Smart City Security: Fake News Detection in Consumer Electronics

Siyou Guo
Shandong University of Technology

Qilei Li
Queen Mary University of London

Mingliang Gao, Guisheng Zhang, Gwanggil Jeon
Shandong University of Technology

***Abstract*—The rapid and widespread adoption of consumer electronics and technology, such as smartphones, smart speakers, and IoT devices, has fundamentally changed how we access and consume information. This technological revolution has not only enhanced our daily convenience but has also driven the development of smart cities and brought significant advancements in urban living. However, this digital convenience comes with challenges, notably the production of fake news. Through various consumer electronics platforms, false information can be quickly produced and spread, undermining public trust and social order. Although fake news-detecting technology has advanced rapidly benefitting from deep-learning techniques, it often fails to consider feature interactions. To address this issue, we propose an Enhanced Feature Interactions Network (EFI-Net) for fake news detection. Specifically, the EFI-Net introduces an Efficient Additive Learning (EAL) module to enhance feature interaction for language models at different scales. Experiments were conducted using the ARG fake news detection dataset, and the proposed network achieves an accuracy of 88.9% on English fake news and 78.7% on Chinese fake news, which outperforms the state-of-the-art (SOTA) method by a large margin. This work has substantial implications for consumer electronics and technology. Users can benefit from more reliable information filtering and verification by integrating EFI-Net into various consumer electronic platforms.**

THE DEVELOPMENT OF SMART CITIES hinges on the seamless integration of consumer electronics technology. From smartphones and smart home appliances to a vast network of sensors and IoT devices, these advancements promise greater convenience in urban management and an enhanced quality of life for residents. This interconnected ecosystem, however, presents unique challenges, particularly in the realm of cybersecurity and data privacy, with the spread of fake news emerging as a critical threat [1], [2].

Digital Object Identifier 10.1109/MCE.YYYY.Doi Number

Date of publication DD MM YYYY; date of current version DD MM YYYY

The latest developments in consumer electronics for smart cities are closely linked to integrating the Internet of Things, wireless communication technologies, and various information-physical systems. These

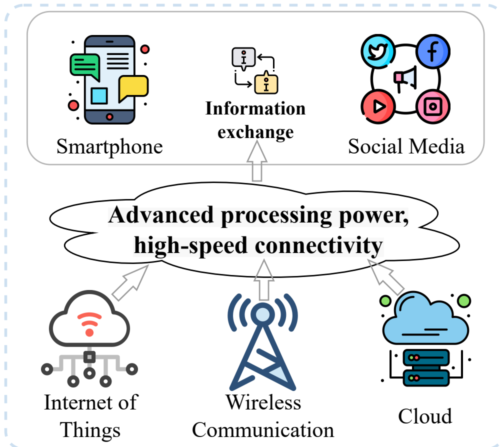

Figure 1. Illustration of information production and dissemination from consumer electronic devices empowered by high-speed connectivity.

developments aim to improve the quality of life, infrastructure, and economic and social progress through seamless digital services. As shown in Figure 1, the latest generation of smartphones, with their advanced processing power, high-speed connectivity, and integrated social media platforms, serve as primary conduits for information consumption and dissemination. This reliance on interconnected systems, while enabling real-time monitoring, analysis, and optimization of urban functions, also introduces vulnerabilities that can be exploited for malicious purposes.

A prime example of this vulnerability is the proliferation of fake news [3], [4], [5]. Fake news refers to the deliberate presentation of false or misleading claims as news, where these claims are misleading [6]. Fake news can rapidly propagate through the interconnected network of consumer electronic devices, impacting public perception and behavior. This is particularly concerning in the context of smart cities, which rely heavily on real-time information for critical decision-making. The spread of disinformation through social media feeds on smartphones or manipulated data on urban dashboards can lead to public safety hazards, economic disruption, and erosion of trust in government institutions.

Current state-of-the-art fake news detection approaches primarily focus on social context and content analysis [7]. Social-context-based methods leverage the dynamics of information diffusion to analyze propagation patterns, user reactions, and social network structures. This approach is particularly relevant to consumer electronics, as it can utilize data from social media interactions and device usage patterns. Content-based methods, on the other hand, scrutinize the information itself and inconsistencies with established facts. Advanced natural language processing techniques, which often employ pre-trained language models like BERT [8], are used to discern linguistic cues indicative of fake news. Despite these advancements, existing methods often struggle to capture the complex interplay between textual features and broader contextual information. This limitation hinders their ability to accurately assess the veracity of news, particularly in the face of increasingly sophisticated disinformation campaigns.

Addressing this challenge requires a multi-pronged approach. Firstly, future consumer electronics should directly incorporate advanced fake news detection mechanisms into their operating systems. Secondly, research efforts should focus on developing more sophisticated algorithms that can effectively capture the nuances of language and context. Finally, raising public awareness about the dangers of fake news and promoting media literacy is crucial to empower individuals to evaluate the information they encounter on their devices critically. Tackling fake news in consumer electronics ensures smart cities deliver enhanced convenience, improved quality of life, and efficient urban management while safeguarding security, privacy, and citizen well-being.

To address the above-mentioned considerations, we propose an Enhanced Feature Interactions Network (EFI-Net) for fake news detection. The EFI-Net aggregates small and large language models and introduces an Efficient Additive Learning (EAL) module [9] to enhance feature interaction. The incorporation of EFI-Net into consumer electronics offers a solution for enhancing the reliability and security of information access. For instance, EFI-Net can detect and filter fake news on smartphones and smart speakers. In IoT applications, EFI-Net provides a mechanism for authenticating data transmitted to clients. The main contributions of this paper are outlined as follows:

- We provide a comprehensive review of consumer electronics technologies for smart cities, including their roles and weaknesses. We emphasize the dangers of fake news and the limitations of existing detection methods.
- We propose a fake news detection framework that leverages the Large Language Model (LLM) and

machine learning techniques to identify and classify deceptive content. We introduce an Efficient Additive Learning (EAL) module into the feature layer to enhance feature interactions.

- We conduct comprehensive experiments on publicly available datasets, and the results demonstrate that the proposed model outperforms the SOTA methods in fake news detection.

The rest of this article is organized as follows. The "Related Work" Section reviews current methods for fake news detection. The "Fake News Detection System" Section delves into the details of the proposed model. The "Experimental Results" Section covers the experimental setup, the datasets employed, and the performance evaluation of the proposed model. Lastly, the "Conclusion and Future Work" Section presents the main findings of this study and suggests potential future directions.

RELATED WORK

Fake news detection has been considered one of the important topics by researchers. Manual fact-checking is a common strategy to mitigate the adverse effects of fake news dissemination [10]. However, it is insufficient when dealing with the vast amount of generated information. To solve these problems, deep learning models have been employed to identify fake news automatically. EANN_T [11] learns effective signals using auxiliary adversarial training, aiming at removing event-related features as much as possible. Zhang *et al.* [12] employed publisher emotions and social emotions to detect fake news. Zhu *et al.* [13] proposed an entity debiasing framework (ENDEF) to address entity bias in training fake news detection models. Recently, LLMs have been proven to have excellent ability in multiple classification and reasoning tasks. Several researchers [14], [15] emphasize the role of LLMs in fake news detection. For instance, Teo *et al.* [16] combined conventional machine learning techniques with large language models to address the challenges posed by manually crafted features, which is time-consuming and less effective. Hu *et al.* [15] considered that LLMs are suboptimal at veracity judgment but proficient at content analysis. Wu *et al.* [17] utilized LLM to rephrase news, introduce stylistic diversity, and extract content-centered guidelines from LLM to debunk fake news. Luvembe *et al.* [18] optimize feature interaction by ensuring inter-modal semantic consistency through bidirectional complementary attention between modalities. Rani *et al.* [19] integrated sophisticated fake news detection models into smartphones, smart speakers, and IoT devices to enhance the reliability and trustworthiness of the information consumed by users. Wang *et al.* [20] demonstrated that effective fake news detection mechanisms can mitigate the spread of misinformation, thus helping to maintain social order and trust. For the consumer electronics industry, integrating these technologies can enhance the value and reliability of their products [21]. In a broader sense, advancing fake news detection contributes to the resilience of democratic processes and the integrity of public discourse, underscoring its critical role in modern society.

FAKE NEWS DETECTION SYSTEM

We propose a framework named EFI-Net to address the challenge of fake news detection in consumer electronics. The architecture of EFI-Net is shown in Figure 2. It consists of three key components: (a) Efficient Additive Learning, (b) LLM Reasoning Rationale, and (c) Prediction.

The news text and the rationales generated by the LLM are encoded using a pre-trained Small Language Model (SLM) [8] to capture their semantic representations. Subsequently, Efficient Additive Learning (EAL) takes these semantic representations as input and transforms them to generate enhanced features. These enhanced features are expected to be more informative for fake news detection. The LLM Reasoning Rationale component utilizes a dedicated block named "News-Rationale Collaboration". This block combines the semantic representation of the news text with the representations obtained after encoding the LLM-generated rationales. It aims to extract new features that capture the reasoning and context provided by the rationales, and it seeks to achieve a deeper understanding of news content.

Efficient Additive Learning

To enhance feature interactions, we introduce an Efficient Additive Learning (EAL) module into the feature layer of the baseline ARG framework [15]. The structure of EAL is shown in Figure 3.

The EAL module enhances feature interaction at the feature layer through the specific "query-key" interaction mechanism. Compared to attention mechanisms, the EAL utilizes the interactions between queries and keys to learn global context information without explicitly computing key-value interactions.

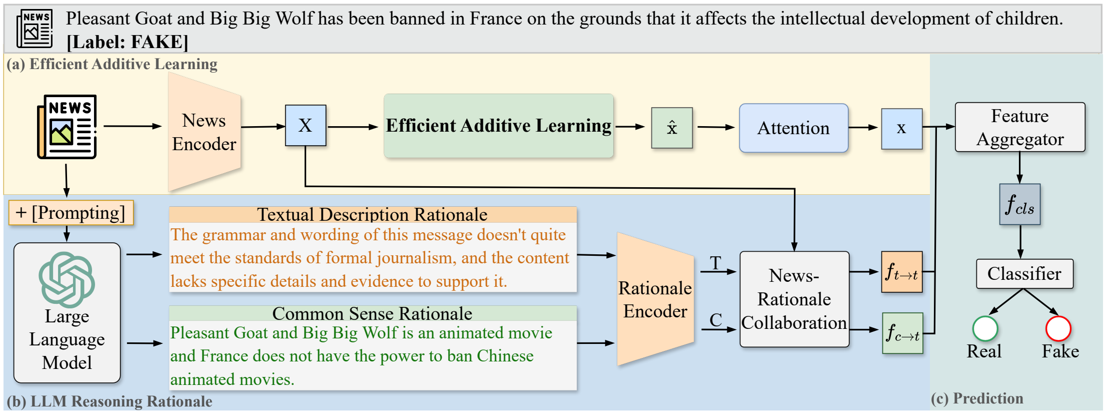

Figure 2. Illustration of Enhanced Feature Interactions Network (EFI-Net) for fake news detection.

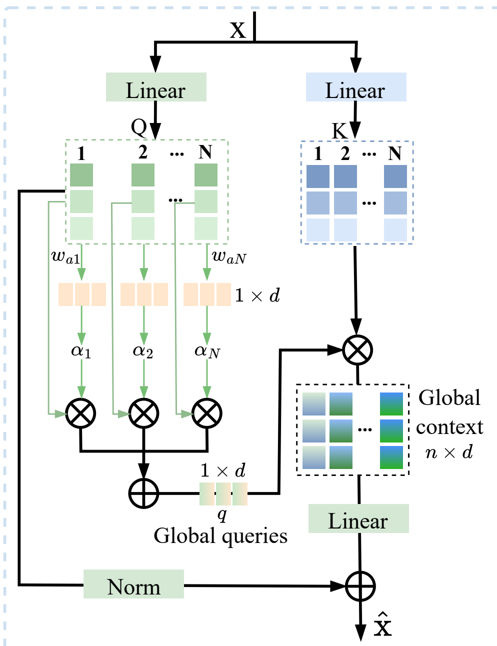

Figure 3. Illustration of the Efficient Addition Learning (EAL) module.

The input embedding matrices are projected onto the query matrices (Q) and key matrices (K) separately to represent the input features from different perspectives. The query matrix captures the "what" and the key matrix captures the "how" of the information within the semantic representation.

To understand the relative significance of each element in the query matrix, EAL multiplies the query matrix (Q) with a vector of learnable parameters (w_a) and then divides the square root of the embedding dimension (d). The results α can be considered a weighting system, assigning importance scores to each element in the query.

The query matrix is multiplied by w_a and pooled to obtain a single global query vector. The vector (q) represents the importance of all query elements and is used for subsequent contextual information learning.

The global query vector (q) is multiplied element by element with the key matrix (K) to obtain the global context matrix. This matrix contains information from each key and has the flexibility to learn correlations in the input sequence. The global context matrix is linearly transformed to obtain a final weighted output. The output contains rich contextual information and is used for subsequent downstream tasks.

Large Language Model Reasoning Rationale

Large Language Models (LLMs) are complex neural networks trained with large amounts of textual data to generate and understand natural language text. Although LLMs do not perform as well as specially fine-tuned SLMs (*e.g.*, BERT [8]) in independently judging news truthfulness, the multifaceted explanations generated by LLMs can improve the judgments of SLMs [15].

As shown in Figure 2 (b), the common sense reasoning and textual description reasoning capabilities of LLM are used to verify the authenticity of news content. The LLM uses prompt learning to generate truthful judgments about news. To facilitate thorough information exchange between rationales and news, the News-Rationale Collaboration block incorporates a dual cross-attention to enhance feature interactions.

Prediction and Loss Function

As shown in Figure 2, we aggregate the rationale-aware news vectors $f_{t\to x}$ and $f_{c\to x}$ and the news vector

Table 1. Comparison results with the SOTA methods. The best results are highlighted in bold.

Methods	Model	English				Chinese			
		macro F1	Accuracy	$F1_{real}$	$F1_{fake}$	macro F1	Accuracy	$F1_{real}$	$F1_{fake}$
$EANN_T$ [11]	SLM	0.763	0.864	0.918	0.608	0.754	0.756	0.773	0.736
Publisher-Emo [12]	SLM	0.766	0.868	0.920	0.611	0.761	0.763	0.784	0.738
ENDEF [13]	SLM	0.768	0.865	0.918	0.618	0.765	0.766	0.779	0.751
SuperICL [14]	LLM+SLM	0.736	0.864	0.920	0.551	0.757	0.759	0.779	0.734
GPT-3.5-turbo [15]	LLM	0.702	0.813	0.884	0.519	0.725	0.734	0.774	0.676
Baseline [15]	LLM+SLM	0.790	0.878	0.926	0.653	0.784	0.786	0.804	0.764
EFI-Net (Ours)	LLM+SLM	**0.802**	**0.889**	**0.933**	**0.671**	**0.785**	**0.787**	**0.804**	**0.767**

x to make the final decision. For the news content labelled $y \in \{0,1\}$, these vectors are aggregated with varying weights:

$$f_{\text{cls}} = w_1 \cdot \text{x} + w_2 \cdot f_{c \to x} + w_3 \cdot f_{t \to x}, \tag{1}$$

where w_1, w_2, and w_3 are learnable parameters ranging from 0 to 1. The fusion vector f_{cls} is input into the MLP classifier to finalize the veracity of the news:

$$L_{\text{ce}} = f_{\text{ce}}(\text{MLP}(f_{\text{cls}}), y), \tag{2}$$

where f_{ce} denotes the cross-entropy classification loss.

The model analyses each rationale and predicts its usefulness y_t for determining the truthfulness of the news. Based on these predictions, the model adjusts the importance of each rationale before making its final prediction. Assuming that rationales leading to correct judgments are more useful, we use the correctness of the judgments as the rationale usefulness labels. The loss formula for determining usefulness is as follows:

$$\hat{y}_t = \text{sigmoid}(\text{MLP}(f_{x \to t})), \tag{3}$$

$$L_{\text{et}} = f_{\text{ce}}(\hat{y}_t, y_t), \tag{4}$$

The overall loss function is calculated as the weighted sum of the previously mentioned loss terms:

$$L = L_{\text{ce}} + \lambda (L_{\text{ec}} + L_{\text{et}}), \tag{5}$$

where λ is hyperparameter.

EXPERIMENTAL RESULTS

To validate the effectiveness of the proposed module, we used the ARG [15] dataset. This dataset includes a Chinese subset and an English subset. This dataset has been preprocessed by removing duplicates and segmenting temporal data. These steps are taken to prevent over-performance of the SLM due to data leakage. We employ Adam as the optimizer and conduct a grid search to determine the best learning rate. A grid search was performed for the weights of the loss functions L_{ce}, L_{ec}, and L_{et} in EFI-Net. The "News-Rationale Collaboration" incorporates a four-head transformer block. We implemented the framework using PyTorch [22] and trained it on two 3090 Ti GPUs in parallel.

In this work, four evaluation metrics are used. Accuracy represents the proportion of correctly classified instances of fake and real news. $F1_{\text{real}}$ and $F1_{\text{fake}}$ are employed to evaluate the accuracy in identifying real news and fake news. The macro F1 [23] calculates the F1 scores (*i.e.*, $F1_{\text{real}}$ and $F1_{\text{fake}}$) for each category and then averages these scores.

The experimental results are shown in Table 1. It shows that the EFI-Net model achieves the best results across four key metrics on both Chinese and English news datasets. EFI-Net scores an 88.9% detection accuracy rate on the English fake news dataset and 78.7% on the Chinese dataset. Specifically, the EFI-Net outperforms the ENDEF [13] and GPT-3.5-turbo [15] in all the evaluation metrics. The reason lies that the EFI-Net uses both large language models (LLM) and small language models (SLM), while ENDEF uses SLM-only and GPT-3.5-turbo [15] uses LLM-only. EFI-Net combines the advantages of LLMs and SLMs. LLMs offer richer semantic understanding and generative capabilities, while SLMs excel at capturing language features specific to certain domains. This synergy enables EFI-Net to achieve a more comprehensive understanding of news content and accurately assess its veracity. SuperICL [14] achieves better results than the LLM-only approach, but it does not consistently exceed the baseline SLM across both datasets. Furthermore, we compared approaches that utilize both large and small language models, *i.e.*, SuperICL [14] and baseline method [15]. The proposed EFI-Net outperforms the two hybrid models. The EAL module enhances feature interaction through a ”query-key” mechanism. This allows the model to better capture the complex relationships between textual features, thereby improving detection accuracy.

CONCLUSION AND FUTURE WORK

Fake news erodes public trust and hinders the development of consumer electronics for smart cities. This article proposes a fake news detection framework aggregating SLM and LLM and introduces the EAL module to enhance feature interaction, which enhances the accuracy of the detection system. Despite advancements in fake news detection for English datasets, performance on Chinese datasets still has room for improvement. Future research directions include exploring the potential of cross-lingual transfer learning and multilingual models to further enhance the accuracy and reliability of Chinese fake news detection and ultimately improve trust in consumer electronics within smart cities.

REFERENCES

1. S. Saeedi, A. C. Fong, S. P. Mohanty, A. K. Gupta, and S. Carr, “Consumer artificial intelligence mishaps and mitigation strategies,” *IEEE Consumer Electronics Magazine*, vol. 11, no. 3, pp. 13–24, 2021.
2. S. Guo, Q. Li, M. Gao, G. Zhang, J. Pan, and G. Jeon, “Deep learning-based face forgery detection for facial payment systems,” *IEEE Consumer Electronics Magazine*, 2024.
3. K. Michael and R. Abbas, “Behind covid-19 contact trace apps: The google–apple partnership,” *IEEE Consumer electronics magazine*, vol. 9, no. 5, pp. 71–76, 2020.
4. G. Zhang, M. Gao, Q. Li, W. Zhai, and G. Jeon, “Multi-modal generative deepfake detection via visual-language pretraining with gate fusion for cognitive computation,” *Cognitive Computation*, pp. 1–14, 2024.
5. Q. Li, M. Gao, G. Zhang, W. Zhai, J. Chen, and G. Jeon, “Towards multimodal disinformation detection by vision-language knowledge interaction,” *Information Fusion*, vol. 102, p. 102037, 2024.
6. A. Gelfert, “Fake news: A definition,” *Informal logic*, vol. 38, no. 1, pp. 84–117, 2018.
7. E. Aïmeur, S. Amri, and G. Brassard, “Fake news, disinformation and misinformation in social media: a review,” *Social Network Analysis and Mining*, vol. 13, no. 1, p. 30, 2023.
8. J. D. M.-W. C. Kenton and L. K. Toutanova, “Bert: Pre-training of deep bidirectional transformers for language understanding,” in *Proceedings of naacL-HLT*, vol. 1, 2019, p. 2.
9. A. Shaker, M. Maaz, H. Rasheed, S. Khan, M.-H. Yang, and F. S. Khan, “Swiftformer: Efficient additive attention for transformer-based real-time mobile vision applications,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023, pp. 17 425–17 436.
10. N. Capuano, G. Fenza, V. Loia, and F. D. Nota, “Content-based fake news detection with machine and deep learning: A systematic review,” *Neurocomputing*, vol. 530, pp. 91–103, 2023.
11. Y. Wang, F. Ma, Z. Jin, Y. Yuan, G. Xun, K. Jha, L. Su, and J. Gao, “Eann: Event adversarial neural networks for multi-modal fake news detection,” in *Proceedings of the 24th acm sigkdd international conference on knowledge discovery & data mining*, 2018, pp. 849–857.
12. X. Zhang, J. Cao, X. Li, Q. Sheng, L. Zhong, and K. Shu, “Mining dual emotion for fake news detection,” in *Proceedings of the web conference 2021*, 2021, pp. 3465–3476.
13. Y. Zhu, Q. Sheng, J. Cao, S. Li, D. Wang, and F. Zhuang, “Generalizing to the future: Mitigating entity bias in fake news detection,” in *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2022, pp. 2120–2125.
14. C. Xu, Y. Xu, S. Wang, Y. Liu, C. Zhu, and J. McAuley, “Small models are valuable plug-ins for large language models,” *arXiv preprint arXiv:2305.08848*, 2023.
15. B. Hu, Q. Sheng, J. Cao, Y. Shi, Y. Li, D. Wang, and P. Qi, “Bad actor, good advisor: Exploring the role of large language models in fake news detection,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 38, no. 20, 2024, pp. 22 105–22 113.
16. T. W. Teo, H. N. Chua, M. B. Jasser, and R. T. Wong, “Integrating large language models and machine learning for fake news detection,” in *2024 20th IEEE International Colloquium on Signal Processing & Its Applications (CSPA)*. IEEE, 2024, pp. 102–107.
17. J. Wu, J. Guo, and B. Hooi, “Fake news in sheep’s clothing: Robust fake news detection against llm-empowered style attacks,” in *Proceedings of the*

30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 2024, pp. 3367–3378.

18. A. M. Luvembe, W. Li, S. Li, F. Liu, and X. Wu, “Caf-odnn: Complementary attention fusion with optimized deep neural network for multimodal fake news detection,” *Information Processing & Management*, vol. 61, no. 3, p. 103653, 2024.
19. P. Rani and J. Shokeen, “Fnnet: a secure ensemble-based approach for fake news detection using blockchain,” *The Journal of Supercomputing*, pp. 1–38, 2024.
20. S. Wang, X. Xu, X. Zhang, Y. Wang, and W. Song, “Veracity-aware and event-driven personalized news recommendation for fake news mitigation,” in *Proceedings of the ACM web conference 2022*, 2022, pp. 3673–3684.
21. W. Wan, W. Sun, Q. Zeng, L. Pan, and J. Xu, “Progress in artificial intelligence applications based on the combination of self-driven sensors and deep learning,” in *2024 4th International Conference on Consumer Electronics and Computer Engineering (ICCECE)*. IEEE, 2024, pp. 279–284.
22. A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga *et al.*, “Pytorch: An imperative style, high-performance deep learning library,” *Advances in neural information processing systems*, vol. 32, pp. 8024–8035, 2019.
23. D. Chicco and G. Jurman, “The advantages of the matthews correlation coefficient (mcc) over f1 score and accuracy in binary classification evaluation,” *BMC genomics*, vol. 21, pp. 1–13, 2020.

Siyou Guo is currently working toward an M.S. degree with the School of Shandong University of Technology, Zibo, China. Contact him at 23504030565@stumail.sdut.edu.cn.

Qilei Li is working toward a Ph.D. with the University of Queen Mary University of London, London, E1 4NS, United Kingdom. He contributed equally to this work with Siyou Guo. Contact him at q.li@qmul.ac.uk.

Mingliang Gao is an associate professor at the Shandong University of Technology. He is the first corresponding author of this article. Contact him at mlgao@sdut.edu.cn.

Guisheng Zhang is working toward an M.S. degree with the School of Shandong University of Technology, Zibo, China. Contact him at 22504030001@stumail.sdut.edu.cn.

Gwanggil Jeon is a professor at Shandong University of Technology, Zibo, China. He is the second corresponding author of this article. Contact him at ggjeon@gmail.com.