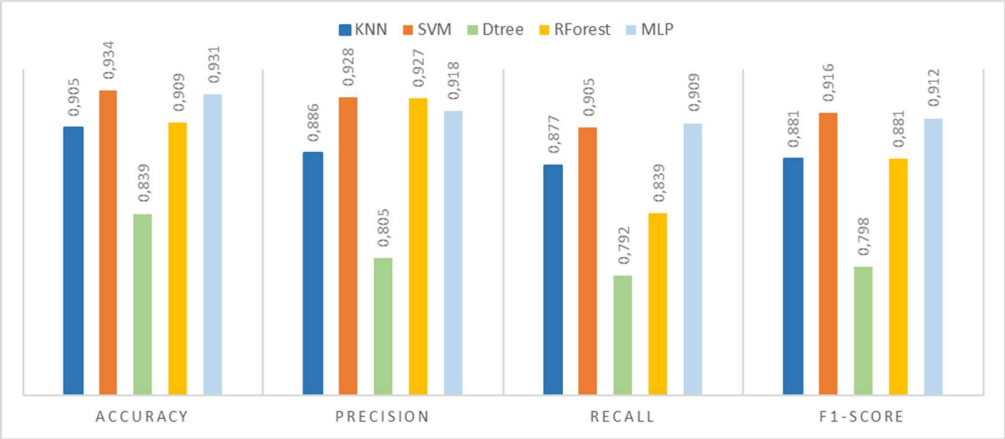

Fig. 2: Experiment one - Comparison between the different classification methods.

test samples, employing the same methods used during training, followed by the application of the fine-tuned BERT model for sample classification. Upon the identification of any threats within the samples, an alert is generated.

IV. Dataset & Pre-Proccesing

The experiments utilize a dataset originally introduced in [8], comprising tweets labelled as terrorism-related or not. Specifically, the terrorism-related tweets were sourced from a study on **ISIS's Twitter activity** [22], featuring over 13,000 tweets by pro-ISIS supporters globally, post the November 2015 Paris Attacks. These tweets, exclusively in English, were analysed for their text content, excluding metadata like usernames or follower counts. Conversely, the non-terrorism tweets were gathered from discussions around the **FIFA World Cup 2018** [23], amassing over 400,000 tweets via the Tweepy API. Enhancements to the original dataset include recent additions covering events from 2023-2024, such as the European Athletics Championship and UEFA Champions League Matches, and a selection from the **CyberTweets dataset** [24] focusing on general discussions about cyber network attacks, adding 21,000 tweets to the corpus. This latter addition includes 7,000 tweets on cyber network attack discussions. Initially totalling around 520,000 entries, continual updates have expanded the dataset to approximately 530,000 elements, ensuring a comprehensive and evolving database for analysis.

Within the architecture of the proposed framework, dedicated to identifying malicious content, he pre-processing stage plays a crucial role. This stage is essential for refining textual data, significantly enhancing the classification model's efficiency and accuracy in detecting and categorizing threats. It lays the groundwork for optimal model performance in threat detection.

At the heart of this phase is the transformation of text data into a standardized and unbiased format, essential for accurate classification. Techniques such as converting text to lowercase, removing URLs, hashtags, emoticons, special characters, anonymizing or completely removing names of people, tokenizing, and eliminating stop words are meticulously applied. This ensures the text is not only purified and uniform but also respects privacy and reduces biases. Such detailed data preparation is vital for the fine-tuning process, enabling the model to concentrate on relevant features and patterns without being swayed by irrelevant personal information or presentation styles. This rigorous optimization of data, including the anonymization of sensitive information, is indispensable for ensuring the model's high precision in threat identification and its robustness against the complexities of textual data.

V. Experiments

Two distinct experiments were conducted. The first experiment evaluated the efficacy of various machine learning classifiers to select the most effective one for comparison with our fine-tuned BERT model. Given the absence of existing literature on the analysis of the datasets presented, we adhered to the methodology outlined in [9], with the addition of a new classifier for comparative analysis. The classifiers considered included: Support Vector Machines (SVM) [25], *K-Nearest Neighbours (KNN)* [26], *Decision Trees* [27], *Random Forests* [28], and *Multilayer Perceptron (MLP)* [29].

Upon determining the most effective machine learning classifier in the initial experiment, we proceeded with a second experiment to directly compare it with our fine-tuned BERT model. This comparison was specifically used to evaluate the performance enhancements offered by the fine-tuned BERT over the previously identified top-performing classifier. The experimental results unequivocally demonstrated the superior capabilities of the fine-tuned BERT model, showcasing its enhanced efficiency, greater accuracy, and improved handling of complex linguistic data. This phase conclusively proved the advantages of leveraging fine-tuned BERT technology in surpassing the limitations of conventional machine learning classifiers in our application domain.

A. Experiment one

In this experiment, the performance metrics employed include *accuracy*, *precision*, *recall*, and *F1-score*. The following symbols are introduced to clarify the experimental setup:

- T and N denote the datasets for terrorism and non-terrorism, respectively.

- T_{train} and T_{test} represent the training and testing subsets of the terrorism dataset.
- N_{train} and N_{test} represent the training and testing subsets of the non-terrorism dataset.

Adhering to the methodology described in [9], we randomly allocated 50% of the dataset to the training set (T_{train}) and the remaining 50% to the test set (T_{test}) within the terrorism dataset (T). Similarly, for the non-terrorism dataset (N), we adopted a comparable approach to form training (N_{train}) and testing subsets (N_{test}), ensuring balanced data distribution with no overlap between the training and testing sets. To address the balance and comprehensiveness of the data, the dimension of the non-terrorism dataset was expanded to include 25,000 elements, thereby enhancing the robustness of our analysis. The classifiers were then trained on these augmented positive and negative training sets, with their performance subsequently evaluated on the test sets. This entire procedure was repeated 30 times to introduce variability and ensure the reliability of our findings.TABLE I details the size of each dataset utilized in this study.

TABLE I. DIMENSION OF THE DATASET FOR THE EXPERIMENT ONE.

T_{train}	T_{test}	N_{train}	N_{test}
6715	6714	12500	12500
13429		25000	

TABLE II. RESULTS OF EXPERIMENT ONE.

	Accuracy	*Precision*	*Recall*	*F1-score*
KNN	0.905	0.886	0.877	0.881
SVM	**0.934**	**0.928**	**0.909**	**0.916**
Dtree	0.839	0.805	0.792	0.798
RForest	0.909	0.927	0.839	0.881
MLP	0.931	0.918	0.905	0.912

TABLE II and Fig. 2 show the results. The standout performer is the Support Vector Machine (SVM), which leads in all metrics, notably achieving an accuracy of 93.4% and an F1-score of 91.6%, marking it as the most effective model in balancing relevance and completeness of predictions. The Decision Tree (Dtree) model trails behind, recording the lowest scores, which might indicate its simplicity or unsuitability for the dataset's complexity. In contrast, the Random Forest (RForest) model, an ensemble method, substantially improves upon the Decision Tree's limitations, yet its lower recall compared to precision suggests a conservative stance towards positive class predictions. The Multilayer Perceptron (MLP), a neural network approach, showcases strong performance with metrics closely trailing those of the SVM. It indicates the neural network's competitive edge in handling the task, with a slight edge in handling both precision and recall. Despite the close competition, SVM's slight lead in performance metrics, especially in precision and recall, sets it apart, underpinning its selection for the second experiment.

B. Experiment two

The second experiment focused on directly comparing the fine-tuned BERT model against the SVM classifier, which emerged as the best performer from the initial set of experiments. Starting from the splits of the experiment one, the dataset was partitioned into 70% for training, 15% for validation, and the remaining 15% for testing purposes. This structured approach allowed us to meticulously tune and validate both models under equitable conditions, ensuring a fair and accurate comparison of their performance capabilities.

The validation set played a crucial role in this experiment. For the fine-tuned BERT model, it facilitated the selection of the optimal set of hyperparameters, ensuring that the model was adequately adjusted to the specifics of our dataset, thereby maximizing its performance. Concurrently, the validation set served as a critical tool for fine-tuning the SVM classifier. By methodically adjusting its parameters based on validation set performance, we were able to refine the SVM model, ensuring that it operated at its highest possible efficiency within the context of our data.

The validation phase thus served a dual purpose: it not only provided a benchmark for selecting the best-performing model parameters but also acted as a preliminary gauge of each model's effectiveness before proceeding to the final evaluation stage. This meticulous approach to validation and parameter selection was instrumental in enhancing the comparability of the two models, setting the stage for a rigorous and insightful final assessment based on the testing set.

he hyperparameters selected for fine-tuning BERT are as follows:*batch_size=16,max_length=512,scheduler_gamma=.95,num_epochs=20,patience=5*.

Finally, we would like to underline that utilizing a GPU, we were able to complete the fine-tuning process for the BERT model in just 15 minutes, highlighting the impact of specialized hardware in expediting the training phase and facilitating rapid deployment.

TABLE III. RESULTS OF EXPERIMENT TWO.

	Accuracy	*Precision*	*Recall*	*F1-score*
BERT	0.973	0.968	0.97	0.969
SVM	0.937	0.951	0.94	0.94

TABLE III shows the comparative results between the two classification methods. The comparison between the fine-tuned BERT model and the SVM classifier underscores a significant advancement in performance metrics across the board for the BERT model. With an accuracy of 97%, BERT outstrips the SVM by a notable margin, reinforcing the effectiveness of deep learning techniques in complex classification tasks. This high accuracy indicates that the BERT model is exceptionally reliable in identifying the correct categories for the given inputs. Precision, a metric indicative of the quality of positive predictions, is nearly perfect for BERT at 96.8%, surpassing the already impressive 95.1% of the SVM. This suggests that when BERT predicts a post to be malicious, it is highly likely to be correct, a crucial feature for operational settings where false positives can have significant implications. Recall, or the ability to capture all relevant instances, stands at 97% for BERT, demonstrating its superior capability to identify and classify all potential threats accurately. This is critical in threat intelligence, where missing a malicious post could have dire consequences. The SVM, while still performing admirably, falls short in this regard with a 94% recall, indicating BERT's enhanced sensitivity to detecting nuanced or subtle indicators of malicious content. The F1-score, which balances precision and recall, is nearly identical for BERT at 96.9%, signifying a well-rounded model that does not sacrifice one measure for the improvement of another. In comparison, the SVM's F1-score of 94% is strong but highlights the incremental benefits brought

by the fine-tuning of BERT, particularly in its ability to generalize and accurately classify across a diverse set of samples.

Overall, these results clearly demonstrate the superiority of the fine-tuned BERT model over traditional machine learning methods like SVM for this specific task. The advancements in accuracy, precision, recall, and F1-score not only highlight the efficacy of deep learning models in processing and classifying complex textual data but also underscore the potential for these technologies to significantly enhance threat intelligence frameworks.

VI. Conclusions

In this work, THINT tool has been presented. THINT is a framework employing deep learning to detect malicious content within Open-Source Intelligence (OSINT) data. Central to THINT is a fine-tuned BERT model, marking a paradigm shift from conventional machine learning to cutting-edge deep learning in NLP for security applications, particularly in scrutinizing social media for threats. This shift resulted in a notable 3% to 4% increase in overall performance and a significant speed enhancement in processing compared to traditional classifiers. The superior performance of the fine-tuned BERT model hinges on its nuanced comprehension of language, crucial for distinguishing harmful content accurately. Moreover, its fast-processing capability is vital for real-time threat detection, addressing the rapid proliferation of online content. Looking ahead, enhancing THINT could involve data augmentation for greater model resilience, experimenting with various optimization techniques for better convergence, and exploring alternative transformer models for broader NLP tasks in threat detection. Additionally, more comprehensive experiments will be conducted to explore various state-of-the-art techniques. These advancements not only demonstrate technological strides against digital terrorism but also pave the way for future innovations in applying deep learning and NLP to safeguard the digital sphere.

Acknowledgment

This research is supported by the STARLIGHT "Sustainable Autonomy and Resilience for LEAs using AI against High priority Threats" H2020 project. The project has received funding from the European Union's Horizon 2020 under grant agreement No 101021797.

References

[1] Petrosyan, A. "Number of internet and social media users worldwide as of January 2024". January 2024 – Statista.com

[2] Berger, J. M., Morgan, J. (2015) "The ISIS Twitter Census: Defining and describing the population of ISIS supporters on Twitter" The Brookings Project on U.S. Relations with the Islamic World, Analysis Paper No. 20.

[3] Weimann, G. "Terrorism in Cyberspace: The Next Generation." Columbia University Press. 2015.

[4] Davey, J., Ebner, J. "The Christchurch Attacks: Livestream Terror in the Viral Video Age." The Institute for Strategic Dialogue (ISD). 2019.

[5] Awan, I. "Cyber-Extremism: Isis and the Power of Social Media". Society, 54(2), 138-149. DOI: 10.1007/s12115-017-0114-0. 2017.

[6] Conway, M. "Determining the Role of the Internet in Violent Extremism and Terrorism: Six Suggestions for Progressing Research". Studies in Conflict & Terrorism, 40(1), 77-98. DOI: 10.1080/1057610X.2016.1157408. 2017.

[7] Weimann, G. "Terrorist Migration to the Dark Web". Perspectives on Terrorism, 10(3). 2016.

[8] San Biagio, M., Cipolla, M., La Mattina, E., Morreale, "Revealing terrorist threats in social media". ICISNA 2023.

[9] San Biagio, M., Simoncini, S., La Mattina, E., Morreale, V. "MARPLE: A Framework for Social Media Threat Intelligence," 2024 International Conference on Artificial Intelligence, Computer, Data Sciences and Applications (ACDSA), Victoria, Seychelles, 2024, pp. 1-6, doi: 10.1109/ACDSA59508.2024.10467738.

[10] Ferrara, E., Varol, O., Davis, C., Menczer, F., & Flammini, A. "The rise of social bots". Communications of the ACM, 59(7), 96-104. 2016.

[11] Torres-Soriano, M. R. "The Caliphate's Virtual Stability: The Jihadist Information Warfare and Its Challenge to Western Strategies to Counter Online Terrorism." Studies in Conflict & Terrorism, 43(10), 854-872. 2020.

[12] Bodine-Baron, E., Helmus, T. C., Magnuson, M., & Winkelman, Z. "Examining ISIS Support and Opposition Networks on Twitter." RAND Corporation. 2020.

[13] Ferrara, E., Wang, W.-Q., Varol, O., Flammini, A., Galstyan, A. "Predicting Online Extremism, Content Adopters, and Interaction Reciprocity." EPJ Data Science, 5(1), 17. 2016.

[14] Hussain, M. N, Tokdemir, S., Agarwal N., Al-Khateeb, S. "Analyzing Disinformation and Crowd Manipulation Tactics on YouTube." Frontiers in Communication, 4:56. 2019.

[15] Weimann, G. "Terror on Facebook, Twitter, and Youtube". Brown Journal of World Affairs, 16(2), 45-54. 2010.

[16] Braddock, K., & Horgan, J. "Towards a guide for constructing and disseminating counter-narratives to reduce support for terrorism". Studies in Conflict & Terrorism, 39(5), 381-404. 2016.

[17] Latonero, M., & Kift, P. "On digital passages and borders: Refugees and the new infrastructure for movement and control". Social Media + Society, 4(1), 2056305118764432. 2018.

[18] Jain, P. N., Vaidya, A. "Analysis of Social Media Based on Terrorism - A Review". Vietnam Journal of Computer Science 08(9). 2021.

[19] Devlin, J., Chang, M.W., Lee, K. "BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding". DOI:10.48550/ARXIV.1810.04805. 2018.

[20] Sun, C., Qiu, X., Xu, Y., & Huang, X. "How to Fine-Tune BERT for Text Classification?" In Proceedings of the China National Conference on Chinese Computational Linguistics (CCL) (pp. 194-206). Springer, Cham. 2019.

[21] Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, & Jamie Brew "HuggingFace's Transformers: State-of-the-art Natural Language Processing". CoRR, abs/1910.03771. 2019.

[22] Kaggle – https://www.kaggle.com/datasets/fifthtribe/how-isis-uses-twitter .

[23] Kaggle - https://www.kaggle.com/datasets/rgupta09/world-cup-2018-tweets .

[24] V. Behzadan, C. Aguirre, A. Bose, W. Hsu – Corpus and Deep Learning Classfier for Collection of Cyber Threat Indicators in Twitter Stream. Big Data, 5002-5007. (2018).

[25] Cortes, C., Vapnik, V.. "Support-vector networks", Machine Learning, 273-297, 1995.

[26] Cover, T. M., Hart, P. E.. "Nearest neighbor pattern classification". IEEE Transactions on Information Theory. 13 (1): 21–27, 1967.

[27] Breiman, L., Friedman, J., Stone, C. J., Olshen, R.A. "Classification and Regression Trees", Group, 37(15):237–251, 1984.

[28] Ho, T. K. "Random Decision Forests". In ICDAR pp. 278–282, 1995.

[29] Murtagh, F. "Multilayer perceptrons for classification and regression". Neurocomputing Volume 2, Issues 5–6, Pages 183-197. 1991