1st Workshop on Hidden Information, Steganography, Privacy, and Emerging Risks (WHISPER 2026)

Lorenzo Cazzaro[1], Daniela Gallo[2,3], Angelica Liguori[2,*], Luca Caviglione[5], Giuseppe Manco[4] and Salvatore Orlando[1]

[1]*University of Luxembourg, Luxembourg*
[2]*Institute for High Performance Computing and Networking - Italian National Research Council, Rende, Italy*
[3]*University of Salento, Lecce, Italy*
[5]*Institute for Applied Mathematics and Information Technologies - Italian National Research Council, Genova, Italy*
[1]*Ca' Foscari University of Venice, Italy*

Abstract

Artificial Intelligence (AI) systems are increasingly integrated into a wide range of domains, reshaping the landscape of information security. From synthetic content generation to intelligent system deployment, AI technologies now influence every layer of communication, data processing, and multimedia production. This evolution brings both opportunities and risks. On one hand, AI can enhance defenses by detecting hidden information, tracing content origins, and identifying covert channels; on the other hand, it can empower adversaries to design stealthier data-hiding methods, generate highly realistic steganographic media, and optimize side-channel attacks. Moreover, even AI can be watermarked or targeted by data exfiltration attacks, being a fundamental asset for companies. In this context, the 1st Workshop on Hidden Information, Steganography, Privacy, and Emerging Risks (WHISPER 2026), co-located with ITASEC & SERICS, aims to bring together researchers and practitioners from both academia and industry to explore the intersection of AI, steganography, digital watermarking, and privacy. The workshop will focus on emerging risks, defensive and adversarial strategies, and the dual role of AI, as both a potential threat and a powerful safeguard, in the protection, tracing, and authentication of information.

Keywords

Information-hiding, watermarking, steganography, privacy

1. Preface

Artificial Intelligence (AI) systems are increasingly integrated into a wide range of domains, including sensitive and safety-critical tasks, reshaping the landscape of information security. From synthetic content generation to intelligent system deployment, AI technologies now influence every layer of communication, data processing, and multimedia production. This evolution brings both opportunities and risks. On the one hand, AI can enhance the design of information-hiding techniques and detect when these techniques are used; on the other hand, it introduces novel attack vectors and privacy challenges. For defensive purposes, AI models can detect fingerprints left by steganographic methods or monitor subtle variations in system parameters to expose covert channels that evade traditional defenses. Conversely, adversaries are leveraging AI for offensive strategies. Generative models can produce steganographic media that are statistically indistinguishable from benign content, while AI algorithms can uncover and optimize new side channels, enabling stealthy data exfiltration. This "cat-and-mouse" dynamics is a critical area of focus for the cybersecurity community.

1st Workshop on Hidden Information, Steganography, Privacy, and Emerging Risks (WHISPER 2026), co-located with the 2nd Joint National Conference on Cybersecurity (ITASEC & SERICS 2026), February 09–13, 2026, Cagliari, Italy
*Corresponding author.
lorenzo.cazzaro@uni.lu (L. Cazzaro); daniela.gallo@icar.cnr.it (D. Gallo); angelica.liguori@icar.cnr.it (A. Liguori); luca.caviglione@ge.imati.cnr.it (L. Caviglione); giuseppe.manco@icar.cnr.it (G. Manco); orlando@unive.it (S. Orlando)
0000-0001-6479-2949 (L. Cazzaro); 0009-0009-3245-7738 (D. Gallo); 0000-0001-9402-7375 (A. Liguori); 0000-0001-6466-3354 (L. Caviglione); 0000-0001-9672-3833 (G. Manco); 0000-0002-4155-9797 (S. Orlando)
© 2026 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

CEUR Workshop Proceedings ceur-ws.org ISSN 1613-0073
published 2026-11-14

CEUR-WS.org/Vol-4329/preface-whisper26.pdf

Within this evolving context, ensuring the protection, integrity, and authenticity of information has become a central cybersecurity challenge. The widespread adoption of generative AI further amplifies the need for developing reliable mechanisms to watermark and trace digital content, including ML models themselves. The first edition of "WHISPER - Workshop on Hidden Information, Steganography, Privacy, and Emerging Risks" aims to bring together researchers and practitioners to explore cutting-edge developments in covert communications, digital watermarking, and steganography, with emphasis on the use of AI.

The scope of the workshop includes, but is not limited to, the following topics:

- Watermarking schemes, usages, and challenges, including model and content watermarking for intellectual property protection.
- Novel cloaking mechanisms within models and/or datasets with robustness and fairness properties.
- New threat models leveraging information-hiding techniques.
- Secure watermarking of large language models, generative models, and traditional machine learning models.
- Deep learning solutions for detecting and generating covert communication.
- Adversarial attacks and defenses in data hiding and watermarking.
- Data-exfiltration attacks to AI models, including membership inference attacks and privacy leakages.

By fostering collaboration among academia, industry, and public institutions, the workshop seeks to promote discussion on cutting-edge developments and open research challenges in information-tracing and information-hiding, emphasizing also its intersection with AI and its security implications.

The paper submission deadline was December 18, 2025, followed by acceptance notifications on January 12, 2026, and the camera-ready submission deadline on February 23, 2026. The workshop took place on February 9, 2026.

The workshop accepted both original submissions and previously published papers. It received six papers, four of which were original contributions. Each submission underwent a peer-review process and was evaluated by two members of the Program Committee. Following the review process, four papers were accepted for presentation, of which two are included in these proceedings.

2. Keynote

The program for WHISPER 2026 included a keynote by Professor Mauro Barni (University of Siena), titled "Watermarking in the Age of Generative AI: lessons learned from the past and current challenges."

3. Program Committee

- Giacomo Benedetti, IMATI-CNR, Italy
- Antonio Emanuele Cinà, University of Genoa, Italy
- Massimo Guarascio, ICAR-CNR, Italy
- Wojciech Mazurczyk, Warsaw University of Technology & FernUniversitaet in Hagen, Poland and Germany
- David Megías Jiménez, Universitat Oberta de Catalunya
- Aleksandra Mileva, Goce Delcev University, Republic of Macedonia
- Silvia Lucia Sanna, University of Cagliari
- Diego Soi, University of Cagliari

Acknowledgments

The first edition of WHISPER is supported by Project WHAM! - Watermarking Hazards and novel perspectives in Adversarial Machine learning (B53D23013340006).

Declaration on Generative AI

The authors have not employed any Generative AI tools.