ENDICOTT COLLEGE

IEEE ISDFS'26

International Symposium on Digital Forensics and Security

14th INTERNATIONAL SYMPOSIUM ON DIGITAL FORENSICS AND SECURITY

SYMPOSIUM PROGRAM AND ABSTRACTS

EDITED BY

PROF. DR. ASAF VAROL

PROF. DR. MURAT KARABATAK

PROF. DR. CIHAN VAROL

March 19 – 20, 2026

Endicott College, Boston, MA, USA

Consortium Members:

IEEE Education Society · Wentworth Institute of Technology · Trinity University · UNIVERSITÉ TÉLUQ · IPCA Instituto Politécnico do Cávado e do Ave

14TH INTERNATIONAL SYMPOSIUM ON DIGITAL FORENSICS AND SECURITY

19-20 MARCH 2026

ENDICOTT COLLEGE

SYMPOSIUM PROGRAM AND ABSTRACTS

EDITED BY

PROF. DR. ASAF VAROL

PROF. DR. MURAT KARABATAK

PROF. DR. CIHAN VAROL

BOSTON - MA - US

2026

Preface by Prof. Dr. Asaf Varol, General Coordinator of ISDFS 2026
14th International Symposium on Digital Forensics and Security, Boston, MA, USA, Endicott College, March 19-20, 2026

Dear Participants of ISDFS 2026,

Since its inception in 2013, the International Symposium on Digital Forensics and Security (ISDFS) has grown into a well-recognized international platform that brings together researchers, practitioners, and industry experts working in digital forensics, cybersecurity, information assurance, and emerging security technologies. Hosted by several distinguished universities in the United States, Türkiye, Romania, and Portugal, ISDFS reflects strong international collaboration and academic impact. The symposium provides an interdisciplinary environment where scholars present innovative research findings, discuss emerging cyber threats, and explore new methodologies and technologies aimed at strengthening digital security infrastructures. Through keynote speeches, technical sessions, and collaborative discussions, ISDFS continues to contribute to the advancement of digital forensics and cybersecurity research worldwide. It is therefore our great pleasure to celebrate the successful completion of the **14th International Symposium on Digital Forensics and Security (ISDFS 2026)**. For more than a decade, this symposium has been organized without interruption, and even the global challenges caused by the **COVID-19 pandemic** did not break this tradition, demonstrating the resilience and dedication of the ISDFS community.

Over the years, the ISDFS conference has grown steadily in terms of academic quality, number of submissions, and international participation. Between **2016 and 2022**, the conference received technical sponsorship from different IEEE sections. During these years, ISDFS was successfully organized in several countries including **Türkiye, Romania, Portugal, Lebanon, and the United States**.

During the last **four years (since 2023)**, the symposium has been organized consecutively in different states across the **United States**, hosted by various universities under the umbrella of the **IEEE Education Society**. This collaboration has significantly strengthened the scientific visibility and global reputation of the conference.

The countries, host institutions, and dates of ISDFS conferences organized since 2013 are summarized in **Table 1.**

Table 1. Chronological Overview of the International Symposium on Digital Forensics and Security (ISDFS)

Year	Symposium	Location	Date
2026	14th International Symposium on Digital Forensics and Security (ISDFS 2026)	Endicott College, Boston, MA, USA	March 19–20, 2026
2025	13th International Symposium on Digital Forensics and Security (ISDFS 2025)	Wentworth Institute of Technology, Boston, MA, USA	April 24–25, 2025
2024	12th International Symposium on Digital Forensics and Security (ISDFS 2024)	Trinity University, Texas, USA	April 29–30, 2024
2023	11th International Symposium on Digital Forensics and Security (ISDFS 2023)	University of Tennessee at Chattanooga, Tennessee, USA	May 11–12, 2023
2022	10th International Symposium on Digital Forensics and Security (ISDFS 2022)	Maltepe University, İstanbul, Türkiye	June 6-7, 2022

2021	9th International Symposium on Digital Forensics and Security (ISDFS 2021)	Fırat University, Elazığ, Türkiye	June 28–29, 2021
2020	8th International Symposium on Digital Forensics and Security (ISDFS 2020)	Arab Open University, Beirut, Lebanon	June 1–2, 2020
2019	7th International Symposium on Digital Forensics and Security (ISDFS 2019)	Instituto Politecnico do Cavado e do ave, Barcelos, Portugal	June 10-12, 2019
2018	6th International Symposium on Digital Forensics and Security (ISDFS 2018)	Fırat University, Antalya, Türkiye	March 22–25, 2018
2017	5th International Symposium on Digital Forensics and Security (ISDFS 2017)	The University of Medicine, Pharmacy, Sciences and Technology of Tirgu Mures, Mures, Romania	April 26–28, 2017
2016	4th International Symposium on Digital Forensics and Security (ISDFS 2016)	University of Arkansas at Little Rock, USA	April 25–27, 2016
2015	3rd International Symposium on Digital Forensics and Security (ISDFS 2015)	Gazi University, Ankara, Türkiye	May 11–12, 2015
2014	2nd International Symposium on Digital Forensics and Security (ISDFS 2014)	Sam Houston State University, Texas, USA	May 12–13, 2014
2013	1st International Symposium on Digital Forensics and Security (ISDFS 2013)	Fırat University, Elazığ, Türkiye	May 20–21, 2013

The **technical sponsorship of the IEEE Education Society** has provided significant strength and international visibility for the conference. In this regard, we would like to express our sincere appreciation to **Prof. Dr. Arnold Pears**, President of the IEEE Education Society (2025–2027), and **Prof. Dr. Diana Andone**, Conference Chair, for their valuable support and encouragement.

One of the distinctive characteristics of the ISDFS conference is that it is governed and organized by an **international consortium** of universities and institutions. The consortium members for **ISDFS 2026** include:

- Association of Software and Cyber Security of Türkiye
- Istanbul Technical University (Türkiye)
- Maltepe University (Türkiye)
- Fırat University (Türkiye)
- Gazi University (Türkiye)
- IEEE Education Society (USA)
- Endicott College (USA)
- Wentworth Institute of Technology (USA)
- Trinity University (USA)
- University of Tennessee at Chattanooga (USA)
- Sam Houston State University (USA)
- University of Tartu (Estonia)
- Polytechnic Institute of Cavado and Ave (Portugal)
- Singidunum University (Serbia)
- TELUQ University (Canada)

The consortium membership is periodically updated. I would like to extend my heartfelt gratitude to all consortium members for their invaluable contributions and continuous support of the ISDFS conference.

All submitted papers undergo a **rigorous peer-review process**. Each submission is evaluated **by at least two independent reviewers**. Following the reviewers' recommendations, the **ISDFS Executive Committee** performs a final evaluation of the manuscripts. Papers that comply with **the IEEE publication standards and formatting requirements** are submitted for approval to IEEE and subsequently published in the **IEEE Xplore Digital Library**, one of the largest scientific digital libraries in the world. Through this process, the research presented at ISDFS becomes accessible to researchers and practitioners worldwide. Furthermore, the indexing of ISDFS by Scopus significantly enhances the international recognition and academic impact of the conference.

We would like to express our sincere appreciation to **Dr. Sara Quay, Provost of Endicott College**, and **Dr. Mark Herlihy, Dean of the School of Social Sciences, Communication, and Humanities,** for approving the organization **of ISDFS 2026 at Endicott College**. Special thanks are also extended **to Dr. KC (Kyungseok) Choo**, Associate Dean and Professor at Endicott College and Director of the Homeland Security Graduate Program, who served as **the Local Organizing Chair** and led the local organizing team with great dedication.

We are also grateful to the distinguished members of the **Honorary Committee**, including:

- Dr. Lori Mann Bruce, President of the University of Tennessee at Chattanooga
- Dr. Jerold Hale, Provost of the University of Tennessee at Chattanooga
- Dr. Kumar Yelamarthi, Dean of the College of Engineering and Computer Science
- Dr. Edibe Sözen Çetintaş, Rector of Maltepe University
- Prof. Dr. Fahrettin Göktaş, Rector of Fırat University

The **Microsoft CMT service** was used for managing the peer-review process for this conference. This service was provided free of charge by Microsoft, including all expenses related to **Azure cloud services**, **software development, and technical support**. Without this support, managing the **499 papers submitted to ISDFS 2026 through the CMT system** would not have been possible with such efficiency. We would like to extend our sincere thanks to **Microsoft** for this valuable contribution.

Out of **the 499 submitted papers, 197 papers** were accepted for presentation, resulting in an acceptance rate of approximately **39%** for ISDFS 2026.

Organizing an international conference is a demanding and complex process that requires the dedicated efforts of many individuals. The successful completion of ISDFS 2026 is the result of the collaborative work of numerous researchers, reviewers, students, and staff members working behind the scenes. I would like to express my sincere appreciation to all authors and participants who contributed to the success of this conference.

ISDFS 2026 is honored to host two distinguished **Keynote Speakers**. Our first keynote speaker, **Christophe Landries** from the **Belgian Federal Police / Directorate of International Police Cooperation / Joint Cybercrime Action Taskforce (J-CAT),** delivered the keynote titled:

"*AI-Powered Crime, Immersive Platforms, and the future of Criminal Investigations*"

Our second keynote speaker, **Supervisory Special Agent Peter Gomez** from the New **England Regional Computer Forensics Lab (NERCFL)**, delivered the keynote titled:

"The FBI Role in Global Cybercrime Disruption: Partnerships, Capacity Building, and Enforcement."

We sincerely thank both keynote speakers for their valuable contributions.

Finally, I would like to express my special appreciation to **Dr. Murat Karabatak** from **Fırat University** and **Dr. Cihan Varol** from **Sam Houston State University**, who have been actively involved in the organization of ISDFS since its establishment in 2013. Their dedication, tireless efforts, and continuous support have played a crucial role in the sustained success and international growth of the ISDFS symposium.

Once again, I would like to thank all participants and contributors and wish everyone a productive and inspiring conference.

March 19, 2026

Prof. Dr. Asaf Varol
General Coordinator, ISDFS

ORGANIZERS

HONORARY BOARD

Dr. Sara Quay	Provost, Endicott College, Beverly, MA, US
Dr. Mark Herlihy	Dean of Social Sciences, Communications & Humanities, Endicott College, Beverly, MA, US
Dr. Lori Mann Bruce	Chancellor, The University of Tennessee at Chattanooga, TN, US
Dr. Jerold Hale	Provost, The University of Tennessee at Chattanooga, TN, US
Dr. Kumar Yelamarthi	Dean of the College of Engineering and Computer Science, The University of Tennessee at Chattanooga, TN, US
Prof. Dr. Edibe Sözen	Rector of Maltepe University, TR
Prof. Dr. Fahrettin Göktaş	Rector of Fırat University, TR

General Chair of ISDFS

Prof. Dr. Asaf Varol	University Tennessee at Chattanooga, TN, US

Term Chair of ISDFS 2026

KC (Kyungseok) Choo, Ph.D.,	Associate Dean, School of Social Sciences, Communication, & Humanities, Professor and Director, Endicott College, Beverly, MA, US

Publication Chair

Prof. Dr. Cihan Varol	Sam Houston State University, TX, US

Treasurer

Prof. Dr. Murat Karabatak	Association of Software and Cyber Security of Türkiye Fırat University, Elazığ, TR

The IEEE Sponsors

Prof. Dr. Arnold Pears	IEEE Education Society (2025-2027 Future President)
Prof. Dr. Diana Andone	IEEE Education Society Conference Chair, US

IEEE Signatory Person

Prof. Dr. Asaf Varol	The University of Tennessee at Chattanooga, TN, US

Technical Program Contacts

KC (Kyungseok) Choo, Ph.D.,	Endicott College, Beverly, MA, US
Prof. Dr. Arnold Pears,	IEEE Education Society (2025-2027 Current President), Sweden
Prof. Dr. Diana Andone,	IEEE Education Society Conference Chair, Romania
Prof. Dr. Nizamettin Aydın,	IEEE Türkiye Section Chair, Türkiye
Assit. Prof. Dr. Rebeka Rocco	Criminal Justice Department, Endicott College, Beverly, MA, US
Assoc. Prof. Dr. Mark Beaudry,	Criminal Justice Department, Worcester State University, Adjunct Faculty, Homeland Security Graduate Program, Endicott College
Eric LaForte,	Special Agent, Mark Jenkins digital Forensics Laboratory, US Coast Guard.

Graphic Design

Prof. Dr. Murat Karabatak	Firat University, Elazig, TR
Mahmut Kaplan	Fırat University, Turkey

Web Design and Coding

Prof. Dr. Asaf Varol	The University of Tennessee at Chattanooga, TN, US
Elif Buse Köprücü	Fırat University, Elazığ, TR

Venue Contact

KC (Kyungseok), Choo	Endicott College, Bevely, MA, US

Consortium Members

Prof. Dr. Murat Karabatak	Association of Software and Cyber Security of Türkiye, TR
Prof. Dr. Şeref Sağıroğlu	Gazi University, Ankara, Turkey
Prof. Dr. Çetin Arslan	Council of Judges and Prosecutors (CJP), Türkiye, Ankara, Turkey
Prof. Dr. Özay Gürtuğ	Beykoz University, Istanbul, TR
Prof. Dr. Fahrettin Göktaş	Fırat University, Elazig, Turkey
Assoc. Prof. Dr. Sedat Akleylek	University of Tartu, Estonia
Prof. Dr. Maria Manuela Cruz-Cunha	Polytechnic Institute of Cavado and Ave, Portugal
Assist. Prof. Dr. Eva Tuba	Trinity University, US
Prof. Dr. Milan Tuba	Singidunum University, Serbia
Prof. Dr. Hamadou Saliah-Hassane	TELUQ University, Quebec, CA and IEEE Education Society
Prof. Dr. Ahad Nasab	The University of Tennessee at Chattanooga, Tennessee, US
Prof. Dr. Erkan Kaplanoğlu	The University of Tennessee at Chattanooga, US
Prof. Dr. Nizamettin Aydın	İstanbul Technical University, Istanbul, TR, IEEE Türkiye Section Chair
Prof. Dr. Arnold Pears	IEEE Education Society (2025-2027 President)
Assist. Prof. Dr. Ashar Neyaz	Wentworth Institute of Technology, Massachusetts, US
KC (Kyungseok), Choo	Endicott College, Bevely, MA, US
Prof. Dr. İhsan Yılmaz	Maltepe Üniversitesi, Istanbul, TR

SCIENTIFIC COMMITTEE MEMBERS

First Name	Middle	Last Name	Organization
Abdulsamet		Haşıloğlu	İstanbul 29 Mayıs Üniversitesi, TR
Ahad		Nasab	The University of Tennessee at Chattanooga, TN, USA
Ahmet		Çınar	Fırat University, TR
Ahmet	Furkan	Aydoğan	Sam Houston State University, US
Aishwarya		Asesh	Adobe
António	M. Rosado da	Cruz	Instituto Politécnico de Viana do Castelo, Portugal
Asaf		Varol	The University of Tennessee at Chattanooga, TN, USA
Ashar		Neyaz	Wentworth Institute of Technology, USA
Ashwin	Viswanathan	Kannan	Amazon Labs, USA
Avinash		Kumar	Sam Houston State University, USA
Ayhan	Osman	Erdem	Gazi University, TR
Ayşe		Güler	Gazi University, TR
Ayşe		Salman	Maltepe Universiy, TR
Aytuğ		Boyacı	National Defense University, TR
Banu		Diri	Yıldız Technical University, TR
Bedri		Özer	Firat University, TR
Bihter		Daş	Firat University, TR
Bilal		Alataş	Firat University, TR
Bisso		Saley	Université Abdou Moumouni, Nijer
Bartók		Roland	Miskolc-Egyetemváros, HUNGARY
Burcu		Koç	Middle East Technical University, TR
Chiamaka	Femi	Adeyinka	SHSU
Constantino		Martins	Polytechnic of Porto – School of Engineering, PT
Dalei		Wu	The University of Tennessee at Chattanooga, US
Derya		Avcı	Fırat University, TR
Ecir	Uğur	Küçüksille	Suleyman Demirel University, TR
Elio	Sancristobal	Ruiz	UNED, Spain
Ensar		Gül	Beykoz University, TR
Erdal		Güvenoğlu	Maltepe University, TR
Erhan		Akbal	Firat University, TR
Erkan		Kaplanoğlu	The University of Tennessee at Chattanooga, US
Erkan		Tanyıldızı	Firat University, TR
Eva		Tuba	Trinity University, US
Fatih		Özkaynak	Firat University, TR
Fatih		Özyurt	Fırat University, TR
Ferhat		Uçar	Firat University, TR
Gabriela		Gonçalves	Interdisciplinary Studies Research Center, PT
Gazi		Aygün	The University of Tennessee at Chattanooga, US
Gheorghe		Sebestyen	Technical University of Cluj-Napoca, Romania
Gökhan		Erdemir	The University of Tennessee at Chattanooga, US
Hacer		Varol	Stephen F. Austin State University, US
Hala		Strohmier	University of South Carolina, USA
Halit		Oğuztüzün	Middle East Technical University, TR
Halit	Hami	Öz	İstanbul Gelişim University, TR
Hamadou	Saliah	Hassane	Teluq University, Canada

Hanumat	Sastry	G	School of Computer Science UPES, India
Hornyák		Olivér	University of Miskolc, HU
Hürevren		Kılıç	Atılım University, TR
Huseyin		Aygün	Montgomery College, USA
Ibrahim		Soğukpinar	Gebze Institute of Technology, TR
Ibrahim		Türkoğlu	Firat University, TR
Janos		Juhasz	University of Miskolc, Hungary
Jayesh		Soni	Florida International University, US
Joao		Vilaça	Polytechnic Institute of Cávado and Ave, PT
Joaquim		Santos	Polytechnic of Porto – School of Engineering, PT
Joaquim	P.	Silva	Polytechnic University of Cávado and Ave, Portugal
Johann	Marquez	Barja	IMEC & University of Antwerp
József		Vásárhelyi	University of Miskolc, Hungary
Jolly		Upadhayaya	Merrimack College, US
Kayhan		Erciyes	Yasar University, TR
Khalid	Zine	Dine	Mohammed V University in Rabat, FSR, Morocco
Khushi		Gupta	University of North Georgia, USA
Lahcen		Tamym	Aix-Marseille University, University of Toulon, Marseille, FR
Lalith		Maddali	BrightEdge Technologies, US
Laszlo		Kovacs	University of Miskolc, HU
Lei		Zhang	University of Maryland Eastern Shore, US
Leonel	Filipe Simões	Santos	Polytechnic of Leiria, Portugal
Luis		Ferreira	Polytechnic Institute of Cávado and Ave, Portugal
Madhuri		Koushik	Netflix
Mandar		Khoje	Moveworks, USA
Manoj		Gupta	JECRC University, Jaipur, IN
Maria Manuela	Cruz	Cunha	Polytechnic Institute of Cávado and Ave, PT
Mehmet	Sıraç	Özerdem	Dicle University, TR
Muhammet		Baykara	Firat University, TR
Miguel	Rodriguez	Artacho	UNED University (Spain)
Mohammed	Chaouki	Abounaima	Sidi Mohamed Ben Abdellah University, Fes Morocco
Muharrem	Tolga	Sakallı	Trakya University, TR
Murat		Aydoğan	Firat University, TR
Murat		Karabatak	Firat University, TR
Murat		Koyuncu	Atılım University, TR
Mustafa		Kaya	Fırat University, TR
Narasimha K.	Karpoor	Shashidhar	Sam Houston State University, USA
Nizamettin		Aydın	İstanbul Technical University, TR
Nuno	Mateus	Coelho	University of Trás os Montes e Alto Douro, PT
Nuno		Lopez	Polytechnic Institute of Cávado and Ave, PT
Nuno		Rodrigues	Polytechnic Institute of Cávado and Ave, PT
Osayomore		Aigbogun	Sam Houston State University, US
Osman	Ayhan	Erdem	Gazi University, TR
Özal		Yıldırım	Firat University, TR
Patrícia	Isabel	Leite	Politécnico do Cávado e do Ave, PT
Paulo	Marques	Teixeira	Polythecnic University of Cavado and Ave, PT
Pedro		Pinto	Instituto Politécnico de Viana do Castelo (IPVC), PT
Phani		Lanka	Sam Houston State University, US
Rabab		Benotsmane	University of Miskolc, HU
Raul		Cordeiro	Sharing University, US

Roland		Bartók	University of Miskolc, HU
Resul		Daş	Edinburgh Napier University, UK
Rui	Miguel Soares	Silva	Instituto Politécnico de Beja
Salah	El	Hadaj	Cadi Ayad University, Morocco
Salim	Jibrin	Danbatta	Üsküdar University, TR
Sandro		Carvalho	Polytechnic Institute of Cávado and Ave, Portugal
Sedat		Akleylek	Tartu University, Estonia
Şengül		Doğan	Firat University, TR
Şeref		Sağıroğlu	Gazi University, TR
Serkan		Gönen	İstanbul Gelişim University, TR
Serkan		Varol	The University of Tennessee at Chattanooga, USA
Sırma	Çekirdek	Yavuz	Yıldız Technical University, TR
Şule	Gündüz	Öğüdücü	Istanbul Technical University, TR
Susana		Nicola	Instituto de Engenharia de Sistemas e Computadores Tecnologia e Ciência, PT
Susan		Anwar	Philander Smith College, US
Timothy		Laryea	Western Illinois University, US
Tuğba		Akyel	Maltepe University, TR
Tuğba		Dalyan	İstanbul Bilgi University, TR
Türker		Tuncer	Firat University, TR
Ümit		Can	Munzur University, TR
Vásárhelyi		József	University of Miskolc, HU
Volkan		Tunalı	University of the West of Scotland, UK
Xiao		Hu	Purdue University, USA
Yingfeng		Wang	The University of Tennessee at Chattanooga, TN, USA
Yu		Liang	The University of Tennessee at Chattanooga, US
Yunus		Santur	Firat University, TR
Yusuf		Özturk	San Diego University, USA
Zisis		Tsiatsikas	University of the Aegean, GR

KEYNOTE SPEAKERS

1ST SPEAKER

Name of Presenter	Christophe LANDRIES
Rank	Specialized Chief Inspector (Detective Informatics/Accountancy)
Current Job Title	J-CAT (Cyber) Liaison Officer
Current Agency/ Institution	Belgian Federal Police / Directorate of International Police Cooperation / Joint Cybercrime Action Taskforce (J-CAT)
Talk Title	AI-Powered Crime, Immersive Platforms, and the future of Criminal Investigations.

Speaker Biography

Since November 2021, I'm working as a cyber liaison officer for the Belgian Federal Police in J-CAT (Joint Cybercrime Action Task-force) based at Europol HQ. Before, I worked for the Federal Computer Crime Unit as an expert in dark web and cryptocurrency investigations. I have more than 17 years of experience in cyber-crime investigations and delivered/developed many trainings in dark web and cryptocurrency related topics for CEPOL, Council of Europe, OCSE, law enforcement agencies, Catholic University of Lille etc. Graduated in accountancy, informatics and obtained a master's degree in tax law in 2019 (thesis: Capital gains on cryptocurrency). In 2024, I successfully completed my master after master in company law (thesis: Criminal liability of legal entities, their directors and agents in connection with the use of cryptocurrencies) and in 2025 a master in meta-verse at the University of Nicosia (online placement at Endicott College developing a law enforcement training package in meta-verse, sex meta-crimes and blockchain investigations). Furthermore, I'm specialized in cyber-crime investigations and in many fields of cryptocurrency such as investigations, blockchain analysis, trading, legislation, new and emerging technologies, taxation etc. Also, I'm specialized in meta-verse and virtual worlds investigations as well as in artificial intelligence in relation to (cyber) crime.

2ND SPEAKER

Name of Presenter	Peter GOMEZ
Rank	Supervisory Special Agent (SSA)
Current Job Title	Lab Director
Current Agency/ Institution	New England Regional Computer Forensics Lab (NERCFL)
Talk Title	The FBI Role in Global Cybercrime Disruption: Partnerships, Capacity Building, and Enforcement

Speaker Biography

SSA Gomez has served in the FBI for over 27 years. He is currently the Lab Director for the New England Regional Computer Forensics Lab (NERCFL) working with federal, state, and local partners to provide digital forensic services to law enforcement partners in Massachusetts, New Hampshire, Rhode Island, and Maine. Prior to this role, SSA Gomez was the FBI's Senior Cyber Liaison Officer at Europol assigned to the Joint Cybercrime Action Taskforce (J-CAT). SSA Gomez has also served as an FBI Law Enforcement Attache working cybercrime and other matters in Australia, New Zealand, Philippines, Turkey, and Canada. SSA Gomez served more than a decade overseas working with multiple international law enforcement and security partners. SSA Gomez served for two years at FBIHQ Cyber Division as a program manager and the National Program Manager for the FBI Cyber Crime Task Force (CCTF). SSA Gomez has served his entire career with the FBI working on multiple task forces in Boston and internationally. Prior to joining the FBI, SSA Gomez was a state prosecutor in Orlando, Florida for three years conducting approximately forty jury trials.

SYMPOSIUM PROGRAM

VENUE : **ENDICOTT COLLEGE, 376 HALE ST, BEVERLY, MA 01915, UNITED STATES**

ZOOM WAC 126 : **CLICK TO JOIN MEETINGS** **MEETING ID: TBD;** **PASSWORD: TBD**

ZOOM WAC 127 : **CLICK TO JOIN MEETINGS** **MEETING ID: TBD;** **PASSWORD: TBD**

MARCH 19, 2026 (THURSDAY)

08:30-09:30 **REGISTRATION**

09:30-10:00 **OPENING CEREMONY (HALLE LIBRARY-THE LITTLE THEATER)**

DR.K.C (KYUNGSEOK) CHOO, ISDFS 2026 TERM CHAIR

ASAF VAROL, GENERAL CHAIR OF ISDFS

PROF. DR. ARNOLD PEARS, CHAIR OF IEEE EDUCATION SOCIETY, KTH, SWEDEN)

DIANA ANDONE, VICE PRESIDENT FOR CONFERENCES & WORKSHOPS

MARK HERLIHY, DEAN OF THE SCHOOL OF SOCIAL SCIENCES, COMMUNICATION & HUMANITIES

10:00-10:45 **KEYNOTE SPEECH (HALLE LIBRARY-THE LITTLE THEATER)**

CHAIR **PROF. DR. ASAF VAROL**

SPEAKER **CHRISTOPHE LANDRIES**

TITLE **AI-POWERED CRIME, IMMERSIVE PLATFORMS, AND THE FUTURE OF CRIMINAL INVESTIGATIONS**

10:45-11:00 **COFFEE BREAK**

11:00-12:30 **SESSION 01 PRESENTATIONS (CHAIR: DR. ASAF VAROL)**

(WAX ACADEMIC CENTER 125 - IN PERSON)

PAPER ID 1: FORENSIC INVESTIGATION OF SDRSHARP AND RTL-SDR DONGLE IN EAVESDROPPING RADIO COMMUNICATIONS
ERASMUS MFODWO, NARASIMHA SHASHIDHAR, CIHAN VAROL AND AHMET FURKAN AYDOGAN

PAPER ID 75: WEAK ENFORCEMENT AND LOW COMPLIANCE IN PCI DSS: A COMPARATIVE SECURITY STUDY
SOONWON PARK AND JOHN D. HASTINGS

PAPER ID 165: FAKE-OT: AN OPEN ARCHITECTURE FOR THREAT DETECTION IN OT ENVIRONMENTS USING MQTT-BASED HONEYPOTS
HEBERT SILVA, FELIPE CARVALHO AND TIAGO DEMAY

Paper ID 268: An Explainable AI Study of Phishing Detection Model Degradation Across Real-World Events
Michael Ivanicki and Brian Robert Callahan

Paper ID 277: Evaluating Security Policy Compliance in Infrastructure as Code Generated by Large Language Models
Ryo Hase, Ye Wang, Toshiaki Koike-Akino, Jing Liu, Kieran Parsons and Jumpei Hato

Paper ID 350: Porting and Evaluating Return-Oriented Programming Defenses Implemented by the OpenBSD Operating System
Jenna Esposito, Aaila Arif, Raul Cortinas and Brian Robert Callahan

11:00-12:30 Session 02 Presentations (Chair: Sundeep Bobba)

(Wax Academic Center 126 – Zoom Online)

Paper ID 23: An Information Asymmetry Game for Trigger-based DNN Model Watermarking
Chaoyue Huang, Gejian Zhao, Hanzhou Wu, Zhihua Xia and Asad Malik

Paper ID 87: Design and Analysis of Energy Efficient Ascon 80PQ
Ananthraj Rao Kekuda, Anirudh G P, Ashwin Thomas and Srinivasa V S Sarma D

Paper ID 256: Comprehensive Review and Experimental Study of Holiday-Aware Multimodal Crime Prediction
Sapna V. M, Vishruth S. Megur, Vishwanath Bhavimani, Somesh, Vignesh Madivala and Prasad B Honnavalli

Paper ID 274: An Investigation on the application of Pareto Principle to Windows Based System Resource Usage
Mihai Lazarescu and Sieteng Soh

Paper ID 333: High Accuracy Is Not Enough: Epistemic Bias in Machine Learning Task Formulation
Grazia Garzo and Alessandro Palumbo

Paper ID 410: Artificial Intelligence-Driven Predictive Models for Identifying Risk Factors of Chronic Diseases
Shaiful Mahmud, Khaleel Khan Mohammed, Vasu Raj Jain and Sarthak Anandkumar Shah

Paper ID 454: Real-Time Vision-Based Human–Robot Interaction Framework for Low-Cost Embedded Robotic Arms
Wewage Anjana Pruthuvi Dep and Vishimi Embuldeniya

11:00-12:30 Session 03 Presentations (Chair: Karthik Pappu)

(Wax Academic Center 127 – Zoom Online)

Paper ID 19: AI-Powered Android Immune System: A Hybrid Static and Semantic Model for Malware Neutralization
Mohd Fozla Rabby

Paper ID 56: SQL vs. NoSQL Comparative Analysis of Database Management Systems for Scalability and Performance
Ravi Sankar Susarla and Viswa Bharath Kolla

Paper ID 225: Research Study on AI-Driven Weapon and Violence Detection in Smart Cities
Thangavel Murugan, Anoud Salem Alkarbi, Maryam Ahmed Alzeyoudi, Noura Ali Matar Alketbi, Afra Khaleifah Alteneiji and N. Nasurudeen Ahamed

Paper ID 265: Enhancing Detection of Bloodstain Patterns and Footprint Impressions using Deep Learning
Thangavel Murugan, Maryam Saif Abdulla Saeed, Hind Abdulla Asman Balmur Al Ameri, N Nasurudeen Ahamed, Aysha Salem Hamdan Obaid Alkaabi, Priyan Malarvizhi Kumar, Shamsa Rashid Salem Balhas Alshamsi and Alia Ali Salem Alwali Almazrouei

Paper ID 296: Evidence-Quality Telemetry for Cloud Incident Response: Detecting Gaps, Drift, and Integrity Failures in Large-Scale Observability Pipelines
Chalapathi Koneni and Sanjay Lokula

Paper ID 396: BloodTrace: Where Drops Become Decision
Dharanidharan Ranganathan, Devkumar Barot, Qian Liu and Oluwasola Mary Adedayo

Paper ID 434: Akura: Adaptive Sinhala Learning for Early Childhood using Gesture, Voice, Emotion and Handwriting
N. W. P. G. T. Mihiran, R. M. D. S. Rathnayake, H. A. D. T. Piyathilaka, N. W. P. G. T. T. Rahul, Dinuka Wijendra and Jenny Krishara

12:30-14:00 Lunch

14.00 – 14.45 **Keynote Speech (Halle Library-The Little Theater)**

Chair **KC (Kyungseok) Choo**

Speaker **Peter Gomez**

Title **The FBI Role in Global Cybercrime Disruption: Partnerships, Capacity Building, and Enforcement**

15:00-16:30 **Session 04 Presentations (Chair: Hala Strohmier)**
(Wax Academic Center 125 - In Person)

Paper ID 156: Intelligent Data Integration and Analytics for Automotive Aftermarket Retail: An Informatica-BigQuery-Tableau Framework
Rambabu Tangirala

Paper ID 200: Performance Evaluation of Post-Quantum Cryptography in IoT: A Case Study on MQTT over TLS under Network Constraints
Emerson Costa Santos, Tiago Augusto Orcajo Demay Cordeiro and Hebert de Oliveira Silva

Paper ID 282: AskCMMC.ai for AI-Driven CMMC 2.0 Compliance Automation
Hala Strohmier Berry, Aaryan R Londhe, Abdur Rahman Khan and Ronit Pawar

Paper ID 285: Stack Buffer Overflow Risk Analysis for PX4-Controlled Commercial UAVs
Hala Strohmier Berry, Chris Clark, Kadin DeMesme, Marco Paolo Saavedra and Travis Samatov

Paper ID 387: Exploration of Bottlenecks and Optimizations in Privacy-Preserving Machine Learning Inference
Bingyu Liu, Salem Othman, Leonidas Deligiannidis, and Robbie McCue

Paper ID 398: On Impact of Ensemble Strategies in Tabular Data Regression
Ivona Brajevic, Una Tuba and Milan Tuba

15:00-16:30 **SESSION 05 PRESENTATIONS (CHAIR: SHIVA KUMARA)**

(WAX ACADEMIC CENTER 126 – ZOOM ONLINE)

PAPER ID 166: NORMALIZING FLOW FOR FINANCIAL ANOMALY DETECTION: THE IMPACT OF LOSS FUNCTION
AMIR RASHIDI, RUOBIN QI AND RASHIDA HASAN

PAPER ID 241: QUANTUM-SAFE COUNTERMEASURES: MITIGATING DETECTOR-BLINDING ATTACKS IN QUANTUM KEY DISTRIBUTION SYSTEMS
WAIL ZITA AND MALEK MALKAWI

PAPER ID 245: TOKEN-BASED AUTHENTICATION SYSTEM FAILURES IN CLOUD ENVIRONMENTS: A CASE STUDY OF THE MICROSOFT STORM-0558 INCIDENT
HAMZA BENMILOUD AND MALEK MALKAWI

PAPER ID 257: **BIG DATA-DRIVEN AI MODELS FOR NETWORK ANOMALY DETECTION AND CYBER THREAT FORECASTING**
AMIT KUMAR MESHRAM AND VIKRANT SIKARWAR

PAPER ID 263: ARTIFICIAL INTELLIGENCE-BASED ANTI-MONEY LAUNDERING SOLUTIONS: ENHANCING DETECTION ACCURACY IN HIGH-VOLUME FINANCIAL DATA
AMIT KUMAR MESHRAM AND VIKRANT SIKARWAR

PAPER ID 317: PERFORMANCE ANALYSIS OF A CLOUD-NATIVE WEB APPLICATION DEPLOYED ON KUBERNETES
DHARMENDRA AHUJA

PAPER ID 363: SAFEKID SCAN: EARLY DETECTION OF DIGITAL ADDICTION IN MINORS
ANJANA H.H.H, WEERAKKODI Y.P, GIMHANI J.M.K.P, NETHMINI M.A.N, JENNY KRISHARA AND POORNA PANDUWAWALA

15:00-16:30 **SESSION 06 PRESENTATIONS WAPAIS + ISDFS (CHAIR: NUNO LOPES)**

(WAX ACADEMIC CENTER 127 – ZOOM ONLINE)

PAPER ID 192: NEUROSHIELD: A TEMPORAL SPIKE TIMING ATTACK DETECTION FRAMEWORK FOR STDP BASED NEUROMORPHIC SYSTEMS IN MULTI-CLOUD ENVIRONMENTS
NAGA SUJITHA VUMMANENI, SUNDEEP BOBBA AND AKHIL PEDDI

PAPER ID 193: EVALUATING THE VULNERABILITY OF DEEPFAKE IMAGE DETECTION MODELS TO ADVERSARIAL MANIPULATIONS
PRAKRITI SHAKYA, KATYA MKRTCHYAN AND RASHIDA HASAN

PAPER ID 195: ZERO TRUST ARCHITECTURE FOR 5G-ENABLED MOBILE CLOUD COMPUTING (MCC)
URVISH PANDYA AND SWETA PANDYA

PAPER ID 235: THE NEED FOR STANDARDIZED EVIDENCE SAMPLING IN CMMC ASSESSMENTS: A SURVEY-BASED ANALYSIS OF ASSESSOR PRACTICES
LOGAN THERRIEN AND JOHN HASTINGS

PAPER ID 273: ADAPTINDEX: ADAPTIVE INDEX SELECTION FOR IOT VECTOR DATABASES
CHANDRASHEKHAR MEDICHERLA, CHAITANYA KULKARNI, VISWANATHAN RANGANATHAN, MILAN PARIKH AND VINAY SONI

PAPER ID 339: MALWARE CLASSIFICATION ON PE FILES USING DEEP NEURAL NETWORKS
DANIEL VILACA AND LUIS FERREIRA

PAPER ID 409: SELECTIVE MEMORY SHARING IN MULTI-AGENT LLM TEAMS VIA AGENTGYM-RL
BHAVUK JAIN, GUNJAN JAIN AND HARDEO K. THAKUR

PAPER ID 488: COMPARATIVE STUDY OF SYMBOLIC EXECUTION TOOLS APPLIED TO VULNERABILITY DETECTION
ANDRE CARDOSO, OSCAR RIBEIRO AND NUNO LOPES

16:30-16:45 Coffee Break

16:45-18:30 Session 07 Presentations (Chair: Milan Tuba)

(Wax Academic Center 125 - In Person)

Paper ID 242: AI-Enhanced Cybersecurity Risk Assessment for Smart Grid Infrastructures Using NIST Framework
Hala Strohmier Berry, Bhargav Palaskar, Soumik Mitra and Yash N Dasri

Paper ID 391: A Systematic Study of Security and Privacy in Large Language Models
Bingyu Liu

Paper ID 392: Trustworthy and Reliable Machine Learning for Healthcare
Bingyu Liu

Paper ID 397: Transfer Learning for Malware Detection Using RGB Binary Visualization: A Comparative Study
Eva Tuba, Ivona Brajevic, Adis Alihodzic, Ana Trisovic and Milan Tuba

Paper ID 407: TRADES-Based Defense Against Adversarial Attacks in Medical Image Classification
Kily Jasso, Audrey Tollett, Ira Tuba and Eva Tuba

Paper ID 450: Tamper Detection in CT DICOM Images for Digital Forensics and Clinical Integrity
Hala Strohmier Berry, William Carroll, Jessica Brown and Sydney Halupa

16:45-18:30 Session 08 Presentations (Chair: Ashar Neyaz)

(Wax Academic Center 126 - In Person)

Paper ID 220: Hardening the OSv Unikernel with Efficient Address Randomization: Design and Performance Evaluation
Alex Wollman and John Hastings

Paper ID 299: Volatile Memory Forensics of Tails OS in a Virtualized Environment
Nurettin Senol, Jayden Thai, Semih Cal and Ahmet Aydogan

Paper ID 372: Governance and Auditability of AI-Driven Retail Decision Pipelines in Cloud-Native Architectures
Prithvi raj Veluchamy

Paper ID 384: Enhancing Phishing Url Detection with Machine Learning Algorithms
Weijie Pang, Ashar Neyaz and Raymond Eichner

Paper ID 395: From Binary Vulnerability Detection to CWE Classification: A Hierarchical Prompting Study
Obinna Okeke, Sushant Nepal, Venkat Sai Suman Lamba Karanam and Yan Wu

PaperID 443: High-Performance Distributed Deep Learning Using Adaptive Parallelism and Dynamic Workload Scheduling
Pavan Kumar Boyapati and Siva Teja Reddy Kandula

16:45-18:30 Session 09 Presentations (Chair: Sanjoy Mukherjee)

(Wax Academic Center 127 – Zoom Online)

Paper ID 7: Two Pillars of Banking Intelligence: A Comparative Analysis of AI Techniques for Fraud Prevention and Churn Mitigation
Sreenivasulu Gajula

Paper ID 50: Optimizing Resistive Losses in Transmission Systems via Voltage Level Selection
Kendall D. Standridge-Monroe and Asaf Varol

Paper ID 90: Deterministic LLMs: A Practical Forensic Framework for Verifiable and Reproducible Local LLM Inference
Joel D. Molina

Paper ID 121: An Effective System for Medical Image Diagnosis Using Deep Convolutional Networks (CNNs) in Healthcare Sector
Sanjoy Mukherjee

Paper ID 146: Enhancing Early Diabetes Screening Through Machine Learning and Explainable AI
Jahnavi Anilkumar Kachhia

Paper ID 148: Secure and Privacy-Preserving Healthcare Federated Learning via Differential Privacy Mechanisms
Maunik K Shah and Munir Rajesh Mehta

Paper ID 270: Unsupervised Baseline Clustering and Incremental Adaptation for IoT Device Traffic Profiling
Sean M. Alderman and John D. Hastings

Paper ID 393: Detecting Fileless Malware through Memory Forensics with Recurrent Neural Networks
Noah Preston and Ajay Kumara Makanahalli Annaiah

19:00-21:00 Gala Dinner

March 20, 2026 (Friday)

09:30-11:00 Session 10 Presentations (Chair: Fatih Özkaynak)

(Wax Academic Center 125 - In Person)

Paper ID 4: BICIR: A Blockchain-Based Interoperability Model for Cross-National Cyber Incidents
Osayomore O. Aigbogun and Cihan Varol

Paper ID 114: From Monolithic Middleware to Cloud-Native Microservices: A Performance-Driven Modernization Study
Sauhard Bhatt

Paper ID 167: Operationalizing Data Economy in Data Spaces: A Token-Based Reference Implementation
Muhamed Turkanovic´, Martin Ferenec, Teo Lah and Vid Kersic

Paper ID 371: Romance Scam Reporting and Support: Help-Seeking Timing, Trust, and Escalation
LD Herrera

Paper ID 400: AI-Driven Adaptive Training Architecture for Post-Disaster Structural Damage Assessment
Hüseyin Deniz, Vahdettin Cem Baydoğan, Bahar Demirel and Fatih özkaynak

09:30-11:00 Session 11 Presentations (Chair: Deepak Suram)

(Wax Academic Center 126 – Zoom Online)

Paper ID 142: Explainable and Human-Centric Approaches in Audio Steganalysis: A Review
Sarah Rahim and Guhanathan Poravi

Paper ID 173: A Comparative Analysis of URL Features for Machine Learning Phishing Detection
Chukwunalu Asuai and Yusuf Moshood

Paper ID 183: Explainable Risk Decision Systems Using Artificial Intelligence Models for Payment Fraud Detection and Identification

Deepak Reddy Suram

Paper ID 338: XGBoost-Enhanced Recurrent Hybrid Models for Wearable Inertial Navigation in GNSS-Denied Environments

Vedant Singh, Yash Sinha, Tushar Swami, SK Hithasree and Nagalakshmi S R

Paper ID 355: Version-Controlled Decentralized Firmware Integrity Verification with On-Chain Rollback Protection for Cyber-Physical Systems on Ethereum

Maruf Farhan, Usman Butt, Madhuki Rajapakshe and Rejwan Bin Sulaiman

Paper ID 365: QSignature 1.0: A Dynamical Regime Classification Framework for Causal Time Series Data

Ahmad Muhammad, Salim Jibrin Danbatta, Muhammad Abubakar Isah, Ibrahim Yahaya Muhammad, Sulaiman Sulaiman Ahmad and Abdelrahman Ghozlan

Paper ID 457: DASTestBed: An Automated Benchmarking Framework for DAST Scanners with Extensible Ground Truth Modeling

Rand Deeb, Alisa Vorobeva and Omar Farshad Jeelani

Paper ID 459: Feature-Equivalence Deduplication and Memoization of HTTP(S) Requests for Web Scanners: Formal Model, Concurrency, and Complexity Bounds

Rand Deeb, Alisa Vorobeva and Omar Farshad Jeelani

09:30-11:00 Session 12 Presentations (Chair: Thangavel Murugan)

(Wax Academic Center 127 – Zoom Online)

Paper ID 98: GenReachAI: An Agentic Generative AI Framework for Automated Reachability Analysis of Enterprise Software Vulnerabilities

Niranjan Pachaiyappan

Paper ID 108: A Systematization of Privacy Threats and Defenses in Modern Agentic Web Browsers

Niranjan Pachaiyappan

Paper ID 158: Modernizing SAP Business Objects Using Enterprise Data Warehousing Principles

Prasanth Sathyapalan

Paper ID 213: Empirical Analysis of AI Confidence Methods in Digital Forensic Standards

Milinda Rambel Stone and Varghese Vaidyan

Paper ID 221: Nowhere to Hide: Comparative Analysis of Residential vs. Cloud Attack Surfaces

Sankalp Lanka, Om Palsule, Rishik Lanka and Phani Lanka

Paper ID 236: Forensic Event Reconstruction of Web Attacks Using Log Decoder and Rule-Based Correlation

Muhammad Nur Yasir Utomo, Hudan Studiawan and Baskoro Adi Pratomo

Paper ID 334: FakeSpeech: LLM-Driven Semantic Manipulation and Voice Cloning for Realistic Deepfake Speech Benchmarking

Alaa Alsaeedi, Amal AlMansour and Amani Jamal

11:00-11:15 Coffee Break

11:15-12:30 SESSION 13 PRESENTATIONS (CHAIR: EVA TUBA)

(WAX ACADEMIC CENTER 125 - IN PERSON)

PAPER ID 160: AI-ENHANCED DIGITAL FORENSICS: A RESEARCH VISION FOR TRUSTWORTHY, EXPLAINABLE, AND HUMAN-CENTERED FORENSIC INTELLIGENCE

NADEEM DAUDPOTA

PAPER ID 311: NEURO-SYMBOLIC GRAPH AUTOENCODERS WITH RARE PATTERN MINING FOR PROVENANCE-BASED ANOMALY DETECTION

ASIF TAUHID, SIDAHMED BENABDERRAHMANE, MOHAMAD ALTRABULSI, AHAMED FOISAL AND TALAL RAHWAN

PAPER ID 340: DIGITAL FORENSICS APPLICATIONS OF ELECTRIC NETWORK FREQUENCY

FATIH YAMAN; ÜMÜHAN ÖZKAYNAK; FATIH ÖZKAYNAK

PAPER ID 430: AI FOR DEVSECOPS OPTIMIZATION: INVESTIGATING THE ROLE OF AI/ML IN PREDICTIVE VULNERABILITY DETECTION DURING CI/CD PIPELINE STAGES

TARUN KALWANI

PAPER ID 449: DATA PROTECTION AND NATIONAL SECURITY IMPLICATIONS OF POST-DISASTER STRUCTURAL DAMAGE ASSESSMENT SYSTEMS

VAHDETTIN CEM BAYDOĞAN, BAHAR DEMIREL, TUBA DEMIR, SEDAT SAVAŞ, FERHAT UÇAR, FATIH ÖZKAYNAK

11:15-12:30 SESSION 14 PRESENTATIONS (CHAIR: HENRY CYRIL)

(WAX ACADEMIC CENTER 126 – ZOOM ONLINE)

PAPER ID 132: AI-DRIVEN SELF-HEALING AND INTELLIGENT QUEUING THROUGH ANOMALY DETECTION IN 5G CELLULAR NETWORKS

HENRY P CYRIL

PAPER ID 186: TRANSFORMER BASED FRAMEWORK FOR IMBALANCED TRANSACTION FRAUD DETECTION IN FINTECH SYSTEMS

SANDEEP SHIVAM, VENKAT NUTALAPATI, TEJAS PRAVINBHAI PATEL, AMIT KUMAR PADHY, MADHUSHREE KUMARI AND RAJESH PURUSHOTHAMAN

PAPER ID 275: TRPO MULTI-AGENT ACTIVE LEARNING FOR EXPLAINABLE DDOS DETECTION IN HEALTHCARE IOMT

OMAR FARSHAD JEELANI

PAPER ID 302: INTEGRATING GENERATIVE AI INTO RETAIL CHECKOUT SYSTEMS: A CASE STUDY IN CLOUD AND APPLICATION INTEGRATION

GOPALAKRISHNAN VENKATASUBBU AND RAJGOPAL DEVABHAKTUNI

PAPER ID 385: ADVERSARIAL ROBUSTNESS OF ML-BASED INTRUSION DETECTION SYSTEMS

DMITRY SIVKOV, ROMAN SAFIULLIN, ALISA VOROBEVA, VIKTORIIA KORZHUK AND OMAR FARSHAD JEELANI

PAPER ID 432: PRODUCTIZING AI-DRIVEN NETWORK SECURITY SYSTEMS: ARCHITECTURE, TRADE-OFFS, AND PRODUCT MANAGEMENT PERSPECTIVES

BALU CHAVAN

PAPER ID 444: AUTOMATING ORGANIZATIONAL CYBER SECURITY POLICY COMPLIANCE AGAINST INDUSTRY STANDARDS USING AGENTIC AI

ROHIT NEGI, SOUMYO V. CHAKRABORTY, AMIT NEGI AND SANDEEP K. SHUKLA

PAPER ID 460: DATA-DRIVEN PREDICTION OF ADVERTISING DIGITAL CAMPAIGN EFFECTIVENESS USING ARTIFICIAL INTELLIGENCE

ABHINAY KUMAR REDDY SEELLA

11:15-12:30 SESSION 15 PRESENTATIONS (CHAIR: DR. SURYA SUNNAM)

(WAX ACADEMIC CENTER 127 – ZOOM ONLINE)

PAPER ID 233: A HYBRID APPROACH TO DETECT ILLEGAL ACTIVITIES IN DARK WEB DATA
AKASH G GAONKAR, ABHINAV V P, ADITYA V BHAT, JENY JIJO AND ABISHEK K

PAPER ID 251: (SEAL) SEQUENTIALLY EVOLVINGALERT LEARNING FOR SMARTCYBER SECURITY
RIONA. V AND R. HEMALATHA

PAPER ID 255: LARGE-SCALE FINANCIAL FORECASTING USING ADVANCED GAI-BASED LARGE LANGUAGE MODELS AND TIME SERIES ANALYSIS
SURYA VEERA BRAHMAJI RAO SUNNAM

PAPER ID 259: DRONE ASSISTED REMOTE WELLNESS MONITORING USING RGB CAMERA
EGEMEN BORA TUNCARSLAN AND IHSAN YILMAZ

PAPER ID 284: YOLO11S OPTIMIZATION FOR MINUTIAE DETECTION
AHSAN UL ISLAM, WADDUWAGE SHANIKA PERERA, ERASMUS MFODWO AND VAN VUNG PHAM

PAPER ID 297: ADVANCING FIRE AND SMOKE DETECTION IN FORENSIC SURVEILLANCE: A STUDY WITH CONTEXT AWARE AUGMENTATION AND OPTIMIZATION
PAUL N. ISIBOR, PAMELA KIRUI, OSAYOMORE O. AIGBOGUN, ISAH MOHAMMED, BRIGHT JIWUEZE AND VAN VUNG PHAM

PAPER ID 377: PREDICTIVE MODELS FOR URBAN AIR QUALITY MANAGEMENT USING AI
DULARA LIYANAGE, NIMASHA VITHANAGE, IMASHA WIJEWARDANE, NIMASHA FERNANDO, DINUKA WIJENDRA AND THAMALI DASSANAYAKE

PAPER ID 389: CROSS-LANGUAGE TRANSFER LEARNING FOR VULNERABILITY DETECTION: DIRECTIONAL ASYMMETRY BETWEEN JAVA AND PYTHON
ADIBA MAHMUD, YASMEEN RAWAJFIH AND FAN WU

12:30-14:00 LUNCH AND CLOSING REMARKS

14:00-16:00 SESSION 16 PRESENTATIONS (CHAIR: SRIKUMAR NAYAK)

(WAX ACADEMIC CENTER 126 – ZOOM ONLINE)

PAPER ID 143: TRUSTFED-HE: TRUSTWORTHY FEDERATED FRAUD ANALYTICS FOR BANKS WITH HOMOMORPHIC SECURE AGGREGATION AND POISONING-RESILIENT TRAINING
SRIKUMAR NAYAK

PAPER ID 178: BHF-GUARD: BREAKING AND HARDENING FINANCIAL ML WITH ADVERSARIAL STRESS TESTS AND CERTIFIED ROBUSTNESS CHECKS
SRIKUMAR NAYAK

PAPER ID 224: ENFORCING ACCOUNTABILITY IN AUTONOMOUS OPS: A ZERO-TRUST MULTI-AGENT FRAMEWORK WITH FORENSIC REASONING LEDGERS
MADHVESH KUMAR AND DEEPIKA SINGH

PAPER ID 247: THE URL NEXTDOOR: A DIGITAL FORENSIC ANALYSIS OF NEIGHBORHOOD APPS
JOSEPH BROWN, ABDUR RAHMAN ONIK AND IBRAHIM BAGGILI

PAPER ID 267: TELEGRAM AS A TRANSFORMATIVE CRIMINAL MARKETPLACE: ANALYZING IDENTITY THEFT DYNAMICS ON SURFACE-ACCESSIBLE PLATFORMS
MANIK KAUR, DAVID MAIMON, MARIO M. KUBEK AND ANU G. BOURGEOIS

PAPER ID 348: INCIDENT-AWARE CI/CD PIPELINES: LEARNING FROM PRODUCTION FAILURES TO PREVENT CERTIFICATE ROTATION DRIFT

YOGESH S. THANVI, LAKSHMI VIDYA PERI AND YOGEESH KUNIGAL GANGAIAH

PAPER ID 351: EXTENDING THE TMMI FRAMEWORK FOR SECURE TESTING OF AI AGENTS

LAKSHMI VIDYA PERI

PAPER ID 402: BEYOND NVD: REGIONAL VULNERABILITY DATABASES FOR GLOBAL COVERAGE

SHASHANK KOGANTI AND SUBHASISH MAZUMDAR

PAPER ID 415: ENTERPRISE-GRADE AI-DRIVEN TEXT ANALYTICS AND INSIGHT EXTRACTION USING TRANSFORMER-BASED NLP MODELS

CHANDRA PRAKASH SINGH

14:00-16:00 SESSION 17 PRESENTATIONS (CHAIR: IMAN VAKILINIA)

(WAX ACADEMIC CENTER 127 – ZOOM ONLINE)

PAPER ID 119: DEEP LEARNING-BASED INTRUSION DETECTION AND CYBERSECURITY FRAMEWORK FOR CONNECTED VEHICLE CAN BUS COMMUNICATION NETWORKS

SHIVA KUMARA AND HENRY P CYRIL

PAPER ID 139 IMPROVED PROSTATE ZONAL SEGMENTATION: ADDRESSING THE PERIPHERAL ZONE CHALLENGE VIA A PERIPHERAL-CENTRIC SWINUNET

ABIDUS SATTAR AZIZ, MD MASUM RANA, YOUSUF ABDULLAH BORNA, MD REZWANUL AKTER PALLAB AND PLATO CHAKMA

PAPER ID 140: VISION TRANSFORMER FOR EPILEPSY SEIZURE PREDICTION AND DETECTION: A NARRATIVE REVIEW

MD MASUM RANA, ABIDUS SATTAR AZIZ, MD REZWANUL AKTER PALLAB AND YOUSUF ABDULLAH BORNA

PAPER ID 211: INFOTAINPC: EXPLORING THE PRIVACY CONCERNS WITHIN IN-VEHICLE INFOTAINMENT DATA SERVICES USING TAM

SAMIUL ALAM

PAPER ID 222: DATABASE FORENSICS READINESS: AN EXAMINATION OF REDIS

MUHAMMAD ABDUL MOIZ ZIA AND OLUWASOLA MARY ADEDAYO

PAPER ID 374: CONTEXTUAL REINFORCEMENT LEARNING FOR LINGUISTIC INTENT-GATED ACCESS CONTROL IN PRODUCTION AI SYSTEMS

KARTHIK PAPPU

PAPER ID 386: TEMPORAL BEHAVIORAL ARCHETYPES OF RANSOMWARE IN ACTIVE DIRECTORY ENVIRONMENTS

PRAJNA BHANDARY AND CHARLES NICHOLAS

PAPER ID 399: DARKMINE: DEEP LEARNING FOR DARK WEB THREAT INTELLIGENCE– ANONYMOUS ATTACK INFRASTRUCTURE ATTRIBUTION AND CRIMINAL ORGANIZATION DETECTION

USHA RATNAM JAMMULA AND NAGA SUJITHA VUMMANENI

PAPER ID 484: DEVSECOPS-DRIVEN SECURITY CONTROLS FOR ERP RELEASE PIPELINES

YOGEESH KUNIGAL GANGAIAH, KARTHIK PAPPU AND YOGESH S. THANVI

SESSION 18 OFF-LINE VIDEO PRESENTATIONS (CLICK FOR ACCESS THE PRESENTATIONS)

PAPER ID 9: THERMAL IMAGING AND CONVOLUTIONAL NEURAL NETS FOR ADVANCED WARNING OF FIRES ON CRITICAL DATA INFRASTRUCTURE

STEVEN MAY AND CIHAN VAROL

PAPER ID 10: VIDEO STREAMING OVER VEHICULAR AD HOC NETWORKS: A REVIEW

OMER MOHAMMED SALIH HASSAN, ISMAIL AMIN ALI AND ASAF VAROL

TABLE OF CONTENTS

PAGE

ABSTRACTS

Paper ID 1

FORENSIC INVESTIGATION OF SDRSHARP AND RTL-SDR DONGLE IN EAVESDROPPING RADIO COMMUNICATIONS

Erasmus Mfodwo; Narasimha Shashidhar; Cihan Varol; Ahmet Furkan Aydogan

Abstract- Software-defined radio applications such as SDRSharp are misused by criminals to capture, demodulate, and eavesdrop on confidential information from radars, military radio systems, satellite transmissions, local and federal police radio communications using an RTL-SDR dongle. This study examines the forensic artifacts generated by SDRSharp and RTL-SDR dongle on Windows 11 Pro, identifying traces left behind in memory and data repositories. Using Magnet AXIOM, we investigated Browser History, Hardware Device Interfaces, Windows Event Logs, Registry, AppData, Downloads, ProgramData, System32, SDRSharp Installer Folder, Prefetch, Recycle Bin, and RAM content. Our findings revealed artifacts, including web links, package installation directories, hardware interfaces, drivers, event logs, configuration files, prefetch files, and memory exhibits, which can aid digital forensic investigators in proving or disproving software-defined radio usage. This study highlights the importance of software-defined radio forensics in modern cybersecurity, law enforcement, and military security, emphasizing the need for forensic methodologies, tools, and expertise.

Paper ID 4

BICIR: A BLOCKCHAIN-BASED INTEROPERABILITY MODEL FOR CROSS-NATIONAL CYBER INCIDENTS

Osayomore O. Aigbogun; Cihan Varol

Abstract- Current approaches to international cyber incident response are hindered by some legal, technical and organizational barriers that prevent real-time cooperation across national borders. While frameworks such as the Budapest Convention and STIX/TAXII aim to standardize data sharing, they tend to lack enforceable mechanisms in terms of jurisdiction control and verifiability. This paper introduces BICIR, a blockchainbased interoperability framework using Hyperledger Fabric and a novel Selective Broadcast Algorithm (SBA) to enable metadatadriven, jurisdiction-aware incident routing. With the implementation of a testbed based on two organizations, the system demonstrated 100% routing accuracy and a 50% reduction in unnecessary broadcast overhead. Results were validated through functional endpoint tests, confirming deterministic behavior and compliance-based data dissemination at the smart contract level. Index Terms−Blockchain, Cyber incidents, Hyperledger Fabric, incident response, interoperability.

Paper ID 7

TWO PILLARS OF BANKING INTELLIGENCE: A COMPARATIVE ANALYSIS OF AI TECHNIQUES FOR FRAUD PREVENTION AND CHURN MITIGATION

Sreenivasulu Gajula

Abstract- The digital banking trend has enhanced fraud activities that are a major threat to financial institutions and their consumers. It is a continuation of a sophisticated banking fraud prevention framework based on the use of Machine Learning (ML) technologies. The suggested system uses robust ML and DL algorithms to identify and stop fraudulent transactions in real-time. Two benchmark datasets, Kaggle Credit Card Fraud Detection and IEEE-CIS Fraud Detection, are used to introduce a CNN-LSTM model and compare it against XGBoost. The CNN-LSTM model consistently beats the baseline models−K-Nearest Neighbors (KNN), Multi-Layer Perceptron (MLP), Logistic Regression (LR), and solo LSTM in every experiment. CNN-LSTM scored 99.78% and 99.87% accuracy and F1-score respectively on CCFD dataset, which is significantly higher than the KNN (88.7% accuracy) and MLP (80.38% accuracy). Likewise, CNN-LSTM achieved their highest accuracy and F1-score of 98.12% on the IEEE-CIS dataset, compared to LR (61% accuracy) and LSTM (86% accuracy). These results point to the usefulness of deep learning specifically CNN-LSTM- in the perfect identification of fraudulent transactions in heterogeneous datasets.

Paper ID 9

THERMAL IMAGING AND CONVOLUTIONAL NEURAL NETS FOR ADVANCED WARNING OF FIRES ON CRITICAL DATA INFRASTRUCTURE

Steven May; Cihan Varol

Abstract- Following a series of fire-related incidents within a corporate setting, this study investigates the application of thermal imaging and machine learning for early fire detection in information technology (IT) environments. Traditional smoke-based fire detection systems often suffer from delayed response times due to the latency in smoke diffusion, especially within server rooms where critical infrastructure requires immediate intervention. Addressing this gap, the research evaluates the use of a compact FLIR Lepton 3.5 thermal camera integrated into a Cat S62 Pro smartphone to detect hazardous heat levels in computing equipment. By training a convolutional neural network (CNN) using the GoogLeNet architecture, the system effectively differentiated between safe and unsafe thermal states based on a calibrated temperature threshold of 150°C. The CNN was trained using images of computer systems with and without a localized heating element simulating an overheated component. Results showed that the model, when trained with the Adam optimizer, achieved high accuracy in identifying potentially dangerous heat conditions, presenting a promising alternative to conventional smoke detection. This proof-of-concept highlights the potential for real-time, automated thermal anomaly detection systems tailored to server rooms, with implications for enhancing fire prevention, minimizing false alarms, and safeguarding critical IT assets.

Paper ID 10

VIDEO STREAMING OVER VEHICULAR AD HOC NETWORKS: A REVIEW

Omer Mohammed Salih Hassan; Ismail Amin Ali; Asaf Varol

Abstract- The future holds great promise for vehicular networks to reach the same level of widespread use as smartphones have today. Future vehicles will link multiple sensors to advanced onboard processing systems, enabling them to connect with surrounding vehicles and roadside infrastructure through the necessary communication channels. The research community now focuses on video streaming in VANETs because it enables safety applications, traffic management, and entertainment services. The dynamic nature of vehicular environments makes it difficult to provide video streaming services with high quality, reliability, and low latency. The survey examines modern technological developments, current obstacles, advanced solutions, and unaddressed research problems in this field through a systematic analysis.

Paper ID 16

R V F (2025): ADDRESSING THE DEFENCE OF HACKING

Junade Ali

Abstract- The defence of hacking (sometimes referred to as the "Trojan Horse Defence" or the "SODDI Defence", Some Other Dude Did It Defence) is prevalent in computer cases and a challenge for those working in the criminal justice system. Historical reviews of cases have demonstrated the defence operating to varying levels of success. However, there remains an absence in academic literature of case studies of how digital forensics investigators can address this defence, to assist courts in acquitting the innocent and convicting the guilty. This case study follows the case of R v F where a defendant asserted this defence and the author worked alongside a police investigator to investigate the merits of the defence and bring empirical evidence before the jury. As the first case study of its kind, it presents practical lessons and techniques for digital forensic investigators.

Paper ID 18

COMPARATIVE ANALYSIS OF CYBERSECURITY LAW AND WARFARE IN THE UNITED STATES (CCPA) AND EUROPE (GDPR): A PRE-2020 VS POST-2020 PERSPECTIVE

Martina Kalu; Cihan Varol

Abstract- This paper provides a comparative analysis of cybersecurity laws and warfare strategies between the United States and the European Union, with a dual emphasis on developments before and after 2020. While international

collaboration on cybersecurity has made progress in recent decades, legal frameworks remain fragmented across regions, leaving critical infrastructure exposed to transnational threats. The study begins by examining foundational cybersecurity policies, doctrines, and terminology developed prior to 2020, particularly highlighting inconsistencies in the definitions of cyberspace and cyber warfare. It then evaluates post-2020 evolutions driven by high-profile incidents, such as SolarWinds and the MOVEit breach, as well as legislative shifts like the California Privacy Rights Act (CPRA) and the EU's Digital Services Act (DSA). Through a detailed comparison of the GDPR and CCPA/CPRA, the paper identifies key divergences in enforcement models, consumer rights, and cross-border applicability. It further assesses national cybersecurity strategies, including the U.S. NIST framework and the EU's ENISA-supported initiatives. Findings suggest that while each region has developed sophisticated mechanisms to address cyber threats, the absence of standardized international legal definitions and doctrines hinders cohesive global defense. The paper concludes by advocating for deeper transatlantic collaboration, harmonized terminology, and adaptive legal frameworks to meet the challenges of an increasingly complex cyber threat landscape.

Paper ID 19

AI-Powered Android Immune System: A Hybrid Static and Semantic Model for Malware Neutralization

Mohd Fozla Rabby

Abstract- We propose an AI-augmented Android defense framework that autonomously converts malicious applications into secure, functional variants using hybrid static analysis, graph-based feature extraction, and advanced large language models (LLMs). Unlike traditional detection-only methods, this framework identifies and localizes malicious code segments, replacing them with benign equivalents. Structural features from APKs, such as AndroidManifest.xml and the complete Method Call Graph (MCG), are analyzed by Gemini 2.5 Flash LLM for semantic interpretation and neutralization. The resulting Java code is recompiled in Android Studio and repackaged into safe APKs. Focusing on APKs under 20 MB ensures efficient processing. Experimental results with real-world malware sam- ples demonstrate successful benignification while maintaining operational integrity, establishing a novel paradigm for AI- enhanced mobile security.

Paper ID 20

VANET Simulators: An Overview and Comparative Analysis

Omer Mohammed Salih Hassan; Ismail Amin Ali; Asaf Varol

Abstract- Intelligent Transportation Systems (ITS) utilize Vehicular Ad Hoc Networks (VANETs) for real-time data exchange, road-safety enhancement, and coordinated operation of connected and autonomous vehicles. The high mobility, frequent topology changes, and heterogeneous communication modes in VANETs necessitate advanced simulation environments for robust development, testing, and performance evaluation. This study reviews approximately 70 recent scholarly works to comparatively analyze five prominent VANET simulators (SUMO, OMNeT++, ns-3, Veins, and MATLAB-based platforms), assessing architecture, protocol support, scalability, and mobility integration. It identifies each tool's strengths, limitations, and appropriate use cases ranging from traffic modeling to network and cross-layer evaluation. Three key open challenges are highlighted: interoperability, more realistic mobility modeling, and AI-enabled hybrid simulation, and the findings are intended to guide the design of more realistic, flexible, and better-integrated VANET simulation frameworks.

Paper ID 23

An Information Asymmetry Game for Trigger-based DNN Model Watermarking

Chaoyue Huang, Gejian Zhao, Hanzhou Wu, Zhihua Xia and Asad Malik

Abstract- As a valuable digital product, deep neural networks (DNNs) face increasingly severe threats to the intellectual property, making it necessary to develop effective technical measures to protect them. Trigger-based watermarking methods achieve copyright protection by embedding triggers into the host DNNs. However, the attacker may remove the watermark by pruning or fine-tuning. We model this interaction as a game under conditions of information

asymmetry, namely, the defender embeds a secret watermark with private knowledge, while the attacker can only access the watermarked model and seek removal. We define strategies, costs, and utilities for both players, derive the attacker's optimal pruning budget, and establish an exponential lower bound on the accuracy of watermark detection after attack. Experimental results demonstrate the feasibility of the watermarked model, and indicate that sparse watermarking can resist removal with negligible accuracy loss. This study highlights the effectiveness of game-theoretic analysis in guiding the design of robust watermarking schemes for model copyright protection.

Paper ID 44

WEAKLY SUPERVISED KNOWLEDGE BASE CONSTRUCTION FOR TELEGRAM GIFT-CARD FRAUD MESSAGES

Chunlan Gao; Yubao Wu

Abstract- Telegram has emerged as a major platform for advertising discounted or resold digital assets such as Amazon, Vanilla, and Walmart gift cards. Many of these promotions are linked to underground commerce, coupon abuse, or the resale of compromised balances. Extracting structured financial information from such loosely formatted Telegram messages is challenging due to inconsistent layouts, frequent use of emojis, and the inherent scarcity of labeled data. To overcome this, we present an end-to-end Weakly Supervised Knowledge Base Construction (KBC) system that transforms raw Telegram chat exports into structured discount triples of the form (brand, original price, discount price, discount rate). The system integrates a tailored pattern-based candidate generation module, probabilistic weak labeling using multiple heuristic cues, and a classifier refinement step leveraging both structural and semantic features. Experimental results demonstrate that features derived purely from message structure achieve strong predictive performance (F1 = 0.89), and that the agreement between automated weak labels and human judgment reaches 80.6%. Brand-level analysis of extracted triples further reveals realistic underground pricing trends, confirming the system's ability to capture semantically meaningful financial relationships without reliance on manual annotation. This proposed framework enables the scalable population of domain-specific knowledge bases for financial-fraud analytics

Paper ID 50

OPTIMIZING RESISTIVE LOSSES IN TRANSMISSION SYSTEMS VIA VOLTAGE LEVEL SELECTION

Kendall D. Standridge-Monroe; Asaf Varol

Abstract- Transmission line efficiency is critical to support the overall power grid and the growing energy demand. This study investigates the optimal transmission line voltage selection to minimize resistive losses while balancing associated economic and physical constraints. To evaluate these trade-offs quantitatively, Microsoft Excel Solver was utilized to develop a model for analyzing power transmission across a variety of transmission powers (10-400MW) and line lengths (60-160km). Constraints were established for minimum conductor sizes allowed at different voltages. The goal of the model is to determine the voltage level that results in the lowest conductor material costs, structural costs, and long-term cost of resistive losses. Additionally, the model selects the conductor type (from available conductors) most suited for the variables at hand.

The results show that while resistive losses are reduced as voltage increases, the most economical solution is often one that balances these costs for the transmission line's distance and load. At lower to moderate power loads (10MW, 50MW, and 150MW), the model selected 230kV as the optimal voltage to minimize conductor material costs, structure costs, and resistive losses. At the highest load evaluated (400MW), the model selected 500kV as the optimal voltage to minimize overall costs. This data provides crucial guidance for transmission planning, as it demonstrates that optimization cannot rely solely on electrical efficiency and must also account for material costs and system constraints.

Paper ID 51

EXPLAINABLE CUSTOMER LIFETIME VALUE FOR PERSONALIZED ENTERPRISE RESOURCE PLANNING STRATEGY AND INTERVENTION

Vinay Singh; Prashant Gupta; DhirajKumar Pathak

Abstract- Logistics, supply chain operations, and enterprise resource planning (ERP) systems are important in the modern competitive digital economy in improving business efficiencies and customer satisfaction. As e-commerce has been developing rather quickly, customer behavior has been getting diversified and complicated, which means that precision marketing and a personal approach are necessary to enhance loyalty and profitability. The study presents the Explainable Customer Lifetime Value (XCLV) model, which aids ERP-based strategic decision-making by accurately predicting and clarifying customer value. The Olist Brazilian E-Commerce Public Dataset is preprocessed with systematic cleaning, Olist dataset Merging Schema, encoding, scaling, and sparse autoencoder-based dimensionality reduction. RFM (Recency, Frequency, Monetary) measures and K-Means clustering with optimal k = 3 are then computed and validated using Elbow and Silhouette analysis. Random Forest, LightGBM, and XGBoost machine-learning models are trained and assessed based on the R2, MSE, RMSE, and MAE metrics. The best predictive performance R2 = 95.64, MSE = 1581.48, RMSE = 39.77, MAE = 22.49 is shown in the matrix-based approach compared with the models of Random Forest, LightGBM, XGBoost, Gradient Boosting, and Decision Tree. Model explainability through LIME, along with residual analysis, is a way that the system remains understandable and trusted. Essentially, this research highlights the impact of combining RFM segmentation, deep feature compression, and explainable machine learning in a way that not only improves the precision of CLV prediction and customer strategy but also e-commerce enterprises can plan their tactics in an effective manner, save valuable customers, and provide the tailored services that are necessary for their sustained development.

Paper ID 56

SQL VS. NOSQL COMPARATIVE ANALYSIS OF DATABASE MANAGEMENT SYSTEMS FOR SCALABILITY AND PERFORMANCE

Ravi Sankar Susarla; Viswa Bharath Kolla

Abstract- As web applications and big data have grown exponentially, choosing a database system that is both scalable and highly-performing has become an important issue. Although traditional relational databases (SQL) and NoSQL databases offer a schema-less environment that enables high bulk write, but with the risk of much higher read latency. SQL databases and NoSQL databases (such as MySQL and MongoDB) have different benefits, but their relative performance at different workloads is not clear. The purpose of the study is to standardize and compare MySQL and MongoDB in a standardized MERN-stack setting in terms of latency, throughput, storage and resource efficiency. The Docker and Kubernetes were used to create a reproducible automated workflow to deploy, configure, test both databases under the same conditions and run specific workloads with the collection of detailed performance metrics. The results reveal that MySQL achieves superior performance in connection latency (0–3 ms vs. 280–495 ms for MongoDB), read operations (32 ms vs. 570 ms), and throughput (>1000 TPS vs. <100 TPS), while MongoDB excels in bulk write performance (439 ms vs. 3013 ms for MySQL) and per-record insert efficiency (0.44 ms vs. 3.01 ms). In this comparison, it is important to note the obvious trade-offs: MySQL is recommended to be used in read-heavy, low-latency, and transactional applications, whereas MongoDB is best used in write-intensive and flexible-schema applications. This discussion offers a good analysis framework to be applied in the choice of database, it is objective in performance trade-offs, and it has provided a repeatable methodology to be applied in academic and real-world applications.

Paper ID 65

A LIGHTWEIGHT HYBRID TEMPORAL CNN-TRANSFORMER ARCHITECTURE FOR REAL-TIME HEALTHCARE ANOMALY DETECTION ON WEARABLES

Akshit Naithani; Vrishin Jain

Abstract- The advancement of wearable health technology has created a critical need for real-time anomaly detection in physiological signals, such as ECG and glucose data, directly on the device. However, the computational constraints of lowpower wearables present a significant barrier to deploying powerful deep learning models. This paper proposes

a novel lightweight hybrid architecture that synergistically combines Temporal Convolutional Networks (TCNs) and a factorized selfattention mechanism to achieve efficient and accurate on-device anomaly detection. The TCN layer provides local feature extraction with a dilated receptive field, while the linear-complexity attention mechanism enables the model to capture essential longrange dependencies in time-series data. A theoretical analysis demonstrates the model's suitability for embedded deployment, with a low parameter count and computational footprint that aligns with the memory and power constraints of microcontrollers. The proposed system architecture outlines a complete pipeline for data preprocessing, on-device inference and alert generation. This work provides a theoretical framework for a new class of embedded intelligence in wearables, offering the potential for instantaneous health alerts, enhanced privacy and reliable operation independent of cloud connectivity, thereby paving the way for more accessible and proactive healthcare monitoring.

Paper ID 75

WEAK ENFORCEMENT AND LOW COMPLIANCE IN PCI DSS: A COMPARATIVE SECURITY STUDY

Soonwon Park; John D. Hastings

Abstract- Although credit and debit card data continue to be a prime target for attackers, organizational adherence to the Payment Card Industry Data Security Standard (PCI DSS) re- mains surprisingly low. Despite prior work showing that PCI DSS can reduce card fraud, only 32.4% of organizations were fully compliant in 2022, suggesting possible deficiencies in enforce- ment mechanisms. This study employs a comparative analysis (qualitative and indicator-based) to examine how enforcement mechanisms relate to implementation success in PCI DSS in relation to HIPAA, NIS2, and GDPR. The analysis reveals that PCI DSS significantly lags far behind these security frameworks and that its sanctions are orders of magnitude smaller than those under GDPR and NIS2. The findings indicate a positive association between stronger, multi-modal enforcement (including public disclosure, license actions, and imprisonment) and higher implementation rates, and highlight the structural weakness of PCI DSS's bank-dependent monitoring model. Enhanced non-monetary sanctions and the creation of an independent supervisory authority are recommended to increase transparency, reduce conflicts of interest, and improve PCI DSS compliance without discouraging card acceptance.

Paper ID 76

SAFEMEDINET: FED AI SYSTEMS FOR PRIVACY-PRESERVING THREAT DETECTION IN HEALTHCARE

Irin Sultana; Syed Mustavi Maheen; Naresh Kshetri; Shriram KS Pandian

Abstract- The explosive growth of digital healthcare data and networked Internet of Medical Things (IoMT) devices has heightened vulnerabilities inside healthcare networks, hence exposing sensitive medical systems to sophisticated cyber assaults. The safeMEDInet framework offers a secure, federated artificial intelligence (AI) architecture that allows decentralized healthcare institutions to cooperatively identify and address problems without disclosing raw patient data. safeMEDInet utilizes federated learning along with privacy-preserving techniques, such as differential privacy, homomorphic encryption, and Byzantine-resilient aggregation, to guarantee confidentiality, integrity, and adherence to regulations in remote settings. The proposed framework integrates a hybrid CNN-LSTM model for spatiotemporal intrusion detection with secure model synchronization and encrypted parameter sharing to ensure robust accuracy against various cyber-attacks, including ransomware, unauthorized access, and distributed denial-of-service (DDoS) intrusions. Empirical assessments utilizing MIMIC-IV, HealthData.gov, and WHO datasets reveal that safeMEDInet achieves a detection accuracy of 96.8% with robust privacy assurances ($\varepsilon = 1.9$) and sustains an accuracy of 88.4% despite 30% Byzantine interference, surpassing traditional federated and centralized systems. The findings confirm safeMEDInet's capacity to guarantee high detection reliability, low processing cost, and mathematical assurance of privacy resilience. This research positions safeMEDInet as a pivotal advancement towards safe, scalable, and ethically governed Healthcare 5.0 ecosystems, incorporating AI-driven privacy, federated cooperation, and blockchain-supported data integrity for next-generation medical cybersecurity.

Paper ID 84

PPAISEC: PRIVACY-PRESERVING AI MODELS IN HEALTHCARE SECURITY - AI FRAMEWORKS SYNTHESIS

Tanzina Sultana; Asura Akter Sunna; Mohammed Majbah Uddin; Naresh Kshetri

Abstract- As artificial intelligence (AI) technologies, particularly generative and collaborative learning models are increasingly integrated into healthcare and other sensitive domains, data privacy, security, and fairness concerns have grown significantly. This paper focuses on a thorough examination of current privacy-preserving AI models, including federated learning (FL), differential privacy (DP), homomorphic encryption, and generative adversarial networks (GANs). Key contributions are reviewed across recent works that explore privacy-preserving mechanisms within domains such as clinical diagnostics, drug discovery, Internet of Medical Things (IoMT), and virtual health systems. Dynamic federated models (e.g., DynamicFL) that adjust model architecture based on computational heterogeneity and encryption-augmented FL architectures are presented to maintain data locality while ensuring equitable performance. GAN-based synthetic data generators (e.g., medGAN, CorGAN) offer alternative solutions to share healthcare data without compromising patient identity and introducing new threats if misused. Across these models, a multi-phase life cycle of threats is identified—spanning data collection, model training, inference, and system integration—highlighting the importance of proactive governance. Information compliance frameworks such as the EU AI Act and the U.S. AI Bill of Rights are counting for standardizing technological implementation in healthcare data management. In this paper, the Privacy-Preserving AI (PPAI) framework is proposed, and existing AI models are examined. Several privacy-preserving layers are included in the PPAI framework, starting with a generative adversarial network for data synthesis, followed by homomorphic encryption for computation on encrypted data, differential privacy for data anonymization, and federated learning for collaborative model training without sharing raw data. Through the coordination of these multi-layered protections to tackle technical, legal, and practical concerns, PPAI improves clinical explainability and exhibits strong privacy protection.

Paper ID 85

COMPARATIVE ANALYSIS OF STEGANOGRAPHY INJECTION AND DETECTION TOOLS/METHODS: TRADITIONAL VS. AI-BASED

Rahaf Alnuaimi; Maryam Almarzooqi; Farkhund Iqbal

Abstract- The paper presents a cross-media benchmarking study of steganography injection and steganalysis tools for images, audio, and video under a unified protocol. Traditional and learning based methods are compared using shared datasets, fixed payload settings, and standardized metrics, and a weighted scoring rubric is introduced to summarize imperceptibility, robustness, security, computational cost, and detection behaviour. Results show that highly imperceptible image embedding (e.g., LSB-based tools with ≈0.01% pixel modification) can still be detected reliably, while aggressive embedding (e.g., F5 with 98.78% pixel modification) produces visible artifacts and remains detectable. For audio and video, metadata, and structure-based indicators enable reliable detection, with the combined audio detection strategy achieving 100% identification of stego files while preserving clean-file accuracy. These results support evidence-based tool selection for digital forensics triage and secure deployment under operational constraints.

Paper ID 87

DESIGN AND ANALYSIS OF ENERGY EFFICIENT ASCON 80PQ

Ananthraj Rao Kekuda, Anirudh G P, Ashwin Thomas and Srinivasa V S Sarma D

Abstract- This paper presents the energy efficient design of ASCON 80pq with performance analysis of cryptographic algorithm for lightweight, post quantum secure applications. ASCON, a lightweight cryptographic scheme, has been recently standardized by national institute of standards and technology (NIST) as part of its Lightweight Cryptography project due to its power efficiency and robust security features. This paper presents the novel secure hardware implementation of ASCON 80pq and compares with its predecessor ASCON 128, balancing power, area, and performance tradeoffs. The implementation of a Verilog RTL analyzes key performance metrics such as throughput, dynamic power consumption, and area efficiency. The power optimization technique, such as clock gating, is demonstrated to significantly reduce dynamic power while preserving performance. The results of FPGA and ASIC

implementations prove potential benefits of ASCON 80pq over its predecessor counterpart and make it more efficient cryptographic solutions for IoT systems.

Paper ID 90

DETERMINISTIC LLMS: A PRACTICAL FORENSIC FRAMEWORK FOR VERIFIABLE AND REPRODUCIBLE LOCAL LLM INFERENCE

Joel Molina

Abstract- Large Language Models (LLMs) have a growing role in legal, investigative, and enterprise workflows, yet the outputs they produce are often difficult to reproduce or verify. Minor variations in hardware, numerical precision, or runtime libraries can cause identical prompts to produce different results, undermining the admissibility of AI-generated evidence in a forensic setting. This paper presents a practical framework for bit-identical LLM inference as well as an independent verification of these AI-generated artifacts. The framework records cryptographic hashes of model artifacts and configuration fields, captures precise environmental metadata, and logs inferencetime parameters and seeds. A separate verification process instantiates a clean runtime, reloads the hashed model, recorded seed, and configuration, and performs bit-for-bit comparisons to produce robust pass/fail forensic reports. This approach is validated across multiple prompts, repeated runs, and adversarial inputs. Across 40 controlled inference and verification runs, the framework achieved a 100% bit-identical reproducibility rate on two open-weight LLM families. The results demonstrate perfect reproducibility under a controlled hardware and software environment, providing a promising groundwork for lightweight, easy-to-use system additions to strengthen the admissibility of AI-generated evidence. This paper also compares this framework with existing AI provenance methods to showcase the benefits of artifact-level determinism for forensic verification. Together, these contributions provide a promising path toward verifiable LLM inference suitable for investigative workflows and institutional deployment.

Paper ID 96

INTERPRETABLE ENSEMBLE LEARNING FOR NETWORK TRAFFIC ANOMALY DETECTION: A SHAP-BASED EXPLAINABLE AI FRAMEWORK FOR EMBEDDED SYSTEMS SECURITY

Wanru Shao

Abstract- Network security threats in embedded systems pose significant challenges to critical infrastructure protection. This paper presents a comprehensive framework combining ensemble learning methods with explainable artificial intelligence (XAI) techniques for robust anomaly detection in network traffic. We evaluate multiple machine learning models including Random Forest, Gradient Boosting, Support Vector Machines, and ensemble methods on a real-world network traffic dataset containing 19 features derived from packet-level and frequency-domain characteristics. Our experimental results demonstrate that ensemble methods achieve superior performance, with Random Forest attaining 90% accuracy and an AUC of 0.617 on validation data. Furthermore, we employ SHAP (SHapley Additive exPlanations) analysis to provide interpretable insights into model predictions, revealing that packet_count_5s, inter_arrival_time, and spectral_entropy are the most influential features for anomaly detection. The integration of XAI techniques enhances model trustworthiness and facilitates deployment in security-critical embedded systems where interpretability is paramount.

Paper ID 97

ANALYZING THE EFFECTS OF PROMPT STYLES ON LARGE LANGUAGE MODEL CHATBOT RESPONSES

Niloofar Kolahchi and Michael W. Totaro

Abstract- The use of Large Language Models (LLMs) continues to increase in educational, technical, and conversational contexts, where prompt phrasing plays a critical role in shaping the output. While prior studies primarily employ prompt engineering to optimize task performance or propose qualitative guidelines, there remains limited empirical understanding of how prompt structure and tone systematically influence the linguistic and reasoning behavior of LLM responses. This study addresses this gap by treating prompt type and tone as controlled experimental variables and quantitatively analyzing their effects. We evaluated four widely used models, such as ChatGPT, Claude,

Google Gemini, and Microsoft Copilot. We also used 108 prompts varied across six prompt types, seven tone settings, and three STEM topics. Objective text metrics were employed to analyze responses, including word count, sentiment polarity, personal pronoun frequency, readability scores, lexical diversity, and chain-of-thought (CoT) indicators. The results show that prompt types such as zero-shot CoT, few-shot, and language/text generation produced longer and reasoned outputs, while code generation yielded concise responses with minimal explanation. Tone strongly affected stylistic behavior. Friendly and role-based tones increased positivity, pronoun usage, and readability scores, while expert and formal tones produced more complex language. Overall, this work provides a systematic and cross-model characterization of prompt-driven response behavior; it offers empirical guidance for designing clearer and more accessible human–AI interactions.

Paper ID 98

GENREACHAI: AN AGENTIC GENERATIVE AI FRAMEWORK FOR AUTOMATED REACHABILITY ANALYSIS OF ENTERPRISE SOFTWARE VULNERABILITIES

Niranjan Pachaiyappan

Abstract- In this paper, we propose a novel agentic AI system, GenReachAI, that performs context-aware reachability analysis on known vulnerabilities in enterprise software ecosystems by leveraging Large Language Models (LLMs), Model Context Protocol (MCP) servers, and Integrated Development Environments (IDEs). As reported by the Open-Source Security and Risk Analysis (OSSRA), modern software development pipelines faced unprecedented challenges in 2024, with 96% of all applications containing an average of 526 open-source components, resulting in nearly 45% false positive rates in Software Composition Analysis (SCA) scanners. The system proposed in this paper aims to mitigate this critical gap by reducing false positive rates by 85% while maintaining 97% recall through three modalities: application-build time in CI/CD pipelines, real-time reachability analysis, and continuous embedded scanning in development centric IDEs. We outline a multi-agent architecture orchestrated by an AI agentic framework with a stateful workflow engine, featuring agents dedicated to software package inventory, en- vironmental context extraction, semantic reachability analysis, and continuous verification. Environmental context awareness includes Web Application Firewall (WAF) rules, Container Or- chestration (Kubernetes) network policies, and service mesh telemetry, which facilitates the accurate assessment of reachable vulnerabilities beyond traditional static analysis engines. We performed empirical validation on 500 production systems containing popular zero-day vulnerabilities, such as Log4Shell and Reach2Shell, demonstrating that 60% of instances flagged by traditional scanners were not reachable due to factors such as network isolation policies. By integrating an LLM-powered agentic engine with semantic code reasoning capabilities and streamlining data access through standardized MCP interfaces, our system enables development teams to focus remediation efforts on verified reachable and exploitable vulnerabilities.

Paper ID 101

EXPLAINABLE RISK DECISION SYSTEMS USING ARTIFICIAL INTELLIGENCE MODELS FOR PAYMENT FRAUD IDENTIFICATION WITH MITIGATION

Dilip Patel

Abstract- The persistently high-dimensional and varied data, extremely unequal class distribution, and dynamic fraudulent behavior make payment fraud detection an extremely challenging challenge. In this paper, the publicly available IEEE-CIS fraud dataset from kaggle is preprocessed extensively by removing the unique identifiers, filling in the missing values, encoding the categorical attributes, and then applying feature-selection methods to reduce the dimensionality and noise. The issue of class imbalance is resolved by using SMOTE to create synthetic minority instances, thus allowing for better learning of the rare fraud patterns. The resulting dataset is split into training and testing sets, where a stacking ensemble model combining XGBoost, LightGBM, and CatBoost is used as the main classifier. The experimental outcomes show that the stacking model proposed achieves an accuracy as high as 99.75% which is significantly different from the accuracies of the individual baseline models such as BERT (91.2%), CNN (85.40%), and AdaBoost (92%) Hence, the stacking model outperforms all individual baseline models. In order to make the model more understandable, the most significant transactional risk indicators are highlighted, and insights into the model are provided by the SHAP and LIME explainable AI algorithms. The combination of interpretability with high-performing stacking provides the integration of a trustworthy and transparent payment fraud detection system.

Paper ID 102

FROM CRACKS TO CROOKS: YOUTUBE AS A VECTOR FOR MALWARE DISTRIBUTION

Iman Vakilinia

Abstract- With billions of users and an immense volume of daily uploads, YouTube has become an attractive target for cybercriminals aiming to leverage its vast audience. The platform's openness and trustworthiness provide an ideal environment for deceptive campaigns that can operate under the radar of conventional security tools. This paper explores how cybercriminals exploit YouTube to disseminate malware, focusing on campaigns that promote free software or game cheats. It discusses deceptive video demonstrations and the techniques behind malware delivery. Additionally, the paper presents a new evasion technique that abuses YouTube's multilingual metadata capabilities to circumvent automated detection systems. Findings indicate that this method is increasingly being used in recent malicious videos to avoid detection and removal.

Paper ID 103

PROOF-BY-ASSET: A BLOCKCHAIN-BACKED MODEL FOR ASSET-BASED AUTHENTICATION

Zoe Elliott; Iman Vakilinia

Abstract- Online services are increasingly vulnerable to large-scale automated account creation, which enables spamming, resource abuse, and denial-of-service attacks. Traditional countermeasures such as CAPTCHAs, email verification, and phone number confirmation are either easily bypassed by adversaries or impose significant usability burdens on legitimate users. As automation tools become cheaper and more accessible, these defenses fail to maintain an effective balance between security, cost, and user experience. To address this challenge, we propose an asset-based authentication framework that leverages blockchain technology to economically deter fraudulent or large-scale automated account creation. Our design introduces two complementary models: an off-chain scheme, where users prove asset ownership without locking funds, and an on-chain scheme, where users temporarily deposit cryptocurrency into a smart contract as collateral for authentication. To further enhance user privacy, we extend our model with a zero-knowledge proof-based scheme that allows users to demonstrate sufficient asset ownership without revealing their blockchain address, balance, or transaction history.

We present the architecture and workflow of these authentication models and analyze their security and privacy properties. Furthermore, we implement a prototype demonstrating the feasibility and practicality of our proposed framework. Evaluation results show that our asset-based authentication approach provides a robust, privacy-preserving, and economically grounded alternative for securing online services against automated and malicious activities.

Paper ID 107

STABILITY ANALYSIS OF SYNCHRONOUS GENERATOR SYSTEMS UNDER VARIABLE LOAD CONDITIONS

Brenden Lippard; Chase Guttu; Asaf Varol

Abstract- Instability in the modern-day power system can lead to disruptions that can have both severe economic and fiscal problems. This paper presents a brief analysis of synchronous generator stability under variable load conditions. The study demonstrates our testing with both MATLAB/Simulink and Python-based simulations and environments to analyze dynamic behavior, signal stability, and transient response. The single-machine infinite-bus (SMIB) model is used as the home for evaluating "damping, inertia", and "excitation control effects". Eigenvalue analysis and "time-domain simulation" results confirm the reliability of open-source simulation tools as educational and analytical resources for power system dynamics.

Paper ID 108

A Systematization of Privacy Threats and Defenses in Modern Agentic Web Browsers

Niranjan Pachaiyappan

Abstract- Agentic web browsers, which leverage cloud-based large language models (LLMs) to automate user tasks, represent a fundamental paradigm shift in web interaction. While promising unprecedented functionality and societal benefits, their architectural design introduces systemic threats to user privacy that move beyond traditional web tracking. This Systematization of Knowledge (SoK) paper delineates the novel privacy risks inherent in current agentic browser architectures, including data collection mechanisms, regulatory compliance gaps, and advanced attack scenarios. We structure our analysis into a taxonomy of privacy violations and formalize the threat model adapting established cybersecurity frameworks (STRIDE, MAESTRO). To address these threats, we propose a multi-layered defense strategy combining AI-driven guards with cryptographic solutions. We analyze the state-of-the-art in Fully Homomorphic Encryption (FHE) for LLM inference and detail practical implementations of on-device PII redaction and behavioral anomaly detection. Furthermore, we examine the regulatory landscape (GDPR, CCPA) and its implications for agentic systems. Our evaluation framework introduces novel privacy metrics, and we provide an implementation roadmap grounded in recent industry case studies. Finally, we identify critical open research problems to guide development of secure, privacy-preserving agentic systems.

Paper ID 113

Fragility of Children's Online Privacy

Abdulbast Abushgra, Adam Keeling, Jesse Caudell, Ada Johnson, Jayden Bowman, and Magan Miller

Abstract- In today's digital environment, the privacy and data security of children have become increasingly critical as internet usage among minors continues to grow. Protecting children from the risks embedded in online platforms is essential, yet many parents remain insufficiently informed about these threats. This study integrates findings from the London School of Economics, McAfee, and King's College London to evaluate both children's and parents' awareness of online risks and their understanding of data protection practices. We also examine U.S. state and federal legislation intended to hold social media platforms, websites, and applications accountable for safeguarding minors' privacy. Results indicate that while children possess a basic awareness of online risks, they lack understanding of more complex data collection methods and privacy safeguards. Moreover, U.S. legal protections remain inadequate, with coverage limited to children under 13 and enforcement mechanisms proving weak. To address these shortcomings, we recommend revising age thresholds in existing legislation, strengthening enforcement, and developing international frameworks to ensure consistent privacy protections for minors worldwide. Additionally, we propose mandating parental control features on internet-enabled devices at the point of purchase.

Paper ID 114

From Monolithic Middleware to Cloud-Native Microservices: A Performance-Driven Modernization Study

Sauhard Bhatt

Abstract- In this work, we examine the important evolution from monolithic legacy middleware to modern, scalable microservice architectures. As organizations increasingly pursue digital transformation, legacy middleware can become a bottleneck, slowing agility and integration. This investigation concerns a systematic migration scheme that can be employed to minimize disruption and maximize system effectiveness. We used a sample of 441 system interaction activities from a simulated real enterprise to investigate performance statistics before and after the modernization process. We use Docker for containerization, Kubernetes for orchestration, and Apache Kafka for event streaming as our main technologies, along with in-house Python scripts for performance telemetry. Measure the effects of modernization. The research considers latency, throughput, and error rates to assess the benefits of modernization. Our results show that message processing latency decreases dramatically and system throughput improves significantly after migration. With a thorough modernization program, companies can increase resilience, scalability, and maintainability. We offer a map for enterprise architects who want to breathe new life into their ageing legacy with minimal disruption, backed by solid empirical evidence that introduces measurable improvements through targeted architectural refactoring.

Paper ID 116

ACHIEVING 60% DIESEL REDUCTION AT MUSIC FESTIVALS THROUGH INTELLIGENT LOAD SEGMENTATION

Robert Owens; Asaf Varol

Abstract- Music festivals typically rely on 50+ diesel generators producing over 1,000 tons of CO2 annually. Prior attempts at battery energy storage system (BESS) implementation have failed due to fundamental power quality incompatibilities. This paper presents a practical hybrid framework achieving 60% diesel reduction using intelligent load segmentation. The research identifies two root causes for BESS failures in festival applications: (1) professional audio systems require Total Harmonic Distortion (THD) below 0.1%, while inverters achieve only 2%—a 20x gap; and (2) bass transients create 10-15x power spikes that trigger inverter protection systems. The solution is load-segmented hybrid operation where BESS handles steady-state applications (production village, camping, vendors) while diesel manages transient-intensive loads (audio amplification, stage lighting, video). Three scaled configurations were developed for festivals of 10K-100K attendees, each achieving 60% diesel reduction while maintaining >99.5% reliability. A detailed analysis of Config B (20-50K attendees) demonstrates 60.7% reduction in diesel consumption, from 41,888L to 16,436L for a 9-day event. The L-Acoustics L2 audio system provides an additional efficiency advantage: 24.7% less power consumption with +6dB SPL compared to legacy K1 systems. The framework enables 2–3-year payback through combined fuel savings and premium service revenue (powered camping, EV charging), achieving 67.3% IRR without subsidies. This approach provides festival organizers with a realistic pathway to substantial emission reductions while improving economics.

Paper ID 119

DEEP LEARNING-BASED INTRUSION DETECTION AND CYBERSECURITY FRAMEWORK FOR CONNECTED VEHICLE CAN BUS COMMUNICATION NETWORKS

Shiva Kumara; Henry p Cyril

Abstract- Modern connected and autonomous cars are more dependent on Controller Area Network Communication (CAN) for operation, which makes them more vulnerable to cyberattacks that might compromise their vital safety features. This study presents a deep learning-based IDS capable of detecting a huge range of CAN bus attacks with excellent accuracy and operational effectiveness. A number of techniques were used in combination to preprocess the OTIDS dataset (enhanced feature engineering, temporal analysis of network communications, and balanced classes) to improve the discriminability of these models. To evaluate model performance, it created and tested three architectures: a Deep Neural Network (DNN), a 1D Convolutional Neural Network (CNN), and a hybrid CNN–Bidirectional Long Short-Term Memory (BiLSTM) Network. Results demonstrated that the best network, a hybrid CNN-BiLSTM Network, achieved 90.18 percent accuracy, indicating that it successfully captured spatial and temporal relationships among CAN-related network data. A Real-time Cybersecurity Framework was implemented to provide real-time threat prediction, threat-level assessment, and analytical visualization. This proposed model has demonstrated great promise for adoption within Cyber Security Infrastructures for Intelligent Vehicles.

Paper ID 120

THORAX DISEASE CLASSIFICATION BASED ON SPATIAL TRANSFORMER NETWORKS AND SQUEEZE-AND-EXCITATION BLOCKS

K Jenni; Nala AlAhmari; G Aniruth; Msrinivas

Abstract- Pricise and early diagnosis of thorax diseases is key to the patient safety. Still, the traditional chest X-ray reading procedures with radiology are time consuming and can make manual errors. On time recognition of thorax diseases can save a patient in healthier way. The proposed work offers a solution with a new deep learning system called ThoraxNext. ThoraxNext uses a special type of neural network based on ConvNeXt to identify significant features, clues in the X-ray images. It also applies Spatial Transformer Networs STNs) to extract on cruicial points in the chest X-ray that are close to hold information about thorax illnesses. In addition, Squeeze-and-Excitation (SE) blocks are also integrated to capture the different patterns of thorax disease in the X-ray images and enhancing model's ability. The hybrid feature of integrating STNs and SE blocks helps system to detect different patterns of thorax diseases more accurately.

Paper ID 121

AN EFFECTIVE SYSTEM FOR MEDICAL IMAGE DIAGNOSIS USING DEEP CONVOLUTIONAL NETWORKS (CNNS) IN HEALTHCARE SECTOR

Sanjoy Mukherjee

Abstract- The great majority of avoidable cases of blindness worldwide are caused by Diabetic Retinopathy (DR), a degenerative eye disease. The early detection by means of automated evaluation of retinal pictures is really helpful in getting the patients' outcomes better and decreasing the clinical burden. This paper describes an innovative efficient and high-quality medical image classification system using MESSIDOR-2 dataset, based on deep learning. The procedure incorporates differing image preparation methods such as CLAHE (Contrast Limited Adaptive Histogram Equalization) which is used to boost the contrast, Gaussian filtering (noise reduction), picture cropping (region refinement) and color normalization (uniform lighting). Data augmentation was also employed to diversify the dataset by flipping, rotating and changing the brightness of the images to ensure that the model was not overfit. The models proposed to be used are the Hybrid CNN-XGBNet and the Inception V3 whose performance is compared with the simple models that include DenseNet201, IncRes-v2-1FT, VGG-16, ResNet-50, Random Forest and Decision Tree in the DR classification. The Hybrid CNN-XGBNet model that uses CNN feature extraction and XGBoost as a classifier achieves the highest accuracy, 96.94, and the optimal accuracy, precision, recall, and F1-score of 99.0. The findings support the excellent generalizability, diagnostic accuracy, and strength of the hybrid model, demonstrating its ability to perform AI-assisted medical image diagnosis in humans such as the detection and screening of DR in real-time.

Paper ID 132

AI-DRIVEN SELF-HEALING AND INTELLIGENT QUEUING THROUGH ANOMALY DETECTION IN 5G CELLULAR NETWORKS

Henry Cyril

Abstract- The growing complexity and openness of 5G Open Radio Access Network (O-RAN) architectures represent important challenges in the process of providing reliable, secure, and low-latency network operations. The traditional threshold-based monitoring and statistical traffic control systems cannot cope with dynamic anomalies and performance degradation. In this paper, the authors suggest an AI-based framework that combines the use of supervised anomaly detection, autonomous self-healing, and intelligent priority-based scheduling to improve the resilience of the 5G O-RAN networks. The NETSLAB 5G O-RAN Intrusion Detection Dataset is used to predict anomalous network traffic using the Random Forest and XGBoost models. Identified abnormalities are further categorized to invoke automated self-remedial measures, such as rate limiting, connection termination, firewall enforcement and generation of alerts. A smart queuing system dynamically assigns an importance to traffic in a network with regard to confidence in an anomaly and the nature of services. According to the results of the experiments, the XGBoost model shows better performance, which is characterized by an accuracy of 98.23, a recall of 99.66, and an ROC-AUC of 0.996, and confirms the usefulness of the suggested framework in the autonomous and adaptive management of 5G networks.

Paper ID 139

IMPROVED PROSTATE ZONAL SEGMENTATION: ADDRESSING THE PERIPHERAL ZONE CHALLENGE VIA A PERIPHERAL-CENTRIC SWINUNET

Abidus Sattar Aziz; Md Masum Rana; Yousuf Abdullah Borna; Md Rezwanul Akter Pallab; Plato Chakma

Abstract- Globally, prostate cancer is the second most widespread cancer among men, underscoring the significance of accurate prostate segmentation in early detection, staging, and treatment planning. Automating the segmentation of these areas remains difficult because of subtle tissue contrasts and irregular boundaries. Most importantly, segmenting peripheral zone is always a challenge due to most of the cancers develop near peripheral zone. This study presents a SwinUNet architecture, designed with a focus on peripheral, to address the challenges of prostate zonal segmentation. Our approach features several novel components, including patch extraction centered on peripheral to guarantee comprehensive visualisation of prostate anatomy, and self-attention mechanisms in the bottleneck layer to capture vital long-range spatial connections necessary for identifying zone boundaries. We also applied volume calibration post-processing to correct systematic estimation biases without compromising segmentation accuracy. Previous research has

frequently en- countered challenges in attaining high accuracy for prostate zonal segmentation, with the peripheral zone being especially troublesome. The majority of existing U-Net based models obtain Dice scores in the peripheral zone that are below 80%, which hampers their clinical utility. Existing methods are substantially outperformed by our approach. The proposed peripheral-centric SwinUNet surpassed traditional U-Net architectures with a pe- ripheral zone Dice coefficient of 0.82 (82%) and a transition zone Dice coefficient of 0.94 (94%). Analysis of the qualitative data showed improved definition of boundaries and a decrease in incorrect positive areas, and attention visualization confirmed that the model concentrates effectively on anatomically relevant structures during segmentation. These results establish a new benchmark for prostate zonal segmentation, with direct clinical implications for enhanced cancer risk assessment, more accurate treatment planning, and automated diagnostic workflows in prostate healthcare.

Paper ID 140

VISION TRANSFORMER FOR EPILEPSY SEIZURE PREDICTION AND DETECTION: A NARRATIVE REVIEW

Md Masum Rana; Abidus Sattar Aziz; Md Rezwanul Akter Pallab; Yousuf Abdullah Borna

Abstract- Epilepsy could be a constant neurological disease stamped by recurrent seizures. World Health Organization (WHO) stated that around 50 million people in the world are affected by epileptic seizure. Affective seizure prediction and discovery are vital for overseeing epilepsy and improving the well-being of those affected. While different techniques have been recommended for this disease, there's a developing demand for tackling the capabilities of vision transformers. Vision transformers, which is a variation of artificial neural networks, have displayed promising results in areas like image classification and object detection in computer vision. As of now, Research is conducted in exploring their potential for predicting and detecting epilepsy seizures. Typically, due to their capacity to capture spatial and temporal designs inside electroencephalogram data. This review points to carefully examine the recent research on utilizing vision transformers for epilepsy seizure prediction and detection. The measure for selecting relevant studies is whether the paper is related to Epilepsy prediction or detection utilizing EEG information or not. Through a comprehensive assessment of relevant works, we dig into the strategies, performance metrics, preferences, and limitations in existing vision transformer models in this setting. We hope that by tackling the capabilities of vision transformers, it is possible to develop more exact and reliable models for estimating and recognizing epileptic seizures. We hope our work will further help researchers in their studies to find better solutions.

Paper ID 141

COMPUTE-OPTIMAL RESOURCE ALLOCATION FOR DISTRIBUTED LARGE LANGUAGE MODEL INFERENCE IN CLOUD-SCALE INTELLIGENT SYSTEMS

Tejas Pravinbhai Patel, Sandeep Shivam, Viswanathan Ranganathan

Abstract- Test-time compute scaling enables LLMs to improve reasoning by allocating additional inference resources. While techniques like best-of-N sampling and tree search show remarkable gains, existing work focuses on single-GPU scenarios, neglecting distributed deployment challenges. We introduce AdaptiScale, a framework for compute-optimal resource allocation in distributed LLM serving with test-time scaling. Our approach combines: (1) a lightweight difficulty estimator achieving 84% accuracy with <2ms latency, (2) hierarchical scheduling balancing parallel sampling across nodes, and (3) elastic scaling adapting cluster topology to workload patterns. Evaluation on GSM8K, MATH, and HumanEval using Mistral-7B on NVIDIA A4000 GPUs shows AdaptiScale achieves 20% higher throughput and 2.1× better cost-efficiency versus static best-of-16 baseline. We achieve 67.2% accuracy on GSM8K (vs 65.1% static baseline) while reducing cost from $0.85 to $0.52 per 1000 requests. Simulated 4-node cluster deployment demonstrates 90% scaling efficiency, validating production viability for cost-effective reasoning-enhanced LLM deployment.

Paper ID 142

EXPLAINABLE AND HUMAN-CENTRIC APPROACHES IN AUDIO STEGANALYSIS: A REVIEW

Sarah Rahim and Guhanathan Poravi

Abstract- Audio steganalysis, the detection of hidden information in audio signals, is increasingly vital for cybersecurity, digital forensics, and malware monitoring. Although machine learning and deep learning have improved detection accuracy, most methods remain opaque, offering limited interpretability and minimal support for human users. This systematic literature review examines explainable and human-centric audio steganalysis, structured around three questions: where do explanations originate, how are explanations presented and who consumes them. We synthesize 26 peer-reviewed studies, finding that most explanations are model-centric, visualization-heavy, and rarely validated with end users. Emerging hybrid approaches combining signal processing with interpretable AI show promise but remain limited. This paper makes three contributions: an explainability-centric taxonomy of audio steganalysis approaches, mapping human roles to explanation types, and identification of research gaps with a forward-looking agenda for task-aware, user-evaluated, and legally robust explainable systems. These findings provide a foundation for designing transparent, trustworthy, and operationally relevant audio steganalysis frameworks.

Paper ID 143

TRUSTFED-HE: TRUSTWORTHY FEDERATED FRAUD ANALYTICS FOR BANKS WITH HOMOMORPHIC SECURE AGGREGATION AND POISONING-RESILIENT TRAINING

Srikumar Nayak

Abstract- Banks can detect fraud more accurately when they learn from patterns observed across institutions, but direct data sharing is restricted by privacy rules and security policies. Federated learning (FL) helps by training a shared model without moving raw data, yet two practical issues remain: (i) client updates can leak sensitive information if aggregation is not protected, and (ii) FL can fail when some clients are malicious (Byzantine behavior or poisoning). To address these problems, we propose TrustFed-HE, a trustworthy federated analytics pipeline that combines multiparty homomorphic secure aggregation with robust optimization and poisoning-resistant training. Each bank trains locally, applies update clipping and anomaly screening, and encrypts updates for secure aggregation, while a proximal stabilization term improves convergence under non-IID bank partitions. Experiments on the Bank Account Fraud Dataset (NeurIPS 2022) show that TrustFed-HE achieves strong clean performance (AUC-ROC 0.958, AUC-PR 0.914, Recall@1%FPR 0.742) and remains reliable under attacks: with 20% malicious clients using label-flip poisoning, AUC-PR drops only from 0.914 to 0.882, while standard FedAvg drops to 0.813. Under Byzantine sign-flip updates (20% clients), TrustFed-HE retains Recall@1%FPR 0.706 versus 0.628 for FedAvg. These results indicate that secure aggregation alone is not enough; robustness controls must be integrated into the training loop to make crossbank fraud analytics safe and dependable.

Paper ID 146

ENHANCING EARLY DIABETES SCREENING THROUGH MACHINE LEARNING AND EXPLAINABLE AI

Jahnavi Anilkumar Kachhia

Abstract- Early diabetes diagnosis is critical to minimize complications in the long term and enhance patient outcomes, but the traditional diagnostic tools tend to diagnose the disease only when serious metabolic alterations have taken place. This paper discusses the research problem which is whether ML models with Explainable AI can deliver accurate and clinically interpretable screening of early diabetes. Based on publicly available Pima Indians Diabetes Dataset, the extensive preprocessing was implemented, such as outlier removal, min-max normalizing, and balancing classes based on SMOTE. A hybrid learning model consisting of the CatBoost, Random Forest, and Decision Tree classifiers with soft voting was created to increase the predictive robustness. The results obtained by the use of accuracy, precision, recall, and F1-score were evaluated as the model performance and this resulted in balanced and reliable results. SHAP and LIME were utilized to describe both the importance of global features and individual predictions in order to support clinical trust and transparency. The findings demonstrate that Explainable AI-powered ensemble learning may effectively aid in the early diagnosis of diabetes. Nevertheless, it is noted that there are limitations, including the bias

in the demographics of the data set that were used and the possibility of generalization, which need to be validated on larger and more real-world populations.

Paper ID 148

SECURE AND PRIVACY-PRESERVING HEALTHCARE FEDERATED LEARNING VIA DIFFERENTIAL PRIVACY MECHANISMS

Maunik K Shah and Munir Rajesh Mehta

Abstract- Diabetes is considered to be one of the most important healthcare issues of their time as it affects over 30 million individuals in the US and costs an annual healthcare more than 327 billion. The rising number of diabetes patients is also accompanied by the increase in diabetes hospital readmissions that adversely influences patient health, health care spending and hospital productivity. The study employs a significant clinical database drawn out of 130 hospitals within the United States to study the problem of the identification of diabetes and early readmission. An 80/20 train-test split was conducted after thorough data treatment which included label encoding, SMOTE balancing and feature standardization. Centralized deep neural network models and federated learning frameworks were both utilized, with the latter having a differential privacy-enhanced variant to protect patient confidentiality. Precision (Prec), recall (Rec), accuracy (Acc), and F1-score were the measures utilized to assess the model's performance. Centralized learning reached the peak in terms of Acc (86.81%) and F1 score (86.75%), whereas federated learning without privacy reached an accuracy of 84.47%, and privacy-preserving federated learning held a competitive 83.02%. The outcomes exceeded those obtained using classical approaches like Random Forest, Logistic Regression, K-Nearest Neighbor, and Support Vector Machine. The importance of this study is that it shows privacy-preserving federated learning as a good, large-scale solution for joint medical research and trials, besides being able to manage the privacy of sensitive diabetic patient data and predicting performance which is the main aspect of healthcare areas such as early readmission interventions.

Paper ID 156

INTELLIGENT DATA INTEGRATION AND ANALYTICS FOR AUTOMOTIVE AFTERMARKET RETAIL: AN INFORMATICA-BIGQUERY-TABLEAU FRAMEWORK

Rambabu Tangirala

Abstract- The automotive parts retail industry faces challenges in managing large-scale, distributed data ecosystems. This research presents an AI-driven enterprise data management framework using Informatica, Google Cloud BigQuery, and Tableau. Our case study examines a global automotive parts provider with 10,000+ retail locations processing 8-10 million daily transactions. The architecture leverages artificial intelligence for real-time synchronization, predictive inventory optimization, and automated data quality management. Implementation results show data processing latency reduced from 24 hours to 15 minutes (99% improvement), query response time decreased from 45 seconds to 2.3 seconds (94.9% improvement), and data quality score improved from 78% to 96%. The framework incorporates federated learning for privacy-preserving analytics while maintaining GDPR and CCPA compliance.

Paper ID 157

DISTRIBUTED ADAPTIVE SPECULATIVE DECODING: ACCELERATING LARGE LANGUAGE MODEL INFERENCE WITH CONTEXT-AWARE DRAFT SELECTION

Tejas Pravinbhai Patel

Abstract- Large Language Models (LLMs) have revolutionized natural language processing, yet their autoregressive generation process remains a critical bottleneck for real-time applications. Speculative decoding has emerged as a promising approach to accelerate inference by leveraging smaller draft models to predict future tokens, which are then verified in parallel by the target model. However, existing methods rely on fixed draft models and single-node execution, limiting their effectiveness across diverse input distributions and scalability in production environments.

We propose Distributed Adaptive Speculative Decoding (DASD), a novel framework that introduces three key innovations: (1) context-aware draft model selection using a lightweight routing network that dynamically chooses optimal draft models based on input characteristics, (2) asynchronous distributed speculation that pipelines draft generation and target verification across multiple GPUs with fault-tolerant rollback mechanisms, and (3) adaptive speculation length that adjusts the number of speculative tokens based on real-time acceptance rates. Through extensive experiments on NVIDIA A4000 GPUs using Llama-3-8B and multiple draft models, we demonstrate that DASD achieves 2.6–2.8× speedup compared to standard autoregressive decoding and 15–20% improvement over fixed-strategy speculative decoding, while maintaining identical output quality. Our approach shows consistent gains across code generation (76% acceptance), factual QA (72% acceptance), and conversational tasks (68% acceptance). The proposed distributed protocol reduces end-to-end latency by 42% in multi-GPU settings and provides graceful degradation under GPU failures. We release our implementation and comprehensive benchmarks to facilitate future research in efficient LLM serving.

Paper ID 158

MODERNIZING SAP BUSINESS OBJECTS USING ENTERPRISE DATA WAREHOUSING PRINCIPLES

Prasanth Sathyapalan

Abstract- In the past, most SAP BusinessObjects (SAP BO) installations were founded on transactional OLTP (Online Transaction Processing)-based schemas, and, over time, the universes have been modified to meet the short-term reporting requirements. This paper analyzes a warehouse-first modernization strategy that transforms the underlying data warehouse through dimensional modelling, with centralized fact table, conformed dimensions and surrogate keys, before the BI tools are rolled out. Two analytical models, legacy OLTP and warehouse-aligned star schema were implemented and compared in Power BI using same dashboards and measures based on the Global Superstore data. The evaluation criteria of performance and stability were as follows: query join complexity, refresh time, visual rendering time, analytical consistency, and high support of advanced time-intelligence calculations. Findings show that star schema decreases the complexity of joins by approximately 60, improves the refresh performance by 30-50 and even take shorter time to visualize than the OLTP model. A centralized measure set and a separate date dimension allow reusable and consistent time-intelligence calculations, whereas the warehouse-first strategy simplifies universes, makes them more maintainable, and provides uniform implementation of business logic and row-level security across platforms. The results of this study prove that modernization of the warehouse-first is highly effective in improving the performance, reliability and stability of the BI.

Paper ID 160

AI-ENHANCED DIGITAL FORENSICS: A RESEARCH VISION FOR TRUSTWORTHY, EXPLAINABLE, AND HUMAN-CENTERED FORENSIC INTELLIGENCE

Nadeem Daudpota

Abstract- The rapid growth of digital data, the increasing sophistication of cyber threats, and the widespread reliance on digital evidence have placed unprecedented demands on digital forensic investigations [1], [2]. Traditional forensic workflows, largely manual, time-intensive, and investigator-driven struggle to scale in environments involving terabytes of heterogeneous data, cloud-native infrastructures, and fileless attacks [1], [3]. At the same time, recent advances in Artificial Intelligence (AI), Machine Learning (ML), and Large Language Models (LLMs) offer powerful opportunities to augment forensic analysis. However, uncritical adoption of AI raises concerns related to explainability, reproducibility, legal admissibility, and ethical use [5], [6].

This paper presents a conceptual research vision and framework for AI-Enhanced Digital Forensics, outlining the foundational goals, design principles, and research directions of the proposed AI-Enhanced Digital Forensics Lab (AEDF-Lab). Rather than replacing human expertise, the proposed approach emphasizes human-centered, explainable, and legally defensible AI that supports forensic triage, analysis, and reporting. The paper introduces three core research directions: (i) AI-assisted forensic triage and artifact prioritization, (ii) explainable AI frameworks for trustworthy forensic decision-making, and (iii) LLM-augmented forensic reporting and chain-of-custody automation. This paper discusses the intellectual merit, anticipated research outcomes, and broader impacts of this initiative, positioning AEDF-Lab as a platform for advancing AI-enhanced digital forensics through a principled research vision.

Unlike existing AI-driven forensic approaches that emphasize automation, the proposed research explicitly prioritizes explainability, reproducibility, and legal defensibility, ensuring alignment with evidentiary and ethical standards.

Paper ID 165

FAKE-OT: AN OPEN ARCHITECTURE FOR THREAT DETECTION IN OT ENVIRONMENTS USING MQTT-BASED HONEYPOTS

Hebert Silva; Felipe Carvalho; Tiago Demay

Abstract- Operational Technology (OT) systems are increasingly interconnected with Information Technology (IT) and cloud services, expanding the attack surface of industrial environments. Lightweight protocols such as MQTT, widely used in Industrial IoT (IIoT) and Cyber-Physical Systems (CPS), lack native security mechanisms, making them vulnerable to reconnaissance, injection, and persistence attacks. This paper presents FakeOT, an open and modular architecture for proactive threat detection in OT environments using MQTT-based honeypots. The architecture integrates deception, data normalization, and analytical correlation across four layers to enable early detection of adversarial behavior and anomalous traffic. A virtualized testbed based on Proxmox, Docker, IPFire, Mosquitto, and Wazuh demonstrates the feasibility of the approach, capturing simulated attack scenarios with low latency and high reliability. Ongoing experimental validation includes planned deployment in a controlled Internet-facing segment to collect real-world telemetry and assess resilience against large-scale bot and scanner activity. The results indicate that FakeOT provides a reproducible, extensible, and scalable framework for cyber defense and intelligence in industrial environments.

Paper ID 166

NORMALIZING FLOW FOR FINANCIAL ANOMALY DETECTION: THE IMPACT OF LOSS FUNCTION

Amir Rashid; Ruobin Qi; Rashida Hasan

Abstract- Financial fraud detection faces persistent challenges due to extreme class imbalance, high dimensionality, and evolving fraudulent behaviors in transaction data. Normalizing flow, a flexible class of generative models, shows promise for anomaly detection but is sensitive to the choice of loss function. This paper systematically evaluates 11 loss function variants, including negative log-likelihood, entropy, variational information bottleneck, spectral, and ensemble variants within a unified experimental framework. Experiments span three financial fraud datasets: Credit Card transactions, PaySim mobile money simulator, and vehicle insurance claims. Results reveal dataset-specific sensitivity patterns where loss function impact varies significantly across real versus synthetic datasets, with improvements ranging from 3.4% to 7.6% in F1 score over baseline. While no single loss function universally outperforms others, simpler regularization techniques often match or exceed complex ensemble variants, providing valuable guidance for optimizing normalizing flow models in financial fraud detection.

Paper ID 167

OPERATIONALIZING DATA ECONOMY IN DATA SPACES: A TOKEN-BASED REFERENCE IMPLEMENTATION

Muhamed Turkanović; Martin Ferenec; Teo Lah; Vid Keršič

Abstract- Data spaces enable trusted, sovereign, and interoperable data sharing among autonomous organizations, yet practical support for sustainable data economy mechanisms remains limited. While existing data space frameworks provide policy-based access and usage control, pricing, payment, and accountable data monetization are yet unoperationalized. This paper addresses this gap by presenting a token-based approach for operationalizing the data economy in data spaces. Building on Eclipse Dataspace Components and our Web3-enabled extension, DSX Engine, we introduce an EVM-compatible blockchain infrastructure that integrates decentralized discovery, ERC–20–based tokenization, and smart–contract–driven pricing, escrow, and clearing mechanisms. The solution supports multiple pricing models and is implemented as a reference architecture. Validation in the DIH Agrifood Data Space demonstrates the feasibility of monetizing datasets under machine-enforceable policies, providing a concrete foundation for economically sustainable data spaces.

Paper ID 169

A Formal Model of Access Token Bifurcation and Integrity-Based Isolation in Windows Kernel Architectures

Jean Rosemond Dora; Ladislav Hluchy

Abstract- Access control in modern operating systems (OS) is managed by complex kernel-mode structures that mediate between user intent and hardware execution. This research paper provides a formal analysis of Access Token Theory (ATT), specifically investigating the mechanics such as "Windows Mandatory Integrity Control (MIC)" and the "User Account Control (UAC)" subsystem. We formalize the "Split-Token" administrative model using a lattice-based security approach and present a deterministic algorithm for token generation. Moreover, we provide a comparative analysis of privilege sets (SePrivileges) and their impact on the system's attack surface. Our findings suggest that while kernel-resident tokens offer strong isolation, the transition states between integrity levels remain the core vector for exploitation.

Paper ID 170

Windows Operating System Credentials - Elevation with Impersonation

Jean Rosemond Dora; Ladislav Hluchy

Abstract- Privilege escalation (PE) within the Windows ecosystem often circumvents traditional User Account Control (UAC) boundaries by exploiting specific security tokens. This paper identifies, analyzes, and formalizes nine critical Windows privileges—termed the "Impersonation Enablers"—that permit a medium-integrity process to transition deterministically to SYSTEM-level integrity. We provide a comparative analysis of these vectors and propose a state-transition model for token theft. Furthermore, we evaluate contemporary Attack Surface Reduction (ASR) mitigation techniques in Windows 11 operating system (OS) and Server 2025 environments, as it boasts complex features that enhance security, flexibility, and performance. We focus more on the SeImpersonatePrivilege privilege, as it allows us to impersonate any token for which we can obtain a reference. It is interesting since the built-in LocalService, Network Service, and the default IIS accounts have it assigned natively (by default).

Paper ID 171

Security Architecture of Kerberos Authentication and Domain Credential Management

Jean Rosemond Dora; Ladislav Hluchy

Abstract- As distributed computing environments transition toward zero-trust architectures, the Kerberos protocol remains the cornerstone of Windows Domain authentication and enterprise Identity and Access Management (IAM). Despite its advanced design based on symmetric-key cryptography, robust and modern attack vectors—including Golden Ticket attacks, Kerberoasting, and AS-REP roasting—have exposed systemic weaknesses in credential handling. This research provides a comprehensive analysis of the Kerberos V5 exchange process. We present a novel mathematical model to quantify the entropy requirements of session keys and propose an expanded Kerberos model utilizing Elliptic Curve Cryptography (ECC), a mathematical concept, to mitigate offline brute-force attacks. Our comparative analysis demonstrates that the proposed framework reduces computational overhead by ≈ 18% while significantly increasing resistance to ticket falsification (forgery).

Paper ID 173

A Comparative Analysis of URL Features for Machine Learning Phishing Detection

Chukwunalu Asuai; Yusuf Moshood

Abstract- Phishing attacks exploit user trust to steal credentials and compromise systems, with blacklist-based defenses struggling to keep pace with rapidly evolving threats. Machine learning offers proactive detection, but existing research shows inconsistent results due to varying feature selection and model choices. In this work, we systematically evaluate

20 binary URL features and six machine learning classifiers using 4,200 verified websites from PhishTank and the University of New Brunswick. The features span lexical patterns, HTML characteristics, and domain properties. Results demonstrate that ensemble methods achieve 91.67% accuracy with 92.74% precision, significantly out performing traditional classifiers at 87-88% accuracy. Statistical validation through 10-fold cross-validation (91.52±0.43% accuracy) and McNemar's test ($p<0.001$) confirms this performance gap is statistically significant and not due to chance. Feature importance analysis reveals that redirect behavior, status bar manipulation, and domain end are the strongest discriminators. We demonstrate that effective phishing detection requires neither extensive computational resources nor complex deep learning architectures, establishing reproducible benchmarks for future research.

Paper ID 176

MACHINE-LEARNING DRIVEN PERFORMANCE MODELING AND OPTIMIZATION OF PROBABILISTIC AND HARDWARE ACCELERATORS

Uday Korat; Mitesh Patel

Abstract- The growing need to achieve high computation capabilities at a lower cost and energy usage have pushed GPU accelerators to be a major part of the current computing system. This paper presents a machine-learning-based model that predicts the GPU's performance using benchmark data from the Kaggle GPU Benchmarks Compilation. The framework uses systematic preprocessing and feature engineering on GPUs, then trains and evaluates four learning models: Random Forest, Gradient Boosting, Optimized Gradient Boosting, and a Multilayer Perceptron (MLP). Experimental results demonstrate that the Gradient Boosting model outperforms the others in terms of predictive performance (R2= 0.9998, RMSE= 45.27, MAE= 25.70), and that five-fold cross-validation, with an average R2= 0.9995, confirms the stability and consistency of the suggested model. More analysis of errors based on the category of GPUs and various levels of performance shows that errors in prediction are very low when it comes to the ensemble-based models. Altogether, the findings suggest that tree-based ensemble learning, combined with feature engineering that leverages GPU-specific features, is an accurate, scalable, and interpretable solution for GPU performance modeling, enabling efficient GPU selection, reduced benchmarking overhead, and informed computational planning.

Paper ID 178

BHF-GUARD: BREAKING AND HARDENING FINANCIAL ML WITH ADVERSARIAL STRESS TESTS AND CERTIFIED ROBUSTNESS CHECKS

Srikumar Nayak

Abstract- Financial fraud models are now used in high-volume screening, but their security is often checked with weak tests that can hide failure modes. This work studies a practical break–harden workflow for tabular fraud detection under time shift and adversarial manipulation. Using a strict future-time split on the IEEE-CIS Fraud Detection dataset, we introduce BHF-Guard, a unified evaluation and training pipeline that (i) standardizes inputs for comparable perturbation budgets, (ii) measures robustness with strong white-box attacks (FGSM/PGD) and complementary black-box/transfer settings, (iii) hardens a differentiable fraud scorer with adversarial training, and (iv) validates robustness claims using simple certification checks (interval bound propagation) and masking diagnostics. On clean future-time test data, BHF-Guard improves operational ranking and probability quality, achieving AUC-PR of 0.592 □} 0.004 and Recall@1%FPR of 0.645 □} 0.006 with lower calibration error (ECE 0.021 □} 0.002) compared to the strongest tuned baseline (LightGBM: AUC-PR 0.573 □} 0.004). Under PGD-40 at $\epsilon = 1.0$ (standardized space), BHF-Guard retains AUC-PR 0.471 and Recall@1%FPR 0.531, while standard neural baselines drop more sharply (TabNet AUC-PR 0.368, MLP 0.351). Certified evaluation further shows non-trivial verified stability at small budgets (CRA 0.781 at $\epsilon = 0.50$), and the reported sanity checks indicate no gradient masking signature. These results support a practical security evaluation template for financial ML that links attack strength, hardening, and verification in one reproducible protocol.

Paper ID 183

EXPLAINABLE RISK DECISION SYSTEMS USING ARTIFICIAL INTELLIGENCE MODELS FOR PAYMENT FRAUD DETECTION AND IDENTIFICATION

Deepak Reddy Suram

Abstract- Fraud in payment is a pressing issue in online financial ecosystems because of the dynamic characteristics of attacks and unevenly distributed transaction data. This paper suggests a risk-conscious and explainable fraud detection model based on supervised learning models, including the Random Forest and Long Short-Term Memory (LSTM) to reinforce predictive accuracy, interpretability, and operational stability. The framework has a powerful preprocessing pipeline that includes outlier handle, feature engineering, Boruta-based feature selection and SMOTEENN balancing to deal with extreme class imbalance. Experimental experiments on a very big set of fraudulent transactions indicate that the Random Forest model has the highest accuracy of 98.30 percent, and the LSTM model has the high accuracy of 99.29 percent, capturing the trends of frauds in sequences. Moreover, Explainable AI methods with LIME increase transparency, as the mechanisms make feature-by-feature explanations about each risky transaction. In general, the suggested framework helps identify fraud in a secure and data-driven way, integrating effective predictive modeling with understandable and reliable decision-making insights necessary in contemporary payment systems.

Paper ID 186

TRANSFORMER BASED FRAMEWORK FOR IMBALANCED TRANSACTION FRAUD DETECTION IN FINTECH SYSTEMS

Sandeep Shivam, Venkat Nutalapati, Tejas Pravinbhai Patel, Amit Kumar Padhy, Madhushree Kumari and Rajesh Purushothaman

Abstract- The high rate of FinTech systems development has greatly exposed risk to more advanced cyber-attacks and proper prediction and identification of fraud cases has become a crucial issue. This paper presents a Transformer-based design in the context of transaction-level fraud detection in FinTech systems that is sensitive to extreme class imbalance (fraudulent transactions are a very small percentage of the overall activity). An end-to-end pipeline of preprocessing that includes feature engineering, robust scaling, Principal Component Analysis (PCA) and Synthetic Minority Oversampling Technique (SMOTE) is used on the IEEE-CIS Fraud Detection dataset. Three deep learning models are considered, including Transformer (proposed) and ELECTRA and LSTM with Attention, which are evaluated in the same experimental conditions. The findings demonstrate that the ELECTRA model is the least sensitive to validation with the highest value of 92.55 per cent, whereas LSTM with Attention is the second highest with a sensitivity of 92.00 per cent, and the proposed Transformer model is its competitor with a sensitivity of 90.95 per cent. With slightly lower overall accuracy, the Transformer has much greater sensitivity to fraudulent transactions (29.69% detection rate), which is more useful in practical FinTech settings where fraud loss is very expensive. These findings support the practicality of attention-based models and DL architectures for identifying and preventing cyberattacks in FinTech settings.

Paper ID 189

AN EXPERIMENTAL STUDY OF LINUX LANDLOCK: FILE-SYSTEM ENFORCEMENT BEHAVIOR AND OVERHEAD

Harsh Deepak Singh; Michael J. Dinneen; Sathiamoorthy Manoharan

Abstract- Safely executing untrusted programs requires con- straining host file-system access without sacrificing performance or deployability. Linux Landlock is a recent kernel mechanism that enables unprivileged, process-scoped file-system access control, but its practical guarantees and overhead remain underexplored. We present an empirical evaluation of Landlock for isolating untrusted user-space programs. Using an adversarial probe suite, we characterize enforcement against unauthorized file accesses and common pathname evasions; using controlled sysbench workloads, we quantify runtime overhead under file-intensive execution. We also assess ruleset expressiveness and deployment ergonomics in practice. Results show that Landlock enforces deny- by-default, path-based confinement with negligible overhead for the evaluated workloads. We further identify explicit boundary conditions, including persistence of pre-opened file descriptors and the lack of mediation for non-file channels. Overall, our findings clarify Landlock's security properties and guide its use as a practical building block for untrusted-code isolation.

Paper ID 190

DESIGN AND DEVELOPMENT OF AN ARTIFICIAL INTELLIGENCE SUPPORTED EDUCATIONAL PORTAL WITH ROLE-BASED AUTHORIZATION STRUCTURE

Songül Karabatak, Muslim Alanoğlu, Murat Karabatak and Beyza Basatoğrul

Abstract- This study addresses the design and development process of an artificial intelligence–supported educational portal with a role-based authorization structure, aimed at reducing the adaptation challenges and experiences referred to as "reality shock" encountered by pre-service teachers during their transition into the teaching profession. The study was conducted within the scope of the TÜBİTAK 1001 research program as part of the project titled "Design of an Instruction Supported by a Digital Assistant for Pre-service Teachers to Overcome Reality Shock and Investigation of its Effectiveness" (Project No. 323K013). The research was carried out as a design- and development-oriented system development study. Accordingly, the web-based educational portal was structured around three primary user roles: Admin, Advisor, and Student. This paper presents a detailed overview of the portal's role-based authorization model, the functions of user panels, and the integration and positioning of the AI-supported digital assistant within the system. The development of the portal is ongoing, and following the completion of the development process, an experimental implementation and evaluation phase is planned to examine the effectiveness of the portal.

Paper ID 192

NEUROSHIELD: A TEMPORAL SPIKE TIMING ATTACK DETECTION FRAMEWORK FOR STDP BASED NEUROMORPHIC SYSTEMS IN MULTI-CLOUD ENVIRONMENTS

Naga Sujitha Vummaneni; Sundeep Bobba; Akhil Peddi

Abstract- Researchers are turning to edge computing and artificial intelligence workloads in the cloud as an energy-efficient solution, when building neuromorphic computing systems that use Spike Timing Dependent Plasticity. However, the unique temporal dynamics of STDP can create brand-new areas of vulnerability that traditional intrusion detection systems can't cover. NeuroShield is one of the comprehensive solutions developed to address these new dangers, and it's been designed to sniff out and neutralise temporal spike timing attacks targeting neuromorphic systems running on Amazon Web Services, Microsoft Azure, and Google Cloud Platform. We've identified five kinds of attacks that exploit the millisecond-scale timing windows that are characteristic of STDP learning rules. Synaptic weight manipulation, spike injection attacks, temporal replay attacks, hijacking of the learning phase, and distributed spike flooding. Our experimental results show that NeuroShield hits its target with a mean detection accuracy of 94.7% and responds in just 2.3 milliseconds, thanks to a brand-new bio-inspired anomaly detection algorithm that digs deep into spike frequencies, inter- vals and synaptic weight changes to separate legitimate neural activity from malicious interference. When compared to existing solutions, NeuroShield manages to outperform them by 23.4% in terms of detection rate, all the while keeping false positives to an absolute minimum, under 2.1%.

Paper ID 193

EVALUATING THE VULNERABILITY OF DEEPFAKE IMAGE DETECTION MODELS TO ADVERSARIAL MANIPULATIONS

Prakriti Shakya; Katya Mkrtchyan; Rashida Hasan

Abstract- Deepfake image detection systems have achieved very high accuracy on widely used benchmark datasets. However, their reliability under adversarial manipulation and dataset shift remains unclear. In real-world deployment, detectors may encounter adversarially perturbed inputs or data from unseen distributions, where standard evaluation results may not reflect true robustness. This work presents a systematic evaluation of deepfake image detectors under adversarial and cross-dataset conditions. We evaluate five representative model architectures spanning convolutional and transformer-based designs across multiple training and testing configurations. Adversarial examples are generated using an iterative gradient-based attack. Experiments are conducted on CelebA and RealVsFake datasets, with additional cross-dataset evaluation using FaceForensics++.

Our results show that while all models achieve near-perfect performance when training and testing data distributions match, performance degrades sharply when evaluated on adversarially attacked images or unseen datasets. Incorporating adversarially perturbed samples during training significantly improves robustness. However, such

robustness does not consistently transfer across datasets. These findings demonstrate that high in-distribution accuracy alone is insufficient for assessing real-world deepfake detection performance and highlight the need for robustness-aware evaluation protocols.

Paper ID 195

ZERO TRUST ARCHITECTURE FOR 5G-ENABLED MOBILE CLOUD COMPUTING (MCC)

Urvish Pandya; Sweta Pandya

Abstract- The convergence of Mobile Cloud Computing (MCC) and Fifth Generation (5G) networks is essential for enabling rapid, low-latency connectivity for emerging and next-gen technologies such as smart, autonomous vehicles, augmented reality, and innovative healthcare. Still, the need for unrestricted and diverse 5G-enabled MCCs contributes to further changes in the MCC environment, including increased security threats and even more varied, distributed characteristics of the trillion-dollar 5G MCC systems globally. It disrupts the standard MCC security and protection perimeter frameworks. Here and now, MCC security perimeter frameworks are entirely out of date and cannot meet the security and protection perimeter needs of the MCC 5G systems.

Zero Trust Architecture (ZTA) can fit this MCC 5G condition. ZTA is focused on continuous authentication, emphasizing 'never trust, always verify,' and offers flexible security protection to meet the varying needs of MCC 5G environments while still protecting 'least-privilege' access. This paper presents and helps describe the primary threats to 5G-enabled Mobile Cloud Computing, and proposes a Zero Trust Architecture (ZTA) security protection framework that defines the technical paradigm needed to enhance confidentiality, integrity, and availability (C-I-A) while maintaining optimal, effective performance (OEP).

Paper ID 196

SECURING UNMANNED AERIAL VEHICLES: ADDRESSING CYBER THREATS AND VULNERABILITIES IN UAVS SYSTEMS

Jaber El Mahjoub; Abderahim Abdellaoui

Abstract- Nowadays, Unmanned Aerial Vehicles (UAVs) have become increasingly wide-spread and commonly used for both civil and military applications. As a result, their security has become paramount, necessitating the implementation of robust and effec-tive countermeasures against a wide range of emerging cyber threats. This article out-lines the current state of UAV security issues along with corresponding mitigations and presents a demonstration of penetration testing applied to UAV systems to evalu-ate their resilience against cyberattacks. More specifically, it delves into the most prevalent vulnerabilities found in UAV systems, notably weak encryption mechanisms, lack of secure communication protocols, and critical vulnerabilities in the firmware. It further highlights previous research findings to outline the threat vectors most commonly exploited in the existing academic and technical literature, including GPS spoofing and signal jamming techniques. Finally, it explores innovative mitigation strategies such as advanced cryptographic methods and AI-based threat detection systems. The paper elaborates on the importance and implications of enhanced safety protocols, considering both the advantages and limitations of current protection approaches.

Paper ID 197

OPTIMIZING POST-QUANTUM CRYPTOGRAPHIC ALGORITHMS FOR RESOURCE-CONSTRAINED DEVICES

Sai Charan Reddy Nevuri

Abstract- This paper presents three novel contributions to post-quantum cryptography (PQC) implementation on resource- constrained devices: (1) the first comprehensive benchmarking of NTRU, U-LP, LP, and BLISS schemes across three hardware tiers (desktop, ARMv7 embedded, ATmega2560 microcontroller); (2) optimized implementations incorporating precomputed NTT tables that reduce polynomial multiplication time by 43% on ARMv7 platforms and uniform noise sampling that eliminates decryption failures in U-LP encryption; and (3) practical deployment case studies demonstrating BLISS signatures achieving 1.8s signing on ATmega2560 and U-LP enabling

zero-failure LoRaWAN firmware updates. Our results establish that with appropriate optimisations, lattice-based cryptosystems achieve acceptable throughput (0.3-2.1s per operation) on ultraconstrained platforms, though memory-security tradeoffs persist. This work provides concrete implementation guidelines for quantum-resistant cryptography in IoT and embedded systems.

Paper ID 198

ENHANCING POST-QUANTUM KEMS: A SECURE AND EFFICIENT TRANSFORMATION

Sai Charan Reddy Nevuri

Abstract- This paper proposes a new generic Encrypt-then- MAC (EtM) transformation for constructing IND-CCA secure Key Encapsulation Mechanisms (KEMs) in the post-quantum setting. The transformation builds on a public-key encryption scheme satisfying one-wayness under plaintext-checking attacks (OW-PCA) and an existentially unforgeable MAC, and formally achieves tight IND-CCA security in the random oracle model via a rigorous security reduction. Unlike the widely used Fujisaki-Okamoto (FO) transformation, our approach eliminates the need for ciphertext re-encryption during decapsulation—a major source of computational overhead and side-channel vulnerability in lattice-based implementations. We instantiate the transformation using the ML-KEM (FIPS 203) subroutine, yielding the MLKEM-EtM variant. Experimental results show that ML-KEMEtM reduces decapsulation cycle counts by 72 − 79% across security levels, with only a modest increase in encapsulation cost (1.8 − 7.3%) and ciphertext size (1.0 − 2.1%). The scheme thus offers a practical and efficient alternative to FO-based KEMs, especially in performance-critical scenarios such as post-quantum TLS handshakes, where decapsulation efficiency i s paramount.

Paper ID 200

PERFORMANCE EVALUATION OF POST-QUANTUM CRYPTOGRAPHY IN IOT: A CASE STUDY ON MQTT OVER TLS UNDER NETWORK CONSTRAINTS

Emerson Costa Santos; Tiago Augusto Orcajo Demay Cordeiro; Hebert de Oliveira Silva

Abstract- The imminent advent of quantum computing poses a significant threat to current asymmetric cryptographic standards, such as RSA and Elliptic Curve Cryptography (ECC), widely used in Internet of Things (IoT) protocols. This paper evaluates the performance impact of migrating to Post-Quantum Cryptography (PQC) in an MQTT (Message Queuing Telemetry Transport) environment secured by TLS. We propose a reproducible benchmark framework based on containerization (Docker) and hardware emulation (QEMU/ARM64) to simulate resource-constrained IoT devices. We compared the handshake latency of standard TLS against a hybrid implementation utilizing NIST-standardized PQC algorithms: ML-KEM-768 (Key Encapsulation) and ML-DSA-65 (Digital Signature). Experiments conducted under controlled network impairments (latency, jitter, and packet loss) revealed that PQC introduces an approximate 4.7x increase in connection establishment time compared to classical TLS, primarily due to computational overhead on the ARM64 architecture. However, message publication latency post-handshake remains negligible. These findings suggest that while PQC is viable for IoT, it necessitates architectural shifts towards persistent sessions to mitigate initial handshake costs.

Paper ID 201

THE INVESTIGATION OF THE RELATIONSHIP BETWEEN PEER RELATIONSHIPS AND INTERNET ADDICTION LEVELS OF HIGH SCHOOL STUDENTS

Merve Özer, Ayşenur Kuloğlu, Müslim Alanoğlu and Murat Karabatak

Abstract- This study aimed to examine the relationship between peer relationships and internet addiction levels among high school students. The sample of the study consisted of 400 high school students studying in the Afşin district of Kahramanmaraş during the 2025–2026 academic year. Data were collected using the Peer Relationships Scale to measure high school students' peer relationships and the Young Internet Addiction Test–Short Form to assess internet addiction. The results indicated that high school students had high levels of peer relationships and moderate levels of internet addiction. It was found that peer relationships differed significantly according to gender and grade level variables, whereas internet addiction did not show a significant difference. The analyses revealed a significant weak

negative relationship between peer relationships and internet addiction, and that peer relationships significantly predicted internet addiction at a low level.

Paper ID 207

ENHANCING INDOOR LOCALIZATION ACCURACY WITH BLUETOOTH LOW ENERGY RSSI SIGNALS ANALYSIS USING MACHINE LEARNING ALGORITHMS

Mitesh Patel; Uday Korat

Abstract- Recently, intelligent settings have begun to demand real-time indoor positioning, particularly where GPS signals are not accessible or distrusted. Although Received Signal Strength Indicator (RSSI) readings are very susceptible to multipath propagation, signal attenuation, and indoor interference, Bluetooth Low Energy (BLE) can be used as an inexpensive option To address these challenges, this study proposes a machine learning–driven indoor localization framework that emphasizes signal stability–oriented preprocessing, robust feature engineering, and a dynamically weighted hybrid ensemble strategy to improve resilience against environmental noise and device heterogeneity. The proposed approach evaluates multiple regression and ensemble learning models using experimentally collected BLE RSSI data. Results show that tree-based and ensemble methods outperform conventional approaches by effectively modeling nonlinear and noisy RSSI–distance relationships. Random Forest achieved a high average R^2of 0.993 with low MAE (0.113) and RMSE (0.352), while XGBoost produced the lowest mean Euclidean positioning error (0.177). Stacking ensembles based on ElasticNet and Ridge further improved prediction stability, surpassing deep learning baselines such as LSTM in overall localization accuracy. The findings demonstrate that combining stability-focused preprocessing with adaptive ensemble learning provides a reliable and accurate BLE-based indoor localization solution, offering improved robustness compared to traditional machine learning and deep learning methods.

Paper ID 211

INFOTAINPC: EXPLORING THE PRIVACY CONCERNS WITHIN IN-VEHICLE INFOTAINMENT DATA SERVICES USING TAM

Samiul Alam

Abstract- Information sharing is seen as a good aspect of any domain space because it helps to do analytics and understand beyond the upfront. Data sharing and data generation are exponential growth due to the massive use of devices that are connected to networks such as smartphones and IoT devices. The automotive industry has also adapted the use of the connected devices to ease its infrastructural framework for implementing connected and autonomous vehicles and providing convenience in modern high-tech vehicles. Most of this is conducted using the in-vehicle infotainment system which is the most interactive part of the vehicle for a user. To avail many services, a user may have to be connected to the interface of the infotainment system using smartphone or smart built-in applications, which require and store information. This study explores the user's willingness to adapt to such environment. Our methodology promotes a research design to survey users who are familiar with this technological environment to be able to provide valuable responds. Our hypotheses, built upon Technology Acceptance Model (TAM), suggest that users gain trust in providers when they understand that providers implemental marginal safeguards. We also found that perceived privacy concerns do not build trust in users towards using systems of infotainment providers. Users tend to adopt the system depending on the influence and benefits of the system.

Paper ID 213

EMPIRICAL ANALYSIS OF AI CONFIDENCE METHODS IN DIGITAL FORENSIC STANDARDS

Milinda Rambel Stone; Varghese Vaidyan

Abstract- AI deployment in digital forensics has outpaced validation standards, creating uncertainty about evidence reliability and legal admissibility. This systematic analysis of 165 NIST and SWGDE standards (2010-2025) reveals that 73\% lack confidence measurement protocols, with only 11.5\% providing AI-specific guidance. NIST AI publications achieve 7.88\% coverage while SWGDE standards lag at 3.65\%. These gaps directly affect practitioners lacking standardized protocols, courts unable to assess admissibility under Daubert requirements, and laboratories

without quality assurance benchmarks. The public dataset and confidence assessment rubric provide baselines for measuring standards development and evidence-based recommendations for addressing validation gaps.

Paper ID 220

HARDENING THE OSV UNIKERNEL WITH EFFICIENT ADDRESS RANDOMIZATION: DESIGN AND PERFORMANCE EVALUATION

Alex Wollman; John Hastings

Abstract- Unikernels are single-purpose library operating systems that run the kernel and application in one address space, but often omit security mitigations such as address space layout randomization (ASLR). In OSv, boot, program loading, and thread creation select largely deterministic addresses, leading to near-identical layouts across instances and more repeatable exploitation. To reduce layout predictability, this research introduces ASLR-style diversity into OSv by randomizing the application base and thread stack regions through targeted changes to core memory-management and loading routines. The implementation adds minimal complexity while preserving OSv's lightweight design goals. Evaluation against an unmodified baseline finds comparable boot time, application runtime, and memory usage. Analysis indicates that the generated addresses exhibit a uniform distribution. These results show that layout-randomization defenses can be efficiently and effectively integrated into OSv unikernels, improving resistance to reliable exploitation.

Paper ID 221

NOWHERE TO HIDE: COMPARATIVE ANALYSIS OF RESIDENTIAL VS. CLOUD ATTACK SURFACES

Sankalp Lanka, Om Palsule, Rishik Lanka; Phani Lanka

Abstract- The expansion of internet connected devices has widened the global attack surface, making it essential to understand how attackers adapt to different networks. Implementing uniform security measures everywhere is not feasible. Commercial cloud providers employ enterprise grade monitoring to defend their networks, while residential homes often lack such protection. This raises critical questions: Is the intensity of attacks the same in both environments, or does the lack of monitoring in homes obscure the true danger? To answer this, this study presents a controlled comparison of cyberattacks on cloud hosted versus residential networks using identical honeypots. We emulated SSH, Telnet, and IoT interfaces to characterize attacker behavior. Fourteen cloud instances were deployed across major providers, including Microsoft Azure, Amazon Web Services (AWS), Google Cloud Platform (GCP), and Oracle Cloud Infrastructure (OCI), alongside three residential nodes exposing identical services on home infrastructure. Over a three-week period, we collected data on connection attempts, credential guessing, executed commands, and malware delivery across 8 services. Residential networks received more attacks than AWS instances and were targeted with specific IoT credentials not seen in the cloud. These findings provide empirical measurements of attacker activity, highlighting the distinct security risks faced in various cloud platforms and home networks.

Paper ID 222

DATABASE FORENSICS READINESS: AN EXAMINATION OF REDIS

Muhammad Abdul Moiz Zia; Oluwasola Mary Adedayo

Abstract- Databases are essential evidence containers in modern digital forensic investigations, storing crucial information like user activity, transaction records, authentication data, and system events. Although digital forensics research has focused on relational databases, the rise of NoSQL databases, such as Redis, has introduced new challenges in digital forensic readiness. Although Redis supports scalability, low latency, and flexible data structures, for improved performance, it provides limited auditing and event logging capabilities that may hinder forensic capabilities. This study aims to evaluate Redis's forensic readiness through normal and malicious CRUD operations and analyzes the digital evidence that can be extracted through its native logging, monitoring, and built-in tools. Through systematic experiments simulating authorized and unauthorized activities under different configurations and access control settings, we find that Redis captures basic operational traces but lacks consistent user attribution, command execution outcomes, session tracking, and error reporting under many default configurations. Based on these findings, we provide

recommendations for improving Redis's forensic readiness by leveraging external tools and methods to capture additional evidence. To the best of our knowledge, this is the first study to systematically evaluate Redis from a forensic readiness perspective by experimentally mapping CRUD operations to the evidentiary artifacts they generate and combining them for a more complete forensic picture.

Paper ID 224

Enforcing Accountability in Autonomous Ops: A Zero-Trust Multi-Agent Framework with Forensic Reasoning Ledgers

Madhvesh Kumar; Deepika Singh

Abstract- As Large Language Models (LLMs) transition from passive chatbots to active agents capable of executing code and managing infrastructure, the security perimeter shifts from user authentication to agent authorization. Current monolithic agent architectures lack the granularity required for high-stakes enterprise environments, creating a liability where errors are indistinguishable from malice. This paper proposes a Zero-Trust Multi-Agent System (ZT-MAS) architecture. By decomposing the AI into role-specific agents (Scout, Analyst, Fixer) governed by a deterministic Supervisor, we enforce "Inter-Agent Zero Trust," preventing hallucination propagation. Furthermore, we introduce the "Forensic Log of Agency," a mechanism that serializes internal Chain of Thought (CoT) reasoning into an immutable ledger, ensuring that every autonomous action is legally and technically auditable. We provide comprehensive implementation details, comparative analysis with existing frameworks, and empirical validation through multiple security scenarios.

Paper ID 225

Research Study on AI-Driven Weapon and Violence Detection in Smart Cities

Thangavel Murugan; Anoud Salem Alkarbi; Maryam Ahmed Alzeyoudi; Noura Ali Matar Alketbi; Afra Khaleifah Alteneiji; N. Nasurudeen Ahamed

Abstract- Smart cities are increasingly using artificial intelligence (AI) for improved public safety in ways that cannot be accomplished with current traditional surveillance techniques, which rely on ongoing surveillance and do not utilize real-time analytical intelligence. Technological innovations in AI and computer vision now enable the automated identification of weapons and violent behavior in public spaces, thereby increasing situational awareness and reducing response times for public safety services. The objective of this document is to develop a comprehensive overview of current AI technologies used to identify weapons and violence, discussing their effectiveness, limitations, and challenges to utilization within the urban surveillance environment. A review of the research literature identified significant research gaps regarding accuracy, scalability, and real-time processing as substantial impediments to the effective implementation of AI-based surveillance technology. As a result of this analysis, an Intelligent Surveillance System conceptual framework that combines real-time video surveillance with AI-enhanced analytics will be presented. In addition, the greater potential of intelligent surveillance to increase public safety operations within the Smart City environment will also be discussed.

Paper ID 228

Microarchitectural Espionage: FPGA-Based Security Analysis of Branch Prediction in RISC-V Out-of-Order Cores

Mahreen Khan; Muhammad Emir Bin Mohd Shahfie; Maria Mushtaq; Renaud Pacalet; Ludovic Apvrille

Abstract- Modern processor microarchitectural optimizations, while enhancing performance, inadvertently introduce side channels that can leak sensitive information through timing variations. This paper presents an FPGA-based security testbed for studying branch predictor side-channel vulnerabilities in open-source RISC-V out-of-order cores. We demonstrate a configurable platform built on the Berkeley Out-of-Order Machine (BOOM) core, adapted for resource-constrained FPGA deployment with customizable branch predictor configurations. Through bare-metal execution and cycle-accurate timing measurements, we implement and evaluate three classes of timing attacks: Conditional Branch Prediction Attacks (CBPA), Indirect Branch Prediction Attacks (IBPA), and a practical smart-lock application attack. Our results show that simplified one-level predictors exhibit deterministic timing separations of 9 to 17 cycles, enabling

perfect secret recovery with 100% accuracy for 16-bit secrets within 500 measurement rounds. We further demonstrate practical attack scenarios, including the extraction of a randomly generated 4-digit smart-lock code, and evaluate the impact of branch predictor complexity on attack feasibility. This work provides an open-source framework for reproducible microarchitectural security research on RISC-V platforms, enabling evaluation of both attacks and countermeasures.

Paper ID 230

NETWORK HARDENING BASED ON COMPANION PLANTING IN IOT ENVIRONMENTS WITH UNKNOWN-VULNERABILITY DEVICES

Shingo Yamaguchi

Abstract- This paper proposes a bio-inspired IoT network hardening method based on companion planting to mitigate botnet threats in environments with unknown-vulnerability devices (UVDs). Unlike conventional methods that treat vulnerabilities as defects to be removed, our approach leverages device diversity to fragment potential infection paths. The core novelty lies in the strategic retention of UVDs within the topology while explicitly restricting their relay functionality, achieving robust security without the connectivity loss inherent in prior exclusion-based methods. Simulation results on scale-free networks demonstrate that this method effectively limits the size of vulnerable components below predefined thresholds, with structural analysis revealing a functional vulnerability rate limit of 50%.

Paper ID 232

DATA-DRIVEN AI TECHNIQUES IN RAMAN SPECTROSCOPY FOR BIOMEDICAL APPLICATIONS: A COMPREHENSIVE REVIEW

Shahbaz Ahmed; Muhammad Raheel Raza; Asaf Varol

Abstract- Raman Spectroscopy has emerged as a significant method for biological diagnosis. It emerged because of its ability to provide molecular-level information without the use of labels. The use of machine learning (ML) techniques, especially clustering and deep learning (DL), has strongly improved the analysis of Raman spectra data, thus enhancing the classification and detection of different cancer variants and microbial infections. This study highlights how supervised learning algorithms, such as Support Vector Machines (SVM) and Random Forests, and unsupervised learning algorithms, such as hierarchical, PCA, and K-means clustering, contribute to interpreting and classifying the complex spectral data more accurately. Deep learning models such as CNN is also used to examine their performance in biomarker identification and pattern recognition. This survey paper shows how diagnostic accuracy is improved by merging Raman Spectroscopy with intelligent models, highlighting some key limitations and future directions, including noisy data, high computational demands, and dependence on labeled data. This work aims to contribute to the continued development of advanced diagnostic techniques based on spectroscopy and intelligent data analysis.

Paper ID 233

A HYBRID APPROACH TO DETECT ILLEGAL ACTIVITIES IN DARK WEB DATA

Akash G Gaonkar, Abhinav V P, Aditya V Bhat, Jeny Jijo and Abishek K

Abstract- The current solutions for dark web monitoring are prone to the problem of "context blindness," as they are based on unimodal analysis that cannot identify the illicit content hidden using evasive text or images. To fill this research gap, this paper proposes a hybrid analysis framework that can automatically scan and classify dark web pages by combining Visual (YOLOv5 + CLIP) and Textual (OCR + DistilBERT) analysis streams. The proposed system was trained on a custom-annotated dataset of 1,500 dark web screenshots and tested on real-world .onion URLs. By using a weighted fusion algorithm to combine the multimodal evidence, the proposed method was able to achieve a classification accuracy of 93.3%, which is substantially better than the baselines (Text-only: 87.7%, Visual-only: 72.7%). These findings clearly show that the combination of contextual information from multiple modalities is an effective way to resolve the ambiguity, providing a powerful and automated solution for law enforcement to identify active digital crimes.

Paper ID 234

DEVELOPMENT OF AN ARTIFICIAL INTELLIGENCE–SUPPORTED MOBILE LEARNING APPLICATION

Songül Karabatak; Muslim Alanoğlu; Ayşenur Kuloğlu; Dilara Sucu

Abstract- In this study, the design and development process of an artificial intelligence–supported mobile learning application developed within the scope of the research project titled "Design of an Instruction Supported by a Digital Assistant for Pre-service Teachers to Overcome Reality Shock and Investigation of its Effectiveness", conducted under the TÜBİTAK 1001 program and numbered 323K013, is addressed. Within this scope, a mobile application was developed as a complementary component to the web-based educational platform, aiming to provide pre-service teachers with learning opportunities independent of time and place. The mobile application was structured to include module-based content delivery, media-supported learning materials, and artificial intelligence–supported digital assistant components. In the study, the architectural structure of the developed mobile application, its user interface and navigation design, as well as the integration of artificial intelligence components into the mobile environment, are explained in detail. The system development process has been completed, and it is planned to proceed to the experimental implementation and evaluation phase to test the effectiveness of the application on pre-service teachers' professional adaptation processes and reality shock.

Paper ID 235

THE NEED FOR STANDARDIZED EVIDENCE SAMPLING IN CMMC ASSESSMENTS: A SURVEY-BASED ANALYSIS OF ASSESSOR PRACTICES

Logan Therrien; John Hastings

Abstract- The Cybersecurity Maturity Model Certification (CMMC) framework provides a common standard for protecting sensitive unclassified information in defense contracting. While CMMC defines assessment objectives and control requirements, limited formal guidance exists regarding evidence sampling, the process by which assessors select, review, and validate artifacts to substantiate compliance. Analyzing data collected through an anonymous survey of CMMC-certified assessors and lead assessors, this exploratory study investigates whether inconsistencies in evidence sampling practices exist within the CMMC assessment ecosystem and evaluates the need for a risk-informed standardized sampling methodology. Across 17 usable survey responses, results indicate that evidence sampling practices are predominantly driven by assessor judgment, perceived risk, and environmental complexity rather than formalized standards, with formal statistical sampling models rarely referenced. Participants frequently reported inconsistencies across assessments and expressed broad support for the development of standardized guidance, while generally opposing rigid percentage-based requirements. The findings support the conclusion that the absence of a uniform evidence sampling framework introduces variability that may affect assessment reliability and confidence in certification outcomes. Recommendations are provided to inform future CMMC assessment methodology development and further empirical research.

Paper ID 236

FORENSIC EVENT RECONSTRUCTION OF WEB ATTACKS USING LOG DECODER AND RULE-BASED CORRELATION

Muhammad Nur Yasir Utomo; Hudan Studiawan; Baskoro Adi Pratomo

Abstract- Cyberattacks on web-based applications have increased in recent years. While preventive measures in web security have been widely explored, post-incident investigations, particularly event reconstruction, remain underexplored despite their pivotal role in understanding attack sequences. The complexity of this problem increases because attackers generally disguise their attacks in the form of encoded URLs. To address this problem, this study proposes an integrated event reconstruction method that combines timeline generation, log decoder, and rule-based event correlation to analyze multi-stage attack activities. The log decoder is used to transform encoded commands into their original semantics to provide clearer interpretation and improve generalization of correlation rules. The experimental evaluation was conducted using public dataset, the AIT Log Data Set V2.0 (FOX scenario), complemented by two real-world datasets to assess practical applicability. The experiment focused on server-side log artifacts from web and system servers. A total of 13 server-side log files were merged into a super timeline, preprocessed to decode hidden commands, and analyzed using correlation rules. The rules were categorized into

reconnaissance, exploitation, and privilege escalation activities. The proposed method achieves 88.4% precision, 82.1% recall, and an F1-score of 85.1% against ground truth records on AIT dataset and demonstrates its ability to reconstruct the attack chain from reconnaissance to privilege escalation in web security incidents. The proposed method also proves its capability across different attack characteristics, ranging from failed to successful exploitation attempts on real-word datasets.

Paper ID 237

GO-SAFEINPUT: A ZERO-DEPENDENCY INPUT SANITIZATION FRAMEWORK FOR DIGITAL FORENSICS AND CYBERSECURITY APPLICATIONS

Ravi Sastry Kadali

Abstract- Input injection vulnerabilities remain a critical threat to digital forensics platforms and cybersecurity infrastructure. The MITRE 2025 CWE Top 25 analysis identifies 41 Known Exploited Vulnerabilities across four injection categories: command injection (20 KEVs), path traversal (10 KEVs), cross-site scripting (ranked #1), and SQL injection (ranked #2). Digital forensics tools that process untrusted evidence files, case management systems handling user inputs, and security platforms analyzing malicious artifacts are particularly vulnerable to these attacks. This paper presents go-safeinput, a unified input sanitization framework for Go applications providing context-aware defenses against all major injection categories with zero external dependencies. The zero-dependency design is critical for forensics tools where software provenance and supply chain integrity are paramount. Our security evaluation demonstrates 100% detection across 530+ attack payloads. Performance analysis shows sub-millisecond latency (0.03--0.15ms) with less than 2% throughput overhead. We present integration patterns for forensics evidence processing pipelines, case management systems, and security analysis platforms, demonstrating practical applicability to the digital forensics and security domain.

Paper ID 239

DESIGN OF A WEB-BASED MODULAR SELF-ASSESSMENT SYSTEM FOR EDUCATION PORTALS

Songül Karabatak; Muslim Alanoglu; Meltem Yurtcu; Suleyman Polat

Abstract- In this study, the design and development process of a digital Self-Assessment System aimed at evaluating learners' knowledge levels, awareness, and readiness is presented. The system was developed within a web-based framework and designed as part of an educational program consisting of eight distinct modules, enabling learners to assess their own competencies through structured assessment instruments. The assessment tests used in the system were examined through pilot implementations based on comprehensive item pools developed for each module. The pilot studies were conducted with groups of 17–20 participants, and the collected data were analyzed using the Test Analysis Program (TAP) in terms of item difficulty, item discrimination, and reliability indices. Based on the analysis results, problematic items were revised or removed from the item pools, and the final sets of assessment items were revised in line with expert feedback and integrated into the system. In addition, the study provides a detailed description of the system architecture and the digital integration approach grounded in item development and analysis processes. The system development process has been completed, and the resulting structure offers an assessment infrastructure suitable for web-based implementations across different educational contexts and user groups.

Paper ID 241

QUANTUM-SAFE COUNTERMEASURES: MITIGATING DETECTOR-BLINDING ATTACKS IN QUANTUM KEY DISTRIBUTION SYSTEMS

Wail Zita; Malek Malkawi

Abstract- Detector-blinding attacks threaten practical Quantum Key Distribution (QKD) systems by exploiting non-ideal behavior in single-photon detectors rather than weaknesses in the underlying protocol. By injecting bright optical power, an adversary can force detectors into an abnormal operating regime and manipulate measurement outcomes while keeping the observed Quantum Bit Error Rate (QBER) within acceptable bounds. This paper proposes a simulation-based, hardwareaware countermeasure that combines optical power thresholding with statistical monitoring

of detector click dynamics using a sliding-window Z-score test. The framework operates entirely at the receiver side and does not require modifications to the QKD protocol. In Monte Carlo evaluation over step, ramp, and pulsed attack profiles, the combined defense achieves TPR = 1.00 and enforces conservative session abortion to prevent silent key compromise.

Paper ID 242

AI-Enhanced Cybersecurity Risk Assessment for Smart Grid Infrastructures Using NIST Framework

Hala Strohmier Berry, Bhargav Palaskar, Soumik Mitra and Yash N Dasri

Abstract- The digitalization of power systems has introduced cyber-physical vulnerabilities in smart grid infrastructures that traditional, rule-based detection mechanisms are insufficient to address. This paper presents an AI-enhanced cybersecurity risk assessment framework that integrates the NIST Cybersecurity Framework (CSF 2.0), NIST SP 800-30, and MITRE ATT&CK for ICS with machine learning-based threat detection. Three classifiers — Logistic Regression, Random Forest, and XGBoost — were trained and evaluated on the Blockchain-Enabled 6G Microgrid Dataset, a multi-layer dataset comprising physical, communication, and blockchain operational features with a severe class imbalance ratio of 29.5:1. Model performance was assessed using accuracy, precision, recall, F1-score, and ROC-AUC, with recall and F1-score prioritized given the asymmetric cost of false negatives in cybersecurity contexts. XGBoost achieved the highest performance (accuracy: 0.993, F1-score: 0.885, ROC-AUC: 0.974), outperforming Logistic Regression and Random Forest through its sequential boosting architecture and scale_pos_weight parameter, which up-weighted the minority attack class during training. A key contribution is the systematic mapping of XGBoost probabilistic output scores to NIST SP 800-30 likelihood tiers, enabling continuous quantitative risk scoring via R= L × I and supporting automated protective responses in real time. A case study of the 2015–2016 Ukrainian power grid attacks further contextualizes the framework's applicability to real-world critical infrastructure scenarios.

Paper ID 243

ChaCha20-E: An Improved ChaCha Algorithm for Secure Data Transmission on IoMT Devices

Jabari Khawla and Abdellaoui Abderrahim

Abstract- The Internet has evolved beyond connecting computers and devices to enabling the connectivity of physical objects, known as the Internet of Things (IoT). This technology is widely adopted across sectors such as healthcare, education, agriculture, and commerce. Since IoT devices handle sensitive information, encryption is essential. However, many existing encryption techniques involve high computational and power consumption, making them unsuitable for resource-constrained environments. This paper proposes ChaCha20-E, an improved version of ChaCha designed to enhance security while reducing computational and memory consumption for Internet of Medical Things (IoMT) devices.

Paper ID 244

LRD-Net: A Lightweight Real-Centered Detection Network for Cross-Domain Face Forgery Detection

Xuecen Zhang; Vipin Chaudhary

Abstract- The rapid advancement of diffusion-based generative models has made face forgery detection a critical challenge in digital forensics. Current detection methods face two fundamental limitations: poor cross-domain generalization when encountering unseen forgery types, and substantial computational overhead that hinders deployment on resource-constrained devices. We propose LRD-Net (Lightweight Real-centered Detection Network), a novel framework that addresses both challenges simultaneously. Unlike existing dual-branch approaches that process spatial and frequency information independently, LRD-Net adopts a sequential frequency-guided architecture where a lightweight Multi-Scale Wavelet Guidance Module generates attention signals that condition a MobileNetV3-based spatial backbone. This design enables effective exploitation of frequency-domain cues while avoiding the redundancy

of parallel feature extraction. Furthermore, LRD-Net employs a real-centered learning strategy with exponential moving average prototype updates and drift regularization, anchoring representations around authentic facial images rather than modeling diverse forgery patterns. Extensive experiments on the DiFF benchmark demonstrate that LRD-Net achieves state-of-the-art cross-domain detection accuracy, consistently outperforming existing methods. Critically, LRD-Net accomplishes this with only 2.63M parameters—approximately 9× fewer than conventional approaches—while achieving over 8× faster training and nearly 10× faster inference. These results demonstrate that robust cross-domain face forgery detection can be achieved without sacrificing computational efficiency, making LRD-Net suitable for real-time deployment in mobile authentication systems and resource-constrained environments.

Paper ID 245

TOKEN-BASED AUTHENTICATION SYSTEM FAILURES IN CLOUD ENVIRONMENTS: A CASE STUDY OF THE MICROSOFT STORM-0558 INCIDENT

Hamza Benmiloud; Malek Malkawi

Abstract- Cloud-based identity systems increasingly rely on token-based authentication mechanisms as a primary component for providing efficient and scalable access control across distributed environments. While these modern systems offer strong cryptographic guarantees, they also introduce new security challenges when supporting controls such as key management, system isolation, and monitoring are insufficient. This paper analyzes the Microsoft Storm-0558 incident as a case study to examine structural weaknesses in modern cloud identity infrastructures and to evaluate mitigation techniques. The study analyzes the incident, evaluates risk management and mitigation strategies, and extracts key lessons learned and recommendations. Two academic studies on token-based authentication monitoring and single sign-on systems are examined to support the proposed mitigation strategies. The findings demonstrate that cryptographic validity alone does not ensure security in authentication systems and that delayed detection significantly increases the operational and geopolitical impact of identity-based attacks. The paper concludes by emphasizing the importance of automated, token-aware monitoring and strict trust boundary enforcement to strengthen cybersecurity operations and prevent similar incidents.

Paper ID 247

THE URL NEXTDOOR: A DIGITAL FORENSIC ANALYSIS OF NEIGHBORHOOD APPS

Joseph Brown; Abdur Onik; Ibrahim Baggili

Abstract- This paper contains a digital forensic analysis of the neighborhood social media applications NextDoor and Neighbors. The research fills a gap in the literature by establishing a hybrid method of forensic analysis based on established techniques. Traffic and API analysis produced a variety of artifacts, including plaintext passwords and tokens. A python script exploiting the tokens was developed and is included here. A detailed depiction of the forensic analysis performed for this paper is presented to aid future investigators, as well as extensive recommendations for future work.

Paper ID 251

(SEAL) SEQUENTIALLY EVOLVING ALERT LEARNING FOR SMART CYBER SECURITY

Riona V; Hemalatha R

Abstract- Security Operations Centers (SOCs) generate extremely large volumes of security alerts daily, many of which are redundant or false positives, leading to alert fatigue and delayed response to critical threats. To address this challenge, this paper proposes Sequentially Evolving Alert Learning (SEAL), an adaptive alert prioritization framework that classifies alerts into Low, Medium, and High severity levels using an ensemble of Random Forest and XGBoost classifiers. The framework incorporates a sequential retraining mechanism in which newly labeled alerts are periodically integrated into the training dataset to mitigate concept drift and adapt to evolving attack patterns. SEAL was evaluated on a dataset comprising 4.3 million actual and simulated security alerts. Experimental results demonstrate an overall classification accuracy of 96.89%, with high recall for critical alerts, indicating effective prioritization performance. A Streamlit-based dashboard is integrated to provide real-time visualization of alert severity distribution and model performance metrics, supporting operational decision-making in SOC environments. The results indicate

that SEAL enhances alert prioritization reliability, adaptability, and scalability in high-volume cybersecurity monitoring systems.

Paper ID 253

BEYOND CHARTS - FINANCIAL PREDICTION WITH LANGUAGE MODEL & TIME SERIES

Sudesh Pawar

Abstract- The nonlinearity and volatility of financial markets remain major challenges in computational finance because of the need to predict stock prices accurately. This paper presents a comparative analysis of machine learning and deep learning models for predicting the S&P 500 by integrating financial time-series systems with a GPT-based language model. The study follows a structured pipeline comprising data cleaning, feature engineering, normalization, systematic training, and validation. ARIMA is used as the baseline model, and XGBoost, Transformer, and GPT-2-based autoregressive models are evaluated against it. The experimental results suggest that the Transformer architecture delivers the most robust performance, with an RMSE of 0.1507 and an R^2 of 0.9773, demonstrating a strong ability to capture time dependent relationships in multivariate financial data. In contrast, although the GPT-2 model shows high explanatory power, it is also volatile in terms of percentage error, with a MAPE of 81.62. These findings highlight the potential of attention-based architectures for predicting financial time series, while also revealing their limitations in terms of computational cost and model extrapolation. The analysis provides a reliable framework that can be adopted in future studies to improve efficiency, stability, and applicability in dynamic market environments.

Paper ID 255

LARGE-SCALE FINANCIAL FORECASTING USING ADVANCED GAI-BASED LARGE LANGUAGE MODELS AND TIME SERIES ANALYSIS

Surya Veera Brahmaji Rao Sunnam

Abstract- Financial forecasting has become a point of scrutiny with the development of AI and the increasing access to financial information but precise stock price prediction is hard because the market is volatile and nonlinear. This research paper provides a hybrid forecasting framework with past S&P 500 data (approximately 200,000 historical records) to assess machine-learning and sophisticated GAI-based deep-learning models. The methodology includes systematic data preprocessing, large-scale feature engineering, processing of missing and infinite values, and Min-Max normalization, followed by time-ordered data splitting for training and evaluation. XGBoost is used to model nonlinear relationships among structured financial data with time-series, and an ALBERT-based Transformer model is employed to capture long-range temporal dependencies using attention mechanisms. As can be seen in the experimental results, ALBERT has a higher performance with an R2 of 0.9930, RMSE of 4.3021, and MAE of 2.1380, which is significantly much higher than XGBoost and conventional deep learning models like LSTM, which has a higher error rate. The strength and precision of the ALBERT model in close correlation of actual price movements is proven by the comparative analysis. It is concluded that GAI-based Transformer models can be an effective and scalable solution to the needs of financial time series forecasting, though at the price of increased computational costs. This work demonstrates the practical applicability of lightweight Transformer architectures for accurate and data-driven financial decision-making systems.

Paper ID 256

COMPREHENSIVE REVIEW AND EXPERIMENTAL STUDY OF HOLIDAY-AWARE MULTIMODAL CRIME PREDICTION

Sapna V. M; Vishruth S. Megur; Vishwanath Bhavimani; Somesh S; Vignesh Madivala; Prasad B Honnavalli

Abstract- Crime prediction is essential for enhancing public safety by identifying patterns that signal potential criminal activity. While traditional models rely on historical crime data and recent approaches incorporate social media sentiment to capture real-time public mood, the influence of holidays and special events remains largely overlooked. Such events can significantly alter human behavior and routine activities, thereby affecting crime patterns. This paper presents a comprehensive review and experimental study of holiday-aware multimodal crime prediction, highlighting

gaps in existing time-series and sentiment-based approaches and evaluating how structured holiday features, combined with temporal and sentiment signals, can improve predictive accuracy.

Paper ID 257

BIG DATA-DRIVEN AI MODELS FOR NETWORK ANOMALY DETECTION AND CYBER THREAT FORECASTING

Amit Meshram; Vikrant Sikarwar

Abstract- The rapid growth of the network infrastructures which rely on big data, the increased complexity of cyberattacks has made it necessary to have a strong demand on intelligent and reliable network anomaly detection systems. The old security methods and traditional machine learning models are usually unable to process bulk traffic across the network and to extract complex temporal variations related to evolving cyber threats. The framework is tested on the CICIDS 2017 dataset, which provides an ideal, detailed view of normal and malicious network activity. Extensive experimental study shows that the BiLSTM model has high detection accuracy of 99.58% and a high precision of 99.45, a high recall of 99.34% and F1 score of 99.54, which outperforms a number of baseline models, including Autoencoder (88.28%), Logistic Regression (96.6%), and LightGBM (91.35%). The high level of accuracy illustrates the efficacy of bidirectional temporal learning in simulating the trend of network traffic sequentially and differentiating between benign and malicious behavior. The findings that the suggested framework is scalable, robust and is applicable in big data systems and is an effective framework in terms of network anomaly proactive detection and effective predictions of cyber threats in a contemporary cybersecurity infrastructure.

Paper ID 259

DRONE ASSISTED REMOTE WELLNESS MONITORING USING RGB CAMERA

Egemen Tunçarslan; Ihsan Yılmaz

Abstract- In remote rural regions and severe weather condi- tions (e.g., snowstorms and floods), emergency response teams may be delayed or unable to reach patients in time. This motivates drone-assisted remote triage workflows where a first- response UAV provides early situational and physiological teleme- try to a command center. We present a practical RGB camera- based wellness monitoring pipeline that estimates physiological indicators from video via remote photoplethysmography (rPPG) and motion-derived signals. Unlike accuracy-only approaches, the proposed system targets safety-critical operation by integrating explicit quality gating and a refusal policy to avoid wrong- but-confident outputs under adverse capture conditions (motion, unstable sampling, intermittent face loss, or illumination drift).

The pipeline consists of face ROI extraction, POS-inspired signal projection, temporal stabilization, and an interpretable Signal Quality Score (SQS) used for selective publishing. In a pilot real-time capture log (N=3448 JSONL rows; effective median capture rate ≈ 3.72 FPS), the system achieves high heart- rate telemetry availability (97.71% coverage) while transparently abstaining on insufficient-evidence segments with explicit reason codes. Since the pilot operates at substantially lower frame rates than typical rPPG settings, we emphasize safety-first availability under strict gating rather than clinical-grade accuracy, and conservatively disable HRV reporting under low-FPS conditions. Failure-mode analysis indicates that dominant rejections are driven by sampling instability and missing respiration evidence. The proposed design supports drone-enabled telehealth triage, where reliability, interpretability, and honest abstention are critical for operational decision-making.

Paper ID 260

LINEAR CRYPTANALYSIS OF BLOCK CIPHER LELBC

Yingjie Zhang; Gang Liu

Abstract- Linear cryptanalysis is one of the fundamental tools for evaluating the security of block ciphers. In this paper, we investigate the security of the lightweight block cipher LELBC against linear cryptanalysis at the distinguisher level. We propose an automatic search framework based on mixed-integer linear programming to identify optimal linear trails

and linear hulls of reduced-round LELBC. By modeling the propagation of linear masks and correlation bounds within a unified MILP formulation, we obtain optimal linear trails for up to 9 rounds in the single-key setting. We further construct the 9-round linear hull by enumerating all linear characteristics sharing the same input and output masks and estimate the corresponding expected linear potential. Our results show that the best linear distinguishers of LELBC achieve correlation bounds comparable to the best known differential and differential-linear distinguishers, providing new insights into the linear security margin of LELBC and complementing existing cryptanalytic results.

Paper ID 262

DARKTRACEUI: A MULTIMODAL FRAMEWORK FOR IDENTIFYING DARK PATTERNS IN WEB AND TOR ECOSYSTEMS

Sapna V M; Karthik Kademani; Prakruthi G P; Likhit Avinash V; Keerthana M; Prasad Honnavalli

Abstract- Dark patterns are deceptive interface design strategies that exploit cognitive biases and mental heuristics to mislead users into unintended actions. These practices have been extensively studied in domains such as e-commerce and social media applications. However, their presence within anonymous networks such as The Onion Router (TOR) remains insufficiently explored. This paper proposes a systematic multimodal framework for the automated identification and categorization of dark patterns across both environments. The approach integrates a fine-tuned BERT-based sequence classification model for textual analysis, achieving 90% accuracy and F1-score, and a YOLOv5m object detection model for visual pattern recognition, achieving a 65% mAP@0.5 score. To address the absence of structured datasets within the TOR ecosystem, a generative artificial intelligence pipeline leveraging Large Language Models (LLMs) with few-shot prompting is introduced for the analysis of unstructured content. The results provide a comparative evaluation of dark pattern intent and impact, contributing toward the development of automated ethical auditing frameworks and strengthening awareness of manipulative practices across both mainstream and anonymous browsing environments.

Paper ID 263

ARTIFICIAL INTELLIGENCE-BASED ANTI-MONEY LAUNDERING SOLUTIONS: ENHANCING DETECTION ACCURACY IN HIGH-VOLUME FINANCIAL DATA

Amit Meshram; Vikrant Sikarwar

Abstract- Money laundering detection has been complicating with the growing pace of digital banking and volume financial transactions. Anti-money laundering (AML) systems based on traditional rules are often not adaptable to evolving laundering patterns, have a high false-positive rate, and are poorly scalable. This work presents a proposed AML detection framework based on artificial intelligence to improve on detection accuracy of a large-scale financial transactions data. There are three models XGBoost, Extra Trees and a hybrid stacking ensemble models are evaluated on the PaySim benchmark synthetic dataset. To increase transparency and regulatory trust, Explainable Artificial Intelligence (XAI) methods are incorporated to explain model decisions and reveal the features of transactions that impact them. The experimental findings indicate that the hybrid stacking model is the most effective with an accuracy of 98.03, and Extra Trees is highly computational, hence can be used to implement the AML on a large scale. The suggested framework demonstrates considerable prospects of the effective and stable AML implementation in the financial system nowadays.

Paper ID 265

ENHANCING DETECTION OF BLOODSTAIN PATTERNS AND FOOTPRINT IMPRESSIONS USING DEEP LEARNING

Thangavel Murugan; Maryam Saif Abdulla Saeed; Hind Abdulla Asman Balmur Al Ameri; N Nasurudeen Ahamed; Aysha Salem Hamdan Obaid Alkaab, Priyan Malarvizhi Kumar; Shamsa Rashid Salem Balhas Alshamsi; Alia Ali Salem Alwali Almazrouei

Abstract- Traditional methods of analyzing blood stain patterns and shoe prints generally involve manually examining the evidence visually. This type of analysis is prone to subjectivity and can also take a long time to complete. In this

study, a deep-learning framework is developed to perform two different tasks: (a) classify bloodstain patterns, (b) make demographic predictions about the footprint impressions produced from shoes. A Vision Transformer (ViT-Base/16) was trained on approximately 2,700 images representing the different blood patterns within the database; the model achieved an overall classification accuracy of 94.2 percent across all classes of bloodstain patterns. A ResNet-18 model was trained on 2,400 images taken from the UNB StepUP-P150 dataset to predict the gender, height range, and age group of people from their footwear; the gender prediction model achieved an overall classification accuracy of 91.6 percent. The developed deep learning systems will produce probability outputs and automated reporting, thereby improving the objectivity, efficiency, and standardization of forensic evidence.

Paper ID 267

TELEGRAM AS A TRANSFORMATIVE CRIMINAL MARKETPLACE: ANALYZING IDENTITY THEFT DYNAMICS ON SURFACE-ACCESSIBLE PLATFORMS

Manik Kaur; David Maimon; Mario Kubek; Anu Bourgeois

Abstract- Identity theft operations have migrated from specialized darknet markets to surface-accessible messaging platforms, transforming how stolen credentials are distributed and monetized. Current notification systems rely on breach disclosures and static darknet scanning, leaving victims exposed during the critical window between credential compromise and fraudulent use.

This study presents a framework for monitoring public Telegram channels combining machine learning-based entity recognition with geographic and institutional attribution to enable proactive victim notification. Applied to 25 of approximately 100 monitored channels throughout 2023, the pipeline processed 473,554 PII records with high precision (89% for credit cards, 94% for SSNs). Analysis reveals distinct patterns in data composition, temporal dynamics, and institutional targeting that inform risk-based prioritization and extend Routine Activity Theory to two-stage identity theft crimes.

Paper ID 268

AN EXPLAINABLE AI STUDY OF PHISHING DETECTION MODEL DEGRADATION ACROSS REAL-WORLD EVENTS

Michael Ivanicki; Brian Callahan

Abstract- Phishing websites continue to evolve rapidly, driven by technological change and large-scale global events, challenging the long-term reliability of machine learning–based detection systems. While prior research has demonstrated strong classification performance on static datasets, limited work has examined how external real-world factors influence model behavior and feature relevance over time. We investigate the degradation of phishing website detection models and connect observed changes in feature importance to real-world phishing trends and events.

Using a dataset of 300,000 benign and malicious webpages collected between 2018 and 2020, Random Forest classifiers were trained on year-specific data and evaluated across different temporal test sets. Model behavior was analyzed using global feature importance measures from Random Forests and local explanations generated via LIME. Explainability results were contextualized using real-world phishing statistics.

Results indicate that while core structural features such as the number of links and input fields remain consistently important across years, other features exhibit notable temporal drift. Changes in locally influential features align with periods of increased phishing activity associated with external factors including the expansion of 5G infrastructure and the COVID-19 pandemic. These findings suggest that shifts in attacker behavior during major technological and societal events can directly affect the decision-making patterns of phishing detection models, highlighting the importance of incorporating contextual and temporal analysis when deploying machine learning systems in dynamic and adversarial cybersecurity environments.

Paper ID 270

UNSUPERVISED BASELINE CLUSTERING AND INCREMENTAL ADAPTATION FOR IOT DEVICE TRAFFIC PROFILING

Sean Alderman; John Hastings

Abstract- The growth and heterogeneity of IoT devices create security challenges where static identification models can degrade as traffic evolves. This paper presents a two-stage, flow-feature- based pipeline for unsupervised IoT device traffic profiling and incremental model updating, evaluated on selected long-duration captures from the Deakin IoT dataset. For baseline profiling, density-based clustering (DBSCAN) isolates a substantial outlier portion of the data and produces the strongest alignment with ground-truth device labels among tested classical methods (NMI 0.78), outperforming centroid-based clustering on cluster purity. For incremental adaptation, we evaluate stream-oriented cluster- ing approaches and find that BIRCH supports efficient updates (0.13 seconds per update) and forms comparatively coherent clusters for a held-out novel device (purity 0.87), but with limited capture of novel traffic (share 0.72) and a measurable trade-off in known-device accuracy after adaptation (0.71). Overall, the results highlight a practical trade-off between high-purity static profiling and the flexibility of incremental clustering for evolving IoT environments.

Paper ID 272

PALM: PROTOTYPE-ALIGNED LABEL MANIFOLD LEARNING FOR MULTI-LABEL CLASSIFICATION WITH PARTIAL ANNOTATIONS

Yuqi Song; Xin Zhang

Abstract- Multi-label classification is a fundamental problem in computer vision, yet existing methods typically rely on fully annotated data, which is costly and often impractical to obtain at scale. In real-world scenarios, images are frequently annotated with only a subset of their true labels, resulting in incomplete and ambiguous supervision. Learning effectively under such partial annotations remains a significant challenge due to false-negative bias and error accumulation. In this paper, we propose PALM, a Prototype-Aligned Label Manifold framework for multi-label classification with partial annotations. PALM embeds visual features and labels into a shared semantic space by modeling each label as a learnable prototype refined through attention-based message passing, enabling explicit modeling of label relationships. To robustly associate images with labels under incomplete supervision, we formulate image–label alignment as an optimal transport problem with partial constraints, producing soft and globally consistent predictions. Furthermore, PALM incorporates manifold regularization to propagate reliable supervision across visually similar instances, along with an uncertainty-aware curriculum pseudo-labeling strategy to safely exploit missing labels while mitigating confirmation bias. We evaluate PALM on three widely used benchmarks—MS-COCO, NUS-WIDE, and CUB-200—under four annotation protocols, including fully annotated labels, partially annotated labels, partial positive labels, and single positive label settings. Experimental results demonstrate that PALM consistently outperforms strong baselines across all settings, with particularly significant gains under severe annotation sparsity. These results highlight the effectiveness and robustness of the proposed framework for scalable multi-label learning with incomplete annotations.

Paper ID 273

ADAPTINDEX: ADAPTIVE INDEX SELECTION FOR IOT VECTOR DATABASES

Chandrashekhar Medicherla; Milan Parikh; Chaitanya Kulkarni; Vinay Soni; Viswanathan Ranganathan

Abstract- Vector databases power critical IoT applications including smart city sensor search, industrial monitoring, and e-commerce product retrieval, yet selecting optimal index structures (HNSW vs IVF-FLAT vs Hybrid) remains a manual expert task. Misconfiguration causes severe failures: a $340K Black Friday outage from memory exhaustion, and a $2.1M manufacturing incident from poor recall. This paper presents AdaptIndex, a lightweight machine learning system that autonomously selects optimal indexes through real-time workload analysis. Our novel contribution is comprehensive feature engineering extracting 18 features across three categories: query patterns (rate, batch size, temporal peaks), data distribution (dimensionality, clustering, sparsity), and performance metrics (latency, memory, cache behavior). Using gradient boosting classification, the system achieves 89% prediction accuracy with only 12MB model size and sub-5ms inference time, enabling edge deployment. Validation on four datasets (SIFT1M, GIST1M,

Synthetic-IoT, Production-Edge) demonstrates 23-31% latency improvement over static configurations. Production deployment across 47 edge gateways spanning 12 geographic regions, serving 5 million daily queries over 90 days, achieved 29% P95 latency reduction, 99.2% uptime, and zero failures. Deployment artifacts including monitoring dashboards and anonymized traces available upon request.

Paper ID 274

AN INVESTIGATION ON THE APPLICATION OF PARETO PRINCIPLE TO WINDOWS BASED SYSTEM RESOURCE USAGE

Mihai Lazarescu; Sie Teng Soh

Abstract- This paper presents an empirical study on the compliance of Windows based systems with the Pareto Principle from the perspective of fundamental resource usage – Central Processing Unit (CPU), Random Access Memory (RAM) and Network Traffic. The aim of the work was to determine if the Pareto Principle is applicable and thus could be used to establish a baseline for a normal state for a running system which, in turn, can be used to detect an anomaly that occurs in the system. We present results from a set of experiments that systematically captured the resource usage under different conditions for four versions of Windows. The results show that in most cases, the CPU, RAM Network Traffic usage was compliant with the Pareto Principle and provides a discussion of the results.

Paper ID 275

TRPO MULTI-AGENT ACTIVE LEARNING FOR EXPLAINABLE DDOS DETECTION IN HEALTHCARE IOMT

Omar Farshad Jeelani

Abstract- Cyber-physical infrastructures and networks of the Internet of Medical Things (IoMT) are becoming more vulnerable to distributed denial-of-service (DDoS) attacks, where high availability demands are matched with a lack of labels, nonportability, and auditability. In this paper, a deployable DDoS detection model is proposed, which consists of a small supervised detector paired with a TRPO-based multi-agent decision layer to organize both label acquisition and stability-conscious explainable feature selection with explicit query and sparsity budgets. The detection problem is expressed in the form of streaming binary classification of flow/window representations of metadata-level telemetry at various network vantage points. CICIoMT2024 is an in-domain healthcare IoMT benchmark experimented and cross-domain pretrained with CICDDoS2019 and optional out-of-distribution stress tested with the CAIDA DDoS Attack 2007 trace. On CICIoMT2024, the entire system attains a different value of AUROC 0.986 and F1 0.957 with a lower rate of false-alarm (0.016) and a shorter median time-to-detect (2.0 s) compared to the supervised and uncertainty-sampling baselines. The gains of label-efficiency analysis are regularly increased irrespective of the budgets; however, the largest enhancement is observed when low-labeling regimes are considered. Pretraining also enhances fixed budget (B=300) performance, with F1 increasing between 0.926 (no pretraining) and 0.939 with combined pretraining, and FAR decreasing between 0.031 and 0.024. The explainable feature selection identifies a small set of features that are mainly characterized by protocol asymmetry, rate, and entropy features, and stability analysis shows that there are controlled feature updates without affecting interpretability.

Paper ID 276

PERFORMANCE ANALYSIS OF FIELD-LEVEL ENCRYPTION ON STRUCTURED SENSITIVE DATA USING ELLIPTIC CURVES

Anxhela Baraj; Jonatan Lerga

Abstract- lliptic Curve Cryptography (ECC) is widely used for secure communication, digital signatures, and key exchange. While its performance is well-studied in standard cryptographic contexts, less is known about its behavior on structured sensitive data at the attribute level in databases. This paper presents a hybrid experimental analysis combining cryptographic benchmarking with lightweight machine learning to evaluate ECC performance across diverse attribute types. Encryption and decryption times were measured using secp256r1 and brainpoolP256r1 curves, with X25519 included for reference key exchange. Data-level features such as length, entropy, and character

composition were analyzed using Random Forest regression to assess their influence on execution time, with the machine learning component used for interpretability rather than predictive deployment. Results show that curve choice significantly affects performance, encryption scales linearly with dataset size, and data-level characteristics have minimal impact. These findings support simplified selective encryption strategies in database systems, reducing the need for data-aware cryptographic tuning while maintaining secure and predictable performance.

Paper ID 277

EVALUATING SECURITY POLICY COMPLIANCE IN INFRASTRUCTURE AS CODE GENERATED BY LARGE LANGUAGE MODELS

Ryo Hase; Ye Wang; Toshiaki Koike-Akino; Jing Liu; Kieran Parsons; Jumpei Hato

Abstract- Infrastructure as Code (IaC) automates cloud resource provisioning, yet developing and maintaining IaC scripts remains challenging due to variations in domain-specific languages across providers. Recent advances in large language models (LLMs) offer promise for automating IaC generation, but the security policy compliance of LLM-generated IaC scripts is as important as deployability. In this work, we empirically evaluate configuration-level policy violations in LLM-generated IaC scripts using the IaC-Eval benchmark and Checkov for security policy assessment. Our results indicate that modern LLMs available as of 2025 exhibit improved syntactic correctness and better alignment with user intent, particularly when using retries and error feedback. However, security policy violations persist in generated IaC scripts, typically ranging from around five to fifteen per script across six difficulty levels defined in IaC-Eval. These results underscore the necessity of rigorous verification before deploying generated IaC scripts.

Paper ID 279

A HYBRID GRAPH-BASED ANALYSIS FRAMEWORK FOR DISCOVERING RELATIONAL FRAUD PATTERNS

Büşra Demir Sezgin; Hakan Burak Emekli

Abstract- Financial fraud detection increasingly requires analytical frameworks that go beyond isolated transaction-level predictions and capture the complex relational structures underlying fraudulent behavior. This study proposes a risk-aware, graph-based analysis framework that addresses the fraud detection problem by modeling relational and sequential transaction chains rather than isolated transactions. The proposed approach integrates continuous-valued fraud risk scores produced by supervised machine learning models into a directed transaction graph constructed based on temporal and behavioral similarities. On this graph structure, shortest-path-based search algorithms are employed to discover potential fraud chains. The distinguishing feature of the proposed framework is the integration of an adaptive algorithm selection mechanism that automatically determines the most suitable search algorithm by considering the structural properties of the transaction graph. Experimental findings demonstrate that the A* algorithm, which is automatically selected by the adaptive mechanism, achieves a path precision of 92.86%, an average risk concentration of 95.49%, and a high-risk path ratio of 83.3%. These quantitative results significantly improve the quality and interpretability of the detected fraud chains from an analyst's perspective. The results indicate that evaluating fraud detection performance solely based on individual classification outputs is insufficient, and that path-level risk density and structural consistency of transaction chains provide more meaningful and explanatory insights. The proposed path-based and risk-oriented analysis framework addresses limitations related to explainability and adaptive decision-making in the existing literature, while offering a practical solution to support analyst-driven decision processes in real-world financial systems.

Paper ID 282

ASKCMMC.AI FOR AI-DRIVEN CMMC 2.0 COMPLIANCE AUTOMATION

Hala Strohmier Berry; Aaryan R Londhe; Abdur Rahman Khan; Ronit Pawar

Abstract- Chieving CMMC 2.0 compliance is challenging for small and mid-sized enterprises. Evidence is dispersed across systems, guidance evolves continuously, and teams must translate policy into actionable technical and procedural controls. Organizations depend on manual interpretation and mapping to NIST SP 800-171. This results in high labor

costs, inconsistent outcomes, and slow readiness cycles. AskCMMC.ai addresses these challenges by automating control understanding and evidence alignment using artificial intelligence. The system employs large language models with retrieval-augmented generation and vector embeddings. These contextualize queries against authoritative standards and organization-specific documents. The system produces traceable rationales with full source attribution. It creates control ID mappings with automatic normalization and preliminary gap analyses through structured evidence retrieval. AskCMMC.ai leverages sophisticated human-in-the-loop safeguards. These include retrieval debug transparency and focus mode for analyst-controlled context. Explicit hallucination prevention constraints and persistent chat audit trails ensure AI reasoning remains interpretable and verifiable. By codifying control semantics and reusing prior assessments, AskCMMC.ai reduces analyst time, improves consistency, and accelerates remediation planning. This paper examines the methodology, system architecture, and evaluation across representative compliance tasks. It focuses on accuracy, latency, and reviewer effort, as well as human-in-the-loop safeguards, data security, and model risk management. Evaluation across 120 compliance-related queries demonstrated 90.8% fully correct responses, with average response latency of 1.95 seconds. AskCMMC.ai reduced analyst task completion time by approximately 81% across representative compliance activities.

Paper ID 284

YOLO11s Optimization for Minutiae Detection

Ahsan Islam; Shanika Perera Wadduwage; Erasmus Mfodwo; Van Vung Pham

Abstract- Fingerprint minutiae are only a few pixels wide and are easily degraded when images are resized or when feature maps are repeatedly downsampled inside a detector, which makes reliable minutiae localization difficult in forensic settings. We propose \textit{YOLO11s Minutiae}, a customized YOLO11s-based, single-pass detector designed for minutiae-scale targets. The key change is an added high-resolution detection stage so minutiae can be detected before fine ridge detail is removed by downsampling. We further strengthen early feature extraction to better encode ridge micro-texture and add a mid-level context pooling module to incorporate local ridge-flow neighborhood cues and suppress ridge-noise false alarms. Training uses class-balanced oversampling, conservative augmentations that preserve ridge geometry, and an AdamW optimizer with a cosine learning-rate schedule. In the Minutiae Leple dataset, YOLO11s provides the best baseline among YOLO11 variants, and the high-resolution customization improves recall and mAP@0.50.Compared to the baseline of YOLO11s (mAP@0.50 = 0.727), the high-resolution P2 detection stage improves mAP@0.50 to 0.735 and achieves the highest recall of 0.745, demonstrating improved sensitivity to structures on the minutiae-scale. Overall, the proposed approach provides an efficient and practical pipeline for minutiae detection in forensic workflows.

Paper ID 285

Stack Buffer Overflow Risk Analysis for PX4-Controlled Commercial UAVs

Hala Strohmier Berry; Chris Clark; Kadin DeMesme; Marco Paolo Saavedra; Travis Samatov

Abstract- This research investigates cybersecurity vulnerabilities of commercial drones by checking the resilience of PX4 autopilot software to memory corruption and denial-of-service attacks. The study evaluates how malformed inputs, message flooding, and process-level interference affect flight stability and mission execution. Experiments are conducted using PX4's SITL environment integrated with QGroundControl and jMAVSim, with autonomous missions executed while adversarial inputs are injected through telemetry interfaces and local process interactions. The system is measured through operational metrics including mission completion rate, autopilot uptime, flight mode stability, and recovery behavior. Results show that while PX4 demonstrates strong resilience to protocol-level buffer overflow and flooding attacks, a critical vulnerability was identified at the process level, the absence of handlers for fatal UNIX signals (SIGSEGV, SIGBUS, SIGFPE, SIGILL) allows any locally exe- cuting process to crash the autopilot with 100% reliability, reducing mission completion rates to zero. A minimal patch installing these handlers eliminates the vulnerability entirely, restoring mission completion to 100% under the same attack conditions. These findings provide actionable information on strengthening cyber resilience of UAV and motivate mitigation strategies that address not only input validation of the application-layer but also the attack surfaces of the operating system.

Paper ID 293

BOUNCED CHECK RISK PREDICTION VIA MULTI-OBJECTIVE HYPERPARAMETER OPTIMIZATION: BALANCING MACRO F1 AND BUSINESS UPLIFT

Kerem Kaya; Emir Çetin Memiş; Seyit Ertuğrul; Hakan Karamanlı; Erkal Bıyıklıoğlu; Yassine Dria; Semen Son-Turan; Nazlı Torağanlı-Karamollaoğlu; Tuna Çakar

Abstract- In this study, the risk of bounced checks is formulated as an imbalanced binary classification and decision-support problem for operational risk screening. We model individual and corporate customers separately to capture segment-specific risk behavior and decision thresholds. We compare strong tree-ensemble and linear baselines (e.g., gradient boosting, random forests, logistic regression). We use multi-objective optimization to balance macro F1 with a utility-based business objective (uplift / net-result) defined relative to an accept-all baseline and choose decision thresholds to maximize expected net outcome under a minimum macro-F1 quality constraint. Gradient-boosted trees (LightGBM and XGBoost) perform best across both segments. In the individual segment, macro F1 ranges from 0.86 to 0.87; on the 3-month holdout, LightGBM achieves 1.23% gross sunk rate and a +2.09% relative net-result increase compared to the accept-all baseline. In the corporate segment, macro F1 ranges from 0.88 to 0.89; on the same holdout, LightGBM achieves 1.48% gross sunk rate and a +34.21% relative net-result increase compared to the accept-all baseline. These results indicate that segment- aware, utility-driven model selection can improve both forecasting quality and financial outcomes.

Paper ID 294

A COMPARATIVE STUDY OF MACHINE LEARNING MODELS FOR MICRO-SEGMENT CREDIT RISK PREDICTION

Kerem Kaya; Emir Çetin Memiş; Seyit Ertuğrul; Hakan Karamanlı; Yassine Dria; Semen Son-Turan; Nazlı Torağanlı-Karamollaoğlu; Tuna Çakar

Abstract- We studied the early prediction of credit risk in the micro segment as a binary classification problem. The dataset was provided by a financial company. The tables were combined with the common identifier feature to create the final dataset consisting of over 150 thousand rows and 92 columns. Since the target distribution was unbalanced, undersampling was applied during the model training phase. In the preprocessing step, date fields were reduced to month-period level, inflation data was added to the dataset on a month-by-month basis. Feature engineering was structured in two main groups: payment history flags for up to 18 months based on previous payment performance history of customers and limit and balance aggregations derived according to credit type and time windows (L3M–L36M). Monetary aggregations were normalized with inflation information to produce various inflation columns. We evaluated Logistic Regression, Random Forest, and Extra Trees, and three gradient-boosting models (XGBoost, LightGBM, CatBoost) for comparative performance. Cross-validation results showed that boosting-based models tended to perform better, and macro F1 scores were obtained at XGBoost 0.81939, LightGBM 0.82157, CatBoost 0.82031; LightGBM was marginally higher than the other boosting models. With RFECV, the optimal number of features for XGBoost was found to be 70, and it was shown that similar performance could be maintained with a more compact feature set.

Paper ID 295

AN LLM-POWERED API TESTING FRAMEWORK BASED ON STRUCTURAL SIMILARITY ANALYSIS

Anıl Sezgin; Tuğberk Kocatekin; Mert Yağcıoğlu

Abstract- In the context of the accuracy, dependability, and security of application programming interfaces (APIs), the importance of API testing is undeniable. Nevertheless, the effort required to comprehensively test APIs is significant, especially when creating valid API test cases. This study seeks to develop a framework that automates the generation of API test cases while ensuring the structural validity and functional correctness of the API test cases. The study aims to reduce the effort required by human testers to generate API test cases. The study proposes a framework that utilizes a large language model (LLM) to generate API test cases based on the concept of structural similarity. The proposed framework utilizes the LLM to generate API test cases while ensuring the functional correctness of the API test cases. To validate the structural similarity between the API test cases generated by the proposed framework, the study introduces the Jaccard similarity-based API test case structural validation mechanism. The study utilizes a multi-domain API ecosystem that contains more than 40 endpoints to validate the effectiveness of the proposed API test case generation framework. The study found that the proposed API test case generation framework achieved a high success

rate of 84% in generating valid API test cases for different API domains. The study found that the proposed API test case generation framework using the LLM achieved a high mean Jaccard similarity coefficient of 0.827, indicating the high level of structural alignment between the API test cases generated by the LLM and the API specifications. Therefore, the study concludes that the proposed API test case generation framework using the LLM is effective in automating the generation of API test cases while ensuring the accuracy, dependability, and security of the API test cases.

Paper ID 296

EVIDENCE-QUALITY TELEMETRY FOR CLOUD INCIDENT RESPONSE: DETECTING GAPS, DRIFT, AND INTEGRITY FAILURES IN LARGE-SCALE OBSERVABILITY PIPELINES

Chalapathi Koneni; Sanjay Lokula

Abstract- High-velocity streams of telemetry are now essential in modern cloud incident response systems but more than three-quarters of respondents say that they make critical decisions with incomplete or inconsistent data. The inherent systemic inability of pipelines improves the investigation capability and prolongs Mean Time to Resolution (MTTR). To mitigate the core vulnerability associated with the evidence-quality telemetry, this research suggested a framework to implement evidence-quality telemetry, and data trustworthiness as the first-class, verifiable system property. The overall aim was to establish and prove a scalable data pipeline architecture that executes rigorous indispensable considerations of completeness, steadiness, purity, and freshness on cloud incident response systems. Multi-regional pipeline A pipeline with Apache Flink processing 15 TB of telemetry every day based on Kubernetes was implemented. It was stress-tested on a set of strict simulation frameworks that were churning Data Gaps, Concept Drift and cryptographic Integrity Failure. The architecture achieved a high Recall of 0.983 and 99.1% success rate in integrity recovery ensuring evidence-quality is met. The operation cost of this performance was a limited operation latency overhead of 90 percent (45 ms) and 10 percent throughput loss, which was considered reasonable. The investigation revealed that under triage accuracy, the Max Integrity Issue Flag had a strong negative correlation (r = -0.7307), and thus data integrity is significant in comparison to simple completeness. The framework offers a scalable and robust engineering blueprint, which will turn best-effort telemetry into an evidence-quality asset of autonomous security decision-making. In the future, efforts will be aimed at counteracting the base overhead of latency by investigating hardware-accelerated hashing in Smart NICs and extending this validation strategy to federated multi-cloud environments.

Paper ID 297

ADVANCING FIRE AND SMOKE DETECTION IN FORENSIC SURVEILLANCE: A STUDY WITH CONTEXT AWARE AUGMENTATION AND OPTIMIZATION

Paul Isibor; Isah Mohammed; Bright Jiwueze; Pamela Kirui; Osayomore Aigbogun; Van Vung Pham

Abstract- Early smoke detection in CCTV footage is important for forensic reconstruction and timely response, but faint, hazy, and distant smoke remains difficult to detect in practice. This paper studies a family of YOLOv11 detectors for fire and smoke and frames the task as a forensic problem, where missing early smoke cues is more harmful than raising additional false alarms. We systematically compare CSP based and no CSP YOLOv11 variants and conduct controlled ablation studies on simple, smoke focused image perturbations using a curated fire and smoke image dataset. Our results show that both architectural design and model tuning strongly affect performance. CSP tuned models consistently improve accuracy over the default YOLOv11 settings, while replacing CSP blocks with plain convolutional layers achieves the highest mAP on the evaluated dataset. However, these no CSP models require longer training due to deeper convolution stacks and weaker gradient flow. Across the evaluated variants, a tuned YOLOv11s no CSP model trained with haze/contrast style augmentations achieves the best overall performance. It attains precision 0.61, recall 0.56, mAP50 0.58, and mAP50:95 0.24, compared to the default YOLOv11s baseline with precision 0.41, recall 0.46, mAP50 0.40, and mAP50:95 0.15. These results indicate that removing CSP and carefully tuning training improves detection stability while maintaining real time performance.

Paper ID 298

EVENT-DRIVEN AGENTIC SOC (ED-ASOC): SUPERVISOR-BASED LLM FRAMEWORK FOR DYNAMIC INCIDENT RESPONSE AND SOAR ORCHESTRATION

Enes Özgözler; Asaf Varol; İhsan Tanrıverdi

Abstract- As cyber threats grow more sophisticated, Security Operations Centers (SOC) face mounting pressure from the gap between alert volumes and the speed of traditional incident response. Analysts struggle with cognitive overload caused by high false-positive rates, leading to alert fatigue and missed critical threats. Current SOAR platforms, despite automating routine tasks, rely on polling-based data retrieval and static playbooks that cannot adapt to polymorphic or zero-day attacks. This paper presents the Event-Driven Agentic SOC (ED-ASOC) framework, built around a Supervisor Agent that enables the shift from scripted automation to autonomous orchestration. By replacing periodic polling with Webhook-based real-time ingestion, the framework cuts detection-to-analysis latency from minutes to sub-second levels. The Supervisor Agent enriches raw telemetry with organizational context through Retrieval-Augmented Generation (RAG), producing adaptive responses that move beyond rigid rules. The agent can execute remediation actions, including IP blocking, endpoint isolation, and user suspension, via defined API schemas, closing the gap between detection and containment. A tiered Human-in-the-Loop (HITL) mechanism ensures that high-impact decisions require analyst approval. Experimental evaluation using FortiSIEM and FortiSOAR in a controlled testbed demonstrates a 73% reduction in Mean Time to Response (MTTR) and 89% decrease in detection latency compared to polling-based baselines, while reducing Tier-1 analyst workload by approximately 45%.

Paper ID 299

VOLATILE MEMORY FORENSICS OF TAILS OS IN A VIRTUALIZED ENVIRONMENT

Nurettin Senol; Jayden Thai; Semih Cal; Ahmet Aydogan

Abstract- Operating systems focused primarily on privacy, such as Tails, are designed to minimize forensic traceability through the elimination of permanent storage and reliance on a non-persistent execution model. This study examines the availability of forensic traces retained in volatile memory, specifically within the non-persistent, virtualized environment of Tails. A controlled experiment involving normal user activity was conducted, followed by live RAM data acquisition and analysis using FTK Imager, Volatility 3, and keyword-based extraction techniques. The results demonstrate that traces related to file system activity, network communications, software execution, and external device interaction persist in temporary memory. These findings highlight Tails' forensic visibility beyond its documented threat model and underscore the evidentiary value of live memory analysis in privacy-focused operating systems.

Paper ID 300

AGENTIC FRAMEWORK FOR CONTINUOUS REVENUE GOVERNANCE ACROSS CRM, CPQ, AND CLM

Sivasai Nadella

Abstract- In this paper, we propose a Multi-Agent Artificial Intelligence architecture for to shape and regulate the end-to-End revenue lifecycle journeys, particularly consolidating the operational divide between CPQ (Configure Price Quote), CRM (Customer Relationship Management) and CLM (Contract Lifecycle Management). With businesses relying more and more on fragmented platforms to collect the revenues earned from their customers, data leaks and lost time between systems has become a substantial drag. This paper proposes a new decentralized architecture in which autonomous agents -- trained to negotiate, to comply and to reconcile -- interplay in order to offer correctness of data and policy enforcement throughout the entire revenue chain. We use a dummy dataset of 469 data instances that spans across complex B2B sales cycles with different price tiers, clauses in contracts and renewal terms. The experimental instantiation uses Python agent systems with TensorFlow as a decision-making logic and simulates a high-velocity sales environment. Results show that multi-agent coordination substantially lowers revenue leak- age and increases contract compliance rates over old-fashion rule based integration techniques. Where revenue governance is seen as a dynamic and multi-stakeholder environment, this approach provides scalable solution for modern enterprises that want to optimize their quote-to-cash velocity in the context of stringent regulatory compliance. We also discuss deployment considerations (auditability, human-in-the-loop approvals, and legacy integration) and outline limitations and future validation on real-world datasets.

Paper ID 302

INTEGRATING GENERATIVE AI INTO RETAIL CHECKOUT SYSTEMS: A CASE STUDY IN CLOUD AND APPLICATION INTEGRATION

Gopalakrishnan Venkatasubbu; Rajgopal Devabhaktuni

Abstract- The retail POS checkout experience has long been a transaction barrier – an untapped well of potential value. Illustrative Case: A Generative AI model for retail. The following case introduces Hierarchical Neural Story Generation (HNSG) job [21] into a high-stake business of checkout systems of medium size retailer "OmniRetail". The goal was to better up-sell to customers and increase the real-time detection of fraudulent orders without adding any latency. This paper describes the architecture design and implementation of a new integration model based on event driven microservices and hybrid cloud. While the AI layer was a relatively simple organic paragraph generator, the technical debt to support running in a very high-availability environment, as well as integration with legacy Point-of Sale (POS) systems added significant time to this project. Our study is based on packet traces collected for 484 checkout terminals over six months, as well as a dataset containing transaction logs and anonymized customer loyalty data, system performance measures. Its architecture was built on Python microservices, an Apache Kafka event bus for decoupling, and a Kubernetes cluster for scalable AI model deployment. The results show that, running under an asynchronous event-driven architecture, the complex generative model can be queried in real-time and potentially improve the upsell conversion rate by 22% and effectively detect emerging fraud patterns with sub-second transactional response time.

Paper ID 311

NEURO-SYMBOLIC GRAPH AUTOENCODERS WITH RARE PATTERN MINING FOR PROVENANCE-BASED ANOMALY DETECTION

Sidahmed Benabderrahmane; Asif Tauhid; Mohamad Altrabulsi; Ahamed Foisal; Talal Rahwan

Abstract- Advanced Persistent Threats (APTs) are sophisticated, long-term cyberattacks that are difficult to detect because they operate stealthily and often blend into normal system behavior. This paper presents a neuro-symbolic anomaly detection framework that combines a Graph Autoencoder (GAE) with rare pattern mining to identify APT-like activities in system-level provenance data. Our approach first constructs a process behavioral graph using k-Nearest Neighbors based on feature similarity, then learns normal relational structure using a Graph Autoencoder. Anomaly candidates are identified through deviations between observed and reconstructed graph structure. To further improve detection, we integrate a rare pattern mining module that discovers infrequent behavioral co-occurrences and uses them to boost anomaly scores for processes exhibiting rare signatures. We evaluate the proposed method on the DARPA Transparent Computing datasets and show that rare-pattern boosting yields substantial gains in anomaly ranking quality over the baseline GAE. Compared with existing unsupervised approaches on the same benchmark, our single unified model consistently outperforms individual context-based detectors and achieves performance competitive with ensemble aggregation methods that require multiple separate detectors. These results highlight the value of coupling graph-based representation learning with classical pattern mining to improve both effectiveness and interpretability in provenance-based security anomaly detection.

Paper ID 313

PERSONALIZED PRODUCT RECOMMENDATION IN E-COMMERCE USING MACHINE LEARNING TECHNIQUES

Navdeep Singh

Abstract- Businesses may boost customer satisfaction, engagement, and conversion rates with personalised product recommendations, which are an essential part of contemporary e-commerce. The proposed paper will suggest a hybrid ensemble machine learning system comprising Random Forest, XGBoost, and LightGBM along with stacking, soft voting, and bagging approaches to create accurate and context-sensitive recommendations. The framework takes advantage of transactional, behavioral and contextual data engineered features and uses SMOTE based balancing and ROC based thresholds to address class imbalance and enhance predictive reliability. An F1-score of 99.40, an accuracy of 97.69, a precision of 99.01, a recall of 98.07 and a combined customer-product dataset demonstrate that the proposed model outperforms traditional ML and DL techniques. Also, explainability analysis which is based on SHAP verifies

that recommendations are based on features including sentiment, product category, interaction type, and product rating which offers transparency. The findings underscore the capability of the framework to provide powerful, intuitively explainable, and expandable individualized suggestions to large-scale e-commerce platforms.

Paper ID 317

PERFORMANCE ANALYSIS OF A CLOUD-NATIVE WEB APPLICATION DEPLOYED ON KUBERNETES

Dharmendra Ahuja

Abstract- Container orchestration is becoming important in cloud-native web applications to realize high-scale and efficient control of resources in the face of variable workloads. The article provides an experimental work on Kubernetes Horizontal Pod Autoscaling (HPA) and a cloud-native e-commerce recommendation application that is deployed on Amazon Elastic Kubernetes Service (EKS). The microservices application is tested in the setting of production-like loading condition to test the impact of the idea of autoscaling on performance. The experiment is based on scalability at scale, response latency, throughput, CPU usage, scaling behavior of pods, and node-hour usage, comparing both statical and autoscaled deployments. Experimental results indicate that HPA doubles throughput between 50 and 100 requests per second, halves peak response latency exceeding 600ms down to about 110ms, slashes average CPU utilization per node in the neighborhood of 90% to 65%, and cuts node-hour usage by almost 50%. Prometheus is used to collect performance data and Grafana is used to visualize performance data to facilitate repeatable analysis. The findings show that Kubernetes HPA is highly useful in enhancing performance scalability and cost-efficiency in practical cloud-native applications.

Paper ID 333

HIGH ACCURACY IS NOT ENOUGH: EPISTEMIC BIAS IN MACHINE LEARNING TASK FORMULATION

Grazia Garzo; Alessandro Palumbo

Abstract- While predictive accuracy is a central evaluation metric in Machine Learning, it does not, by itself, guarantee the validity of a learning task. This paper presents a case study demonstrating that Machine Learning-based classification systems can achieve high performance even when the inferred relationships lack epistemic validity or normative justification. Specifically, we construct an image-based classification task that operationalizes a criminality classifier inspired by Lombrosian physiognomy, a historically influential but scientifically discredited framework. Within that theory, criminality is assumed to be predictable from specific facial characteristics, understood as indicators of an individual's innate predisposition to criminal behavior. Despite relying on synthetic data and artificial labels, the resulting model achieves high predictive accuracy. This outcome highlights a critical limitation of accuracy-driven evaluation: performance may conceal epistemic and normative bias introduced upstream, at the level of task definition and label semantics, thereby lending legitimacy to discriminatory or otherwise unacceptable forms of inference.

Paper ID 334

FAKESPEECH: LLM-DRIVEN SEMANTIC MANIPULATION AND VOICE CLONING FOR REALISTIC DEEPFAKE SPEECH BENCHMARKING

Alaa Alsaeedi; Amal AlMansour; Amani Jamal

Abstract- Recent advances in large language models and voice cloning have enabled deepfakes that alter semantic meaning while keeping the speaker's tone and visual identity consistent. Existing datasets mainly capture surface-level acoustic or prosodic artifacts, overlooking semantic manipulations that are harder to detect. This paper introduces FakeSpeech, an audio-visual deepfake dataset designed to benchmark semantic-level speech manipulation under realistic visual alignment. The dataset contains 970 talking-face clips (485 real and 485 fake) derived from FakeAVCeleb videos. Fake samples were generated by rewriting transcripts using GPT-4.5 and re-synthesizing voices through ElevenLabs, guided by the Phantom Reading technique to preserve lip synchronization. When evaluated with an audio-only baseline, FakeAVCeleb achieved 0.85 accuracy, whereas FakeSpeech dropped to 0.67, indicating that the added speech manipulation significantly increased dataset complexity and detection difficulty. Overall, FakeSpeech bridges the gap between acoustic and multimodal forgeries, offering a realistic benchmark for studying semantic–

prosodic alignment and advancing deepfake detection beyond surface artifacts. Unlike prior benchmarks, we keep identity fixed. The video frames are real and unchanged, and the voice stays the same speaker. Only the spoken content changes. In other words, prior work often studies how a person speaks, while we study what the person says in the same familiar way. This content-level manipulation is therefore harder to detect.

Paper ID 336

GRADIENT-ACCELERATED COSMOLOGICAL INFERENCE: A JAX-BASED FRAMEWORK FOR DIFFERENTIABLE BAYESIAN COMPUTATION IN ASTROPHYSICS

Mayank Jha

Abstract- This paper introduces a novel, fully differentiable computational framework built on JAX for accelerating and stabilising Bayesian inference in high-dimensional cosmological parameter estimation. By implementing automatic differentiation, GPU-accelerated just-in-time compilation, and vectorised probabilistic programming, our library enables exact gradient computations for complex cosmological likelihoods previously infeasible with finite-difference methods. We demonstrate how this supports advanced machine learning inference techniques, including Hamiltonian Monte Carlo, No-U-Turn Sampling, and Stochastic Variational Inference with normalising flows. In a Dark Energy Survey Year 1 3×2pt case study, we show order-of-magnitude improvements in sampling efficiency (achieving 10.3× higher effective samples per second for Ωc compared to traditional MCMC) and enable rapid Fisher matrix estimation without tuning. This work bridges differentiable programming and astrophysical inference, offering a scalable, gradient-based paradigm for next-generation cosmological analysis.

Paper ID 338

XGBOOST-ENHANCED RECURRENT HYBRID MODELS FOR WEARABLE INERTIAL NAVIGATION IN GNSS-DENIED ENVIRONMENTS

Vedant Singh; Nagalakshmi S R; Tushar Swami; Yash Sinha; Sk Hithasree

Abstract- When satellite signals are lost or degraded, GPS based navigation stops working in urban canyons, tunnels, dense forests, and indoors. This makes it hard to keep track of things in everyday situations like fitness tracking and important situations like emergency response. Wearable devices are very popular for tracking health and activity, but they rely heavily on GPS to find their way around. This makes them less useful and less enjoyable in areas where GNSS is not available. To overcome this limitation, this study introduces a GPS-independent Inertial Navigation System (INS) for wearable devices, augmented with Machine Learning (ML) and Kalman-filtered sensor fusion to ensure dependable real-time positioning. Generic deep learning models like LSTM, ResNet-1D, and Transformer architectures were compared to hybrid models that combine advanced regressors like GRU-SVR, GRU-XGBoost, and BiGRU-XGBoost with gated recurrent encoders. According to the experimental results, the recommended XGBoost-Enhanced Recurrent Hybrids have the best positional accuracy, reducing drift from 162 m/km (GRU SVR) to 16 m/km while maintaining an inference latency of less than 5 ms. These findings demonstrate how ensemble-driven recurrent models effectively balance speed and accuracy, enabling GPS-free navigation in areas without GNSS.

Paper ID 339

MALWARE CLASSIFICATION ON PE FILES USING DEEP NEURAL NETWORKS

Daniel Vilaça; Lu´ıs Ferreira

Abstract- Traditional malware detection methods relying on signature-based matching are increasingly ineffective against modern obfuscation techniques such as polymorphism and pack ing. To address this limitation, this study explores the application of Deep Learning for the static analysis of Portable Executable (PE) files, transforming raw binary sequences into 2D grayscale visual representations to establish a high-accuracy foundation for automated malware triage. Two Convolutional Neural Net works (CNN) architectures were implemented and compared, namely a CNN as a baseline model trained from scratch, and a VGG16model using transfer learning from pre-trained ImageNet weights. Both models were evaluated on the Malimg dataset (9,339 samples across 25 families). To validate robustness against class imbalance, a Stratified K-Fold Cross-Validation and a Two Stage Fine-Tuning strategy were employed.

Results indicate that while the baseline CNN suffered from gradient instability and failed to detect minority classes (0 percent recall on Autorun.K), the optimized VGG16 model demonstrated superior stability and generalization. The proposed approach achieved a peak classification accuracy of 99.33 percent, effectively solving the minority class data scarcity problem. These findings validate the VGG16 architecture as a robust visual engine for the proposed multi-stage defense pipeline.

Paper ID 340

DIGITAL FORENSICS APPLICATIONS OF ELECTRIC NETWORK FREQUENCY

Fatih Yaman; Ümühan Özkaynak; Fatih Özkaynak

Abstract- Electric Network Frequency (ENF) analysis has emerged as a valuable tool in digital forensics for verifying the temporal authenticity and integrity of multimedia recordings. This study presents a usability-oriented evaluation of ENF-based forensic analysis, focusing on real grid frequency measurements rather than algorithm-centric performance benchmarks. A scenario-driven methodology is proposed to assess ENF applicability in common forensic tasks such as timestamp verification and temporal consistency analysis. Experimental evaluations are conducted using multiple anonymized ENF reference datasets recorded on the same power grid at distinct, non-overlapping time intervals. Correlation-based matching is employed to examine temporal alignment and discriminative behavior of ENF signals. The results demonstrate that ENF exhibits clear time-dependent characteristics, yielding high correlation for same-interval comparisons and significantly lower correlation across different intervals. These findings confirm the suitability of ENF as an intrinsic temporal fingerprint under appropriate conditions. The study further highlights practical limitations, including ambiguous outcomes arising from weak ENF variation or marginal correlation values, and emphasizes the importance of cautious interpretation. Overall, this work positions ENF analysis as a transparent and interpretable decision-support mechanism that complements existing digital forensic techniques and clarifies the operational boundaries of ENF-based evidence.

Paper ID 348

INCIDENT-AWARE CI/CD PIPELINES: LEARNING FROM PRODUCTION FAILURES TO PREVENT CERTIFICATE ROTATION DRIFT

Lakshmi Vidya Peri; Yogesh Thanvi; Yogeesh Kunigal Gangaiah

Abstract- Modern microservices architectures rely on mutual TLS (mTLS) for inter-service authentication, yet certificate authority (CA) rotation remains a leading source of production incidents due to trust bundle drift across distributed regions. Traditional CI/CD pipelines test against static configurations and lack awareness of failure modes previously observed in production, allowing the same class of incidents to escape repeatedly. This paper presents an incident-aware CI/CD pipeline that closes this gap by integrating production incident signals into automated testing through typed guardrails. We validate the approach on a CA rotation trust drift scenario across a multi-region Kubernetes architecture with nine services in three regions. Experimental evaluation in five paired trials (ten total runs) demonstrates that incident-aware pipelines block 100% of escaped incidents versus 0% for static pipelines ($p = 0.0079$, Fisher's exact test). The mean CI overhead of 20% is not statistically significant ($p = 0.8413$, Mann-Whitney U test), and incident-aware pipelines eliminate all deployment restarts by preventing runtime failures. These results show that encoding production failure modes as CI guardrails effectively prevents incident recurrence without meaningful impact on development velocity.

Paper ID 350

PORTING AND EVALUATING RETURN-ORIENTED PROGRAMMING DEFENSES IMPLEMENTED BY THE OPENBSD OPERATING SYSTEM

Jenna Esposito; Aaila Arif; Raul Cortinas; Brian Callahan

Abstract- Return-oriented programming (ROP) continues to be a serious attack vector nearly a decade and a half after first being disclosed. Techniques to circumvent this style of attack include both hardware- and software-based solutions. CPUs that lack the hardware capabilities must rely solely on software-based solutions.

In 2019, the OpenBSD project introduced a number of software-based anti-ROP mitigations in their copy of the LLVM compiler suite. They claimed significant reductions in unique gadgets in the OpenBSD kernel and libc with zero to minimal drawbacks. Such solutions have not been ported elsewhere nor integrated into upstream LLVM.

In this paper, we port two of OpenBSD's software-based anti-ROP mitigations for x86_64 CPUs, alternative register selection and compile-time instruction rewriting, to the FreeBSD operating system. We then measure their effects on reduction of unique gadgets, binary size, and runtime performance.

We were able to corroborate some the claims of minimal size and runtime penalties; gadget reduction was substantially less than claimed and less effective than claimed. Our investigation demonstrates the potential difficulties of seemingly obvious compiler-based anti-ROP techniques.

Paper ID 351

EXTENDING THE TMMI FRAMEWORK FOR SECURE TESTING OF AI AGENTS

Lakshmi Vidya Peri

Abstract- The Test Maturity Model integration (TMMi) defines five maturity levels and 16 process areas for software test process improvement, but these process areas do not address the testing challenges introduced by AI agents: non-deterministic reasoning, dynamic delegation chains, natural-language attack surfaces, persistent memory corruption, autonomous tool invocation, and multi-agent trust exploitation. Even organizations at the TMMi Level 5 lack process areas, practices, and assessment criteria for agentic systems. This paper extends TMMi by proposing 16 Agent Security Test Process Areas (ASTPAs) distributed across maturity Levels 2 through 5, mirroring TMMi's 5+5+3+3 structural distribution. ASTPAs are grounded in the OWASP Top 10 for Agentic Applications 2026 and the OWASP Agentic AI Threats and Mitigations taxonomy. The extension includes three formal testable security properties, seven Level 4 metrics, and TAMAR-compatible assessment criteria, enabling organizations to evaluate and improve their agent security testing maturity within the established TMMi framework.

Paper ID 355

VERSION-CONTROLLED DECENTRALIZED FIRMWARE INTEGRITY VERIFICATION WITH ON-CHAIN ROLLBACK PROTECTION FOR CYBER-PHYSICAL SYSTEMS ON ETHEREUM

Maruf Farhan; Usman Butt; Madhuki Rajapakshe; Rejwan Bin Sulaiman

Abstract- Firmware update mechanisms represent a critical attack surface in cyber-physical and Internet-of-Things (IoT) systems, as poor-quality firmware introduced into these systems can provide persistent security vulnerabilities. Although digital signature schemes confirm the authenticity of firmware, they do not mitigate so-called rollback attacks, where an attacker re-installs the older vulnerable versions. Existing approaches on blockchain enhance integrity and auditability but frequently have no contract-level method of enforcing version progression nor imposing recurring transaction fees on devices. This paper introduces a decentralized firmware integrity verification framework to implement strict anti-rollback protection by Ethereum smart contracts. The system maintains a tamper-evident on-chain registry of firmware metadata an rejects downgrade attempts through consensus-enforced monotonic version control. To enhance the scalability, a dual-mode verification design is introduced where routine integrity checks can be performed off-chain at zero gas cost, where the on-chain (audited) verification can be used when immutable proof of such a check is required. The framework is applied and tested on the Ethereum Sepolia test network. On Ethereum Sepolia, deployment consumed 913,425 gas, firmware registration averaged 106,878 gas, audited verification averaged 37,052 gas, and measured time-to-first-confirmation averaged 97.71 s (12–408 s). Results indicate that it is an effective rollback prevention technique, has low transaction overhead for firmware registration, and is practical for large-scale deployment of IoT.

Paper ID 356

POST-QUANTUM CRYPTOGRAPHY FOR WEB AUTHENTICATION PROTOCOLS: A SYSTEMATIC REVIEW OF OAUTH 2.0, OPENID CONNECT, AND SAML MIGRATION

Ravindu Dissanayake; Harindu Wijesinghe; Jaith Vindinu; Kulanga Jayasinghe; Kavinga Abeywardena; Amila Senarathne

Abstract- OAuth 2.0, OpenID Connect (OIDC), and SAML rely on classical public-key primitives such as RSA and ECDSA, which are vulnerable to quantum attacks via Shor's algorithm. This systematic review examines migration of these protocols to Post-Quantum Cryptography (PQC) following the 2024 NIST standardization of ML-DSA and ML-KEM. We map cryptographic dependencies across all three protocols, evaluate NIST-standardized algorithms for authentication use cases, and analyze practical migration challenges. Token size explosion, with ML-DSA-65 signatures approximately 52 times larger than ECDSA P-256, represents the dominant implementation barrier, compounded by incomplete JOSE standardization and limited ecosystem maturity. Missing formal security proofs and federation migration frameworks are identified as critical priorities before production deployment.

Paper ID 363

SAFEKID SCAN: EARLY DETECTION OF DIGITAL ADDICTION IN MINORS

Anjana H.H.H; Weerakkodi Y.P; Gimhani J.M.K.P; Nethmini M.A.N; Jenny Krishara; Poorna Panduwawala

Abstract- The lack of awareness of the consequences of the rapid development of social media use among children and adolescents has increased the risk of digital addiction, impacting academic performance, behavior, and mental health. Current detection tools mainly rely on surveys or screen time data, which are subjective and unsuitable for real-time intervention. To overcome these limitations, this paper presents a smart and privacy-enhancing architecture to early identify and manage social media addiction in minors aged 10–18. The system combines multimodal image, caption, and hashtag processing in English and Sinhala to generate addiction risk scores. It also integrates complaint-based risk detection, AI driven decision support, blockchain-based information integrity, and IoT-based behavioral monitoring. Experimental evaluation achieved 95.97% validation accuracy for media analysis, 85.56% for complaint analysis, and 96.97% for IoT-based monitoring. The framework supports accurate risk identification, timely intervention, and safer digital environments for children and adolescents.

Paper ID 365

QSIGNATURE 1.0: A DYNAMICAL REGIME CLASSIFICATION FRAMEWORK FOR CAUSAL TIME SERIES DATA

Ahmad Muhammad, Salim Jibrin Danbatta; Muhammad Abubakar Isah; Ibrahim Yahaya Muhammad; Sulaiman Sulaiman Ahmad; Abdelrahman Ghozlan

Abstract- The causal response of a system to external perturbation encodes its governing dynamical signature. When only the output response \(R(t)\) is observable without knowledge of the input or a parametric model, inferring the system class remains a fundamental challenge. This work introduces two persistence timescale estimators, \(\tau_s\) and \(\tau_u\), which yield two scalar diagnostics, \(\Delta_{su} = (\tau_s - \tau_u)/\tau_u\) and \(R_{su} = \tau_s/\tau_u\). These diagnostics provide a fingerprint of linear time-invariant dynamical systems. Evaluation on 49 canonical systems yields five distinct regions in the \((\Delta_{su}, R_{su})\) plane (QSpace): exponential monotonic, fractional, underdamped, weakly damped, and conservative oscillatory. The framework admits direct physical interpretation: \(\Delta_{su}\) varies monotonically with the damping ratio \(\zeta\), and negative \(R_{su}\) signals centroid reversal in weakly damped systems. Crucially, we applied the framework to real-world forensic data, analyzing 20 malware execution traces from the ECU-MALNETT corpus, which contains 20,500 samples across 58 families, focusing on APT28, APT29, and Lazarus families, Packet-rate analysis reveals distinct behavioral fingerprints: APT29 exhibits weakly damped beaconing \((\Delta_{su}\approx -1.14, R_{su}\approx -0.14)\), APT28 shows extreme oscillatory activity \((\Delta_{su} < -2.0)\), and Lazarus spans multiple regimes, all distinguishable purely from timing signatures without deep packet inspection. The QSignature library implementing these estimators is available on GitHub.

Paper ID 366

AN INTERPRETABLE MULTIMODAL AI FRAMEWORK FOR SEVERITY-AWARE AND GUIDELINE-ALIGNED TREATMENT RECOMMENDATION IN CHRONIC SPONTANEOUS URTICARIA

Samidi Jayawickrama; Ramindu Nimes; Thewan Damnidu; Pradicksha Pradeepraj; Dharshana Kasthurirathna; Samanthi Siriwardana

Abstract- Chronic Spontaneous Urticaria (CSU) is a recurrent inflammatory disorder characterized by fluctuating disease severity and variable therapeutic response, requiring continuous monitoring and guideline-aligned treatment. Current artificial intelligence approaches for urticaria primarily focus on image-based lesion analysis and do not integrate heterogeneous clinical evidence for therapeutic decision support. In this study, we propose an interpretable multimodal clinical decision support framework that jointly analyzes dermatological images, laboratory biomarkers, and clinical text narratives to model disease severity and support drug-class recommendation. The framework employs a task-conditioned gated fusion mechanism to adaptively combine modality-specific information in alignment with established treatment guidelines. The proposed approach was evaluated on 3000 patient-level CSU records using patient-level data partitioning. An image-only EfficientNet-B3 baseline achieved 68% accuracy in drug-class prediction, whereas the multimodal gated fusion model achieved 82.3% accuracy with a macro-averaged F1-score of 0.81. These results demonstrate that integrating multimodal clinical evidence substantially improves therapeutic modeling reliability and supports clinically aligned decision-making in CSU management.

Paper ID 371

ROMANCE SCAM REPORTING AND SUPPORT: HELP-SEEKING TIMING, TRUST, AND ESCALATION

LD Herrera

Abstract- Built on social engineering and identity deception, romance scams often create financial loss and distress. For victims, it can be difficult to know where to go, what information is needed, and what outcomes are realistic. This paper reports results from an anonymous survey of people who were targeted by or experienced a romance scam (completed surveys: n=386), focusing on (1) when and whether victims first reach out for help, (2) perceived difficulty and confidence in navigating support, (3) how trust relates to expectations of assistance, and (4) how loss severity relates to transfer-method complexity. When help was sought, it was often after money was sent rather than earlier in the scam. Many respondents also reported low confidence in finding support and navigational difficulty. Trust in formal organizations is positively associated with expectations of assistance, while perceived reporting difficulty is negatively associated with expectations. Loss severity increases with transfer-method complexity, highlighting the value of time-sensitive harm-reduction guidance. Based on these findings, this paper outlines design requirements and proposes a coordinated model centered on a structured incident record, a minimal evidence bundle, and "report once, route many" workflows that connect victims to platforms, financial institutions, and reporting channels with minimal additional burden.

Paper ID 372

GOVERNANCE AND AUDITABILITY OF AI-DRIVEN RETAIL DECISION PIPELINES IN CLOUD-NATIVE ARCHITECTURES

Prithvi raj Veluchamy

Abstract- This research examines the necessary critical frameworks needed for Artificial Intelligence decision pipeline governance and auditability in cloud-native retail environments. As retail organisations are moving towards automated inventory management, dynamic pricing and personalised customer experiences, the complexity of microservices-based architectures brings complex challenges to transparency. This research makes use of a special data set which consist of four hundred and ninety one instances of retail transactional and operational logs in order to evaluate a proposed auditing framework. The research uses the containerization tools such as Kubernetes for orchestration and special monitoring stacks for tracking the decision lineage. By combining automated logging and policy-as-code engines, the research shows how retail companies can ensure regulatory compliance on a large scale with AI operations. The results indicate that real-time auditability is a major step in reducing the risks of "black-box" decision-making so that automated changes in pricing or stock allocation are traceable and justifiable. Numerical performance metrics are

applied to validate the results, which indicate that accurate lossy fluid-governed pipelines are able to satisfy the low-latency requirements of contemporary cloud native retail infrastructures.

Paper ID 373

INFORMATION SECURITY WEB SYSTEM BASED ON THE ISO 27001 STANDARD: A CASE STUDY IN A CONSTRUCTION COMPANY IN LIMA

Eduardo Giordano Torres Rossi; Rosmery Milagros Pardave Huerta; Ernesto Adolfo Carrera Salas

This project presents the design of a security protocol based on the ISO/IEC 27001 standard, complemented by the development of the SecuPro web system, aimed at automating the management of policies, assets, incidents, and users while promoting regulatory compliance and staff training. The solution, implemented with Node.js and PostgreSQL under modular architecture, was validated by a construction company in Lima through controlled testing and user surveys. The system achieved a 41.67% reduction in operational errors, 92% policy adherence, and an average incident detection time of 57.5 seconds, demonstrating its effectiveness in strengthening ISO 27001-based security management in SMEs.

Paper ID 374

CONTEXTUAL REINFORCEMENT LEARNING FOR LINGUISTIC INTENT-GATED ACCESS CONTROL IN PRODUCTION AI SYSTEMS

Karthik Pappu (DSU)

Abstract- Static policy engines (RBAC/ABAC) cannot adapt to linguistic variability in natural-language access requests. We propose cRL-LIM, a framework integrating contextual reinforcement learning with LLM-based intent parsing: an LLM decomposes requests into structured features processed by a contextual bandit (LinUCB / Thompson Sampling) that fuses semantic intent with runtime context. Across six scenarios totaling 276,000+ episodes, LinUCB achieves 99.5% violation reduction (0.10% vs.20.30% baseline) and maintains this safety guarantee under load stress, adversarial injection, memory poisoning, and policy drift, while static baselines degrade up to 39%. The learned policy generalizes to unseen oracle policies and held-out request types, confirming that it captures security-relevant features beyond rote memorization.

Paper ID 375

MACHINE LEARNING-DRIVEN THREAT DETECTION AND RESPONSE FRAMEWORK FOR MODERN CYBERSECURITY SYSTEMS

Sathesh P.Sivashanmugam

Abstract- With the rapid growth of cyber threats in modern network environments, intelligent and automated intrusion detection systems have become essential for ensuring cybersecurity resilience. The paper suggests intelligent threat detection and response mitigation framework that is based on security data and is assessed by the CICIDS-2017 dataset. The methodology consists of systematic data preprocessing, outlier and duplicate removal, and correlation-based feature selection, which reduces the 78 original features to 45 important network-flow attributes. RobustScaler is used to normalize the features, and Random Forest (RF), eXtreme Gradient Boosting (XGBoost), and a hybrid stacking ensemble model of RF-XGBoost with Logistic Regression as the meta-learner are used for multi-class classification. Compared with the single RF and XGBoost models, the hybrid model outperforms them experimentally in accuracy (99%), precision (99%), recall (99%), and F1-score (99%). To increase transparency and confidence, Explainable Artificial Intelligence (XAI) with LIME is utilized to explain model decisions by finding network features that have a significant impact. The suggested framework integrates high detection accuracy, the ability to respond automatically, and interpretability, which makes it appropriate to scalable and intelligent cybersecurity applications.

Paper ID 376

RANSOMWARE ATTACKS ON AI SYSTEMS: A CROSS-DOMAIN THREAT AND CONTROL ANALYSIS

Yuvaraj Govindarajulu; Behnaz Karimi

Abstract- Ransomware has emerged as one of the most pervasive and economically damaging threats in modern information systems, while artificial intelligence (AI) has become a foundational component of critical enterprise and societal infrastructures. As organizations increasingly depend on predictive, generative, and agentic AI systems, these technologies are becoming high-value targets for extortion-driven cyberattacks; however, existing ransomware research and defense frameworks largely focus on traditional IT environments and do not adequately address AI-specific assets, workflows, and dependencies.

In this paper, we formalize the concept of AI ransomware and analyze attacks targeting AI artifacts, pipelines, and autonomous systems. We propose a cross-domain attack model and an AI ransomware kill chain, illustrated through predictive AI case studies. We examine threats across predictive, generative, and agentic paradigms and introduce a defense framework spanning development-time, deployment-time, and runtime controls with AI-specific incident response.

Paper ID 377

PREDICTIVE MODELS FOR URBAN AIR QUALITY MANAGEMENT USING AI

Dulara Liyanage; Nimasha Vithanage; Imasha Wijewardane; Nimasha Fernando; Dinuka Wijendra; Thamali Dassanayake

Abstract- Air pollution threatens public health in data scarce urban areas like Sri Lanka, where sparse monitoring hinders proactive management. We propose an integrated AI framework: hybrid SARIMAX-Temporal Fusion Transformer for multi-pollutant forecasting, ensemble spatial estimation for gap-filling, CEEMDAN-Seq2Seq for 24-hour AQI risk alerting, GRU for anomaly detection, and XAI for transparency. Validated on Central Environmental Authority data (2019-2024), the model achieves an 81.6% decrease in the value of the RMSE metric for ozone forecasting, as well as an R^2 value of 0.9077 for high-risk AQI prediction, outperforming the baseline methods by 15-81%. The framework is modular in nature, thereby providing policymakers with the ability to use real-time dashboards, thus making Sri Lanka move from reactive to proactive management.

Paper ID 379

ON OPTIMAL POWER ALLOCATION IN DOWNLINK MULTICARRIER NOMA SYSTEMS

Muhammad Nomani Kabir; Yasser Alginahi; Wael Said; Golam Saleh Ahmed Salem

Abstract- This paper investigates optimal power allocation in downlink multicarrier systems employing Non-Orthogonal Multiple Access (NOMA) for energy-efficient communication. We analyze optimization models governing power distribution among multiplexed users and subcarriers and refine key theoretical propositions for both intra- and inter-cluster allocation. The proposed refinements incorporate practical QoS requirements and user fairness considerations while maintaining analytical tractability. Results demonstrate that the enhanced allocation framework improves energy efficiency without compromising system throughput. The models from the propositions' underscores can be used for optimal power control in energy-aware, high-capacity 5G communication networks.

Paper ID 383

MORPHONET-MUGIL: A DEEP LEARNING–BASED APPLICATION FOR AUTOMATED MORPHOMETRIC MEASUREMENT OF MUGIL CEPHALUS

Charlie Marzan

Abstract- Traditional methods for measuring Mugil cepphalus fish in the fisheries industry rely heavily on manual techniques that are labor-intensive, time-consuming, and prone to human error. These limitations can slow research activities and affect the efficiency of fishery resource management. To address this problem, an automated solution

that streamlines morphometric measurement while improving accuracy and consistency was developed. The system is implemented as a user-friendly, responsive web application that enables users to easily capture fish images with their smartphones. Within the application, several image preprocessing techniques are applied to improve image quality and diversity before analysis. To determine the most effective approach for keypoint detection, five deep learning architectures: ResNet50, MobileNetV3, ConvNeXt, NASNetMobile, and EfficientNetV2. Results show that ConvNeXt achieved the best performance, obtaining an Object Keypoint Similarity (OKS) score of 0.9758 in detecting 24 anatomical landmarks of Mugil cephalus fish. The developed MorphoNet-Mugil application demonstrated strong real-world performance, with an average inference time of 5.24 seconds and a mean absolute error of 4.91. These findings suggest that the proposed system provides an efficient and reliable automated alternative to traditional morphometric measurement methods. By integrating deep learning and image preprocessing within a desktop and mobile platform, this work contributes a practical tool that can support researchers and fish farmers in improving data collection and decision-making in aquaculture.

Paper ID 384

ENHANCING PHISHING URL DETECTION WITH MACHINE LEARNING ALGORITHMS

Weijie Pang; Ashar Neyaz; Raymond Eichner

Abstract- Phishing URL detection remains a critical cybersecurity challenge. This study evaluates Logistic Regression and Random Forest classifiers using TF–IDF features extracted from a large dataset of 450,176 labeled URLs. Both models achieved high accuracy, with Logistic Regression reaching 99.787% and Random Forest 99.867%. Confusion matrix analysis indicates that Random Forest further reduces false negatives while maintaining a low false positive rate. Overall, Random Forest demonstrates superior performance, highlighting the advantage of ensemble learning for capturing complex lexical patterns in phishing URLs.

Paper ID 385

ADVERSARIAL ROBUSTNESS OF ML-BASED INTRUSION DETECTION SYSTEMS

Omar Farshad Jeelani; Viktoriia Korzhuk; Dmitry Sivkov; Alisa Vorobeva; Roman Safiullin

Abstract- Machine-learning intrusion detection systems (IDS) are increasingly deployed, yet their robustness to adversarial manipulation of flow features is less studied for gradient-boosting models. This paper evaluates a CatBoost-based IDS trained on the CSE-CIC-IDS2018 dataset under four white-box evasion attacks (FGSM, BIM, PGD, and Carlini–Wagner). We report accuracy and precision/recall/F1 to quantify degradation and discuss practical constraints on feature-space perturbations. On CSE-CIC-IDS2018, accuracy drops from 94.86% (clean) to 63.53% (FGSM), 63.51% (BIM), 59.67% (PGD), and 3.53% (CW), revealing severe robustness gaps.

Paper ID 386

TEMPORAL BEHAVIORAL ARCHETYPES OF RANSOMWARE IN ACTIVE DIRECTORY ENVIRONMENTS

Prajna Bhandary; Charles Nicholas

Abstract- Ransomware detection in enterprise environments often relies on static signatures or family-specific indicators that fail to generalize across rapidly evolving variants. This paper investigates behavior-based modeling of early-stage ransomware activity within Active Directory (AD) environments using empirical data from 156 controlled execution runs spanning multiple ransomware families. Windows Event Logs from the final 20 minutes of execution are temporally aggregated and represented as multivariate event sequences. Dynamic Time Warping (DTW) is used to align execution dynamics, and density-based clustering (HDBSCAN) identifies recurring behavioral structures without family supervision.

Results show that ransomware executions do not form strictly separable family-specific clusters in AD telemetry. Instead, activity converges into a small number of recurring temporal archetypes characterized by multi-stage patterns involving authentication events, system modifications, process creation, and application-level activity. Statistical analysis indicates only moderate association between family label and dominant archetype, highlighting substantial behavioral overlap.

These findings suggest that enterprise ransomware detection may benefit more from modeling shared compromise archetypes rather than relying solely on family-specific signatures, highlighting behavioral archetype detection as a shift from family-bound classification toward convergence-based detection in Active Directory environments.

Paper ID 387

EXPLORATION OF BOTTLENECKS AND OPTIMIZATIONS IN PRIVACY-PRESERVING MACHINE LEARNING INFERENCE

Bingyu Liu; Salem Othman; Leonidas Deligiannidis; Robbie McCue

Abstract- Deep neural network (DNN) services are widely deployed to provide accurate and efficient inference across many domains. In two-party communication scenarios, clients provide private inputs (e.g., images or videos) to cloud servers that execute inference using pre-trained models, leading to substantial risks of data leakage and model exposure. Secure inference techniques based on cryptographic primitives enable privacy-preserving two-party inference without revealing inputs, model parameters, or intermediate activations. However, existing privacy-preserving inference systems incur significant computation and communication overheads, resulting in diverse performance bottlenecks. In this work, we present a systematic analysis of secure inference costs. We formalize the dominant computation and communication cost drivers of homomorphic encryption and secure multi-party computation. We survey state-of-the-art privacy-preserving inference systems and examine how specific techniques and design choices mitigate computational and communication overheads, including SIMD packing, data layout optimization, protocol switching, and cost-model–driven compilation. Our analysis reveals fundamental bottlenecks in both linear and non-linear secure computations, identifies key challenges in achieving efficient and scalable secure two-party inference under realistic threat models, and motivates the need for hybrid system designs. Overall, this work clarifies the security landscape of privacy-preserving inference and provides valuable insights for building practical, scalable, and high-performance secure inference systems.

Paper ID 389

CROSS-LANGUAGE TRANSFER LEARNING FOR VULNERABILITY DETECTION: DIRECTIONAL ASYMMETRY BETWEEN JAVA AND PYTHON

Adiba Mahmud; Yasmeen Rawajfih; Fan Wu

Abstract- Whether vulnerability detection models transfer across programming languages, and whether such transfer is symmetric, remains an open question with direct consequences for securing polyglot codebases. We address this using an ensem- ble of frozen CodeBERT, GraphCodeBERT, and CodeT5 rep- resentations fused through cross-attention, evaluated on 11,060 Java and Python samples spanning 241 CWE types. Across five controlled configurations, we observe a pronounced directional asymmetry: Java-trained models achieve 75.15% F1 on Python (97.99% of the Python baseline), while Python-trained models collapse to 1.76% F1 on Java, an approximately 43× gap that has not been previously reported. Mixed-language training preserves within-language performance for both languages with no accuracy penalty, offering a practical path for polyglot deployment.

Paper ID 390

ROBUST WATERMARKING FOR SATELLITE IMAGERY: LEVERAGING SEMANTICALLY INSIGNIFICANT AREAS

Ngok Hung Pham; Aleksey Maksimov; Victor Fedoseev

Abstract- In this paper, we propose an approach for copyright protection of satellite imagery by embedding digital watermarks in a way that preserves significant spatial information. The key idea is to identify semantically insignificant regions within the image and embed the watermark into those areas. In this study, water-covered regions of the Earth's surface were selected as the target for embedding, as they are considered carrying minimal semantic value in land use tasks. To detect these regions, several deep neural network based models were employed. As baseline watermarking methods, the following algorithms were tested: DWT-DCT, MAX DCT, DWT-DCT-SVD, JAWS, and WAM. Experimental evaluation conducted under the WAVES protocol demonstrated the viability of the proposed approach:

watermarks embedded in semantically insignificant areas were successfully extracted from unaltered images. Under various distortions, the robustness of the watermark decreased compared to full-frame embedding; however, this drawback was offset by the preservation of critical content and the generally lower distortion level introduced by the watermarking process.

Paper ID 391

A Systematic Study of Security and Privacy in Large Language Models

Bingyu Liu

Abstract- Large language models (LLMs) are sophisticated and powerful intelligent systems with versatile functionalities that can be easily modulated through natural language prompts, such as machine translation, question replying, and text summarization. They learn information by analyzing huge amount of textual data, enabling them to perform a wide range of language-related tasks. However, the immense capabilities of LLMs also give rise to significant privacy and security concerns. When processing and generating large volumes of data, these models risk accidentally exposing sensitive information, thereby posing threats to data privacy. This paper aims to discuss the privacy issues associated with LLMs to foster a comprehensive understanding of these challenges. Specifically, it presents a thorough research that systematically examines various forms of data exposure and potential malicious exploitation that may occur within these systems. Finally, the discussion examines the challenges encountered and highlights potential future directions for enhancing privacy protection in LLMs.

Paper ID 392

Trustworthy and Reliable Machine Learning for Healthcare

Bingyu Liu

Abstract- Despite rapid advances in machine learning, widespread clinical adoption in radiology remains constrained by fundamental challenges, including stringent data privacy requirements, the high cost and variability of expert image annotation, and limited robustness of models across heterogeneous scanners, imaging protocols, and patient populations. In addition, severe class imbalance and algorithmic bias inherent to many radiologic applications can undermine diagnostic precision and disproportionately affect underrepresented populations, raising concerns about reliability and equity in clinical deployment. Emerging approaches such as privacy-preserving federated learning, self-supervised and weakly supervised learning, and domain adaptation offer promising strategies to improve scalability and generalizability under real-world conditions. Equally critical are rigorous assessments of algorithmic fairness and the integration of radiology-aware explainable artificial intelligence to enhance transparency, mitigate bias, and foster clinician trust. In this work, we systematically analyze the interconnected challenges of imbalanced datasets, bias, data privacy, and diagnostic precision, and present a unified framework that outlines practical design principles and evaluation strategies for developing robust, fair, and privacy-preserving machine learning systems for clinical radiology.

Paper ID 393

Detecting Fileless Malware through Memory Forensics with Recurrent Neural Networks

Noah Preston; Ajay Kumara Makanahalli Annaiah

Abstract- Fileless malware is an increasingly stealthy cybersecurity threat. threat that executes entirely in volatile memory, leveraging legitimate system utilities to evade traditional signature-based and static analysis defenses. Conventional machine learning approaches typically rely on static features or aggregated behavioral indicators, which fail to capture the temporal and execution order dependencies inherent in fileless attacks.

This paper presents a detection framework integrating memory forensics with a recurrent neural network (RNN)–based temporal modeling. Sequences of forensic and behavioral features extracted from volatile memory capture transient

process activity, memory-resident module behavior, and short-lived network interactions characteristic of fileless attacks.

A comparative evaluation of multiple RNN architectures shows that a Bidirectional LSTM achieves the best performance, reaching an accuracy of 0.93 with a recall of 1.00. Feature analysis identifies process behavior, memory-resident modules, and ephemeral network connections as strong indicators of fileless activity. These results demonstrate that behavioral memory features derived from volatile memory provide a more reliable basis for detecting evasive in-memory threats than static file-based approaches.

Paper ID 395

FROM BINARY VULNERABILITY DETECTION TO CWE CLASSIFICATION: A HIERARCHICAL PROMPTING STUDY

Obinna Okeke; Sushant Nepal; Venkat Sai Suman Lamba Karanam; Yan Wu

Abstract- The problem of fine-grained CWE (Common Weakness Enumeration)-specific vulnerability classification has received limited attention, despite its importance for understanding vulnerability semantics and supporting downstream security analysis. In this paper, we investigate how LLM performance scales from binary detection to multi-class CWE classification using the SecVulEval dataset across 145 CWE categories. Inspired by In-Context Learning (ICL), we propose a multilayer prompting framework, where each layer progressively enriches contextual information to guide subsequent layers toward accurate CWE classification.

Our paper reports the following findings using GPT-4o as the chosen LLM. Across balanced and imbalanced settings, we observe a consistent gap between binary vulnerability detection (vulnerable vs. safe) and exact CWE classification. For example, on a naturally imbalanced setting, the best binary F1 is 0.367 (using plain prompting), while the best exact CWE accuracy is 0.182 (using hierarchical-aware prompting). We further introduce family (hierarchical) accuracy as an evaluation metric, demonstrating that while LLMs can reliably identify the CWE family, they often fail to distinguish between exact CWE subtypes within the same family (i.e., between CWE siblings). This result indicates that many exact CWE classification errors are due to within-family confusion, i.e., "near misses" rather than cross-family misclassifications. These findings demonstrate that hierarchical prompting and family-based evaluation gives a much clearer picture of what LLMs are actually good at and where they get confused, while dealing with complex list of security flaws.

Paper ID 396

BLOODTRACE: WHERE DROPS BECOME DECISION

Dharanidharan Ranganathan; Devkumar Barot; Qian Liu; Oluwasola Mary Adedayo

Abstract- Bloodstain Pattern Analysis (BPA) plays a critical role in reconstructing violent events but remains heavily reliant on subjective expert interpretation. While machine learning approaches have shown promise, many rely on fragile grayscale thresholding or opaque deep learning models that lack forensic interpretability. This paper presents BloodTrace, an interpretable machine learning framework that combines dual-path color-based segmentation, micro-droplet recovery logic, and calibrated CatBoost classification to distinguish gunshot from blunt-impact bloodstain patterns. Using an improved dataset of 905 images with an 80/20 stratified split, the system achieves 95.6% accuracy on held-out test data and maintains approximately 90% accuracy on challenging external images where baseline grayscale methods degrade to 50%. Feature importance analysis confirms alignment with physically meaningful stain characteristics, demonstrating that interpretable, feature-based approaches can provide transparent and reliable decision support for BPA.

Paper ID 397

TRANSFER LEARNING FOR MALWARE DETECTION USING RGB BINARY VISUALIZATION: A COMPARATIVE STUDY

Eva Tuba; Ivona Brajevic; Adis Alihodzic; Ana Trisovic; Milan Tuba

Abstract- Malware detection using deep learning faces challenges in model selection for practical deployment. We systematically compare five transfer learning architectures (VGG16, ResNet50, DenseNet121, MobileNetV2, EfficientNetB0) on the MaleBin RGB malware dataset (12,000+ images through March 2025). Experiments on NVIDIA A100 GPU evaluated accuracy, efficiency, and deployment suitability. DenseNet121 achieved highest accuracy (91.20%, 8M parameters), MobileNetV2 provided optimal edge deployment (90.39%, 3.5M parameters), while ResNet50 and EfficientNetB0 unexpectedly underperformed (77.34%, 71.16%). Directions for practitioners are to deploy DenseNet121 for cloud environments, prioritizing accuracy, and MobileNetV2 for resource-constrained edge devices.

Paper ID 398

ON IMPACT OF ENSEMBLE STRATEGIES IN TABULAR DATA REGRESSION

Ivona Brajevic; Una Tuba; Milan Tuba

Abstract- Ensemble learning methods have demonstrated remarkable success in housing price prediction due to their ability to combine multiple models and achieve superior predictive performance. This study evaluates tree-based ensemble models, Random Forest, Extra Trees, and AdaBoost, alongside a Voting ensemble strategy for housing price prediction on the Boston Housing dataset. The experiments employ a rigorous repeated hold-out evaluation protocol combined with cross-validation on the training set to ensure stable performance estimates across multiple random splits. Additionally, Random Search-based hyperparameter optimization is applied to assess how systematic tuning influences model performance and robustness. Our results demonstrate that Extra Trees consistently achieves strong performance with excellent stability. The Voting ensemble, combining two Extra Trees models and one AdaBoost model, achieves superior overall performance with an MSE of 8.113, representing a 29.8% improvement over the best individual model. Hyperparameter optimization substantially improves all models, with the most pronounced gains observed for AdaBoost, where optimized configurations reduce MSE by approximately 24%. Comparative analysis with gradient boosting methods reported in recent literature reveals that the Voting ensemble achieves substantially superior performance, with approximately 68% lower MSE than gradient boosting approaches. These findings suggest that carefully constructed Voting ensembles with optimized hyperparameters represent a highly effective and reliable approach for housing price prediction problems.

Paper ID 399

DARKMINE: DEEP LEARNING FOR DARK WEB THREAT INTELLIGENCE – ANONYMOUS ATTACK INFRASTRUCTURE ATTRIBUTION AND CRIMINAL ORGANIZATION DETECTION

Usha Jammula; Naga Sujitha Vummaneni

Abstract- The dark web has become the primary infrastructure for adversarial threat actors, enabling anonymous coordination, tool distribution, and attack planning. Current monitoring approaches are largely manual and fail to extract actionable intelligence about threat actor organizations and attack campaigns. This paper introduces DARKMINE, a framework for automated dark web threat intelligence extraction combining domain-adapted natural language processing, heterogeneous graph neural networks, and temporal point process modeling. DARKMINE addresses four key challenges: identifying threat actor personas and organizational structures, attributing attack infrastructure to criminal organizations, predicting attack campaign timing and targets, and constructing threat actor collaboration networks. We validate DARKMINE on 2.3M forum posts, 847K marketplace listings, and 156K IRC logs from 15 underground communities over 18 months. Results demonstrate 85.2% precision in identifying 312 criminal organizations, 78.4% infrastructure attribution accuracy, and 72.1% attack prediction precision at 14-day horizons. Evaluation against 94 confirmed threat actors shows 89.3% identification accuracy with 0.87 overall F1 score. Organizations using DARKMINE predictions achieve 34% reduction in successful attacks compared to reactive baselines.

Paper ID 400

AI-Driven Adaptive Training Architecture for Post-Disaster Structural Damage Assessment

Hüseyin Deniz; Cem Baydoğan; Bahar Demirel; Fatih Özkaynak

Abstract- Post-disaster structural damage assessment demands rapid, consistent, and regulation-aligned decision-making under uncertainty, yet traditional training approaches lack individualized progression tracking and performance- aware feedback. This study proposes an AI-driven adaptive training architecture that integrates modular domain-specific content with a competency-based learner modeling framework and a structured analytics pipeline. Learner interactions— accuracy, response time, error patterns, and repetitions—are transformed into dynamic mastery estimates across predefined competencies. An adaptive recommendation engine uses these states to generate personalized remediation and next-step guidance. The layered architecture ensures scalability, maintainability, and regulation-aware content updates. A formal mathematical representation supports transparent competency tracking and real-time updates. Preliminary pilot deployment demonstrates stable mastery progression, consistent weak-competency identification, and short-term performance improvements following adaptive recommendations, with operational stability under concurrent use. This work bridges post-disaster assessment methodologies with adaptive learning systems, contributing a scalable, performance-driven training infrastructure for safety-critical professional education. Future work includes controlled experimental validation and statistical evaluation of learning gains.

Paper ID 401

Lifecycle-Integrated Security for AI-Cloud Convergence in Cyber-Physical Infrastructure

S M Zia Ur Rashid; Deepa Gurung; Sonam Raj Gupta; Suman Rath

Abstract- The convergence of Artificial Intelligence (AI) inference pipelines with cloud infrastructure creates a dual attack surface where cloud security standards and AI governance frameworks intersect without unified enforcement mechanisms. AI governance, cloud security, and industrial control system standards intersect without unified enforcement, leaving hybrid deployments exposed to cross-layer attacks that threaten safety-critical operations. This paper makes three primary contributions: (i) we synthesize these frameworks into a lifecycle-staged threat taxonomy structured around explicit attacker capability tiers, (ii) we propose a Unified Reference Architecture spanning a Secure Data Factory, a hardened model supply chain, and a runtime governance layer, (iii) we present a case study through Grid-Guard, a hybrid Transmission System Operator scenario in which coordinated defenses drawn from NIST AI RMF, MITRE ATLAS, OWASP AI Exchange and GenAI, CSA MAESTRO, and NERC CIP defeat a multi-tier physical-financial manipulation campaign without human intervention. Controls are mapped against all five frameworks and current NERC CIP standards to demonstrate that a single cloud-native architecture can simultaneously satisfy AI governance, adversarial robustness, agentic safety, and industrial regulatory compliance obligations.

Paper ID 402

Beyond NVD: Regional Vulnerability Databases for Global Coverage

Subhasish Mazumdar; Shashank Koganti

Abstract- Organizations which rely exclusively on the National Vulnerability Database (NVD) from USA face measurable information gaps that compromise their threat assessment and risk mitigation. This paper evaluates whether or not integration of the NVD with three regional vulnerability databases --- China National Vulnerability Database (CNVD), China National Vulnerability Database of Information Security (CNNVD), and Japan Vulnerability Notes (JVN) --- would (a) result in more comprehensive coverage and (b) be feasible. We answer both in the affirmative through analysis of 427,440 vulnerability records spanning January 2015 through August 2025. We demonstrate necessity by identifying totally absent severity ratings in 21,230 Common Vulnerabilities and Exposures (CVE) records in the NVD; feasibility by addressing the challenges of language translation and semantic heterogeneity; and improvement in coverage by showing, for example, that (i) severity ratings missing from 3,976 CVEs in NVD were found to be either Critical or High in the regional databases; (ii) fifteen vulnerabilities, crucial in practice, are missed entirely by NVD but documented in regional databases; (iii) translations of descriptions likely yield threat information

unknown to NVD; and (iv) regional databases are quicker than the NVD in dynamic modifications. Thus, our findings demonstrate that integration is feasible and would make vulnerability management more robust.

Paper ID 407

TRADES-Based Defense Against Adversarial Attacks in Medical Image Classification

Kily Jasso; Audrey Tollett; Ira Tuba; Eva Tuba

Abstract- Deep learning models for medical image classification are increasingly deployed in clinical settings yet remain vulnerable to adversarial attacks that can cause life-threatening misdiagnoses. We evaluate three CNN architectures (LeNet5, MobileNetV1) and three training methods (standard, adversarial, TRADES) against multiple adversarial attacks (FGSM, PGD, C&W). Testing on chest X-ray pneumonia detection with 5,863 images, we demonstrate that TRADES achieves 77.4% robust accuracy under PGD attacks (epsilon=2/255) versus 38.8% for standard training, while maintaining 82.5% clean accuracy—only 2.6% degradation. Critically, simpler architectures demonstrate superior adversarial robustness, with LeNet5 achieving 9.5% better performance than MobileNetV1 (77.4% vs 70.7%). Most importantly for clinical safety, false negative rates representing missed disease diagnoses decrease from 39.2% to 3.6% with TRADES training, a 91% reduction in missed diagnoses. Our findings provide evidence-based guidelines for deploying secure deep learning systems in healthcare environments.

Paper ID 409

Selective Memory Sharing in Multi-Agent LLM Teams via AgentGym-RL

Bhavuk Jain; Gunjan Jain; Hardeo Thakur

Abstract- Multi-agent LLM systems must balance retaining sufficient information for long-horizon reasoning while avoiding context overload that degrades coordination efficiency. If agents share everything, they become inefficient; if they work in isolation, they fail to coordinate. Prior work, including systems such as AgentBoard, typically assumes either full transcript sharing or strictly private memory. In this work, we study memory organisation in LLM-based agent teams by comparing three regimes: private memory, fully shared memory, and selectively shared memory. We introduce a lightweight control framework built on a Planner–Executor–Critic architecture, in which a controller decides what information is written to a shared scratchpad. Using this setup, we analyse how different memory regimes influence task success, communication overhead, and coordination behaviour. Our selective-sharing framework improves estimated task success by 6–12\% while reducing token usage by up to 55\% compared to full transcript sharing, demonstrating that structured memory routing yields more stable coordination than simply increasing context length.

Paper ID 410

Artificial Intelligence-Driven Predictive Models for Identifying Risk Factors of Chronic Diseases

Shaiful Mahmud; Khaleel khan Mohammed; Vasu Jain; Sarthak Anandkumar Shah

Abstract- Diabetes mellitus is a chronic metabolic disease that is a significant global public health concern. Complications may potentially be avoided or postponed with early diabetes diagnosis and treatment. The development of ML and DL has created new possibilities in the analysis of clinical data, allowing to identify the pattern concealed in them and increase the accuracy of diagnosis. This research proposes a diabetes prediction model based on the PIMA Indian Diabetes Dataset through the application of the Machine Learning (ML) and Deep Learning (DL) methods. Random Forest (RF) and Long Short-term Memory (LSTM) are two high-performance models that were implemented to extract nonlinear and temporal relationships in the data. Experimental testing showed that RF had 97.54% accuracy and 96.32 F1-score, whereas LSTM had a steady 98.85% accuracy and 98.20 F1-score. The proposed RF and LSTM models showed a clear superiority when compared to more traditional models, like Naive Bayes, Decision Trees, AdaBoost, and SVM (72-79% accuracy), to the more advanced models, like ANN, CNN-LSTM, and DNN (89-93% accuracy). The proposed framework demonstrates that combining ensemble and sequential learning offers a scalable and accurate solution for early diabetes detection, representing a key contribution to clinical decision support.

Paper ID 415

ENTERPRISE-GRADE AI-DRIVEN TEXT ANALYTICS AND INSIGHT EXTRACTION USING TRANSFORMER-BASED NLP MODELS

Chandra singh

Abstract- The rapid increase in user-created reviews on e-commerce sites like Amazon poses major challenges in deriving actionable information to consumers and sellers. The absence of structured summaries and reliance on manual analysis of large-scale reviews lead to information overload and limited decision support. To overcome this shortcoming, this research suggests an enterprise-level AI-based text analytics model of automated sentiment classification with transformer-based and ensemble learning models. A representative subset of 50,000 Amazon reviews is used and undergoes systematic data cleaning, advanced text preprocessing. Machine learning models like Nu-SVM and a Voting Ensemble are used and evaluated, along with transformer architectures like RoBERTa and DistilBERT. In contrast to the previous researches that analyze the models separately, the proposed framework combines structured preprocessing, classic feature-based models, and contextual transformer models into a single enterprise-oriented evaluation pipeline. Based on the results of the experiment, RoBERTa is the most accurate (80.51%), followed by DistilBERT (80.32%) and Voting Ensemble (79.43) and Nu-SVM (79.19). These results show that, in a typical comparison setting, transformer-based models are able to increase performance relative to conventional baselines. The results prove that the transformer-based contextual representations are more effective than the traditional ones in the context of sentiment discrimination and offer a scalable and robust option to extract automated insights on the enterprise level in large-scale e-commerce settings.

Paper ID 422

PHARMA SUPPLY CHAIN CYBER RISK DASHBOARD: VISUALIZING THREATS AND APPLYING CNSS CONTROLS ACROSS SUPPLIERS, MANUFACTURERS, AND DISTRIBUTORS

Pravallika Gummadivelli; Gahangir Hossain

Abstract- Cyberattacks on pharmaceutical supply chains are increasingly aimed at intellectual property (IP), production integrity and distribution logistics. Cyber threats, such as ransomware and tampering of APIs, are a major operational risk, especially with the global dependence on digitalized manufacturing systems and real-time delivery systems. This paper proposes a dashboard to monitor cyber risks in pharmaceutical supply chains. It uses the Committee on National Security Systems' (CNSS) model of security to map threats and identify actionable controls for suppliers, manufacturers and distributors. The threat-modeling method was structured, combining MITRE ATT&CK maps, vulnerability scores, and CNSS triad control (confidentiality integrity availability). The dashboard prototype was created using Python, Neo4j and Power BI. Real-time CNSS aligned indicators are more effective in detecting supplier-origin compromised products, reduce alert fatigue and allow for better segmentation of risks. The approach highlights gaps in current monitoring--specifically integrity validation and cross-entity visibility--and demonstrates how a unified visualization platform can strengthen risk management for high-value, safety-critical pharma ecosystems.

Paper ID 423

SECURING CLINICAL TRIAL DATA VAULTS: A CIA-DRIVEN ARCHITECTURE FOR CONFIDENTIALITY, INTEGRITY, AND AVAILABILITY

Pardhiv Vasireddy; Gahangir Hossain

Abstract- Clinical trials generate highly sensitive datasets that include protected health information (PHI), protocol data, and derived analytical results. Security breaches or fraudulent manipulation of these data can harm participants, compromise scientific validity, and undermine regulatory trust. This article proposes the design specifications for a Clinical Trial Data Vault (CTDV) security architecture that satisfies confidentiality, integrity, and availability (CIA) requirements while supporting operational needs of trial sponsors. We review relevant regulatory frameworks and technical standards governing clinical trials, applicable health data privacy requirements, and emerging approaches for secure clinical data platforms, including blockchain-enabled transparency. We then identify major threat vectors involving sponsors, contract research organizations (CROs), and external researchers. The proposed CTDV architecture

integrates data de-identification, strong encryption, role- and attribute-based access control, tamper-evident logging, and resilient storage systems. A structured security-control mapping methodology is introduced to evaluate risks across CIA elements and key regulatory obligations. We provide a qualitative assessment of risk reduction achieved through the proposed architecture. Results indicate that the CTDV design enables governed and timely access for authorized researchers while mitigating common attack vectors such as credential theft, database tampering, and insider misuse. The article concludes with practical implementation recommendations and directions for future research to strengthen secure clinical trial data management.

Paper ID 430

AI FOR DEVSECOPS OPTIMIZATION: INVESTIGATING THE ROLE OF AI/ML IN PREDICTIVE VULNERABILITY DETECTION DURING CI/CD PIPELINE STAGES

Tarun Kalwani

Abstract- In this paper, the focus is placed on the integration of Artificial Intelligence and Machine Learning into the CI/CD pipelines of DevSecOps for predictive vulnerability detection. A specially selected dataset of 416 past code commits and vulnerability reports was used for the model building of the Random Forest and Recurrent Neural Network models. The designed AI-based system was seen to reach a detection accuracy of up to 93.8%, greatly outperforming the results yielded by conventional static code analysis solutions yet achieving a latency decrease of over 83%, as measured across repeated CI/CD build cycles in the associated security scan. With predictive models directly integrated into the CI/CD pipeline, the system allows for real-time and context-aware security reviews and the decrease of false positives, ultimately boosting the productivity of developers and the deployment time.

Paper ID 431

LLM-AUGMENTED AGENTIC CONSENSUS SWARMS FOR AUTONOMOUS EDGE SECURITY

Raju Dandigam; Ravi Teja Thutari; Tejaskumar Vaidya

Abstract- Multi-Agent Edge Security: Threat Detection and Mitigation in Resource Constrained Networks Edge comput- ing has introduced several security challenges such as dis- tributed timely threat detection and mitigation under resource- constrained environment. The traditional centralized security solutions are characterized with long detection time, single point of failure and weak to various types of threats. The multi- agent-based intrusion detection and lightweight large language model (LLM) based semantic reasoning have been proposed and explored in the literature. However, there is no research work has been done using decentralized consensus, LLM-based semantic reasoning and Byzantine fault tolerance in a unified resource- constrained edge security framework. In this paper, we pro-posed a LLM-Augmented Agentic Consensus Swarm (LAACS) framework, which enables each edge gateway to be equipped with a lightweight LLM for conducting semantic reasoning with other edge gateways by using a hybrid semantic-numerical consensus protocol. The consensus protocol allows distributed agents to perform collaborative threat detection and validation with the help of numerical confidence scores and learn each other's semantic threat representations based on their knowledge bases and statistics from the anomaly detection results. We have conducted experiments using three real-world edge datasets. The results show that the proposed system achieved an average detection rate of 94.7% at an average consensus latency of 23ms while maintaining a swarm of 50 agents, which outperforms the state-of-the-art distributed intrusion detection system (IDS) by 12.3% in terms of detection rate and 3.2 times in terms of latency. The proposed system achieved a linear scalability of 200 agents and maintained an average detection rate of 89.1% even under a 30% Byzantine attack. The main contributions of this paper are as follows: A LLM-augmented decentralized edge security framework, a hybrid semantic-numerical Byzantine-resilient consensus protocol, the effectiveness of the proposed LAACS system and the evaluation using real-world edge security datasets.

Paper ID 432

PRODUCTIZING AI-DRIVEN NETWORK SECURITY SYSTEMS: ARCHITECTURE, TRADE-OFFS, AND PRODUCT MANAGEMENT PERSPECTIVES

Balu Chavan

Abstract—The high pace of cyber threats development requires a shift in the reactive security strategies towards the proactive ones that are AI-focused. The study discusses productization of artificial intelligence in telecommunication and network security industries with special attention to the strategic challenges associated with product managers. This research pinpoints some of the success factors in commercialization by combining advanced machine learning with security solutions that are ready to be applied in the market. The study makes use of an edited dataset of 477 data instances described as the network traffic patterns, latency indicators, and threat vectors. The analytical frameworks and sophisticated visualization tools Python based are some of the tools used in this research to model the performance of products in a variety of network settings. The results indicate that the process of productizing AI is not simply a technical one but a multi-dimensional task that covers the following aspects: data integrity, user experience design, and scalable infrastructure. The present paper presents a map to help product managers negotiate the challenges of the lifecycle management in a time where digital resilience is a requirement that cannot be achieved without automated defense.

Paper ID 434

AKURA; ADAPTIVE SINHALA LEARNING FOR EARLY CHILDHOOD USING GESTURE, VOICE, EMOTION AND HANDWRITING

N. W. P. G. T. Mihiran; R. M. D. S. Rathnayake; H. A. D. T. Piyathilaka; N. W. P. G. T. T. Rahul; Dinuka Wijendra; Jenny Krishara

Abstract- Early childhood Sinhala language education in Sri Lanka faces persistent challenges such as large class sizes, reliance on rote-based learning, lack of personalized engagement, and limited engaging content, resulting in poor early literacy outcomes. This research presents Akura, a novel multimodal mobile platform that integrates real-time handwriting recognition, gesture-based numeracy, emotionaware lesson adaptation, and speech pronunciation feedback to deliver an interactive and personalized Sinhala learning experience for children aged 3–6. The system is optimized to run entirely on low-end mobile devices and uses on-device intelligence to ensure data privacy and low latency. Each module of Akura's multimodal approach has been designed with child-specific optimizations to dynamically adjust feedback, lesson difficulty, and content in real time according to each child's performance and emotional state, while generating progress reports that summarize learning outcomes and interaction trends for caregivers and educators. The study yielded an 8.1% error rate for the speech pronunciation assessment model and 92.5% accuracy for the emotion classification model. Additionally, the gesture recognition model achieved 93.4% accuracy, and the real-time handwriting recognition module achieved 95% accuracy through active stroke-level validation. To validate real-world applicability, Akura underwent technical validation on low-end mobile devices and an educational impact assessment with 30 preschoolers. By fusing multimodal sensing, Akura targets improved engagement and measurable gains in early Sinhala literacy and numeracy. Overall, addressing the critical gap left by existing static Sinhala applications, this work introduces a novel end-to-end multimodal framework for Sinhala early childhood education that is deployable offline and suitable for resource-constrained settings.

Paper ID 443

HIGH-PERFORMANCE DISTRIBUTED DEEP LEARNING USING ADAPTIVE PARALLELISM AND DYNAMIC WORKLOAD SCHEDULING

Pavan Kumar Boyapati; Siva Teja Reddy Kandula

Abstract- Dynamic and large-scale workloads are essential to the management of modern distributed computing and learning systems, particularly in the event of resource failure or failure in execution. This work introduces a scalable and intelligent system of adaptive workload scheduling in high- performance distributed systems. Experiments with Google 2019 cluster workload dataset, a representation of actual resource consumption, task patterns and failures, utilize a preprocessing pipeline in data inspection, feature engineering, normalization, and class imbalance correction using SMOTEENN. MLP and CatBoost models are trained to predict failures on the task and scheduling readiness with

accuracies of 99.15% and 87.62%, respectively, which is very successful in comparison with MobileNetV2 (70.4%), Support Vector Regression (70%), and TSSAP clustering (53.8%). Findings have shown higher predictive accuracy, scalability and reliability, and therefore the framework is applicable to high-performance computing, cloud data centers and resource-intelligent distributed learning systems.

Paper ID 444

AUTOMATING ORGANIZATIONAL CYBER SECURITY POLICY COMPLIANCE AGAINST INDUSTRY STANDARDS USING AGENTIC AI

Rohit Negi; Soumyo Chakraborty; Amit Negi; Sandeep Shukla

Abstract- Auditing and compliance management are an integral part of a cybersecurity management systems (CSMS). However, the frequency of audits and compliance checks is typically once a year for external audits and twice a year for internal audits. Under audit and compliance management, cybersecurity policy documents defined according to normative references are also reviewed. This review is either performed at the semantic level, which may miss contextual reasoning, or manually, which requires significant cognitive effort and has a very high time complexity. To fill this gap, in this paper we focus on leveraging Artificial Intelligence (AI) in cybersecurity management. We propose the Multi-Agent Cybersecurity Policy Analysis & Validation Workbench (MACAW), an agentic AI framework that leverages Large Language Models (LLMs) and Retrieval Augmented Generation (RAG) with planning and tool-use capabilities to analyze policies, frameworks, standards, and guidelines. Agents understand contextual dependencies and generate adaptive responses. In this experiment, LLM with RAG is used to generate contextual prompts. This results in an automated multi-agent policy analysis tool for compliance audit and assessing the compliance of policies with relevant industry standards for security controls. We benchmark our approach using two sets of cybersecurity policies, one from the transportation sector and another from the financial and banking sector, and security controls from two widely accepted standards - NIST 800-53 as well as ISO 27002:2013 / ISO 27002:2022. Experimental results demonstrate that our agentic framework achieves superior performance in contextual accuracy and automation efficiency, with accuracy (as represented by the F1 score) in the range of 83\% to 99\% against a benchmark of consensus decisions from human cybersecurity audit experts. This has considerable implications for both productivity and security applications.

Paper ID 449

DATA PROTECTION AND NATIONAL SECURITY IMPLICATIONS OF POST-DISASTER STRUCTURAL DAMAGE ASSESSMENT SYSTEMS

Cem Baydoğan; Bahar Demirel; Tuba Demir; Sedat Savaş; Ferhat Uçar; Fatih özkaynak

Abstract- The large-scale digitalization of post-disaster damage assessment systems has enhanced coordination and response efficiency but has simultaneously expanded cybersecurity and privacy risks. Centralized disaster databases frequently store personally identifiable information, geospatial infrastructure mappings, and structural vulnerability indicators, creating an enlarged attack surface and increasing exposure to adversarial exploitation. In this study, disaster data governance is examined through an integrated analytical framework combining data protection law, national security considerations, and formal threat modeling. Systemic risk is formulated as a function of data sensitivity, exposure level, and adversary capability. The attack surface of centralized disaster platforms is mathematically characterized, and integrity and disinformation risks are incorporated into an extended strategic impact model. Based on this framework, a secure architectural design is proposed, integrating role-based and context-aware access control, tiered data visibility, cryptographic identity abstraction, differential privacy for public disclosure, and tamper-evident audit logging. The proposed controls are explicitly mapped to risk-reduction parameters within the formal model. It is argued that disaster data platforms should be conceptualized as components of national resilience infrastructure, requiring a calibrated balance between rapid emergency data sharing and proportional regulatory safeguards.

Paper ID 450

TAMPER DETECTION IN CT DICOM IMAGES FOR DIGITAL FORENSICS AND CLINICAL INTEGRITY

Hala Strohmier Berry; William Carroll; Jessica Brown; Sydney Halupa

Abstract- Ensuring the integrity of medical imaging is critical for accurate clinical decision-making and for maintaining evidentiary reliability in forensic investigations. The increasing connectivity of medical imaging systems and reliance on the DICOM standard have expanded the attack surface for unauthorized image manipulation. This study evaluates deep learning–based methods for detecting tampering in CT DICOM images within a digital forensic and clinical integrity context. Two complementary detection paradigms were examined: an artifact-based manipulation tracing approach and a reconstruction-based anomaly detection approach. Models were trained and evaluated using CT-GAN, Back-in-Time Diffusion (BTD), and Medical Image Tampering datasets, with performance assessed using accuracy, precision, recall, and F1-score metrics. Results indicate that both approaches demonstrate high sensitivity to tampered images; however, performance varies under cross-dataset conditions, highlighting the impact of training distribution and domain shift. The reconstruction-based approach showed improved discrimination when trained on authentic baseline data, while the artifact-based approach achieved strong performance across diverse tampering modalities. Findings suggest these methods can serve as effective screening tools for detecting manipulation in medical images, but threshold calibration and dataset diversity are essential to reduce false positives in clinical workflows. This work contributes operational insights for integrating tamper detection into digital forensic processes and securing medical imaging environments.

Paper ID 453

LEAN-DRIVEN SAP PRODUCTION PLANNING OPTIMIZATION USING MACHINE LEARNING FOR INVENTORY AND THROUGHPUT EFFICIENCY

Mahendrakumar Kalal

Abstract- Inventory management plays a pivotal role in enterprise sustainability according to global sustainability principles. Complex algorithms incorporated into SAP (Enterprise Resource Planning) systems increase efficiency, lower costs, and facilitate sustainable operations. The present paper suggests a SAP-inspired Lean-based model of production optimization to model throughput efficiency of continuous-flow production based on Multi-Stage Continuous-Flow Manufacturing Process dataset (14,088 samples, 116 features). The propose framework encompasses systematic preprocessing such as data cleaning and Yeo-Johnson power transformation, and refined feature engineering such as Lean stability metrics, SAP-oriented measures and indicators, statistical descriptions, and rolling and lag variables. SelectKBest with Mutual Information Regression feature selection reduces the data to 100 informative predictors. The tuned Random Forest is selected as the best model with the highest performance (R2 = 85.91, RMSE = 1.4033, MAE = 1.1055, and MAPE = 1.53), outperforming XGBoost model in predictive accuracy and stability. SHAP and LIME provide global and local interpretability, while scenario simulation and Monte Carlo analysis assess robustness. The results have shown that explainable machine learning combined with the Lean principles provides a scalable, interpretable, and sustainable SAP-integrated production planning system.

Paper ID 454

REAL-TIME VISION-BASED HUMAN–ROBOT INTERACTION FRAMEWORK FOR LOW-COST EMBEDDED ROBOTIC ARMS

Wewage Dep; Vishmi Embuldeniya

Abstract- Low-cost robotic manipulation remains inaccessible to many small and medium-sized enterprises (SMEs) due to hardware cost and programming complexity. Conventional robotic control approaches rely heavily on inverse kinematic modelling and geometric calibration, limiting their practicality in resource-constrained settings.

This paper presents a real-time vision-based human–robot interaction framework for controlling a 4-DOF robotic arm using a monocular RGB camera and embedded wireless actuation. Human joint angles are geometrically extracted using MediaPipe pose estimation, filtered via a discrete Kalman estimator, and directly mapped to servo actuation without inverse kinematics.

The perception–processing–actuation pipeline operates at approximately 10 FPS under bounded Wi-Fi delay. Experimental results demonstrate a mean latency of 57.3 ms and inter-frame angular variation below 4 degrees across all joints. The results confirm stable real-time imitation control using a lightweight, embedded-compatible architecture suitable for educational and SME deployment.

Paper ID 457

DASTestBed: An Automated Benchmarking Framework for DAST Scanners with Extensible Ground Truth Modeling

Rand Deeb; Alisa Vorobeva; Omar Jeelani

Abstract- Dynamic Application Security Testing (DAST) tools assess web applications at runtime, but their evaluation is often unreliable because test environments vary and ground truth is usually not machine readable. This makes fair comparison difficult and weakens evidence for tool selection. This paper presents DASTestBed, an automated containerized benchmarking framework for repeatable DAST experiments with explicit scope and auditable artifacts. Benchmark scope is modeled as structured vulnerability records in YAML, where each record enumerates concrete request instances used as atomic scoring units. The framework orchestrates scanners and targets in isolated containers, collects raw artifacts, parses heterogeneous outputs through an adapter parser interface, normalizes findings into a common schema, matches them to indexed request instances using location aware and variant aware logic, and deduplicates repeated alerts before scoring. In addition to effectiveness metrics, DASTestBed records operational telemetry, including CPU usage, memory usage, and HTTP request volume. We demonstrate the full pipeline on the Damn Vulnerable Web Application, DVWA, by integrating Nuclei and OWASP ZAP. Using one fixed profile per tool under a validated and indexed scope of 103 request instances, Nuclei produced 15 true positives and 0 false positives, with precision 1.00, recall 0.15, and F1 score 0.25. Under the same scope, OWASP ZAP produced 100 true positives and 42 false positives, with precision 0.70, recall 0.97, and F1 score 0.82. ZAP achieved higher coverage but required more resources, reaching 91% CPU usage, 3.2 GB memory usage, and 19,501 HTTP requests, compared with Nuclei at 0.92% CPU usage, 937.3 MB memory usage, and 11,305 HTTP requests. These results show that DASTestBed supports tool agnostic, scope aware benchmarking and enables practical comparison of detection outcomes together with operational cost.

Paper ID 459

Feature-Equivalence Deduplication and Memoization of HTTP(S) Requests for Web Scanners: Formal Model, Concurrency, and Complexity Bounds

Rand Deeb; Alisa Vorobeva; Omar Jeelani

Abstract- Web scanners and web-measurement crawlers often issue redundant HTTP and HTTPS requests due to overlapping discovery paths, repeated modules, retries, and concurrent workers. We present a request-level deduplication and memoization mechanism placed at the dispatch boundary. The method defines deterministic canonicalization, policy-driven feature-equivalence identity, and canonical serialization used to derive deduplication keys. Memoization uses bucketed storage with collision-bounded verification through two-stage feature checks, and replay is gated by a strict validity predicate combining store-local monotonic time-to-live, context binding, and safety guards for method and session-bearing requests. For concurrency, we provide at-most-one concurrent dispatch per key during a miss window using in-flight claims and a multi-process lease variant with storage-level effectively-once persistence via idempotent writes. We prove key stability, collision-aware replay verification (in exact and compact modes), and dispatch/storage properties under standard atomicity assumptions, and derive explicit CPU, memory, and persistent-store complexity bounds for lookup and save operations.

Paper ID 460

DATA-DRIVEN PREDICTION OF ADVERTISING DIGITAL CAMPAIGN EFFECTIVENESS USING ARTIFICIAL INTELLIGENCE

Abhinay Kumar Reddy Seella

Abstract- Background: Predicting the effectiveness of digital advertising campaigns is still a difficult task, and nonlinear relationships between engagement, cost, and the attributes of a campaign can be complex and difficult to predict. Objective: This paper create an artificial intelligence model that is data-driven to forecast the effectiveness of digital advertising campaigns, and to better estimate ROI and assist in decision making. Methodology: Boosting-based regression models, such as LightGBM, XGBoost, and CatBoost, are applied and tested with the help of the Marketing Campaign Performance Dataset (Kaggle) and standard regression metrics, as well as, k-fold cross-validation. In contrast to other previous studies that address one activity at a time, such as the click-through rate or revenues forecasting, this research hypothesizes to create a single, interpretable campaign ROI forecasting framework. Interpretation of model predictions is done using explainable AI (SHAP). Results: It has been experimented that the CatBoost Regressor has the highest performance with an R^2 of 0.9817 (98.17%), as well as the lowest error measures. Conclusion: The suggested framework presents an effective integration of machine learning and explainable AI to deliver precise, understandable, and usable insights to streamline the digital advertising campaigns.

Paper ID 463

A CONTROLLED COMPARATIVE STUDY OF MONGODB SECURITY UNDER REST AND GRAPHQL APIS WITH PROXY-BASED PROTECTION

Siddhi Jhade; Sapna VM; Shruti Kumari; Tejas Gowrish; Ullas Girish; Prasad Honnavalli

Abstract- The widespread adoption of NoSQL databases such as MongoDB and modern APIs such as REST and GraphQL is directly related to their ability to provide flexible and scalable solutions for the rapid adoption of data driven applications. However, vulnerabilities associated with the design of APIs and the manner in which inputs are handled may create security issues that jeopardise the integrity of the data and can adversely affect the performance of systems that rely on data-driven methodologies. Although they are widely used, very little empirical work has been done directly comparing the security implications of accessing MongoDB via REST and GraphQL under controlled and identical backend conditions.

This paper presents a controlled comparative study of MongoDB security when accessed with the use of REST and GraphQL APIs, with a lightweight reverse-proxy mitigation framework. REST and GraphQL servers that were intentionally vulnerable were deployed on a shared MongoDB backend, and injection and denial-of-service (DoS) attack workloads of semantic equivalence were built against both interfaces. The proposed proxy incorporates an ML-based injection detection model and a DoS control mechanism in the form of a time-out mechanism. The classifier achieved high detection accuracy with consistent cross-validation performance, while the proxy successfully blocked injection attempts and significantly reduced DoS impact with minimal latency overhead. The findings indicate security trade-offs between REST and GraphQL architectures and illustrate the efficiency of lightweight, layered defenses in enhancing MongoDB deployments.

Paper ID 464

PERFORMANCE AND COMPARATIVE ANALYSIS OF A MULTI-STAGE TRANSFER LEARNING FRAMEWORK FOR BRAIN DISEASE CLASSIFICATION USING CLAHE-ENHANCED 3D-RENDERED MRI IMAGES

Isaac Angelo M. Dioses; Jesusimo L. Dioses Jr.; Alexander Hernandez

Brain tumor detection using magnetic resonance imaging (MRI) plays a critical role in early diagnosis and treatment planning. However, manual analysis of MRI images can be time-consuming and prone to human error. This study proposes a deep learning framework for brain MRI classification that integrates Contrast Limited Adaptive Histogram Equalization (CLAHE) preprocessing with a Multi- Stage Transfer Learning (MSTL) strategy. The proposed framework evaluates three convolutional neural network architectures, MobileNetV2, ResNet50, and EfficientNet-B0, to analyze their performance in classifying 3D-rendered brain MRI images into tumor categories. CLAHE was applied to enhance image contrast and improve the visibility of structural patterns before training. The MSTL framework

progressively fine-tunes pretrained models through multiple stages, enabling better adaptation of learned features to the MRI dataset. Experimental results demonstrate that all three models achieved high classification performance. Among the evaluated architectures, ResNet50 achieved the highest accuracy of 99.06%, followed by EfficientNet-B0 at 98.90% and MobileNetV2 at 98.12%. Training curves and confusion matrix analysis further confirmed stable convergence and strong classification capability across the models. The novelty of this study lies in combining CLAHE-based MRI enhancement with a progressive transfer learning framework to improve deep learning performance in medical image classification. The proposed approach may support AI-assisted diagnostic systems for automated brain tumor detection and improve the efficiency of clinical decision-making processes.

Paper ID 467

A SWOT Analysis of Mobile Privacy and Security Education for K–12 Students

Tayiba Raheem; Mary Nusrat; Gahangir Hossain

Abstract- The presence of mobile devices has transformed the social landscape for K-12 students, yet this constant connectivity exposes young users to significant privacy and security risks. While adolescents are often termed “digital natives”. Research indicates they frequently lack the technical proficiency to navigate complex data ecosystems, such as managing app permissions and understanding third-party data collection. Furthermore, traditional safety approaches relying on "protectionist” surveillance and parental control applications often erode trust between adults and teens, potentially suppressing the development of necessary risk-coping strategies. This paper utilizes a SWOT (Strengths, Weaknesses, Opportunities, and Threats) analysis framework to evaluate the strategic viability of integrating mobile privacy education into K-12 curricula. The analysis reveals that while high student engagement with mobile technology serves as a primary Strength, the reliance on restrictive monitoring constitutes a significant Weakness. Threats include the escalating frequency of cyber incidents in schools, such as data breaches. However, Opportunities exist to implement “joint family oversight” models where parents and teens collaborate to manage online safety. The study concludes that effective K-12 cybersecurity education must shift from a compliance-based surveillance model to a resilience-based approach, empowering students with the agency to actively manage their own digital privacy.

Paper ID 472

An RFID-Assisted Machine Learning Approach for Robust Face Recognition

Aminul Islam; Atia Farzana Chowdury; Abdullah Al Mamun; Nur Hossain Bhuiyan

Abstract- Face recognition is one of the most widely deployed biometric technologies, extensively used in security, surveillance, and access-control systems. Conventional vision-based methods, however, often face significant performance degradation under challenging conditions such as low illumination, facial occlusion, pose variation, adverse weather, and privacy constraints, limiting their reliability in real-world environments. To address these limitations, this paper proposes an RFID-assisted machine learning framework for robust face recognition that reduces dependency on visual facial data and enhances authentication reliability in security-critical environments. The proposed system integrates Radio Frequency Identification (RFID) sensing with a Support Vector Machine (SVM) classifier to perform identity recognition under non-line-of-sight and low-visibility conditions. RFID signals are carefully preprocessed to enhance signal quality, and discriminative features are extracted for supervised classification. Experimental evaluations conducted on a dataset of 4,535 RFID samples using supervised SVM classification using cross-validation demonstrate that the proposed approach achieves 98% recognition accuracy, significantly improving robustness and reliability compared to traditional camera-based systems, particularly in visually challenged or obstructed environments. The framework provides a contactless, cost-effective, and scalable solution for secure authentication, highlighting the potential of combining RFID technology with machine learning to enhance face recognition performance. This work paves the way for practical deployment of non-image-dependent biometric systems in real-world security-critical scenarios.

Paper ID 476

AUTOMATED EXTRACTION OF BROWSER HISTORY FROM COMPUTER HARD DISKS

Sakar Joshi, Sankardas Roy

Abstract- In the context of a digital forensics investigation, a user's browser activity may reflect their behavior, and the browser history can be a valuable artifact for the investigator to reconstruct the timeline of the relevant events. However, the extraction of browser history from a seized hard disk is currently a slow, error-prone process, partly due to the fact that popular browsers (e.g., Chrome vs. Firefox) differ in architecture. We propose an automated approach for extracting browser history from a computer hard disk or its image. The implemented solution is a Python pipeline that extracts the web history from a disk or its image while ensuring data integrity. The pipeline consists of a custom image handler, filesystem analysis routines, and parsers for popular browsers. Our tool identifies user profiles, locates browser history databases, performs the analysis, and finally exports the results in multiple formats (CSV and JSON). For validation, images from popular CTF challenges as well as live computer hard disks were used, which demonstrates improvements over manual methods. The results confirm that automated browser history extraction is possible.

Paper ID 484

DEVSECOPS-DRIVEN SECURITY CONTROLS FOR ERP RELEASE PIPELINES

Yogeesh Kunigal Gangaiah; Karthik Pappu; Yogesh Thanvi

Abstract- Enterprise Resource Planning (ERP) security is commonly treated as a periodic assessment that runs outside the CI/CD flow, leaving configuration drift, privilege escalation, and business-rule violations undetected until a manual audit cycle. This paper presents ERP-ReleaseGuard, a DevSecOps pattern that embeds three ERP-domain controls directly into the release pipeline: (i) configuration-snapshot drift detection for role-based access control (RBAC) matrices, (ii) runtime guest-access probing, and (iii) a normalized risk score that produces a deterministic go/no-go gate. We evaluated the pattern on a Kubernetes testbed comprising four FastAPI microservice stubs that simulate an ERP accounts and an HR landscape with three fault classes at two severity levels. In a controlled experiment with 42 runs (7 scenarios ×3 seeds ×2 pipelines), the guarded pipeline detects 94.4% of injected faults compared to 0% for the baseline (Fisher's exact test, $p < 0.001$). Impact-based test selection reduces test effort by 58.3% (TCRR) without increasing leakage, and the risk-score gate achieves a precision of 1.00 with a recall of 0.94.

Paper ID 488

COMPARATIVE STUDY OF SYMBOLIC EXECUTION TOOLS APPLIED TO VULNERABILITY DETECTION

Andre Cardoso; Oscar Ribeiro; Nuno Lopes

Abstract—The massive growth in web and social media apps led to the exposure of multiple security vulnerabilities, caused by the lack of knowledge of developers, which led to the emergence of attacks like Denial of Service and Cross-Site Scripting. Testing for vulnerabilities is crucial in software development and requires automated approaches. Symbolic Execution is a Static Application Security Testing technique used to identify vulnerabilities in software. This paper compares two symbolic execution tools, KLEE and Owi, focusing on their performance and accuracy in vulnerability detection. After experimenting with both tools and analysing the results, KLEE demonstrated better performance and more accurate results in detecting vulnerabilities.

Paper ID 490

AI-DRIVEN MALWARE DEFENSE: TRANSFORMER MODEL FOR REAL-TIME DETECTION AND THREAT ANALYSIS

Viktoria Meskova; Natalia Valencia; Gustavo A. Chaparro-Baquero; Alexander Perez-Pons

Abstract- Real-time analysis is a critical component for identifying and understanding threats caused by malicious software. Over the years, various approaches have been developed to identify and mitigate harmful behaviour, moving beyond traditional signature-based techniques. In particular, neural networks have been widely used to enhance

detection capabilities. However, accurate malware analysis remains a significant challenge, as threats continuously evolve. This research proposes a machine learning–based framework for malware analysis and classification, leveraging features extracted from Portable Executable (PE) files. The proposed methodology incorporates a Transformerbased model trained on Speakeasy dataset to classify encoded API call sequences as benign or malicious. An attention-based explanation method, complemented by a GPT-driven interpretability mechanism, is introduced to enhance understanding of the model's predictions. The study also highlights interpretability benefits derived from Explainable AI techniques, as well as remaining limitations in detecting unseen malware families.

Paper ID 491

EVALUATING THE RELIABILITY OF GENERATIVE AI IN SOFTWARE ENGINEERING REFACTORING TASKS

Zhala Othman; Prof. Dr. Murat KARABATAK

Abstract- Software refactoring is essential as a maintenance method intended to simplify the structure of software code without changing any output functionality. Classical static analysis tools give consistent, rule-based advice; however, they have no semantic context to consider when specifying the necessary optimization in complex patterns. Lately, Large Language Models (LLMs), e.g., GPT-4, have appeared as a promising tool for this purpose, but they are not yet used in practice because of their notorious lack of reliability — specifically, hallucinations and functional regression. Contributions: In this paper, we introduce and assess a self-optimizing programming agent (SPA) that marries GPT-4 generativity with a deterministic safety verification loop. It leverages the Radon library for static complexity analysis and a dynamic module testing stack to ensure functional correctness. We present an experimental study in which the SPA relatively successfully refactored a highly complex real-world Python legacy code, reducing its Cyclomatic Complexityfrom 8 to 1 (87.5%) and increasing the maintainability index up to its theoretical maximum value. Crucially, the automatic feedback loop of the system was able to identify and correct initial flaws in the model and to keep its behavior unchanged. This work shows that LLMs can be harnessed to effectively address technical debt if they are integrated into a robust, self-correcting QA process.

Paper ID 492

SHORTCUTPROBE: IDENTIFYING BACKDOOR SAMPLES VIA INTERNAL STABILITY ANALYSIS

Mahaveer Prasad; Ajit Kumar Yadav; Rajeev Kumar; Victor Fedoseev

Abstract- Backdoor attacks compromise the trustworthiness deep neural networks by inserting hidden malicious behaviors during training. As a result, the model performs well on clean inputs but produces attacker-controlled predictions when a trigger is present, making poisoned samples hard to detect because they look normal at the output level. To address this problem, we propose ShortcutProbe, an inference-time detection framework that examines how the model internally makes decisions. Our key idea is that backdoor samples rely on shortcut pathways, focusing on a small set of critical neurons, while clean inputs depend on widely distributed features. To capture this difference, we use two complementary analyses. First, Active Neuron Dropout is applied to measure Prediction Shift Uncertainty, which quantifies how predictions change under controlled neuron perturbations. Second, a gradient-guided progressive masking strategy computes sudden confidence drop by tracking how prediction confidence drops as important channels are removed. By combining these signals, ShortcutProbe effectively distinguishes clean samples from shortcut-driven backdoor inputs. Experiments show that the proposed method achieves state-of-the-art detection performance with significantly reduced false positives.

Paper ID 497

AN INVESTIGATION OF THE RELATIONSHIP BETWEEN PRE-SERVICE TEACHERS' ARTIFICIAL INTELLIGENCE

Ahsen Aslantaş; Songül Karabatak; Murat Karabatak

Abstract- This study aimed to examine the relationship between pre-service teachers' artificial intelligence (AI) awareness levels and their AI anxiety levels. The research was designed using a quantitative relational survey model.

The study group consisted of 210 pre-service teachers enrolled at a faculty of education. Data were collected using the Artificial Intelligence Awareness Scale and the Artificial Intelligence Anxiety Scale. Descriptive statistics and Pearson product–moment correlation analysis were employed to analyze the data. The findings indicated that pre-service teachers' AI anxiety levels were low to moderate, whereas their AI awareness levels were relatively high. Correlation analysis revealed a statistically significant negative relationship between AI awareness and AI anxiety. These results suggest that as pre-service teachers' awareness of artificial intelligence increases, their anxiety toward AI tends to decrease. The findings highlight the importance of integrating AI literacy into teacher education programs to enhance awareness and reduce technology-related anxiety. The study provides implications for the design of pre-service training programs aimed at supporting both cognitive understanding and affective adaptation to emerging technologies in education.

Paper ID 499

DEEP LEARNING-BASED CLASSIFICATION OF CUTTING TOOL WEAR IN DRY TURNING OF TI-6AL-4V USING THERMOGRAPHY

Busra Tan Saatci, Mustafa Ulas, Turan Gurgenc, Engin Unal

Abstract- Monitoring cutting tool wear is a critical requirement for optimizing machining efficiency and ensuring surface integrity, especially when machining difficult-to-cut materials such as titanium alloys. This study proposes a novel, non-contact, deep learning-based approach to classifying cutting tool wear during dry turning of Ti-6Al-4V using infrared thermography. A comprehensive dataset of approximately 56,000 thermal images was compiled and labeled according to ISO flank wear standards. The performance of three transfer learning architectures—GoogLeNet, VGG16, and Inception V3—was evaluated to classify wear stages (low, medium, high). Inception V3 achieved the highest classification accuracy (99.79%) thanks to its superior feature extraction capabilities, while GoogLeNet also offered a highly competitive accuracy (99.53%) with a 70% reduction in training time, highlighting its suitability for real-time industrial applications. The results demonstrate that thermal patterns, stemming from titanium's low thermal conductivity, serve as reliable indicators of wear progression. This automated monitoring framework offers significant advantages for Industry 4.0 applications in the aerospace and automotive sectors, providing a cost-effective and scalable solution for predictive maintenance and zero-defect production. These results facilitate the transition from scheduled tool changes to condition-based maintenance in highvalue titanium machining.

SPONSOR

SOFTWARE AND CYBER SECURITY ASSOCIATION

Elazig – TURKEY

www.softcybersec.org **softcybersec@gmail.com**