



RU-AI: A Large Multimodal Dataset for Machine-Generated Content Detection

Liting Huang*
liting.huang@student.uts.edu.au
University of Technology Sydney
Sydney, New South Wales, Australia

Zhihao Zhang*
zhihao.zhang11@unsw.edu.au
University of New South Wales
Sydney, New South Wales, Australia

Yiran Zhang
yiran.zhang1@students.mq.edu.au
Macquarie University
Sydney, New South Wales, Australia

Xiyue Zhou
xzho9805@uni.sydney.edu.au
The University of Sydney
Sydney, New South Wales, Australia

Shoujin Wang
shoujin.wang@uts.edu.au
University of Technology Sydney
Sydney, New South Wales, Australia

Abstract

The recent generative AI models' capability of creating realistic and human-like content is significantly transforming the ways in which people communicate, create and work. The machine-generated content is a double-edged sword. On one hand, it can benefit the society when used appropriately. On the other hand, it may mislead people, posing threats to the society, especially when mixed together with natural content created by humans. Hence, there is an urgent need to develop effective methods to detect machine-generated content. However, the lack of aligned multimodal datasets inhibited the development of such methods, particularly in triple-modality settings (e.g., text, image, and voice). In this paper, we introduce RU-AI, a new large-scale multimodal dataset for robust and effective detection of machine-generated content in text, image and voice. Our dataset is constructed on the basis of three large publicly available datasets: Flickr8K, COCO and Places205, by adding their corresponding AI duplicates, resulting in a total of 1,475,370 instances. In addition, we created an additional noise variant of the dataset for testing the robustness of detection models. We conducted extensive experiments with the current SOTA detection methods on our dataset. The results reveal that existing models still struggle to achieve accurate and robust detection on our dataset. We hope that this new data set can promote research in the field of machine-generated content detection, fostering the responsible use of generative AI. The source code and datasets are available at <https://github.com/ZhihaoZhang97/RU-AI>.

CCS Concepts

• **Information systems** → **Multimedia content creation**.

Keywords

Multimodal datasets, Machine generation, Content detection

*Both authors contributed equally to this research.



This work is licensed under a Creative Commons Attribution 4.0 International License (CC BY 4.0).

WWW Companion '25, Sydney, NSW, Australia.

© 2025 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 979-8-4007-1331-6/25/04
<https://doi.org/10.1145/3701716.3715306>

ACM Reference Format:

Liting Huang, Zhihao Zhang, Yiran Zhang, Xiyue Zhou, and Shoujin Wang. 2025. RU-AI: A Large Multimodal Dataset for Machine-Generated Content Detection. In *Companion of the 16th ACM/SPEC International Conference on Performance Engineering (WWW Companion '25)*, April 28-May 2, 2025, Sydney, NSW, Australia. ACM, New York, NY, USA, 4 pages. <https://doi.org/10.1145/3701716.3715306>

1 Introduction

The recent advent of large generative AI models has significantly impacted various sectors of human society, including industry, research, and education. Deploying generative AI models appropriately can save professionals' time for substantive matters, assist scientists in discovering new medicine, and provide students with customised feedback. However, generative AI models' capabilities may be misused for generating discriminate or disrespectful content, spreading misinformation, and conducting target scams. These challenges indicate the pressing need for robust mechanisms and models to verify data originality, particularly for multimodal content. Previous studies [8] have explored manual methods and automated approaches for detecting machine-generated content. Researchers [1] have found that humans can identify machine-generated content, but the quality of results is largely influenced by an individual's background, experience and preference. Additionally, manually identifying machine-generated content is labour-intensive and difficult to scale. Therefore, it is crucial to explore automated approaches to detect machine-generated content across different modalities. However, the lack of aligned multimodal datasets constrains the development of such methods. The majority of existing datasets for machine-generated content detection [26] are constructed within a single modality, thereby neglecting valuable cross-modal information that could enhance the detection of machine-generated content. Although there exists several multimodal datasets for machine-generated content detection [12], they primarily focus on human faces and voices, which significantly limits their scope of applicability. Only one dataset [20] includes machine-generated images and text, but lacks synthesised voice. Combining voice and images with text helps to reduce ambiguity, while integrating all three modalities can minimise information loss. This approach can be effectively applied to combat fraud and disinformation.

To address these problems, we introduce RU-AI, a large-scale multimodal dataset for robust detection of machine-generated content in text, image and voice. The real portion of our multimodal




Figure 1: Data collection process for all datasets. It shows the text collection (left), the image collection (middle), and the voice collection (right). Text is generated through rephrasing caption text. Images are created by referencing both caption text and real images. Voices are synthesized using caption text. All collected data is then matched with their original data pairs.

dataset are collected from three large publicly available image caption datasets: COCO [16], Flickr8K [23] and Places205 [24] with their human voices. Machine-generated data is produced from generative AI models, including five models for text, image, and voice generation. This yields total of 1,475,370 instances each for both real and machine-generated data. We also construct a noise-augmented variant of our dataset across all modalities to evaluate model robustness. Furthermore, to demonstrate the effectiveness of our dataset, we develop and train baseline detection models on RU-AI and conduct extensive experiments to assess their performance on our dataset. Our **key contributions** can be summarised as follows:

- We introduce the first large-scale three-modal (text, image, voice) dataset, comprising 245,895 human-generated instances and 1,229,475 corresponding machine-generated instances. These machine-generated instances are produced using five current state-of-the-art (SOTA) generative models.
- We introduce a variant of the proposed dataset specifically designed for model robustness testing. This version incorporates additional noise across all modalities, with three distinct types of noise applied to each modality.
- We develop baseline detection models and evaluate their performance across all modalities. Our experimental results and analysis reveal that current SOTA triple-modality models struggle to accurately verify content originality and lack robustness against noise.

2 RU-AI Dataset

We introduce RU-AI, a new large-scale dataset that features multi-modal real and human-generated content, accompanied by machine-generated content. It encompasses three modalities: real-world images, image-describing caption text, and spoken voices for captions, providing more than 245k real and machine-generated aligned instance pairs for each modality. As shown in Figure 1, image and text instances are collected from COCO [16], Flickr8K [23] and Places205 [24], while voice instances are sourced from SLS Corpora¹, which form the real and human-generated part of our dataset. The machine-generated portion is produced by five SOTA generative models specific to each modality.

2.1 Data Construction

Text. Since COCO and Flickr8K incorporate five captions for each image, we extract the longest caption for each image to ensure the most comprehensive description. For Places205, which includes

a single caption for each image, we sort the images’ captions in descending order and extract the first 12,000 image-caption pairs. It is worth noting that Places205 has a large number of captions without punctuation, while only a small number of non-punctuated captions are present in Flickr8K and COCO. To maintain consistency across datasets, we use FullStop [7] to add punctuation for all non-punctuated captions. The text data is then equally distributed among five Large Language Models (LLMs) for rephrasing. We select both open-source and proprietary LLMs based on different architectures. The rephrased generations are stored with the original captions in a structured data format identified by unique IDs.

Image. To ensure the quality and consistency of machine-generated images across all three datasets, we exclude images under 224 pixels in width or height and resize those over 640 pixels to a maximum of 640 pixels while maintaining their original aspect ratios. We apply five diffusion models with different parameter sizes and fine-tuning techniques to generate images, each contributing equally to the overall dataset. Image captions, positive-negative prompt pairs, and original images are used as seeds for these models during inference to ensure that the semantic meaning and realistic scenes are captured in image generation. The generated images maintain the same aspect ratio as the original images and are named after the original image IDs with the model names attached.

Voice. The voices for captions are cloud-sourced from a group of individuals, with each individual being assigned random captions for inference; therefore, we merge the spoken caption voices of each individual, identified by their unique user ID, into one-minute long voice files. For YourTTS [5], the machine-generated voices are synthesised by using the captions along with their corresponding one-minute human voices as references. For other models, the voices are synthesised directly from the caption text using default settings. All five models contribute equally to the generated voice data, and these voices are named and saved similarly to the images, with the original IDs and model names attached.

Data Augmentation. Based on the collected multimodal data, a noise-augmented variant is created towards robustness evaluation. For each modality, we introduce three different types of noise, with each type applied to 20% of the data. As a result, the final dataset consists of 60% for noise data and 40% clean data. For text, the noises added at the word level are: 1) randomly removing a single character from words; 2) randomly shuffling two characters within words; 3) randomly replacing a single character in words. These alterations are applied to 20% of the words in a sentence, and the augmented sentences adhere to the specified noise distribution. For

¹<https://sls.csail.mit.edu/>

Table 1: The number of machine-generated data from different datasets in each modality of RU-AI.

Modality	Model	Dataset		
		Flickr8K	COCO	Place
Text	ChatGPT ²	1,610	24,272	23,736
	Llama-2-13B [21]	1,518	23,121	23,488
	RWKV-5-7B [14]	1,628	23,898	23,780
	Mixtral-8x7B [9]	1,599	23,752	23,858
	PaLM-2 [2]	1,535	24,221	23,879
Image	StableDiffusion-1.5 [19]	1,610	24,272	23,736
	StableDiffusion-1.5-Hyper [18]	1,518	23,121	23,488
	AbsoluteReality-1.8 ³	1,628	23,898	23,780
	epiCRealism-VAE ⁴	1,599	23,752	23,858
	StableDiffusion-XL [17]	1,535	24,221	23,879
Voice	YourTTS [5]	1,610	24,272	23,736
	XTTS-2[4]	1,518	23,121	23,488
	EfficientSpeech [3]	1,628	23,898	23,780
	VITS [13]	1,599	23,752	23,858
	StyleTTS-2 [15]	1,535	24,221	23,879
Total / modality		7,890	119,264	118,741

images, three types of noise are added to randomly selected images through scikit-image [22], following the specified noise distribution: 1) Gaussian noise; 2) Salt and Pepper noise; 3) Poisson noise. For voices, three different noises are added with a 15dB Signal-to-Noise Ratio through Audiomentations [11], following the specified noise distribution: 1) Gaussian White noise; 2) Cafe noise; 3) Rain noise.

2.2 Data Processing

We employ a regular expression method to remove unrelated LLM-generation from the output to maintain the quality of the text data. Additionally, as generative models may produce empty sentences or blank images due to safety and alignment-related censorship, we remove these to ensure consistency across our dataset. To ensure cross-modality alignment, we perform a data-matching process that retains the minimal overlapping subset across all three modalities. Therefore a matched data instance is formed by aligning the image, text, and voice from different modality models corresponding to a specific ID. We implement this approach in our collected dataset, resulting in 245,895 aligned real-and-AI data pairs. As shown in Table 1, the final processed data for each modality includes 7,890, 119,264 and 118,741 pairs from Flickr8K, COCO, and Places205.

2.3 Data Characteristics

Our dataset includes 245,895 pairs of real and machine-generated data from text, image, and voice modalities. Table 2 shows the average length of captions from different LLMs and their average similarity between the reference and generated text, indicated by average ROUGE-L and BERTScore. Table 3 shows the synthesised voice characteristics from different models, including average voice duration, average amplitude, and average spectral flatness. The average amplitude range refers to the average normalised audio scale, while the average spectral flatness is a measure of the sound’s noise-like level. Since all images in Flickr8K are unlabelled with object instances, we employ pre-trained YOLOv5 [10] to annotate the unlabelled images. The data characteristics are based on the

²<https://openai.com/index/chatgpt/>

³<https://huggingface.co/Lykon/AbsoluteReality>

⁴<https://huggingface.co/emilianJR/epiCRealism>

Table 2: Statistics of generated text data. "Avg. Len." is the average number of words per sentence, "Avg. RL" is the average ROUGE-L score, and "Avg. BS" is the average BERT Score. Reference is the original text.

Text Source	Avg. Len.	Avg. RL	Avg. BS
ChatGPT	22.97	0.3698	0.9167
Llama-2-13B	26.96	0.2795	0.8852
RWKV-5-7B	21.07	0.6411	0.9464
Mixtral-8x7B	19.28	0.5662	0.9411
PaLM-2	17.23	0.4615	0.9278
Reference	21.66	N/A	N/A

Table 3: Statistics of synthesised voice data. "Avg. Len." is the average duration of the voices in seconds, "Avg. AR" is the average amplitude ranges, and "Avg. BS" is the average spectral flatness. Reference is the original voices.

Voice Source	Avg. Len.	Avg. AR	Avg. SF
YourTTS	7.338	1.830	0.1264
XTTS-2	7.842	1.826	0.1140
EfficientSpeech	5.980	1.416	0.02997
VITS	7.536	1.831	0.09121
StyleTTS-2	7.563	1.001	0.03474
Reference	8.700	1.423	0.03605

original annotations and the YOLOv5 annotations. Our dataset can be divided into 20 different categories, with "Person" being the largest category, and there are 1,040 images that do not belong to any category and are classified as "Other".

3 Proposed Baseline

Based on the RU-AI dataset, we propose a unified classification model capable of determining the origin of input data from different modalities. Our model utilises a multimodal embedding model to project the features from different modalities into a common space through high-dimensional embedding sequences. These embedding sequences are then fed into a multilayer perceptron network to distinguish the origin of the input data.

3.1 Implementation Details

We use two pre-trained SOTA multimodal encoder models with their large settings as the baseline multimodal embedding models: ImageBind [6] and LanguageBind [25]. ImageBind encodes all input data into 1,024-dimensional embedding sequences, while LanguageBind encodes them into 768-dimensional embedding sequences. Each multimodal embedding model is connected to a separate multilayer perceptron network, consisting of a single 256-dimensional hidden layer. The input layers of these networks match the dimensions of their respective embedding models, 1,024-dimensional for ImageBind and 768-dimensional for LanguageBind. These networks are used for the classification of the input data as either real or machine-generated one. Specifically, the output is calculated from probability vectors $\vec{z} = (z_r, z_m)$ using the softmax function $\sigma(\vec{z})_i$. The instance is classified as real if $z_r > z_m$; otherwise, it is classified as machine-generated. To analyse the effectiveness of different modalities in our data, we train and evaluate our classification model with different combinations of the data. The training

Table 4: Results of the proposed classification model with ablation studies. "LB" is LanguageBind, "IB" is ImageBind, "A" is Accuracy, "P" is Precision, "ALL" is all three modalities.

Model	Modality	Original			Noise Augmentation		
		A@ALL	P@ALL	F1@ALL	A@ALL	P@ALL	F1@ALL
LB + MLP	Text	69.76	69.78	69.75	65.18	62.46	68.70
	Image	76.49	85.95	72.93	62.98	86.95	45.32
	Voice	68.89	64.09	73.43	68.36	63.66	73.07
	Text+Image	77.09	76.91	77.16	64.63	71.76	57.82
	Text+Voice	77.69	81.56	76.23	74.84	77.65	73.54
	Image+Voice	82.26	78.89	83.23	78.20	91.34	74.12
	All	84.20	84.20	82.67	80.02	78.02	80.74
IB + MLP	Text	60.70	57.36	67.96	64.21	61.42	68.21
	Image	67.75	71.03	65.03	65.54	63.51	68.04
	Voice	65.77	63.61	68.28	66.93	63.39	70.87
	Text+Image	71.48	75.18	69.23	68.58	72.41	65.73
	Text+Voice	75.39	69.05	78.90	73.05	67.14	77.06
	Image+Voice	81.25	82.84	80.78	79.54	75.99	80.88
	All	80.63	79.19	81.10	76.42	72.97	78.11

samples are randomly selected from 80% of each data combination, while the remaining 20% of the samples are used as testing sets. We maintain the same size of data throughout the experiment and distribute the modalities equally. The performance is evaluated using Accuracy, Precision, and F1 metrics on both original and noise-augmented data. All models were trained for 5 epochs, using the Adam optimizer with cross-entropy loss and a learning rate of 1e-4, with the embedding model frozen. We follow the embedding models' original feature extraction techniques in all three modalities for consistency across datasets.

3.2 Results

The classification results from our proposed model are presented in Table 4. To evaluate the effectiveness of each modality, we conduct ablation studies using single and mixed modalities during training. In general, baseline models that utilise LanguageBind as the embedding model outperform those incorporating ImageBind. The best F1 score achieved by LanguageBind multilayer perception (MLP) is 84.20, while the best F1 score from ImageBind MLP is 81.10. This may be attributed to LanguageBind's enhanced cross-modality architecture, which binds all modalities through language, leveraging its rich semantics to provide more comprehensive information despite a smaller embedding size. For both baseline models, training with multi-modality data generally yields better performance compared to single-modality data. Since training with mixed-modality data introduces greater complexity compared to single-modality data, which could lead to better model convergence. We also observe that image and voice data tend to contribute more significantly to performance improvements. This may be caused by the richer features inherent in image and voice data compared to text. The performance trends are similar between the original data and the noise-augmented data across modalities. However, all baseline models exhibit decreased performance when evaluated on noise-augmented data, with an average drop of 2-4%. This decrease could be attributed to the added noise disrupting data features from both human and machine sources, which can subsequently lead to a degradation in model performance during evaluation.

Our baseline results from SOTA multimodal pre-trained models indicate that machine-generated content detection remains a challenging task, particularly for noise-augmented data. To mitigate the

potential misuse of generative models, future work should focus on developing more robust and reliable detection methods.

4 Conclusion

We propose RU-AI, a comprehensive large-scale multimodal dataset designed for machine-generated content detection. Unlike existing datasets, which are limited to single or dual modalities, RU-AI aligns text, image, and voice data with their corresponding machine-generated counterparts. To ensure data diversity, we integrate three large open-source datasets and employ five generative models for each modality, both with and without data augmentation. Experimental results from SOTA baseline models show that unified classification models struggle to effectively detect machine-generated content on our dataset, especially when augmented with noise. This highlights the persistent challenges in the era of generative AI. By making the RU-AI dataset publicly available, we hope to foster further research in this dynamic and rapidly evolving area.

References

- [1] Ehab A. AlBadawy, Siwei Lyu, and Hany Farid. 2019. Detecting AI-Synthesized Speech Using Bispectral Analysis. In *CVPR Workshops*. 104–109.
- [2] Rohan Anil and et al. 2023. PaLM2 Technical Report. *CoRR* abs/2305.10403 (2023).
- [3] Rowel Atienza. 2023. EfficientSpeech: An On-Device Text to Speech Model. In *ICASSP*. 1–5.
- [4] Edresson Casanova, Kelly Davis, and et al. 2024. XTTS: a Massively Multilingual Zero-Shot Text-to-Speech Model. In *INTERSPEECH*. 4978–4982.
- [5] Edresson Casanova and et al. 2022. YourTTS: Towards Zero-Shot Multi-Speaker TTS and Zero-Shot Voice Conversion for Everyone. In *ICML*. 2709–2720.
- [6] Rohit Girdhar, Alaaeldin El-Nouby, and et al. 2023. ImageBind One Embedding Space to Bind Them All. In *CVPR*. 15180–15190.
- [7] Oliver Guhr, Anne-Kathrin Schumann, and et al. 2021. FullStop: Multilingual Deep Models for Punctuation Prediction. In *SwissText*.
- [8] Ganesh Jawahar, Muhammad Abdul-Mageed, and et al. 2020. Automatic Detection of Machine Generated Text: A Critical Survey. In *COLING*. 2296–2309.
- [9] Albert Q. Jiang and et al. 2024. Mixtral of Experts. *CoRR* abs/2401.04088 (2024).
- [10] Glenn Jocher and et al. 2021. *ultralytics/yolov5: v6.0 - YOLOv5n 'Nano' models, Roboflow integration, TensorFlow export, OpenCV DNN support*.
- [11] Iver Jordal, Araik Tamazian, and et al. 2024. *iver56/audiomentations: v0.38.0*.
- [12] Hasam Khalid, Shahroz Tariq, and et al. 2021. FakeAVCeleb: A Novel Audio-Video Multimodal Deepfake Dataset. In *NeurIPS Datasets and Benchmarks*.
- [13] Jaehyeon Kim, Jungil Kong, and et al. 2021. Conditional Variational Autoencoder with Adversarial Learning for End-to-End Text-to-Speech. In *ICML*. 5530–5540.
- [14] Yuhui Li, Fangyun Wei, and et al. 2024. EAGLE: Speculative Sampling Requires Rethinking Feature Uncertainty. In *ICML*. 28935–28948.
- [15] Yinghao Aaron Li, Cong Han, and et al. 2023. StyleTTS 2: Towards Human-Level Text-to-Speech through Style Diffusion and Adversarial Training with Large Speech Language Models. In *NeurIPS*, Vol. 36. 19594–19621.
- [16] Tsung-Yi Lin, Michael Maire, Serge J. Belongie, and et al. 2014. Microsoft COCO: Common Objects in Context. In *ECCV* (5). 740–755.
- [17] Dustin Podell, Zion English, and et al. 2024. SDXL: Improving Latent Diffusion Models for High-Resolution Image Synthesis. In *ICLR*.
- [18] Yuxi Ren, Xin Xia, and et al. 2024. Hyper-SD: Trajectory Segmented Consistency Model for Efficient Image Synthesis. *CoRR* abs/2404.13686 (2024).
- [19] Robin Rombach, Andreas Blattmann, and et al. 2022. High-Resolution Image Synthesis with Latent Diffusion Models. In *CVPR*. 10674–10685.
- [20] Rui Shao, Tianxing Wu, and Ziwei Liu. 2023. Detecting and Grounding Multimodal Media Manipulation. In *CVPR*. 6904–6913.
- [21] Hugo Touvron, Louis Martin, and et al. 2023. Llama 2: Open Foundation and Fine-Tuned Chat Models. *CoRR* abs/2307.09288 (2023).
- [22] Stéfan van der Walt, Johannes L. Schönberger, and et al. 2014. scikit-image: image processing in Python. *PeerJ* 2 (2014), e453.
- [23] Peter Young, Alice Lai, and et al. 2014. From image descriptions to visual denotations: New similarity metrics for semantic inference over event descriptions. *Trans. Assoc. Comput. Linguistics* 2 (2014), 67–78.
- [24] Bolei Zhou, Àgata Lapedriza, and et al. 2014. Learning Deep Features for Scene Recognition using Places Database. In *NIPS*. 487–495.
- [25] Bin Zhu, Bin Lin, and et al. 2024. LanguageBind: Extending Video-Language Pretraining to N-modality by Language-based Semantic Alignment. In *ICLR*.
- [26] Mingjian Zhu, Hanting Chen, and et al. 2023. GenImage: A Million-Scale Benchmark for Detecting AI-Generated Image. In *NeurIPS*. 77771–77782.