



**HAL**  
open science

# Scalable Framework for Classifying AI-Generated Content Across Modalities

Anh-Kiet Duong, Petra Gomez-Krämer

► **To cite this version:**

Anh-Kiet Duong, Petra Gomez-Krämer. Scalable Framework for Classifying AI-Generated Content Across Modalities. Defactify4 @ AAI 2025, AAI, Jan 2025, Philadelphia, United States. 10.48550/arXiv.2502.00375 . hal-05040148

**HAL Id: hal-05040148**

**<https://hal.science/hal-05040148v1>**

Submitted on 18 Apr 2025

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

# Scalable Framework for Classifying AI-Generated Content Across Modalities

Anh-Kiet Duong<sup>1,\*</sup>, Petra Gomez-Krämer<sup>1</sup>

<sup>1</sup>*L3i Laboratory, La Rochelle University, Avenue Michel Crépeau, 17042 La Rochelle Cedex 1 - France*

## Abstract

The rapid growth of generative AI technologies has heightened the importance of effectively distinguishing between human and AI-generated content, as well as classifying outputs from diverse generative models. This paper presents a scalable framework that integrates perceptual hashing, similarity measurement, and pseudo-labeling to address these challenges. Our method enables the incorporation of new generative models without retraining, ensuring adaptability and robustness in dynamic scenarios. Comprehensive evaluations on the Defactify4 dataset demonstrate competitive performance in text and image classification tasks, achieving high accuracy across both distinguishing human and AI-generated content and classifying among generative methods. These results highlight the framework's potential for real-world applications as generative AI continues to evolve. Source codes are publicly available at <https://github.com/ffyytt/defactify4>.

## Keywords

AI-generated content classification, incremental learning, pseudo-labeling, contrastive learning

## 1. Introduction

The proliferation of generative artificial intelligence (AI) technologies has introduced new societal challenges, particularly in distinguishing between human-generated and AI-generated content [1]. The increasing sophistication of these models allows them to produce text and images that are often indistinguishable from human creations, raising concerns over their misuse in spreading misinformation, generating fake news, and creating deceptive media. Beyond this, the diversity of generative AI models, each employing distinct architectures and techniques highlights the importance of not only identifying AI-generated content but also classifying it by its source. Such classifications enable deeper forensic insights and are critical for trust-building in applications like content moderation, digital forensics, and fact-checking.

A critical aspect of this challenge is the need to differentiate between the various generative methods themselves. This distinction is essential for assessing the relative difficulty of detecting different approaches and understanding which models are more or less susceptible to detection. Without this understanding, it would be difficult to evaluate the effectiveness of detection systems and their ability to keep pace with emerging AI technologies [2].

Compounding these challenges is the relentless growth in the number and variety of generative models. The landscape of generative AI is far from static, with new methods continually emerging and pushing the boundaries of realism and creativity. This rapid evolution necessitates the development of adaptable classification frameworks capable of integrating new generative models without requiring costly and exhaustive retraining [3]. Addressing this issue is crucial for building sustainable and scalable detection systems.

To address these challenges, this paper proposes a novel approach tailored for classifying and detecting AI-generated content in both text and image domains. The key contributions of this work are:

---

*De-Factify 4: 4rd Workshop on Multimodal Fact Checking and Hate Speech Detection, co-located with AAAI 2025. 2025 Philadelphia, Pennsylvania, USA*

\*Corresponding author.

✉ [anh.duong@univ-lr.fr](mailto:anh.duong@univ-lr.fr) (A. Duong); [petra.gomez@univ-lr.fr](mailto:petra.gomez@univ-lr.fr) (P. Gomez-Krämer)

🌐 <https://ffyytt.github.io/> (A. Duong)

🆔 0009-0005-0230-6104 (A. Duong); 0000-0002-5515-7828 (P. Gomez-Krämer)



© 2024 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

- A method that adapts to new generative models by incorporating their features into the classification pipeline without retraining
- The evaluation of contrastive learning for incremental learning, particularly in classifying AI, human-generated, and various generative methods
- The use of pseudo-labeling to familiarize the model with augmentations in the test set, to enrich the training data, and to enable long-term learning that allows adaptation to small changes in generative models in real-world scenarios

In recent years, advancements in generative models have significantly enhanced the capabilities of AI in creating text and images. Large language models (LLMs) such as GPT-3 [4], GPT-4 [5], and Mistral [6] have revolutionized text generation, enabling coherent and contextually rich outputs across diverse applications, but also raising challenges in detecting AI-generated text. Similarly, generative models like Stable Diffusion [7], DALL-E [8], Midjourney [9] have achieved remarkable success in producing high-quality, photorealistic images. However, the increasing realism of such outputs poses risks for misinformation and underscores the need for robust detection systems. Recent studies have explored ensemble methods combining multiple models [10, 11] or treating fake content as anomalies [12], achieving notable success in feature separability. However, these approaches often fall short in scalability when faced with the addition of new labels. Techniques like ArcFace have addressed some of these limitations, offering improved adaptability and performance in scenarios with evolving generative models. The remaining details and related works are provided in Appendix 5.1.

This paper is organized to provide a comprehensive overview of the proposed method and its evaluation. Section 2 introduces the method, detailing its key components, including perceptual hashing, similarity measurement, and pseudo-labeling, while Section 3 presents the experimental setup, datasets, implementation details, and results to highlight the method’s effectiveness. The paper concludes in Section 4 with insights and future directions, with additional materials, including detailed related work, dataset descriptions, and supplementary experiments, provided in the Appendix (Sec. 5).

## 2. Method

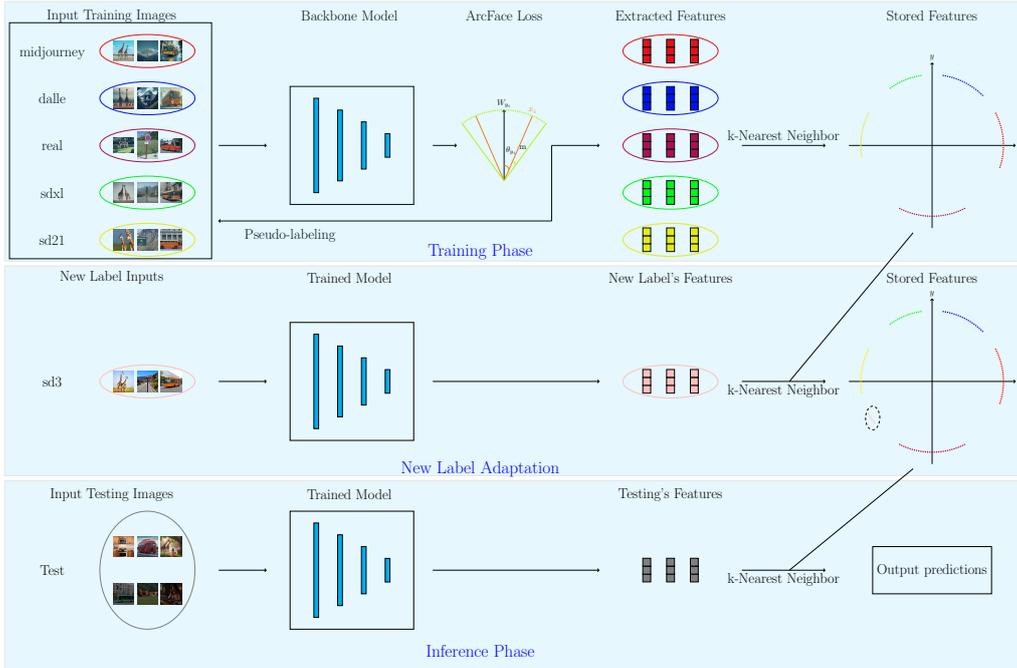
The proposed method comprises three main components: perceptual hashing (Sec. 2.1), similarity measurement and comparison (Sec. 2.2), and pseudo-labeling (Sec. 2.3). Each component is designed to address specific challenges in distinguishing between human-generated and AI-generated content, as well as classifying content from different generative models.

In Figure 1 we illustrate the overall framework, which consists of three stages: training, new label adaptation, and inference. In the training stage, we train the model using the ArcFace loss (Eq. (3)) and apply pseudo-labeling to augment the dataset. Features are extracted from the input data and stored for later use, with pruning techniques for  $k$ -nearest neighbors employed to minimize storage requirements [13]. In the new label adaptation stage, features are extracted from the new data and seamlessly integrated into the existing feature storage. Finally, in the inference stage, features are extracted from incoming data and compared with the stored features to determine the most similar label. For text we adopted BART Large [14], and for image we used Swin Transformer V2 Base [15] as backbone model. This design ensures scalability and adaptability to new generative models while maintaining robust performance.

### 2.1. Perceptual hashing

Perceptual hashing is key to our method, offering a compact and efficient data representation for comparison. We train a model with ArcFace loss [16] to enhance feature discrimination by enforcing class margins. The trained model outputs high-dimensional feature vectors, which serve as perceptual hash digests for the input samples.

Once trained, the model extracts features from all samples, converting them into high-dimensional feature vectors, referred to as hashing digests. These vectors are stored in a database and used for



**Figure 1:** Overall framework of the proposed method, which consists of three stages, *i.e.*, training, new label adaptation, and inference.

similarity comparisons. This approach enables efficient comparison and retrieval, leveraging the discriminative capability of the learned representations to distinguish between human-generated and AI-generated content, as well as between different generation models.

## 2.2. Similarity measurement and comparison

To compare features, we employ a  $k$ -nearest neighbor ( $k$ -NN) approach, inspired by solutions from prior competitions [17, 18]. This non-parametric method allows us to measure the similarity between samples effectively, leveraging the local neighborhood structure within the feature space. Specifically, we calculate the similarity metric using the cosine similarity, which is well-suited for high-dimensional feature spaces and ensures scale-invariant comparisons. The simplicity and adaptability of  $k$ -NN make it particularly suited for our scenario, where the inclusion of new labels or generative models is anticipated.

When new labels appear, our method avoids retraining the entire model. Instead, features from a few representative samples generated by the new AI model are extracted and then appended to the existing feature set. Seamlessly integrating the new class into the similarity-based comparison process. This approach ensures a scalable and efficient adaptation mechanism, maintaining the system's flexibility to accommodate evolving AI-generated content without significant computational overhead.

## 2.3. Pseudo labeling

Pseudo-labeling is employed as a crucial component in our approach to address the limitations of a small labeled training dataset while improving the model's adaptability and robustness. By leveraging unlabeled data, pseudo-labeling effectively expands the training set, where high-confidence predictions of the model are used as surrogate labels [19]. This augmentation strategy not only increases the diversity of the training samples but also familiarizes the model with potential augmentations present in test data, thereby enhancing generalization performance.

In deployment scenarios, pseudo-labeling demonstrates its practicality by exploiting the abundance of unlabeled data derived from user inputs. These inputs, typically available in substantial volumes and without explicit labels, provide a valuable resource for iterative model training and fine-tuning.

Furthermore, as AI-generated content evolves with updates to generative models over time, pseudo-labeling allows the model to adapt dynamically to these changes. This adaptability ensures robust performance in real-world applications, where continuous learning from new and diverse data is essential to maintaining high accuracy and relevance.

Similar to the approach proposed in [17], we utilize a dynamic pseudo-labeling mechanism to enhance our training process. Specifically, at each training epoch, the model is used to predict labels for the test set, and only the top  $p_{pseudo}$  percent of predictions with the highest confidence probabilities are selected as pseudo-labeled data. These high-confidence samples are then added to the training process to further refine the model.

In subsequent epochs, the pseudo-labeling step is repeated, and a new set of top  $p_{pseudo}$  percent high-confidence predictions is selected. This iterative process helps to progressively eliminate potentially incorrect pseudo-labels, ensuring that only reliable samples contribute to the training. Importantly, the original labeled training data remains the primary focus during training, as it provides a reliable foundation with ground-truth labels. The pseudo-labeled samples constitute only a small fraction of the training data, serving as a complementary augmentation to improve model adaptability. The training continues in this manner until the model converges, leveraging the evolving predictions to adapt effectively and maintain robustness against noisy pseudo-labels.

### 3. Experiments

We conduct experiments to evaluate the effectiveness of our proposed method on both image and text datasets from the Defactify4 competition. The experiments are designed to test the model’s ability to distinguish between human-generated and AI-generated content, to classify content across different generative methods, and to adapt to new labels. Detailed descriptions of the datasets (Sec. 5.3), implementation (Sec. 3.2), and results (Sec. 3.3) are provided in the following subsections.

#### 3.1. Dataset

Our experiments utilize two benchmark datasets: **Defactify4-Image** and **Defactify4-Text**, each designed to evaluate AI vs. human classification and method-specific categorization. Both datasets comprise training and testing splits, with the testing sets further divided into unaltered data (Test 1) and augmented data (Test 2) to assess model robustness and generalization. **Defactify4-Image** consists of six classes, where one class represents real images derived from the COCO dataset [20], and the remaining five classes correspond to outputs from different AI image-generation models. Meanwhile, **Defactify4-Text** comprises seven classes, including one class of human-written text and six classes generated by various text-generation models. Detailed descriptions of the datasets are provided in Appendix 5.3.

#### 3.2. Implementation

For the backbone model used to extract features from images, we adopt a Swin Transformer V2 Base model [15] with an image size of  $256 \times 256$  and a window size of 8, pre-trained on ImageNet [21]. For text data, we use a pre-trained BART Large model [14] with a maximum token length of 512. Training is done for 100 epochs using Adam optimizer [22] with a learning rate of  $1 \times 10^{-4}$  and a batch size of 32. In addition to the primary inputs, we incorporate auxiliary data that we found to be significant. For text data, we include the text length in terms of the number of characters and words as additional features. For image data, we leverage the image size as a supplementary input. After feature pooling from the backbone and concat with auxiliary data, we apply a fully connected layer with 512 nodes to reduce the dimensionality of the features, using the Parametric ReLU (PReLU) function as activation layer.

For  $k$ -nearest neighbors ( $k$ -NN), we use the implementation provided by scikit-learn [13] with default parameters and adopt cosine as the metric, enabling faster computations and pruning to reduce the

number of stored features. In the pseudo-labeling process, we select the top  $p_{pseudo} = 5\%$  of predictions with the highest confidence probabilities for each predicted label  $\hat{y}$ , and assign  $\hat{y}$  as the label for the corresponding data points.

Data augmentation is applied to both image and text data to enhance robustness. For text, we introduce random starting and ending positions in sequences and inject random meaningless strings at arbitrary points within the data. For image data, we employ augmentations including horizontal flip, Gaussian noise, image compression, and random brightness/contrast adjustments. To ensure compatibility with the input pipeline, resizing operations are performed while preserving the aspect ratio of the original image.

### 3.3. Results

In this section, we present the results obtained from evaluating our proposed method on both image and text datasets. The competition for each data type is divided into two tasks: Task A focuses on distinguishing between AI-generated and human-produced content, while Task B classifies the content across different methods, with human-produced as one of the categories.

**Table 1**

Leaderboard of the **Defactify4-Image** task

Team Name	Task A	Task B
SeeTrails	0.8334	0.4986
NYCU	0.8329	0.491
random.py	0.8326	0.4936
Xiaoyu	0.8316	0.4888
TAHAKOM	0.8305	0.4816
SKDU	0.83	0.4864
Nitiz	0.8152	0.4193
OAR	0.7996	0.2726
RoVIT	0.759	0.4222
dakiet (our method)	0.833	0.4935

**Table 2**

Leaderboard of the **Defactify4-Text** task

Team Name	Task A	Task B
Sarang	1	0.9531
tesla	0.9962	0.9218
SKDU	0.9945	0.7615
Drocks	0.9941	0.627
Llama_Mamba	0.988	0.4551
AI_Blues	0.9547	0.3697
NLP_great	0.9157	0.1874
Osint	0.8982	0.3072
Xiaoyu	0.803	0.5696
Rohan	0.7546	0.4053
dakiet (our method)	0.9999	0.9082

In Table 2 and Table 1, we present the leaderboards (evaluated on Test 2) for the text and image tasks, respectively, in the competition. Our proposed method demonstrates competitive performance compared to other teams. Specifically, we achieved the 2nd place in Task A and the 3rd place in Task B for both text and image datasets. This highlights the effectiveness of our method, which not only achieves high accuracy but also incorporates the ability for continual learning through pseudo-labeling. More importantly, our method has the potential to expand the number of labels, accommodating new classes in the future. This is crucial given the increasing diversity of AI-generated models and outputs. The capability to seamlessly add new labels ensures our approach remains adaptable and robust as the landscape of generative AI continues to evolve.

## 4. Conclusion

This work presents a robust and scalable framework for the detection and classification of AI-generated content across text and image domains. By integrating perceptual hashing, similarity measurement, and pseudo-labeling, our method addresses the critical challenges posed by the rapid evolution of generative AI. Experimental results on the Defactify4 dataset validate the effectiveness of our approach, which achieves competitive performance while maintaining adaptability to new generative models. Importantly, the proposed method provides a flexible solution for continual learning, accommodating the growing diversity of AI-generated content without the need for retraining. This adaptability ensures relevance in real-world applications, where the landscape of generative AI is dynamic and ever-expanding. Future work includes exploring techniques to convert high-dimensional floating-point vectors into compact binary representations while maintaining classification performance.

## References

- [1] Y. Wang, Y. Pan, M. Yan, Z. Su, T. H. Luan, A survey on ChatGPT: AI-generated contents, challenges, and solutions, *IEEE Open Journal of the Computer Society* (2023).
- [2] U. Ojha, Y. Li, Y. J. Lee, Towards universal fake image detectors that generalize across generative models, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 24480–24489.
- [3] L. Lin, N. Gupta, Y. Zhang, H. Ren, C.-H. Liu, F. Ding, X. Wang, X. Li, L. Verdoliva, S. Hu, Detecting multimedia generated by large AI models: A survey, *arXiv preprint arXiv:2402.00045* (2024).
- [4] T. Brown, B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, et al., Language models are few-shot learners, *Advances in neural information processing systems* 33 (2020) 1877–1901.
- [5] J. Achiam, S. Adler, S. Agarwal, L. Ahmad, I. Akkaya, F. L. Aleman, D. Almeida, J. Altenschmidt, S. Altman, S. Anadkat, et al., Gpt-4 technical report, *arXiv preprint arXiv:2303.08774* (2023).
- [6] A. Q. Jiang, A. Sablayrolles, A. Mensch, C. Bamford, D. S. Chaplot, D. d. l. Casas, F. Bressand, G. Lengyel, G. Lample, L. Saulnier, et al., Mistral 7b, *arXiv preprint arXiv:2310.06825* (2023).
- [7] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, B. Ommer, High-resolution image synthesis with latent diffusion models, in: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2022, pp. 10684–10695.
- [8] A. Ramesh, M. Pavlov, G. Goh, S. Gray, C. Voss, A. Radford, M. Chen, I. Sutskever, Zero-shot text-to-image generation, in: *International conference on machine learning*, Pmlr, 2021, pp. 8821–8831.
- [9] MidJourney, Midjourney: An independent research lab exploring new mediums of thought, <https://www.midjourney.com>, 2023. Accessed: 2025-01-02.
- [10] T. A. Mohamed, M. H. Khafgy, A. B. ElSedawy, A. S. Ismail, A proposed model for distinguishing between human-based and chatgpt content in scientific articles., *IEEE Access* (2024).
- [11] H. Abburi, M. Suesserman, N. Pudota, B. Veeramani, E. Bowen, S. Bhattacharya, Generative ai text classification using ensemble llm approaches, *arXiv preprint arXiv:2309.07755* (2023).
- [12] H. Khalid, S. S. Woo, OC-FakeDect: Classifying Deepfakes Using One-class Variational Autoencoder, in: *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2020, pp. 2794–2803.
- [13] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, et al., Scikit-learn: Machine learning in python, *the Journal of machine Learning research* 12 (2011) 2825–2830.
- [14] M. Lewis, Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension, *arXiv preprint arXiv:1910.13461* (2019).
- [15] Z. Liu, H. Hu, Y. Lin, Z. Yao, Z. Xie, Y. Wei, J. Ning, Y. Cao, Z. Zhang, L. Dong, et al., Swin transformer v2: Scaling up capacity and resolution, in: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2022, pp. 12009–12019.
- [16] J. Deng, J. Guo, N. Xue, S. Zafeiriou, Arcface: Additive angular margin loss for deep face recognition, in: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2019, pp. 4690–4699.
- [17] M. S. A. Toofanee, S. Dowlut, M. Hamroun, K. Tamine, V. Petit, A. K. Duong, D. Sauveron, Dfu-siam a novel diabetic foot ulcer classification with deep learning, *IEEE Access* (2023).
- [18] S. Jeon, 1st place solution to google landmark retrieval 2020, *arXiv preprint arXiv:2009.05132* (2020).
- [19] P. Cascante-Bonilla, F. Tan, Y. Qi, V. Ordonez, Curriculum labeling: Revisiting pseudo-labeling for semi-supervised learning, in: *Proceedings of the AAAI conference on artificial intelligence*, volume 35, 2021, pp. 6912–6920.
- [20] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, C. L. Zitnick, Microsoft coco: Common objects in context, in: *Computer Vision–ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part V* 13, Springer, 2014, pp. 740–755.

- [21] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, L. Fei-Fei, Imagenet: A large-scale hierarchical image database, in: 2009 IEEE conference on computer vision and pattern recognition, Ieee, 2009, pp. 248–255.
- [22] D. P. Kingma, Adam: A method for stochastic optimization, arXiv preprint arXiv:1412.6980 (2014).
- [23] F. Wang, X. Xiang, J. Cheng, A. L. Yuille, Normface: L2 hypersphere embedding for face verification, in: Proceedings of the 25th ACM international conference on Multimedia, 2017, pp. 1041–1049.
- [24] T. Wu, L. Luo, Y.-F. Li, S. Pan, T.-T. Vu, G. Haffari, Continual learning for large language models: A survey, arXiv preprint arXiv:2402.01364 (2024).
- [25] I. DeAndres-Tame, R. Tolosana, P. Melzi, R. Vera-Rodriguez, M. Kim, C. Rathgeb, X. Liu, A. Morales, J. Fierrez, J. Ortega-Garcia, et al., Frcsyn challenge at cvpr 2024: Face recognition challenge in the era of synthetic data, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2024, pp. 3173–3183.
- [26] S.-Y. Wang, O. Wang, R. Zhang, A. Owens, A. A. Efros, CNN-Generated Images Are Surprisingly Easy to Spot... for Now, in: 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2020, pp. 8692–8701.
- [27] I. C. Camacho, K. Wang, A Comprehensive Review of Deep-Learning-Based Methods for Image Forensics, *Journal of Imaging* 7 (2021) 69.
- [28] F. Marra, D. Gragnaniello, D. Cozzolino, L. Verdoliva, Detection of GAN-Generated Fake Images over Social Networks, in: 2018 IEEE Conference on Multimedia Information Processing and Retrieval (MIPR), 2018, pp. 384–389.
- [29] L. Li, Face X-Ray for More General Face Forgery Detection, 2020.
- [30] Y. Li, S. Lyu, Exposing DeepFake Videos By Detecting Face Warping Artifacts, 2018.
- [31] X. Xuan, B. Peng, W. Wang, J. Dong, On the generalization of GAN image forensics, 2019. arXiv:1902.11153.
- [32] P. Esser, S. Kulal, A. Blattmann, R. Entezari, J. Müller, H. Saini, Y. Levi, D. Lorenz, A. Sauer, F. Boesel, et al., Scaling rectified flow transformers for high-resolution image synthesis, in: Forty-first International Conference on Machine Learning, 2024.
- [33] D. Podell, Z. English, K. Lacey, A. Blattmann, T. Dockhorn, J. Müller, J. Penna, R. Rombach, Sdxl: Improving latent diffusion models for high-resolution image synthesis, arXiv preprint arXiv:2307.01952 (2023).
- [34] J. Betker, G. Goh, L. Jing, T. Brooks, J. Wang, L. Li, L. Ouyang, J. Zhuang, J. Lee, Y. Guo, et al., Improving image generation with better captions, *Computer Science*. <https://cdn.openai.com/papers/dalle-3.pdf> 2 (2023) 8.
- [35] G. Team, Gemma (2024). URL: <https://www.kaggle.com/m/3301>. doi:10.34740/KAGGLE/M/3301.
- [36] Qwen2 technical report (2024).
- [37] AI@Meta, Llama 3 model card (2024). URL: [https://github.com/meta-llama/llama3/blob/main/MODEL\\_CARD.md](https://github.com/meta-llama/llama3/blob/main/MODEL_CARD.md).
- [38] A. Young, B. Chen, C. Li, C. Huang, G. Zhang, G. Zhang, H. Li, J. Zhu, J. Chen, J. Chang, et al., Yi: Open foundation models by 01. ai, arXiv preprint arXiv:2403.04652 (2024).
- [39] Q. Cui, Q.-Y. Jiang, X.-S. Wei, W.-J. Li, O. Yoshie, Exchnet: A unified hashing network for large-scale fine-grained image retrieval, in: *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part III* 16, Springer, 2020, pp. 189–205.
- [40] Y. Shen, X. Sun, X.-S. Wei, Q.-Y. Jiang, J. Yang, Semicon: A learning-to-hash solution for large-scale fine-grained image retrieval, in: *European Conference on Computer Vision*, Springer, 2022, pp. 531–548.

## 5. Appendix

### 5.1. Related work

In this section, we review related work on three key areas. First, we discuss loss functions used for classification tasks (Sec. 5.1.1). Next, we explore advancements in text generation with large language models (LLMs), which have significantly advanced the field of natural language processing, and the detection of AI-generated text (Sec. 5.1.2). Finally, we present recent progress in image generation models and the detection of AI-generated images (Sec. 5.1.3).

#### 5.1.1. Loss functions for classification

We begin by analyzing the conventional Softmax loss, a widely used loss function in classification tasks:

$$L_{\text{Softmax}} = -\frac{1}{N} \sum_{i=1}^N \log \frac{e^{W_{y_i}^T x_i + b_{y_i}}}{\sum_{j=1}^n e^{W_j^T x_i + b_j}} = -\frac{1}{N} \sum_{i=1}^N \log \frac{e^{\|W_{y_i}\| \|x_i\| \cos(\theta_{y_i x_i}) + b_{y_i}}}{\sum_{j=1}^n e^{\|W_j\| \|x_i\| \cos(\theta_{j x_i}) + b_j}} \quad (1)$$

where  $y_i$  is the class label of the  $i^{\text{th}}$  sample,  $\theta_{j x_i}$  is the angle between the weight  $W_j^T$  and the feature  $x_i$ ,  $n$  is the total number of classes, and  $N$  is the number of samples. Models trained with the Softmax loss are limited to a fixed set of classes and require retraining whenever new labels are added. This limitation comes from the fully connected layer, which has a fixed number of nodes corresponding to the predefined classes.

One workaround for models trained with Softmax loss is to remove the final fully connected layer and use the extracted feature embeddings. However, as shown in [16], this method still has limitations. While Softmax can generate separable feature embeddings, it often leads to unclear decision boundaries, making it ineffective for robust classification without further adjustments. In contrast, the Large Margin Cosine Loss (LMCL) explicitly introduces a margin between the closest classes, leading to better-defined decision boundaries and more robust separability.

By applying  $l_2$  normalization to both the weights and features, ensuring  $\|W_j\| = \|x_i\| = 1$ , introducing a scaling factor  $s$ , and setting the bias  $b_j = 0$ , as described in [23], the original Softmax loss function (1) is reformulated as:

$$L_{\text{NormFace}} = -\frac{1}{N} \sum_{i=1}^N \log \frac{e^{s \cos \theta_{y_i}}}{\sum_{j=1}^n e^{s \cos \theta_{j i}}} = -\frac{1}{N} \sum_{i=1}^N \log \frac{e^{s \cos \theta_{y_i}}}{e^{s \cos \theta_{y_i}} + \sum_{j=1, j \neq y_i}^n e^{s \cos \theta_{j i}}}. \quad (2)$$

As mentioned earlier, the Softmax loss function suffers from limitations when it comes to handling the addition of new classes and ensuring clear decision boundaries. Since the embedding features are distributed around each class center on the hypersphere, an additive angular margin penalty,  $m$ , between the feature vector  $x_i$  and the corresponding class weight  $W_{y_i}$  can be introduced. This margin enhances both the intra-class compactness and inter-class separability, encouraging better feature discrimination. In other words, the model learns to increase the angular distance between different classes while simultaneously reducing the distance within the same class.

$$L_{\text{ArcFace}} = -\frac{1}{N} \sum_{i=1}^N \log \frac{e^{s(\cos(\theta_{y_i} + m))}}{e^{s(\cos(\theta_{y_i} + m))} + \sum_{j=1, j \neq y_i}^n e^{s \cos \theta_{j i}}} \quad (3)$$

This reformulation, known as ArcFace [16], incorporates the additive angular margin  $m$  directly into the decision boundaries, as shown in Equation (3). ArcFace has established itself as a state-of-the-art approach for classification tasks, especially in scenarios where new labels may emerge outside the training set. While newer loss functions may surpass ArcFace in specific scenarios, its widespread use

in competitions and ease of implementation have solidified its position as a benchmark choice for its combination of performance, versatility, and reliability [18].

### 5.1.2. Detection of AI-generated text

Text generation has witnessed rapid advancements with the development of large language models (LLMs) such as GPT-3 [4], GPT-4 [5], and Mistral [6]. These models are designed to generate coherent and contextually relevant text, leveraging vast pretraining datasets to produce outputs that closely resemble human-written content. Applications of these models include creative writing, summarization, and conversational agents, showcasing their ability to adapt across diverse linguistic tasks.

One notable advantage of LLMs lies in their capacity to generalize across tasks with minimal fine-tuning, often performing well with zero-shot or few-shot learning. This flexibility has enabled their widespread adoption in various fields, but it has also introduced challenges in distinguishing between human-written and AI-generated text. As models become increasingly realistic, research has focused on methods to identify synthetic text, emphasizing the importance of robust detection frameworks to mitigate misuse and ensure trust in AI-generated content [24]. Recent studies have explored ensemble methods that aggregate output probabilities from multiple backbone models trained with softmax loss, achieving competitive results in text classification tasks [10, 11]. However, while effective, these approaches exhibit limited scalability when faced with a growing number of labels, underscoring the need for alternative methods capable of adapting to the dynamic nature of generative AI.

### 5.1.3. Detection of AI-generated images

Recent years have seen remarkable progress in image generation, driven by models such as Stable Diffusion [7], DALL-E [8], Midjourney [9]. These generative models are capable of creating high-quality, photorealistic images from textual prompts, offering unprecedented control over the content and style of the generated visuals. Their applications range from digital art and design to content creation and entertainment, making them invaluable tools in creative industries.

These models operate by learning to generate images that align with textual descriptions, often leveraging latent space representations to ensure fine-grained control over image attributes. The rapid evolution of these models has enabled the creation of visually compelling and contextually accurate outputs, but it has also raised ethical and practical concerns. The ability to generate highly realistic images poses risks for misinformation and deceptive media. Consequently, the development of reliable classification systems to differentiate between real and AI-generated images has become a pressing need to ensure the ethical use of such technologies [25].

Deep learning generated fake images are typically created by generative neural network (GAN) models, but can also be created using autoencoders [26]. In most cases, the creation of fake images consists of replacing a person or a face in an existing image or video with another person or face. As videos contain more information than images, most methods apply to the detection of fake videos [27]. However some methods have been proposed for images. The first methods for fake image detection focus on the detection of images generated by a specific GAN [28]. Common CNN models detect spatial cues such as artifacts on facial boundaries [29], or traces left by the GAN [30, 31, 26]. Khalid and Woo [12] propose a one-class variational autoencoder model to detect fake images as anomalies. However, these methods cannot adapt to new AI image generation models.

## 5.2. Additional experiments

In this section, we present additional experiments to evaluate the adaptability of our method when integrating a new label without retraining the model.

Table 3 shows the performance of our method on the **Defactify4-Image** dataset. In both tests, both the Softmax loss and our proposed framework (using ArcFace loss) perform well and have relatively similar results. Our proposed method performs slightly better in `Test 2`, and both methods achieve perfect accuracy in `Test 1`. Using the pre-trained model without fine-tuning results in relatively low

**Table 3**

Performance results on the **Defactify4-Image** dataset for Task 1 (AI-generated vs. Human-produced) and Task 2 (Classification of different methods). a11 indicates using the entire dataset with the proposed method (using ArcFace loss), while a11-x refers to the dataset excluding samples from the class labeled as x. And swinv2's training data refers to using a pre-trained Swin Transformer V2 [15] model to extract features from images without additional training.

Testing data		Training data									
		swinv2's training data	Defactify dataset								
			all-dalle-sdxl	all-midjourney	all-dalle	all-real	all-sdxl	all-sd21	all-sd3	Softmax	all
Test 1	Task A	0.98	0.98	1.00	1.00	0.98	0.99	1.00	0.99	1.00	1.00
	Task B	0.90	0.92	1.00	0.99	0.94	0.98	0.99	0.96	1.00	1.00
Test 2	Task A	0.7651	0.7817	0.8295	0.8306	0.8034	0.8278	0.8293	0.8250	0.8303	0.8330
	Task B	0.3170	0.3598	0.4506	0.4766	0.3920	0.4672	0.4637	0.4383	0.4870	0.4935

performance. This is because the model was trained on a broad dataset and does not perform well on a fine-grained dataset for specific tasks. Additionally, AI-generated images may have a different distribution compared to the images the model was trained on, causing difficulties in performance. When removing two labels from the data, the model experiences a significant drop in performance, though it still outperforms the pre-trained Swin Transformer V2 Base model. When removing one label for training, the model still maintains acceptable performance, except when excluding the "real" label. The reason for this is that the gap between AI-generated images and human-produced images is relatively large, so training exclusively on AI-generated images causes the model to become unfamiliar with real image data. However, when removing one AI-generated label, the model performs slightly lower but still remains highly competitive compared to the Softmax method. This demonstrates the potential for future expansion of the number of labels in our proposed method.

**Table 4**

Performance results on the **Defactify4-Text** dataset for Task 1 (AI-generated vs. Human-produced) and Task 2 (Classification of different methods). a11 indicates using the entire dataset with the proposed method (using ArcFace loss), while a11-x refers to the dataset excluding samples from the class labeled as x. And BART's training data refers to using a pre-trained BART [14] model to extract features from input text without additional training.

Testing data		Training data									
		BART's training data	Defactify dataset								
			all-yi-llama	all-Human	all-gemma2	all-mistral	all-qwen2	all-llama	all-yi	all-gpt4o	Softmax
Test 1	Task A	0.98	0.99	0.98	1.00	1.00	0.99	1.00	1.00	1.00	1.00
	Task B	0.94	0.95	0.95	0.96	0.96	0.95	0.95	0.96	0.95	0.96
Test 2	Task A	0.9970	0.9978	0.9950	0.9980	0.9980	0.9970	0.9978	0.9971	0.9947	0.9963
	Task B	0.7946	0.8190	0.8350	0.8738	0.8692	0.8206	0.8462	0.8433	0.8268	0.9049

Table 4 shows the performance of our method on the **Defactify4-Text** dataset. In both tests, similar to the image task the Softmax method and our proposed method perform well and have relatively similar results. Our proposed method performs slightly better in Test 2. Using the pre-trained model without fine-tuning results in relatively low performance. When removing one label for training, the model still maintains acceptable performance. However, compared to the image dataset, the performance drops when using only the pre-trained BART Large model or excluding the "human" label is less noticeable here. This suggests that AI-generated text has a distribution that is not drastically different from the original wide dataset, and human-generated text does not deviate much from the other labels. This indicates that text generation models are performing better than image generation models, as the distribution of outputs between human and AI-generated text is less distinct. However, more data is needed to fully validate this observation. When removing two labels from the data, the model experiences a significant drop in performance, though it still outperforms the pre-trained model.

In both Table 3 and Table 4, when removing one AI-generated label the performance of the model decreases in both tasks. However, the model performs slightly lower but still remains highly competitive compared to the Softmax method. This demonstrates the potential for future expansion of the number of labels in our proposed method.

### 5.3. Dataset detail

The experiments are conducted on two benchmark datasets: **Defactify4-Image** (Sec. 5.3.1) and **Defactify4-Text** (Sec. 5.3.2). Each dataset includes training and testing splits tailored for two tasks: AI vs. human classification and method-specific categorization. The datasets are carefully designed with a mix of real and AI-generated data, incorporating augmentations to test the model’s robustness and generalization capabilities. Below, we provide detailed descriptions of each dataset.

#### 5.3.1. Dataset of AI-generated images

The **Defactify4-Image** dataset is a benchmark designed to evaluate the ability to distinguish between real and AI-generated images. It comprises seven data categories (6 classes and captions), including captions and various image sources. Among these, real images selected from the COCO dataset [20] are represented by the `coco_image` class. While the other five categories (`sd3_image`, `sd21_image`, `sdxl_image`, `dalle_image`, `midjourney_image`) are generated using specific AI models: Stable Diffusion (v3 [32], v2.1 [7], XL [33]), DALL-E [34], and MidJourney [9], respectively. Captions act as the input prompts for these generative models and correspond to the caption of real images in the dataset.

**Training Data** The training set consists of 42,000 images across six classes, with 7,000 samples per class. Each class corresponds to one of the five generative models and the real image class (`coco_image`). All images share the same caption within their index, for instance, `sd3_image[i]`, `sd21_image[i]` are generated from the same `caption[i]` of `coco_image[i]`. Where `record[i]` refers to the  $i^{\text{th}}$  sample of the *record* in the dataset.

**Testing Data** The dataset provides two distinct test sets to assess model performance. `Test 1` comprises 9,000 images, each paired with its original caption, representing unaltered outputs from generative models or real images. While `Test 2` consists of 45,000 images where augmentation techniques have been applied, enabling evaluation of the model’s robustness and generalization across various transformations.

#### 5.3.2. Dataset of AI-generated text

The **Defactify4-Text** dataset is a benchmark designed to evaluate the ability to distinguish between human-written and AI-generated text. It comprises eight columns: `prompt`, `Human_story`, `gemma-2-9b`, `mistral-7B`, `qwen-2-72B`, `llama-8B`, `yi-large`, and `GPT_4-o`. The `prompt` column contains the instruction, while `Human_story` is a human-written text corresponding to the prompt. The remaining columns represent outputs from various generative models, including `gemma-2-9b` [35], `mistral-7B` [6], `qwen-2-72B` [36], `llama-8B` [37], `yi-large` [38], and `GPT_4-o` [5], based on the provided prompt.

**Training Data** The training set consists of 51,248 text samples across seven classes, with 7,321 samples per class. Each class corresponds to one of the generative models or the human-written text class (`Human_story`). Similar to **Defactify4-Image** all samples share the same prompt within their index.

**Testing Data** For evaluation, the dataset provides two separate test sets. `Test 1` includes 10,983 samples, each associated with its original prompt and the corresponding generated or human-written text. These samples represent the raw outputs from the generative models or human authors. On the other hand, `Test 2` contains 10,963 samples where various augmentation techniques have been applied to the text data, allowing for the assessment of model robustness and generalization under various transformations.

## 5.4. Discussion

The proposed method demonstrates strong scalability, particularly in its ability to adapt to new labels and modalities. However, the approach has certain limitations, primarily the need to store feature representations. Specifically, the extracted features consist of high-dimensional floating-point vectors, which can impose significant memory requirements as the dataset size increases. For instance, in our implementation, each sample in the training set requires  $32 \times 512 = 16,384$  bits to store, given that each floating-point number occupies 32 bits, and 512 is the length of the vector. This storage demand becomes challenging with large-scale datasets.

Recent studies have highlighted strategies to address this issue by converting high-dimensional floating-point vectors into shorter binary vectors [39, 40]. These approaches significantly reduce storage requirements but often come with a trade-off in terms of performance, making them less suitable for competitive tasks such as those in this challenge. Exploring methods to balance memory efficiency and accuracy remains an important area for future research.