Synthetic Realities and Artificial Intelligence-generated Contents

Daniel Moreira, Sébastien Marcel, and Anderson Rocha

May 2024

Welcome to the IEEE Security & Privacy Special Issue on Synthetic Realities and Artificial Intelligence-generated Contents! In this edition, we delve into the topic of Synthetic Realities, where Generative Artificial Intelligence (GAI) is revolutionizing the construction of narratives, blurring the boundaries between fact and fiction, for the good and the bad. Indeed, content created or enabled by GAI spans a wide spectrum of usage and intentions, from fostering positive experiences, such as entertainment, training, and education, to more questionable utilization, such as deception, propaganda, and manipulation.

With the advent and maturity of GAI techniques, much has changed in Forensics, Security, and Privacy. The way researchers and experts have been doing forensics and security over the past decades is continuously challenged with each new version of powerful AI content generators. The synthetic content ranges from audio, image, and video to text and their combinations, coming from prominent models such as ChatGPT, LaMDA, ImageGen, StableDiffusion, Sora, and Gemini, among others.

This Special Issue seeks to understand the required changes in the way Forensics, Security, and Privacy experts operate, including how to deal with autogenerated fake and synthetic data (e.g., text, images, videos, and 3D content), how much autogeneration methods are "shaping" new realities that do not exist, and what it means for our society. The call presented the following important questions: What are the possible new applications for Forensics, Security, and Privacy? What are the threats and challenges? Forensic aspects should include any topics related to post hoc investigation practices after the occurrence of events regarding created content (e.g., generated fake news or deepfakes and how to detect them). Security aspects should include topics related to how such

contents might affect our lives in terms of document authenticity and deception. Privacy should be related to aspects of how GAI methods deal with our data and their adequacy for the data protection frameworks in different regions of the globe.

Regardless of the intent, as the capabilities of GAI techniques continue to advance, experts from the Forensics, Security, and Privacy disciplines face an ever-evolving landscape. With each iteration of GAI-driven content creation becoming increasingly sophisticated and readily available to the public, the daunting spread of synthetic digital material that is indistinguishable from authentic material, whether audio, image, video, or text, demands swift adaptation and innovation.

This Special Issue serves as a platform for authors and their contributions to explore the multifaceted challenges and opportunities of Synthetic Realities within Forensics, Security, and Privacy. It, thus, examines the burgeoning environment of GAI, detailing its impact, technological advancements, and ethical quandaries. Synthetic Realities provide innovative solutions and opportunities for immersive experiences across various sectors, including education, health care, and commerce. However, these advancements also usher in substantial challenges.

The herein-presented works probe the evolving landscape, revealing shifts in the experts' methodologies and strategies for detecting and mitigating auto-generated content. Moreover, the reported efforts are not just about addressing challenges; they are also about harnessing the potential of Synthetic Realities to drive smarter solutions and analyzing the societal implications of GAI-enabled realities.

Seven high-quality papers were accepted for publication, conveying a 28% acceptance rate. This was the outcome of a rigorous peer-reviewed selection of manuscripts from a diverse pool of 25 worldwide submissions. The accepted authors and their institutions represent a wide variety of nationalities, attesting to the endeavor's international aspect.

Among the published material, Sun et al. [1] contribute "Unleashing Malware Analysis and Understanding with Generative AI", where they employ GAI techniques to dissect technical logs of detected malware behavior and generate insightful reports to readily inform human technicians and help them plan mitigation more quickly.

In another noteworthy contribution, Pastor-Gallindo et al. [2] present "Large-Language-Model-Powered Agent-based Framework for Misinformation and Dis-

information Research: Opportunities and Open Challenges". The authors propose a research framework to guide the generation of agent-based social networks for the study and simulation of misinformation and disinformation spread. Additionally, they elucidate the open challenges within this critical domain.

Tariang et al. [3], in turn, offer valuable insights in "Synthetic Image Verification in the Era of Generative Artificial Intelligence: What Works and What Isn't There Yet". This work provides a comprehensive overview of approaches for detecting and attributing the source of synthetic images. The authors shed light on the evolving topic of synthetic image analysis by critically examining the approaches' strengths, weaknesses, and directions for future development.

Next, Yavuz [4] discusses the problem of deepfakes and how they have rapidly developed, been misused, and been democratized in a handful of years. "A Multidisciplinary Look at History and Future of Deepfake with Gartner Hype Cycle" introduces a multidisciplinary study (technical, legal, and societal) and speculates on the future of deepfakes using the Gartner Hype Cycle.

De Cristofaro [5] contributes "Synthetic Data: Methods, Use Cases, and Risks" and provides a gentle introduction to synthetic data, discussing their use cases, the privacy challenges that are still unaddressed, and their inherent limitations as an effective privacy-enhancing technology.

In another work about the deepfake problem, "Deepfake Detection in Super-Recognizers and Police Officers", Meike Ramon et al. [6] explore the relationship between deepfake detection performance (DDP) and face identity processing (FIP) skills, comparing Super-Recognizers (SR) and non-SR police officers. Using videos from the Deepfake Detection Challenge, they find no significant links between DDP and FIP, leading to a very interesting finding.

Last but not least, closing this Special Issue, Maiano et al. [7] present "Human Versus Machine: A Comparative Analysis in Detecting Artificial Intelligence-generated Images", in which they discuss how humans compare to machines in the hard task of detecting Synthetic Realities. They highlight the challenges to automated detectors and reveal human detection biases, strengths, and weaknesses. Moreover, they introduce a dataset of generated human faces and compared the performance of automatic detectors with humans, making their article a must-read.

These contributions represent a significant fraction of the diverse issues related to Synthetic Realities. Each article contributes uniquely to improving our understanding of the implications of Synthetic Realities to Forensics, Security, and Privacy, adding greatly to the discourse and paving the way for future

advancements in the field.

As a significant outcome, this Special Issue shows the dual-edged nature of Synthetic Realities and advocates for interdisciplinary research, informed public discourse, and collaborative efforts to harness their benefits while mitigating risks. It also contributes to the discourse on the responsible development and application of AI and synthetic media in modern society.

In addition to advancing research in the detection of Synthetic Realities, a concerted effort is required from academia, governments, industry stakeholders developing media synthesizers, and nongovernmental organizations. Together, they must raise awareness among the general population about the existence and increasing dissemination of such technologies. Knowledgeable individuals can better discern falsifications and mitigate associated risks, thus fostering a more mature and aware level of credulity in digital media. Collaborative initiatives promoting media technological literacy and critical thinking skills are essential in navigating the complexities of Synthetic Realities and safeguarding against their potential misuse. Through collective action and informed engagement, societies can effectively harness emerging technologies' benefits while protecting against unintended consequences.

In closing, we express our heartfelt gratitude to all of the authors who contributed their insightful research to this Special Issue. We also want to thank the reviewers who dedicated their time and expertise to carefully evaluating the submitted manuscripts. Their invaluable feedback and constructive criticism played a crucial role in ensuring the quality and relevance of the selected articles.

Lastly, we thank the readership for their interest in this edition. We hope the insights and perspectives shared within these pages will inspire further exploration and dialogue on Synthetic Realities applied to Forensics, Security, and Privacy.

References

[1] Y S Sun, Z-K Chen, Y-T Huang, and M C Chen. Unleashing malware analysis and understanding with generative AI. *IEEE Security & Privacy*, 22(3), 2024.

[2] J Pastor-Galindo, P Nespoli, and J A Ruipérez-Valiente. Large-language-model-powered agent-based framework for misinformation and disinforma-

tion research: Opportunities and open challenges. *IEEE Security & Privacy*, 22(3), 2024.

[3] D Tariang, R Corvi, D Cozzolino, G Poggi, K Nagano, and L Verdoliva. Synthetic image verification in the era of generative artificial intelligence: What works and what isn't there yet. *IEEE Security & Privacy*, 22(3), 2024.

[4] C Yavuz. A multidisciplinary look at history and future of deepfake with Gartner hype cycle. *IEEE Security & Privacy*, 22(3), 2024.

[5] E De Cristofaro. Synthetic data: Methods, use cases, and risks. *IEEE Security & Privacy*, 22(3), 2024.

[6] M Ramon, M Vowels, and M Groh. Deepfake detection in super-recognizers and police officers. *IEEE Security & Privacy*, 22(3), 2024.

[7] L Maiano et al. Human versus machine: A comparative analysis in detecting artificial intelligence-generated images. *IEEE Security & Privacy*, 22(3), 2024.

The Editors

Daniel Moreira is an assistant professor in the Department of Computer Science at Loyola University Chicago, Chicago, IL 60611 USA. His research interests include media forensics, machine learning, computer vision, and biometrics applied for the greater good. Moreira received a Ph.D. in computer science from the Universidade Estadual de Campinas, Brazil. He is an associate editor of IEEE Transactions on Information Forensics and Security and Elsevier's Pattern Recognition journals. He was a former elected member of the IEEE Information Forensics and Security Technical Committee, 2021–2023 term, and an appointed member of the IEEE Signal Processing Society Education Center Editorial Board, 2022–2023 term. He is a Member of IEEE. Contact him at dmoreira1@luc.edu.

Sébastien Marcel is a senior researcher at the Idiap Research Institute, 1920 Martigny, Switzerland, where he heads the Biometrics Security and Privacy group; a professor at the School of Criminal Justice, University de Lausanne; a lecturer at the Ecole Polytechnique Fédérale de Lausanne; and the director

of the Swiss Center for Biometrics Research and Testing. His research interests include face recognition, speaker recognition, vein recognition, attack detection (presentation attacks, morphing attacks, deepfakes), template protection, and generative artificial intelligence for biometrics. Marcel received a Ph.D. in signal processing from Université de Rennes I in France at CNET, the research center of France Telecom (now Orange Labs). He is an associate editor of IEEE Transactions on Biometrics and Identity Science. He was an associate editor of IEEE Signal Processing Letters, an associate editor of IEEE Transactions on Information Forensics and Security, a guest editor of the IEEE Transactions on Information Forensics and Security Special Issue on Biometric Spoofing and Countermeasures, and coeditor of the IEEE Signal Processing Magazine Special Issue on Biometric Security and Privacy. He is also the lead editor of the Springer Handbook of Biometrics Anti-Spoofing (editions 1–3). He is a Senior Member of IEEE. Contact him at marcel@idiap.ch.

Anderson Rocha is a full professor of artificial intelligence and digital forensics at the Institute of Computing, University of Campinas, Campinas, SP 13.083-852, Brazil. His research interests include artificial intelligence, digital forensics, and reasoning for complex data. Rocha received a Ph.D. in computer science from the University of Campinas. He is an elected affiliate of the Brazilian Academy of Sciences and the Brazilian Academy of Forensic Sciences. He is a three-term elected member of the IEEE Information Forensics and Security Technical Committee (IFS-TC) and a former chair of this committee; in 2023, he was elected again as the IFS-TC chair for the 2025—2026 term. He is a Microsoft Research and a Google Research Faculty Fellow. In 2016, he was awarded the Tan Chin Tuan Fellowship, a recognition promoted by the Tan Chin Tuan Foundation in Singapore. Since 2023, he has been an Asia Pacific Artificial Intelligence Association fellow. He is ranked among the top 2% of research scientists worldwide, according to PLoS One/Stanford and Research.com studies. He is a LinkedIn Top Voice in Artificial Intelligence for continuously raising awareness of Al and its potential impacts on society at large. He is a Fellow of IEEE. Contact him at anderson.rocha@ic.unicamp.br.